EQUITABLE LABS



COMMITMENT

The Equitable Labs Commitment

What we believe about AI shapes what we build.

by Katie Grillaert, Paul Crafer, Ed Vocke and Lindsey Washburn

FIRST EDITION · © 2026 EQUITABLE LABS, PBC. ALL RIGHTS RESERVED.



01 · PROBLEM SHARED

THE BARGAIN

The benefit of AI does not arrive on its own.

Every generation is offered a technology that will fix something, and asked to give something up in return. AI arrives with a particular bargain: if you want the benefit, hand over the data that describes your people, your patients and your borrowers, and trust whoever holds it to use it well.

Most of the world is being asked to accept that bargain without seeing the terms.

We are not against the technology. We use it every day, and some of the most careful people we know are building it. What we object to is the assumption that benefit arrives on its own, and that the institutions serving most of the world's people will catch up later. They will not catch up later. Capability compounds in the places that already have it, and the distance shows up in decisions taken without the people those decisions land on.

This is a short document about three things: who is in the room when those decisions are made, what evidence they are entitled to, and what it would take for a regulator in Apia, a clinic in Kigali or a bank in Stockholm to ask the same questions of a system as a firm in San Francisco.

We are a small practice, and a commitment changes nothing by itself. What follows is what we hold, written down so that we can be held to it.

WHAT WE HOLD

01 Start with the person the system lands on

Before we ask what a system does, we ask who it does it to. A decision that lands on a person is a moment in somebody's week: a loan declined, a claim flagged, a patient triaged. Those moments are where risk becomes real, and they are the least represented in the rooms where systems are designed. An assessment that cannot name the person on the receiving end of the decision is not finished.

02 Evidence, not assertion

This field runs on confidence. We prefer proof. A governance program that cannot show what it measured, on what evidence, and what happened when the answer was unwelcome is a document rather than a control. So we build the artifact as well as the argument: the test, the record, the method somebody else can run without us in the room.

03 Intelligence can be shared without the data moving

The data that matters most is the data that cannot move: patient records, loan books, case files, the things an institution holds on behalf of people who never agreed to anything. Learning together has always seemed to require pooling, and pooling is exactly what the people the data describes cannot consent to. It does not have to be that way. We work on methods that let institutions improve a shared model while each keeps custody of what it holds, because that is the only version of collaboration everybody can say yes to.

WHAT WE HOLD

04 The capability gap is the equity gap

Adoption is not a matter of enthusiasm. It takes people, money and time that a small regulator, a municipal government, a cooperative or a rural clinic does not have. Equity in AI adoption is therefore a distribution of capacity: the ability to specify what you need, to test what you are sold, to run it afterwards, and to refuse it. We work on the capacity side, with the institutions that have the least of it and the most to gain.

05 Harm is local, and so is the evidence

Risk looks different in Apia and Kigali than it does in Brussels. A model that is safe against one benchmark can be useless in a language with few speakers, or dangerous where a family shares one phone. Global frameworks are necessary and not sufficient. The people who can tell you how a system fails are usually the ones living with the failure, and they are rarely asked. We treat their account as evidence, not as consultation.

06 Say what we are for, in public, and report on it

We are a Delaware public benefit corporation. That is a legal duty before it is a badge: our charter names the benefit we exist to produce, and our board has a duty to weigh it against return. We will report on what we did about it, including the parts that did not work, and we keep the promises in this document specific enough to be checked.

None of us set out to start *this* company.

Katie and Paul met on the same piece of work: drafting the certification schemes that providers and deployers of high-risk AI systems are measured against at ForHumanity. Writing an audit criterion is a specific discipline. You cannot write one without deciding what evidence would satisfy it, and that decision is where a regulation either becomes a control or stays a paragraph.

Ed joined when the work turned to the hardest question in the EU AI Act. With a consortium of thirteen authors we wrote our bid for the European Commission's tender on modeling and evaluating sociotechnical risk in general-purpose AI: how a model's behavior combines with the context it is placed in, and how you would assess that before the harm rather than after it.

Lindsey came to the work from the public sector, where she had seen how easily a strong policy can stall between adoption and implementation. Her focus was on what it takes to make governance usable: bringing the right people into the room, giving them clear responsibilities, and helping institutions act on what they learn.

Then the work moved. Through the UNDP's AI Trust and Safety Re-imagination Programme we all worked on AI safety with teams in countries where the assumptions behind European and American frameworks do not hold. In Samoa and in Rwanda, and in the other places the program took us, the questions were different and more concrete. Whose language does the model not speak? What does safety mean where a household shares one phone? What oversight looks like when the state is the buyer and the community is the subject. None of it could be answered from Brussels, and most of it had already been answered locally by people nobody had asked.

We came back with the conviction the practice is built on. The gap in AI is not a gap in ambition or in talent. It is a gap in capacity: the ability to specify a system, to test it, to refuse it, and to keep the evidence afterwards. The institutions with the least of that capacity serve the most people, and they are the ones being asked to accept terms they had no part in writing.

So we started Equitable Labs as a public benefit corporation, with a mandate that describes the benefit we exist to produce. Alongside it we started Goose FL, its technology practice, to build the privacy-preserving AI infrastructure that lets institutions learn from one another without moving their records.

Another advisory practice on its own would not change much. One that can write the standard, test against it, and *build* the thing it recommends might.



THE PEOPLE

Four of us, four different ways into the same problem



Katie Grillaert

AI RESEARCH · ALGORITHM ETHICS · AI ADOPTION · WORKSHOPS & LECTURES

Katie works where regulation becomes something an organization can actually do. Seven years in data science taught her how the systems are built; the years since have gone into making them accountable. She helped draft ForHumanity's EU AI Act certification schemes for providers and deployers, leads a regulatory sandbox on Article 9 (risk management) of the EU AI Act, and is a certified auditor of AI and algorithmic systems. She advises a project at UNRWA, the United Nations agency for Palestine refugees, on the ethics of general-purpose AI in humanitarian settings, and publishes on how risk-averse financial algorithms shut people out.



Paul Crafer

AI PRODUCT MANAGEMENT · AI GOVERNANCE · BUSINESS CASE ASSESSMENT · WORKSHOPS & LECTURES

Paul came to AI governance after twenty years leading technology teams, which is why he asks what a control costs to run rather than what it says. He drafts EU AI Act certification schemes with ForHumanity and keeps a working liaison with the European Data Protection Board, the Austrian and Dutch data protection authorities and Swedac, the Swedish accreditation body. He works on AI policy and governance projects with organizations such as GIZ, the German development agency, and ACET, the African Center for Economic Transformation. He runs the workshops where regulators, engineers and affected communities have to argue about the same evidence, and he writes about what governance literacy requires in practice. He works from Stockholm.

THE PEOPLE



Ed Vocke

AI SYSTEMS DESIGN & DEVELOPMENT | DATA SCIENCE | TECHNOLOGY RESEARCH | WORKSHOPS & LECTURES

Ed has spent a decade turning frontier AI research into systems that survive contact with a regulator. He co-authored 'Applying Privacy-Enhancing Technologies to LLMs in Critical Infrastructure Contexts' (IEEE CARS 2025). He builds audit infrastructure and secure deployments for institutions that cannot ship their data anywhere, and works on the measurement problem underneath: what a model's behavior means once it is in front of a loan officer, a clinician or a border official. He came to AI through early work on language models in 2015, and writes on the fairness of anti-money-laundering and counter-terrorist-financing models in fintech.



Lindsey Washburn

AI GOVERNANCE | PUBLIC BENEFIT | AI FOR PUBLIC & SOCIAL SECTOR | STAKEHOLDER ENGAGEMENT | WORKSHOPS & LECTURES

Lindsey builds the public-sector and multilateral side of the work. Before that she led enterprise AI governance engagements with senior government officials, designed a secure and responsible technology lab, and ran an enterprise-wide AI governance program for a US city government. Earlier she was chief operating officer of the International Council of Ophthalmology, where she grew the budget from two to eighteen million dollars and negotiated the collaborations behind a WHO global action plan.

THE PRACTICE

What we do about it

Strategic AI advisory

We work with regulators, banks, health systems and funds on the questions they cannot answer from inside: what a system is for, what it puts at risk, and what evidence would settle the matter. Retained, with scope and hours agreed per initiative.

Applied research and pilots

Small, finite pieces of work that produce something checkable: an assessment method, a benchmark, a pilot that either works or is written up as not working.

Lectures and workshops

Governance and AI literacy for the people who will have to run this afterwards: boards, regulators, clinicians, engineers. We teach the material we use.

Technology

Goose FL is the technology practice of Equitable Labs and builds privacy-preserving AI infrastructure, so institutions can learn together without pooling the records. We advise on it, and we say plainly that we have a stake in it.

OUR CHARTER, IN PART

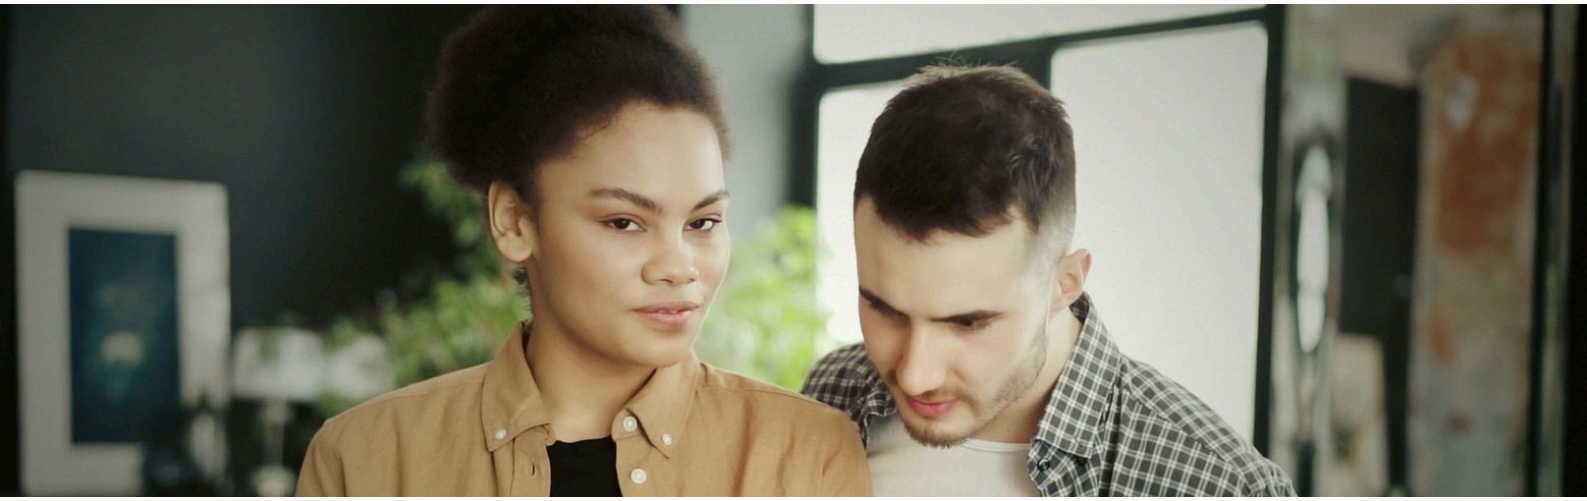
The Corporation exists to improve global equity in AI adoption and its benefits by designing, developing, and implementing technologies that prioritize digital rights and advance AI as a public interest technology. The specific public benefit to be promoted is the improvement of fairness and reliability in automated decision systems affecting underserved populations, including individuals and communities that face disproportionate barriers or risks when interacting with such systems.

This includes, in part, the development and deployment of privacy-preserving technologies that enable institutions to collaborate on risk detection and other public-interest use cases without compromising individual data or privacy.

CHARTER OF EQUITABLE LABS, PBC

Bring us the decision you cannot answer alone.

One paragraph is enough. Tell us what is not working and who it lands on. We answer within two working days, and we say plainly whether we are the right people for it.



hello@equitable-labs.com · equitable-labs.com

© 2026 EQUITABLE LABS, PBC. A DELAWARE PUBLIC BENEFIT CORPORATION.