

# Earth-Cosmos Binary

Seth D. Baum and Andrea Owe  
Global Catastrophic Risk Institute  
<https://gcrinstitute.org>

*Preprint currently under peer review. This version dated 26 September 2022.*

## Abstract

What should be done with the cosmos? This is an ethical question of profound importance. Furthermore, prospects for advanced technology mean that this is a question that humans may someday need to answer. This paper proposes one type of answer, the Earth-Cosmos Binary (ECB), in which Earth and its vicinity are preserved for humans and other existing Earth-life while the rest of the cosmos is allocated for the radical optimization of moral value. The ECB constitutes a balance between the moral view of preserving Earth and the rest of the cosmos in its current form and the moral view of converting the cosmos into something of greater moral value. Arguably, Earth and its vicinity is the portion of the cosmos most worth preserving due to the special nature of Earth. The paper considers a variety of arguments for and against the ECB, finding a strong but not definitive case in favor of it. The paper also presents variants of the ECB which may have additional moral value under certain conditions.

**Keywords:** ethics, preservation, replacement, cosmos, Earth, space settlement

## 1. Opening Concepts

### 1.1 Optimizing The Cosmos

*The Big Question:* What should be done with the cosmos?

*The cosmic genie:* Suppose humans uncover a genie that will grant exactly one wish. The genie cannot violate the laws of physics, and therefore is not magic, but it can do everything that is possible within the laws of physics. In particular, it is able to alter the makeup of the accessible portion of the cosmos as fast as is physically possible and for as long as is physically possible. The genie can (among other things) send spacecrafts to distant exoplanets, terraform, and rearrange planets into Dyson swarms. One further constraint: the genie cannot do its own moral reasoning. Humans can ask the genie to follow some specific moral framework (e.g., “maximize expected utility”), but they cannot ask the genie to come up with its own moral framework (e.g., “do whatever is good or right”). In this situation, what should humans wish for?

The cosmic genie is a thought experiment, but actual technological breakthroughs may approximate it. Some existing scholarship on potential future artificial intelligence (AI) and nanotechnology paints a picture of major technological breakthrough with implications of potentially astronomical proportion (e.g., Drexler, 2013; Bostrom, 2014). It is not known if or when this or other advanced technology will be developed. However, it is at a minimum plausible that humans may, in the not too distant future, have the capacity to substantially alter the accessible portion of the cosmos. In practice, actual human technology may not be as powerful as a hypothetical genie, but it may nonetheless be able to do a lot. The Big Question may be more than just a hypothetical.

Questions of what should be done are fundamentally a matter of ethics. As with other ethics questions, there is likely to be debate and disagreement on The Big Question. Hints of disagreement can be found in existing literature, for example in debates on whether or not it is morally acceptable to terraform other planets (Sparrow, 1999; Fogg, 2000; McKay & Zubrin, 2002; Schwartz, 2013), in calls

to use the cosmos to maximize some notion of welfare of humans and/or other sentient beings (Bostrom, 2008), or of life (Mautner, 2020) or ecological flourishing (Owe, 2022), and in broader literature debating the ethics of outer space (Schwartz & Milligan, 2016; Smith & Abney, 2019; Schwartz et al., 2022) and the distant future (Tonn, 2002). However, The Big Question has not yet been studied in much detail.

This paper outlines one type of answer to The Big Question. Here we are interested in the general structure of the answer and not in particular details. That includes details about specific ethical frameworks: utilitarianism, ecocentrism, etc. The proposed structure can be compatible with a range of ethical frameworks; this is one of its attractions. Analysis of specific ethical frameworks and other details is left for future research.

## 1.2 Replacement of Humans and Other Earth-Life

*Radical optimization:* An extreme alteration of some entity so as to make this entity optimized according to some optimization criterion. The entity can be an object (e.g., a rock) or a set of objects (e.g., the rings of Saturn) and can also include other phenomena such as electromagnetic radiation. Radical optimization entails rearranging the entity so that a certain criterion is optimized: for example, rearranging a planet into a giant mind for some combination of information processing and experiencing subjective well-being. Radical optimization can conceivably be done at any scale, from the microscopic to the astronomical. The optimization criterion can be a singular value (e.g., happiness) or a plurality of values aggregated in some way (e.g., an aggregate of happiness and biodiversity).

The radical optimization of the cosmos is one type of answer to The Big Question. This raises the question of what that means for life on Earth. First consider humans. The following concerns a type of situation that could potentially occur under certain circumstances. This situation can be important even without considering the cosmos, but it is also central to certain answers to The Big Question.

*The Difficult Question:* Given the opportunity to replace humans with something that is, in all morally relevant respects, superior, such that all humans would die, should this opportunity be taken?

It is plausible, and perhaps very likely, that human beings do not constitute the optimum configuration of matter according to a variety of moral standards. Humans are an artifact of evolution, not a product of moral optimization. To the extent that humans are optimized for anything, it would be their own biological (i.e., reproductive) success within their rather broad ecological niche. If humans were to wish for the cosmic genie to maximize, for example, some notion of welfare, the genie may proceed to replace humans with something that is more capable of producing or experiencing welfare. Ditto for other conceptions of moral good. In other words, the genie may radically optimize human bodies into something of greater moral value.

The Difficult Question is difficult because it involves an extreme divergence of major moral intuitions. One intuition is that moral value should be optimized. If there is opportunity to make something better, then it is good to do so. A related intuition is that the human species is not a morally relevant category, meaning that being a member of the human species does not automatically confer moral value. Work in environmental philosophy has long advocated against *speciesism*, meaning the practice of valuing one species over another for no reason other than its status as a particular species (Singer, 1977; Curry, 2011). The favored species is typically humans, in which case the speciesism is *ethical anthropocentrism*.<sup>1</sup> Research in transhumanism and human enhancement considers that human

---

<sup>1</sup> Ethical anthropocentrism, the view that humans are morally superior, can be distinguished from ontological anthropocentrism, the view that humans occupy a central or privileged place in the universe (Owe & Baum, 2021). For example, ethical anthropocentrism might claim that humans are morally superior to other animals, whereas ontological anthropocentrism might claim that humans are not animals. A common instance of ontological anthropocentrism is the use of the word “animals” to mean “animals other than humans”, which goes against the scientific finding that humans are members of the animal kingdom.

nature can be improved upon, and that it may be moral to do so (Persson & Savulescu, 2012), even if it results in human extinction (Shiller, 2017). In the astrobiological context, it has been argued that it may be optimal for human civilization to sacrifice itself to advance the moral value of an extraterrestrial civilization, or vice versa (Baum, 2010).

A countervailing intuition is that humanity, despite all its flaws, is good, and it would be a major loss for humanity to be gone; indeed human extinction may be a catastrophe of immense proportion. A related intuition is that it would be unfair and overly demanding to ask humans to cause the death of themselves, their friends and family, and everyone else they care about. Indeed, some perspectives argue that optimization should not be the primary moral objective. Non-consequentialist moral theories (e.g., deontology and virtue ethics) are commonly focused on other concepts besides optimization; even within the scope of consequentialism, some frameworks call for doing good without necessarily optimizing goodness (Jamieson & Elliot, 2009). Finally, there is the intuition that there is value in preserving things in more-or-less their current form, such as is common in the domain of environmental preservation (Rolston, 2020). Preservation does not require complete stasis, but it does entail maintaining conditions within some range: there cannot be *too much* change. The preservation of humanity still permits humanity to change, but it does not permit humanity to be made extinct.

In some scenarios, the distinction between preservation and replacement can be blurry. For example, it has been proposed that humans may someday merge with their technological creations into a new form of entity that transcends current human biology but retains a certain cognitive and cultural humanness (Kurzweil, 2005). This may occur, for example, via human minds being uploaded into computers (Hanson, 2016). If radical optimization involves phenomena such as these, then answering The Difficult Question may depend on how these hybrid scenarios are evaluated. For purposes of this paper, we set these possibilities aside to focus on the (arguably more probable) scenarios in which radical optimization clearly involves human replacement.

*The Difficult Question-Variant:* Given the opportunity to replace Earth-life with something that is, in all morally relevant respects, superior, such that all Earth-life would die, should this opportunity be taken?

The Variant is a non-anthropocentric articulation of The Difficult Question. To a large extent, the difficulty of The Difficult Question is not about humans per se, but instead is about the living beings that exist at the time when the question is faced.<sup>2</sup> One can reject ethical anthropocentrism and still be troubled by The Difficult Question. Or, one could not value humans at all, instead valuing other species, and be troubled by the Variant. The trouble comes from the dilemma of replacing everything currently good and familiar with something new and different and better.

### 1.3 A Reconciliation

Two possible answers to The Difficult Question are as follows:

*Preservation:* The view that humans and other Earth-life should be preserved, even under the possibility of radical optimization.

*Replacement:* The view that humans and other Earth-life should be replaced if doing so brings greater moral value, especially via radical optimization.

This paper presents an answer to The Big Question that, contingent on details discussed below, may constitute an optimal balance between preservation and replacement:

---

<sup>2</sup> The Variant is *biocentric*, meaning that it favors life over non-life. Some environmental philosophy research has advocated for biocentrism (Taylor, 1986). Other research has argued for also valuing at least some abiotic entities, such as art (Budd, 1995) and mountains and rivers as parts of ecosystems (Rolston, 2020). The Variant could be rewritten in a non-biocentric form without loss of generality.

*The Earth-Cosmos Binary (ECB):* Preserve Earth, plus some surrounding vicinity of Earth, for humans and other Earth-life. Radically optimize the rest of the accessible cosmos for some conception of moral value. A sketch appears in Figure 1.



**Figure 1.** Earth-Cosmos Binary. The boundary geometry shown here is an illustrative sketch.

The ECB has three variables: (1) where to set the boundary between (a) Earth and its surrounding vicinity and (b) the rest of the cosmos; (2) in what way to optimize the rest of the accessible cosmos, or what optimization criteria to pursue; and (3) what to do on Earth and in its surrounding vicinity. These variables can be set in many ways. Therefore, the ECB is not a singular answer to The Big Question, but instead is a class of answers that all have the same ECB structure. This paper comments on how the three variables can be set, but it does not offer singular answers for how they should be set. Instead, the focus is on the merits of the ECB structure. Details of how to set the three variables is left for future research. There are also variants of the ECB, discussed in Section 3.

Though the ECB has not previously been analyzed in depth, various aspects of it have gotten some prior attention. As is often the case (Camarena, 2020), these ideas were first explored in science fiction. The earliest discussions of ECB themes that we are aware of are in science fiction. In *Diaspora* (Egan, 1997), Earth is left for a variety of biological and enhanced humans, while the superior disembodied and computerized intelligences exist in simulated realities running on hardware located throughout the Solar System. In *Robot* (Moravec, 1998), it is suggested that the superior beings (in this case AIs and enhanced humans) should be allowed to pursue their goals away from Earth, leaving regular humans to pursue their own ambitions on Earth. In *Accelerando* (Stross, 2005), the situation is reversed: Earth and the inner solar system are occupied by advanced AIs while humans are pushed further out into the cosmos.

In the academic literature, the idea of the ECB can be traced to the concept of “planetary parks”, which are regions of planets that are intentionally preserved in their original form when humans settle on other portions of the planets (Cockell & Horneck, 2004, 2006). De Garis (2005) proposes that advanced AI might—but would not necessarily—replace humans with something that is in some respects superior, and that humans will be deeply divided over whether to build such AI. This divide essentially maps to differing perspectives on The Difficult Question. Bostrom (2014, 219-220) considers using AI to turn almost the entire accessible universe into “hedonium”, which is an arrangement of matter that optimizes pleasurable experience. Under this proposal, the Milky Way galaxy would be preserved for humans and Earth-life. Bostrom advocates for this proposal, making an argument similar to what is below described as the moral compromise argument. Finally, Holt (2021) considers a humankind so technologically mature that it could colonize outer space and then sever all ties with Earth. At that point, the remaining Earth biosphere could be preserved or replaced with something else. Holt argues that Earth’s biosphere should be preserved and that humans should not sever ties with it, on grounds that it would be in our own best interest and that the Earth biosphere is of great moral value of its own.

## **2. Evaluating the Merits of the Earth-Cosmos Binary**

Whether the ECB should be pursued depends on how various arguments and factors are resolved, important examples of which are surveyed in this section.

### **2.1 The Moral Compromise Argument for the ECB**

A major attraction of the ECB is that it offers a compromise—potentially a very strong compromise—between preservation and radical optimization. Essentially, the ECB divides the cosmos into two parts: one for preservation and another for radical optimization. This does entail sacrifices for both sides: there is no radical optimization within Earth’s vicinity and there is no opportunity for Earth-life to expand beyond Earth’s vicinity.<sup>3</sup> The ECB nonetheless accommodates the moral intuitions in favor of preserving humans and other Earth-life while also accommodating, to some degree, the moral intuitions in favor of radical optimization.

The structure of the cosmos suggests that the ECB could offer a very high degree of accommodation to both preservation and radical optimization. In other words, neither side would sacrifice much. Earth constitutes a very small portion of the cosmos. The Sun is just one of an estimated  $10^{11}$  to  $10^{12}$  stars in the Milky Way; the Milky Way is just one of an estimated  $10^{11}$  to  $10^{12}$  galaxies in the universe (ESA, 2022). The entire universe may not be accessible, but even if only a small portion of it is, that could still constitute a much larger region than Earth. The ECB could allocate the entire Solar System or potentially even the entire Milky Way to preservation and still leave almost all of the accessible cosmos for radical optimization. Thus, the loss for the replacement view could be extremely small.

Furthermore, Earth-life has thus far done very little with any region outside the Solar System. There have been just five interstellar probes: Voyager 1, Voyager 2, Pioneer 10, Pioneer 11 and New Horizons. Aside from these probes, the only other engagement between Earth-life and extrasolar space has been via artifacts of extrasolar space (mainly electromagnetic radiation) reaching Earth, such as for stargazing and astronomy. The ECB could limit Earth-life’s opportunity for future space missions and astronomy, but perhaps not to an extreme degree. Regardless, other facets of life on Earth could continue more-or-less in their current form, perhaps with some substantial but non-radical improvements from the cosmic genie or the advanced technology.

### **2.2 The Radical Optimization Argument Against the ECB**

It can be argued that Earth and its vicinity should not be preserved, and instead that it should be radically optimized. In other words, radical optimization is more important than preservation such that no compromise should be made. This argument draws parallels to environmentalism on Earth: Perhaps the important part is not what something is or was at a given time, but how good it can be. For example, one might see moral value in natural ecosystems due to their complexity, richness, and diversity of life. In this case, it might be better to turn a desert into a forest than to preserve the desert. Similar arguments have been made in favor of terraforming lifeless planets (McKay, 2009; Grinspoon, 2016) and spreading Earth-life into space (Margulis & West, 1997; Mautner, 2020). The logic here is sound: given that moral value should be optimized, if radically optimizing Earth and its vicinity would bring about more moral value, then it would be morally good to do so.

It is indeed correct that from this perspective, the ECB is morally suboptimal. Arguments against speciesism and anthropocentrism are relevant in this context. We find the arguments against speciesism and anthropocentrism to be absolutely compelling: it is wrong to favor a species (human or other) due to the mere fact of it being that species. We likewise have some sympathy for the radical optimization argument. Still, we find there is a compelling case for the ECB instead of full radical

---

<sup>3</sup> The lack of opportunity for Earth-life to expand beyond Earth’s vicinity assumes that Earth-life would have no role in radical optimization.

optimization for reasons discussed throughout this paper. However, the case is not definitive: we cannot definitively reject the radical optimization argument.

### **2.3 The Cosmic Preservation Argument Against the ECB**

It can be argued that no radical optimization of the cosmos should be pursued, whether via the ECB or other approaches. This is essentially to say that the cosmos should be preserved in roughly its current form, analogous to arguments for environmental preservation on Earth. The cosmos is inherently dynamic and cannot be preserved in its current form in perpetuity. Nonetheless, it can be preserved from rapid and radical alteration by Earth-originating civilization. Note that this is a much more extensive form of preservation than has been discussed previously. Sections 1.3 and 2.1 focus on the preservation of Earth-life. The cosmic preservation argument calls for the preservation of the entire cosmos, including its vast abiotic portions.

A case for cosmic preservation can be made based on the moral value of processes that are “natural” in the sense of being unaltered by technological civilization.<sup>4</sup> Environmental ethics research sometimes argues that there is moral value in preserving and continuing Earth’s natural evolutionary history, and that this provides a basis for protecting existing biodiversity (Rolston, 1988). A similar argument has been made for the moral value of astronomical processes as a reason for humans to preserve extraterrestrial environments as they are (Rolston, 1986; Milligan, 2015).<sup>5</sup> This argument sees moral value in natural processes even where no life is involved, as these processes are the creative forces underlying everything that exists in the universe, including life on Earth (Rolston, 1986). Some further argue that it would be geocentric to impose Earth-based conceptions of moral value onto other planets and celestial places (Milligan, 2015; Schwartz, 2018).

We are skeptical of the cosmic preservation argument. The basic logic checks out: if the entire cosmos should be preserved in (more or less) its current form, then there would be no point in the ECB or other radical optimization because the cosmos is already in its optimal form. However, we doubt that the entire cosmos should be preserved in its current form. To the contrary, we believe that, for at least a large portion of the cosmos, any value of preserving natural cosmic processes is likely to be outweighed by the value of turning the cosmos into something better. It seems clear that Earth, with all its life and vitality, is of much greater moral value than other known planets.<sup>6</sup> It is almost unthinkable to suggest that it would be good for Earth to be converted into another dead planet. Conversely, it seems clearly good to convert other planets into Earths, or more generally to convert the cosmos into something of greater moral value.

### **2.4 The Optimality Argument for the ECB**

It is possible that both the radical optimization argument and the cosmic preservation argument are incorrect, and that the ECB may actually be the optimal configuration of the cosmos, at least for some form of the ECB with some specification of its three variables.

The optimality of the ECB could follow from the tension between radical optimization and preservation. Perhaps there is value to preservation, but this is only one finite value to be weighed against other finite values. The improvements created through radical optimization could be one of these other values. What to do may depend on whether, for that something, the value of preservation

<sup>4</sup> There is a different sense in which everything is “natural”: everything consists of matter and follows the laws of nature; nothing is supernatural.

<sup>5</sup> Other arguments for the moral value of abiotic extraterrestrial environments have been made, such as due to the aesthetic quality of extraterrestrial landscapes (Milligan, 2015). These arguments may provide less basis for cosmic preservation due to the possibility of radically optimizing these conceptions of moral value. Perhaps a cosmic genie could radically improve the aesthetic quality of extraterrestrial landscapes.

<sup>6</sup> It remains possible that the search for extraterrestrial intelligence (SETI) will identify planets that are similarly full of life and vitality.

exceeds the value of radical optimization, as it is not possible to both preserve and radically optimize the same thing. The value of preservation may be strongest for Earth due to its biosphere and general planetary dynamism. It could be the case that the value of preservation exceeds the value of radical optimization on Earth but not in the rest of the cosmos. In that case, the optimality argument would be correct and the ECB would be optimal.

### **2.5 The Hybridization Argument Against the ECB**

The case for the ECB hinges on there being value to the spatial separation between preservation and radical optimization. Alternatively, there could be more value to a hybrid structure in which radically optimized entities and existing forms of Earth-life coexist alongside each other.

A case for hybridization can be made on grounds that Earth-life and radically optimized entities would be of mutual benefit. Perhaps Earth-life could provide functional assistance to radically optimized entities, though this may be unlikely if the radically optimized entities are sufficiently superior. If nothing else, perhaps radically optimized entities might enjoy having Earth-life around, such as in the role of a pet or a specimen of historical interest. Likewise, perhaps Earth-life would benefit from exposure to radically optimized entities, such as to receive help in addressing various problems on Earth.

We are skeptical of the hybrid structure. The removal of the spatial separation between preservation and radical optimization implies that Earth is not substantially preserved. It is difficult to imagine much value to preserving Earth-life without also preserving Earth. For example, if the radically optimized entities might enjoy Earth-life as pets, they might get even more enjoyment out of Earth as something akin to a wildlife preserve. The idea of radically optimized entities interacting with Earth does show that there could be value to versions of the ECB in which the boundary is somewhat blurry. Nonetheless, we believe that it would be better to maintain some form of spatially separated binary structure.

### **2.6 The Infinite Value Argument Against the ECB**

It is sometimes proposed that the possibility of infinite value in the cosmos poses severe challenges for moral evaluation (Dubey & Laguzzi, 2021; Wilkinson, 2022). This issue derives from the common mathematical supposition that  $\infty+x=\infty$  for any finite or infinite  $x$ . If this is the case, then as long as moral value extends into the infinitely distant future, it does not matter what happens along the way. This is an issue of conceptual importance for all moral theories that aggregate value across space and time.

The issue of infinite value is of particular relevance to scenarios involving expansion into outer space, including but not limited to the ECB and other potential answers to The Big Question. Given the finite habitability of Earth, something elsewhere in the cosmos is needed to achieve infinite value. Current physics suggests that infinite value may be impossible, such as due to the heat death of the universe, but the matter has not been conclusively resolved. As long as there is some nonzero probability of infinite value, then the mathematical issue persists. In that case, there may be no moral distinction between the ECB and the alternatives, as long as they each have some nonzero probability of resulting in infinite value. Resolution of the infinite value argument requires technical analysis beyond the scope of this paper.

### **2.7 The Political Compromise Argument for the ECB**

If humans discover a cosmic genie or invent advanced technology, then they may disagree amongst themselves about what to wish for or do with it. Indeed, disagreement on matters of ethics is an enduring feature of human populations. Political compromise entails identification of a position that all relevant human parties are willing to agree with. The ECB may not resolve all political disagreement.

In particular, there could still be disagreement over how to set the three ECB variables. Nonetheless, given the difficulty of The Difficult Question, it is plausible that the ECB could be a major point of political compromise.

Political compromise can be morally valuable as a matter of conflict resolution. Disagreement over The Big Question could result in violence, perhaps even large-scale war. The worst-case scenario could be a war so large that it destroys human civilization, perhaps even resulting in Earth-originating civilization never doing anything further with the cosmos. If advanced technology is developed in time scales of decades or perhaps even centuries, it may be done in a world with nuclear weapons. The world currently possesses an estimated 13,400 nuclear weapons (Kristensen & Kile, 2020). The organization Global Zero has presented a plan for full nuclear disarmament by 2045 (Global Zero, 2022), though recent events such as the Russian invasion of Ukraine suggest a more sluggish pace for disarmament (Notte & Bidgood, 2022). The potential for nuclear war to constitute a civilization-ending catastrophe is a matter of ongoing research (Baum, 2015; Scouras, 2019). Even if nuclear war would not cause such an extreme catastrophe, it would still constitute a major harm. Ditto for large-scale wars fought with other weapons. Even small-scale violence would be a loss.

Given the astronomical stakes of The Big Question, it is not unreasonable to suppose that there could be war. Prior research has proposed that there could be war between (A) rival sides on the question of whether or not to build advanced technology in the first place (de Garis, 2005) and (B) rival groups each seeking to build advanced technology on their own terms (Miller, 2012). The ECB could help with both scenarios. Regarding (A), debate over whether to build advanced technology may hinge on the threat it poses to humans; this is the case in the scenario of de Garis (2005) and other work such as Joy (2000). The ECB is a way to build advanced technology without threatening humans. Regarding (A), research in AI ethics has documented a strong but not universal orientation toward anthropocentrism in both the overall field of AI ethics (Owe & Baum, 2021) and in specific projects oriented toward advanced AI (Fitzgerald et al., 2020). This suggests that many developers of advanced technology may have an affinity for preservation of humans, with some favoring radical optimization. Perhaps all developers could agree on the ECB.

Political compromise can be morally valuable as a matter of procedural justice: no participating party is forced to accept something they object to. This is another reason to favor the ECB, insofar as it is indeed a point of political compromise. However, there are limits to this value. A just process arguably should account for the perspectives of all parties affected by a decision, not only those participating in the decision. Decisions on The Big Question could affect everything that exists on Earth and the entire accessible cosmos into the distant future. It is physically impossible for all affected parties to participate. The best that can be done may be some sort of proxy representation (Tonn, 1996; O’Neil, 2001). The task is further complicated by the fact that The Big Question involves decisions on which types of future individuals to create, such as future humans vs. future radically optimized entities. All of this reduces the extent to which a consensus among humans involved in handling the cosmic genie or developing advanced technology could be procedurally just.

## **2.8 Handling Uncertainty About The Big Question**

There are several sources of uncertainty surrounding the ECB as an answer to The Big Question.

One source of uncertainty is in the implementation of potential answers to The Big Question. Nothing like radical optimization has ever been done before, and it is not currently known how well it would work. Indeed, discussions of radical optimization sometimes propose that it may fail, resulting in zero value. In one type of scenario, an advanced AI system interprets its instructions differently than humans intended (as computers are wont to do). For example, humans may attempt to instruct the AI to maximize happiness, but instead the AI converts the accessible cosmos into smiley faces (Loosemore, 2014). An entire cosmos of smiley faces may be of zero moral value. That would be a catastrophic loss



relative to the current state of affairs in the cosmos, which at least includes humans and other Earth-life.

There is also uncertainty about the politics of The Big Question. Because technology akin to the cosmic genie does not yet exist, there has not yet been significant debate about what to do with it. The extent of disagreement over what to do remains uncertain. Perhaps disagreement would be intense or even catastrophically violent. Or, perhaps disagreement would be more limited.

A different source of uncertainty is on how in principle to answer The Big Question. This is an instance of “moral uncertainty”, meaning uncertainty about which moral values should be used (MacAskill et al., 2020). Any disagreement over how to answer The Big Question could indicate that humanity has not fully thought through all aspects of The Big Question and may have not yet identified the correct answer. Indeed, human morality has changed significantly over the years and may continue to change in the future (Danaher, 2021); this suggests that current understandings of morality may be flawed, and likewise that people today should hold some degree of moral uncertainty.

Given uncertainty about The Big Question, there is value to approaches that are crafted with this uncertainty in mind. The ECB does well in this regard, but other approaches are also worth noting.

The ECB is attractive as a hedge against some of the uncertainty described above. Even if the implementation of radical optimization fails, resulting in zero moral value where it is implemented, the ECB will still at least have moral value from Earth and its vicinity.<sup>7</sup> The ECB could reduce political disagreement, which could avoid catastrophic violence. If political disagreement would otherwise be limited, then the ECB does not add much value, but it would not make things worse. With respect to moral uncertainty, the ECB is attractive because it implements two moral concepts, preservation and radical optimization, and provides a high degree of accommodation to both, as discussed in Section 2.1.

An alternative approach to answering The Big Question is to abstain from providing an immediate answer and instead arrange a process for answering it at some later time, especially after having given The Big Question more careful consideration. Prior literature has proposed this as an “AI nanny” (Goertzel, 2012), in which an AI guides the process for more careful consideration, or a “long reflection” (Ord, 2020), in which humans lead the process. Such an approach does not constitute an argument for or against the ECB per se. Instead, it can be taken to be an argument against attempting to implement or even study the ECB or other answers to The Big Question at this time. Given the opportunity, the idea of postponing decisions on The Big Question may have some merit. Evaluating the merits of postponement is beyond the scope of this paper. What can be said is this: If there is opportunity to postpone, and if postponement would be a wise course of action, then current analyses of The Big Question, such as this paper, would at least provide a head start. Alternatively, if decisions cannot or should not be postponed, then current analyses of The Big Question are that much more important and urgent.

It is possible that even after extensive inquiry into The Big Question, a definitive answer would not be reached. Indeed, philosophers have been studying ethics for millennia and consensus is nowhere in sight. In the absence of a definitive answer, it may be necessary for humanity to make its best guess on what to do. That could mean pursuing a specific course of action across the cosmos such as the ECB. Alternatively, it could mean abstaining from acting. It could be argued that it would be inappropriate to alter the cosmos despite uncertainty about how it should be altered. Doing so may be incautious or hubristic.<sup>8</sup> We are skeptical of this argument. Arguably, certainty should be as high as possible, but

---

<sup>7</sup> This specifically applies to scenarios in which the wrong thing is radically optimized, as in the smiley face example. There could be other scenarios in which the implementation failure results in a failure to maintain the Earth-cosmos boundary, in which case the value on Earth and its vicinity could also be destroyed.

<sup>8</sup> Under certain interpretations of the precautionary principle, no major transformative actions should be taken unless it is certain that they would be good (Sunstein, 2002). Human hubris has arguably caused major problems such as climate change (Sadler-Smith & Akstinaite, 2021); pursuing the ECB or other radical optimization of the cosmos could be

virtually all decisions have some uncertainty, especially important large-scale decisions. It is essential to be able to make these decisions despite the persistence of *some* uncertainty. The optimal timing of decision-making given lingering uncertainty and the potential to reduce it is an important topic beyond the scope of this paper.

### 3. Variants of the Earth-Cosmos Binary

The essence of the ECB is to answer The Big Question by preserving some region and radically optimizing some other region. The ECB preserves Earth plus some vicinity of it, but other geometries are possible for the boundary between preservation and radical optimization. The ECB also radically optimizes the rest of the cosmos according to one fixed optimization criterion, but other geometries of radical optimization are also possible. Other geometries are explored here. The geometries are given terminology under the following format: “[geometry of preserve boundaries] [geometry of radical optimization]”. This terminology format is also used in “Earth-Cosmos Binary”: the preserve boundary is between Earth and the cosmos, while the radical optimization geometry is a binary between the preserve and the rest of the cosmos. At the end of the section, an expanded terminology is introduced for variants in which the moral values of radical optimization are determined by extraterrestrial beings.

#### 3.1 Earth-Cosmos Ternary and Further Recursion

If and when the ECB is established, it may begin with radical optimization in the region just on the other side of the boundary from Earth. That would be the closest region from Earth where radical optimization could be performed and so it makes for a likely starting place. The radical optimization may then expand outwards from there in an increasingly macro fashion.

Suppose that during the process of outward expansion, new concepts or techniques for radical optimization are discovered. For example, if the process is led by an advanced AI, then that AI may learn new concepts as it proceeds. Indeed, current AI systems are already capable of some degree of learning, though presumably much less learning than this sort of advanced AI. By this point, there has already been some radical optimization within some region immediately surrounding the Earth-plus-vicinity preserve. The initial radical optimization has proceeded according to concepts or techniques that have since been found to be suboptimal. Moral value could potentially be increased by going back and redoing these regions, that is by re-radically optimizing the initial radical optimization.

The prospect of re-radical optimization raises the same issues as the radical optimization of Earth. It would be a second iteration of The Difficult Question. Should the initial radical optimizer replace itself? Likewise, the matter could be resolved by a second iteration of the ECB. Instead of re-radical optimizing the initial band of radical optimization, it could be left in its initial form. A new boundary would be set. On the outer side of the boundary, radical optimization would proceed using the new and improved concepts or techniques. The result would be an *Earth-Cosmos Ternary*: three concentric value domains as sketched in Figure 2a. Following the same logic, there could be subsequent iterations producing an *Earth-Cosmos Quaternary*, an *Earth-Cosmos Quinary*, or even indefinite recursion in an *Earth-Cosmos Infinary* as sketched in Figure 2b. Evaluation of the merits of these various structures could proceed in a similar fashion as described above.

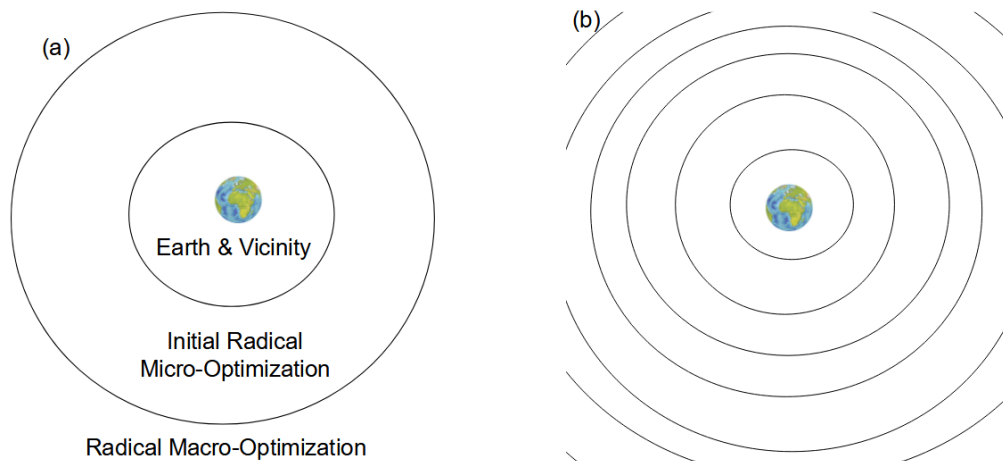
#### 3.2 Earth-Cosmos Pienaries

The issue of political compromise speaks to a deeper underlying issue: the human population is morally diverse, with different people favoring different moral values (Haidt et al., 2003), including those favoring a plurality of moral values. It is possible that moral diversity would persist even after an extended process following a postponement of answering The Big Question—for comparison,

---

similarly hubristic, causing major unforeseen problems.

philosophers have been studying morality for millennia and continue to lack consensus on basic questions, while various nations and political factions continue to maintain adversarial relations.

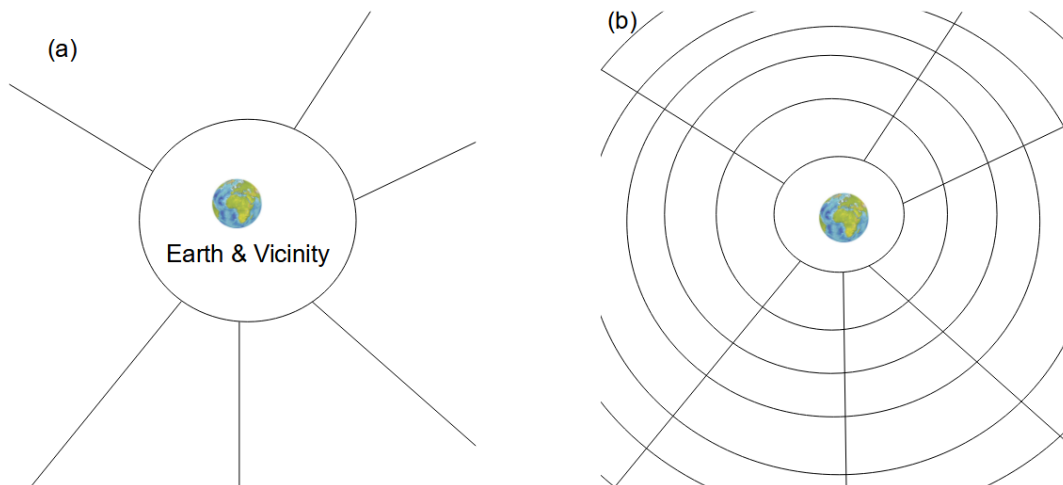


**Figure 2.** (a) Earth-Cosmos Ternary; (b) Earth-Cosmos Infinary. Boundary geometries shown here are illustrative sketches.

Instead of insisting on one consensus answer to The Big Question, a multitude of answers could be pursued. Under this approach, Earth and its vicinity could remain morally and politically divided in the same way that it currently is. The rest of the cosmos would be divided into segments, each of which would be radically optimized according to a different optimization criterion corresponding to a different set of moral values. The exact distribution of segments could be set according to some moral or political process to be determined through future research or by the participants in the process. The boundaries segment between segments may be set radially in a geometry similar to pie slices, henceforth termed a *pienary*. Figure 3a sketches an *Earth-Cosmos Bipienary*, meaning an Earth-Cosmos Binary overlaid with pienary geometry. Figure 3b sketches an *Earth-Cosmos Infipienary*, meaning an Earth-Cosmos Infinary overlaid with pienary geometry. Pienaries may be politically advantageous by avoiding the need to reach global consensus. Likewise, evaluation of the merits of pienaries may be similar to evaluation of political compromise for the ECB (Section 2.7). An additional issue is on the allocation of the cosmos to different factions. As Figure 3 illustrates, there is no requirement that each faction receive an equal allotment.

### 3.3 Indigenous Cosmic Cheese

At present, Earth is the only known inhabited planet. However, the search for extraterrestrial life/intelligence/etc. continues. It remains possible that, elsewhere in the cosmos, there are other living and/or intelligent beings. More generally, there may be other places of special moral significance in the way that Earth may be of special moral significance. These places may likewise be worth preserving in the same way that the ECB preserves Earth. Indeed, several scholars have argued for a variety of moral values present in extraterrestrial places (Rolston, 1986; Milligan, 2015; Lupisella, 2016; Schwartz, 2018). ECB designs that do not account for this may destroy such places via radical optimization. Doing so may be morally suboptimal in the same way that radically optimizing Earth may be morally suboptimal. There may additionally be a case for preserving these places analogous to the moral compromise argument for preserving Earth (Section 2.1).



**Figure 3.** (a) Earth-Cosmos Bipienary; (b) Earth-Cosmos Infipienary. Boundary geometries shown here are illustrative sketches.

Within the human population, people who live in a place where they have deep history are known as Indigenous. Similarly, it can be said that all humans are indigenous to Earth, and furthermore that any extraterrestrial populations are indigenous to wherever in the cosmos they are from. On Earth, many Indigenous human populations have been treated quite poorly, including genocide. It is now commonly recognized that poor treatment of Indigenous human populations is a mistake. If and when humanity expands into outer space, there may be opportunity to avoid repeating this mistake.

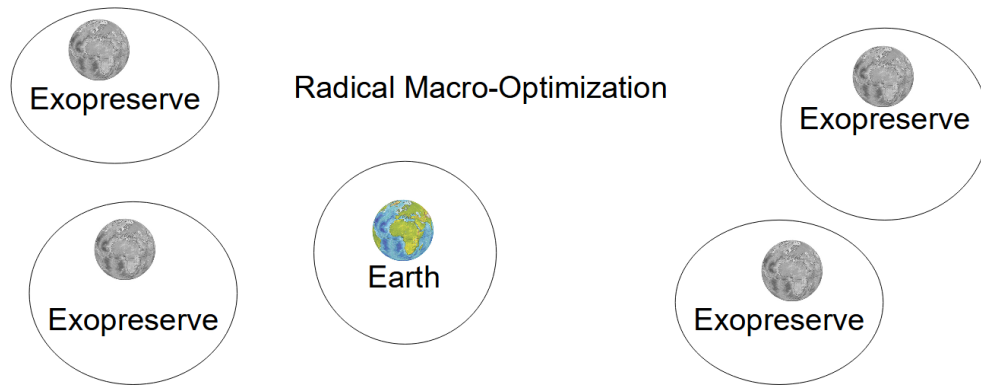
Therefore, instead of radically optimizing the entire cosmos outside of Earth’s vicinity, select extraterrestrial locations could be preserved. Such preserves, which may be termed *exopreserves*, could resemble national parks on Earth, and planetary parks and solar system heritage zones as discussed by Cockell and Horneck (2004, 2006) and Oman-Reagan (2016a, b) or much larger. The exopreserves would protect populations and any other sources of moral value that are indigenous to that location. The result would be a geometry resembling Swiss cheese (Feinstein et al., 2016), with holes for preserves including but not limited to Earth. Therefore, this design is termed *Indigenous Cosmic Cheese*.

Indigenous Cosmic Cheese specifically refers to the boundary between the regions that are preserved and the regions that are radically optimized. Grammatically, it is analogous to “Earth-Cosmos”. If the radical optimization uses a single standard of value, as in the ECB, then the result is an Indigenous Cosmic Cheese Binary, as sketched in Figure 4. Alternatively, if the radical optimization uses multiple standards of value, then other geometries can result. Figure 5 sketches an Indigenous Cosmic Cheese Infipienary. The exopreserves in Figures 4 and 5 all preserve exoplanets and their vicinities, though exopreserves could also focus on preserving other things besides exoplanets.


### 3.4 Extraterrestrial Value Sources

The Indigenous Cosmic Cheese Infipienary sketched in Figure 5 illustrates an attribute of all frameworks considered thus far: all radical optimization emanates from Earth. There is a sense in which this is *ethically geocentric* in the sense that it places Earth at the moral center of the cosmos (Fogg, 2000). Ethical geocentrism can be distinguished from empirical geocentrism, which situates Earth at the geometric center of the cosmos. Arguably, ethical geocentrism is morally wrong because it favors Earth values for no reason other than the fact that the values are from Earth. The issue here is analogous to *ethical anthropocentrism*, in which humans are morally favored over other species for no

reason other than the fact of them being human.<sup>9</sup> Arguably, both ethical anthropocentrism and ethical geocentrism constitute biases that can result in morally suboptimal outcomes.



**Figure 4.** Indigenous Cosmic Cheese Binary. Boundary geometries are illustrative sketches.



**Figure 5.** Indigenous Cosmic Cheese Infipienary. Boundary geometries are illustrative sketches. The blue-green planet is Earth; the grayscale planets are exopreserves.

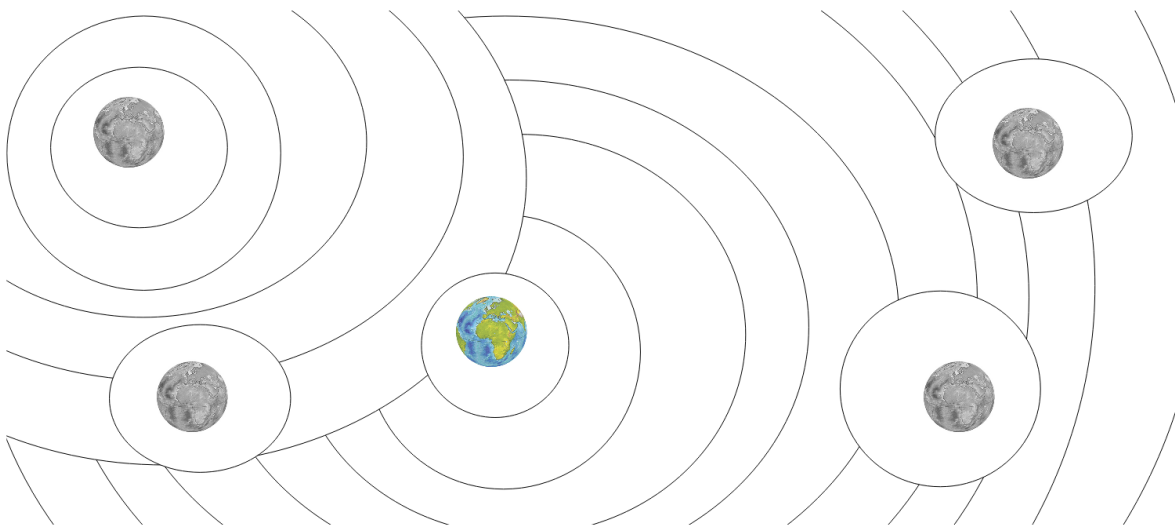
An unbiased pursuit of moral value in the cosmos should be open to the possibility that other (non-Earth-originating) civilizations may have moral values that equal or exceed that of humans and other Earthlings (Baum, 2010; Vakoch, 2014). Other civilizations may also have values that complement Earth values, such that the best values would be some combination of both. A non-ethically-geocentric answer to The Big Question would consider extraterrestrial values on equal footing with Earth values. Let us assume that when humans initially answer The Big Question, no extraterrestrial civilizations are involved in the decision. This could be because no contact with extraterrestrial civilizations has yet been made, as is currently the case. Under this assumption, humans would answer The Big Question in a way that accounts for the possibility of later encountering extraterrestrial civilizations and considering their values. It is plausible that extraterrestrial civilizations would not be encountered until Earth-originating radical optimization spreads across the cosmos.

If one extraterrestrial location is found to have moral values that are superior to all others, the non-preserve cosmos could be re-radically optimized according to these values. The effect would be an initial wave of radical optimization emanating from Earth across the accessible cosmos, followed by a

<sup>9</sup> See footnote 1 above for elaboration on ethical anthropocentrism.

second wave of radical optimization emanating from the extraterrestrial location. This process is sketched in Figure 6. The extraterrestrial value source in Figures 6 is assumed to be an exoplanet, though value sources could be other objects than planets.<sup>10</sup>

Accommodation of extraterrestrial value sources requires new terminology. The terminology format used thus far, “[geometry of preserve boundaries] [geometry of radical optimization]”, has nothing to indicate the source of moral value used in radical optimization. Instead, the terminology has assumed that radical optimization emanates from Earth. That assumption is no longer valid. To accommodate extraterrestrial value sources, the terminology is expanded to the format “[geometry of preserve boundaries] [source of value] [geometry of radical optimization]”. Thus, “Earth-Cosmos Binary” as sketched in Figure 1 becomes *Earth-Cosmos Earth Binary*, while “Indigenous Cosmic Cheese Infipienary” as sketched in Figure 5 becomes *Indigenous Cosmic Cheese Earth Infipienary*. Figure 6 shows the transition from *Indigenous Cosmic Cheese Earth Infinary* to *Indigenous Cosmic Cheese Exoplanet Infinary*, for some unspecified exoplanet.



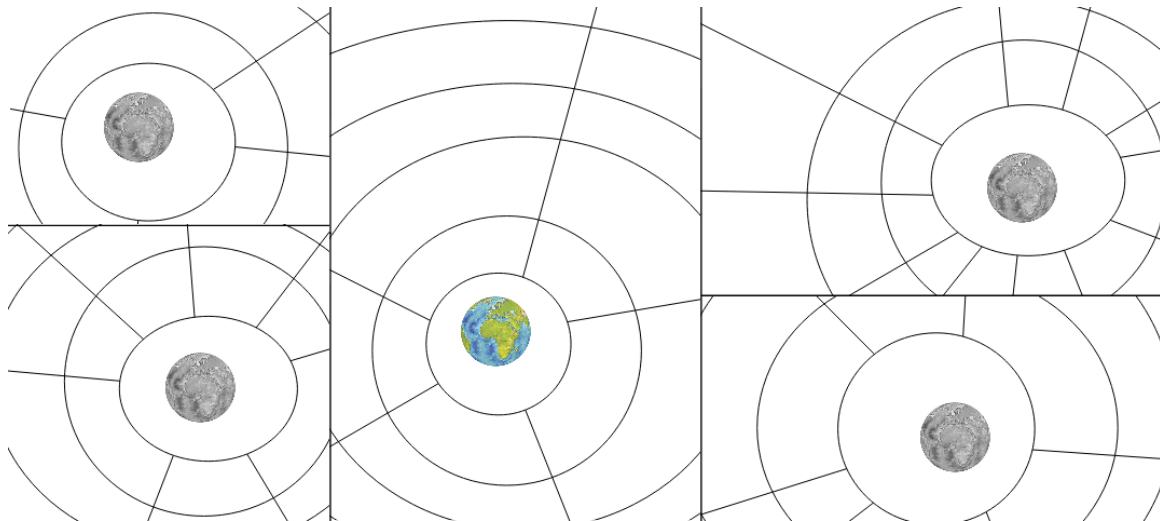
**Figure 6.** Transition from Indigenous Cosmic Cheese Earth Infinary to Indigenous Cosmic Cheese Exoplanet Infinary. Boundary geometries are illustrative sketches. The blue-green planet is Earth; the grayscale planets are exoplanets.

The Indigenous Cosmic Cheese Exoplanet Infinary sketched in Figure 6 recognizes that extraterrestrial values may be superior to those of Earth. However, it still assumes that there is one superior value framework that should be pursued across the accessible non-preserve cosmos. A case can be made for a more pluralistic approach advancing a diversity of values originating from across the cosmos, similar to how the Earth-Cosmos Pienary accommodates a diversity of values originating from Earth. Pienaries allocate slices of the cosmos emanating from a single value source, be it Earth or somewhere else. New geometry is needed to accommodate multiple value sources. Figure 7 sketches what can be called an *Indigenous Cosmic Cheese Multisource Infipienary*, in which infipienaries emanate from multiple value sources, each of which is also preserved. Each infipienary is allocated an adjacent region of the cosmos.

Extraterrestrial value sources raise novel issues. Politically, there are issues deriving from the relative power of Earth-originating radical optimization and extraterrestrial civilization(s). Perhaps the Earth-originating radical optimization would be more powerful, perhaps due to it being a radical optimization using the resources of a substantial portion of the cosmos. In that case, allocating some

<sup>10</sup> The same applies for the value sources in Figure 7, to be introduced below.

portion of the cosmos to extraterrestrial value source(s) involves the Earth-originating radical optimization deferring to the extraterrestrials. For this to occur, humans may need to build this deference into their answer to The Big Question. This requires a certain generosity and high-mindedness from humans. It is unclear whether humans answering The Big Question would accept this.



**Figure 7.** Indigenous Cosmic Cheese Multisource Infipienary. Boundary geometries are illustrative sketches. The blue-green planet is Earth; the grayscale planets are exoplanets.

Alternatively, perhaps the extraterrestrial civilization(s) would be more powerful, perhaps due to them being older and more technologically advanced. In that case, the extraterrestrial civilization(s) may overturn Earth-originating radical optimization by force. No human generosity is needed for this. However, this begs the question of why such an advanced civilization had not already radically optimized the accessible cosmos in the first place. The Fermi Paradox is relevant here: if some extraterrestrial civilization has already radically optimized the accessible cosmos, then current human astronomy presumably would have already detected it. Or, perhaps Earth is already contained within something akin to an exopreserve that is designed to hide the radical optimization from Earthlings.

A third political possibility is for Earth-originating radical optimization and extraterrestrial civilization(s) to have roughly equal power. This may be exceedingly unlikely, requiring an unusual coincidence. Nonetheless, in this type of scenario, there may need to be political negotiation similar to negotiation between humans on their initial answer to The Big Question.

Extraterrestrial value sources also raise moral issues. One is the issue of how to compare human and extraterrestrial values. It is plausible that extraterrestrial values could be superior to human values, such that the optimal approach is the Indigenous Cosmic Cheese Exoplanet Infinary, in which radical optimization switches from human to extraterrestrial values. This entails a decision of whether to make this switch, which raises issues of how to make that decision. The switch would presumably be made if the extraterrestrial values are found to be superior to the human values. It is not obvious how to establish criteria for judging the relative merits of novel value systems. Finally, the Indigenous Cosmic Cheese Multisource Infipienary entails a division of the cosmos into separate units. There is the moral question of how the division should be done, as well as the political question of what divisions would be accepted by the relevant parties.

#### **4. Answering The Big Question**

The ECB and its variants offer one compelling set of answers to the Big Question. The optimal answer will depend on details beyond the scope of this paper, including the selection of specific ethical



frameworks; the extent of political disagreement about The Big Question, including the risk of that disagreement escalating to catastrophic war; the viability of postponing the answering of The Big Question to give it more careful deliberation; the geometric structure of the cosmos as it relates to the placement of boundaries; the existence of extraterrestrial value sources; and the particulars of alternatives to the ECB and its variants. Such matters could be pursued in future research.

## Acknowledgements

Thanks to Anders Sandberg, Ben Goertzel, David Brin, and David Pearce for helpful input. Thanks to Joseph Gottlieb, Karim Jebari, Dakota Norris, and McKenna Fitzgerald for comments on an earlier version of this draft.

## References

- Arias-Arévalo, P., Martín-López, B., Gómez-Baggethun, E. 2017, Exploring intrinsic, instrumental, and relational values for sustainable management of social-ecological systems, *Ecology and Society*, 22, 4, 43. <https://doi.org/10.5751/ES-09812-220443>
- Baum, S.D. 2010, Universalist ethics in extraterrestrial encounter, *Acta Astronautica*, 66, 3-4, 617-623.
- Baum, S.D. 2015, Winter-safe deterrence: The risk of nuclear winter and its challenge to deterrence, *Contemporary Security Policy*, 36, 1, 123-148.
- Berry, P. M., Fabók, V., Blicharska, M., Bredin, Y. K., Llorente, M. G., Kovács, E., Geamana, N., Stanciu, A., Termansen, M., Jääskeläinen, T., Haslett, J. R., Harrison, P. A. 2018, Why conserve biodiversity? A multi-national exploration of stakeholders' views on the arguments for biodiversity conservation, *Biodiversity and Conservation*, 27, 7, 1741-1762. <https://doi.org/10.1007/s10531-016-1173-z>
- Bostrom, N. 2008, Letter from Utopia, *Studies in Ethics, Law, and Technology* 2, 1, 1-7.
- Bostrom, N. 2014 *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.
- Bruskotter, J. T., Nelson, M. P., Vucetich, J. A. 2015, Does nature possess intrinsic value? An empirical assessment of Americans' beliefs, *The Ohio State University*. DOI: 10.13140/RG.2.1.1867.3129
- Budd, M. 1995 *Values of art*. Penguin.
- Camarena, M.E. 2020, Predictions of science fiction that came true, in Froehlich, A. (Eds.), *Outer Space and Popular Culture*, Springer, pp.129-144.
- Cockell, C.S., Horneck, G. 2004, A planetary park system for Mars, *Space Policy* 20, 4, 291-295. <https://doi.org/10.1016/j.spacepol.2004.08.003>
- Cockell, C.S., Horneck, G. 2006, Planetary parks – Formulating a wilderness policy for planetary bodies, *Space Policy*, 22, 256-261. doi:10.1016/j.spacepol.2006.08.006
- Coeckelbergh, M., Gunkel, D. 2014, Facing animals: A relational, other-oriented approach to moral standing, *Journal of Agricultural and Environmental Ethics*, 27, 5, 715-733. <http://dx.doi.org/10.1007/s10806-013-9486-3>
- Curry, P. 2011 *Ecological Ethics. An Introduction* (2nd ed.), Polity Press.
- Danaher, J. 2021, Axiological futurism: The systematic study of the future of values, *Futures*, 132, 102780.
- Dasgupta, P. 2021, *The Economics of biodiversity: The Dasgupta review*. UK Government, HM Treasury. [https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment\\_data/file/962785/The\\_Economics\\_of\\_Biodiversity\\_The\\_Dasgupta\\_Review\\_Full\\_Report.pdf](https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/962785/The_Economics_of_Biodiversity_The_Dasgupta_Review_Full_Report.pdf)
- de Garis, H. 2005 *The Artilect War*, Etc Press.
- Drexler, K.E. 2013 *Radical Abundance: How a Revolution in Nanotechnology Will Change Civilization*, PublicAffairs.



- Dubey, R.S., Laguzzi, G. 2021, Equitable preference relations on infinite utility streams, *Journal of Mathematical Economics*, **97**, 102542.
- Egan, G. 1997 *Diaspora*, Millenium.
- ESA, 2022, How many stars are there in the universe? [https://www.esa.int/Science\\_Exploration/Space\\_Science/Herschel/How\\_many\\_stars\\_are\\_there\\_in\\_the\\_Universe](https://www.esa.int/Science_Exploration/Space_Science/Herschel/How_many_stars_are_there_in_the_Universe)
- European Union, 2010, Directive 2010/63/EU of the European Parliament and of the Council on the protection of animals used for scientific purposes, *Official Journal of the European Union*, **L**, **276**, 33–79. <https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2010:276:0033:0079:en:PDF>
- Feinstein, J.F., Morley, S., Yang, H. 2016, Abstract Swiss cheese space and classicalisation of Swiss cheeses, *Journal of Mathematical Analysis and Applications*, **438**, **1**, 119-141.
- Fitzgerald, M., Boddy, A., Baum, S.D. 2020, 2020 survey of artificial general intelligence projects for ethics, risk, and policy, Global Catastrophic Risk Institute Technical Report 20-1.
- Fogg, M.J. 2000, The ethical dimensions of space settlement, *Space Policy*, **16**, **3**, 205-211.
- Global Zero, 2022, Reaching zero. <https://www.globalzero.org/reaching-zero>.
- Goertzel, B. 2012, Should humanity build a global AI nanny to delay the singularity until it's better, *Journal of Consciousness Studies*, **19**, **1-2**, 96-11.
- Grinspoon, D. 2016 *Earth in human hands: Shaping our planet's future*, Hachette Book Group.
- Haidt, J., Rosenberg, E., Hom, H. 2003, Differentiating diversities: Moral diversity is not like other kinds, *Journal of Applied Social Psychology*, **33**, **1**, 1-36.
- Hanson, R. 2016 *The Age of Em: Work, Love, and Life When Robots Rule the Earth*, Oxford University Press.
- Holt, L.A. 2021, Why shouldn't we cut the human-biosphere umbilical cord?, *Futures*, **133**, 102821. <https://doi.org/10.1016/j.futures.2021.102821>
- Jamieson, D., Elliot, R. 2009, Progressive consequentialism, *Philosophical Perspectives*, **23**, 241-251.
- Johansson-Stenman, O. 2018, Animal welfare and social decisions: Is it time to take Bentham seriously?, *Ecological Economics*, **145**, 90-103. <https://doi.org/10.1016/j.ecolecon.2017.08.019>
- Joy, B. 2000, Why the future doesn't need us, *Wired*, **8**, **4**, 238-263.
- Kristensen, H.M., Kile, S.N. 2020, World nuclear forces, 2020. In *SIPRI Yearbook 2020*, Oxford University Press, pp. 325-393.
- Kurzweil, R. 2005 *The Singularity Is Near: When Humans Transcend Biology* Cover of The Singularity is Near, Viking.
- Loosemore, R.P.W. 2014, The maverick nanny with a dopamine drip: Debunking fallacies in the theory of AI motivation, *Association for the Advancement of Artificial Intelligence 2014 Spring Symposium*.
- Lupisella, M. 2016, Cosmological theories of value: Relationalism and connectedness as foundations for cosmic creativity, in J.S.J. Schwartz & T. Milligan (Eds.), *The ethics of space exploration*, Springer International Publishing Switzerland, pp. 75-91.
- MacAskill, W., Bykvist, K., Ord, T. 2020 *Moral Uncertainty*, Oxford University Press.
- Mautner, M.N. 2020 *Seeding the universe with life: Securing our cosmological future*. Galactic ecology, astroethics and directed panspermia, Legacy Books.
- McKay, K.P. 2009, Planetary ecosynthesis on Mars: Restoration ecology and environmental ethics, in Bertka, C.M. (Ed.), *Exploring the origin, extent, and future of life*, Cambridge University Press, pp. 245-260.
- McKay, K.P., Zubrin, R. 2002, Do indigenous Martian bacteria have precedence over human exploration?, in *On to Mars: Colonizing a New World*, Apogee Books Space Series, pp. 177-182.
- McShane, K. 2016, Is biodiversity intrinsically valuable? (And what might that mean?), in J. Garson,

- A. Plutynski, & S. Sarkar (Eds.), *The Routledge handbook of philosophy of biodiversity*, Routledge, pp. 155-167.
- Miller, J.D. 2012 *Singularity Rising: Surviving and Thriving in a Smarter, Richer, and More Dangerous World*, BenBella.
- Milligan, T. 2015 *Nobody owns the moon: The ethics of space exploitation*, MacFarland and Company.
- Moravec, H. 1998 *Robot: Mere Machine to Transcendent Mind*, Oxford University Press.
- Notte, H., Bidgood, S. 2022, What would Russia's break with the West mean for nuclear arms control? War on the Rocks, 14 February. <https://warontherocks.com/2022/02/what-would-russias-break-with-the-west-mean-for-nuclear-arms-control>
- O'Neill, J. 2001, Representing people, representing nature, representing the world, *Environment and Planning C: Government and Policy*, 19, 4, 483-500.
- Ord, T. 2020 *The Precipice: Existential Risk and the Future of Humanity*, Hachette.
- Owe, A. 2022, Greening the universe: The case for ecocentric space expansion, in Schwartz, J.S.J., Billings, L., Nesvold, E. (Eds.). *Reclaiming Space: Progressive and Multicultural Visions of Space Exploration*, Oxford University Press, in press.
- Owe, A., Baum, S.D. 2021, Moral consideration of nonhumans in the ethics of artificial intelligence, *AI & Ethics*, 1, 4, 517-528.
- Pascual, U., Gould, R., Chan, K.M.A. (Eds.) 2018, Sustainability challenges: Relational values, *Current Opinion in Environmental Sustainability*, 35, A1-A8, 1-132. <https://www.sciencedirect.com/journal/current-opinion-in-environmental-sustainability/vol/35>
- Persson, I., Savulescu, J. 2012 *Unfit for the future: The need for moral enhancement*, Oxford University Press.
- Rolston, H. III. 1986, The preservation of natural value in the solar system, in Hargrove, E. (Ed.), *Beyond spaceship Earth: Environmental ethics and the solar system*, Sierra Club Books, pp.140-182.
- Rolston, H. III. 1988 *Environmental ethics. Duties to and values in the natural world*, Temple University Press.
- Rolston, H. III. 2020 *A New Environmental Ethics. The Next Millennium for Life on Earth*, 2<sup>nd</sup> Ed., Routledge.
- Sadler-Smith, E. & Akstinaite, V. 2021, Human hubris, anthropogenic climate change, and an environmental ethic of humility, *Organization & Environment*, August. <https://doi.org/10.1177%2F10860266211039000>
- Schwartz, J.S.J. 2012, On the moral permissibility of terraforming, *Ethics & The Environment*, 18, 2.
- Schwartz, J.S.J. 2018, Where no planetary protection policy has gone before, *International Journal of Astrobiology*, July, 1-9. DOI: 10.1017/S1473550418000228
- Schwartz, J.S.J., Milligan, T. (Eds.) 2016 *The Ethics of Space Exploration*, Springer.
- Schwartz, J.S.J., Billings, L., Nesvold, E. (Eds.) *Reclaiming Space: Progressive and Multicultural Visions of Space Exploration*, Oxford University Press, in press.
- Scouras, J. 2019, Nuclear war as a global catastrophic risk, *Journal of Benefit-Cost Analysis*, 10, 2, 274-295.
- Shiller, D. 2017, In defense of artificial replacement, *Bioethics*, 31, 5, 393-399.
- Singer, P. 1977 *Animal Liberation*, Granada.
- Smith, K.C., Abney, K. (Eds.) 2019, Human colonization of other worlds, *Futures*, 110, 1-66.
- Sparrow, R. 1999, The ethics of terraforming, *Environmental Ethics*, 21, 3, 227-245.
- Stross, C. 2005 *Accelerando*, Orbit.
- Sunstein, C. R. 2002, The paralyzing principle, *Regulation*, 25, 32-37.
- Taylor, P. 1986 *Respect for nature*, Princeton University Press.
- Tonn, B. 1996, A design for future-oriented government, *Futures*, 28, 5, 413-431.

Tonn, B.E. 2002, Distant futures and the environment, *Futures*, 34, 2, 117–132.

United Nations. 1992, Convention on biological diversity, <https://www.cbd.int/doc/legal/cbd-en.pdf>

Vakoch, D.A. (Ed.) 2014 *Extraterrestrial altruism: Evolution and ethics in the cosmos*, Springer.

Wilkinson, H. 2022, Infinite aggregation and risk, *Australasian Journal of Philosophy*, in press, <https://doi.org/10.1080/00048402.2021.2013265>.