

# SPACE COLONIZATION AND EXISTENTIAL RISK

JOSEPH GOTTLIEB

Texas Tech University

joseph.gottlieb@ttu.edu

Forthcoming in the *Journal of The American Philosophical  
Association*

## Abstract

Ian Stoner (2017) has recently argued that we ought not colonize Mars because (i) doing so would flout our *pro tanto* obligation to not violate the Principle of Scientific Conservation, and (ii) there is no countervailing considerations that render our violation of the principle permissible. Here I remain agnostic on (i). Instead, my primary goal is to challenge (ii): there *are* countervailing considerations that render our violation of the principle permissible. As such, Stoner has failed to establish that we ought not colonize Mars. I close with some thoughts on what it would take to show that we *do* have an obligation to colonize Mars, and related issues concerning the relationship between the way we discount our preferences over time and projects with long time-horizons like space colonization.

**keywords:** Space Colonization, Mars, Existential Risk, The Principle of Scientific Conservation

## 1 The Argument from Conservation

Space colonization has been standardly the province of science fiction. This is still true today, but less so, especially with the advent of commercial space operators like SpaceX. If we were to colonize space—that is, set up a self-sustaining (Earth-independent), hopefully permanent, colony for human life in space—an obvious target is Mars. And eventually, we will be able to colonize Mars. But a further question is: ought we? Ian Stoner (2017) argues that the answer is *no*. His argument chiefly turns on what he calls The Principle of Scientific Conservation, or 'PSC' for short:

**PSC** Destructive or significantly invasive investigation of an object of scientific interest is only morally permissible when (a) significantly invasive investigation does not threaten the scientific or non-scientific values instantiated in that object and (b) no adequate alternatives to significantly invasive investigation are available. (2017: 342)

At its most general, Stoner's argument goes like this:

**P1** We have a *pro tanto* moral obligation not to violate the PSC.

**P2** Colonizing Mars would violate the PSC.

**P3** There is no countervailing and weightier moral considerations that would override our obligation not to violate the PSC in colonizing Mars.

**P4** If there is no countervailing and weightier moral considerations that would override our obligation not to violate the PSC in colonizing Mars, then we ought not colonize Mars.

∴ We ought not colonize Mars.

I contend that Stoner’s argument is unsound. I’ll grant P1, P2, and P4. Instead, I target P3. So my thesis is this: there *is* a countervailing and weightier moral consideration that would override our obligation not to violate the PSC in colonizing Mars. This consideration is, of course, one point *in favor* of colonizing Mars. So by way of closing, I’ll explore, in a programmatic way, how one might argue that we do in fact have an obligation to colonize Mars, and the prospects and problems for such an argument.

Before we press on, three dialectical clarifications are in order. First: Stoner also musters a related argument for his conclusion. This is based on (what he calls) *the tread lightly principle*: “[w]hen we visit wilderness, we have a moral obligation to conduct ourselves in such a way that the impact of our presence is indistinguishable, after a suitable period of time, from counterfactual worlds in which we didnt visit that wilderness” (ibid: 348). Stoner’s argument from this principle is not interestingly different from his argument based on the PSC—both employ similar pro tanto principles—so my case against the PSC-argument, if successful, will cut against this other

argument as well.

Second: Stoner is right when he says that just because compelling reasons to contaminate Mars *could* exist does not mean we have a challenge to the pro tanto constraints asserted by the PSC. We need, as Stoner adds, “to identify an actual value that outweighs it” (ibid: 337, fn. 4). He considers five, but I will argue that the full force of one of them—viz. colonizing Mars to help ensure survival of the human species—remains under-appreciated by Stoner’s own analysis. So what will emerge is not simply a disagreement about what is more valuable on the whole—humans interests and our long-term survival versus the conservation and preservation of Mars’ natural state—but whether, and to what degree, colonization of Mars can help us survive long-term.<sup>1</sup>

Third: while Stoner is not explicit on this point, I assume that he takes the PSC to hold for perpetuity. This means that PSC remains in force even when we have exhausted all scientific research of the object in question. In this way, for Stoner, Mars will end up like the Great Pyramids. Having studied the pyramids for over a hundred years, it’s not clear how much is left to learn about them. But that does not mean that we can now turn them into hotels. This point matters, since one might wonder whether we can preserve Mars for scientific study for some period of time, and then once that time has elapsed, pursue colonization without being beholden to the PSC. Of course,

---

<sup>1</sup>Stoner focuses on Mars, but I assume that his argumentative strategy can be applied to any planet or celestial body that is of scientific interest but also a target for colonization.

one *could* go this route, but if this is all Stoner himself intends, it is hard to see how his thesis has any bite. For then what Stoner grants is permissible is what we were planning on doing all along. We are not yet technologically ready to colonize Mars. (We are only approaching being technologically ready for humans to *visit* Mars.) Consequently, on the present proposal, we would start colonizing *exactly when we would have anyway*—in twenty-five years, say—and *right up until we are ready*, we can tread lightly in our scientific investigations, building on our knowledge of the past 50 years and learning all we can about Mars without violating the PSC.

Further evidence that Stoner believes we should never colonize Mars comes from his use of the tread lightly principle. Using this principle to argue that we ought not colonize Mars only makes sense if Stoner intends it to show that we should never colonize Mars. After all, if such a principle were to be violated by our presence on Mars at all, it would be violated whenever our presence began. So for both of these reasons, I'll assume that Stoner's official thesis is that we should *never* colonize Mars, and that's the thesis I'll chiefly address. However, if Stoner's suggestion is instead that we can colonize Mars only once all possible scientific investigation has been completed, even if we are able to colonize Mars well before such investigations are completed, I'm inclined to think this is wrong too. Given certain motivations, it is permissible to colonize Mars as soon as possible.<sup>2</sup>

---

<sup>2</sup>Many thanks to an anonymous referee for the query that prompted this clarification.

## 2 Mars and Existential Risk

Assume that we have at least a *prima facie* obligation to ensure the human race's long-term survival. I'll assume this without argument.<sup>3</sup> This claim, and the arguments produced by consideration of it, are part-and-parcel of Stoner's preferred *practical-controversy-centered* mode, which relies on beliefs about specific cases, or pre-theoretic principles, to derive (perhaps surprising) conclusions in applied ethics (ibid: 335-336). Let's also assume, at least for now, that if we colonize Mars, we will optimize the chance of the human race surviving long-term. This value—viz. the value of the human race surviving long-term—is a countervailing and weightier moral consideration that overrides our obligation not to violate the PSC in colonizing Mars. By way of motivation, consider the following case:

**Cure:** Suppose that a new pyramid is found by Giza in Egypt. It is, naturally enough, an object of immense scientific and archaeological curiosity. But suppose that, amongst its many valuable features, there lies deep within its structure, a vial of blue elixir locked in a sarcophagus. This elixir for whatever reason provides our best shot to cure cancer, but it will hinge on billions of dollars of future investment. Moreover, to obtain the vial, the pyramid and all its other contents *must* be destroyed.

Let me make five preliminary points about CURE. First: we have a case where the PSC is germane. The pyramid is an object of scientific interest, and its possible destruction is in question. Second: we have a case where the *pro tanto*

---

<sup>3</sup>Stoner never denies we have an obligation to preserve humanity, and at one point seems to agree that we have such an obligation (ibid: 339). I will briefly return to our (purported) obligation to extend human life in the conclusion.

obligation that is the PSC can be compared against a potentially weighty moral consideration, viz. our chance to cure cancer, given the paucity of other options. Third: we have a case where the intended end—actually curing cancer and saving present and future lives—is not guaranteed by taking the destructive action, but just made significantly more likely. Fourth: destroying the pyramid would ‘only’ save *a lot* of lives. Cancer is not the sort of thing that is an existential threat to the human race. Nonetheless curing it would be of monumental moral importance. Finally, fifth: describing the elixir as the ‘best shot’ to cure cancer is deliberate. This leaves open that there might be other cures for cancer. Likewise, there might be other ways to promote species-survival.

I take it that the following is obvious: in CURE, we have identified a countervailing and weightier moral consideration that would override our obligation not to violate the PSC with respect to the pyramid, one that permits us to obtain the elixir as soon as we are able. The baseline point I want to advance is this. If in CURE we have identified a value that overrides our obligation not to flout the PSC with respect to the pyramid, then we have identified a value that overrides our obligation not to flout the PSC in colonizing Mars: the value in optimizing the chance that the human race survives long-term. In what follows I will defend this claim, looking at objections by examining variations on CURE.

The first thing to note is that, in CURE, we are dealing with something that is arguably of *lesser* value. In CURE, we destroy the pyramid because

it will save lives from cancer, and our destroying the pyramid is the only way to do so. But in colonizing Mars, the idea is that we are not just trying to save lives, but save the human race. Cancer is, to put it mildly, serious. But it is not an existential threat to humanity. So, if the verdict in CURE is that we have identified a consideration that overrides the PSC, then surely in treating Mars as a hedge against existential threats we will have identified a consideration that overrides the PSC.

It's tempting to think that in CURE we are talking about saving current lives, whereas with Mars we are concerned with ensuring that future humans exist, and obligations concerning the former aren't cleanly transferable to obligations concerning the latter. This isn't a real difference though; we also cure cancer to prevent future people from getting it, and, depending on how things shake out, Mars could be a home for current people too. We might imagine a case where a previously undetected asteroid threatens the planet, with only ten years to prepare. On such a scenario, colonizing Mars would be about saving current lives as much as it would be about having a home for future generations.

It is also important to note that Mars is a relatively attractive target for a second home. To see this, it helps to compare it with other potential colonization targets like the moon. A day on Mars is roughly equivalent to ours, being just over 24 hours long. By comparison, a day on the moon is equivalent to roughly 28 Earth days. This alone would make the adjustment there a challenge. Because of this relatively short day, the day-night temper-

ature swings are less dramatic. By comparison, the moon has on average a 572 degrees Fahrenheit day-night swing. And a summer day on Mars can get up to 70 degrees Fahrenheit at the equator. This makes living on the surface, while no means a picnic, certainly easier. Mars has an atmosphere, albeit a thin one. Still, with an atmosphere, we can get some radiation protection, which makes up somewhat for Mars only have magnetic fields at its southern hemisphere. An atmosphere also means we can pressurize structures using air from outside. Again by comparison, the moon has an extremely weak magnetic field and virtually no atmosphere. And perhaps most importantly, Mars has abundant water, not just as vapor in its (thin) atmosphere, but in abundance at the poles and perhaps elsewhere in the form of ground ice.<sup>4</sup> I belabor these points because if Mars is going to serve as a backup planet—a place to guard against existential threats—then clearly it needs to be a place on which *we can exist*. And while it's far from perfect, it is better on the habitability-scale relative to the other options in the solar system.

Of course, one might grant that Mars would provide *a* hedge, but insist that there are better options that would not require colonization, and so would present less of a risk for objects of scientific value. Stoner makes something like this claim when he says:

The range of species-level threats addressed by a Mars colony is

---

<sup>4</sup>For evidence that recurring slope lineae (RSL) are due to extant water flows on Mars, see Ojha *et al* (2015). Unsurprisingly, this explanation is not uncontroversial; for an alternative explanation of RSL that does not appeal to extant water flows, but granular-flow, see Dundas *et al* (2017).

relatively narrow (York 2002). A Mars colony would not insure against large-scale threats to the solar system, such as nearby supernovae, invading extraterrestrials, or an early expansion of the sun. Nor would it insure against threats we pose to ourselves, such as war and environmental destruction. We carry these threats to ourselves everywhere we go, and we would carry them with us to Mars.

A Mars colony would only insure against externally imposed large-scale environmental threats specific to Earth. A colony on Mars would be unmolested by, for example, a Chicxulub-scale asteroid or comet strike on Earth. But is a Mars colony the best way to hedge against this risk?

...[N]ote that while it's relatively easy to imagine an asteroid or comet impact knocking civilization back a few hundred years, its genuinely difficult to imagine a *sapiens-extincting* impact. Contra Niven, Chicxulub didn't kill the dinosaurs because they lacked a space program; it killed them because they lacked blankets (2017: 339, emphasis in original).<sup>5</sup>

Stoner is technically right that a Mars colony would not insure against all species-level threats. And he is also technically right that the dinosaurs did

---

<sup>5</sup>Stoner is referring to a quip, attributed to Arthur C. Clark by fellow science fiction author Larry Niven, that the dinosaurs became extinct because they didn't have a space program.

not need a space program to survive Chicxulub. Yet all of this misses the point. In the case of nuclear war, it is true that such a war could happen on Mars, just like it could happen anywhere. The point though is that if a nuclear war broke out on Earth, it *need not* break out on Mars too. (Of course, for nuclear war to be a genuine existential threat, it would have to be global, and induce a ‘nuclear winter’ scenario.) In the case of dinosaurs, the comparison is not between blankets and a space program that could harm Mars; if that were the choice, it would be an easy one. Rather, the choice is between (something like) a massive series of bunkers, filled with provisions to last thousands of years and a space program. In that case, the choice is less easy.

Stoner also seems to be assuming that colonizing Mars is the end-goal, and that only if Mars *itself* is the best option to save us against existential threats would colonizing Mars be warranted. But this is false. Even if *Mars* is not the best hedge, it might enable us to get to places that *are*. This is the basic idea behind ‘bootstrapping’, and it is a core element in most visions of an interplanetary future; Mars—or some other celestial body—is a pit stop that makes getting to another target easier. There are different ways this can be done. Perhaps we’ll realize that Mars cannot be fully terraformed, or that for some other reason it is not the right place for humanity to hang its collective hat on. This does not mean that colonizing Mars to save humanity no longer presents us with a consideration that would override our obligation not to flout the PSC. For it might be the case that we *must* colonize Mars *so that*

we can find a true, better second home—a home that *would* hedge against a wider range of existential risks. For example, maybe Mars is crucial because, with its abundant water ice, it can become a propellant depot via *in situ* resource development. The water ice would then be used to produce rocket fuel.<sup>6</sup> Mars then becomes not just a colony, but a way-station. Using Mars in this way may not require full-scale colonization with millions of people, but it will likely require *some* permanent presence, and in this way runs into the same PSC-style worries Stoner invokes. On the other hand, those worries might be lessened, since with fewer people, there will presumably be less risk of wide-spread contamination and other issues.

This does, of course, beckon a question of efficiency: if the idea is to use Mars largely as a propellant depot, might asteroids or shallow moons, which have gravity wells that are naturally more amenable to space travel, be more viable? (There is an inverse relationship between the depth of an object’s gravity well and the energy acquired to access and escape that object.) It depends, since accessibility is not the only relevant factor. Phobos and Deimos (the moons of Mars) have shallow gravity wells, and thus will fair better than Mars itself in terms of accessibility, but their overall viability as a propellant depot will also depend on their native resources, like water ice. The same comparison holds for other targets, such as closer near-Earth asteroids. Of course, such a move wouldn’t disarm any possible appeal to the PSC, since

---

<sup>6</sup>It is quite common for something like this to be part of plans for intra-solar system development. For just two recent examples, see Mueller *et al* (2012), and Musk (2017).

these are objects of scientific interest in their own right, where a human presence might have a negative impact. Also, asteroids and the relevant moons will have significantly lower gravity and atmospheric protection from radiation than Mars, both of which would be harmful to human inhabitants over the long-term.

This point aside, I think the intuitions here are fairly clear. To see this, consider a variation of CURE:

**Ingredient:** Suppose that a new pyramid is found by Giza in Egypt. It is, naturally enough, an object of immense scientific and archaeological curiosity. But suppose that, amongst its many valuable features, there lies deep within its structure, a sarcophagus that contains an essential ingredient to what we have good reason to believe is the cure for cancer. This ingredient will not on its own cure cancer, but it is an essential component for finding what will eventually be the cure. To develop the cure and find the other ingredient will require billions of dollars of future investment, and much challenging work. Moreover, to obtain this essential ingredient, the pyramid and all its other contents *must* be destroyed.

Unlike in CURE, in INGREDIENT, we would not destroy the pyramid to get the cure for cancer. We would instead do so to get something that is very important to such a cure. Nonetheless, I venture most would ultimately decide that we have found something that overrides the PSC, and thus makes it permissible for us to destroy the pyramid. I think INGREDIENT parallels the bootstrapping model above in the relevant respects. We are not destroying the pyramid because the pyramid contains the cure to cancer. Likewise, on the current story, we are not colonizing Mars because colonizing Mars is itself

the final end-goal in our attempt to protect against existential risk. Instead, we colonize Mars because it is part—a key part—of a broader series of goals, or a stepping-stone to some further (potentially final) destination.

Now granted, although part of the point of using Mars as a bootstrap is to make more distant colonization targets within our grasp, we should be careful not to put too much argumentative weight on this point in the present context. Distance can make colonization more challenging. And as we will discuss in greater detail shortly, the likelihood of success is relevant when adjudicating whether the pro tanto obligation not to violate the PSC can be overridden. Nonetheless, the basic point here still holds: if one thinks that in *INGREDIENT* we have found a consideration that overrides the PSC, then one should think that we have a consideration that overrides the PSC when it comes to colonizing Mars, even if Mars is just used for bootstrapping purposes.

Stoner (2017: 340) also discusses investing in detection and re-direct capabilities for near-Earth objects, along with seed banks and more robust Earth-based knowledge preservations methods, as alternatives to colonizing Mars. He suggests that it is only because such options are omitted that cost-benefit analysis favor Mars as being how we should discharge our obligation to protect the human race (*ibid*). The problem is that for Stoner's objection to stick, it would have to be the case that such solutions would be more instrumentally useful (with respect to survival) than a colony on Mars. Stoner gives us no reason to think that they would be, and it is hard to see how

they could be. Near-Earth object detection capabilities are a good thing, but the scope of threats they can handle is limited, viz. near-Earth objects. This is true of terrestrial based solutions more generally, as the efficacy of things like well-stocked ‘refuges’ has been questioned (Beckstead 2015). By comparison, simply in virtue of its spatial distance to Earth, a colony on Mars can hedge against more, like global pandemics. Stoner might concede that asteroid-detection and other terrestrial measures are not as instrumentally useful, but given the weight of the PSC, it would be better on balance to opt for a solution that while less effective with respect to safeguarding the species, does not threaten objects of scientific interest. While I doubt Stoner would actually take this line in the end, if he did all it would reveal is that Stoner and I disagree about the strength of the intuition favoring the PSC. He thinks it’s quite strong. I don’t. The burden though is clearly on Stoner to show that we ought not colonize Mars, so a lack of consensus is a more serious problem for him than it is for me.

At the same time, since the costs of settling Mars are significant, any lack of consensus surrounding the chances that a Mars colony actually succeeds matters too. Notice that the point here isn’t how Mars fares relative to other colonization targets in the the Solar System. By this measure, Mars likely is our best bet. We might imagine as an alternative a large space station parked at a Lagrange point between here and the moon. This option would not run into the PSC, but it would be far more spatially-limited than Mars, and it would arguably require continual resources from Earth, which would

defeat the purpose. Another option that has been proposed is colonizing Venus via aerostat habitats, where at 30 miles above the surface, we can enjoy one atmosphere of pressure, only slightly lower gravity than Earth, and a temperature climate.<sup>7</sup> This is all certainly a plus, but there is a big downside; Venus' planetary resources, unlike Mars', are trapped in a hellish, lava-filled pressure-cooker. This becomes a problem in the long-term, since anything we needed might have to be either brought with, or constructed at 30 miles up. But even if Mars is indeed our best colonization target, and thus our best hedge against species-level threats, it still might not be a good hedge. And that's the worry. It might simply be the best of a bad lot. Perhaps Mars is *so bad* of a hedge that, *even if* it is the best hedge, violating the PSC by colonizing it would not be worth it.

This is an important concern. It is also somewhat difficult to assess, turning as it does on both an unsettled empirical question (just how much more likely are we to survive if we colonize), and an unsettled moral question (at just what increase in probability of success is it permissible to violate the PSC). But headway is possible. Consider the following further case:

**Risky Cure:** Suppose that a new pyramid is found by Giza in Egypt. It is, naturally enough, an object of immense scientific and archaeological curiosity. But suppose that, amongst its many valuable features, there lies deep within its structure, a vial of blue elixir locked in a sarcophagus. This elixir for whatever reason provides our best shot to cure cancer, but it will hinge on billions of dollars of future investment. Obtaining the vial will require

---

<sup>7</sup>See, e.g. NASA's High Altitude Venus Operational Concept or 'HAVOC' project: <https://sacd.larc.nasa.gov/smab/havoc/>.

many years of work and much financial investment, along with the destruction of the pyramid and all of its other contents. On top of this, the project is highly risky given its costs: there is only a 30% chance of success. That is, there is a 70% chance that we will destroy the pyramid and *not* get the elixir.

The 30% chance of success in RISKY CURE is meant to track a conservative estimate of the chance of success of a Mars colony. So if the correct verdict in RISKY CURE is that we should go ahead and attempt to get the elixir—and I submit that most would say it is—then we ought to attempt to colonize Mars, despite the (relatively) low chance of success.

How should we think about ‘success’ and ‘failure’ when considering space colonization? Trivially, one factor relevant to success is the continuity of the colony itself—that it does what it was primarily designed to do, viz. that it serve as a conduit for the continued existence of the human race. In this respect, it could be that any space settlement without an indigenous extant breathable atmosphere will have a low chance of success and a high chance of failure, simply because to breath we will be reliant on technology that is liable to malfunction.<sup>8</sup> Call this kind of success factor a *foundational* success factor. It is foundational because it concerns the very existence of the colony.

Success and failure does not end here, though. Another factor is the well-being of the colonists themselves. The well-being of the colonists includes not just their physical health, but their personal liberties. This latter issue crops

---

<sup>8</sup>Absent terraforming, this is a reality that anyone considering colonizing Mars must face, but it is also not a reality that is limited to colonizing Mars. We don’t (alas) know of any true ‘Earth-twins.’ So this is a problem for colonizing *anything*.

up do the natural constraints put on settlers by the limits of a space colony—limits that could encourage authoritarianism (Cockell 2016a; Milligan 2016), discourage dissent and worker’s rights (Cockell 2016b; Schwartz 2016), and discourage reproductive rights (Szocik *et al* 2018).<sup>9</sup> These kinds of success factors are *non-foundational* success factors. To call them ‘non-foundational’ is not to say that they are not important, but only to say that in the first instance they concern the *conditions* of the colony, not the colony’s existence. It is true that some non-foundational factors *could* impact the existence of the colony under extreme circumstances. Authoritarianism could cause the colonists to rebel and damage the colonies’ critical support systems. The point is simply that a colony could be successful by one core measure (it being a conduit for the human race’s continued existence beyond Earth), and unsuccessful by another (it being a place to which few would want to vacation).

We need to be clear-eyed about what it would take for a colony to succeed given these factors. If we are not, the 30% estimate in RISKY CURE will look too high. Yet if we are, the 30% estimate might look reasonable, and maybe too low. To see this, consider another foundational success factor: reproduction. We don’t yet know whether we can reliably reproduce in reduced gravity (Szocik *et al* 2018). But this is something we can game out beforehand. It has been estimated that minimum viable human population for an

---

<sup>9</sup>These worries would be exacerbated in cases where colonization targets are so distant that it requires generation or ‘worldships’ (Schwartz 2018). Mars is far, but not *that* far.

extra-terrestrial colony should include about 5000 individuals, assuming we want that colony to survive for the long-term (Impey 2015). For a Martian colony to be a hedge against existential risk, it needs to be self-sustaining, and thus capable of internally replenishing its own population. This won't be achieved right away. In the short-term, immigration from Earth will be the primary means of establishing a sustainable population. But if turns out that humans cannot reproduce in low gravity (Mars, recall, it at .38 G), we can design 'reproduction centers' with a full G of spin gravity, perhaps in Martian orbit. (I presume that if we are technologically capable of building a Martian colony, we are technologically capable of building such an orbital center.) Of course, this is less than ideal in numerous ways, but it is hard to see how that alone should weaken the strength of our obligation to colonize Mars, especially when the primary reason for colonizing Mars is to preserve the human species. Much of colonizing Mars will be thrilling and enjoyable, but to assume it will be a pleasure cruise through-and-through is to stack the deck against a Martian colony ever being counted as a success.

With respect to non-foundational factors, we should reject the idea that the success of a space colony requires freedoms analogous to those found on Earth. Liberty and freedom may look different beyond Earth, and this is just a fact that any space settler will have to recognize. As Tony Milligan (2016: 12-13) notes, we essentially face a dilemma: if the colony is too illiberal then establishing it might not be a worthwhile, but if it is too liberal in its freedoms it is unlikely to survive. So while violent forms of disobedi-

ence might not be advisable—lest we risk catastrophic consequences from damaged infrastructure—there are acceptable forms of civil disobedience in space, and engineering methods that can be used to mitigate against violent disobedience, in case it happens (e.g. Cockell 2016b). Indeed, as Schwartz (2016) points out, even low-level employees in the air production industry can justifiably strike, given that the duty to provide air (and other essential services) is born by the wider community.<sup>10</sup>

A final objection one might press is that Mars is an object of greater scientific significance than the new pyramid, and for this reason, the analogy does not hold, as our obligation not to violate the PSC is now more pressing. Perhaps Mars is of greater interest. I doubt it matters, though. For this is not the only difference between the cases. Another difference, recall, is this: in CURE, we are saving a lot of lives. In colonizing Mars, whether as a bootstrap or not, we are trying to save the human race. That’s a difference that makes a difference.<sup>11</sup> Consider another variant on CURE:

**Martian Cure:** Suppose that deep within Olympus Mons—the 72,000 foot volcano on the surface of Mars—is a vial of blue elixir locked in bejeweled box, left by an ancient alien civilization from a distant galaxy. The civilization inscribed some details of its history on the box. This elixir provides our best shot to cure cancer, though it will hinge on billions of dollars of future investment. However, there is a further catch: the bejeweled box—our only

---

<sup>10</sup>Schwartz’s discussion concerns Lunar labor relations, but there is no reason this could not be applied to Martian labor relations.

<sup>11</sup>Relatedly, Bostrom (2003) argues from utilitarian principles that we ought not delay colonization because of delay has a massive opportunity costs, viz. the loss of billions of potential lives worth living not being realized.

other piece of evidence of this civilization—*must* be destroyed to obtain the vial.

As in CURE, in MARTIAN CURE we have our best chance to cure cancer. Yet now we are contemplating destroying an object of arguably greater scientific interest than Mars: our only other piece of evidence of an ancient alien civilization. That this is of greater scientific significance is underscored by the fact that whereas we have others planets to study, without the box we only have one other piece of evidence for aliens, viz. the vial. Even so, in MARTIAN CURE it is arguable that we have a consideration that would override the PSC with respect to the box. If that's right, then that the pyramid is of lesser scientific significance than Mars cuts no ice.<sup>12</sup>

### 3 When to Colonize

Stoner's argument fails. There is a consideration that warrants us overriding the PSC when it comes to colonizing Mars: the opportunity it provides as a hedge against existential risk. On its own, this gets us close to the claim that we *do* have an obligation to colonize Mars. I'll close by exploring in a programmatic fashion what would take us all the way to such a claim.

A sketch of a reasonably compelling argument for our having an obligation to colonize Mars proceeds in two stages. The first stage looks like this:

---

<sup>12</sup>For two competing accounts of scientific significance, see Kitcher (2001) and Brown (2010).

**M1** We have an obligation to ensure the long-term survival of the human race.

**M2** If we have an obligation to ensure the long-term survival of the human race, then we have an obligation to perform the actions and undertake the projects that, amongst all the feasible actions and projects, will make the long-term survival of the human race most likely.

∴ We have an obligation to perform the actions and undertake the projects that, amongst all the feasible actions and projects, will make the long-term survival of the human race most likely.

Given the conclusion of the first stage, here's the second:

**M3** Colonizing Mars is amongst the feasible projects that will make the long-term survival of the human race most likely.

∴ We have an obligation to colonize Mars.<sup>13</sup>

The 'feasible' clause in M2 and M3 allows that there might be targets for colonization that are better than Mars insofar as our long-term survival goes. For instance, suppose that Proxima Centauri *b*—an exoplanet in the habitable zone orbiting Proxima Centauri, the star closest to the Sun at 4.22 light years away—has a rocky surface, liquid water, a breathable atmosphere, and a temperate climate. The existence of Proxima Centauri *b* would not show

---

<sup>13</sup>For a similar argument that we ought to colonize space (though not necessarily Mars), see Schwartz (2011).

that M3 is false, because colonizing Proxima Centauri *b* is currently not feasible, and may never be. (If it ever is in the far future, this might be due precisely to our having colonized Mars first.) What counts as feasible is largely an empirical matter, but there are also moral and pragmatic constraints too. Suppose one thought that to ensure humanities' long-term survival we should refrain from expanding our population beyond Earth (as would be the case if we colonized Mars), and instead focus on contracting our population. The problem here is, if not moral, certainly practical; Malthusian-style solutions are not something many will be running to get behind.

Now I have tried to show that Stoner's case against colonizing Mars fails by motivating M3, and arguing that *if* M1 is true, then the value of the human race surviving long-term is a countervailing and weightier moral consideration that overrides our obligation not to violate the PSC. To go further though and establish that we do in fact have an obligation to colonize Mars, we will need to do at least four things. First, we will need to provide a robust defense of M1. This is a non-trivial task. Saying exactly what's wrong with human extinction takes work (Finneron-Burns 2017), as does saying what moral reasons we have, if any, for promoting the long-term survival of humanity (Frick 2017). In fact, if an obligation to extend humanity is simply an obligation to future generations, then the non-identity problem might render M1 incoherent (Parfit 1984).<sup>14</sup>

---

<sup>14</sup>Although the applicability of the non-identity problem to space colonization is not entirely obvious either. Schwartz (2019) suggests that it may not be applicable, since the relevant obligation here is only that future persons continue to come into existence, not

Second, while M2 seems fairly innocuous, M3 will also require a more detailed defense—certainly more detailed than has been provided here. For instance, to fully defend M3, it would be ideal to have an empirically grounded full assessment of both the costs and benefits of colonizing Mars with regards to protecting against existential threats, weighed against the costs and benefits of the other suite of activities we will want to pursue, like reducing our stockpile of nuclear weapons and implementing broad climate change mitigation (and adaption) strategies.

Third: the present argument elides the difference between prima facie obligations and all-things-considered obligations (Ross 1930). Prima facie obligations can be overridden; it is an obligation such that, absent conflicting obligations, would be an all-things considered obligation. So, to make the case that we have an all-things-considered obligation to colonize Mars, we would have to show that M1 still holds after balancing it with any other conflicting prima facie obligations. I have attempted that with the PSC, but there might be other considerations.

Finally, and related to our first point, in saying that ‘we’ have an obligation to colonize Mars, the argument’s conclusion technically leaves open exactly what this obligation entails for present day society, and relatedly,

---

that future persons have certain levels of well-being. For my part, it is hard to make sense of having an obligation to ensure that future persons continue to come into existence absent having an obligation to ensure that this future existence is one worth having. In any event, the non-identity problem is not the only issue that renders obligations to future generations questionable; for an independent argument that no such obligations exist, see, e.g., Golding (1972).

just how strongly this obligation weighs on us today. In CURE we are dealing with something (cancer) that is an actual, present, eminent threat. So we are inclined to say that it is permissible, if not obligatory to destroy the pyramid for the cure now, or as soon as we are able. By contrast, while it is possible for us to discover a planet-killing asteroid tomorrow, and it is possible for a global nuclear war to break out overnight, these threats are not actually present in this way. So even granting an obligation to colonize Mars *simpliciter*—something Stoner denies—we might still wonder whether this obligation requires us to act as soon as we are technologically able, especially if that time comes before we have exhausted all possible non-invasive scientific investigations of Mars.

This is a fair point, yet protecting against possible existential threats is not exactly a can to be kicked down the road. The reason we would colonize Mars as soon as we are technologically able is precisely to avoid doing it when our collective backs are against the wall. We would not want to *start* our colonization efforts only when we *do* know of a planet-killing asteroid, and thereby risk having not nearly enough time to complete our efforts. You don't buy health insurance only when death is at your door. The same thinking goes for the species.

That said, this kind of forward-thinking does not exactly track how we in fact tend to behave. Consider: colonizing Mars is a project with a long time-horizon. Regardless of how we do it, a full-fledged (i.e. self-sufficient)

colony on Mars will take a while, likely decades from start to finish.<sup>15</sup> Humans, though, are not good at waiting, and we tend to care more about what happens to us, now, than what happens to future generations. This tendency is at play in both hedonic and non-hedonic valuing. When the value of an activity is hedonic, we exhibit near bias: we prefer painful activities to be distant and pleasurable activities to be near. When the value of an activity is non-hedonic, we do so in a way that biases the perceived social distance between you and who else might benefit from those activities.<sup>16</sup> Space colonization will have both hedonic and non-hedonic value; it will be fun and exciting (though also dangerous and perhaps terrifying), but it will also be valuable given its importance to future generations. Yet now the time-scales involved present a further problem. We assign a positive discount rate to future well-being: the value of someone's well-being in (say) 50 years is only worth a fraction of the well-being of a present person. So the question we are inevitably faced with is this: why would I start colonizing space now when the well-being of those who would be *actual* space colonists, living on Mars, is worth so little to me relative to my present peers? The thinking goes: 'sure, maybe we can visit Mars, but I'd rather spend the vast amount of money

---

<sup>15</sup>The exact timeframe depends on a number of factors, including the minimum viable human population. One potentially limiting factor is that minimum-energy launch windows for a Mars transit occur every two years. In addition, if colonization includes terraforming, the time-scales involved typically run from anything between 1000 and 100,000 years, depending on the feasibility of the proposed method. For a study that supports the longer-end estimate, see McKay & Marinova (2001).

<sup>16</sup>For a recent discussion of time-biases, see Greene & Sullivan (2015). For a discussion of social-biases, and how they interact with discounting the welfare of future generations, see Greene (ms).

it would take to colonize Mars on present concerns and present people.’ The idea here is akin to the one we have when we prefer a quick weekend getaway to a local winery in two weeks to a 3-week trip to Bali in ten years.

One complication with this is that, even if we assume that the weekend winery getaway and the trip to Bali are equally likely to occur, this assumption may not hold for colonizing Mars. Perhaps institutional support will erode over time, as when one presidential administration decides to no longer continue the commitments of a past one. (Witness, for example, Nixon’s cancellation of the Apollo program.) In fact, according to Jacobs & Matthews (2012), when it comes to present-directed policy bias—i.e. how citizens weigh short-term outcomes more heavily than long-term outcomes in policy choice—the primary factor is uncertainty surrounding the processes of policy causation and political commitments over the long-run.

I don’t know whether such reasoning, and the temporal and social biases that underpin it, is rational. In one respect it doesn’t matter. These biases are a pervasive feature of human psychology and as such present an inherent and forceful impediment to getting space colonization off the ground. We will always tend towards buck-passing, even if we know that we ought not.<sup>17</sup> In another sense though, it does matter. For if this reasoning *is* rational, then even if there is an obligation to colonize Mars, this obligation isn’t necessarily one that must be borne by us in present day. How to approach resolving this matter and related issues is, however, a problem for another day.<sup>18</sup>

---

<sup>17</sup>Compare this with our inaction regarding mitigating climate change.

<sup>18</sup>For helpful comments and discussion, many thanks to Alexis Elder, Bob Fischer,

## References

- [1] Beckstead, Nick. (2015). ‘How much could refuges help us recover from a global catastrophe?’ *Futures*. 72: 36-44.
- [2] Bostrom, Nick. (2003). ‘Astronomical Waste: The Opportunity Cost of Delayed Technological Development.’ *Utilitas*.15 (3): 308 314.
- [3] Cockell, Charles. Ed. (2016a). *Dissent, Revolution and Liberty Beyond Earth*. Springer.
- [4] Cockell, Charles. (2016b). ‘Disobedience in Outer Space.’ In Charles Cockell (ed.), *Dissent, Revolution and Liberty Beyond Earth*. Springer.
- [5] Dundas, Colin M., *et al.* (2017). ‘Granular flows at recurring slope lineae on Mars indicate a limited role for liquid water.’ *Nature Geoscience*. 10 (12): 903-907.
- [6] Finneron-Burns, E.M. (2017). ‘What’s wrong with human extinction?’ *Canadian Journal of Philosophy* 47 (2-3): 327-343.
- [7] Frick, Johann. (2017). ‘On the survival of humanity.’ *Canadian Journal of Philosophy* 47 (2-3): 344-367.
- [8] Golding, M.P. (1972). ‘Obligations to Future Generations.’ *The Monist* 56 (1):85-99.
- [9] Greene, Preston & Meghan Sullivan. (2015). ‘Against Time Bias.’ *Ethics* 125 (4): 947-970.
- [10] Greene, Preston. (No Date). ‘Time Bias, Social Bias, and Future Generations.’ (ms.)
- [11] Jacobs, A., & Matthews, J. (2012). ‘Why Do Citizens Discount the Future? Public Opinion and the Timing of Policy Consequences.’ *British Journal of Political Science*. 42(4), 903-935.
- [12] Impey, Chris. (2015). *Beyond: Our future in space*. W.W Norton & Company.

---

Jessica Gottlieb, Judith Gottlieb, Saja Parvzian, Jeremy Schwartz, Joel Velasco, and three anonymous referees for *J-APA*.

- [13] Kitcher, Philip. (2001). *Science, Truth, and Democracy*. Oxford: Oxford University Press.
- [14] McKay, Christopher, & Marinova, Margarita. (2001). The Physics, Biology, and Environmental Ethics of Making Mars Habitable. *Astrobiology* 1 (1): 89-109.
- [15] Milligan, Tony. (2016). ‘Constrained Dissent and the Rights of Future Generations.’ In Charles Cockell (ed.), *Dissent, Revolution and Liberty Beyond Earth*. Springer.
- [16] Mueller, Robert; Gill, Tracy; Wiley, Larson; & Brink, Jeffrey. (2012). ‘Effective Utilization of Resources and Infrastructure for a Spaceport Network Architecture.’ AIAA SPACE 2012 Conference & Exposition, AIAA SPACE Forum, <https://doi.org/10.2514/6.2012-5233>.
- [17] Musk, Elon. (2017). Making Humans a Multi-Planetary Species. *New Space*. 5 (2): 46-61.
- [18] Ojha, Lujendra *et al.* (2015). ‘Spectral evidence for hydrated salts in recurring slope lineae on Mars.’ *Nature Geoscience*. 8: 829-832.
- [19] Parfit, Derek. (1984). *Reasons and Persons*. Oxford University Press.
- [20] Ross, W.D. (1930). *The Right and the Good*. Clarendon Press.
- [21] Schwartz, James. (2011). ‘Our Moral Obligation to Support Space Colonization.’ *Environmental Ethics*. 33 (1): 67-88.
- [22] Schwartz, James. (2016). ‘Lunar Labor Relations.’ In Charles Cockell (ed.), *Dissent, Revolution and Liberty Beyond Earth*. Springer.
- [23] Schwartz, James. (2018). ‘Worldship Ethics: Obligations to the Crew.’ *Journal of the British Interplanetary Society* 71: 53-64.
- [24] Schwartz, James. (2019). ‘Space Settlement: What’s the Rush?’ *Futures*. <https://doi.org/10.1016/j.futures.2019.02.013>
- [25] Stoner, Ian. (2017). ‘Humans Should not Colonize Mars.’ *Journal of the American Philosophical Association*. 3 (3): 334-353.
- [26] Szocik, Konrad *et al.* (2018). ‘Biological and social challenges of human reproduction in a long-term Mars base.’ *Futures*. 100: 56-62.