The freedom of future people

Andreas T Schmidt (University of Groningen)

The Freedom of Future People

Andreas T Schmidt, University of Groningen, comments welcome: a.t.schmidt@rug.nl;

Abstract: What happens to liberal political philosophy, if we consider not only the freedom of present but also future people? In this article, I explore the case for long-term liberalism: freedom should be a central goal, and we should often be particularly concerned with effects on long-term future distributions of freedom. I provide three arguments. First, liberals should be long-term liberals: liberal arguments to value freedom give us reason to be (particularly) concerned with future freedom, including freedom in the far future. Second, longtermists should be liberals, particularly under conditions of empirical and moral uncertainty. Third, long-term liberalism plausibly justifies some restrictions on the freedom of existing people to secure the freedom of future people, for example when mitigating climate change. At the same time, it likely avoids excessive trade-offs: for both empirical and philosophical reasons, long-term and near-term freedom show significant convergence. Throughout I also highlight important practical implications, for example on longtermist institutional action, climate change, human extinction, and global catastrophic risks.

1 Introduction

What does liberal political philosophy imply about our duties towards future people? The existing discussion on liberal freedom and future people often diagnoses a tension: to protect future people's interests requires interfering with people's options now, for example by restricting access to gas-guzzling cars and cheap meat. Liberal views of freedom, however, seem to speak against interference. Several authors thus conclude that liberal views clash too much with our duties towards future people and defend a republican or perfectionist view instead (Coeckelbergh, 2021; Fragnière, 2016; Pinto, 2021).[1]

This tension picture, however, misses something important: what about the freedom of future people? For example, the German Constitutional Court recently overturned the government's proposed climate bill. It argued that too much emissions reduction was shifted towards the future, which would violate a constitutionally mandated 'intertemporal protection of freedom'.[2] Or consider this preamble:

> 'We the People of the United States, in Order to form a more perfect Union, establish Justice, insure domestic Tranquillity, provide for the common defense, promote the general Welfare, and *secure the*

---

[2] (Bundesverfassungsgericht, 2021). The German constitution protects 'freedom of development' and 'freedom of the person' (Article 1 and 2, GG) as constitutional rights with an 'eternity clause' (Article 79 (3)).

Blessings of Liberty to ourselves and our Posterity, do ordain and establish this Constitution for the United States of America.' [my emphasis]

In this article, I argue that liberals should heed the call to 'secure the Blessings of Liberty to ourselves and our Posterity'. Specifically, I explore the case for:

> *Long-term liberalism*: freedom should be a central goal in institutional design, public policy and social norm-setting; hereby we should often be particularly concerned with *ex ante* effects on long-term future distributions of freedom.

I defend long-term liberalism in three steps.

In section 3, I argue that liberals should be *long-term* liberals. Liberal arguments to value present freedom give us reason to value future freedom too (3.1). And sometimes, long-term distributions of freedom should weigh particularly heavily (3.2 and 3.3).

In section 4, I argue that longtermists should be long-term *liberals*. Roughly, longtermists hold that improving the *ex ante* value of the long-term future should be a key moral priority, where the long term can include people thousands or even millions of years from now. Giving longtermism a 'liberal slant' makes it more plausible under conditions of empirical and moral uncertainty and helps capture some intuitions in intergenerational ethics.

In section 5, I return to the 'tension worry'. I argue that long-term liberalism can justify restricting freedom now to safeguard the freedom of future people, as seems plausible in cases like climate mitigation. At the same time, it likely avoids excessive trade-offs: for both empirical and philosophical reasons, good long-term distributions typically preserve short-term freedom too.

In section 6, I briefly discuss population axiology. As a generic view, long-term liberalism appears surprisingly robust despite widespread disagreement and befuddlement in population ethics.

The article's original contributions are both theoretical and practical.[3] I explore the philosophical case for long-term liberalism but also highlight practical implications along the way, for example on longtermist institutional action, climate change, human extinction, and global catastrophic risks.

I should admit upfront that I only give the 'initial case' for long-term liberalism. Not all questions and potential objections will be answered. Still, I think the case will be strong enough – and its implications important enough – for political philosophers to want to probe it in more detail.

2   What is long-term liberalism?

The position I explore in this article – and whose implications I consider – is:

> *Long-term liberalism*: freedom should be a central goal in institutional design, public policy, and social norm-setting, and we should often be particularly concerned with *ex ante* effects on long-term future distributions of freedom.

I leave some terms vague for now, because I defend long-term liberalism as a generic view that allows for different takes on longtermism (section 2.2) and population ethics (section 6). But note already that '*ex ante*' here means what effects one should expect relative to one's evidence.

---

[3] In philosophy, existing work on future generations typically focuses on a particular liberal principle like Mill's Harm Principle rather than liberal freedom itself (Skagen Ekeli, 2006; Wallack, 2004), on republican theories of freedom (Beckman, 2016; Katz, 2017; Smith, 2013), or on some other view of freedom like perfectionism (Coeckelbergh, 2021; Hannis, 2015). (Vercelli, 1998) is the first academic source I could find that invokes future option-freedom, arguing that sustainable development should preserve the diversity of future option sets. (Ferretti, 2023) appeared while writing this article. Its topic overlaps with mine, but I reach different conclusions, focus on a more comprehensive set of values to support liberal freedom (rather than only on respect), discuss all our obligations towards future people rather than just environmental duties, discuss how to distribute freedom, and explicitly consider the long-term future.

Moreover, '*ex ante*' is not committed to one particular decision-theory and allows for expected value theory but also other decision-theories.

Let me first say what I mean by 'liberalism' and 'freedom'.

2.1 What is liberalism?

'Liberalism' here refers to freedom-based liberalism, that is any view in political philosophy that (i) adopts a liberal view of freedom and (ii) accords great (albeit *pro tanto*) value to socio-political freedom in the justification of institutions, public policy, and social practices (and the norms that underlie them).[4] Of course, 'liberalism' can refer to many other ideas and positions, but I don't discuss them here.[5]

Starting with (i), what makes a theory of freedom liberal? I here understand liberal freedom as *option-freedom*: freedom, essentially, is about which options are open to a person and which ones foreclosed (Pettit, 2003). Theories of liberal freedom typically start with an account of 'options' (or 'specific freedoms and unfreedoms'). Some theorists hold that I am free to do a specific act, if and only if I am not subject to any interpersonal constraints with respect to that act. Other theorists, including myself, defend the Ability View which I assume here: I am free to do an act if and only

---

[4] I here discuss liberalism with a broader scope that also applies to informal social practices (and the norms that govern them). My arguments could also be run with a narrower scope that applies to formal institutions and public policy but not to informal social practices or only a small subset of these.

[5] For example, I set aside 'public reason liberalism' (Vallier, 2018) but also 'practical liberalism', that is, a tradition committed to particular institutions, such as specific individual rights and the rule of law, without invoking any justificatory theory for them (although 'freedom-based liberalism' justifies some form of practical liberalism).

if I am able to do that act.[6] The absence of interpersonal constraints is necessary but insufficient, freedom also requires actual abilities. While I assume the Ability View here, most arguments in this article could also be run, *mutatis mutandis*, with the 'absence of interpersonal interference view'.

Liberalism concerns not just specific freedoms and unfreedoms but *overall* freedom (Carter, 1999; Steiner, 1983). Overall freedom comes in degrees and is typically thought to be a function of the specific freedoms (and sometimes unfreedoms) available to a person. A rich literature in economics and philosophy seeks to measure a person's overall freedom. I here go along with this literature and assume that freedom can be measured. Accordingly, I ignore worries one might have that freedom is not the kind of thing that could ever be measured or compared.

The simplest measure counts the number of specific freedoms available to a person (Pattanaik & Xu, 1990). But economists and philosophers offer more sophisticated measures too. The details are beyond my scope, but notice two choice points.

First, when comparing option sets, it matters how we individuate their members. Having two different wine cellars gives you more freedom than having two bottles of wine. How to individuate options is beyond my scope, so I just assume that, other things equal, the more options one has, the more liberal freedom one has and that there is some way to individuate options consistently across option sets.[7]

Second, most, but not all, authors argue that overall freedom depends not only on the quantity but also the quality of the options. Some also hold that their variety matters as such, others deny this. Again, I do not develop a criterion to assess the quality of options but simply assume that a plausible

---

[6] See (Kristjánsson, 1996; Miller, 1983; Shnayderman, 2013; Steiner, 1994) for the absence of interpersonal interference view and (Kramer, 2003; Parijs, 1997; Schmidt, 2016; Sen, 1988, 1999) for the Ability View.

[7] Carter, for example, individuates specific freedoms and unfreedoms in terms of the spatio-temporal regions one can occupy with one's body (Carter, 1999).

measure captures that, other things equal, the better an option is according to some plausible criterion of value, the more it contributes to overall freedom.[8]

Note that the above, generically liberal understanding is not meant to cover *all* theories of freedom. I exclude other philosophical views on freedom, such as moralised theories – including libertarian rights-based theories – and neo-republican views.[9]

Freedom-based liberalism (i) adopts a liberal theory of freedom and (ii) accords great (albeit *pro tanto*) value to liberal freedom in the justification of institutions, public policy, and social norms. Recent theorists cash out (ii) by assessing institutions and policies in terms of their impact on the *interpersonal distribution* of overall freedom (Carter, 1999, Chapter 3; Schmidt, 2014).[10] For example, according to van Parijs, a universal basic income will facilitate a better distribution of freedom for all (Parijs, 1997). But what criterion should liberals use to assess outcomes in terms of interpersonal distributions of freedom?

Answering the distributive question is complicated. First, there might be different ways to accord value to freedom. Some theorists recently suggested that freedom has both *telic* and *deontic* value (Garnett, 2018, pp. 551–552; Intropi, 2022). For example, freedom has telic value in virtue of being good for individual wellbeing and societal progress. The right response to telic value is to *reduce*

---

[8] See (Schmidt, 2022) and (Pattanaik, 2018) for an overview of these and other questions around freedom measurement.

[9] See (Schmidt, 2022) for these distinctions. Still, my implications matter for republicans too, as they often hold that option-freedom is a necessary but insufficient condition for freedom (one also requires the absence of domination) (Pettit, 2003; Schmidt, 2018a).

[10] This broad approach is dominant in much recent theorizing and endorsed – or at least gestured towards – by theorists such as (Carter, 1999; Cohen, 2011; Ferretti, 2023; Hees, 2012; Kramer, 2003; Parijs, 1997; Rawls, 1971; Schmidt, 2014, 2016; Sen, 1988; Spencer, 1873; Steiner, 1994).

unfreedom and *promote freedom*. Freedom's *deontic* value, in contrast, it not about promoting the good but about respecting other persons as purposive beings who set their own ends. Accordingly, we should not interfere with other agents, even if doing so might be good for them in terms of wellbeing or even option-freedom. The appropriate response to deontic value is not to promote but to *respect freedom* (Intropi, 2022).

While little has been written on deontic value, I gather there are two approaches to account for it.[11] Some authors argue that option-freedom captures our telic reasons to promote freedom but that we also need a second, moralized theory of freedom to capture its deontic value (Bader, 2018). Others think that we can do with one theory of freedom but capture its deontic value in our ranking of interpersonal distributions, for example, by placing greater weight on *equal* or *sufficient* freedom for all. I do not have space to discuss moralized views, so I assume the second route.[12]

The distributive question is further complicated, because we are considering distributions across generations, including distributions with variable populations. One contentious question should be whether we have reason to bring people into existence because they would have high levels of freedom. Do we not only have a duty to 'make people free' but also to 'make free people'?

To make my task feasible, I will assume two separate criteria for ranking outcomes in terms of distributions of freedom.

First, for our telic reasons to promote freedom, I assume that a plausible ranking (a) considers how much aggregate freedom there is but also (b) weighs freedom gains to those who have less as

---

[11] I assume that freedom has telic value that we ought to promote, which is a widely shared view, and allow that freedom might also have deontic value and then consider the implications of both. (For the record, I am inclined to think that freedom's telic value – once spelled out properly – captures what is intuitively plausible about deontic value.)

[12] Thanks to Ian Carter for helpful discussions.

contributing more to the 'betterness' of a distribution than equally-sized freedom gains to those who have more overall freedom (similar to a Pigou-Dalton transfer). To capture (a) and (b), I assume prioritarianism which ranks outcomes by a weighted sum of individual overall freedom with weights that decrease for marginal gains depending on how much freedom an individual already has (I assume any freedom gain, no matter the level, still has positive marginal value).[13] I assume, for now, that we have telic reason to bring people with high levels of freedom into existence.

For the deontic value of freedom, I assume the following. First, I go along with Maria Paola Ferretti who argues that to respect persons, we must guarantee every person with 'a certain measure of freedom', including future people (Ferretti, 2023). I interpret this as a sufficiency level. Second, I also assume that we have deontic reason to abstain from courses of action that foreseeably involve severe freedom violations and abstain from actions that would severely and unjustifiably reduce people's freedom, including the freedom of future people (including reductions above the

---

[13] I prefer prioritarianism over unweighted Totalism, because additional freedom likely has greater telic value for those with less freedom. I also prefer it over 'greatest equal freedom' (as suggested for example by (Norman, 1987; Rawls, 1971, pp. 60; 302; Spencer, 1873, p. 35)): some inequality seems acceptable, if necessary for far more aggregate freedom; moreover, prioritarianism avoids 'levelling down' (see Parfit, 1997). Steiner would argue that such problems do not arise, because everyone should have *equal freedom* and the aggregate is fixed anyway (Steiner, 1987, 1994). Carter convincingly argues that the aggregate can vary (Carter, 1999, Chapter 9.4), which becomes even more apparent for variable populations. I also prefer prioritarianism over pluralist distributive egalitarianism, because it captures freedom's decreasing marginal value, does not require non-separable intergenerational comparisons nor gives any reasons for intergenerational levelling down. For my argument, however, the differences do not matter very much. Moreover, one could additionally assume a negative egalitarian weighing for freedom inequalities within groups, societies, or generations.

sufficiency threshold). For example, we have respect-based reason not to emit greenhouse gases for minor benefits now when this would result in large reductions of future people's freedom. Third, I assume our respect-based duties are directed towards other persons rather than being 'impersonal duties'. For future people, those could be actual persons or 'types' of future persons (or persons *de dicto*) (see, for example, (Kumar, 2018)). Finally, I assume what is called the Asymmetry in population ethics: we have respect-based reason not to bring people into existence with very low levels of freedom but no respect-based reason to bring people with high levels of freedom into existence.

These stipulations strike me as somewhat plausible. But, primarily, I make them to make my task in sections 3 to 5 possible. In section 6, I discuss other options and permutations in distributive ethics and population axiology.

2.2   What is longtermism?

A recent paper provides the following definition of longtermism for political philosophy:

> '*Institutional longtermism*: when evaluating the moral choiceworthiness of institutions, we should often give significant weight to their expected long-term effects.' (Schmidt & Barrett, 2023, p. 6)

Institutional longtermism comes in different strengths, depending on how 'significant weight' is filled in. Under *strong* longtermism, positively affecting the long-term future is *the* key priority and typically outweighs concern for the near term. Under *weak* longtermism, positively affecting the long-term future is *a* key priority but not necessarily the most important one.

Schmidt and Barret also distinguish longtermism by *how far* into the future it reaches (Schmidt & Barrett, 2023, p. 2). On one interpretation, we include the entire future of humanity (including all 'humanity-originating intelligent life' of moral value). For example, effects on people's wellbeing a

million or billion years from now matter. Call this 'long longtermism'.[14] Longtermists typically assume our moral reasons are 'scale-sensitive' and become weightier the greater the number of individuals affected by our actions (some even assume that reasons scale linearly). Effects on the long-term future thus weigh so heavily, because there are so many expected future people. Greaves and MacAskill provide estimates that range from 10^{14} future people to scenarios – involving artificial persons or space settlement – with estimates between 10^{18} to 10^{45} (Greaves & MacAskill, 2023). Given its enormous expected size, the far future will weigh heavily indeed.

Of course, 'long longtermism' is contested and hotly debated.[15] But, following Schmidt and Barrett, I here also include *medium longtermism* which only considers value within the next several hundreds to thousands of years. For example, maybe there is a high baseline extinction risk that humanity can't push down low enough for longer timelines to make sense.[16] Or maybe, given epistemic constraints, our long-term effects all 'wash out' in expectation (Tarsney, 2023). Note that even

---

[14] Longtermism is incompatible with a positive rate of pure time preference in an exponential discounting framework but compatible with uncertainty about such a discount rate (Weitzman, 1998) and with a positive discount rate in a non-exponential, hyperbolic framework (Mogensen, 2019).

[15] For example, is longtermism problematic, because we might be essentially clueless about our long-term impact (Greaves, 2020; Greaves & MacAskill, 2023, sec. 7; Mogensen, 2020), because longtermism might violate anti-aggregative or other deontological duties (Curran, 2023; Heikkinen, 2022; Unruh, 2023), because it might make problematic assumptions in population ethics (Greaves & MacAskill, 2023, sec. 6; Tarsney & Thomas, 2020; Thomas, 2022), because the expected value of humanity's long-term future of humanity might be net-negative (MacAskill, 2022, Chapter 9) or because it might rely on a fanatical decision-theory (Kosonen, 2023)?

[16] Cowen defends such a view in a talk (Cowen, 2023), and climate models sometimes discount future people's interests for the same reason (Stern 2006). Also see (Thorstad, 2023).

medium-term effects will be weighty, as expected future people over the next thousands of years still outnumber people alive today by some orders of magnitude.

So, longtermism comes in different flavours:

	Strong ('*the* key priority')	**Weak** ('*a* key priority')
Long (all future value)	a)	b)
Medium (value over the next hundreds to thousands of years)	c)	d)

My arguments in this article are compatible with all four options, including all the 'shades in between', as weak vs strong and medium vs long come in degrees. Still, in what follows, I implicitly assume b) or d).[17]

Of course, after reading this article, some readers might still reject all forms of longtermism. But even for them some arguments will be relevant: they might agree with my arguments that the freedom of future people matters morally (section 3.1) but disagree with sections where that is cashed out in a (broadly) longtermist framework.

3 Why liberals should be longtermists

In this section, I argue that liberals should be *long-term liberals*: liberals should care about future freedom too (section 3.1); and, sometimes, they should be *particularly* concerned with long-term distributions of future freedom (sections 3.2 and 3.3).

---

[17] I elsewhere argue that animal liberal freedom matters too (Schmidt, 2015, 2018b). But I here exclude non-human animals to keep things manageable (see (Browning & Veit, 2023; MacAskill, 2022, pp. 208–213)).

3.1 Liberals should care about future freedom

Liberals should care about future freedom, because the reasons to value freedom now apply to future freedom too.

Consider how and why freedom matters for liberals. Freedom might have intrinsic value or instrumental value. Liberals can but need not believe that freedom has intrinsic value. Freedom can still be robustly important as a social value, even if it only matters instrumentally. The first structural reason is that, even if only instrumentally valuable, freedom still has *non-specific* value: 'A phenomenon, *x*, has non-specific instrumental value iff *x*, without regard to the nature of its specific instances, is a means to some other valuable phenomenon *y*.' (Carter, 1999, p. 44) Consider an analogy: even if money only has instrumental value, 'we do not value money only as a means to buying the latest Mozart recordings or as a means to eating a bar of chocolate, but also as a means to satisfying whatever our future desires turn out to be.' (Carter, 1999, p. 34). Here are several arguments why freedom has *non-specific* value.

First, across time, people change in what they like, enjoy and value (Carter, 1999, pp. 50–54). At the same time, we underestimate how much we change over time (Quoidbach et al., 2013). Option-freedom affords our future selves the necessary options to pursue what they want, like and care about. Following this argument, liberals should focus on freedom across longer timeframes rather than just freedom at one point in time (Carter, 2013; Schmidt, 2017).[18]

Second, before you commit to a long-term option – a career, life partner, house and so on – it makes sense to first do some exploring to gather information about yourself and your options. Using the language of computer science, one should first 'explore' options before 'exploiting' them (Christian & Griffiths, 2016).

---

[18] The latter point suggests that when ranking interpersonal distributions, distributions of lifetime freedom are more plausible than 'time-slice' distributions. I ignore these niceties here.

Third, John Stuart Mill famously advocated 'experiments in living' (Mill, 1859). Such experiments teach us about the world and ourselves. And when we take up new hobbies, see new places, learn new skills, or have novel experiences, we also experience time to pass more slowly and enjoy our lives more (Howell et al., 2011).

Fourth, autonomy and agency require option-freedom. For example, Tom Hurka argues that agency requires that we reject options, and without freedom we cannot reject options (Hurka, 1987). More empirically, psychologists argue that without choices we develop feelings of learned helplessness and fail to develop feelings of agency and personal control (Peterson & Seligman, 1984).

Finally, Tim Scanlon argues that having some level of choice is typically required for self-expression and individuality (Scanlon, 1988).

So, freedom has non-specific instrumental value for individuals. The second structural reason that freedom is robustly important lies in it being a *social value*. Ethicists engaged in axiology often enquire whether some good has intrinsic value at some fundamental level. For example, stripping away all empirical confounding conditions, is a world with more freedom intrinsically better? Political philosophers, in contrast, see freedom as a *social value* and ask: to what extent would societies be better if its members have freedom? To answer the latter question, we must consider all the different ways in which freedom might be valuable for individuals across longer timescales and for whole societies – and we must do so without stripping away empirical reality. There is a good case freedom is an important social value in this sense.

First, people differ in their likes, personalities, and conceptions of the good. Given this diversity – both natural and cultural – free societies provide better conditions for people to do what they want, value, and like.

Second, in free societies, I can benefit from other people's freedom. For example, in free and diverse societies, individuals encounter better conditions to learn from other people's experiments

and conceptions of the good. Moreover, given widespread and reasonable disagreement, it seems rational to be somewhat uncertain about morality, axiology and political philosophy. And moral uncertainty provides a *pro tanto* reason not to force people into particular conceptions of the good which we cannot confidently judge to be good. At the same time, free societies allow for experimentation and diversity and thus generate better epistemic conditions to somewhat reduce empirical and maybe even moral uncertainty.

Finally, for social relations to be *respectful*, individuals must at least sometimes respect other people's choices. Accordingly, some option-freedom and non-interference is necessary for social conditions in which people can respect each other as agents.[19]

Overall, there is a strong case that liberal freedom has non-specific value, particularly as a social value.[20] This case, I now argue, applies to the freedom of future people too.

Absent any countervailing arguments, the above case for present freedom should apply to future freedom too. It is likely that future individuals will change over time, benefit from exploration and experiments in living, and require freedom for agency, personal control, self-expression, and

---

[19] Some might claim that 'deontic' reasons to 'respect freedom' entail only negative duties of omission, like non-interference (Intropi, 2022). Even if that were so, negative claims and duties for individuals can generate positive duties at the collective and institutional level. For example, positive actions are required to establish and enforce the laws and norms that protect individuals against disrespectful interference. (See Shue's seminal argument that even 'purely negative rights' entail positive institutional duties (Shue, 1980)). Moreover, beyond deontic, we might also have telic reason to prefer societies with respectful relations.

[20] Some empirical evidence also strongly suggest societies with more freedom also have higher wellbeing and life satisfaction (controlling for other factors) (Bavetta et al., 2014; Haller & Hadler, 2004; Veenhoven, 2000). Moreover, nearly all countries with very high levels of income, quality of life, and development are in some sense liberal (unless they are tax havens, very resource-rich per capita or Singapore), which is more indirect evidence. A conclusive empirical case, however, is beyond my scope.

respectful social relations. These reasons have applied to past societies and apply to us now. Therefore, and absent any opposing considerations, they should apply to future people too.

We can even go further. Some reasons appear to amplify *across generations* and *longer timeframes*.

First, free societies generate better conditions for experimentation and learning. Now, with longer time frames, such positive externalities compound, because more people can profit from such experiments and insights. Moreover, imagine such experiments come with short-term 'costs'. Longer timeframes provide more time for such experiments to pay off. Compare an analogy: should we invest in exploratory and costly research that only rarely generates results? Such investments can be efficient, if we have a long time to utilize fundamental insights afterwards, but not if they must pay off short-term.

Second, epistemic arguments and the diversity argument amplify *across generations*. Over the last two hundred years, societies have been changing at a speed much faster than before. If this trend continues even a little bit, future people are likely to want to lead very different lives from us, giving us all the more reason to preserve their freedom to do so.

So, liberal arguments imply we should value the freedom of future people. You can accept this part of my overall argument even if you entirely reject longtermist ideas. But the next two sections explore arguments for the longtermist claim that we should sometimes be *particularly* concerned with our effects on long-term distributions of freedom.

3.2 The Scale Argument

When one has significant *ex ante* effects on long-term future distributions of freedom, those effects should typically be vast, because the expected size of the long-term future is so vast. For our telic reasons to promote freedom, I assumed a prioritarian criterion. Now, if there are actions in institutional design, public policy or social change that have significant *ex ante* effects, those are likely to have significant effects on the *ex ante* long-term priority-weighted sum of freedom. So,

long-term liberalism should 'sometimes be particularly concerned with … *ex ante* effects on long-term future distributions of freedom'. This, in a nutshell, is the Scale Argument for long-term liberalism.

While this argument works most straightforwardly for freedom's telic value, it applies to its deontic value too. For that, I assumed that we aim for outcomes where individuals have sufficient freedom and abstain from courses of action that foreseeably lead to severe freedom violations or that would severely and unjustifiably reduce future people's freedom. Now, some actions in institutional design, public policy or social change might have significant *ex ante* effects on people falling above or below the sufficiency line in the long-term future or they might cause long-term freedom violations and freedom reductions. Such *ex ante* long-term effects might weigh more heavily than near-term effects, because they might affect so many more people. Accordingly, even the deontic value of freedom can give us reason to be particularly concerned with our *ex ante* long-term impact.

The Scale Argument encounters a practical challenge: are there ever any actions through which we can *predictably* and significantly affect the long-term future? I now survey examples where it seems we can and show how those apply to long-term distributions of freedom too. To this end, distinguish two ways of positively affecting the long-term future (MacAskill, 2022): (i) you can increase the probability that humanity, conditional on its future existence, sets out on a more valuable *long-term trajectory*, or (ii) you can increase the expected *size* of the long-term future, primarily by extending how long humanity will exist.

(i) Trajectory-change

Trajectory-change interventions seek to make it more likely that humanity, conditional on its future existence, follows a valuable long-term trajectory. I now briefly survey how public policy, social norms, and institutional design might affect probabilities of trajectory-change and how they would affect distributions of freedom too.

First, persistent climate change could have long-term effects. CO_2 can stay in the atmosphere for thousands of years. And with more weather irregularities, droughts, floods, and more, climate change might require continuous adjustments whose costs could add up to be huge over the long term. Climate change mitigation could free up those resources for other projects and thereby expand future freedom.

Second, catastrophic events might be so bad that humanity is substantially worse off for long periods or even never reaches its civilizational potential. For example, a nuclear winter or bioengineered pandemic might not kill everyone, but humanity might persist for a long time in numbers too low or in conditions too bad for valuable long-term trajectories to get off the ground. Its overall freedom would be lower too.

Third, the potential development of human-level *artificial general intelligence* (AGI) comes with grave risks. The ends pursued by a powerful AGI might not align with those humans have or should have ('the alignment problem') and such an AGI might exercise far-reaching control over humanity ('the control problem'). Or an authoritarian country might use AI to entrench its power and values ('the entrenchment problem'). AI-entrenchment and takeover would likely involve bad long-term distributions of freedom. And in misalignment scenarios, human freedom might not feature as an important goal for AGI. At the same time, if developed and used well, AI might increase our individual and collective capabilities and vastly increase our freedom. So, plausibly, there is a wide spread of trajectories some of which have very high and some very low levels of freedom.

Fourth, some longtermists worry about negative institutional lock-ins in which undesirable institutions get locked in for many years. The most dramatic is a 'totalitarian lock-in' in which authoritarian power gets locked in long-term. Bryan Caplan, for example, worries that centralising political power globally will increase the risk of a global totalitarian lock-in (Caplan, 2011). Some consider AI-entrenchment the biggest risk for totalitarian lock-in: if an authoritarian leader can use artificial superintelligence, they might be able to lock in their power far beyond their own borders

and even beyond their lifetime (Greaves & MacAskill, 2023, sec. 4.3). Totalitarian lock-ins would be bad news for future freedom.[21]

Long-term liberalism implies actions that promote positive trajectory-change. Promoting the telic value of freedom, high-freedom trajectories will, *ex ante*, have a higher priority-weighted aggregate of freedom than low-freedom trajectories. Respecting the deontic value of freedom, we have reason to prevent that people fall below any freedom sufficiency line, which gives us reason to prevent low-freedom trajectories involving totalitarian lock-in, persistent climate change, and catastrophic risks that humanity does not recover from, and more. Moreover, we arguably have respect-based reason to abstain from actions that, *ex ante* and unjustifiably so, lead to unnecessary freedom violations or significantly lower rather than higher freedom. Accordingly, we have reason to abstain from actions, like bio-engineering dangerous viruses, that make high-freedom trajectories less likely (even beyond sufficiency levels).

(ii) Prolonging humanity's future

The central type (ii) action reduces humanity's extinction risk. Some sources of this risk are 'natural', such as supervolcanoes, asteroids and comets. But most total extinction risk comes from man-made risks, such as misaligned artificial intelligence, (bioengineered) pandemics, nuclear war, climate change and 'unknown risks' from future technological developments (Ord, 2020). Plausibly, there are things we can do to lower humanity's total extinction risk.

First, many suggested interventions target particular extinction risks, such as climate change mitigation, nuclear arms control or disarmament, preventing future pandemics, technical AI safety

---

[21] My list of trajectory change interventions is incomplete. For example, perhaps a 'great stagnation' – with stagnating growth and innovation and a declining population – is another long-term low-value trajectory and an extinction risk factor (MacAskill, 2022, Chapter 7).

and AI governance solutions, and much more.[22] Note that even 'targeted' interventions typically require institutions, social norms, or collective action, making them important priorities for long-term liberals.

Second, we might also focus on *risk factors* (Ord, 2020). Risk factors contribute to total extinction risk without causing extinction itself. For example, a great power conflict might not itself end humanity but still increase the risk of destructive artificial intelligence, catastrophic bioweapons or accidental nuclear war. Or disruptive climate change might raise the risk that other extinction causes are triggered. Public policy, and social and institutional change plausibly have non-negligible *ex ante* impact on risk factors. For example, developments that undermine good governance, lead to the capture of democratic institutions, or even to the rise of fascist movements arguably make society less well-prepared to lower catastrophic risk or handle catastrophes when they arise. Conversely, productive international cooperation might reduce the risk of nuclear war, engineered pandemics, climate change, a great power war, and more.[23]

So, plausibly, lowering total extinction risk has significant *ex ante* effects on the value of the long-term future. Such action would become a priority for long-term liberals too, because humanity's premature extinction would massively curtail human freedom. This connection is mostly overlooked in the literature, but Bruce E. Tonn provides a rare exception: "The worst case for maintaining options for future generations is pre-mature human extinction. In this case, there are

---

[22] The precise proposals are beyond my scope. See, for example, (Ord, 2020) for practical proposals on all of those risks, and see (Hilton & McIntyre, 2022) on reducing the risk of a nuclear war, (Esvelt, 2022) on preventing bioengineered pandemics, and (Hilton, 2022) on reducing AI risk.

[23] Of course, often it will seem plausible that we should do *something* about extinction risks and risk factors but, at first glance, not know exactly what. In such cases, there often is great *ex ante* value in doing more research to find out what, if anything, we can do.

no options at all for future generations, no societies to evolve, no lives to live." (B. E. Tonn, 2009, p. 432)

If we can lower total extinction risk – which seems plausible – doing so is an important priority for long-term liberals, as doing so would vastly increase the *ex ante* priority-weighted total sum of freedom. The liberal case for extinction risk reduction is strongest and most straightforward for the telic value of freedom. I assumed earlier that the deontic value, in contrast, does not give us reason to bring people with very high freedom into existence. We then only have much weaker 'respect-based' reason to reduce extinction risk, namely if extinction risk scenarios involve some periods of insufficient freedom before full extinction.

Overall then, because of the scale of potential impact, long-term liberals should often be particularly concerned with effects on long-term distributions of freedom. While empirical details are beyond my scope, it seems plausible that there are things we can do to affect total extinction risk and the probabilities of long-term trajectory change. At the very least, liberals should advocate that we try harder to find out whether we can.

3.3 The Level Argument

The expected size of the future is the primary reason why long-term effects weigh so heavily. But there might be a secondary one: if we play our cards right, future people might have much higher *levels* of freedom than we do today. If so, we have *pro tanto* reason to nudge humanity towards high-freedom trajectories and to make it more likely that humanity persists long enough to achieve those levels. I elsewhere argue that great increases in freedom are still possible and, *ex ante*, valuable (Schmidt, 2023). Here is a superficial summary of three arguments.

First, industrialisation set off an unusual time for humanity, as the graph shows below.

Figure 1

GDP per capita and our capabilities more generally have greatly expanded in the medium-term past, and with it individual freedom. And society and technology might continue to change in ways that will expand future freedom. A simple inductive argument is that when people proclaim that 'this is as good as it gets', we should be sceptical simply because of how much societies have changed, and typically improved, since industrialisation.

Second, some candidates of technological progress have the potential to expand freedom. For example, automation and AI might 'liberate us from work' or at least from a lot of work people have to do but rather wouldn't. This would greatly increase the time people are free to use in ways they find enjoyable and valuable. And medical and scientific progress is likely to increase our capabilities by reducing mortality and morbidity.

Third, we are unlikely to have reached the pinnacle of social liberalisation. For example, zooming out a little, societies have only recently made significant strides towards reducing the unfreedom imposed on people based on their gender. Moreover, more liberal sexual norms and norms around

family, dating and more have also only recently emerged. Further expansions in freedom seem possible.[24]

Future average freedom might flatline or collapse. But, if humanity plays its cards right, it might also be substantially higher than it is today. Our reasons to promote freedom would then give us reason to nudge humanity away from low-freedom and towards high-freedom trajectories.[25] Our reason here will be primarily telic, although arguably we also have respect-based reason not to actively pursue actions that unjustifiably lower the probability of future high-freedom scenarios.

4   Why longtermists should be liberals

I have explored the case for long-term liberalism arguing that *liberals* should be long-term liberals. I now explore the case from the reverse angle: why should *longtermists* be liberals? In other words, if you have sympathies for longtermism, why should you care about long-term future distributions of freedom?

---

[24] One objection here is the supposed 'paradox of choice': too much choice makes us worse off overall (Schwartz, 2009). I elsewhere argue that this objection fails (Schmidt, 2023): freedom being non-specifically and socially valuable means the evidentiary burden for the paradox is high. Theorists like Schwartz do not meet this burden: the invoked experimental psychology evidence is too weak (and often does not replicate) and cannot be extrapolated to other choice contexts let alone aggregate effects across whole lives and societies. And, if anything, aggregate trends support the opposite conclusion.

[25] Of course, the Level Argument does not hold in all conditions, particularly given my axiological assumptions. For example, when we can either increase the near-term freedom of those with little freedom (like the global poor) or increase the (potentially) very high freedom of future people, our prioritarian weightings might prioritise the former, depending on empirics and on which prioritarian weights we choose.

In section 3.1, I already made part of that case: longtermists should be liberals, because of the case I presented for valuing both current and future freedom. But I now also make a case more specific to longtermism: I argue that giving longtermism a 'liberal slant' (i) makes it more plausible under conditions of empirical and moral uncertainty and (ii) helps capture some intuitions in intergenerational ethics.

(i)   Uncertainty and broad action

Efforts to positively affect the long-term future come up against two sources of uncertainty: predicting long-term effects is extremely difficult, and sometimes impossible (*empirical uncertainty*); second, we do not agree on and are uncertain about morality and axiology, so how can we be confident such efforts will actually be good long-term (*moral uncertainty*)?

In response to empirical uncertainty, we have seen some examples of longtermist actions that seem *ex ante* positive despite our uncertainty. Additionally, a thriving research programme seeks to build formal decision-models (like decision-making under risk, ambiguity and uncertainty) and long-run forecasting techniques to reduce our uncertainty regarding long-term impact.[26] In response to *moral* uncertainty, an active research programme develops formal models for decision-making under moral uncertainty (MacAskill et al., 2020). Such efforts will likely reduce hurdles to long-term impact. But beyond such technical answers, another response to uncertainty adopts broad heuristics and desiderata (Brighton & Gigerenzer, 2012). I now discuss three such desiderata to guide longtermist actions and argue that being liberal helps longtermists meet those desiderata.[27]

First, *robustness*: other things equal, longtermist actions are preferable when they are robustly valuable. Given empirical uncertainty, such actions are preferable, when they are valuable across

---

[26] See, for example, (Heal & Millner, 2014; Karger et al., 2023; Mogensen & Thorstad, 2022; Ord et al., 2010; Spangenberg, 2018; Tetlock & Gardner, 2015; B. Tonn & Stiefel, 2013).

[27] MacAskill arrives at similar but not identical desiderata (MacAskill, 2022, pp. 226–227).

plausible long-term scenarios. For example, effective public institutions and global cooperation seem valuable across many scenarios, including when unforeseen catastrophic risks arise. Robustness also has a moral dimension: given moral uncertainty, actions are preferable when they are valuable – or at least not very disvaluable – across different credible moral views. For example, fundamentalist Christian theocracy might have some small but non-zero chance of being the best political system. But theocracies run graver moral risks than liberal democracies, because on many credible normative views theocracy is far more wrong than liberal democracy.

When longtermist actions positively affect freedom, they typically (but not necessarily) perform better on robustness. Freedom itself seems 'morally robust': rather than requiring one conception of the good or one moral theory, freedom is valuable for many different reasons and from many ethical viewpoints. Moreover, given its pluralistic justification, and its value being non-specific, freedom is valuable across different empirical conditions. For example, if humans start desiring different things, freedom will still have non-specific value.

Second, *option-value* and *reversibility*: for several reasons, it is often wise to keep some options open and avoid actions that are hard or impossible to reverse. First, future generations will differ from us, and it's hard to know exactly how. Second, some chosen paths into the future are only valuable if particular empirical predictions come true. But given empirical uncertainty, we need flexibility to switch to a different path when our predictions don't pan out. Third, with long enough timeframes, future people are likely to know more and have better capabilities. For example, with more knowledge, computing power, and data, they might face less empirical uncertainty. Moreover, they will have 'temporally local' knowledge and thus better understand what solutions work for the problems they face. Applying a 'temporal subsidiarity principle', we should leave future people enough flexibility for action. Finally, free discourse and empirical and philosophical research might also reduce moral uncertainty for future generations. Accordingly, we should not lock in contested values that later turn out mistaken.

When longtermist actions positively affect long-term freedom, they typically help preserve option-value and reversibility and thus account for future generations being different, preserve adaptability across different and unforeseen developments, preserve some intergenerational subsidiarity, and place limits on how far societies and institutions can lock in contested values that pose moral risks under moral uncertainty.

Finally, *empowering the future*: by taking sensible actions now, we might even increase what future people know and what capabilities they will have. For example, societies can invest in building knowledge through education and research, pass on effective public institutions, build up capabilities through capital accumulation and capabilities for collective action, or improve the conditions for societies 'to get better at getting better' (see (Barrett, 2020)).

Longtermist actions that promote long-term freedom also increase future capabilities. If future people indeed are in a better position to make wise choices, such increased capabilities will typically yield better outcomes. Moreover, free societies likely provide better conditions for societies to progressively improve than societies with little freedom: they allow for more experimentation, both for individuals and groups, are more adaptable, and allow for freer discourse and the exploration of new ideas. This should provide better conditions for more empirical knowledge and improve moral insight across time, thereby helping future people make wiser choices.

Overall, and other things equal, longtermists should prefer actions that bring about good long-term distributions of freedom.[28] This obviously applies to *broad longtermist actions* which are actions that improve the general conditions for long-term value, for example constitutional design, investment

---

[28] Under moral uncertainty, we have a strong heuristics-based default reason to prefer freedom. Moral uncertainty might of course still make 'non-liberal' calls on specific policy issues. For example, perhaps the downside moral risks of allowing factory farming far outweighs its potential moral upsides (also see (Barrett & Schmidt, 2023) for similar issues arising for moral uncertainty and public reason liberalism).

in education and research, and securing liberal rights. Yet even *targeted actions* that focus directly on a long-term benefit or risk can be pursued in more and less 'liberal' ways. For example, other things equal, we should prefer energy policies that do not empower anti-liberal authoritarian governments and prefer pandemic protocols that are less likely to hollow out liberal rights.

(ii)   Non-instrumental concerns

Above I gave 'instrumental' arguments for longtermists to be liberals. But 'going liberal' also helps capture some *non-instrumental* intuitions in intergenerational ethics. For example, Andreas Bengtson defends a *pro tanto* duty not to treat future generations paternalistically (Bengtson, 2019). And Axel Gosseries and Jonathan Hoffmann argue that we have a duty to respect the 'sovereignty' of future generations (Gosseries, 2017; Hoffmann, 2022, pp. 190–193). I do not claim that long-term liberalism perfectly captures their normative concerns, after all Bengtson, Gosseries, and Hoffmann each adopt normative frames different from mine and from each other. But my suggestion is that long-term liberalism partly captures the intuitions behind their claims: preserving and promoting the freedom of future people will limit how far we treat future generations paternalistically and interfere with their sovereignty.[29]

Overall, longtermists should be long-term *liberals*: there is both a good general case to value future freedom and specific arguments as to why a 'liberal slant' makes longtermism more plausible.

I now return to the 'tension worry' with which I started this article.

---

[29] (Ferretti, 2023; Vercelli, 1998) also note the connection between anti-paternalism and the freedom of future people.

5   Intergenerational trade-offs

We started with an intuitive tension: protecting the future might require restricting individual freedoms today which creates tension between liberalism and our duties towards future people. Long-term liberalism now eases this tension: a concern with freedom itself can justify actions that benefit the future. For example, while climate change mitigation might constrain some of our options now, those might be necessary to protect the freedom of future people.[30] This, I take it, is an important and plausible takeaway.

However, might long-term liberalism not justify too much? With billions to trillions of future lives and their freedom at stake, wouldn't long-term liberalism swamp any concern with near-term freedom? This would be odd indeed: rather than protecting our freedom, liberalism would endorse interferences for the sake of a long and glorious future. Call this the *Swamping Objection*.

Whether swamping happens depends, *inter alia*, on the theoretical shape of longtermism. I here remain neutral between 'weak' or 'strong' and 'medium' or 'long longtermism'. Moreover, longtermism can build in constraints and prerogatives of various strengths to prevent swamping. And under medium longtermism the expected number of future people is less astronomical than under long longtermism. Still, even under weak medium longtermism, one might find it odd that *liberalism* should give so much weight to the far future. However, I think there are good reasons to think we need not be too worried about swamping.

First, while 'swamping' current freedom seems intuitively problematic, constrains on current freedom like in the climate example seem intuitively plausible.

---

[30] While decarbonisation will require freedom interferences, it might also increase near-term freedom, because better air quality would reduce mortality and morbidity. Which effect is stronger is an interesting question beyond my current scope.

Second, in many other cases, the following is plausible:

> *General Convergence*: being concerned with bringing about good long-term distributions of freedom typically gives one strong (but defeasible) reasons to preserve and promote near-term freedom.[31]

The main argument for General Convergence is that (un-)freedom now makes future (un-)freedom more likely – for at least three reasons.

First, social practices and institutions often display *persistence* (or *path dependence*).[32] Some institutions deliberately build in persistence, like supermajority requirements for constitutional change or even constitutional 'eternity clauses' (Araújo & Koessler, 2021). But practices and institutions often persist simply because they are functional or because participants have an interest to keep them going. Through persistence, freedom and unfreedom can be 'sticky'. Therefore, the best route to long-term freedom is likely via near-term freedom rather than an autocratic detour. Promises like 'we'll have a few years of a repressive dictatorship of the proletariat and then we'll transition to free communist heaven!' have not turned out well so far.

Second, (un-)freedom now might reinforce itself and thereby increase how much freedom future people will have. I earlier argued that, ideally, longtermist actions should empower future generations and that liberal institutions now – for example through good epistemic conditions for free discourse and experimentation – might increase long-term capabilities. Unfortunately, *unfree* institutions can reinforce themselves too. For example, extractive economic systems might enable

---

[31] General Convergence holds that the best long-term options will typically be pretty good short-term options that do not 'sacrifice' current people for the long term. It does not imply that the *best* long-term option will also be the *best* short-term option.

[32] See (Mahoney, 2000; Pierson, 2000) on path dependence. See (Abad & Maurer, 2021; Giuliano & Nunn, 2020) on persistence and (Sevilla, 2022) for a critical review.

an economic elite to influence political and legal institutions, creating a cycle through which economic systems become even more extractive (Acemoglu & Robinson, 2008, 2012).

Finally, the most drastic reinforcement scenario would be *lock-ins*. For example, current or future authoritarian dictatorships might use increasing AI capabilities – and other technological advances – to lock in anti-liberal values. Building and defending liberal values and institutions now will help reduce this risk.

However, perhaps General Convergence is only plausible for typical issues in politics, economics and policy but not for *catastrophic risks*. As argued earlier, reducing extinction risk and trajectory-changing catastrophic risk should be important priorities for long-term liberalism. Yet given the high stakes involved, one might worry that such risk reduction would routinely override near-term concerns. Whether this is the case in principle depends on theoretical commitments around population ethics, decision-theory and more. Practically speaking, however, I think long-term liberals need not worry too much about swamping here either. In addition to General Convergence, the following is also plausible:

> *Risk Convergence*: most measures to reduce catastrophic risk not only preserve or promote *ex ante* long-term future freedom but are likely to preserve or even promote *ex ante* near-term freedom too.

The first, more empirical argument for Risk Convergence runs like this: (a) when catastrophic risks materialise, they tend to not only affect wellbeing and cause harm, they also massively reduce freedom; (b) such risks are high enough to materialise in the near term; and (c) there are cost-effective measures to significantly reduce such risks. With some bridge premises, including liberalism, we can conclude that reducing catastrophic risk is a good investment to protect near-term freedom.

As mentioned before, premise (a) seems plausible: catastrophic events affect not only wellbeing or safety but freedom too. For example, the COVID-19 pandemic curtailed freedom through

lockdowns, increased global poverty, slower economic progress, and by cutting short many people's lives. Premise (b) seems intuitively plausible too: such risks are high enough to materialise in the near term. I write this in 2022/2023 just coming out of a pandemic and with the Russian president threatening nuclear war. Moreover, while AI risk scenarios have often been dismissed as science fiction in the past, the topic of AI risk has recently become more mainstream. Nearly every week brings new AI applications or achievements that would have seemed impossible just a few years ago. Aggregating different risks, some researchers put total existential risk as being significant *for this century*, roughly 1 in 6 in Toby Ord's case (Ord, 2020). Finally, premise (c) is plausible too: there are cost-effective measures, particularly at the institutional and policy level, to reduce near-term catastrophic risk. For example, Thornley and Shulman argue that catastrophic risk reduction meets the cost-benefit standards in US policy even if we only focus on near-term benefits (Shulman & Thornley, 2023). Overall, while more could be said to defend individual premises, reducing catastrophic risk seems a good investment to increase short-term freedom too.

The second argument is slightly more 'philosophical': (a) reducing humanity's near-term extinction risk makes our current options more valuable; (b) the more valuable our options, the more overall freedom we have (other things being equal); therefore, (c) reducing humanity's near-term extinction risk increases our overall freedom now.[33]

To defend (a), we could plug in a recent argument by Sam Scheffler. Imagine humanity had lost the power to procreate. Scheffler argues that many of our central activities and projects – building good institutions, doing fundamental research, writing novels, and more – would greatly lose value and meaning. Accordingly, humanity continuing after our deaths (an 'afterlife') seems necessary for us to find value and meaning in many of the central things we do and value (Scheffler, 2013, 2018).

---

[33] Thanks to Andreas Mogensen for helpful discussion of this point.

Together with premise (b), this implies that reducing humanity's existence risk makes us more free in virtue of the better options we have.[34]

The third argument for Risk Convergence draws on the ethics of risk. Philosophers here ask 'what makes risk impositions wrong, including in cases where the bad outcome never materialises?'. For example, if you force me to play Russian roulette and I get the empty chamber, you still did something wrong. Ferretti answers that risk impositions reduce overall freedom (Ferretti, 2016). Overall freedom is a function not only of my specific freedoms but also of the probability with which I can exercise them. The more likely it is I can exercise my options, the more freedom I have (other things being equal). For example, imagine my partner's New Year's resolution is to play Russian roulette with me in my sleep at some point this year. I now have less overall freedom, as it has become less likely that I can successfully exercise my future options. This argument has an interesting implication when applied to global catastrophic risk: the more global catastrophic risk we are exposed to, the lower our overall freedom in virtue of the lower probability with which we can exercise our options. Conversely, by reducing global catastrophic risk, we make people more free (necessarily and other things being equal).

6   Population axiology

For telic reasons to promote freedom, I so far assumed prioritarianism. For potential deontic reasons, I assumed sufficientarianism along with a reason to abstain from actions that unjustifiably reduce future people's freedom. As mentioned before, other assumptions and permutations in

---

[34] This argument of course depends on the plausibility of Scheffler's story and the generality with which it holds, discussion of which is beyond my scope. Moreover, it will not hold when annual extinction risk is already so low and an 'afterlife' so likely that further risk reductions would not improve our options.

distributive ethics and population axiology are possible for both telic and deontic reasons. So, how does long-term liberalism fare once we descend the population axiology rabbit hole?

Distinguish two different questions for political philosophy:

> *Population axiology proper*: Which order gives us the correct betterness ranking of outcomes with variable populations in terms of properties that are intrinsically valuable (like wellbeing)?

> *Population heuristics*: Which order of outcomes with variable populations gives us a plausible *heuristic* or *governing principle* for institutions, public policy and social norm-setting?

In this article, I only considered freedom's instrumental value. So, if freedom is not among the properties that make outcomes intrinsically better, I don't need a separate population axiology for it.[35] However, finding the correct *population heuristic* for freedom might be just as difficult, as that will depend on what the correct population axiology proper is. Unfortunately, 'solving population axiology' is beyond my current scope (and, frankly, abilities). Worse still, it is literally impossible for a population axiology to avoid all implications commonly thought to be counterintuitive (Arrhenius, 1999, 2000). Moreover, generic freedom-based liberalism itself does not uniquely imply one population axiology. Liberals have not written about population ethics in any detail and diverse views, intuitions, and authors abound within the liberal tradition (just compare John Stuart Mill,

---

[35] If freedom *is* intrinsically valuable, an interesting discussion could be had about what population axiology freedom should have. However, its results might ultimately not impact long-term liberalism's shape that much. First, even if freedom has intrinsic value, the 'bulk' of its value is likely instrumental. Second, many arguments for freedom having intrinsic value see it as a *constituent* component of something else of intrinsic value, for example autonomy. Accordingly, comparing distributions of freedom in isolation would not always track its intrinsic value. For example, outcome O_1 might contain a lot of freedom but very little autonomy, because autonomy's other ingredients are missing, whereas O_2 contains less freedom but more autonomy.

Herbert Spencer, John Rawls, Amartya Sen, Hillel Steiner, Nancy Hirschmann, and G.A. Cohen).

Accordingly, I don't defend one population heuristic here. Instead, I now first argue that long-term liberalism is surprisingly robust across different population heuristics.

(i)   Long-term liberalism is surprisingly robust across different population heuristics

I earlier distinguished actions that (i) set humanity on a better long-term trajectory conditional on there being future humans and those that (ii) extend humanity's future, primarily by preventing humanity's extinction.

Type (i) actions are an important priority across different population heuristics for freedom.

First, *Totalism, Averagism, Perfectionism*; I earlier argued that there is potentially a broad spread in how free or unfree people would be in different long-term trajectories. Totalism ranks outcomes by the aggregate freedom contained therein and Averagism ranks outcomes by their average freedom. Perfectionism places great value on reaching very high levels of freedom. If we can affect the probability of different trajectories, both Totalism and Averagism would view that as an important priority given the numbers involved. Similarly, if we can make long-term trajectories with high peaks of freedom more likely, Perfectionism would judge that as important.

Second, *wide person-affecting views*; wide person-affecting views in population ethics rank outcomes by comparing persons – or types or ranks of possible persons – across outcomes instead of ranking outcomes impersonally. Some versions of wide-person affecting views also endorse longtermist trajectory-change actions, as their value does not primarily consist in changing the number of future people but in improving the lives of types of future people (or of people *de dicto*).[36]

Finally, how would axiologies that attach negative weights to *inequalities* in freedom affect long-term liberalism? A concern with *intergenerational* inequalities might make us less longtermist, if average

---

[36] See (Thomas, 2022) on person-affecting views and longtermism more generally.

future freedom is much higher than current freedom. But if average future freedom is lower than near-term freedom – say because of totalitarian lock-in, climate catastrophes or AI takeover – such a concern would make us more longtermist. A concern with *intragenerational* inequalities, I think, is more likely to make us more longtermist. Bad long-term trajectories, such as AI takeover or totalitarian lock-in, would likely contain great and long-lasting inequalities in freedom within future generations. Much more could be said on this, of course.[37] But, overall, a concern with inequalities in freedom more likely supports rather than weakens the case for long-term trajectory change.

Take type (ii) actions next. A (partly) Totalist population heuristic gives us the most straightforward case for reducing extinction risk, as doing so would increase the number of expected people.[38] Other views can support extinction risk reduction via the Level Argument. If average future freedom would be much higher than short-term freedom or if it reached 'high peaks', Averagism and Perfectionism can support reducing extinction risk to increase the chance that humanity gets there (even if they would disagree on how comparatively important this will be).

(ii)   *Population axiology under moral uncertainty*

Given that there is widespread disagreement in population ethics – and no obvious way out – the above response was to perform a rudimentary 'sensitivity analysis'. Alternatively, one could proceed with moral uncertainty about population axiology. This makes sense for liberalism: sections 3.1 and 4 argued that long-term freedom is particularly important as a non-specific social value under

---

[37] Also see (Schmidt & Juijn, 2023) on inequality and longtermism.

[38] This case amplifies with a 'fanatical' decision-theory, where even tiny marginal reductions in extinction risk can swamp other higher-probability changes with smaller payoffs. But most longtermist priorities can do without fanaticism (Kosonen, 2023), particularly so in institutional contexts where probabilities of difference-making are much higher than for individuals (Schmidt & Barrett, 2023).

conditions of moral disagreement and uncertainty. So, what would population axiology proper under moral uncertainty imply for a population heuristic for freedom?

Population axiology under uncertainty would give some weight to Totalism (Greaves & Ord, 2017). While the Repugnant Conclusion is counterintuitive, alternatives to Totalism do not fare (much) better (Greaves, 2017; Zuber et al., 2021). At the same time, under uncertainty, we should give positive credences to other intuitions and views in population axiology. For example, some philosophers adopt 'the Asymmetry' and believe it is neither better nor worse to bring happy people into existence but bad (or wrong) to bring unhappy people into existence. Philosophers have struggled to make the Asymmetry work (McMahan, 2009) and intuitions behind the Asymmetry are mostly rejected by people outside of philosophy (Caviola et al., 2022; Spears, 2020). Still, under moral uncertainty, we should still assign a non-zero credence to the Asymmetry and, of course, to the intuition that we should avoid the Repugnant Conclusion. One proposal to account for those credences is to adopt Totalism but then add a 'critical range' *CR* (Greaves & Ord, 2017; MacAskill, 2022, pp. 186–187). The lower bound of *CR* is strictly greater than a 'zero level of wellbeing' for our population axiology proper. Any additional life with a wellbeing level below *CR* makes an outcome worse, even if such a level is above zero. If its wellbeing level is within the critical range *CR*, the additional life leaves an outcome equally as good. If its level is above *CR*, it makes an outcome better. *CR*'s levels can be adjusted depending on one's credences surrounding Totalism, the Asymmetry, and the Repugnant Conclusion.

Whether this approach to moral uncertainty is plausible is beyond my scope. But if it is, such an axiology would likely support a population heuristic for freedom that contains a Totalist element plus a neutral range. Such a heuristic would support long-term liberalism.[39]

The 'population axiology of freedom' is an unexplored field. And population ethics itself is riddled with disagreement and befuddlement. But long-term liberalism as a generic view appears somewhat robust across different views in population axiology and might come out well if we remain morally uncertain about population axiology.

7 Conclusions

In this article, I have explored the case for:

> *Long-term liberalism*: freedom should be a central goal in institutional design, public policy and social norm-setting; hereby we should often be particularly concerned with *ex ante* effects on long-term future distributions of freedom.

I explored the case in three steps.

First, I argued that liberals should be *long-term* liberals (section 3). Both current and future liberal freedom are non-specifically valuable, particularly when conceptualised as a social value. Sometimes, liberals should be particularly concerned with effects on long-term distributions of freedom because of their scale and because future levels of freedom might be very high.

Second, I argued that longtermists should be long-term *liberals* (section 4). Giving longtermism a 'liberal slant' makes it more plausible under conditions of empirical and moral uncertainty and helps capture some intuitions in intergenerational ethics.

---

[39] A plausible 'zero level' of freedom might already be positive, because, unlike measures of wellbeing, most suggested individual freedom measures do not allow negative values. Under moral uncertainty, we should then probably adopt an even higher level for the critical range.

Third, I argued that long-term liberalism can justify restricting freedom now to safeguard the freedom of future people, as seems plausible in cases like climate change mitigation (section 5). At the same time, it likely avoids excessive trade-offs: for both empirical and philosophical reasons, long-term and near-term freedom show significant convergence.

Finally, I suggested that long-term liberalism remains plausible as a generic view despite much disagreement and befuddlement in population ethics.

Throughout I also highlighted important practical implications, for example on longtermist institutional action, climate change, human extinction, and global catastrophic risks. If liberals want to get serious about securing 'the Blessings of Liberty to ourselves and our Posterity', they have their work cut out for them.

Bibliography

Abad, L. A., & Maurer, N. (2021). History Never Really Says Goodbye: A Critical Review of the Persistence Literature. *Journal of Historical Political Economy*, *1*(1), 31–68. https://doi.org/10.1561/115.00000002

Acemoglu, D., & Robinson, J. A. (2008). Persistence of Power, Elites, and Institutions. *American Economic Review*, *98*(1), 267–293. https://doi.org/10.1257/aer.98.1.267

Acemoglu, D., & Robinson, J. A. (2012). *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. Crown.

Araújo, R., & Koessler, L. (2021). *The Rise of the Constitutional Protection of Future Generations* (SSRN Scholarly Paper ID 3933683). Social Science Research Network. https://papers.ssrn.com/abstract=3933683

Arrhenius, G. (1999). An Impossibility Theorem in Population Axiology with Weak Ordering Assumptions. *Philosophical Studies*, *49*, 11–21.

Arrhenius, G. (2000). An Impossibility Theorem for Welfarist Axiologies. *Economics and Philosophy*, *16*(2), 247–266. https://doi.org/10.1017/s0266267100000249

Bader, R. M. (2018). Moralizing Liberty. In D. Sobel, P. Vallentyne, & S. Wall (Eds.), *Oxford Studies in Political Philosophy Volume 4* (p. 0). Oxford University Press. https://doi.org/10.1093/oso/9780198813972.003.0006

Barrett, J. (2020). Social Reform in a Complex World. *Journal of Ethics and Social Philosophy*, *17*(2), Article 2. https://jesp.org/index.php/jesp/article/view/900

Barrett, J., & Schmidt, A. T. (2023). Moral Uncertainty and Public Justification. *Philosophers' Imprint*, *forthcoming*.

Bavetta, S., Navarra, P., Maimone, D., & Patti, D. (2014). *Freedom and the Pursuit of Happiness: An Economic and Political Perspective*. Cambridge University Press.

Beckman, L. (2016). Power and future people's freedom: Intergenerational domination, climate change, and constitutionalism. *Journal of Political Power*, *9*(2), 289–307. https://doi.org/10.1080/2158379X.2016.1191159

Bengtson, A. (2019). On the Possibility (and Acceptability) of Paternalism Towards Future People. *Ethical Theory and Moral Practice*, *22*(1), 13–25. https://doi.org/10.1007/s10677-018-9969-4

Brighton, H., & Gigerenzer, G. (2012). How Heuristics Handle Uncertainty. In P. M. Todd, G. Gigerenzer, & A. R. Group (Eds.), *Ecological Rationality: Intelligence in the World* (pp. 33–60). Oxford University Press.

Browning, H., & Veit, W. (2023). Longtermism and Animals. In J. Barrett, D. Thorstad, & H. Greaves (Eds.), *Essays on Longtermism*. Oxford University Press.

Bundesverfassungsgericht. (2021). *Bundesverfassungsgericht—Presse—Verfassungsbeschwerden gegen das Klimaschutzgesetz teilweise erfolgreich*. https://www.bundesverfassungsgericht.de/SharedDocs/Pressemitteilungen/DE/2021/bvg21-031.html

Caplan, B. (2011). The totalitarian threat. In N. Bostrom & M. M. Cirkovic (Eds.), *Global Catastrophic Risks* (pp. 504–518). Oxford University Press.

Carter, I. (1999). *A measure of freedom*. Oxford University Press.

Carter, I. (2013). Distributing freedom over whole lives. In A. Gosseries & P. Vanderborght (Eds.), *Arguing about justice: Essays for Philippe Van Parijs* (pp. 135–143). Presses universitaires de Louvain.

Caviola, L., Althaus, D., Mogensen, A. L., & Goodwin, G. P. (2022). Population ethical intuitions. *Cognition*, *218*, 104941. https://doi.org/10.1016/j.cognition.2021.104941

Christian, B., & Griffiths. (2016). *Algorithms to Live By: The Computer Science of Human Decisions*. HarperCollins UK.

Coeckelbergh, M. (2021). *Green Leviathan or the Poetics of Political Liberty: Navigating Freedom in the Age of Climate Change and Artificial Intelligence*. Routledge. https://doi.org/10.4324/9781003159490

Cohen, G. A. (2011). Freedom and Money. In M. Otsuka (Ed.), *On the Currency of Egalitarian Justice, and Other Essays in Political Philosophy* (pp. 166–200). Princeton University Press.

Cowen, T. (2023, January 12). *Tyler Cowen on Effective Altruism*. https://www.youtube.com/watch?v=ZzV7ty1DW_c

Curran, E. (2023). Longtermism and the claims of future people. In J. Barrett, D. Thorstad, & H. Greaves (Eds.), *Essays on Longtermism*. Oxford University Press.

Esvelt, K. (2022). *Delay, Detect, Defend: Preparing for a Future in which Thousands Can Release New Pandemics* (Geneva Paper 29/22; Geneva Papers Research Series). Geneva Centre for Security Policy. https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22

Ferretti, M. P. (2016). Risk imposition and freedom. *Politics, Philosophy & Economics*, *15*(3), 261–279. https://doi.org/10.1177/1470594X15605437

Ferretti, M. P. (2023). Present Risks, Future Lives: Social Freedom and Environmental Sustainability Policies. *The Journal of Ethics*. https://doi.org/10.1007/s10892-022-09415-4

Fragnière, A. (2016). Ecological Limits and the Meaning of Freedom: A Defense of Liberty as Non-Domination. *De Ethica*, *3*(3), 33–49. https://doi.org/10.3384/de-ethica.2001-8819.163333

Garnett, M. (2018). Coercion: The Wrong and the Bad. *Ethics*, *128*(3), 545–573. https://doi.org/10.1086/695989

Giuliano, P., & Nunn, N. (2020). Understanding Cultural Persistence and Change. *The Review of Economic Studies*, *rdaa074*. https://doi.org/10.1093/restud/rdaa074

Gosseries, A. (2017). Generational Sovereignty. In I. González-Ricoy & A. Gosseries (Eds.), *Institutions for Future Generations* (pp. 98–116). Oxford University Press.

Greaves, H. (2017). Population Axiology. *Philosophy Compass*, *12*(11), e12442. https://doi.org/10.1111/phc3.12442

Greaves, H. (2020). *Evidence, cluelessness, and the long term*. https://forum.effectivealtruism.org/posts/LdZcit8zX89rofZf3/evidence-cluelessness-and-the-long-term-hilary-greaves

Greaves, H., & MacAskill, W. (2023). The Case for Strong Longtermism. In J. Barrett, D. Thorstad, & H. Greaves (Eds.), *Essays on Longtermism*. Oxford University Press.

Greaves, H., & Ord, T. (2017). Moral Uncertainty About Population Axiology. *Journal of Ethics and Social Philosophy*, *12*(2), 135–167. https://doi.org/10.26556/jesp.v12i2.223

Haller, M., & Hadler, M. (2004). Happiness as an Expression of Freedom and Self-Determination. In W. Glatzer, S. Von Below, & M. Stoffregen (Eds.), *Challenges for Quality of Life in the Contemporary World: Advances in quality-of-life studies, theory and research* (pp. 207–231). Springer Netherlands. https://doi.org/10.1007/978-1-4020-2903-5_14

Hannis, M. (2015). *Freedom and Environment: Autonomy, Human Flourishing and the Political Philosophy of Sustainability*. Routledge. https://doi.org/10.4324/9781315772769

Heal, G., & Millner, A. (2014). Uncertainty and Decision Making in Climate Change Economics. *Review of Environmental Economics and Policy*, *8*(1), 120–137. https://doi.org/10.1093/reep/ret023

Hees, M. van. (2012). *Legal Reductionism and Freedom*. Springer Science & Business Media.

Heikkinen, K. (2022). *Strong longtermism and the challenge from anti-aggregative moral views* [Global Priorities Institute Working Paper]. https://globalprioritiesinstitute.org/karri-heikkinen-strong-longtermism-and-the-challenge-from-anti-aggregative-moral-views/

Hilton, B. (2022). Preventing an AI-related catastrophe—Problem profile. *80,000 Hours*. https://80000hours.org/problem-profiles/artificial-intelligence/

Hilton, B., & McIntyre, B. (2022). How you can lower the risk of a catastrophic nuclear war. *80,000 Hours*. https://80000hours.org/problem-profiles/nuclear-security/

Hoffmann, J. (2022). *Democracy, Justice, & the Long Term – Designing Institutions for the Future* [Ph.D. dissertation]. University of Warwick.

Howell, R. T., Chenot, D., Hill, G., & Howell, C. J. (2011). Momentary Happiness: The Role of Psychological Need Satisfaction. *Journal of Happiness Studies*, *12*(1), 1–15. https://doi.org/10.1007/s10902-009-9166-1

Hurka, T. (1987). Why Value Autonomy? *Social Theory and Practice*, *13*(3), 361–382.

Intropi, P. (2022). Freedom's values: The good and the right. *Theoria*, *88*(6), 1144–1162. https://doi.org/10.1111/theo.12431

Karger, E., Rosenberg, J., Jacobs, Z., Hadshar, R., Gamin, K., Smith, T., McCaslin, T., Thomas, S., & Tetlock, P. E. (2023). *Forecasting Existential Risks—Evidence from a Long-Run Forecasting Tournament* (#1; FRI Working Paper). Forecasting Research Institute. https://forecastingresearch.org/xpt

Katz, C. (2017). Neorepublicanism and the Domination of Posterity. *Ethics, Policy & Environment*, *20*(3), 294–313. https://doi.org/10.1080/21550085.2017.1374034

Kosonen, C. (2023). *Tiny probabilities and the value of the far future* [GPI Working Paper No. 1-2023]. https://globalprioritiesinstitute.org/tiny-probabilities-and-the-value-of-the-far-future-petra-kosonen/

Kramer, M. H. (2003). *The Quality of Freedom*. Oxford University Press.

Kristjánsson, K. (1996). *Social Freedom: The Responsibility View*. Cambridge University Press.

Kumar, R. (2018). Risking Future Generations. *Ethical Theory and Moral Practice*, *21*(2), 245–257. https://doi.org/10.1007/s10677-018-9880-z

MacAskill, W. (2022). *What We Owe the Future*. Hachette UK.

MacAskill, W., Bykvist, K., & Ord, T. (2020). *Moral Uncertainty*. Oxford University Press.

Mahoney, J. (2000). Path Dependence in Historical Sociology. *Theory and Society*, *29*(4), 507–548.

McMahan, J. (2009). Asymmetries in the Morality of Causing People to Exist. In M. A. Roberts & D. T. Wasserman (Eds.), *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem* (pp. 49–68). Springer Netherlands. https://doi.org/10.1007/978-1-4020-5697-0_3

Mill, J. S. (1859). *On Liberty* (1979th ed.). Penguin.

Miller, D. (1983). Constraints on Freedom. *Ethics*, *94*(1), 66–86.

Mogensen, A. L. (2019). *The only ethical argument for positive δ?* [GPI Working Paper]. https://globalprioritiesinstitute.org/andreas-mogensen-the-only-ethical-argument-for-positive-delta-2/

Mogensen, A. L. (2020). Maximal Cluelessness. *The Philosophical Quarterly*, *71*(1), 141–162. https://doi.org/10.1093/pq/pqaa021

Mogensen, A. L., & Thorstad, D. (2022). Tough enough? Robust satisficing as a decision norm for long-term policy analysis. *Synthese*, *200*(1), 36. https://doi.org/10.1007/s11229-022-03566-5

Norman, R. J. (1987). *Free and equal: A philosophical examination of political values*. Oxford University Press.

Ord, T. (2020). *The Precipice: Existential Risk and the Future of Humanity*. Hachette UK.

Ord, T., Hillerbrand, R., & Sandberg, A. (2010). Probing the improbable: Methodological challenges for risks with low probabilities and high stakes. *Journal of Risk Research*, *13*(2), 191–205. https://doi.org/10.1080/13669870903126267

Parfit, D. (1997). Equality and priority. *Ratio*, *10*(3), 202–221.

Parijs, P. V. (1997). *Real Freedom for All: What (if Anything) Can Justify Capitalism?* Clarendon Press.

Pattanaik, P. K. (2018). Individual Freedom and Welfare Economics. *Journal of Quantitative Economics*, *16*(1), 1–12. https://doi.org/10.1007/s40953-017-0107-2

Pattanaik, P. K., & Xu, Y. (1990). On ranking opportunity sets in terms of freedom of choice. *Recherches Économiques de Louvain/Louvain Economic Review*, *56*(3/4), 383–390. https://doi.org/10.1017/S0770451800043955

Peterson, C., & Seligman, M. E. (1984). Causal explanations as a risk factor for depression: Theory and evidence. *Psychological Review*, *91*(3), 347–374.

Pettit, P. (2003). Agency-Freedom and Option-Freedom. *Journal of Theoretical Politics*, *15*(4), 387–403. https://doi.org/10.1177/0951692803154003

Pierson, P. (2000). Increasing Returns, Path Dependence, and the Study of Politics. *The American Political Science Review*, *94*(2), 251–267. https://doi.org/10.2307/2586011

Pinto, J. (2021). Freedom and ecological limits. *Critical Review of International Social and Political Philosophy*, *24*(5), 676–692. https://doi.org/10.1080/13698230.2019.1698147

Quoidbach, J., Gilbert, D. T., & Wilson, T. D. (2013). The End of History Illusion. *Science*, *339*(6115), 96–98.

Rawls, J. (1971). *A Theory of Justice*. Harvard University Press.

Scanlon, T. M. (1988). The Significance of Choice. In S. M. McMurrin (Ed.), *The Tanner Lectures on Human Values (Vol. 8, pp. 149-216)*. University of Utah Press.

Scheffler, S. (2013). *Death and the Afterlife*. Oxford University Press.

Scheffler, S. (2018). *Why Worry About Future Generations?* Oxford University Press.

Schmidt, A. T. (2014). *Freedom and its distribution* [DPhil]. University of Oxford.

Schmidt, A. T. (2015). Why animals have an interest in freedom. *Historical Social Research*, *40*(4), 92–109.

Schmidt, A. T. (2016). Abilities and the Sources of Unfreedom. *Ethics*, *127*(1), 179–207. https://doi.org/10.1086/687335

Schmidt, A. T. (2017). An unresolved problem: Freedom across lifetimes. *Philosophical Studies*, *174*(6), 1413–1438. https://doi.org/10.1007/s11098-016-0765-5

Schmidt, A. T. (2018a). Domination without Inequality? Mutual Domination, Republicanism, and Gun Control. *Philosophy & Public Affairs*, *46*(2), 175–206. https://doi.org/10.1111/papa.12119

Schmidt, A. T. (2018b). Persons or Property – Freedom and the Legal Status of Animals. *Journal of Moral Philosophy*, *15*(1), 20–45. https://doi.org/10.1163/17455243-46810070

Schmidt, A. T. (2022). Freedom in Political Philosophy. In *Oxford Research Encyclopedia of Politics*. Oxford University Press. https://doi.org/10.1093/acrefore/9780190228637.013.2022

Schmidt, A. T. (2023). *Do we have too much choice?* [Unpublished manuscript].

Schmidt, A. T., & Barrett, J. (2023). Longtermist Political Philosophy: An Agenda for Future Research. In J. Barrett, D. Thorstad, & H. Greaves (Eds.), *Essays on Longtermism*. Oxford University Press.

Schmidt, A. T., & Juijn, D. (2023). Economic inequality and the long-term future. *Politics, Philosophy & Economics*, *online first*, 1470594X231178502. https://doi.org/10.1177/1470594X231178502

Schwartz, B. (2009). *The Paradox of Choice*. Harper Collins.

Sen, A. (1988). Freedom of choice: Concept and content. *European Economic Review*, *32*(2), 269–294.

Sen, A. (1999). *Development as freedom* (1st. ed.). Knopf.

Sevilla, J. (2022). *Persistence: A critical review* (What We Owe the Future - Supplementary Material). https://drive.google.com/file/d/1CiO8Rg1Jw9QHZToTvQlFOJV8dLD62r7u/view

Shnayderman, R. (2013). Social Freedom, Moral Responsibility, Actions and Omissions. *Philosophical Quarterly*, *63*(253), 716–739.

Shue, H. (1980). *Basic Rights: Subsistence, Affluence, and U.S. Foreign Policy*. Princeton University Press.

Shulman, C., & Thornley, E. (2023). Tradeoffs between longtermism and other social metrics: Is longtermism relevant to existential risk in practice? In J. Barrett, D. Thorstad, & H. Greaves (Eds.), *Essays on Longtermism*. Oxford University Press.

Skagen Ekeli, K. (2006). The principle of liberty and legal representation of posterity. *Res Publica*, *12*(4), 385–409. https://doi.org/10.1007/s11158-006-9021-8

Smith, P. T. (2013). The Intergenerational Storm: Dilemma or Domination. *Philosophy and Public Issues - Filosofia E Questioni Pubbliche*, *3*(1).

Spangenberg, J. H. (2018). Looking into the future: Finding suitable models and scenarios. In M. Düwell, G. Bos, & N. van Steenbergen (Eds.), *Towards the Ethics of a Green Future: The Theory and Practice of Human Rights for Future People*. Routledge.

Spears, D. (2020). The Asymmetry of population ethics: Experimental social choice and dual-process moral reasoning. *Economics & Philosophy*, *36*(3), 435–454. https://doi.org/10.1017/S0266267119000221

Spencer, H. (1873). *Social Statics; Or, The Conditions Essential to Human Happiness Specified, & the First of Them Developed*. D. Appleton & Company.

Steiner, H. (1983). How Free: Computing Personal Liberty. *Royal Institute of Philosophy Supplements*, *15*, 73–89.

Steiner, H. (1987). Capitalism, Justice and Equal Starts. *Social Philosophy and Policy*, *5*(01), 49–71.

Steiner, H. (1994). *An Essay on Rights*. Wiley.

Stern, N. (2006). *The Economics of Climate Change: The Stern Review*. Cambridge University Press.

Tarsney, C. (2023). The epistemic challenge to longtermism. *Synthese*, *201*(6), 1–37.

Tarsney, C., & Thomas, T. (2020). *Non-additive axiologies in large worlds* [GPI Working Paper]. https://globalprioritiesinstitute.org/christian-tarsney-and-teruji-thomas-non-additive-axiologies-in-large-worlds/

Tetlock, P. E., & Gardner, D. (2015). *Superforecasting: The Art and Science of Prediction*. Crown/Archetype.

Thomas, T. (2022). The Asymmetry, Uncertainty and the Long Term. *Philosophy and Phenomenological Research*, *forthcoming*.

Thorstad, D. (2023). Existential risk pessimism and the time of perils. *Philosophy and Public Affairs*, *forthcoming*.

Tonn, B. E. (2009). Obligations to future generations and acceptable risks of human extinction. *Futures*, *41*(7), 427–435. https://doi.org/10.1016/j.futures.2009.01.009

Tonn, B., & Stiefel, D. (2013). Evaluating Methods for Estimating Existential Risks. *Risk Analysis*, *33*(10), 1772–1787. https://doi.org/10.1111/risa.12039

Unruh, C. (2023). Against a moral duty to make the future go best. In J. Barrett, D. Thorstad, & H. Greaves (Eds.), *Essays on Longtermism*. Oxford University Press.

Vallier, K. (2018). Public Justification. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2018). https://plato.stanford.edu/archives/spr2018/entries/justification-public/

Veenhoven, R. (2000). Freedom and happiness: A comparative study in forty-four nations in the early 1990s. In E. Diener & E. M. Suh (Eds.), *Culture and Subjective Well-being* (pp. 257–288). MIT Press.

Vercelli, A. (1998). Sustainable Development and the Freedom of Future Generations. In G. Chichilnisky, G. Heal, & A. Vercelli (Eds.), *Sustainability: Dynamics and Uncertainty* (pp. 171–187). Springer Netherlands. https://doi.org/10.1007/978-94-011-4892-4_9

Wallack, M. (2004). The minimum irreversible harm principle—Green inter-generational liberalism. In Y. Levy & M. Wissenburg (Eds.), *Liberal Democracy and Environmentalism: The End of Environmentalism?* (pp. 167–178). Taylor & Francis Group. http://ebookcentral.proquest.com/lib/rug/detail.action?docID=200474

Weitzman, M. L. (1998). Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate. *Journal of Environmental Economics and Management*, *36*(3), 201–208. https://doi.org/10.1006/jeem.1998.1052

Zuber, S., Venkatesh, N., Tännsjö, T., Tarsney, C., Stefánsson, H. O., Steele, K., Spears, D., Sebo, J., Pivato, M., Ord, T., Ng, Y.-K., Masny, M., MacAskill, W., Lawson, N., Kuruc, K., Hutchinson, M., Gustafsson, J. E., Greaves, H., Forsberg, L., … Asheim, G. B. (2021). What Should We Agree on about the Repugnant Conclusion? *Utilitas*, *33*(4), 379–383. https://doi.org/10.1017/S095382082100011X