



RETHINK  
PRIORITIES

# Five Years of RP

Past, Present, and Future

# Table of Contents

Executive Summary..... 4

Background..... 8

Outcomes & Impacts ..... 13

Organizational Updates..... 4

Lessons Learned..... 14

Future Plans ..... 15

Funding Needs ..... 18



# Executive Summary

## Key Accomplishments (2018-2023)

In five years, RP has published over 125 pieces of research, completed more than another 100 research projects, provided different grantmakers with consultation, influenced tens of millions of dollars in funding, fiscally sponsored nine projects, and drove forward the promising field of invertebrate welfare.

Specific highlights of our work include:

- Collaborating with dozens of European Union (EU) animal advocacy organizations to work on setting medium-term policy strategies for farmed animal welfare.
- Providing expert consultation to the Chilean government as they considered a bill (which has advanced to the next legislative stage) to recognize animals as sentient.
- Contributing significantly to burgeoning fields, such as invertebrate welfare—including work related to shrimps and insects (see more [here](#))
- Completing the Moral Weight Project to try to help funders decide how to best allocate resources across species.
- Producing 23 reports commissioned by Open Philanthropy answering their questions about global health and development issues and interventions.
- Conducting over 200 tailored surveys and data analysis projects to help several organizations that build communities of people working on global priorities.
- Launching projects such as [Condor Camp](#) and fiscally sponsoring organizations like [Epoch](#) and [Apollo Research](#) via our Special Projects team, which provides operational support.
- Setting up an Artificial Intelligence (AI) Governance and Strategy team.
- Growing from a two-person operation in 2018 to a team that will soon include 75 RP employees, 30 contractors, and 25 staff of fiscally sponsored projects.



## Mistakes and Challenges

We believe that some of RP's **past projects** failed because we did not adequately consider the project's probability of success, its potential value, and the resources required. For example, our 2018 PriorityWiki project would have involved a large volunteer coordination effort we weren't well-placed to execute, and it is unclear how valuable it would have been even if successful. Our biggest early mistake was not building a plan for each project's path to influence and not putting enough resources into measuring our impact. We initially relied too much on producing research and hoping that it would be impactful just by existing.

Thus far, we have not publicly shared as much of our existing internal impact tracking as we had initially intended due to time constraints.

While we did some preparation prior to scaling in 2022, we would have liked to have established more robust project management systems beforehand. Our project timelines are not always predictable, as it is difficult to determine when you should stop researching due to diminishing returns. We think there are several cases where we spent too long working on a piece of research and would've had more impact by releasing the work earlier.

# Executive Summary (continued)

## Future Plans

### Animal Welfare

The team plans to continue developing decision-relevant research, including investigating policy approaches and the cost-effectiveness of current interventions to help farmed animals, and likely uncovering new interventions to help neglected animals (e.g. fishes and crustaceans) and approaches to advance wild animal welfare. Further funding would help the team in developing a live tracker of the cost-effectiveness of farmed animal interventions, producing data sets on farmed animals, developing a strategy for legal protection of aquatic animals in the EU,

### AI Governance and Strategy

The team aims to bridge technical and policy worlds as well as to serve as a talent pipeline and improve coordination among relevant actors. They will focus on four areas: compute governance, the intersection of China and AI governance, lab governance, and U.S. regulations. With funding, the team could hire additional staff to tackle questions about the impacts of and ideal changes to the 2022 semiconductor export controls, what plausible U.S. regulations on frontier AI systems would be beneficial and counterproductive, and what holes prominent governance proposals leave open and how those can be patched.

### Existential Security

The team will focus on helping launch entrepreneurial projects that aim to reduce existential risk, by identifying and developing the most promising project ideas, and then finding and supporting founders to make those projects happen. The team currently has significant unused management capacity with funding as the key bottleneck for them making new hires.

### Survey and Data Analysis

The team aims to conduct systematic research exploring critical strategic questions. One such area of focus centers on the impact of message framings of existential long-term risks among distinct groups. The team also aims to examine public opinion by assessing the levels of support or opposition for specific policies.

### Worldview Investigations

The team plans to make several contributions to the broader conversation about global priorities, including a cross-cause cost-effectiveness model, surveys of key stakeholders about the inputs to the model, and a series of reports on the cruxes.

## Changes

We now have better project management systems and spend much more time thinking through how to communicate our research and ensure each piece has a higher chance for impact.

We are focusing more on measuring our impact, and hired a Chief Strategy Analyst last year.

In addition to researching neglected areas (e.g. invertebrate welfare), we're now also working on some areas that are comparatively less neglected (e.g. global health and development). As our capacity has grown, we are now better positioned to investigate interventions that may have previously been overlooked for reasons that are unrelated to their cost-effectiveness.

We've established a Special Projects team to help support promising initiatives.



# Get Involved

**Funding needs:** While RP has worked hard to build a funding base, some of our plans will not occur without funding before early 2024 (see [funding gaps](#)). We invite you to donate [here](#).

**Ask us anything:** You can also engage with our work by participating in an “Ask Me Anything” with Co-CEO Peter Wildeford that will be **announced by July 19** in our newsletter and social media, by sharing this post with potential collaborators, or by contacting Director of Development [Janique Behman](mailto:Janique.Behman@rethinkpriorities.org) at [Janique.Behman@rethinkpriorities.org](mailto:Janique.Behman@rethinkpriorities.org).



# I. Background

## About Us

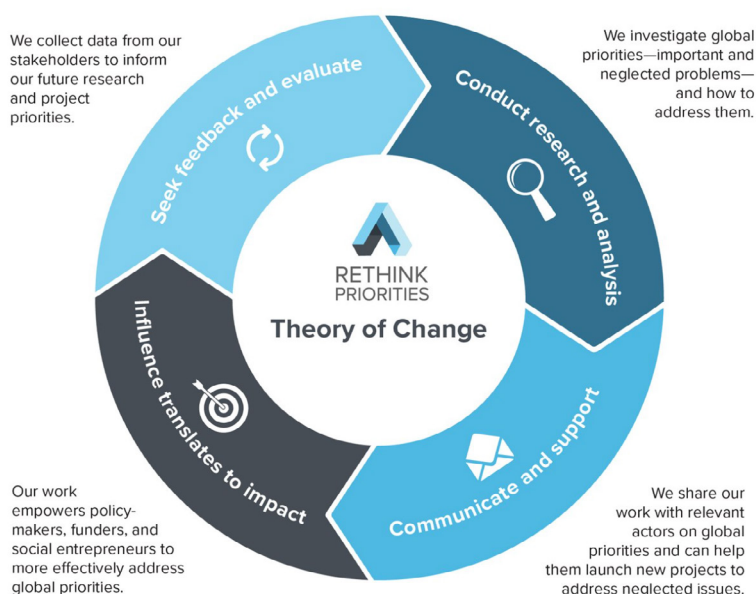
In 2018, [Marcus A. Davis](#) and [Peter Wildeford](#) co-founded Rethink Priorities with the aim of leveraging empirical research to tackle some of the world's most pressing challenges.

RP is now a research and implementation group that works with foundations and impact-focused non-profits to identify pressing opportunities to make the world better, figures out strategies for working on those problems, and does that work.

### RP focuses on important and neglected cause areas, we work on:

- Animal welfare, including wild and farmed animals
- Global health and development, including climate change
- AI governance and strategy
- Existential security, and other work to safeguard a flourishing long-term future
- Understanding and supporting the communities focused on the above issues

## Our Theory of Change



### RP functions as:

- A **consultancy** doing commissioned work in response to demands from organizations aligned with our mission.
- A **research institute** driven by research agendas we set according to our own priorities.
- A **think tank** aiming to inform public policy to improve the world.
- An **accelerator** for new priority projects addressing key problems.

## II. Outcomes & Impacts

Ultimately, RP aims to positively impact others now and into the long-term future. Below is a select list of outcomes from our research, communications, influence, and support efforts.

### Research

Since launching, RP has identified gaps in important, neglected, and tractable cause areas. Our staff have progressed in researching these gaps, drawing conclusions, and attempting to determine next steps. While much of this work is confidential, all of our publicly available research reports can be found in the [database](#) on our website. This report focuses on how we think some of our research is contributing to building newer fields and the impacts of our behind-the-scenes work. Since our Animal Welfare team was the first to scale, we currently have more examples of our work in this area.

### Invertebrate welfare

We've explored sentience, welfare concerns, and potential solutions in policy and industry collaboration for invertebrates, such as insects and shrimps. Key outcomes include:

- The Albert Schweitzer Foundation and ProVeg used our research in an EU consultation on authorizing insects to be fed to chickens and pigs.
- Eurogroup for Animals produced a position paper on the use of insects as food and feed that was informed by our non-public research on insect farming and welfare.
- We participated in requests for information processes with the U.S. Department of Agriculture (USDA) and the U.S. Agency for International Development (USAID) and in public consultations with governmental authorities in the U.K.
- Some industry professionals reached out following our presentation on the black soldier fly at the 2022 Insects as Food & Feed Conference. We also shared our findings at the 2022 Insect Welfare Symposium and in scientific journals.
- Peter Singer cited RP's research on insects and shrimps in his new book, [Animal Liberation Now](#). Singer has stated that, amongst other reasons, he updated his 1975 book in part because “scientists have gained more evidence that suggests the sentience of fish and some invertebrates.”
- We presented our shrimp research to animal advocates from across Latin America and have been invited to speak at universities in Chile, Spain, and the U.K. and to speak about this work on an 80,000 Hours podcast.





## AI Governance and Strategy

We have established a multidisciplinary team to tackle strategic questions around the governance of artificial intelligence. The staff work to reduce catastrophic risks related to AI by conducting research and strengthening the field of AI governance. They have addressed a number of questions, and we believe they are now in a position to engage policymakers in Washington, D.C. and elsewhere. Currently, their four main workstreams are:

- Compute governance
- China
- Lab governance, and
- U.S. regulations.

See more about the team's work and plans in their recent post [here](#).

## Moral Weight Project

Building on our research into invertebrate sentience, our [Moral Weight Project](#) improved understanding of both capacity for welfare and intensity of valenced experiences across species. The project also considered the moral implications of those possible differences. Our research laid the foundation for cross-species cost-effectiveness analyses that inform decisions regarding how many resources funders and organizations should tentatively allocate toward helping each of these

## Communications



## Policy and Advocacy

**Chile:** The Environmental Commission of the Chamber of Deputies invited our Director of Animal Welfare to provide expert consultation as they considered a proposed law to recognize animals as sentient individuals. The Commission decided to advance the bill to the next legislative stage, and cited RP in its report.

**EU:** Ahead of negotiations on farmed animal protection in 2023, our researchers met with dozens of European animal advocacy organizations to discuss policy strategies for the upcoming years. These collaborative conversations directly shaped how the organizations collectively allocated over **\$1M USD** in their budgets.

## Message Testing

### **Existential security:**

We tested different ways to help a U.K. policy think tank communicate effectively about the long-term future. Through our work, they learned that they could increase the likelihood of success of their policy proposals for existential risk reduction by incorporating language like “securing a future for our children’s children’s children” into their communications.

### **Artificial intelligence:**

We recently ran a nationally-representative survey of U.S. public opinions on AI that aimed to replicate and extend other recent polls. The [findings](#) suggest that people are cautious about AI and favor federal regulation. However, respondents perceive other risks, such as nuclear war, as more likely to cause human extinction.

**Animal welfare:** RP has also conducted polls and surveys for some of the largest farmed animal welfare organizations such as The Humane League, Mercy for Animals, and Animal Equality. By providing these advocates with up-to-date information, we aimed to help them develop strategic legislative and corporate policy proposals and messaging to increase their chances of success.



## Influence

Since its inception in 2021, our Global Health and Development team has completed 23 commissioned reports for five organizations. Open Philanthropy (OP) commissioned 78% of these reports. Most projects entailed looking into a specific topic or intervention like hypertension or improving weather forecasting. We believe that our research likely increased the scope of impactful programs considered in OP’s Global Health and Wellbeing portfolio. (Read about how our team works [here](#) and view the publicly available research [here](#).)

Our Animal Welfare team as well as our AI Governance and Strategy team have also helped to answer questions and influence strategy at OP and elsewhere.

Our [analyses of cost-effective cultured meat technologies](#) have informed funders and advocates in the alternative protein space.

We’ve collaborated closely with the [Wild Animal Initiative](#), as we’ve explored tractable interventions to support species such as fishes and rodents.

# Support

## Special Projects

In 2022, we launched the Special Projects program to provide promising initiatives with operational support to accelerate their development. Our support includes hiring, payroll/finance, event planning, communications, and additional operational tasks. Noteworthy examples of projects from our incubation program are:

- [Epoch](#) is a research initiative that forecasts the development of transformative AI. By informing stakeholders about the timing of advancements and potential avenues for influencing AI progress, Epoch aims to facilitate effective governance.
- [Condor Camp](#) is a training and mentorship program to help talented Brazilian university students develop skills to address pressing global challenges. For more about RP's broader existential security work, please click [here](#).

- [Existential Risk Alliance](#) is a summer fellowship program for aspiring existential risk researchers.
- [Apollo Research](#) is a research project aiming to detect and prevent deception in AI.

## Fellowship program

RP has hosted 23 aspiring researchers as visiting fellows, some of whom have gone on to join our team as staff or to work in high-impact areas elsewhere, including at Open Philanthropy, Founders Pledge, the Centre for Effective Altruism, 80,000 Hours, and in U.S. state-level elected government.



### III.

## Organizational Updates

Last year (2022), RP roughly tripled in size. Since 2018, we've grown from a start-up with two co-founders to an established organization with 130 people involved, including our employees, people working for fiscally sponsored projects, and contractors. To scale in a healthy way, we needed to make a variety of organizational changes, including:

- Developing HR and people and project management systems.
- Strengthening our Board of Directors by adding five new people with expertise in finance, operations, and the cause areas in which we work.
- Expanding our leadership team to include new strategy and director-level roles.
- Refining our internal processes for high-level decision-making.
- Establishing a Justice, Equity, Inclusion and Diversity (JEID) committee.
- Adding staff focused on areas like training and professional development.
- Commissioning an external review of our practices and procedures.
- Completing an internal risk assessment and mitigation plan.

\*While we are not aware of any major or legal risks, we are mindful of the need to (1) diversify our funding sources, (2) to communicate strategically as we reach out to new audiences, (3) to further professionalize our research standards, (4) to continue taking infohazards seriously, and (5) to work to retain key personnel.

# IV. Lessons Learned

## Mistakes

In a separate **post**, our co-founders Marcus and Peter reflected (amongst other topics) on the lessons they have learned since they laid out the **key strategic principles** for RP in 2018.

They believe that our biggest early mistake was not building clear plans for how each project would influence decision makers and not putting enough resources into measuring said influence. While we have done internal impact tracking, we have not shared it as publicly as we had planned simply due to time constraints. Moving forward, we plan to share more updates about our plans, strategies, and outcomes.

## Changes and Challenges

RP has also changed in ways that we think make sense given our growth and the evolving landscape in which we work. As mentioned above, we're no longer just doing research. We're also **incubating** other organizations and doing more **policy and advocacy**.

Even within the research realm, we realize that simply producing quality reports is not enough to have an impact. The way that we try to have an impact is through influencing and supporting the work that other organizations and people are doing. So, our relationships with our stakeholders is integral to our success. We are grateful for the relationships that we have developed with organizations like Open Philanthropy and GiveWell, who have entrusted us to explore questions that could shift millions of funding dollars.

Looking forward, we are cognizant of the need to diversify our base of funders and to initiate relationships with additional audiences. Our heavy reliance on a few donors for funding limits the types of projects we can take on. Furthermore, if the support for our research projects and interest in acting upon our advice from a specific major donor changes, this would severely limit the impact we can have.

We currently stress the importance of engaging various other philanthropists, foundations, and organizations, exploring new funding models, and diversifying our fundraising approaches. This could help us establish a sustainable funding pool, reduce dependence on a single funder, and expand our reach to a broader audience.

# V. Future Plans

We envision the future of RP to involve sustainable growth to increase our impact while maintaining our core principles. Working in alignment with our current mission, we aim to clarify and formalize our theories of change, refine our operating model, and actively scan for policy opportunities.

## Animal Welfare

### The team plans to continue developing decision-relevant research, including

- Investigating potentially promising policy approaches to help farmed animals and the possibilities to enhance substitution effects of alternative proteins
- Looking into the cost-effectiveness and other assessments of current interventions to help farmed animals
- Investigating and likely uncovering new tractable interventions to help numerous but neglected animals, including fishes and crustaceans used in food production
- Developing and assessing alternative approaches to advance wild animal welfare, with special attention to the tractability of this problem



### Further funding will help the Animal Welfare team in:

- Forecasting theories of change in animal advocacy, producing crucial strategic insights for people in the space
- Developing a live tracker of the cost-effectiveness of farmed animal interventions, starting with corporate outreach campaigns
- Producing country-level and historical data sets of the numbers of different animals slaughtered for food production and animals alive on farms at any time (including aquatic and invertebrate animals), which will help with:
- Improving resource allocation decisions by grantmakers, animal organizations, and evaluators
- Spotting trends of the different industries and guiding further research efforts to understand such trends
- Developing a strategic roadmap to increase the likelihood of getting the best possible legal protection for aquatic animals, including crustaceans, in the EU
- Producing a farmed insect welfare ask that could be promoted by others, seizing a historical opportunity window to help the farmed animals most commonly killed for food and feed production
- Identifying tractable methods of controlling populations of wild animals living in urban areas, as part of further efforts that may unlock the potential tractability of advancing wild animal welfare in the short term
- Similarly, identifying the quickest agricultural pesticides to control wild insect populations
- Incubating a new project to implement novel interventions to help highly neglected but numerous groups of animals
- Organizing regular meetings and other events to facilitate strategic decision making by leaders in the animal space and to improve the value and impact of research

## AI Governance and Strategy Team

The team is now focused on rapid-turnaround outputs informing near-term decisions by governments and labs in their four key focus areas: compute governance, the intersection of China and AI governance, lab governance, and U.S. regulations. The staff aims to bridge technical and policy worlds, and to sit at the intersection of strategic considerations and concrete decisions. They also aim to serve as a talent pipeline and improve coordination among relevant actors.



In each of the focus areas, there are currently significant windows of opportunity, with more demand from decision-makers for work by the team than can be fulfilled by the current staff alone. The team has sufficient management and operations capacity to hire additional staff with further funding. For example, the team could tackle additional crucial questions about the impacts of and ideal changes to the 2022 semiconductor export controls, what plausible U.S. regulations on frontier AI systems would be beneficial and counterproductive, and what holes prominent governance proposals leave open and how those can be patched.



## Existential Security

The Existential Security Team is focused on helping launch entrepreneurial projects that aim to reduce existential risk, by identifying and developing the most promising project ideas, and then finding and supporting founders to make those projects happen. We are currently focused on projects related to AI safety—examples in this area include info/cybersecurity at top AI labs, facilitating people's transition from AI capabilities research to AI safety research, and field building for AI policy.

Our team currently has significant unused management capacity—funding is the key bottleneck for us making new hires. Across our entire organization everything we do is running into funding concerns and the core reason we can't do more is lack of access to funding. New hires would enable us to more quickly investigate the most effective ways to help launch more projects, and ultimately get more projects off the ground sooner, as well as bring additional entrepreneurial expertise into the team.

## Worldview Investigations Team

The team plans to make several contributions to the broader conversation about global priorities.

Among the planned contributions are:

- A cross-cause cost-effectiveness model, which will allow users to compare interventions like corporate animal welfare campaigns with work on AI safety, the Against Malaria Foundation with attempts to reduce the risk of nuclear war, biosecurity projects with community building, and so on.
- Surveys of key stakeholders about the inputs to the model.
- A series of reports on the cruxes, focusing particularly on animals' moral weights, the expected value of the future, and the tension between making a difference and doing good in expectation.
- In addition, the team will continue to work with policymakers and advocacy organizations to implement moral weight considerations into their decision making, amplifying their impact by publishing their findings in venues that facilitate the growth and recognition of this emerging field.

## Survey and Data Analysis Team

The Survey and Data Analysis Team aims to conduct systematic research exploring critical strategic questions. One such area of focus centers on the impact of message framings of existential risks and longtermism among distinct groups.

The team also aims to examine public opinion by assessing the levels of support or opposition for specific policies advocated by authors and organizations in our network. Through comprehensive analyses, the team aims to determine the level of support for different policy proposals and identify any significant variations among demographics or interest groups. These efforts will provide valuable insights into attitudes and perspectives regarding existential risk, helping to develop strategies and recommendations based on evidence.

## Global Health and Development

With additional funding, the team would increase the impact of their current work, build their emerging research agenda, and engage mainstream funders and stakeholders.

- **Increase the impact of current work:** The majority of the team's time is spent on commissioned work for major high-impact grantmakers. In some cases, our work doubles the number of investigations of underexplored global issues considered in the grantmaker's pipeline per year. Additional funding would allow the team to prepare and make public over 30 existing research reports which they think could be additionally impactful with appropriate dissemination.

- **Build their own research agenda:** The team has a long list of research ideas that they currently do not have the capacity to explore alongside commissioned work. They want to build a line of research into under-researched cause areas in global health, development, and climate change and cross-cutting topics aimed at improving funders' approaches to effectively identify and evaluate cause areas and interventions.
- **Influence mainstream global health and development funders:** Additional funding could increase the team's capacity and enable them to develop new and existing relationships with funders and policymakers. Their aim would be to identify impactful and decision-relevant research and ultimately target the agendas of a wider set of organizations with the goal of influencing them toward more effective interventions and neglected cause areas.



## VI. Funding Needs

The past year has seen many changes in the communities in which we operate, leading to the loss of funding previously promised to RP and others and a decrease in available resources. While we have managed to sustain our current operations, realizing our plans for the following year hinges on securing additional funding before early 2024.

We have often made the greatest impact when we had the freedom to explore new avenues of research. However, this level of flexibility can only be achieved through unrestricted funding. Currently, just 25% of our revenue is unrestricted, limiting our capacity for exploration and innovation.

Today, we invite you to join us in overcoming this limitation and bringing our plans to fruition. **Your donations will accelerate change when it matters most and enable new research that responds swiftly to emerging challenges.**

Such donations represent true counterfactual support, as we do not anticipate securing the full amount of funding that we could productively spend from institutional funders alone. Additionally, given our size, institutional funders often expect us to secure a significant portion of our funding from individual donors.

*To achieve our impact goals and drive positive change, we have identified specific funding needs for each of our cause areas in 2024:*

Team	Amount (in USD)
Animal welfare	3.5M
AI governance and strategy	3.9M
Existential security	2.5M
Global health and development	2.6M
Surveys and data analysis	1.5M
Worldview investigations	1.1M
<b>Total (through EOY 2024)</b>	<b>15.1M</b>

---

## By donating to Rethink Priorities, you will:

- Put important research in front of influential decision-makers—be they policy makers or funders—during critical moments in the current conversations on AI governance, animal welfare policy, climate change, and other topics.
- Influence new grant-making organizations, and drive additional funding towards important causes.
- Lead to the creation of new, high-potential altruistic projects.
- Provide researchers with the freedom to focus on generating impactful insights, in the knowledge that RP's operational infrastructure and long-term sustainability are secure.

## Other ways you can engage with our work

**Share** this document with individuals and organizations who may be interested in collaborations or commissioned work.

**Contact** Director of Development [Janique Belman](mailto:Janique.Belman@rethinkpriorities.org) at [janique@rethinkpriorities.org](mailto:janique@rethinkpriorities.org) with your questions about engagement with or donations to RP.

**Participate** in our upcoming “Ask Me Anything” with Co-CEO Peter Wildeford that will be **announced by July 19** in our newsletter and social media.

# \$10M+

of funding influenced

# 125

pieces of research

# 9

fiscally sponsored projects

# 5

years of RP

***Will you join us and  
fund this future?***

[Click here to donate now.](#)

---



This report is a project of Rethink Priorities—a research and implementation group that works with foundations and impact-focused non-profits to identify pressing opportunities to make the world better, figures out strategies for working on those problems, and does that work. It was written by Rachel Norman and Sarina Wong with substantial input from Marcus A. Davis, Peter Wildeford, Kieran Greig, Abraham Rowe, Janique Behman, and Hannah Tookey. Thanks to other members of the Development team for their helpful feedback.