

Annual report 2022 /

_Legal Priorities Project

Alfredo Parra_
Christoph Winter_





Table of Contents

Summary	2
Introduction	3
Review of 2022	5
Team	5
Board	5
Chart	6
Strategic research	7
Brief overview	7
Academic papers	7
Blog posts	9
Other research	9
Overall assessment	10
Field building	11
Legal Priorities Summer Institute 2022	11
Writing competition	13
Multidisciplinary Forum on Longtermism and the Law	13
Global Priorities Research Across the Disciplines	15
Other achievements	15
Statistics	17
Overall assessment	18
Key organizational successes	18
Issues and mistakes	19
Other potential issues and mistakes	20
Priorities for 2023	21
Methodology for updating our priorities	21
Strategic research	22
Selected ongoing research activities	22
Other selected ongoing research activities	24
Future research	24
Field building	29
Policy priorities for 2023	30
Other priorities	30
How to support us	31



Summary

Note: A private version of this report with additional confidential updates is available for major donors and close collaborators. If you fall into this category and would like to access the extended version, please get in touch with us at hello@legalpriorities.org.¹

This report presents a review of the Legal Priorities Project's work in 2022, including key successes, statistics, bottlenecks, and issues. We also offer a short overview of our priorities for 2023 and briefly describe our methodology for updating our priorities. You can learn more about how to support our work at the end of the report.

In 2022:

- Our research output per FTE was high to very high: With only 3.6 FTE researchers, we had 7 peer-reviewed papers (2 journal articles and 5 book chapters) accepted for publication, 3 papers currently under review, and one book under contract with Oxford University Press. We also added 7 papers to our Working Paper Series (for a total of 18), published a new chapter of our research agenda, and published 6 shorter pieces in online forums. We also spent significant time on reports aimed at informing our prioritization on artificial intelligence and biosecurity in particular (which we plan to publish in Q2 of 2023) and ran a writing competition on “Improving Cost-Benefit Analysis to Account for Existential and Catastrophic Risks” with a judging panel composed of eminent figures in law. Based on our experience, our research output was much higher than typical legal academic research groups of similar size.
- Beyond academic research, we analyzed ongoing policy efforts, and our research received positive feedback from policymakers. Relevant feedback and discussions provided valuable insight into what research would support decision-making and favorable outcomes, which we believe improved our prioritization.
- We experimented with running several events targeting different audiences and received hundreds of applications from students and academics at top institutions worldwide. Some participants have already reported significant updates to their plans as a result. Feedback on our events was overwhelmingly positive, and we gained valuable information about the different types of programs and their effectiveness, which will inform future events.
- Team morale remained high, including during stressful developments, and our operations ran smoothly.

¹ Many thanks to Suzanne van Arsdale, José Villalobos, Matthijs Maas, Laurens Prins, Mackenzie Arnold, Cullen O’Keefe, Jade Leung, Eric Martínez, and John Bliss for reviewing different parts of this report.



In 2023:

- Our research will increasingly focus on reducing specific types of existential risk based on concrete risk scenarios, shifting more focus toward AI risk. While this shift started in 2022, AI risk will become more central to our research in 2023. We will publish an update to our research agenda and theory of change accordingly.
- We will continue to publish research of various types. However, we will significantly increase our focus on non-academic publications, such as policy/technical reports and blog posts, in order to make our work more accessible to policymakers and a wider audience. As part of this strategy, we will also launch a blog featuring shorter pieces by LPP staff and invited researchers.
- We would like to run at least one, but ideally two, flagship field-building programs: The Legal Priorities Summer Institute and the Summer Research Fellowship.
- We will seek to raise at least \$1.1m to maintain our current level of operations for another year. More optimistically, we aim to increase our team size by 1–3 additional FTE, ideally hiring a senior researcher with a background in US law to work on risks from advanced artificial intelligence.

Introduction

The Legal Priorities Project is an independent, global research and field-building project founded by researchers at Harvard University. We conduct strategic legal research that mitigates existential risk and promotes the flourishing of future generations, and we build a field that shares these priorities.

1. **Strategic Research:** We conduct strategic research that generates decision-relevant insights for key actors (such as policymakers, judges, and litigators), focusing on mitigating existential risk and promoting the flourishing of future generations. Some research questions tackled in 2022 included:
 - How do people ordinarily understand and use the term “existential risk,” and how might such usage inform the regulatory efforts of lawmakers?
 - To what extent do people across cultures endorse legal protection for (far) future generations?
 - How feasible are strategies to delay or restrain technological development to reduce potential risks from emerging technologies (especially advanced AI)?



- What are different theories of change governing the impacts of increasingly transformative AI? What is the role of legal interventions under different scenarios and assumptions?
 - Which existing rules of international law address existential risk, and what are the obligations they impose on States?
 - Which legal strategies (such as litigation) should one prioritize to reduce global catastrophic and existential biological risk?
 - What can legal advocates working to mitigate existential risk learn from the leading models of legal and social change?
2. **Field Building:** We aim to strategically raise awareness about our concern for the long-term future and build a field that shares our priorities. We want to bring together individuals at every level of seniority and support the most effective ways to use the law to protect future generations and mitigate existential risk. Field-building initiatives carried out in 2022 included:
- A 1-week, in-person Summer Institute in Oxford for law and policy students and early-career professionals.
 - A writing competition on cost-benefit analysis applied to global catastrophic and existential risk.
 - An academic conference on Law & Longtermism co-organized with the University of Hamburg.
 - A workshop on global priorities research in law co-organized with the Global Priorities Institute at Oxford University.

Our teammates are affiliated with top law schools and private research centers worldwide, including Harvard University, Massachusetts Institute of Technology (MIT), New York University (NYU), Instituto Tecnológico Autónomo de México (ITAM), Strathmore University, Cambridge University, Oxford University, and OpenAI.



Review of 2022

Team

Throughout 2022, we had 6.1 FTE² on our core team, including employees and highly-involved contractors. As of January 1, 2023, we are a core team of 6.7 FTE. In 2022, we hired:

- Laurens Prins as Head of Legal Strategies;
- Mackenzie Arnold as Head of Strategy & Community Building (currently on leave completing a judicial clerkship);
- Matthijs Maas as Senior Research Fellow in Law & AI; and
- José Jaime Villalobos as International Law & Policy Fellow.

Additionally:

- John Bliss is doing a 1-year residency at LPP as Senior Research Fellow.
- Noam Kolt joined as Research Affiliate.
- Judge William Garfinkel joined as Advisor.
- Jonas Schuett transitioned from LPP to GovAI but remained a Research Affiliate and board member at LPP.
- Jade Leung joined our board of directors.

In 2022, we spent about \$1.1mm. We spent significant time fundraising and raised ~\$1.3mm. Between LPP's inception in late 2019 and the end of 2022, we applied for 14 grants (ten fully granted, three partly granted, and one pending).

Board

Our board of directors consists of

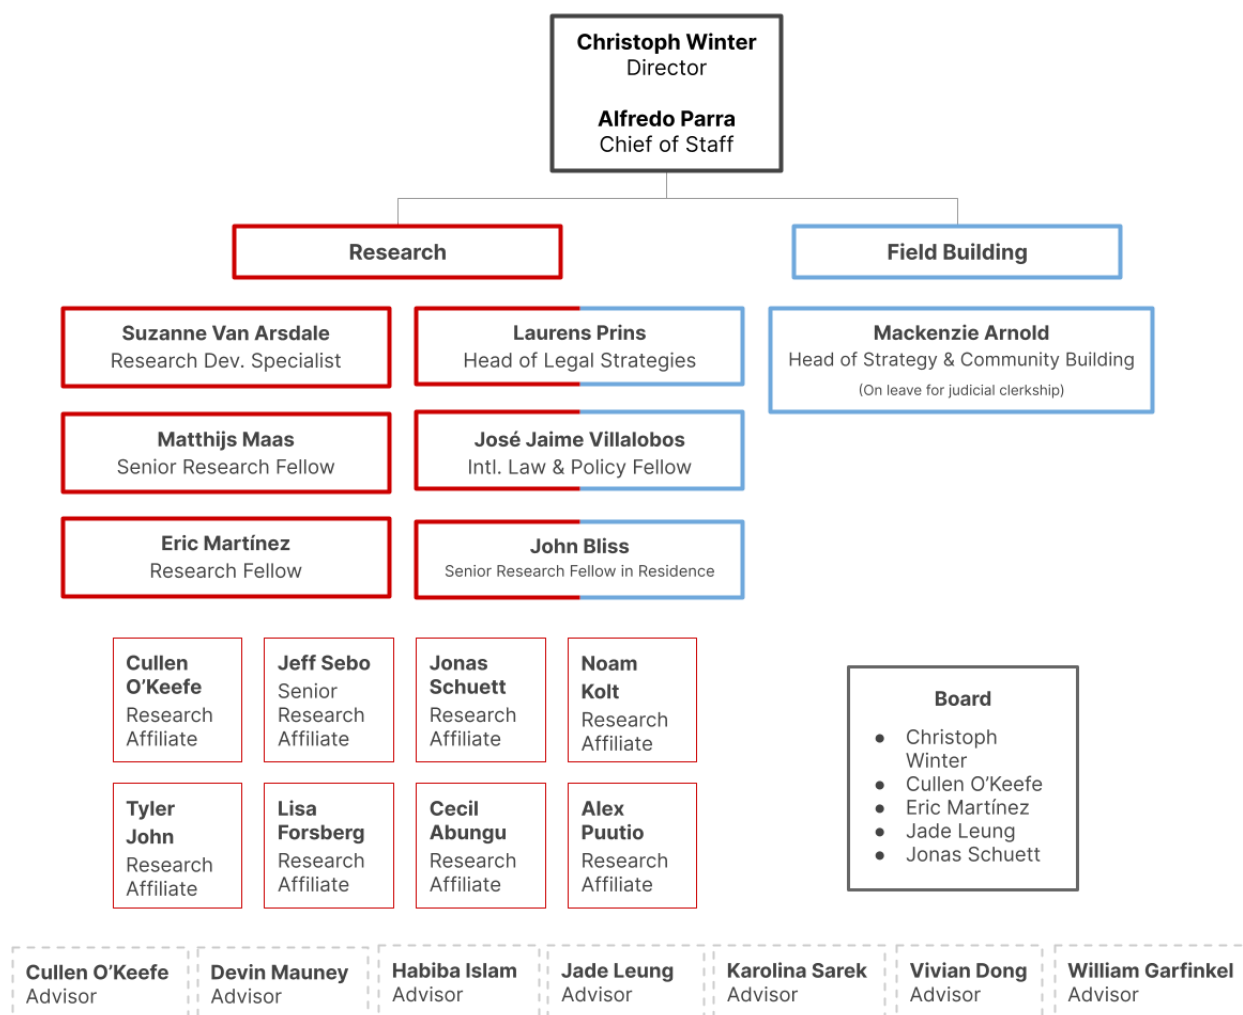
- Christoph Winter (LPP, ITAM, Harvard)
- Cullen O'Keefe (OpenAI)

² Research: 3.6 FTE; Field Building: 1.1 FTE; Operations & Management: 1.4 FTE



- Eric Martínez (LPP, MIT)
- Jade Leung (OpenAI)
- Jonas Schuett (GovAI)

Chart





Strategic research

Brief overview

We had a productive year despite the small size of our research team (3.6 FTE in 2022). Throughout 2022, we had 7 peer-reviewed papers (2 journal articles and 5 book chapters) accepted for publication, 3 papers under review,³ and one book under contract with Oxford University Press. We also added 7 papers to our Working Paper Series (for a total of 18), and published a new chapter of our research agenda as well as 6 shorter pieces in online forums, including a symposium on Verfassungsblog. We also made substantial progress on a few larger projects we aim to publish in 2023, and spent a significant amount of time on reports aimed at informing our prioritization on artificial intelligence and biosecurity in particular (which we plan to publish in Q2 of 2023). Below we list individual research items, including academic papers, blog posts, and other research, before we evaluate our research output for 2022.

Academic papers

- **The ordinary meaning of existential risk**
Eric Martínez and Christoph Winter
Under review
- **The intuitive appeal of legal protection for future generations**
Eric Martínez and Christoph Winter
Forthcoming in J. Barrett, D. Thorstad, & H. Greaves (Eds.), Essays on Longtermism, Oxford University Press
- **Value alignment for advanced artificial judicial intelligence**
Christoph Winter, Nick Hollman, and David Manheim
American Philosophical Quarterly
- **Existential advocacy**
John Bliss
Forthcoming in Georgetown Journal of Legal Ethics
- **Cross-cultural perceptions of rights for future generations**
Eric Martínez and Christoph Winter
Forthcoming in K. Tobia (Ed.), Cambridge Handbook of Experimental Jurisprudence, Cambridge University Press
- **Three lines of defense against risks from AI**

³ Of which one has already been accepted in 2023.



Jonas Schuett
Under review

- **Extreme climate change**
Daniel Bertram
New upcoming chapter of the LPP Research Agenda
- **Catastrophic risk, uncertainty, and agency analysis**
Alasdair Phillips-Robins
1st prize of LPP's writing competition
- **Catastrophic uncertainty and regulatory impact analysis**
Daniel Farber
2nd prize of LPP's writing competition
- **Catastrophic risk review**
Michael Livermore
2nd prize of LPP's writing competition

Two of the papers we worked on primarily in 2021 were accepted for publication in 2022:

- **Protecting future generations: A global survey of legal academics**
Eric Martínez and Christoph Winter
Accepted at multiple generalist US law reviews (still under review at others)
- **Experimental longtermist jurisprudence**
Eric Martínez and Christoph Winter
Forthcoming in S. Magen and K. Prochownik (Eds.), Advances in Experimental Philosophy of Law, Bloomsbury Academic

In addition, Matthijs Maas finished two chapters he primarily worked on before joining LPP, both of which were accepted for publication in late 2022:

- **Military artificial intelligence as a contributor to global catastrophic risk**
Matthijs Maas, Kayla Lucero-Matteucci, and Di Cooke
Forthcoming in S. J. Beard, Lord M. Rees, C. Richards, and C. Ríos-Rojas (Eds.), The Era of Global Risk, Open Book Publishers
- **From evaluation to action: Ethics, epistemology and extreme technological risk**
Lalitha Sundaram, Matthijs Maas, and S. J. Beard.



Forthcoming in C. Rhodes (Ed.), Managing Extreme Technological Risk, World Scientific

Blog posts

- **Paths untaken: The history, epistemology and strategy of technological restraint, and lessons for AI**
Matthijs Maas
Verfassungsblog symposium on “Longtermism and the Law”
- **In Dubio Pro Futura: A novel decision rule to preserve the long-term interests of humanity**
Christoph Bublitz and Christoph Winter
Verfassungsblog symposium on “Longtermism and the Law”
- **Is legal longtermism common sense?**
Eric Martínez and Christoph Winter
Verfassungsblog symposium on “Longtermism and the Law”
- **Bridging legal longtermism and animal law**
Andrew Stawasz and Jeff Sebo
Verfassungsblog symposium on “Longtermism and the Law”
- **The rise of the constitutional protection of future generations**
Renan Araújo and Leonie Koessler
Verfassungsblog symposium on “Longtermism and the Law”
Based on LPP Working Paper No. 7–2021
- **Strategic perspectives on long-term AI governance, parts 1 and 2**
Matthijs Maas
Effective Altruism Forum

Other research

Here, we list a few other notable research projects and successes.

- We made significant updates to our research agenda published in early 2021. We hope to share these updates in early 2023.
- Laurens Prins surveyed biosecurity experts to identify the most promising legal interventions to address biorisks. He is summarizing the results in an internal report.



- Matthijs Maas signed a contract with Oxford University Press to publish a monograph on Architectures of AI Governance based on his PhD thesis (more below).
- Christoph Winter finished a book on utilitarianism and criminal law based on his PhD that we expect to be published in 2023.
- José Villalobos finished an internal report on long-term impact litigation. José started this project during the Summer Research Fellowship 2021.
- We analyzed and conducted research for various ongoing policy efforts, such as the Maastricht Principles on the Human Rights of Future Generations, the UK Wellbeing of Future Generations Bill, and the UN Declaration on Future Generations. This research is not (yet) public. We are particularly excited about the upcoming Maastricht Principles, which consolidate existing international obligations related to the mitigation of existential risk (e.g., forecasting, early warning systems, and disaster strategies; reduction of risks from nuclear weapons, biological weapons, unaligned artificial intelligence, and totalitarianism) into a single document.
- We received positive feedback from policymakers in multiple jurisdictions about the usefulness of our research for their efforts.

Overall assessment

We are pleased with the quantity and quality of our research output per FTE in 2022, especially given that some teammates only joined the organization in 2022 and that we organized and participated in many events. Our current team works very well together and is willing to reprioritize quickly based on new insights thanks to a strong, shared commitment to our core values. Based on our experience, our research output was much higher than typical legal academic research groups of similar size and seniority. We believe this represents significant progress toward establishing longtermist themes in the legal academic literature. We also identified some actionable legal levers that inform decisions in various existential risk domains (such as AI development and biosecurity) and will prioritize this more highly in 2023.

Our main bottleneck has been hiring senior legal researchers (at least at the postdoc level but ideally higher) who are already familiar with longtermism and/or existential risk. We announced three open positions in early 2022 but made only one hire. Despite receiving many applications from highly qualified individuals, most applicants are still very new to our priorities, which limits the opportunities we can offer them.

Discussions with policymakers and others who are more closely integrated with policymaking processes have been very beneficial from numerous perspectives, helping us prioritize better toward Q3 and Q4 of 2022. While we are uncertain whether we should have started these conversations earlier, we are confident that we should



increase these in the future to improve relevant feedback mechanisms (see “[Methodology for updating our priorities](#)”). To read more about our future projects, see “[Future research](#).”

Field building⁴

Legal Priorities Summer Institute 2022

In 2022, we introduced a new flagship summer program: The [Legal Priorities Summer Institute](#). The goal of the Institute is to reach a larger audience with less previous exposure to legal prioritization and to offer them an intensive and immersive learning experience. This program was partly motivated by our experiences and evaluation of previous events and by Open Philanthropy’s [survey](#) pointing out the strong impact of in-person events on fostering engagement.

The one-week program attracted 473 applicants (of these, 200 applicants (42%) were from the [top 20 law schools worldwide](#), and 113 (24%) were from the top 5; 402 (85%) had very little or no previous exposure to longtermism; for more on application statistics, see “[Statistics](#)”). We admitted [35 talented law students and graduates](#) from 18 countries and five continents. This resulted in an extraordinarily diverse cohort. When asked: “From 1 to 10, how likely is it that you would recommend the Institute to a colleague with similar interests to yours?” the average response was **9.1**.

Four months after LPSI, what positive outcomes can you trace back to the event?

I started a working group on law and AI as a way to start exploring AI use and risks of using LLMs in the law—while not directly AI safety, it’s been a step to even start introducing many lawyers to it.

I started a fellowship at LPP, joined initiatives shared in the Slack channel, and connected with other members of the community.

I got looped into a very exciting project right after LPSI related to bio policy.

LPSI assisted me in getting a research assistant position with a barrister and professor; I readjusted my career path and goals to pursuing postgrad research in law and tech before going into practice/politics; and I made contacts that I’ve reached out to for advice in career planning.

I participated in Mexico EA Fellowship and worked on diversity issues; I continued working on algorithmic discrimination research; and I made meaningful connections with others from LPSI.

⁴ You can read more about some of our field-building efforts since 2020 [here](#).



I declined a job offer from a corporate law firm and am currently seeking policy-oriented work instead; I chose to write my master's dissertation on a topic that I hope has the potential to be impactful, relating to decision-making in collective action problems.

I met a handful of people I may try to collaborate with on projects in the future, and I was introduced to some work that I had not previously known about.

Other testimonials:⁵

I have taken many actions already as a result of LPSI. These include a research collaboration with 5 colleagues and career planning. I started to draft job possibilities and the pros and cons of which of them. I have also noticed that I had many limiting thoughts that were making me less confident in my capabilities. I was able to reflect on those with colleagues during the summer institute, which helped me greatly.

The event has given me the confidence and awareness I needed to shift from the career path I was on of training at a corporate law firm to pursuing jobs in policy research and advocacy.

Thanks for all the hard work you all put into this. I had a great time—having five days set aside just to think about these issues and connect with people without worrying about anything else was invaluable. To have it all so well organized and thought through, especially with very generous, responsive, and flexible organizers, was fantastic.

LPSI gave me an opportunity to meet an amazing group of people. I met a group of people from the Global South, which has motivated all of us to bring forth Global South perspectives and collaborate on ways to bring these perspectives to the global avenue. Getting feedback and giving feedback to people was also a very valuable experience. It is impressive how much clarity you can get on ideas when you share them with people. I also formed a network of people who are now my friends—we keep in touch and give each other feedback and support. I am grateful for this.

LPSI was a highlight of my year! It really did help me rethink my career strategy, and I've decided to go by field and cast a broader net for roles. So I'm reaching out to folks in AI safety companies to get a sense of what might match / what I could grow into / what shortcomings I might need to address.

⁵ Submitted via an anonymous feedback form right after the event.



LPSI 2022 participants

Writing competition

This year, we also ran a writing competition on “Improving Cost-Benefit Analysis to Account for Existential and Catastrophic Risks.” The judging panel was composed of eminent figures, including Prof. Benjamin Eidelson (Harvard Law School), Prof. Arden Rowell (University of Illinois), and Rich Theroux (retired Senior Economist and Branch Chief at the OIRA). We awarded prizes to six papers that contributed to our understanding of how cost-benefit analysis can tackle critical questions surrounding future generations and existential and catastrophic risks, with prizes ranging from \$2,500 to \$30,000. We published the three winning entries in our Working Paper Series. The winning entry by Alasdair Phillips-Robins was inspired by his summer fellowship at LPP in 2021, which explored ways to represent the interests of future generations in the US democratic process. The submissions were anonymized and LPP staff did not review them for the purposes of the competition.

Multidisciplinary Forum on Longtermism and the Law

In June, we organized the first-ever conference on Longtermism and the Law, co-hosted with Amal Sethi (University of Hamburg), who was inspired to run the event after engaging with longtermism during our 2021 Summer Research Fellowship. 30 participants from around the world discussed legal strategies to tackle issues relevant to the long-term future, including pandemics and artificial intelligence. Selected submissions were published as a blog symposium on Verfassungsblog, an influential peer-reviewed academic forum on constitutional law and policy. Participant feedback



was excellent. When asked: “From 1 to 10, how likely is it that you would recommend the forum to a colleague with similar interests to yours?” the average response was **8.9**.

Eight months after the Forum, what positive outcomes can you trace back to the event?

I met several prominent people in the law and longtermist space that I had never met before. The conversations I had with them made me think more deeply about how best to positively impact the long-term future, and I have stayed connected with some of them after / as a direct consequence of the event. I also published a piece that directly grew out of the blog post and presentation for the event. Presenting and getting feedback on my post helped me think better about how best to frame longtermist ideas to a general audience.

We have received valuable feedback on our work on "Negative Human Rights as a Basis for Long-term AI Safety and Regulation" and incorporated it into the resulting paper, which was now accepted for publication in the Journal of Artificial Intelligence Research.

The two pieces I published in Verfassungsblog became full papers, which are now submitted.

Good impact—I submitted a grant and published a paper (and more in the pipeline).

While finishing my PhD project in 2022, I have been thinking a lot about the concepts of longtermism and legal longtermism. Although I did not shift my research fundamentally (as I had already almost finished it at the time of the Forum), I included some interesting and focused ideas I got during the Forum in my introduction and some of my chapters and footnotes.

This event helped connect me to some AI governance researchers, which has been helpful.

Other testimonials:⁶

The event was really great; I enjoyed it so much! Lots of brilliant people and plenty of time to connect and discuss. I have no complaints, only praise.

I received many good ideas and criticism, food for thought for further research, and helpful changes and clarifications that I could include in my research.

I think everything was excellent. I especially appreciated having a lot of time for discussion and socializing.

⁶ Submitted via a feedback form right after the event.



The event was very useful for figuring out which people might be interesting partners for collaboration and whose work I'd like to know more about.

Overall, I was very impressed with the organization. Everyone was easy to talk to, and it felt like I had an opportunity to speak to everyone without it being too small. Very excellent, in my view.



Participants of the Multidisciplinary Forum on Longtermism and the Law

Global Priorities Research Across the Disciplines

We co-organized a workshop at the Global Priorities Institute on global priorities research in law as part of a larger workshop on “Global Priorities Research (GPR) Across the Disciplines” that included psychology, history, and law. We benefited greatly from discussions with GPR researchers in other fields, representing the dialogue we hope to preserve as the field grows.

Other achievements

Here are a few other notable successes in field building:

- We launched the Law & Longtermism Network. It had roughly 300 members at the end of 2022.
- Our newsletter had ~1,000 subscribers at the end of 2022.
- We paid small stipends to four student initiatives to organize reading groups at ITAM on topics related to legal longtermism and effective altruism. At least two group organizers who weren't already familiar with longtermism are now doing paid direct EA work in Latin America. One became a “Next Generation Fellow on



Future Generations” at the United Nations Foundation and a “Future Generations Global Ambassador” at the Office of the Future Generations Commissioner for Wales. Given the small effort required on our end, we are excited about the returns of such investments.

- Multiple teammates gave talks on law and longtermism. Some examples include:
 - Cullen O’Keefe, *Governing artificial intelligence*, Harvard Law School
 - John Bliss, *Existential advocacy*, Harvard Law School and UC Berkeley School of Law
 - Christoph Winter, *Global priorities research in law*, Oxford University
 - Eric Martinez, *Protecting future generations: A global survey of legal academics*, Multidisciplinary Forum on Longtermism and the Law, University of Hamburg
 - John Bliss, *Legal activism in the age of existential risk*, Multidisciplinary Forum on Longtermism and the Law, University of Hamburg
 - Matthijs Maas, *Paths untaken: The history, epistemology and strategy of technological restraint (...and lessons for AI)*, Multidisciplinary Forum on Longtermism and the Law, University of Hamburg
 - Matthijs Maas, *Strategic perspectives on long-term AI governance*, Centre for the Governance of AI
 - Laurens Prins, *Developing an impact litigation strategy to limit catastrophic biorisk*, Legal Priorities Summer Institute
 - José Villalobos, *LPP’s prioritization of international law projects*, Legal Priorities Summer Institute, Oxford
- Christoph Winter taught longtermism-relevant courses at ITAM, including on “Law and Artificial Intelligence” and “Law & Effective Altruism,” with an average evaluation of 4.8/5.
- Multiple team members provided ongoing mentoring to law students and professionals working on longtermist issues, including at the Leverhulme Centre for the Future of Intelligence and Magnify Mentoring.
- Matthijs Maas contributed to organizing the Niels Bohr Open World Conference 2022, a high-level interdisciplinary academic conference that focused on trade-offs around scientific openness and the potential for misuse or danger in domains such as the nuclear arms race and AI.



- We organized two informal dinners during the EA Global conferences we attended (in Cambridge and DC). Attendees were primarily other EA Global participants with a background in law and policy.



Law & Longtermism dinner with attendees of EA Global Bay Area 2023

Statistics

Our programs and open positions attracted significant interest from students and academics at different levels of seniority in 2022.

- **Legal Priorities Summer Institute:** We received 473 applications (227 undergrads, 131 master's, 29 JDs, 65 PhDs/JSDs, and 21 others). Of these, 200 applicants (42%) were from the top 20 law schools worldwide and 113 (24%) were from the top 5 (including 94 from Oxford). 402 applicants (85%) were either new or had very little previous exposure to longtermism. 35 participants attended, of which 14 (40%) were either new or had very little previous exposure to the ideas of longtermism.
- **Senior Research Fellows:** In early 2022, we advertised one visiting and three permanent positions for senior researchers. We received 79 applications for these positions (3 full professors, 7 associate professors, 28 assistant professors,⁷ 1 adjunct professor, 19 postdocs, 14 PhD researchers, and 7 others). We made only one offer and hire (Matthijs Maas).
- **Forum on Longtermism & the Law:** 67 researchers applied to participate in the Forum (1 emeritus professor, 5 full professors, 7 associate professors, 12

⁷ This number also includes the UK and Australian equivalent of “Assistant Professor” in the U.S. and elsewhere, namely “Lecturer” and “Senior Lecturer” (20).



assistant professors,⁸ 4 postdocs, 14 PhD researchers, 3 LLM students, 2 undergrads, and 19 others). We admitted 29 participants.

- **Expression of interest (research positions only):** 59 researchers filled out our expression of interest form (2 full professors, 1 associate professor, 2 assistant professors, 2 lecturers, 4 postdocs, 14 PhD researchers, 34 others). While we did not make any offers or hires based on these submissions, the new connections have resulted in various collaboration opportunities and additional leads for future positions.

Overall assessment

We managed to organize high-quality events with very limited staff capacity (1.1 FTE). We attracted significant interest from people from highly prestigious institutions worldwide, especially given how young LPP still is. We are particularly excited about the fact that some former staff and participants of our programs have already gone on to (co-)organize their own programs and initiatives, such as the ILINA Fellowship in Nairobi (Kenya) by Cecil Abungu, the Condor Camp in Cusco (Peru) by Renan Araújo, the EA Mexico Fellowship in Mexico City by Claudette Salinas, the Forum on Law and Longtermism in Hamburg (Germany) by Amal Sethi, and a working group on law and AI by Janna Tay. Similarly, other former team members and program participants have gone on to pursue impactful career opportunities related to legal longtermism after encountering these ideas through LPP (or at least having the opportunity to explore these ideas in depth for the first time after encountering them elsewhere). These include Jonas Schuett and Leonie Koessler (both now at the Centre for the Governance of AI), Renan Araújo (Research Manager at Rethink Priorities), Kaustubh Chaturvedi (Open Philanthropy grantee researching impact litigation strategies in South and Southeast Asia), and Taís Fernanda Blaúth (LPP grantee researching existential risk in international law).

The Summer Institute was a successful experiment that we would like to repeat and possibly scale up. However, we paused other outreach and field-building efforts after September due to staff capacity. We will prioritize solving this bottleneck in early 2023.

Key organizational successes

Meta success

- We continued to build a well-functioning organization with highly-qualified and dedicated people willing to reprioritize anytime based on new evidence.

Research

⁸ Of which 5 “Lecturers” or “Senior Lecturers”.



- Our research output per FTE was high to very high.
- We learned a lot from the feedback from policymakers and policy-integrated researchers, which improved our prioritization.
- Our research informed some policy efforts, most notably the upcoming Maastricht Principles on the Human Rights of Future Generations.

Field building

- We managed to run high-quality events with very limited staff capacity and caught the interest of many people outside the existing longtermist community.
- Program participants have gone on to pursue impactful opportunities in the legal longtermist space, including receiving grants to carry out research or to launch their own initiatives (see [statistics](#) and [overall assessment](#)).
- Several former (junior) LPP teammates (who started engaging with longtermism at LPP) secured impactful jobs at organizations such as GovAI and Rethink Priorities.

Organizational health

- Team morale at the end of 2022 was very high.
- Our operations ran very smoothly.
- We supported the productivity and health of our team with professional development resources, workplace accommodations, and leave policies. We were able to support five teammates who had newborn children in 2022 by offering them generous parental leave or other types of support.
- We took stock of our internal policies and implemented new ones, most notably a security policy to minimize cybersecurity (and other) risks.

Issues and mistakes

- Hiring has been challenging due to a lack of senior researchers familiar with longtermism or existential risk.
- Funding for our Summer Institute became available only in late April, so we organized the event with extremely tight deadlines.
- The results of our writing competition might not have the expected impact due to timing and coordination difficulties.



- We hired two part-time employees we let go shortly after, one during the trial period and one due to a mutually recognized lack of optimal fit with the scope of responsibilities.
- Discussions over LPP's theory of change extended over several months, sometimes making it difficult for teammates to connect their goals to the organization's goals.
- We didn't have sufficient capacity to follow up more closely with participants and other applicants of our Summer Institute after the event.
- We tried to publish an updated research agenda before our Summer Institute in September but delayed the publication repeatedly, causing some stress to the authors.
- Our website lacks functionalities that would make it easier for our team to make essential updates easily.

Other potential issues and mistakes

- We have been largely inactive on social media to focus on other priorities.
- Despite the Forum on Longtermism & the Law being very well evaluated, we think it might have been too broad in scope, covering many longtermist areas. In the future, we will consider organizing events on more concrete questions, problems, and/or risks, such as the legal regulation of catastrophic risks from large language models.



Priorities for 2023

Below, we outline our priorities for the year ahead. To put our new priorities into context, we first sketch out our general approach to *updating* which activities to prioritize. For a more detailed description of the criteria we use to prioritize among research projects, see Chapter 3 (“Methodology”) of our [Research Agenda](#).

Methodology for updating our priorities

As an early-stage organization, we continue to explore various paths to impact and learn about our strengths, weaknesses, and comparative advantages before committing heavily to one strategy. We are aware that circumstances can change quickly and want to remain flexible to adapt our strategy as needed. Three factors in particular influence how we update our priorities:

1. Evaluation of our activities	2. Changes in the landscape	3. New insights from the prioritization community
Examples		
<ul style="list-style-type: none">• <i>Direct external feedback</i>, e.g., from key actors, editors, experts on our research, and participants of our field-building and outreach events on a predefined set of evaluation criteria.• <i>Indirect external feedback</i>, e.g., quality and quantity of applications to our various positions and events.• <i>Internal feedback</i> through 360 feedback reviews and transparent feedback based on our shared values in team meetings and frequent 1-on-1 calls with team members.	<ul style="list-style-type: none">• Changes regarding the expected development of transformative technologies, e.g., due to recent advances in AI.• Changes in funding opportunities.• Other organizations (successfully) tackling overlapping priorities.• Emerging opportunities for potentially large impact.	<ul style="list-style-type: none">• Normative insights, e.g., updates on the relative importance of different approaches to longtermism based on population-ethical considerations.• Empirical insights, e.g., greater strategic clarity on the relative importance of different governance bodies to mitigate AI risk.



Strategic research

We have several ongoing research projects that we aim to finish in 2023. That said, we expect to spend most of our research time on future work based on our updated theory of change.

Selected ongoing research activities

States must mitigate existential risk under international law

José Villalobos, Matthijs Maas, and Christoph Winter

This paper argues that States have an obligation to mitigate existential risk under international law. We discuss how existing legal frameworks address existential risk, how they interact with each other, and the practical implications of such interactions. We hope this will become a reference for relevant actors in national governments to understand their obligations to reduce existential risk.

Handbook on approaches to transformative AI governance

Matthijs Maas

This comprehensive report aims to provide strategic clarity in the transformative AI (TAI) governance community. It provides a mapping and in-depth profiles of 15 distinct approaches to TAI governance.⁹ For each profile, the handbook will provide examples of work or writing in that space, discuss its shared theory of victory and intermediate goals, survey historical cases in support (and as possible counterevidence) of the approach, and survey its various strengths and drawbacks. Components of this report have already received extensive input and positive feedback throughout the community, including at Rethink Priorities, GovAI, CSER, and elsewhere.

To supplement this handbook, we also plan to publish two ‘Foundation Reports’ that provide key background and reference materials: (1) *Disentangling definitions in transformative AI governance* and (2) *Transformative AI governance: A literature review*.

Experimental studies on the effectiveness of legal norms to reduce catastrophic risk

Eric Martínez, Christoph Winter, and Piotr Bystranowski

These studies aim to examine the efficacy of current and potential legal norms in addressing catastrophic risk. In these studies, real-world judges are presented with cases involving catastrophic risk to humanity. Judges are divided into groups, each tasked with resolving the same cases but provided with different laws. The goals of the

⁹ For a primer, see: Maas, Matthijs M. ‘[Strategic Perspectives on Transformative AI Governance: Introduction](#)’. EA Forum, 2 July 2022.



study are to test novel approaches regarding their ability to determine (1) which laws are most effective in reducing catastrophic risks across a variety of scenarios, including pandemics, artificial intelligence, and climate change, and (2) the likelihood that the judiciary will intervene in executive action in catastrophic risk situations.

Architectures of AI governance (book under contract with Oxford University Press, to be delivered August 2023)

Matthijs Maas

Most existing longtermist/TAI governance literature is limited to blogs, essays, reports, and a few journal articles. “Architectures of AI Governance,” based on Matthijs’ PhD dissertation, will be the first academic book from a highly accredited publisher offering a comprehensive case for the importance and urgency of TAI governance. We hope the book will be a valuable resource to incoming scholars and researchers who want to pursue work on AI governance both within and outside academia. Matthijs expects to spend 5–10% of his time finishing this project in 2023.

Foundational precaution: Existing legal requirements mandate safety in large language model development

Matthijs Maas and Christoph Winter

This project will focus on identifying existing legal principles, doctrines, and laws that could apply to (the regulation of) LLM development and deployment under short-timeline, prosaic AGI scenarios. The paper will investigate (1) whether or how application-level approaches to regulating the downstream impacts of AI systems suffice to address AI risks and failure modes that appear at the R&D and training stages; (2) what kind of improved safety assessments and evaluations would patch any gaps at these stages; and (3) whether or which existing legal principles enshrined in many jurisdictions, such as the precautionary principle, might already impose obligations for AI labs to adopt such best-practice safety evaluations, as well as conduct proactive alignment research. Finally, it will review (4) a number of downsides or risks from imposing such widespread precautionary obligations on LLM developers.

Legal strategies to reduce biorisk

Laurens Prins

This internal report will outline legal interventions to reduce global catastrophic and existential biological risk. It is based on a series of interviews with bio experts, which enabled us to identify key risk factors, ten intermediate objectives that could contribute to addressing these risk factors, and the key actors that would need to be engaged to reduce biorisk. The report explores findings from a subsequent survey of bio experts with a variety of backgrounds (e.g., science, public health, pharmaceuticals, law, policy,



and grantmaking), which asked them to prioritize the identified risk factors, intermediate objectives, and actors based on their relative potential to reduce biorisk.

Existential advocacy

John Bliss

This book presents a socio-legal study of the movement arising in response to existential risk. It details how advocates in this field approach core theoretical issues in the legal and social change literature. By comparing these advocates to other movements that have been central to the socio-legal canon, the book offers recommendations regarding the efficacy and ethics of “existential advocacy.”

The technology triad: Disruptive AI, regulatory gaps and value change

Jeroen Hopster and Matthijs Maas

This paper (submitted) draws on concepts from moral philosophy and technology law to develop an analytical framework for how particular disruptive technologies, especially advanced AI technologies, can challenge and overturn established moral and legal systems (potentially disrupting situations of value- or legal-technological lock-in).

Other selected ongoing research activities

Additionally:

- We are launching a **blog** in early 2023 featuring shorter pieces both by LPP staff and invited researchers.
- We are in conversations to publish **encyclopedia entries** on topics related to legal longtermism.
- We are experimenting with writing **short, accessible summaries** of our long-form papers.

Future research

As we get more information and our strategic situation changes, so should our priorities. Recent years have seen major changes and developments that call for a reassessment of our theory of change and will shift the balance of our future research. These developments include but are not limited to:

- A **significant shortening of expected timelines** to transformative AI based on recent breakthroughs and analyses from various lines of evidence.¹⁰

¹⁰ See also [Wynroe et al. \(2023\)](#).



- **Significant public, academic, and policymaker attention for AI risk:** Since the deployment of ChatGPT, Bing Chat/Sydney, and GPT-4, legal academia and policymakers are more motivated to take AI seriously. In some cases, this might translate to greater serious attention on misalignment risks, creating opportunities for productive AI governance. In other cases, these actors may now be more likely to act on AI in ways that recognize its power but do not take alignment seriously, which puts a greater premium on legal interventions that avoid bad outcomes or lock-in of hazardous policies.
- **Growing confidence around effective interventions to reduce AI risk:** recent years have seen greater convergence in relevant research communities on what would be robust interventions in the AI risk space (e.g., safety evaluations or compute governance). This can provide a starting point for LPP to find ways to contribute specific legal insights to AI risk reduction.
- **Greater public and policymaker wariness towards longtermism:** Conversely, since the FTX crisis, legal academia and policymakers are less receptive to work explicitly focusing on longtermism, which translates into less conducive conditions for our work aimed at establishing (legal) longtermism.
- **LPP success in the field of general protection for future generations:** Simultaneously, LPP has already achieved a notable amount of success in developing the case that legal protection for future generations is common sense and in advocating for this in various institutions and fora; as a result, we expect to hit diminishing returns on research insights and reputational gains in this space.

These changes have implications for our overall organizational strategy. They will be reflected in our priority setting in 2023 and beyond in at least the following ways:

- **Greater focus on existential risk:** LPP has always emphasized a focus on medium-range existential risk as a core part of our portfolio.¹¹ However, the above developments further increase the importance of our work on existential risks relative to our work on other long-term trajectory changes or broad longtermism.
- **Greater focus on AI risk:** The shifts we describe above increase the importance of our work on AI risk relative to our work on other existential risks.¹²

¹¹ See for instance our 2021 Research Agenda ([Winter et al. 2021](#)), our 2022 ‘Writing Competition: Improving Cost-Benefit Analysis to Account for Existential and Catastrophic Risks’ ([LPP, 2022](#)), and academic work on the ‘Ordinary Meaning of Existential Risk’ ([Martínez and Winter, 2022](#)), ‘Existential Advocacy’ ([Bliss, 2022](#)), and Jonas Schuett’s contributions to the 2021 Survey on AI existential risk scenarios’ ([Clarke et al. 2021](#)), amongst others.

¹² While our previous focus on existential risks in the medium-term still holds, our new focus on shorter AI timelines is an update we did not anticipate 2.5 years ago.



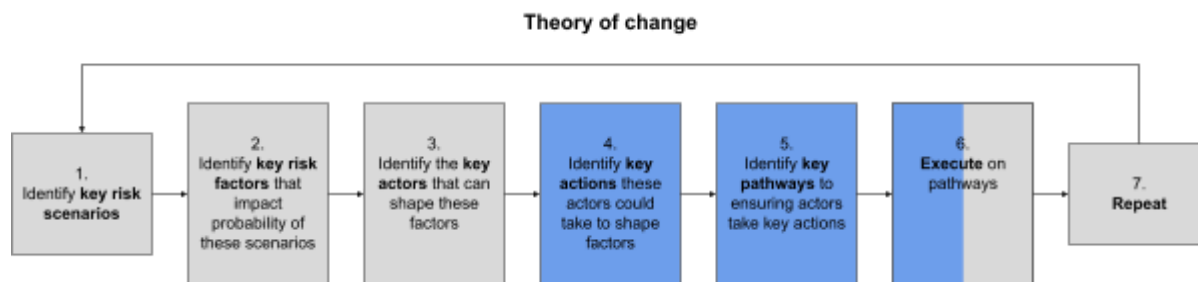
- **Greater focus on specific AI interventions:** The strategic shifts above increase the premium on delivering specific and action-relevant legal insights for AI governance in the near-term (e.g., to facilitate implementation of AI alignment at leading labs) compared to longer-range, indirect, or more theoretical academic research into AI governance.
- **Greater focus on US law:** Under short timelines, TAI/AGI is more likely to be developed by a US-based lab. This increases the relative importance for LPP to investigate ways to shape the US legal and regulatory environment affecting AI development (e.g., compute governance, Section 230, First Amendment rights, Major Questions Doctrine), shifting the center of gravity for our research. However, this does not mean that we will cease investigating relevant legal doctrines and regulatory levers in other jurisdictions¹³ or under international law. Instead, we will keep a close eye on ways in which legal doctrines and regulatory levers in other jurisdictions may offer leverage over various stages of the AI development cycle.
- **Shift towards providing insights through reports and primers:** Under short timeline scenarios, there is a greater premium on shifting LPP’s research outputs towards formats (such as reports) that can be more well-tailored to key decision-makers and which can be delivered faster compared to traditional academic venues. This obviously will not translate into any relaxation of rigor in our work, our internal project vetting procedures, or our engagement with pursuing relevant expert review and input.
- **Shift away from “field-building inspired research”:** In the past two years, we focused in part on what we refer to as “field-building inspired research” (e.g., law professor surveys, encyclopedia entry on legal longtermism). This research was based on a different theory of change and strategic situation, including longer AI timelines. It will therefore be somewhat decreased in future work. To avoid any potential confusion, this does not mean moving away from field-building activities (such as fellowships, etc) in general.
- **Shift away from research on general legal mechanisms to protect future generations:** We do not plan to start any new major research projects that focus on general legal protection offered to future generations, as we expect to hit diminishing returns. That said, given our history of success in this area, we may engage in opportunistic research and policy projects that fit with our overall theory of change.

Much of our future work will depend on the final updates to our theory of change. Given our increasing focus on AI risk, we will publish an applied theory of change specifically

¹³ For instance, one ongoing research paper that is nearing completion, ‘States Must Mitigate Existential Risk under International Law’, does focus on charting the US’ (and other states’) obligations around various technologies under international law.



for AI risk in Q2/2023. That said, future projects will be inspired by the following section, which, for the purposes of this report, has been (over)simplified and should be interpreted in light of the more general implications stated above.



As indicated by the highlighted blue boxes, we believe LPP has a comparative advantage in carrying out specific research projects that address Step 4 (and, to a lesser extent, Steps 5 and 6). Our outreach and field-building activities are crucial for encouraging more people to work particularly on Steps 4–6 at LPP or elsewhere. Similarly, collaborating closely with other organizations is especially important for us to gain more clarity on Steps 1–3 and 5–6.

Zooming in on Step 4 of our theory of change, with this step, we aim to (1) identify possible key actions (and omissions), (2) evaluate which of these actions the key actors should take (or avoid), and (3) determine which key actions they are legally allowed or obliged to take (or avoid) in order to mitigate key risk scenarios. This analysis involves considering both what well-designed laws and policies to reduce a risk scenario might look like and what actors might be able to do based on existing laws and policies.

By clarifying the legality of key actions or how one might achieve (previously identified proxy) goals to reduce existential risk most effectively while meeting existing legal obligations, this step involves some “traditional” legal work. Other work we aim to conduct may be less traditional and involves different methods to evaluate the effects of (potential) laws and policies. Among other things, we aim to experiment with the following concepts and methodologies to (1) increase our understanding of the effects of law on existential risk and (2) increase our capabilities to evaluate possible legal measures more quickly in future policy windows that we may not be able to anticipate now:

- **Weighing Cost-Benefit Analysis by the Most Important Case Hypothesis:** Cost-benefit analysis (CBA) is a robust decision rule used by different governments worldwide and enjoys bipartisan support in US agencies. However, in the context of existential risk scenarios, determining the benefits of different actions is complex, yet tends to be dominated by the effect that a given law has on the probability of existential risk and, more precisely, on the probability of a previously identified specific risk scenario materializing (see Step 1). In other words, the evaluation of the law would likely depend on how that law affects the



most important case (i.e., the specific risk scenario) it can influence—we call this the “Most Important Case Hypothesis.”¹⁴

- **Using forecasting to predict law’s effects on the most important cases:** Assuming that the overall benefits of a law or policy are largely determined by its suitability for reducing the probability of an existential risk scenario, we believe it would be helpful to experiment with forecasts of the effects of (potential) laws and policies on specific risk scenarios. While the Most Important Case Hypothesis simplifies CBA by collapsing complex effects into the most important among them,¹⁵ forecasting the effects of (potential) laws on specific risk scenarios is by no means an easy task. Especially in the early stages, we expect to experiment with different approaches to gain relevant insights quickly—some of which might not work, or the hypothesis may fail entirely for other reasons. Collaboration with forecasters is also crucial to succeed.
- **Using *ex ante* expert/ordinary meaning analysis to understand how laws might be interpreted and applied in the most important cases:** While there are many factors that forecasting might consider, one approach to understanding how a law might affect a risk scenario is to learn how policymakers and judges would likely apply said law. This would partly depend on how they interpret crucial legal terms in said law, such as “existential risk”.¹⁶ Experimental methods could help us understand how policymakers and judges will likely interpret crucial terms, both in abstract and in concrete risk scenarios, which could inform (1) how lawmakers should define these terms in the first place¹⁷ and (2) forecasters predicting how laws containing those terms might impact existential risk scenarios. Depending on a variety of factors, one might

¹⁴ To offer a tentative definition, the *Most Important Case Hypothesis* states that the bulk of a law’s net societal impact may be determined not by its average suitability across cases but rather by how well it addresses a small number of extremely important cases that are affected by it. For clarity: When we refer to a ‘case’ here, we do not mean to imply that each scenario must necessarily involve a legal ruling (‘case’), but rather that it relates to a specific scenario where the legal architecture in place delivers particularly high-impact outcomes (whether or not that is in the form of specific legal cases).

¹⁵ Note that this is based on longtermist CBA, as opposed to the application of CBA based on traditional discount rates. At this stage of the analysis, we are concerned exclusively with what laws and policies would be net positive from a longtermist perspective, rather than with what methods one should use to maximize the strength of the argument in practice. Regarding the latter, [Thornley & Shulman \(2023\)](#) argue that one should use and appeal to traditional CBA.

¹⁶ Of course, many other factors would be at play. This is only one example.

¹⁷ For instance, it could offer information on whether lawmakers should include a of list existential risks examples to avoid later over-expansive interpretations, conceptual slippage, and/or muddling down of the concept, or whether they should emphasize the inclusion of low probability scenarios or a minimum harm threshold in order to anchor its proper scope of application.



favor surveying experts and/or ordinary people. In the context of laypeople, we call this approach “*ex ante* ordinary meaning analysis”¹⁸.

These are just some examples we aim to explore to evaluate different proposals. However, we currently expect to spend the majority of our research capacity on somewhat more traditional legal analysis of actions that LPP or other actors have already identified as promising.

As these updates show, revisiting the theory of change can be crucial, especially for early-stage organizations such as LPP that are still exploring where they can contribute the most. We look forward to working with our team and the associated communities to refine these priorities and identify the best ways to continue improving our theory of change and our related work.

Regardless of whether you have a legal background, we encourage you to reach out with questions you think LPP should prioritize. If you are not sure whether your suggestion falls under “legal research,” even if broadly interpreted, consider consulting this document, where we lay out different types of questions LPP can help with, including some that other legal organizations are unlikely to cover, using terms that are understandable for both lawyers and non-lawyers.

Field building

We would like to run at least one, but ideally two, field-building programs:

- **Legal Priorities Summer Institute**: An intensive, week-long program for younger participants (typically law students) to introduce them to the ideas of legal longtermism. Depending on funding, staff capacity, and venue availability, we could increase the number of participants from ~35 to ~100. Last year, the number of participants was limited by a combination of venue size and staff capacity despite having funding to invite more participants. Our current best guess is that we will run the event with ~35 participants again this year due to staff capacity and, if it is again successful and we increase staff capacity, substantially increase the number of participants in 2024.
- **Summer Research Fellowship**: A longer (8–10 weeks) research-focused program for 5–10 researchers already familiar with longtermism. Unlike previous years, fellows will primarily work on bespoke topics defined by LPP

¹⁸ In our recent paper ([Martínez & Winter, 2022, pg. 5](#)), we define “*ex ante* ordinary meaning analysis” as being “focused not only on how judges can and should interpret legal provisions once they are drafted, but also on how lawmakers can draft a legal provision using words that will best guide judges (and the public) into recovering their intended meaning and legislative aims.” The rationale behind this concept is that how a layperson might understand various terms is legally significant, as judges tend to interpret a term according to its ordinary meaning—that is, how a typical or reasonable person generally understands and uses a given word or concept.



teammates and partner organizations. Having fewer fellows will allow LPP staff to supervise and collaborate with them more closely.

Additionally, we are considering:

- Publishing a career guide for law students, with advice on finding impactful jobs from a legal longtermist perspective;
- Hosting a podcast showcasing success stories and careers of people with a positive impact in law and policy, highlighting what lessons one might draw from this; and
- Hosting smaller in-person events in Cambridge (US) in the second half of 2023 for those who engage closely with the ideas of legal longtermism.

Policy priorities for 2023

In the past year, we gained valuable insights from collaborations and discussions with actors working closely in national and international law and policy. For instance, we increasingly appreciate the impact of a diverse set of minor details (such as the idiosyncrasies or temperaments of key decision-makers) and developed our arguments and tone to effectively communicate without having a polarizing effect or poisoning the well.

We plan to continue having these discussions, although we believe we are better positioned to provide relevant actors within the policy space with decision-relevant research rather than setting out to engage with and convince policymakers ourselves. Given that our direct policy experience remains limited, we will continue to engage more directly in policy efforts only if the downside risks (e.g., polarization, information and attention hazards, lock-in of suboptimal policies) are low and potential benefits are high. This would require us to continue (1) coordinating with other relevant stakeholders to avoid the unilateralist curse and (2) collaborating exclusively with existing and ongoing policy efforts rather than launching our own policy initiatives.

Other priorities

- We are currently constrained by the size of our team. We would like to increase our operational capacity and hire more support staff for our outreach and field-building efforts. If resources allow, we would also like to grow our research team. In particular, we are considering hiring a (senior) researcher with expertise in **U.S. law focusing on risks from advanced AI** and/or a (senior) researcher with expertise in **legislative drafting**.
- We are generally interested in hiring more senior researchers but likely won't spend too much time actively pursuing this (i.e., no formal hiring rounds).



- We are discussing the possibility of offering **consultancy-type research services** (either by us directly or by commissioning projects to researchers in our network), mainly to organizations focusing on the risks of advanced AI systems.
- We will continue our fundraising efforts in early 2023 and aim to raise at least \$1.1m, per our reserves policy.
 - We will also continue to diversify our sources of funding. So far, we have relied primarily on the support of large grantmakers in the effective altruism community, some of which have scaled down their grantmaking recently. We would like to avoid relying on a handful of or even a single major donor. However, we also have to recognize that working on neglected topics makes this a difficult enterprise (almost) by definition.
- Since we are a fully remote team, we aim to organize 2–3 in-person team meetings per year, usually around conferences that multiple teammates attend anyway.
- We plan to improve our website to extend its functionalities and simplify its maintenance.

How to support us

If you are excited about our work, the best way to support us is by **donating to us** (via bank transfer, PayPal, or Venmo, as well as in crypto, stocks, and through DAFs). Feel free to contact us to learn more about how your donation could support our work under different growth scenarios. Donations are tax-deductible in the US, but you can also get in touch with us if you'd like to make a tax-deductible donation of over \$10k outside of the US. These are other ways to support us:

- Share this page with anyone interested in supporting us
- Subscribe to our newsletter and follow us on social media (Facebook, Twitter, LinkedIn)
- Attend and share our events (which we announce in our newsletter and on social media)
- Apply to work with us
- Join our Law & Longtermism Network on Slack
- Offer pro bono advice to us
- Send us any feedback on how we can improve



Legal Priorities
Project

www.legalpriorities.org