Leadership

Philip Tetlock – President, Chief Scientist
Josh Rosenberg – CEO
Ezra Karger – Director of Research

Internal Operations

Astrid Sorflaten – Head of Operations
Christina Aguila – Operations Associate
Tim Francis – Operations Associate

Research

Bridget Williams – Senior Research Manager
Simas Kucinskas – Research Manager
Jessie Barker – Senior Researcher
Jordan Canedy – Senior Researcher
Zach Jacobs – Data Analyst
Connacher Murphy – Data Analyst
Otto Kuusela – Data Analyst
Rebecca de Castro – Data Analyst
Houtan Bastani – Data Analyst
Alexa Pan – Research Analyst
Sam Glover – Research Analyst

17
person
team

Aims

Primary aim: Advance the science of forecasting for the public good

Everything FRI does is anchored in improving how society anticipates the future, especially where the stakes are highest

Improve Decision Making

Improve decision-making on "high-stakes" issues

The institute's primary objective is to develop and rigorously test methods that help governments, non-profits and other key actors make better calls on matters of broad social importance

Produce high-quality forecasting questions

Crafting clear, resolvable questions is a core research line because good questions are the foundation of accurate predictions

Improve Profundity

Stress-test forecasting techniques across domains

By systematically checking robustness, the institute ensures its tools work not just in theory but in varied real-world settings

Invent methods for "unresolvable" or extremely uncertain questions

FRI pushes the methodological frontier so that even topics once thought impossible to forecast (e.g., existential risks) become tractable

Improve Applicability

Translate research into practical tools for policymakers and non-profits

Large-scale experiments and collaborations turn academic insight into decision aids that can be deployed where they matter most

First Gen → 2nd Gen

First Gen: First-generation forecasting research—spearheaded by FRI Chief Scientist Philip Tetlock and coauthors—focused on establishing a rigorous standard for prediction accuracy.

2nd Gen: The next generation of work aims to channel this approach into real-world relevance. We work with policymakers and nonprofit organizations to design practical forecasting tools, and test them in large experiments.

Application

FRI's strategic approach involves translating first-generation forecasting research, which established stringent standards for prediction accuracy, into practical, real-world relevance.

This is achieved through direct collaboration with policymakers and non-profit organizations to design actionable forecasting tools and to test these tools through large-scale experiments.

This ensures the applicability and effectiveness of their methodologies in real-world contexts.

Governments | **Non-profit organisations** | **Policymakers**

4 core research areas

Producing High-Quality Forecasting Questions

This involves the **meticulous development** of questions about complex, long-run topics that are amenable to rigorous forecasting.

This area is fundamental because the accuracy and utility of any forecast are highly dependent on the **clarity, specificity, and resolvability** of the initial question posed.

Novel Methods for Resolving Unresolvable Questions

The institute actively explores and develops new approaches to tackle forecasting challenges that are inherently difficult or where traditional methods have proven insufficient.

This demonstrates a commitment to pushing the methodological frontier. Ezra Karger, FRI's Director of Research, is specifically tasked with developing methods for forecasting "unresolvable questions", highlighting this as a core, specialized capability within the institute.

what
does this
look like?

Testing the Robustness of Forecasting Techniques

This area focuses on rigorously evaluating how well different forecasting techniques perform across a variety of domains and contexts.

This ensures the generalizability, reliability, and external validity of their methods, which is paramount for real-world application in diverse scenarios.

Using Forecasting Tools to Help Organizations Make Better Decisions

FRI is committed to translating its forecasting research into practical, actionable tools that can directly assist policymakers and non-profit organizations in their decision-making processes. This underscores their strong applied focus and commitment to achieving tangible impact.

Research projects & methodologies

FRI conducts its research through various projects, often involving large-scale experiments and expert elicitation surveys, reflecting its empirical and practical orientation Recent and ongoing projects include:

Existential Risk Persuasion Tournament (XPT)

From June to October 2022, a forecasting tournament brought together 169 participants—80 subject matter experts and 89 "superforecasters"—to assess the likelihood of various existential risks by 2030, 2050, and 2100

Participants were incentivized to discuss their reasoning, revise their forecasts, and collaborate in teams, creating a dynamic, interactive setting

The project exemplified a hybrid model that blended (1) structured expert elicitation with (2) the competitive dynamics of forecasting tournaments and (3) elements of deliberation

A central feature was the direct comparison between superforecasters and traditional domain experts—an application of Philip Tetlock's foundational work on predictive accuracy under uncertainty.

Initial results showed a stark difference in risk perception. Superforecasters estimated a median 9.05% chance of a global catastrophe (defined as 10% or more of humanity dying) by 2100, while domain experts gave a median estimate of 20%.

This divergence underscores how differing epistemological frameworks can lead to markedly different assessments of existential threats.

AI Risk Adversarial Collaboration

This project brought together participants with fundamentally differing views on the risks posed by Artificial Intelligence

The core task was to collaboratively identify "near-term cruxes"—forecasting questions that would resolve by 2030 and were expected to lead to the largest changes in beliefs about the risk of existential catastrophe by 2100

Participants were strategically divided into two main groups: an "AI skeptic" group, primarily composed of superforecasters known for their accuracy in short-run forecasts, and an "AI concerned" group, consisting entirely of domain experts.

They made conditional forecasts, predicting how their long-term AI risk beliefs would shift if a specific near-term event (the crux) occurred. Discussions were facilitated via online forums and moderated one-on-one video calls to explore reasoning and disagreements.

This project utilizes "adversarial collaboration," a sophisticated methodology designed not just to seek consensus, but to systematically uncover fundamental disagreements and identify the specific information that could bridge belief gaps.

It represents a novel method for dissecting "unresolvable questions" by examining underlying assumptions and conditional dependencies, which is crucial for complex, contested topics like AI risk

AI Conditional Trees

This project describes a methodology for generating highly informative forecasting questions, particularly concerning AI risk, through specialized interviews with both domain experts and highly skilled forecasters.

This process successfully generated 75 new AI forecasting questions with resolution dates spanning from 2030 to 2070.

This project directly addresses one of FRI's primary research areas: "producing high-quality forecasting questions."

The core of this methodology lies in testing how "informative" these generated questions are. Informative, in this context, is defined by how much knowing the answer to a short-term forecasting question (about the next few decades) would influence a participant's forecast of the risk of AI causing human extinction by 2100.

By focusing on the "Value of Information (VOI)," FRI aims to optimise the utility of forecasting questions for policy and planning, ensuring that the questions posed are indeed the most impactful for informing critical decisions. This moves beyond simply asking questions to strategically crafting them for maximum foresight utility.

ForecastBench

ForecastBench is a dynamic, continuously updated benchmark specifically designed to measure the accuracy of Machine Learning (ML) systems on a constantly evolving set of 1,000 forecasting questions about future events.

A key design principle is to avoid any possibility of data leakage by only including questions for which no known answer exists at the time of submission.

This project directly contributes to the research area of "testing the robustness of forecasting techniques," specifically for AI systems.

It signifies FRI's proactive engagement with the rapidly evolving field of artificial intelligence and its potential role in forecasting, aligning with the broader trend of developing hybrid human-machine forecasting capabilities.

Expert Forecasts of Nuclear Risk

This study systematically assesses expert beliefs about the probability of a nuclear catastrophe by 2045, coinciding with the centenary of the Hiroshima and Nagasaki bombings.

The research involved 110 domain experts and 41 expert forecasters (referred to as "superforecasters"), utilizing a combination of in-depth interviews and structured surveys.

NATO-Russia tensions were frequently identified as the most likely trigger for nuclear exchange, and crisis communication networks and failsafe reviews were the most popular policies for risk reduction.

Participants were asked to (1) predict the likelihood of nuclear conflict, (2) explain the underlying mechanisms driving their predictions, and (3) forecast the potential impact of specific tractable policies on reducing nuclear catastrophe risk.

A notable distinction was found in risk estimation: superforecasters estimated the overall nuclear risk at 1%, while nuclear experts estimated it at 5%.

Similar to the Existential Risk Persuasion Tournament, this project leverages comparative expert elicitation to gain nuanced insights into high-stakes, low-probability events.

Publications and Knowledge Dissemination

FRI and its affiliates are prolific in disseminating their research findings through a variety of channels, including academic papers, comprehensive reports, detailed white papers, insightful editorials, and informative blog posts

This multi-format approach ensures broad reach to both academic and policy audiences.

Academic Papers

"Belief updating in AI-risk debates: Exploring the limits of adversarial collaboration" (2025) directly stems from their AI Risk Adversarial Collaboration project.

"Subjective-probability forecasts of existential risk: Initial results from a hybrid persuasion-forecasting tournament" (2025) presents initial findings from the Existential Risk Persuasion Tournament.

"ForecastBench: A Dynamic Benchmark of AI Forecasting Capabilities" (2025) introduces the ForecastBench project and its implications for AI forecasting.

"Project Improbable" (2024) highlights research into a particularly challenging aspect of forecasting—improving judgment around low-probability events.

"Identifying Good Forecasters via Adaptive Cognitive Tests" (2024) contributes to the understanding of how to better select and train individuals for forecasting tasks.

"The Forecasting Proficiency Test" (2024) introduces a new standardized tool for assessing forecasting skill, linking to broader efforts to evaluate and credential forecasting ability.

"Long-Range Subjective-Probability Forecasts of Slow-Motion Variables in World Politics" (2023) builds on long-standing research into the limitations of expert judgment in making long-range predictions.

"Improving Judgments of Existential Risk" (2022) explores strategies for improving forecasts, generating better questions, crafting clearer explanations, and informing policy around existential threats.

"Reciprocal Scoring: A Method for Forecasting Unanswerable Questions" (2021) introduces a novel method for addressing forecasting problems where ground truth may be permanently unknowable—relevant to many existential risk contexts.

Key Reports & White Papers

These documents often provide more direct policy relevance and detailed project outcomes:

"Can Humanity Achieve a Century of Nuclear Peace?" (Oct 2024) directly stems from the Expert Forecasts of Nuclear Risk project and highlights key contributions to understanding nuclear risk.

"Conditional Trees" (Aug 2024) presents a method for generating informative questions about complex topics, directly linked to the AI Conditional Trees project.

"Roots of Disagreement on AI Risk" (Mar 2024) provides deeper insights from the AI Risk Adversarial Collaboration project, exploring the potential and pitfalls of adversarial collaboration.

"Forecasting Existential Risks" (Jul 2023) offers comprehensive findings from the Existential Risk Persuasion Tournament, focusing on long-run forecasting accuracy.

Editorials, Blog Posts & Other Writing

FRI also engages in broader public discourse through opinion pieces and analyses

"A Better Crystal Ball: The Right Way to Think About the Future" by Scoblic, P., and Tetlock, P. (Foreign Affairs, Nov/Dec 2020) showcases Philip Tetlock's influence on foreign policy and strategic thinking