

# Director of Research Strategy

Location: Remote, with many calls in Pacific Time.

Preference for candidates based in the US who are eligible for a security clearance

Travel: likely ~10% travel

Commitment: Part-time (minimum 14 hours/week) to full-time

Salary range: \$140k–200k/yr for full time (prorated for part-time) + benefits (for full time)

## Key Summary

Atlas Computing is looking for a Director of Research Strategy to ensure we are maximizing our impact toward a better future in the face of rapid advancements in AI capabilities.

Atlas is a nonprofit focused on closing the gap between think tanks and incubators by systematically creating new organizations where none yet exist—organizations designed to secure society against the most severe risks posed by advanced AI. [Field Strategists](#) are at the center of that effort. The Director of Research Strategy ensures they're solving the right problems.

This is not a traditional research role. It is a high-context, judgment-heavy position that blends risk forecasting, strategic prioritization, and research management in close partnership with the CEO and Director of Field Operations. You will spend roughly two hours per week with each Field Strategist, challenging assumptions, stress-testing threat models, suggesting variations on potential interventions, and ensuring solutions target the risks that matter most.

## As Director of Research Strategy, success in the first 6 months looks like:

- You've helped Field Strategists focus on problems that leading forecasters of AI capabilities and funders agree are important (not just tractable or flashy). Each FS (and the CEO) can defend their problem selection to skeptical experts.
- The CEO spends less time on strategic prioritization, without losing confidence that the team's problem selection is world-class.
- You have built or strengthened relationships with 5–10 senior strategic thinkers in the AI safety and AI resilience ecosystems who serve as transitive sounding boards for Atlas's problem selection.
- At least one Field Strategist has materially changed course based on your strategic input.
- (stretch) You are maintaining a living strategic map of neglected risks, defense-dominant equilibria, and intervention opportunities, which you update as AI capabilities advance.
- (stretch) Your understanding of the risk landscape based on Atlas's progress, and you're communicating this to other expert forecasters.
- (stretch) You share ownership for Atlas's successes in fundraising, recruiting new field strategists, and/or growing a network of experts (primarily through your writing and community engagement).

Note that we expect this first cohort to focus primarily on starting new efforts in the United States that could scale globally, rather than creating efforts starting outside the US.

## What you'll do in this role...

- **Set Strategic Direction(s) for Atlas** - Work with the CEO to set top-level priorities for the Field Strategist team, grounded in analysis and/or expert opinion of expected impact
- **Coach and Challenge Field Strategists** - Spend ~2 hours per week per Field Strategist reviewing their work, questioning assumptions, suggesting variations, and ensuring proposed solutions address the most important dimensions of the risks they target.
- **Forecast Risks and Advise the CEO** - Develop forward-looking threat models that anticipate how AI capabilities will reshape the risk landscape over the next 2–5 years; ensure Atlas is positioned to address them before they become urgent.
- **Outreach** - Represent Atlas's strategic perspective to funders, advisors, and peer organizations.

## You may be a fit for this role if you have...

- A track record in strategic foresight, macrostrategy, or future risk forecasting—ideally at an organization respected by major x-risk funders (e.g., Forethought, GPI, FHI, Coefficient Giving, RAND, CSER, FLI).
- Deep familiarity with the landscape of catastrophic and existential risks from AI, including both well-studied threats (misalignment, bioweapons) and less-explored ones (AI-enabled coups, epistemic degradation, power concentration).
- Experience questioning people's assumptions and asking the right questions—a research-manager sensibility rather than a pure researcher or pure operator orientation.
- Demonstrated ability to reason about AI trajectory, safety, and defense-dominant technological equilibria.
- Strong networks in the x-risk, AI safety, or strategic foresight communities; ability to code-switch and speak as a peer to experts from very different ecosystems who hold very different views without cutting off future collaboration opportunities
- Comfort forming and revising hypotheses quickly, and strong calibration on confidence levels.

Meaningful experience in at least one of these areas:

- Macrostrategy, cause prioritization, or global priorities research.
- AI timelines modeling or quantitative forecasting.
- Existential risk analysis or scenario planning.
- AI governance or AI policy research.
- Technology roadmapping or multi-year strategic planning.
- Biosecurity or pandemic preparedness strategy.
- Intelligence community analysis or national security foresight.
- Another area you can convincingly argue is relatively similar to these.

## At Atlas, we also expect everyone to have...

- a practice of reviewing the state-of-the-art after having an idea, coupled with an urge to prototype, test, and iterate.
- a founder mindset, with a high degree of ownership over outcomes (not process) and the ability to personally make and externally drive progress with little oversight.
- experience coordinating and aligning stakeholders to work toward a common goal.
- strong analytical, problem-solving, and communication skills.
- comfort with uncertainty and interest in shaping our culture, as we're still a very small organization.

## Bonus points if you...

- strongly hold a conviction that AI is going to reshape the face of civilization fairly soon, and you want to join a hard-working team focused on making that go as well as possible.
- have published work on macrostrategy, cause prioritization, or AI governance that has influenced the field.
- have experience at organizations like Forethought, GPI, FHI, Coefficient Giving, CSER, or similar.
- have a strong track record of accurately forecasting risky AI capabilities
- thrive in a high-trust, low-structure environment.
- are deeply committed to improving the human condition with technology...
  - and understanding how different people might differently define “improving”, “human”, “condition”, and “technology” in that sentence.
- have work authorization to work in the US (and ideally also would be eligible to hold a US security clearance)

Here's a supplemental living document if you have more questions:

[More information for potential hires](#)

Our recent blog posts also contain helpful context:

<https://blog.atlascomputing.org/>

Instructions to apply can be found at [atlascomputing.org/jobs](https://atlascomputing.org/jobs)