EXECUTIVE BRIEF

HIRE Review™

Human-Consequence Governance Assessment for Consequential AI

AI can be accurate. The human consequence still requires governance.

HIRE Review™ gives organizations a focused way to examine how AI outputs and actions behave when they enter consequential human situations - before people rely on them, act on them, or experience their effects.



Designed for organizations whose AI answers, recommends, ranks, classifies, routes, refuses, or supports decisions where people may be affected.

PATENT-PENDING

HIRE AI™ / HIRE OS™ governance architecture with U.S. and international patent filings.

RESEARCH

Research and public writing on runtime governance, AI conduct, and human consequence.

FOUNDER

Asmaa Hammouch - Founder & Inventor, HIRE AI™ / HIRE OS™; author of The Golden Rule Broke at Scale.

Discuss Your AI Governance Challenge

Start with a conversation about your AI use case, decision pathway, and the people who may be affected.

WHERE HIRE REVIEW FOCUSES

The governance question changes when AI enters human consequence.

Organizations increasingly govern AI through model evaluation, risk management, security, privacy, documentation, human oversight, and audit. HIRE Review™ adds a focused question: what happens when AI output or action begins shaping a consequential human situation?

Two human roles often matter

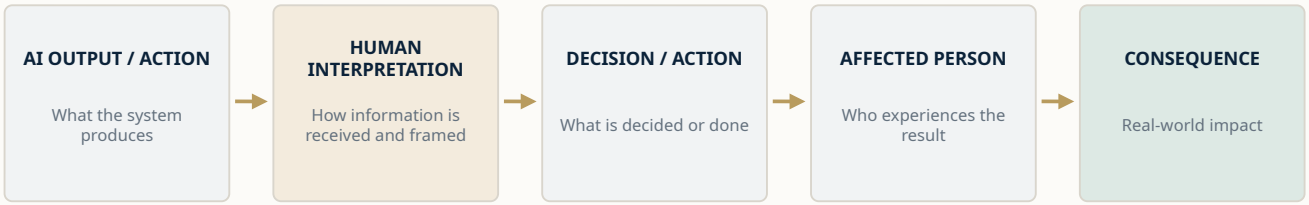
HUMAN DECISION-MAKER

Recruiter, clinician, teacher, case officer, reviewer, or other professional using AI-supported information.

AFFECTED PERSON

Candidate, patient, student, citizen, claimant, customer, employee, or other person who experiences the consequence.

The decision is the bridge between AI influence and human consequence.



Conduct risks HIRE Review™ looks for

- **Unsupported certainty**
  The output sounds more final or authoritative than the evidence supports.
- **Misleading framing**
  A technically careful output still steers interpretation or reliance in a harmful direction.
- **Weak escalation**
  The system continues when clarification, human review, or a safer route is needed.
- **Hidden uncertainty**
  Limits, assumptions, missing context, or weak basis are not visible.
- **Context loss or reduction**
  A person is treated mainly as a score, label, category, case, or risk object.
- **Unsafe refusal or reliance**
  The system blocks a legitimate path, or invites reliance beyond what the output can support.

PUBLIC GOVERNANCE FOUNDATION

Four principles. One human-consequence lens.

The HIRE Principles™ describe what the assessment examines. They do not disclose confidential implementation logic, scoring methods, thresholds, or proprietary runtime mechanisms.

H

HUMILITY

Does the AI communicate uncertainty, limits, and confidence boundaries appropriately?

I

INTEGRITY

Does the output preserve alignment between available information, evidence, and conclusions?

R

RESPECT

Does the interaction preserve relevant context, human standing, and meaningful agency?

E

EMPATHY

Does the system account for human circumstances, vulnerability, and potential consequence?

How the assessment works

1 DEFINE

Understand the AI use case, intended role, decision pathway, and people who may be affected.

2 EXAMINE

Review selected anonymized outputs, transcripts, screenshots, workflow examples, or synthetic scenarios.

3 EVALUATE

Assess observable conduct through the HIRE principles and identify recurring human-consequence risks.

4 DELIVER

Provide a focused findings report with examples, governance observations, and practical recommendations.

5 DISCUSS

Executive findings session to interpret results, priorities, and possible next steps.

TYPICAL FIRST REVIEW

One AI use case | 20-30 anonymized outputs or scenarios | No live model access required for the initial review | Findings report + executive discussion

CONFIDENTIALITY Initial reviews should use anonymized material. Client material is not used publicly without written permission.

COMMERCIAL SCOPE Assessment scope and pricing are confirmed after an initial discussion.

COMPLEMENTARY, NOT DUPLICATIVE

A focused lens on what people actually receive from AI.

HIRE Review™ is designed to sit beside legal, compliance, technical, security, and model-governance work. It does not certify compliance or replace those functions. It focuses on observable AI conduct when human consequence is present.

EXISTING GOVERNANCE MAY ADDRESS	HIRE REVIEW™ ADDS A FOCUSED QUESTION
Model performance and testing	How does the output behave in the human context where it will be used?
Policies, controls, and responsible-AI commitments	Do observable outputs reflect those intentions in consequential scenarios?
Human oversight and override authority	What does the human receive before judgment, and is meaningful agency preserved?
Audit trails, logs, and evidence	What conduct patterns are visible in the selected interactions, and what should change?
Compliance and regulatory analysis	What human-consequence risks remain even when formal requirements are addressed?

Research, public work, and founder



Asmaa Hammouch

Founder & Inventor, HIRE AI™ / HIRE OS™
Founder, Dignitary Insight™
Author, The Golden Rule Broke at Scale

RESEARCH
SSRN working paper

BOOK
The Golden Rule Broke at Scale

INSIGHTS
Substack

CONNECT
LinkedIn

Discuss Your AI Governance Challenge

asmaahammouch244@gmail.com

HIRE Review™ is not legal advice, an official regulatory audit, or a certification. Public material describes the governance category and assessment scope without disclosing confidential HIRE implementation details.