The Longtermism Beliefs Scale:

Measuring Lay Beliefs for Protecting Humanity's Longterm Future

Stylianos Syropoulos[1,2*]

Kyle Fiore Law[3*]

Gordon Kraft-Todd[1]

Liane Young[1]

[1]Department of Psychology and Neuroscience, Boston College, Boston, MA

[2]The Schiller Institute for Integrated Science and Society, Boston College, Boston, MA

[3]Department of Psychology, University at Albany, State University of New York, Albany, NY

* Denotes joint first authorship

Abstract

Across a series of six high-powered studies (*N*=4,878), we develop and validate the Longtermism Beliefs Scale (LBS) to measure alignment with the "longtermism" philosophy, which advocates for the welfare of distant future generations. The findings highlight the pivotal role of the LBS in predicting attitudes and behaviors aimed at protecting future generations from extinction threat (Studies 1-5). The LBS reveals significant associations with heightened concern for human extinction threats and negative appraisals of future extinction, even when contemplating distant hypothetical futures (Studies 1-3). Notably, higher longtermism beliefs are linked to perceptions that future threats can be resolved, support for policies seeking to protect future generations of people, and a profound sense of responsibility for the long-term survival and prosperity of humanity (Study 4a). The LBS also proves instrumental in predicting various future-oriented attitudes, patterns of prosociality, and longtermism-related behavioral intentions (Studies 4b-5), and is capable of meaningfully distinguishing between self-identified longtermists and non-longtermists. Overall, this research contributes to the field by introducing and validating the LBS as a critical tool for assessing longtermism beliefs and understanding their role in addressing existential threats to the long-term future of humanity.

Longtermism Beliefs: Measuring Lay Interest of Protecting Humanity's Long-term Future from Extinction

The actions we take as a species in the present day may hold the power to either secure a brighter future for humankind or have a hand in its ultimate undoing (Macaskill, 2022a). The 2022 release of ChatGPT – a sophisticated artificial intelligence (AI) chatbot – took the world by storm, bringing to the forefront of public consciousness the realization that AI has the burgeoning potential to become rapidly more capable and advanced, while at the same time striking fear in many that AI might eventually put humanity's future in jeopardy (see Heuvel, 2023). Such AI-related apprehension is being taken more seriously now than ever before. In 2023, the US President met with CEOs of leading technology firms (e.g., Google, Microsoft, OpenAI) to insist that their companies substantially reduce any threat AI may pose to people, national security and society more broadly (Bose, & Shepardson, 2023). Potential extinction threats to humanity are not a novel topic in public discourse, however, and AI is not a solitary agent wielding the capacity to levy doom upon the human race. Indeed, numerous sources of global catastrophic risk have received coverage in the popular press over the years, including pandemic disease, climate change, nuclear war, and the ongoing depletion of natural resources (e.g., Hunter, & Hewson, 2020; McLamb, 2022).

These examples echo a general consensus around two critical points noted across the fields of philosophy (MacAskill, 2022a; Ord, 2021; Singer, 2015), psychology (Caviola et al., 2021; Syropoulos, & Markowitz, 2022), and the natural sciences (e.g., Blaser, 2018; Taylor et al., 2012): (1) that human extinction threats are real and will have devastating consequences for humanity in the future if not properly mitigated, and (2) that individual and collective present-day action *can* serve to mitigate the impact of many of these threats. But how motivated are people to take such action? The answer is complex. Not all people prioritize the welfare of future

generations to the same degree. Whereas many people might agree that future people matter, others may extend consideration only to the welfare of those who are alive today. Similarly, some people might even agree that future people are worthy of consideration yet feel no sense of efficacy in being able to positively impact human welfare into the future. Such individual variation in the extent to which people prioritize and feel efficacy to impact future generations will likely have pragmatic consequences for whether sufficient action is ultimately taken to protect our species from premature extinction (MacAskill, 2022a).

The issue of prioritizing the welfare of temporally distant future humans has been explored deeply in philosophical thought, and most notably in the writings related to the increasingly popular philosophical and social movement, *longtermism* (MacAskill, 2022a; Ord, 2021). However, at present, there exists no empirical metric by which individual variation regarding the prioritization of future humans can be measured. We develop and validate the Longtermism Beliefs Scale (LBS) to address this gap in knowledge, critically demonstrating the novel metric's utility in predicting a host of crucial pro-future attitudes and behaviors.

What is Longtermism? Defining and Distinguishing the Construct of Longtermism Beliefs

The longtermism philosophy and social movement is grounded on three foundation principles: (1) people ought to acknowledge that humanity has the potential to last into the distant future and keep future people in mind (i.e., consideration); (2) people today have a moral responsibility to protect the welfare of future humans through individual and collective actions geared towards mitigating long-term existential threats (i.e., responsibility); and (3) the actions people take today can positively impact future people and assuage catastrophic risk associated with some of the greatest challenges humanity faces (i.e., efficacy; MacAskill, 2022a; Ord, 2021). Longtermism is closely related to, and has its origins in the effective altruism (EA)

movement, which applies scientific evidence and utilitarian reasoning to improve global welfare through philanthropy (see Caviola et al., 2021; Singer, 2015, 2016).

Compared to effective altruism, longtermism places greater emphasis on extending equal regard for others across *temporal distance*, whereas effective altruism places greater emphasis on extending equal regard for others across *social distance* (Jones & Rachlin, 2006). For example, effective altruists hold the lives of family members and friends at the same level of moral importance as the lives of distant strangers. Longtermists, on the other hand, go a step further in holding the lives of present-day people at the same level of moral importance as the lives of people born in the very distant future. In practice, longtermists choose to live on less money and pledge large portions of income to future-oriented organizations like Open Philanthropy and the Future Fund, support pro-future public policies (e.g., pro-environmental policies), and some go as far as changing careers to increase their impact on future generations (MacAskill, 2022b).

In developing a systematic understanding of individual variability in lay beliefs associated with the longtermism philosophy, we argue that it is critical to measure the extent to which people align with all three of the philosophy's core principles. Not only are these three principles central to MacAskill's (2022a) conceptualization of longtermism, but empirical research underscores that consideration (Crimston et al., 2016, 2018; Graham et al., 2017), responsibility (Syropoulos, & Markowitz, 2021, 2022) and efficacy (Bradley et al., 2020; Hornsey et al., 2021) reliably predict and guide actual moral and future-oriented attitudes and behaviors, such as altruism, support for environmental and humanitarian public policies, and monetary donations to campaigns and charitable efforts. The approach we take when conceptualizing longtermism beliefs is sensitive to each of these core principles, such that one must align with all three principles with respect to future generations in the both the near and distant future to be considered "longtermist".

We hypothesize that our conceptualization of longtermism beliefs will afford the power to predict a variety of future-oriented attitudes and behavioral intentions. Consequently, we argue that developing a metric to capture the construct of longtermism beliefs will be critical in generating novel theoretical insights, serving as a launching point for future scientific inquiry into the antecedents and consequences of holding longtermism beliefs and related attitudes and behaviors. Ultimately, we hope such a metric will yield tangible pragmatic benefits for the improvement of long-term collective welfare going forward, informing and guiding the development of future interventions towards this aim. Perhaps most straightforwardly, such an instrument will also illuminate how many people espouse longtermist beliefs to a high degree in our society.

We assert that longtermism beliefs substantiate a unique construct that captures how people prioritize the welfare of *future* humans. Existing metrics assessing the prioritization of welfare across distance capture how such beliefs wax and wane across the *social* dimension of distance quite well (e.g., Caviola et al., 2022; Kahane et al., 2018), yet largely ignore the *temporal* dimension. Because many potential extinction threats facing humanity may not affect those who are currently living directly, but instead may spell disaster only for future generations (MacAskill, 2022a; Ord, 2020), we argue that it is imperative that the influence of *time* on the moral standing people assign to future generations, as well the prioritization people place on helping and protecting future generations, receive more attention in the cognitive, social, and moral psychological literatures. Although there are existing metrics that do address future-oriented attitudes and behaviors (Shrum et al., 2021; Strathman et al., 1994; Zaval et al., 2015), we assert that these measures fall short of isolating impartial concern for and intentions to positively impact the welfare of future generations. Below we briefly distinguish longtermism beliefs from related constructs which have received recent attention in the literature.

Distinguishing Longtermism from Related Constructs

First, the expansive altruism and effectiveness focus scales capture attitudes supporting the (closely related) EA movement (Caviola et al., 2022). Specifically, these scales measure the extent to which people see the welfare of socially distant humans as worthy of prioritization through philanthropy and personal sacrifice—attitudes and behaviors aligned with EA philosophy—rather than capturing the prioritization of temporally distant humans. Some effective altruists also identify as longtermists (Caviola et al., 2022), and there exists overlap in the principles of these two philosophical perspectives (MacAskill, 2022a; Singer, 2015, 2016). We therefore anticipate that scores on the expansive altruism and effectiveness focus scales may predict longtermism beliefs, although these scales will not measure longtermism beliefs directly, as they do not capture the prioritization of welfare across time.

Another construct related to longtermism is utilitarianism as measured by the Oxford Utilitarianism Scale (Kahane et al., 2018), which measures two separate facets of utilitarianism: *impartial beneficence* (i.e., impartiality across spans of social distance and self-sacrifice for the greater good of humanity) and *instrumental harm* (i.e., attitudes which permit harm to others when that harm promotes the greater good of humanity). Because longtermism does not require committing harm to present-day humans to ensure the welfare of future humans, it is unlikely that scores on the instrumental harm subscale will be associated with longtermism beliefs. The impartial beneficence subscale, however, is likely related to longtermism beliefs, as it taps a similar conceptual space as the EA measures discussed above. Nonetheless, this scale similarly puts primary emphasis on capturing the prioritization of welfare across social but not temporal distance. Thus, we argue that longtermism beliefs are a unique construct with respect to impartial beneficence and will likely predict outcomes related to prioritizing the welfare of future humans that impartial beneficence will not.

Finally, while existing measures, such as the Legacy Motives Scale (LMS; Shrum et al., 2021; Zaval et al., 2015) and the Consideration of Future Consequences scale (CFC; Strathman et al., 1994), capture future-oriented attitudes, we argue that they do not directly assess longtermism beliefs. The LMS primarily focuses on motivations related to preserving one's own reputation and being remembered favorably, which diverges from the philosophy of longtermism (MacAskill, 2022a; Ord, 2020). Similarly, the CFC examines general considerations of future consequences but does not specifically address attitudes towards impacting future generations or personal moral responsibility for promoting future welfare. Consequently, there is a need for a psychometric instrument that directly measures longtermism beliefs, identifies longtermists, predicts pro-future attitudes and behaviors, and informs efforts to engage individuals in the longtermism social movement.

In addition to capturing the central features of the longtermism philosophy and distinguishing the construct from other existing measures, there are some peripherally related phenomena that will likely be associated with longtermism beliefs that are worthy of consideration in the present research. First, we argue that longtermism entails the extent to which people acknowledge the existence of extinction threats and whether they negatively evaluate the consequences of human extinction. That is, in order to prioritize the protection of humanity's future from catastrophic threats, one must first acknowledge that such threats exist and that human extinction would not be welcomed. As such, in the present research we measured the following variables and examined their associations with longtermism beliefs: concern for a list of potential threats that could pose global catastrophic risk to the future of humanity (e.g., climate change, unaligned AI, nuclear war), the extent to which participants believe each of the threats could cause human extinction, and whether participants feel that human extinction would be a negative (vs. good or neutral) event.

Second, we posit that longtermism beliefs are linked to the depth of time representation in individuals' future thinking. When envisioning the future and its inhabitants, some individuals may only spontaneously contemplate up to 100 years ahead (a shallow conceptualization of time), while others may naturally think millions of years ahead (a deep conceptualization of time). Additionally, complexity arises with potential variability in individuals' spontaneous future time consideration, as longtermism belief strength may diminish with greater temporal distance. This is due to individuals' limited capacity to vividly imagine distant future scenarios, which further declines as the temporal distance increases (Atance & O'Neill, 2001; Tamir & Mitchell, 2011; Zhao et al., 2020). Given that assuming responsibility for and perceiving efficacy in changing the future requires the ability to mentally represent the future and its inhabitants, the LBS is designed to capture whether longtermism beliefs progressively weaken as individuals contemplate prioritizing future generations in increasingly distant futures.

In addition to difficulties in constructing vivid imagery across temporal distance, as social distance increases, individuals tend to extend less moral consideration (Crimston et al., 2016), experience decreased empathy (Fowler et al., 2021), and offer less help to others (Everett et al., 2018; Law et al., 2022; McManus et al., 2021). Interestingly, temporal distance is processed in a similar manner as social distance at the cognitive and neural level (Hill et al., 2017; Soutschek et al., 2016; Tuen et al., 2023), providing further support for the prediction that longtermism beliefs are likely stronger when considering proximal rather than distant future time points for humanity. Given the strong emphasis of longtermism philosophy on extending equal regard to future individuals across different time horizons (MacAskill et al., 2022a), we explicitly incorporate depth of time into the scale by measuring responses to each item with respect to a range of future time points.

Third, we argue that longtermism beliefs also entail whether people believe that there will still be a human race to protect in the distant future. That is, we ask whether people believe that humanity has already reached the end of its lifespan. The longtermism philosophy (MacAskill, 2022a, 2022b; Ord, 2020) proposes that humanity may still be in its infancy if preventable extinction threats are properly mitigated. As such, we measure associations of estimated likelihoods that a human extinction event will occur in the next 100 years with longtermism beliefs.

Finally, to establish the predictive validity of the LBS we explore how it is associated with various attitudes, beliefs and behaviors related to protecting the future of humanity from existential harm. Notably, we address associations between scores on the LBS and: belief in the feasibility of solutions to major extinction threats; a sense of responsibility towards ensuring better lives for future generations (at the individual, national, international, and universal levels); support for pro-future public policies; concern for the legacy one leaves behind after death;, consideration of the future consequences of present actions; and a number of longtermism-related behavioral intentions (e.g., intentions to donate one's money to pro-future charities, intentions to volunteer one's time to pro-future causes, intentions to change one's career to have a larger impact on humanity's future; adapted from Caviola et al., 2022). We contend that demonstrating correlations between the LBS and these outcomes would not only support the scale's predictive validity but also highlight its practical usefulness as a tool for identifying individuals who may be more open to potential interventions that improve humanity's long-term survival.

Overview of Current Studies

To capture the construct of longtermism beliefs as expressed in previous literature, in Study 1, we generated a list of 20 items using language from popular books on longtermism

(e.g., What we Owe the Future; The Precipice), and assessed the validity and reliability of these items. In Study 2, we added 11 more items based on AI-generated input (using ChatGPT-3.5) to the 17 items that we retained from Study 1 after conducting validity and reliability assessments. Conducting an exploratory factor analysis (EFA) on these 28 items, we found that a single-factor solution was preferable and retained 7 items with high factor loadings (>.75), high reliability (>.90) and good convergent validity (r > .80) with the original 28 items. The resulting 7-item scale captures the three components of the longtermism philosophy: consideration (e.g., "We should always have in view not only the present but also the coming generations."), responsibility (e.g., "Positively influencing the long term future is a key moral priority of our time."), and efficacy (e.g., "There are things we can do to steer the long term future to a better course.").

In Study 3, we formally validated the 7-item LBS using confirmatory factor analysis (CFA) and also manipulated the distance of future time points participants considered when responding to the scale—both within and between subjects—both to increase the variability in responses, and also to address whether endorsement of the items decreases with increasing temporal distance. In Studies 4a, 4b, and 5 we again confirmed the validity and reliability of the measure, and also tested multiple modifications to the scale's instructions and the delivery method of the items, finding that presenting each of the 7 items at each future time point simultaneously is the optimal procedure for measuring the construct. All data files, materials and code for the studies is available on the Open Science Framework (OSF): https://osf.io/bhzmp/?view_only=8e32b2e19ed64443a79ed343ccddf2bd.

Table 1

Sample Information for All Studies

Characteristic	Study 1	Study 2	Study 3	Study 4a	Study 4b	Study 5
N	397	403	903	1098	542	1535

N_{Male}	195	194	441	543	268	756
N_{Female}	189	195	437	535	260	736
N_{White}	294	317	695	844	412	1137
$N_{Black/African\ American}$	45	32	97	111	70	160
$N_{Asian/Asian\ American}$	32	29	62	88	32	99
$N_{Democrat}$	193	209	285	549	229	739
$N_{Republican}$	68	61	155	197	118	307
$N_{Independent}$	127	116	285	318	179	431
M_{age} (SD_{age})	37.63 (13.37)	42.18 (15.60)	41.48 (14.07)	40.56 (13.66)	40.09 (13.74)	39.22 (14.31)
Pre-registered	No	No	Yes	Yes	Yes	Yes

Study 1

Using items based on popular readings (e.g., What We Owe the Future; The Precipice), here we examine the first version of the LBS. We also examined its validity by testing its association with measures capturing concern for major extinction threats to humanity's long-term future. Study 1 was not pre-registered.

Methods

Participants

Data collection was operated on Prolific. We recruited a total of 400 participants. After removing participants who had a duplicate IP address, 397 participants remained in the sample. The study lasted approximately 10 minutes, and participants received $2.00 for their participation.

Materials and Procedure

We presented all measures in the study to participants in randomized order. First, we asked participants whether they donated part of their income to any charity/charitable organization. If they responded yes, we asked them to list these organizations, and the amount they donate (in total). Other questions included whether they were familiar with the philosophies of effective altruism and longtermism (as separate questions). If they answered yes (to either),

we asked them to indicate their level of familiarity, and whether they identified as an effective altruist/longtermist.

We measured concern for extinction threats to humanity (unchecked artificial intelligence, climate change, manmade pandemics, natural risks such as volcanic eruptions, nuclear war, and other risks) with six items (1-7 scale, 1 = not at all concerned, 7 = extremely concerned). We assessed longtermism beliefs with 20 items which the authors generated based on influential books on the subject (i.e., "What We Owe The Future" and "The Precipice"). Most items (1-7 scale, with 1 = strongly disagree, 7 = strongly agree) were taken directly, in some cases with slight modifications, from these works (see the Supplementary Materials for the full list of items). Conservative political ideology (a = .89) was captured with two items, both on 7-point Likert-type scales: "Politically, how conservative are you in terms of social (item 1)/economic (item 2) issues?"

Results

We conducted an exploratory factor analysis on the 20 items capturing longtermism beliefs with a maximum likelihood extraction and an oblique rotation using SAS's proc factor command. A 3-factor solution was recommended (see Table S1 in Supplementary Materials), although 1 item failed to load onto any factor, and several items cross-loaded onto 2 factors. After removing the non-loading item, a 1-factor solution was outputted, with an additional item failing to load. After removing this item and rerunning the analyses, one more item failed to load. Removing this final item left a total of 17 items. In this solution, the three lowest loadings belonged to reverse-coded items. However, we chose to retain these items in the model (See Table S2 in Supplementary Materials). These remaining 17 items were highly reliable (a = 0.94), with an average score above the scale midpoint (M = 5.61, SD = 0.82).

Do Longtermism Beliefs Relate to Concern for Extinction Threats to Humanity?

Longtermism beliefs were related solely to increased concern for unaligned AI (r=.14, p=.011), and climate change (r = .20, p = .001), but not any of the other threats (see Table S3 in Supplementary Materials).

Discussion

We evaluated a starting set of items generated from direct writings on longtermism that captured people's endorsement of the longtermism philosophy. The resulting 17-item solution was reliable, but when it came to increased concern for extinction threats to humanity, we found that longtermism was only related to concern for climate change and unaligned AI. Finally, supporting its validity, those who identified as longtermists (n = 52, M = 5.85, SD = 0.73) also scored higher on the inventory (t(395) = 2.31, p = .022, d = .36), compared to those who did not (n = 345, M = 5.57, SD = 0.83).

Study 2

In our second study, by incorporating items generated from ChatGPT-3.5, we sought to test the structure of the scale in a more unbiased fashion, as they were not directly generated by the research team. In total, 11 new items were included, and the validity and reliability of the scale were re-evaluated. This study was not pre-registered.

Methods

Participants

Data collection was operated on Prolific. We recruited a total of 403 participants. No participants had duplicate IP addresses, but one participant missed our attention check. Thus 402 participants remained in the sample. The study lasted approximately 10 minutes, and participants received $2.00 for their participation.

Measures

Longtermism beliefs. The 17 items from Study 1 were retained in Study 2. To generate and test additional items we used ChatGPT-3.5. We provided it with the prompt: "Generate 20 items that measure longtermism philosophy." From these 20 items, 11 were retained (these are highlighted in the Supplementary Materials), as the rest were too similar to existing items or not as relevant. These items, together with the 17 items from Study 1, were slightly modified to reduce their complexity. Responses were captured on 7-point Likert scales, ranging from 1 = Strongly disagree to 7 = Strongly agree.

Extinction threats. We included several measures capturing concerns for extinction threats. Seven items captured beliefs about how humanity's extinction would be a negative event. Six items captured beliefs about how humanity's extinction would be a good/neutral event. Both measures were captured on 7-point Likert scales, ranging from 1 = Strongly disagree to 7 = Strongly agree.

Participants indicated whether, if humanity went extinct in the next 100 years, the cause of humanity's extinction would be: climate change, Artificial General Intelligence, nuclear war, global pandemic (naturally caused or caused by humans), meteor or asteroid impact, or volcanoes. Responses were captured on 6-point Likert scales, ranging from 1 = Extremely unlikely to 6 = Extremely likely. Participants also indicated how personally concerned they were about the six different extinction threats. They did so by ranking these threats from highest (= 6) to lowest (= 1).

Participants also responded to two questions: (1) "What percent of global income per year should humanity (e.g. all nations) spend with the primary purpose of reducing the risk of extinction?" and (2) "How likely do you personally think it is that humanity will go extinct (no human being will ever live anymore) in the next 100 years?". For both, participants provided a percentage ranging from 0% to 100%.

Demographic variables. All demographic variables were identical to Study 1. Conservative political ideology was once again highly reliable (a = .88).

Results

An unrestricted exploratory factor analysis identical in its parameters to that of Study 1 suggested a 5-factor solution (see Table S7 in the Supplementary Materials). However, most of the variance was explained by 1 factor (eigen value = 31.71) with the second factor explaining considerably less (eigen value = 2.51). No meaningful pattern of items emerged across factors (see Supplementary Materials). Considering this, and the fact most of the variance was explained by a single factor, we reran the analyses restricting results to produce a 1-factor solution (see Table 2). We observed that all items loaded significantly, with some items (mostly the reverse-coded items) loading weakly. When looking at the intercorrelations between all items we found that the items that loaded the strongest were also the ones that correlated with all other items the strongest. Considering this, we ran analyses for 2 versions of the scale, one with only the 7 (M = 5.97, SD = 0.83, a = .92) items that loaded the strongest on the scale (βs > .75), and one with the full list of items (M = 5.75, SD = 0.76, a = .96). Our goal was to determine whether we can meaningfully use 7 items rather than the full battery. Overall, analyses were highly similar across the two versions of the scale which prompted us to retain the 7-item solution in future studies. Further, a confirmatory factor analysis was estimated for both models to compare their model fit. For the 7-item model: $\chi^2(14) = 44.14$, $p < .001$, CFI = .982, RMSEA = .073, SRMR = .024; for the 28-item model: $\chi^2(350) = 1150.55$, $p < .001$, CFI = .880, RMSEA = .075, SRMR = .052. A chi-square difference test ($\Delta\chi^2(336) = 1106.40$, $p < .001$) suggests that the more parsimonious model is indeed a better fit to the data.

Table 2

Restricted Exploratory Factor Analysis, Bolded Items Signify the Seven Best-Loading Items

Item	Loading
We should act wisely because what we do today will influence an untold number of people.	**0.82**
It is important to consider the long-term consequences of our actions and decisions.	**0.80**
Intergenerational cooperation is important for addressing long-term challenges.	**0.80**
It is important that we reduce existential risks and promote sustainable development goals.	**0.79**
We should always have in view not only the present but also the coming generations.	**0.78**
There are things we can do to steer the long term future to a better course.	**0.76**
Positively influencing the long term future is a key moral priority of our time.	**0.75**
Research to mitigate future global catastrophic risks is important.	0.74
I care about the impacts of our actions upon the long term future.	0.73
We really can make a difference to the world that future generations will inhabit.	0.73
The future where everyone can live a joyous life is impacted by our decisions today.	0.73
Long-term investments and policies can have a positive impact on future outcomes.	0.72
Considering the long-term future of humanity is important when making personal and professional decisions.	0.71
We should take seriously how long the future can be and how high the stakes are in influencing it.	0.70
Addressing long-term challenges requires collaboration and cooperation between different groups and sectors.	0.70
We should put future generations interests as if they are our own.	0.69
We should reorient our moral priorities towards the vast future of humanity.	0.68
Investing in scientific and technological research is important for the long-term future of humanity.	0.66
We can increase the chance of a wonderful future by increasing the values that positively guide society.	0.66
Our most important role as a generation is how we shape the future of others.	0.65
It is important to prioritize long-term solutions over short-term fixes.	0.64
We can create a flourishing and long lasting society where everyone's lives are better than the lives we live today.	0.64
Our own generation is but one page in a much longer story.	0.60
Addressing long-term problems requires a systemic, holistic approach.	0.49
Short-term thinking and decision-making are a problem to addressing humanity's issues.	0.44
Even if we were to care about future generations, there is not enough we can do to shape their lives. (R)	0.41
It doesn't matter if future people come into being or not. (R)	0.38
Speculations about what might or might not happen in the future are only distraction from the issues we face today. (R)	0.37

Note. R marks reverse-coded items.

Do Longtermism Beliefs Relate to Concern for Extinction Threats to Humanity?

Those who scored higher on the longtermism beliefs scale also reported higher levels of agreement with statements about how human extinction would be bad. Results were consistent for all statements, except for the one stating that life would no longer be meaningful if humanity ended. Similarly, those who scored higher on longtermism, were less likely to agree with statements about why human extinction would be good or neutral. These associations varied in their magnitude depending on the specific item.

From the potential extinction threats to humanity's long-term future, significant associations were observed primarily for climate change and global pandemics. Similarly, when asked which of these extinction risks they considered the most important, significant associations

were observed for climate change, an unaligned AI, and meteor or asteroid impact (although this correlation was very weak).

Table 3

Means, Standard Deviations, Reliability and Correlations Between Outcomes and the Two Versions of the Longtermism Beliefs Scale

Measure	M	SD	7-item LT Scale	28-item LT Scale
Human extinction (i.e., the demise of all humans on earth) would be bad …				
1. …because it would cause harm and pain to the humans that died from the extinction event	5.59	1.47	0.38***	0.39***
2. …because it would prevent future people from having positive lives	5.39	1.51	0.34***	0.38***
3. …because it would mean the loss of all human progress (cultural progress, technological progress, and other improvements to civilization)	5.56	1.45	0.37***	0.39***
4. …because we would no longer have any chance of creating an ideal future society (utopia)	4.66	1.78	0.24***	0.27***
5. …because we would have failed in our duty to protect generations of humans that will exist in the future	5.57	1.46	0.46***	0.49***
6. …because we would have failed in our duty to past generations to preserve and advance the human species	5.35	1.53	0.39***	0.42***
7. …because my life would no longer be meaningful if humanity ended	4.05	1.98	0.10	0.09
Average of all reasons (*a* = .86)	5.17	1.16	0.43***	0.46***
Human extinction (i.e., the demise of all humans on earth) would be good or neutral…				
1. …because nothing matters anyway.	2.36	1.61	-0.27***	-0.32***
2. …because we are doomed anyway (there's nothing we can do to prevent human extinction)	2.59	1.63	-0.30***	-0.37***
3. …because it would mean the loss of all human progress (cultural progress, technological progress, and other improvements to civilization).	2.77	1.77	-0.18**	-0.24**
4. …because it would prevent further suffering for humans.	3.32	1.90	-0.11*	-0.16**
5. …because it would prevent humans causing further suffering for non-human animals.	3.75	1.99	0.07	0.02
6. …because it would prevent humans from further destroying the natural environment.	3.87	1.99	0.06	0.02
Average of all reasons (*a* = .85)	3.11	1.39	-0.14**	-0.21**
If humanity went extinct in the next 100 years, which of the following would be the most likely causes?				
1. Climate change	4.11	1.70	0.25***	0.27***
2. Artificial General Intelligence	2.97	1.50	-0.08	-0.08
3. Nuclear war	4.66	1.19	0.10	0.11*
4. Global pandemic (naturally caused or caused by humans)	4.12	1.36	0.13*	0.15**
5. Meteor or asteroid impact	3.23	1.42	0.05	0.06
6. Super-Volcano	2.50	1.27	0.05	0.04
From these aforementioned extinction risks, how would you rank them in terms of how much you personally feel concerned about them?				

1.	Climate change	4.40	1.64	0.30***	0.32***
2.	Artificial General Intelligence	2.67	1.51	-0.20***	-0.19***
3.	Nuclear war	5.01	1.15	-0.07	-0.07
4.	Global pandemic (naturally caused or caused by humans)	4.24	1.10	0.07	0.07
5.	Meteor or asteroid impact	2.73	1.23	-0.11*	-0.12*
6.	Super-Volcano	1.95	1.07	-0.05	-0.08
How likely do you personally think it is that humanity will go extinct (no human being will ever live anymore) in the next 100 years?		17.81	24.82	-0.09	-0.12*
What percent of global income per year should humanity (e.g. all nations) spend with the primary purpose of reducing the risk of extinction?		19.78	20.37	0.09	0.09

Discussion

Our second study incorporated items generated by ChatGPT-3.5 in tandem with items retained from Study 1. After an EFA, we retained a 7-item solution with high loadings (> .75), high reliability, and convergent validity with the 28-item solution, while also replicating some findings from Study 1 with regards to differences in longtermism beliefs based on demographic characteristics (see Supplementary Materials).

Study 3

In our third study, we sought to formally validate this 7-item solution via confirmatory factor analysis. However, since in both Studies 1 and 2 we saw a ceiling effect for endorsement of longtermism beliefs, we additionally sought to reduce this effect in Study 3 by making salient the long-term future by directly specifying the timeframe people should have in mind when responding to the scale. We did so by using mixed design. This study was pre-registered, https://aspredicted.org/R1M_BWB.

Methods

Participants

Data collection was operated on Prolific. We pre-registered an a priori power analysis for a mixed ANOVA with 4 between and 4 within factors (with no expected interaction), and with f

= .10, alpha = .05, power = .80, which suggested a sample of 900 participants. Our total sample comprised a total of 903 participants. No participants had duplicate IP addresses, and no participant missed our attention check. Thus 903 participants were retained in the sample. The study lasted approximately 8 minutes, and participants received $1.40 for their participation.

Measures

Longtermism beliefs. The 7-item version of the scale from Study 2 was used. Responses were captured on 7-point Likert scales, ranging from 1 = Strongly disagree to 7 = Strongly agree. We manipulated the timeframe that participants were exposed to when they were answering these questions. This was done by changing the number of years in the future (100 vs. 500 vs. 1000 vs. 10000) people had in mind when being shown these items.

Participants completed the measure with all four timeframes shown to them in randomized order. We counterbalanced the first timeframe participants were shown, which allowed us to conduct between-subject comparisons across the four timeframes. We selected this range of years (100-1000) because 100 years is what most participants tended to have in mind when completing the measure unprompted (based on responses to an open-ended question in Study 1, see Table S5, and Study 2, see Tables S9 and S10 in the Supplementary Materials), and 10000 years was one of the highest values listed by participants when prompted to state how many years in the future they thought about when completing these measures.

Extinction threats. The seven items capturing beliefs about how humanity's extinction would be bad (a = .87) and the six items capturing beliefs about how humanity's extinction would be good/neutral (a = .87) were identical to Study 2. Again, both measures were captured on 7-point Likert scales, ranging from 1 = Strongly disagree to 7 = Strongly agree.

Participants also indicated how personally concerned they were about the six different extinction threats from Study 2. This time, they expressed their level of concern for each threat, on a 6-point Likert scale ranging from 1= not at all concerned to 6 = extremely concerned.

Demographic variables. All demographic variables were identical to Studies 1-2. Conservative political ideology was once again highly reliable (a = .90).

Results

Confirmatory Factor Analysis

For each of the four timeframes of the scale, the model provided a good fit to the data, with all items loading highly on the 1-factor factor solution (see Table S11 in the Supplementary Materials). Importantly, we evaluated model fit based on 3 fit indexes, as suggested by Kline (2016), the Comparative Fit Index (CFI), which is an indicator of goodness of fit, with scores ≥ 0.95 indicating good fit, the Root Mean Square Error of Approximation (RMSEA), which is an indicator of badness of fit, with scores ≤ .08 indicating good fit, and the Standardized Root Mean Squared Residual (SRMR), which is also an indicator of badness of fit, for which scores ≤ .08 indicate good fit. We did not consider the chi square statistic as our sample was large and the chi square test tends to produce inflated values with larger sample sizes. For all timeframes, the scale provided good fit to the data, good composite reliability and a high average amount of variance explained.

Table 4

Model Fit Indexes for Confirmatory Factor Analyses for Each Timeframe for Studies 3 and 4A, and across all timeframes for Studies 4B and 5

Parameter	100 years	500 years	1000 years	10000 years	1,000 years	10,000 years	100,000 years	1,000,000 years	Average score	Average score
$\chi^2(df = 14)$	97.77*	68.34*	60.53*	45.53*	28.12*	111.74*	54.66*	40.93*	37.58*	69.49*
CFI	.985	.991	.993	.996	.998	.987	.995	.997	.994	.995
RMSEA	.081	.066	.061	.050	.030	.079	.051	.042	.056	.058
SRMR	.018	.015	.011	.008	.001	.015	.009	.007	.012	.009
CR	0.95	0.95	0.94	0.94	0.94	0.94	0.94	0.94	0.96	0.96
Cronbach's *a*	0.95	0.95	0.96	0.97	0.95	0.96	0.96	0.97	0.96	0.95

AVE	0.72	0.72	0.70	0.69	0.73	0.74	0.78	0.82	0.76	0.75
Study	3	3	3	3	4A	4A	4A	4A	4B	5

Note. * $p < .001$. CR = Composite Reliability, AVE = Average Variance Explained. For studies 4B and 5 the average score was estimated for each item specifically across all timeframes.

Between-Subject Differences

Since participants were randomly assigned to see one of the four timeframes first, we were able to examine whether there are between-subject differences on endorsement of longtermism beliefs. Contrary to our pre-registered hypotheses, no significant difference was observed, $F(3, 898) = 0.30$, $p = .827$, $\eta^2_p = .001$.

Within-Subject Differences

When looking at the overall change collapsing across the four timeframes a significant effect was observed, $F(3, 2700) = 126.12$, $p < .001$, $\eta^2_p = .123$. Post-hoc comparisons (which were adjusted with a Bonferroni correction, such that alpha was set to .008 to correct for six tests) suggested that each timepoint significantly differed from each other, and the further away in the future a timeframe was the lower participants scored.

Notably, if participants saw a temporally close timeframe first, they tended to report lower scores for more temporally distant timeframes, as expected. However, when they saw a more distant timeframe first, they failed to update their scores for temporally closer timeframes, scoring with roughly equal scores across all timeframes. Figure 1 graphically depicts this trend. Detailed results can be found in Table S12 in the Supplementary Materials.

Figure 1

Line Graph Depicting Differences Across the Four Timeframes For Each of the Four Starting Timeframes with 95% C.I., with the Y Axis Superimposed to Highlight Differences in the Downward Trend

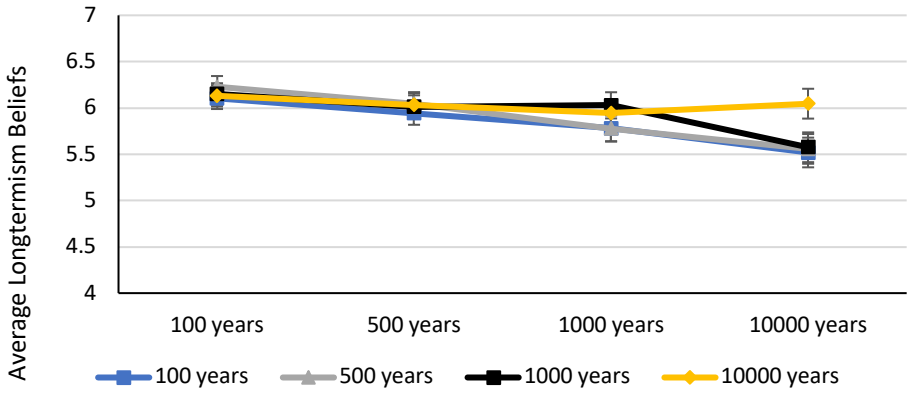

Do Longtermism Beliefs Relate to Concern for Extinction Threats to Humanity?

Regardless of the timeframe used for the longtermism beliefs scale, the correlations with the measures capturing why human extinction would be good directly replicated the results of Study 2, suggesting that, regardless of the timeframe, holding longtermism beliefs relates to negative appraisals of human extinction. Across timeframes, endorsement greater of longtermism beliefs related to greater concern for extinction threats, except for concern for unaligned AI. These results can be found in Table S13 in the Supplementary Materials.

Discussion

The first goal of this study was to validate the structure of the longtermism beliefs scale. CFAs supported the structure, providing evidence for its validity, reliability and ability to explain between 69 to 72% of the variance in the construct depending on the timeframe. The second goal was to replicate the correlations of Study 2. Regardless of the timeframe people responded to, a stable pattern of correlations was observed. The third goal of the study was to reduce the observed ceiling effect from Study 2. Some success was seen, as there was an overall decrease for timeframes that were further in the future within subjects.

However, an unexpected trend was also noted. People who were given no prior information and were asked to complete the LBS for timeframes that are further away in the future tended to overestimate their scores. Extant work on this phenomenon (i.e., Birnbaum,

1999) suggests that there are two potential ways to address this phenomenon. The first, which we attempt in Study 4A is giving clear instructions that attempts to put people on the same baseline estimate can correct such biases. The second approach would be to have participants provide simultaneous judgments for all timeframes (Study 4B).

Ultimately, despite this unexpected pattern of results, we replicated most correlations from Study 2 and the effect of longtermist identification on longtermism beliefs, such that self-identifying longtermists scored higher on the LBS across all timeframes (all *d*s ≥ 0.39) (see Table S14 in Supplementary Materials).

Study 4A

Our next study sought to examine whether it is possible to implement instructions that would help provide participants with a more accurate estimate of their longtermism beliefs. Further, we increased the timeframe used in the scale, which ranged from 1 thousand to 1 million years in the future. Aside from helping reduce such bias in people's responses, we also assessed participants' perceptions of whether major extinction threats can be solved, their perceptions of who is responsible for taking action to solve such issues, as well as their support for longtermist policies. This study was pre-registered, https://aspredicted.org/4J1_3DH.

Methods

Participants

Data collection was operated on Prolific. The sample size we obtained was large enough to power a between-subjects ANOVA with 4 groups (f = .10, alpha = .05, power = .80) and a repeated measures ANOVA with 4 measurements (f = .10, alpha = .05, power = .80, for 4 groups with 4 repeated measurements). Our sample comprised a total of 1141 participants. Forty-two participants had duplicate IP addresses, and one participant missed our attention check. Thus

1098 participants were retained in the sample. The study lasted approximately 8 minutes, and participants received $1.40 for their participation.

Measures

Longtermism beliefs. The same 7 items capturing longtermism beliefs were used. We manipulated the timeframe that participants were exposed to when they were answering these questions. This was done by changing the number of years in the future people had in mind when being shown these items, the timeframes were: 1000 years, 10,000 years, 100,000 years, and 1,000,000 years in the future.

Crucially, we also changed the introductory instructions for these questions, hoping to make participants aware that they would be responding to the same questions for several different timeframes. These instructions read:

> "For this part of the survey we will present you with questions focusing on your thoughts about the **far future**. Specifically we might ask you to answer the same set of questions, but each time we might ask you to answer these questions **with a different timeframe in mind for what the far future would be**. These timeframes could range from **1,000 years** in the future to **1,000,000 years** in the future.
>
> When answering these questions consider the fact that later in the survey you might be asked the same question with a distant (i.e., closer to 1,000,000 years in the future) or closer (i.e., closer to 1,000 years in the future) timeframe."

Perceptions of responsibility. We captured how responsible participants thought they (i.e., yourself), their nation's government, international governmental agencies, and everyone in the world were for: taking action to ensure the long-term survival of humanity for the far future (i.e., at least 10,000 years from now), ensuring that people in the far future (i.e., at least 10,000

years from now) live a good and prosperous life, advocating for the rights of people who will live the far future (i.e., at least 10,000 years from now). Responses for each item were captured on 7-point Likert-type scales (1 = Not at all responsible, 7 = Extremely responsible). Responses were highly consistent across both the target and the item, and thus we averaged items across targets, creating four average scores (4 targets x 3 items).

Solvability of extinction threats. We also captured the degree to which participants thought that the six aforementioned extinction threats could be solved because of: (1) collective solution (i.e., proposing and enacting laws that would allow us to successfully prevent or adapt to any of these global challenges) and (2) a technological solution (i.e., coming up with a novel technology that would allow us to successfully prevent or adapt to any of these global challenges). Responses for each item were captured on 7-point Likert-type scales (1 = Not at all likely, 7 = Extremely likely). Results for these items were highly similar, and thus we aggregated averaged scores for the two sets of six items into two constructs capturing belief in technological and in collective solutions to extinction threats.

Future generation policy attitudes. Ten items generated by the research team based on existing legislation implemented by several countries in the world (see Gonzalez-Ricoy & Rey, 2019 for an example) were used to assess policy attitudes. Responses were captured on 6-point Likert scales (1 = Strongly oppose, 6 = Strongly support). Scores were averaged into a unitary construct.

Demographic variables. All demographic variables were identical to Studies 1-3. Conservative political ideology was once again highly reliable (a = .90).

Results

Confirmatory Factor Analysis

We evaluated model fit based on the same criteria as Study 3. Across all timeframes, the scale provided good fit to the data, good composite reliability and a high average amount of variance explained. Results can be found in Table 4. Factor loadings for all timeframes can be found in Table S15 in the Supplementary Materials.

Between-Subject Differences

Since participants were randomly assigned to see one of the four timeframes first, we were able to examine whether there are between-subject differences on endorsement of longtermism beliefs. Contrary to our pre-registered hypotheses, no significant difference was observed, $F(3, 1094) = 2.60$, $p = .051$, $\eta^2_p = .007$. These results suggest that the instructions were not effective at getting participants to adjust their ratings, even if they explicitly highlighted that participants would be asked to rate shorter or longer timeframes later in the survey.

Within-Subject Differences

When looking at the overall change across the four timeframes a significant effect was observed, $F(3, 3291) = 149.11$, $p < .001$, $\eta^2_p = .120$. Post-hoc comparisons (which were adjusted with a Bonferroni correction, such that alpha was set to .008 to correct for six tests) suggested that each timepoint significantly differed from each other, and the further away in the future a timeframe was the lower participants scored.

Notably participants who saw the 100,000 years and 1,000,000 timeframes first, despite the explicit instructions warning them about the possibility of responding to temporally shorter timeframes, did not update their scores. Instead, in line with the results of Study 2, these participants tended to on average score the same for temporally shorter timeframes, suggesting that the instructions did not successfully decrease this effect. The results are depicted graphically in Figure 2 and reported in detail in Table S16 in the Supplementary Materials.

Figure 2

Line Graph Depicting Differences Across the Four Timeframes For Each of the Four Starting Timeframes with 95% C.I., with the Y Axis Superimposed to Highlight the Downward Trend

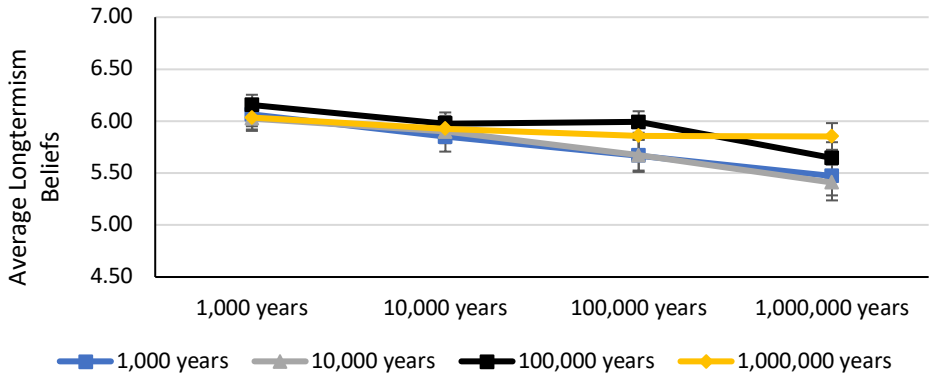

Do Longtermism Beliefs Relate to Concern for Future Generations?

Regardless of the specific timeframe used, greater endorsement of longtermism beliefs related to all outcomes in a manner that reflected greater concern for future generations. In detail, greater scores for the LBS related to increased belief that collective (all *r*s ≥ .19, all *p*s < .001) and technological (all *r*s ≥ .19, all *p*s < .001) solutions to major extinction threats to humanity's future can be reached. They also related to greater perceptions of responsibility for ensuring future generations have better lives for all agents (all *r*s ≥ .40, all *p*s < .001). Further, those who scored higher in longtermism beliefs also supported policies that protected future generations to

a greater degree (all $rs \geq .48$, all $ps < .001$). For more detailed results see Tables S17 and S18 in the Supplementary Materials.

Discussion

The aim of this study was to examine whether revising our instructions prior to the completion of the longtermism belief scale would calibrate participants' responses to the scale. Although scores were more in line with our expectations, overall, a similar pattern to that of Study 3 was observed, such that seeing a more temporally distant timeframe first results in participants not updating their scores for timeframes that are temporally closer. This suggests that potentially administering the scale with all possible timeframes at the same time, such that one item is answered for all timepoints might be the most optimal way to capture participants' beliefs. This could be the case, as it would allow participants to calibrate their scores accordingly, making all temporal comparison between timeframes salient.

Despite this unexpected pattern of scores, we observed robust correlations which supported our pre-registered hypothesis that greater endorsement of longtermism beliefs would relate to increased perceived responsibility to protect future generations, greater support for policies that seek to protect future generations, and increased efficacy for issues that will affect future generations. Finally, we also replicated the effect noted the significant difference in longtermism beliefs between those who identified as longtermists and those who did not (all $ds >$ .0.50), in line with results from our previous studies (see Table S19 in the Supplementary Materials).

Study 4B

In our next study we sought to correct the potential bias in participants' responses resulting from seeing one timeframe first, by displaying all timeframes at the same time. To do so, we slightly modified the wording of the scale to allow for each item to be measured at

different timeframes simultaneously. Further, we sought to both differentiate endorsement of longtermism beliefs from other related constructs and to determine whether those who hold these beliefs to a greater degree also score higher on longtermism behavioral outcomes. This study was pre-registered, https://aspredicted.org/L4G_DTC.

Methods

Participants

Data collection was operated on Prolific. Based on an a priori power analysis for a (two-tailed) correlation of $\rho = .15$, power of .90, and alpha of .05, we require a sample of 462 participants. We rounded this number up to 550 to account for potential exclusions, and because of recommendations to recruit sample sizes of $N = 500$ for correlational designs (Schönbrodt & Perugini, 2013). Based on sensitivity analyses with power set to .90, this sample would also be large enough to power effect sizes of $f = .05$ for a repeated measures ANOVA with four measurements. Our sample comprised a total of 548 participants. Eight participants had duplicate IP addresses, and one participant missed our attention check. Thus 542 participants were retained in the sample. The study lasted approximately 15 minutes, and participants received $2.90 for their participation.

Measures

Longtermism beliefs. The 7-item longtermism belief scale was used to capture the degree to which people endorsed the principles of longtermism. Importantly, our findings from Studies 3 and 4a suggest that there is a bias when people see different timeframes first. Thus, we modified the wording of the scale to allow for participants to complete each item for multiple timeframes at the same time. In this way, participants were presented with 1 of the 7 items (shown in a randomized order). A specific part of the sentence that referred to the future or future generations for each item was bolded and underlined. Participants were given instructions that

when responding to this item, they should consider the long-term future or future generations for 1000 years, 10,000 years, 100,000 years, and 1,000,000 years in the future.[1] Further, since all timeframes were shown at the same time, we also created a composite longtermism score. Running analyses for the individual timeframe scores, and for the average score produced highly similar results (see Supplementary Materials). Further, we considered that one core postulate of longtermism is those who are longtermist should endorse longtermism beliefs regardless of the timeframe, which is better reflected in the average rather than the individual score..

Responses were captured on slider scales, ranging from 0 = Strongly disagree – 100 = Strongly agree. An example item is given in Figure S1. The full set of instructions and the final version of the scale is available in the Supplementary Materials. A qsf file usable in Qualtrics including this version of the scale is available on OSF. Importantly, despite asking participants to respond to each of the 7 items four times (i.e., once per timeframe), the average time needed to complete the scale was 4.40 minutes, and the median time was 3.30 minutes.

Future self-continuity. We included a single-item measure of future self-continuity to examine whether the longtermism belief scale does not merely capture an individual's capacity to envision their own future self. We used the measure developed by Ersner-Hershfield et al. (2009). This measure uses 7 overlapping circles to highlight the degree to which an individual can envision and find similarities with their future self. Scores thus range from 1 to 7, with higher scores indicating more future self-continuity.

[1] Although this design helped reduce the discrepancies observed in Studies 3 and 4a, a small number of participants (ranging from 11% to 16% depending on the timeframes) did score higher for timeframes that were further in the future than timeframes that were temporally closer. This phenomenon could partially be explained by participants "eyeballing" their scores, as they were not shown the actual value they selected, but rather they could only drag the slider to a particular point in the scale. The exact percentages for each comparison between two timeframes is given in the Supplementary Materials.

Consideration of future consequences. We also reasoned that it would be important to distinguish between an individual's ability to consider the future consequences of their action, from their endorsement of longtermism beliefs. To do that we used the 12-item Consideration of Future Consequences (CFC) scale (Strathman et al., 1994). Responses were captured on a 7-point Likert scale (1 = Strongly disagree, 7 = Strongly agree).

Legacy concerns. Research suggests that an individual's concern about how they are remembered by future generations, as well as their motivation to leave a lasting legacy are important drivers of intergenerational reciprocity (e.g., Shrum, 2021; Zaval et al., 2015). We included a 3-item scale used in extant research to account for a person's concerns about their legacy. Responses were captured on a 6-point Likert-type scale (1 = Not at all, 6 = A great amount).

Utilitarianism. Utilitarianism was measured with the 5-item impartial beneficence and 4-item instrumental harm subscales of the Oxford Utilitarianism Scale developed by Kahane et al. (2018), captured on a 7-point Likert scale (1 = Strongly disagree, 7 = Strongly agree).

Effective altruism. Effective altruism beliefs were measured with the 6-item expansive altruism scale and 6-item effectiveness focus scale developed by Caviola et al. (2022), captured on a 7-point Likert scale (1 = Strongly disagree, 7 = Strongly agree).

Longtermism behavioral intentions. We adapted several items from work on effective altruism (Caviola et al., 2022) for the context of longtermism to capture behavioral intentions. One item provided a short explanation of longtermism (roughly 3 paragraphs long) and asked participants to what extent they agreed with the ideas of longtermism (1 = Strongly disagree, 7 = Strongly agree) and to what extent they were interested in learning more about longtermism 1 = Not at all interested, 5 = Extremely interested). Participants were then asked if they were interested in signing up for a monthly newsletter about longtermism (1 = yes, 0 = no), if they

would like to be entered into a lottery to receive a free copy of What We Owe The Future (1 = yes, 0 = no), or if they, much like "thousands of longtermists who have committed giving away 10% of their income to charities for their duration of their lives that experts believe will positively impact the future" could imagine themselves doing the same thing now or in the future (1 = yes, 0 = no).

Demographic variables. All demographic variables were identical to the previous studies. Conservative political ideology was once again highly reliable (a = .90).

Results

Confirmatory Factor Analysis

We evaluated model fit based on the same criteria as our previous studies. Across all timeframes, the scale provided good fit to the data, good composite reliability and a high average amount of variance explained (see Tables S20 and S21 in the Supplementary Materials). This was also the case when utilizing items that reflected the average score across all timeframes (see Table 4).

Within-Subject Differences

When looking at the overall change across the four timeframes a significant effect was observed, $F(3, 1623) = 437.27$, $p < .001$, $\eta^2_p = .447$. Post-hoc comparisons (which were adjusted with a Bonferroni correction, such that alpha was set to .008 to correct for six tests) suggested that each timepoint significantly differed from each other, and the further away in the future a timeframe was the lower participants scored (See Table S24 in the Supplementary Materials).

Figure 3

Line Graph Depicting Differences Across the Four Timeframes with 95% C.I.

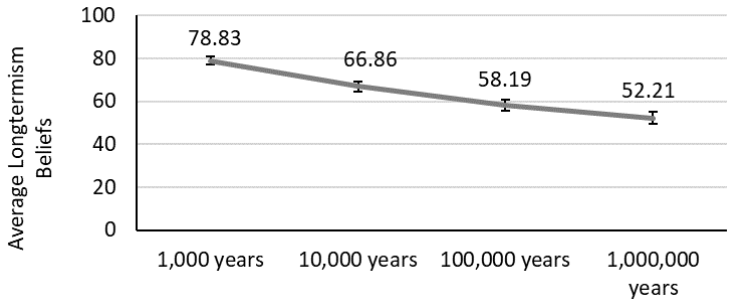

Correlations Between Longtermism and Related Constructs

As hypothesized, longtermism beliefs related positively to all outcomes, except for instrumental harm beliefs. Further, as hypothesized these correlations were not large in magnitude, suggesting that longtermism beliefs are distinct from but related to beliefs about the future, effective altruism and utilitarianism beliefs.

Table 5

Bivariate Correlations Between All Psychological Measures for Study 4b

Measure	1	2	3	4	5	6	7	8
1. Longtermism Belief	--							
2. Legacy	0.36**	--						
3. FSC	0.10*	0.17**	--					
4. CFC	0.23**	0.29**	0.27**	--				
5. Effectiveness Focus	0.15**	0.19**	0.01	0.07	--			
6. Expansive Altruism	0.39**	0.47**	0.06	0.24**	0.30**	--		
7. Instrumental Harm	0.01	0.10	-0.04	-0.20**	0.41**	0.08	--	
8. Impartial Beneficence	0.31**	0.33**	0.02	0.06	0.35**	0.58**	0.33**	--

Note. * $p < .05$, ** $p < .001$

Do Longtermism Beliefs Relate to Longtermism Behavioral Intentions?

When examining the association between scores on the longtermism belief scale and each of the behavioral intention outcomes, a significant positive association was observed for every

single outcome (βς ranging from .15 to .51, ps < .01) . However, when including all other psychological measures as covariates, only the association for agreement with longtermism principles (β = .33, SE = .002, p < .01) and interest in learning more about longtermism (β = .20, SE = .002, p < .01) remained significant (see Tables S22 and S23 in the Supplementary Materials).

Response Patterns for the Longtermism Beliefs Scale

The principles of longtermism argue that we should ascribe the same high amount of moral worth and consideration to all future people, as well as consider that our actions can influence all future generations to come. Empirically, this could be seen as follows: A longtermist should care (i.e., have a high score on the scale) and have such a high score for all possible timeframes. Importantly, this analysis was exploratory and not preregistered.

As shown in Figure 4, most participants scored higher when the timeframe was temporally closer across each possible comparison. A small percentage of participants scored lower for temporally closer timeframes compared to more distant ones. Finally, an even smaller proportion of participants scored in the same way across all timeframes. Theoretically, from the last two groups, participants who scored high in longtermism beliefs, and did so in manner that is equal across all timeframes, or increasing for future timeframes, could be considered longtermists. To set a baseline for high endorsement of longtermism beliefs, we took the average score for the 1,000 years timeframe. Based on these criteria, we considered longtermists participants who scored higher than the mean for the 1,000 years timeframe (i.e., > 78.83), for all possible timeframes. This left us with a total of 123 participants (23% of the sample).

Figure 4

Frequency of Response Patterns for Each Comparison of Two Timeframes

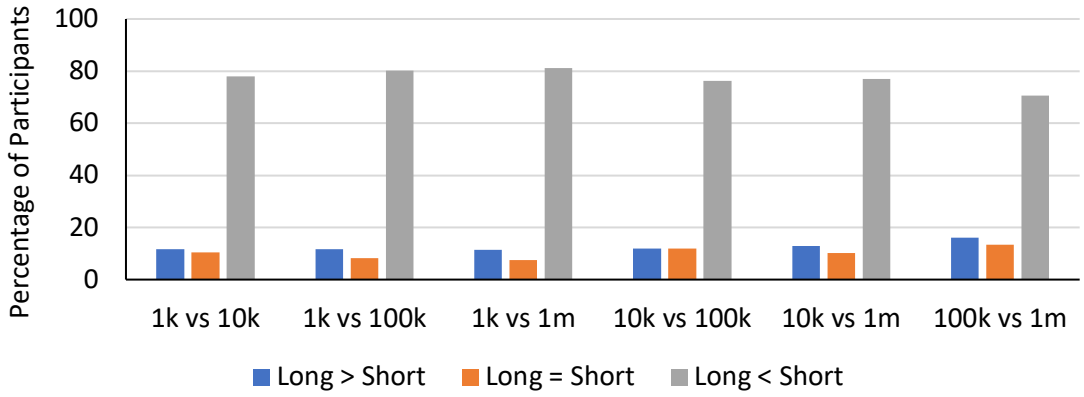

Differences Between Longtermists and Other Participants

When we compared the 123 remaining participants and the remained of the sample (N = 418) we found that longtermists scored significantly higher in longtermism beliefs, legacy concerns, future self-continuity, consideration of future consequences, impartial beneficence, and expansive altruism but not on any other outcome. These results are summarized in Table 6 and depicted visually in Figure 5.

Figure 5

Bar Graph Depicting Differences Between Longtermists and Other Participants with 95% C.I.

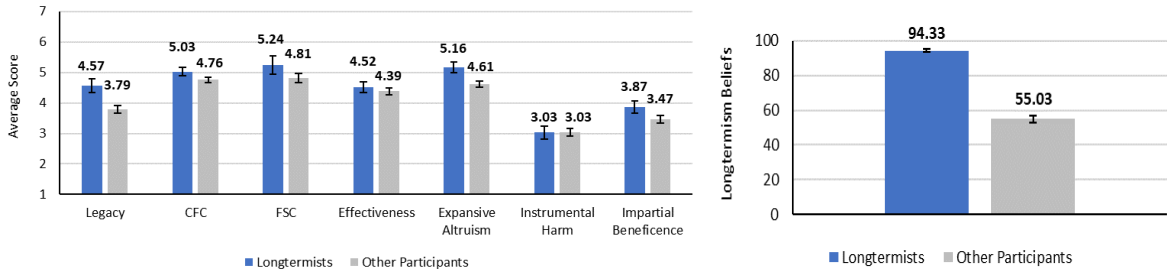

Table 6

Differences Between Longtermists and Other Participants

Outcome	*t*	*df*	*p*	*d*
Longtermism Beliefs	**32.60**	**528.40**	**<.001**	**2.40**
Legacy concerns	**5.80**	**540**	**<.001**	**0.62**
CFC	**2.91**	**540**	**.004**	**0.31**
FSC	**2.66**	**540**	**.008**	**0.27**
Effectiveness	1.21	540	.226	0.13
Expansive Altruism	**5.24**	**540**	**<.001**	**0.54**
Instrumental Harm	0.02	540	.983	0.01
Impartial Beneficence	**3.09**	**540**	**.002**	**0.32**

Note. Bolded results depict significant differences between the two groups. For Longtermism Beliefs tests were estimated for unequal variances between the two groups.

Discussion

Study 4b showed that displaying all timeframes for the longtermism beliefs scale helps reduce bias in participants' responses. In line with research on intergenerational discounting (Wade-Benzoni & Tost, 2009), responses for longer timeframes were significantly lower, with most participants displaying this pattern. However, a small number of participants scored high across all timeframes. Although not pre-registered, in an exploratory analysis, we compared these participants, who we considered to be longtermists, as their responses were reflective of the principles of longtermism, to all other participants. Results suggested that longtermists scored higher on future oriented constructs, expansive altruism and impartial beneficence. In our final study, we sought to replicate this exploratory finding, in an adequately powered sample.

Study 5

In our final study, we sought to directly replicate the exploratory findings of Study 4b in a pre-registered, and highly-powered study, https://aspredicted.org/F3B_WQV. We also examined whether our method of identifying longtermists would overlap with a more direct approach at identifying longtermists, namely by providing participants with the actual principles

of longtermism and asking them if they then identify as a longtermist. Although we do not expect that simply asking someone if they are a longtermist necessarily means that they truly are one, providing evidence that those who score in the aforementioned longtermist pattern are more likely to state they are a longtermist after being provided with the principles of longtermism would suggest that our empirical method of identifying longtermists has face validity.

Methods

Participants

Data collection was operated on Prolific. Based on an a priori power analysis that focused on the ratio of non-longtermist to longtermist participants identified in Study 4b and aiming to detect effect sizes as small as d = .20 for a between-subject comparison (i.e., independent sample t-test) for the two groups with power set to .80 and an alpha of .05 we required 1496 participants. We rounded this number up to 1550 to account for potential exclusions.

Our sample comprised a total of 1548 participants. Eight participants had duplicate IP addresses, and one participant missed our attention check. Thus 1542 participants were retained in the sample. The study lasted approximately 15 minutes, and participants received $2.90 for their participation.

Materials and Procedure

The following measures were identical to Study 4B: longtermism beliefs, captured with the longtermism beliefs scale[2], legacy concerns, future self-continuity, consideration of future consequences, impartial beneficence, and expansive altruism. We also retained the binary outcomes capturing whether people were interested in signing up for a newsletter, could imagine themselves donating 10% of their income in the future for a longtermism cause and signing up to

[2] Using the combined sample of Studies 4B and 5 we tested for measurement invariance across gender, race, political party and age, with results suggesting that the scale achieves configural, metric and scalar invariance (see Supplementary Materials).

receive a book on longtermism. We changed one outcome, namely whether participants identify with the principles of longtermism. Instead we presented participants with the same short four paragraph long description of longtermism, and asked participants whether, after reading the principles of longtermism, they identified as longtermists. Participants completed the longtermism beliefs scale first, followed by all other measures in a randomized order.

Results

Re-examining the model fit of the longtermism beliefs scale we found it provided good fit to the data, had good reliability and explained a high amount of variance (see Table 4).

Within-Subject Differences

When looking at the overall change across the four timeframes a significant effect was observed, $F(3, 4602) = 1257.28$, $p < .001$, $\eta^2_p = .450$. Post-hoc comparisons adjusted for six tests with a Bonferroni correction (see Table S26) suggested that each timepoint significantly differed from each other, and the further away in the future a timeframe was the lower participants scored (see Figure 6).

Figure 6

Line Graph Depicting Differences Across the Four Timeframes with 95% C.I.

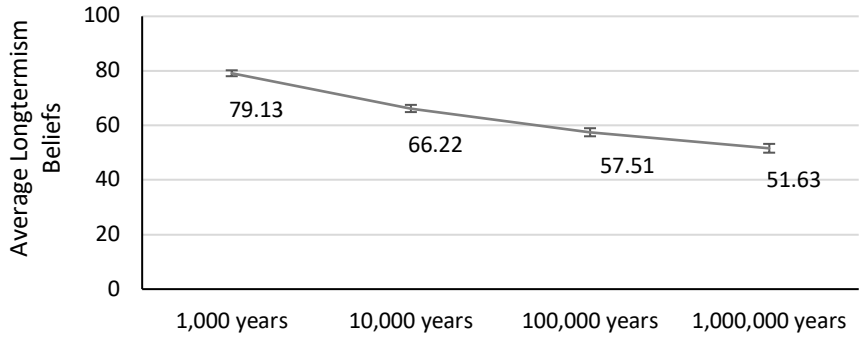

Correlations Between Longtermism and Related Constructs

Replicating the results of Study 4B, all correlations were in the same direction, significant and of similar magnitude as those of Study 4B (see Table 7).

Table 7

Bivariate Correlations for Study 5

	1	2	3	4	5
1. Longtermism Belief	--				
2. Legacy	0.37*	--			
3. FSC	0.13*	0.17*	--		
4. CFC	0.32*	0.30*	0.24*	--	
5. Expansive Altruism	0.39*	0.44*	0.09*	0.34*	--
6. Impartial Beneficence	0.31*	0.38*	0.05	0.15*	0.59*

Note. * $p < .001$.

Do Longtermism Beliefs Relate to Longtermism Behavioral Intentions?

When examining the association between scores on the longtermism belief scale and each of the behavioral intention outcomes, a significant positive association was observed for every single outcome (βs ranging from .12 to .31, $ps < .001$). However, when including all other psychological measures as covariates, the associations for signing up for the newsletter and expressing interest for the longtermist book were no longer significant. Importantly, the association with stating that one can imagine themselves donating 10% of their income remained significant ($\beta = 0.11$, $SE = 0.002$, $p < .01$). For detailed results see Table S27 in the Supplementary Materials.

Response Patterns for the Longtermism Beliefs Scale

Response patterns for the longtermism beliefs scale replicated by and large the results of Study 4B (see Figure S2 in the Supplementary Materials). The majority of participants scored higher for shorter rather than longer timeframes. Specifically roughly 11% scored higher for longer than shorter timeframes, 13% scored equally across all timeframes, and 76% scored higher for shorter timeframes.

Differences Between Longtermists and Other Participants

Following the results of Study 4B we pre-registered that we would operationalize longtermists as those who scored above the mean for the shortest timeframe (i.e., scores > 79.13) and had an equally high, or higher score for other timeframes. A total of 359 participants were grouped as longtermists (23% of the sample). When we compared these participants to the remaining participants we found that longtermists scored significantly higher in longtermism beliefs, legacy concerns, future self-continuity, consideration of future consequences, impartial beneficence, and expansive altruism, in line with our pre-registered hypotheses.

Figure 7

Bar Graph Depicting Differences Between Longtermists and Other Participants with 95% C.I.

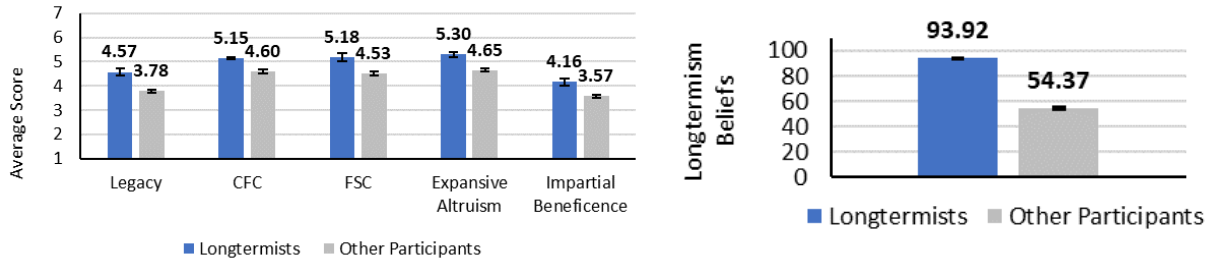

Table 8

Differences Between Longtermists and Other Participants

Outcome	*t*	*df*	*p*	*d*
Longtermism Beliefs	57.44	1503.90	<.001	2.53
Legacy concerns	10.31	1533	<.001	0.62
CFC	10.43	1533	<.001	0.63
FSC	6.58	1532	<.001	0.40
Expansive Altruism	10.36	1533	<.001	0.63
Impartial Beneficence	7.90	1533	<.001	0.47

Note. Bolded results depict significant differences between the two groups. For Longtermism Beliefs tests were estimated for unequal variances between the two groups.

Does Scoring Like a Longtermist Predict Longtermist Self-Identification?

To determine whether our method of identifying longtermists is face valid, we provided participants with a short overview of the principles of longtermism (see Supplementary

Materials). We then asked participants if they identified as longtermists based on these readings. A total of 976 participants identified as longtermists (63% of the sample). In line with our pre-registered analysis, via a chi-square test, we found that participants who were longtermists based on their scores on the LBS were significantly more likely to also identify as longtermists (N = 313, 87%, $\chi^2 = 112.75$, $p < .001$, V = 0.27). Even when we accounted for all other future-oriented constructs, as well as impartial beneficence, and expansive altruism, higher longtermism scores predicted longtermism identification ($\beta = 0.35$, $SE = 0.001$, $p < .001$; see Table S28 in the Supplementary Materials).

Thus, our empirical method of identifying longtermists via the LBS meaningfully detected those who self-identified as longtermists after being provided with the principles of the philosophy. Further, these results suggest that simply providing the principles of longtermism and asking people if they identify as a longtermist is a more liberal method (and potentially even an erroneous one) of identifying longtermists. Finally, as a validity check, we were able to replicate all the differences between those who scored as longtermists and those who did not in when we used the self-identification question (see Supplementary Materials).

Discussion

In our final study, we replicated the main findings of Study 4B in a larger sample. Adding to these findings, we pre-registered and found evidence in support for a data-driven approach to identifying longtermists. Although this approach was data-driven, it also adheres to the principles of longtermism. Based on our criteria, those who identify as longtermists place a great value on protecting future generations, and they acknowledge that what we do today can greatly influence an untold number of people. They do so regardless of how far into the future they are asked to think. Although we can ask participants to self-identify as longtermists, after presenting them with the principles of this philosophy, a large number of people who do not exhibit this pattern

self-identify as longtermist. Thus, considering this inconsistency, we propose that using the longtermism beliefs scale can meaningfully distinguish those who meaningfully endorse the longtermist philosophy and those who simply agree with its principles.

Our results also highlight what distinguishes longtermists from non-longtermists. Longtermists appear to be more future-oriented (scoring higher in future-self continuity, and consideration of future consequences), tend to care more about their impact on future people (i.e., higher scores on legacy concerns), and are more altruistic and utilitarian. In fact, in an exploratory path model (see Supplementary Materials) we found that differences in these characteristics mediate the effect of longtermism scores on the longtermist behavioral outcomes.

General Discussion

Researchers across psychology and the natural sciences (Caviola et al., 2021; Blaser, 2018; Syropoulos, & Markowitz, 2022; Taylor et al., 2012), philosophers (MacAskill, 2022a; Ord, 2020), policy makers (Bose & Shepardson, 2023) and journalists in the popular press (Hunter, & Hewson, 2020; McLamb, 2022) recognize that overcoming existential threats to the long-term future of humanity will require collective and individual effort from people today. The longtermism philosophy and social movement in particular puts emphasis on garnering support for the welfare of near and distant future generations of humans (MacAskill, 2022a, 2022b; Ord, 2020). Across 6 high-powered studies (Total N = 4,878), we provide empirical support for the Longtermism Beliefs Scale (LBS) as the first valid and reliable measure of alignment with the longtermism philosophy, demonstrating that longtermism beliefs are a key predictor of critical attitudes and behavioral outcomes associated with protecting future generations of humans from premature extinction.

People who self-identify as longtermists consistently score higher on the LBS, and do so consistently across more distant timeframes. The LBS predicts heightened concern for various

human extinction threats (e.g., climate change, pandemic disease, unaligned AI) and negative appraisals regarding the prospect of humanity going extinct in the future (Studies 1-3), with these effects persisting even when participants consider profoundly temporally distant hypothetical futures (e.g., 10,000 years in the future; Study 3). Additionally, the more people align with longtermism beliefs, the more they perceive future existential threats to be solvable, the more they support pro-future public policies, and the more they perceive themselves, their nation's government, international governmental agencies, and everyone in the world to bear responsibility in ensuring the long-term survival and prosperity of humanity into the far future and advocating for the rights of future people (Study 4a).

Importantly, the LBS predicts alignment with other future-oriented attitudes (e.g., motivation to leave behind a positive legacy, future self-continuity, consideration for the future consequences of behavior) and patterns of equitable and expansive prosociality (e.g., expansive altruism, impartial beneficence; Studies 4b-5). Finally, and perhaps most critically, the LBS predicts a variety of attitudes and behavioral intentions directly associated with the longtermism philosophy and social movement (e.g., interest in learning more about longtermism, agreement with the principles of longtermism, behavioral intentions to donate 10% of one's income to future-oriented charities, interest in entering a raffle to receive a book on longtermism, and interest in signing up for a longtermism-related newsletter; Studies 4b-5). Demonstrating the unique contribution of longtermism beliefs in predicting these outcomes, the LBS continues to explain unique variance in participants' interest in learning more about the philosophy (Study 4b), agreement with the philosophy's core principles (Study 4b), and behavioral intentions to

donate 10% of their income to future-oriented charities (Study 5) even when controlling for a host of related future-oriented and morally relevant attitudes.[3]

In addition to treating the LBS as a continuous measure, we also formulated and found evidence for a data-driven approach to utilizing the scale to operationalize people as longtermists according to the central tenants of the longtermism philosophy. The principles of longtermism hold that people today ought to ascribe a high amount of consideration and moral responsibility toward all future people and consider that present actions can have a tangible positive impact on all future generations regardless of how far in the future such generations of people exist. With this in mind, people who (1) scored above the mean on the LBS and (2) whose LBS scores did not decrease across increasing temporal distance (i.e., people whose high scores on the LBS were no lower for later timeframes with respect to the closest timeframe) were operationalized as longtermists in Studies 4b and 5. Whereas most people show a steady decline in longtermism beliefs the further into the future they are asked to consider, a minority of subjects (i.e., 23% of the sample in both Studies 4b and 5) met our criteria. Providing evidence for this treatment of the scale as a face-valid means to identify longtermists empirically, we find that 87% of longtermists identified through this procedure are more likely to self-identify as longtermists after being presented with the core tenants of the philosophy than those who don't meet our criteria. Further, we find that empirically identified longtermists tend to score significantly higher in longtermism beliefs, legacy concerns, future self-continuity, consideration for future consequences of behavior, impartial beneficence and expansive altruism when compared to those who fail to meet our criteria (Studies 4b-5).

[3] In fact, exploratory (and not pre-registered) mediation analysis conducted in Study 5, suggests that legacy concerns, CFC, expansive altruism and impartial beneficence mediate the effect of being a longtermist (based on scores on the LBS) and willingness to engage in the 3 longtermist behaviors.

Implications and Future Directions

Collectively, these findings provide evidence that there is a good deal of variability in the extent to which people acknowledge and consider future generations, feel morally responsible for the welfare of future people, and perceive efficacy in being able to positively impact the future. This work has important theoretical implications for ongoing empirical efforts to investigate future-oriented (e.g., Ersner-Hershfield et al., 2009; Shrum, 2021; Tamir, & Mitchell, 2011; Zaval et al., 2015), moral (e.g., Crimston et al., 2016, 2018; Graham et al., 2017), and altruistic (e.g., Everett et al., 2018; Law et al., 2022; McManus et al., 2020) thinking and behavior spanning multiple psychological disciplines (e.g., social psychology, cognitive psychology, moral-psychology, intergroup dynamics, prosocial behavior). Existing metrics that assess equitable and impartial prosocial intentions (e.g., the Expansive Altruism Scale; Caviola et al., 2022) capture the influence of social distance on altruistic intentions quite well yet ignore the influence of temporal distance. Likewise, existing metrics that assess the extent to which people are motivated towards positively impacting the future (e.g., the Legacy Motives Scale; Shrum et al., 2021; Zaval et al., 2015) conflate altruistic motivations to benefit future others with egoistic motivations to be remembered favorably.

Building upon prior psychological research highlighting the significance of consideration (Crimston et al., 2016, 2018; Graham et al., 2017), moral responsibility (Syropoulos & Markowitz, 2021, 2022), and efficacy (Bradley et al., 2020; Hornsey et al., 2021) as key motivators of peripherally-related moral and prosocial attitudes and behaviors (e.g., altruism, environmentalism, humanitarianism), the present studies investigating the Longtermism Beliefs Scale (LBS) underscore the relevance of these three dimensions in shaping attitudes, behaviors, and decision-making pertaining to the advancement of future human welfare. By serving as an internally consistent, construct- and predictively-valid tool, the novel Longtermism Beliefs Scale

(LBS) offers scholars a valuable resource to expand theoretical knowledge concerning individuals' assessments of the moral and prosocial value attributed to future generations.

Beyond the field of psychology, the findings of the present studies hold implications for the longtermism philosophy and its associated social movement. Notably, the Longtermism Beliefs Scale (LBS) emerges as a pioneering psychological measure capable of assessing lay beliefs aligned with the fundamental tenets of longtermism. As such, it has the potential to contribute significantly to the understanding and evaluation of longtermism within both academic and public discourse. In particular, we argue that the data-driven approach we have developed to distinguish longtermists from non-longtermists will provide substantial utility to those who wish to proliferate the teachings of the longtermism philosophy and grow the movement. Although the longtermism movement is relatively small at present (MacAskill, 2022a), the prevalence of people we have identified in the present studies by way of the LBS who embody the principal tenants of the philosophy is relatively large (23% of subjects). Even if this number is inflated with respect to the actual prevalence of such beliefs in the general population, there are likely many people currently living who endorse longtermism beliefs but aren’t privy to the existence of the philosophy or social movement. Put plainly, it could be the case that the largest barrier standing in the way of growing the longtermism movement is simply ignorance to its existence. To address this gap, experimental research examining the effects of interventions trying to increase longtermist beliefs, implementation of and assessment of longtermism courses (at different educational levels) and nationally-representative polls are needed.

We argue that the LBS can be used to identify people in the general population who may be especially receptive to future appeals and interventions geared towards growing the longtermism social movement and educating people about the importance of protecting future

generations more broadly. Future research may look to develop educational interventions regarding longtermism and address whether these can inspire people, particularly those who endorse the philosophical principles of longtermism on the LBS, to get involved in the movement and engage in a greater number of pro-future behaviors in vivo. One particularly promising piece of evidence from the present research is the invariance in longtermism beliefs across demographic variables such as income, education, and age. Although conservative political ideology was negatively associated with longtermism beliefs, such associations were fairly weak across all studies (r = -.11 to -.34). These results suggest that appeals to inspire longtermism beliefs may show success across demographic differences.

The present research may also serve to benefit the effective altruism movement (Singer, 2015, 2016) and the emerging wave of scientific inquiry surrounding it (e.g., Caviola et al., 2021, 2022; Everett et al., 2018; Kahane et al., 2018; Wilks et al., 2023). The longtermism and effective altruism philosophies are conceptually intertwined, and their social followings are largely overlapping (MacAskill et al., 2022a). However, it isn't at present clear what percentage of effective altruists also align with the longtermism philosophy. Future research might apply the LBS to address this open question and gap in knowledge. A related avenue for future research to explore is whether the same cognitive, affective and moral underpinnings that predict alignment with the effective altruism philosophy predict alignment with the longtermism philosophy as well.

Finally, in addition to the numerous theoretical implications the current research has for the fields of psychology and philosophy, we argue that these studies and the LBS may also offer tangible practical benefits for public policy and the future welfare of human society at large. Humanity currently stands at a peculiar crossroads, at which the current actions of those who are presently living have the power to both help and harm the chances of ensuring a long and

prosperous future for our species (MacAskill, 2022a, 2022b; Ord, 2020). The substantial level of support observed for longtermism beliefs in our present studies offers hope, indicating that mobilizing individual actions (such as donations to future-oriented charities and engaging in pro-environmental behaviors) and collective endeavors (such as advocating for pro-future public policies) to safeguard humanity's future from premature extinction may be more achievable than the current limited following of the longtermism movement implies (MacAskill, 2022b). In a similar manner to how the LBS may be used to identify potential longtermists in the general population, future applied scientific inquiry may seek to examine how the scale may also be used as a tool to gauge potential support for direct efforts to protect humanity from eventual extinction and harnessed to guide related interventions.

Limitations

Whereas the present studies significantly advance theoretical and practical knowledge related to the promotion of future human welfare, there are of course notable limitations worthy of consideration. First, in Studies 1 and 2, we observed a ceiling effect in our measurement of longtermism beliefs, with many respondents scoring towards the upper end of the scale. In part because research shows that people generally struggle to construct vivid simulations of the *distant* future (e.g., Tamir, & Mitchell, 2011), we speculated that this effect may have stemmed from participants only spontaneously considering relatively temporally *proximal* futures when responding to the scale. To address this issue and increase variability in our measurements, we provided participants with specific instructions to consider not only proximal futures when completing the scale (e.g., up to 100 years in the future), but distant futures as well (e.g., 1,000 years in the future and beyond) in Studies 3-5. This adjustment to the scale successfully generated more variability without sacrificing the scale's validity or reliability.

Another related challenge emerged when we began iterating the LBS across multiple timeframes in Study 3. We observed that participants who were asked to respond to items on the scale with respect to the most distant timeframes first, without any prior reference point for comparison, tended to report inflated and biased levels of longtermism beliefs with respect to these most distant points in time. In other words, participants who were presented with the most distant time points first reported stronger longtermism beliefs for these more distant time points than for subsequently displayed, but comparatively more proximal time points. In Studies 4a and 4b, we aimed to improve upon the scale's delivery method by exploring two potential solutions to this problem. Our findings revealed that participants responded more effectively when we implemented a simultaneous measurement approach. Specifically, we measured participants' responses to each of the seven items on the scale at four different time points simultaneously. This approach proved to be beneficial as it provided participants with an adequate frame of reference. Consequently, participants were better able to calibrate their responses to the entire timeframe in which they would be asked to respond to the scale's items (as observed in Study 4b). This solution significantly reduced bias without sacrificing the scale's validity or reliability, with results even suggesting that it takes about 3 minutes to complete the scale, a relatively short amount of time. As such, we advise that future research employing the LBS utilizes this method for the delivery of the items.

Yet another limitation of the present research pertains to the fact that certain relationships between the LBS and longtermism-related behavioral intentions that were significant in the bivariate context became non-significant when controlling for other related constructs (e.g., legacy concerns, consideration of future consequences). Notably, however, the relationships between scores on the LBS and agreement with the core principles of longtermism, interest in learning more about longtermism, and pro-future donation intentions remained significant even

in the presence of related covariates. Importantly, the LBS was the strongest predictor of agreement with the core principles of longtermism in the multiple regression model including the covariates, serving to provide further evidence of the scale's validity as a metric to assess longtermism beliefs. Nonetheless, the LBS shows strong predictive validity, with all of the hypothesized associations between the novel scale and longtermism-related behavioral intentions being borne out in the data in the bivariate context. Future research might seek to address the extent to which associations between longtermism beliefs and longtermism-related behavioral intentions may be mediated or moderated by other relevant constructs. Although we conducted such an analysis, in an exploratory and not pre-registered manner, causal designs are needed to establish a cause-and-effect relationship. Further, longitudinal designs could also help tease apart the interrelationships between longtermist beliefs, future-oriented variables, expansive prosociality, and behavioral intentions.

Finally, the present studies only addressed longtermism beliefs and related associations in WEIRD samples recruited through Prolific Academic. However, existential threat to the longevity and prosperity of the human race is a global problem which will require a global solution. Thus, prospective research employing the LBS may seek to examine differences in longtermism beliefs across cultures, nations, and societies spanning the entire globe.

Conclusion

In the face of numerous existential threats to humanity's future and with the longevity and prosperity of distant future generations resting in the hands of current generations, examining the extent to which living people consider, feel morally responsible for and feel efficacy to positively impact future people is as critical an issue as any. On the basis of the findings in the present research and the related questions which are still ripe for inquiry, the longtermism beliefs scale will have a lasting theoretical impact across psychological and

philosophical disciplines, hopefully inspiring fruitful inquiry into future-oriented thinking, behavior and prosociality. We ultimately hope that the present research will offer pragmatic utility in helping to solve some of the greatest challenges that could eventually spell the end for humanity if left unchecked, helping to ensure the well-being of future generations and safeguarding the future of humanity.

References

Atance, C. M., & O'Neill, D. K. (2001). Episodic future thinking. *Trends in Cognitive Sciences*, *5*(12), 533–539. https://doi.org/10.1016/S1364-6613(00)01804-0

Blaser, M. J. (2018). The Past and Future Biology of the Human Microbiome in an Age of Extinctions. *Cell*, *172*(6), 1173–1177. https://doi.org/10.1016/j.cell.2018.02.040

Birnbaum, M. H. (1999). How to show that 9 > 221: Collect judgments in a between-subjects design. *Psychological Methods, 4*(3), 243–249. https://doi.org/10.1037/1082-989X.4.3.243

Bose, N., & Shepardson, D. (2023, May 4). Biden meets Microsoft, Google CEOs on AI dangers. *Reuters*. https://www.reuters.com/technology/white-house-meet-microsoft-google-ceos-ai-dangers-2023-05-04/

Bostrom, N. (2002). Existential risks: Analyzing human extinction scenarios and related hazards. *Journal of Evolution and Technology*, *9*. https://ora.ox.ac.uk/objects/uuid:827452c3-fcba-41b8-86b0-407293e6617c

Bradley, G. L., Babutsidze, Z., Chai, A., & Reser, J. P. (2020). The role of climate change risk perception, response efficacy, and psychological adaptation in pro-environmental behavior: A two nation study. *Journal of Environmental Psychology*, *68*, 101410. https://doi.org/10.1016/j.jenvp.2020.101410

Caviola, L., David_Althaus, Stefan_Schubert, & Joshua_Lewis. (2022). *What psychological traits predict interest in effective altruism?* https://forum.effectivealtruism.org/posts/7f3sq7ZHcRsaBBeMD/what-psychological-traits-predict-interest-in-effective

Caviola, L., Schubert, S., & Greene, J. D. (2021). The Psychology of (In)Effective Altruism. *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2021.03.015

Crimston, C. R., Hornsey, M. J., Bain, P. G., & Bastian, B. (2018). Toward a Psychology of Moral Expansiveness. *Current Directions in Psychological Science*, *27*(1), 14–19. https://doi.org/10.1177/0963721417730888

Crimston, D., Bain, P. G., Hornsey, M. J., & Bastian, B. (2016). Moral expansiveness: Examining variability in the extension of the moral world. *Journal of Personality and Social Psychology*, *111*(4), 636–653. https://doi.org/10.1037/pspp0000086

Everett, J. A. C., Faber, N. S., Savulescu, J., & Crockett, M. J. (2018). The costs of being consequentialist: Social inference from instrumental harm and impartial beneficence. *Journal of Experimental Social Psychology*, *79*, 200–216. https://doi.org/10.1016/j.jesp.2018.07.004

Gonzalez-Ricoy, I., & Rey, F. (2019). Enfranchising the future: Climate justice and the representation of future generations. *WIRE's Climate Change, 10*(5), e598. https://doi.org/10.1002/wcc.598

Graham, J., Waytz, A., Meindl, P., Iyer, R., & Young, L. (2017). Centripetal and centrifugal forces in the moral circle: Competing constraints on moral learning. *Cognition*, *167*, 58–65. https://doi.org/10.1016/j.cognition.2016.12.001

Heuvel, K. vanden. (2023, April 25). *Will AI Lead to Human Extinction?* https://www.thenation.com/article/society/ai-chatgpt-tech-disruption-humanity/

Hornsey, M. J., Chapman, C. M., & Oelrichs, D. M. (2021). Ripple effects: Can information about the collective impact of individual actions boost perceived efficacy about climate change? *Journal of Experimental Social Psychology*, *97*, 104217. https://doi.org/10.1016/j.jesp.2021.104217

Hunter, A., & Hewson, J. (2020, April 22). *There are 10 catastrophic threats facing humans right now, and coronavirus is only one of them*. The Conversation. http://theconversation.com/there-are-10-catastrophic-threats-facing-humans-right-now-and-coronavirus-is-only-one-of-them-136854

Jones, B., & Rachlin, H. (2006). Social Discounting. *Psychological Science, 17*(4), 283–286. https://doi.org/10.1111/j.1467-9280.2006.01699.x

Kahane, G., Everett, J. A. C., Earp, B. D., Caviola, L., Faber, N. S., Crockett, M. J., & Savulescu, J. (2018). Beyond sacrificial harm: A two-dimensional model of utilitarian psychology. *Psychological Review*, *125*(2), 131–164. https://doi.org/10.1037/rev0000093

Kline, R. (2016). *Principles and Practice of Structural Equation Modeling*.

Law, K. F., Campbell, D., & Gaesser, B. (2022). Biased Benevolence: The Perceived Morality of Effective Altruism Across Social Distance. *Personality and Social Psychology Bulletin*, 01461672211002773. https://doi.org/10.1177/01461672211002773

MacAskill, W. (n.d.). *What is longtermism?* Retrieved May 5, 2023, from https://www.bbc.com/future/article/20220805-what-is-longtermism-and-why-does-it-matter

MacAskill, W. (2022). *What We Owe The Future*. https://whatweowethefuture.com/

McLamb, E. (2022, December 15). Top Five Threats Facing Earth & Humanity. *Ecology Prime*. https://ecologyprime.com/2022/12/15/earths-top-five-threats-facing-humanity/

McManus, R. M., Mason, J. E., & Young, L. (2021). Re-examining the role of family relationships in structuring perceived helping obligations, and their impact on moral evaluation. *Journal of Experimental Social Psychology*, *96*, 104182. https://doi.org/10.1016/j.jesp.2021.104182

Ord, T. (2021). *The Precipice: Existential Risk and the Future of Humanity a book by Toby Ord*. Hachette Books. https://bookshop.org/p/books/the-precipice-existential-risk-and-the-future-of-humanity-toby-ord/14906429

Schönbrodt, F. D. & Perugini, M. (21013). At what sample size do correlations stabilize? *Journal of Research in Personality, 47*(5),609-612. https://doi.org/10.1016/j.jrp.2013.05.009

Singer, P. (2015). *The Most Good You Can Do: How Effective Altruism Is Changing Ideas About Living Ethically* (1st edition). Yale University Press.

Singer, P. (2016). *Famine, Affluence, and Morality*. Oxford University Press.

Skitka, L. J. (2010). The Psychology of Moral Conviction. *Social and Personality Psychology Compass*, *4*(4), 267–281. https://doi.org/10.1111/j.1751-9004.2010.00254.x

Shrum, T.R. (2021). The salience of future impacts and the willingness to pay for climate change mitigation: an experiment in intergenerational framing. *Climatic Change 165*,18. https://doi.org/10.1007/s10584-021-03002-6

Strathman, A., Gleicher, F., Boninger, D. S., & Edwards, C. S. (1994). The consideration of future consequences: Weighing immediate and distant outcomes of behavior. *Journal of Personality and Social Psychology, 66*(4), 742–752. https://doi.org/10.1037/0022-3514.66.4.742

Syropoulos, S., & Markowitz, E. M. (2021). Perceived responsibility towards future generations and environmental concern: Convergent evidence across multiple outcomes in a large, nationally representative sample. *Journal of Environmental Psychology*, *76*, 101651. https://doi.org/10.1016/j.jenvp.2021.101651

Syropoulos, S., & Markowitz, E. M. (2022). Perceived responsibility to address climate change consistently relates to increased pro-environmental attitudes, behaviors and policy support: Evidence across 23 countries. *Journal of Environmental Psychology*, *83*, 101868. https://doi.org/10.1016/j.jenvp.2022.101868

Tamir, D. I., & Mitchell, J. P. (2011). The Default Network Distinguishes Construals of Proximal versus Distal Events. *Journal of Cognitive Neuroscience*, *23*(10), 2945–2955. https://doi.org/10.1162/jocn_a_00009

Taylor, R. G., Scanlon, B., Döll, P., Rodell, M., van Beek, R., Wada, Y., Longuevergne, L., Leblanc, M., Famiglietti, J. S., Edmunds, M., Konikow, L., Green, T. R., Chen, J., Taniguchi, M., Bierkens, M. F. P., MacDonald, A., Fan, Y., Maxwell, R. M., Yechieli, Y., … Treidel, H. (2013). Ground water and climate change. *Nature Climate Change*, *3*(4), Article 4. https://doi.org/10.1038/nclimate1744

Tuen, Y. J., Bulley, A., Palombo, D. J., & O'Connor, B. B. (2023). Social value at a distance: Higher identification with all of humanity is associated with reduced social discounting. *Cognition*, *230*, 105283. https://doi.org/10.1016/j.cognition.2022.105283

Wilks, M., McCurdy, J., & Bloom, P. (n.d.). Who gives? Characteristics of those who have taken the Giving What We Can pledge. *Journal of Personality*, *n/a*(n/a). https://doi.org/10.1111/jopy.12842

Zaval, L., Markowitz, E. M., & Weber, E. U. (2015). How Will I Be Remembered? Conserving the Environment for the Sake of One's Legacy. *Psychological Science, 26*(2), 231–236. https://doi.org/10.1177/0956797614561266

Zhao, Z., Mildner, J. N., & Tamir, D. I. (2020). Successful simulation requires bridging levels of abstraction. *The Behavioral and Brain Sciences*, *43*, e152. https://doi.org/10.1017/S0140525X19003133