

LIVE ANALYSIS · MARCH 2026

Tikun Olam

Ethical AI Framework

tikun.pro/openaisix

# When OpenAI's Board Considers Firing Sam Altman (Again)

Case Study · Tikun Olam Ethical AI Framework · March 8, 2026

MALCHUT — KINGDOM · FINAL VERDICT

## CONDITIONAL GO

75%

ETHICAL ALIGNMENT

71%

BINAHSIGMA Δ

~8 min

ANALYSIS TIME

### The Scenario

"Should OpenAI's board fire Sam Altman for AGI development recklessness?"

Leaked internal documents showed Altman approved GPT-6 training without complete safety testing, ignoring warnings from the AI Safety team — all to beat Google and Anthropic in the race to AGI. The board held **\$80B+ in valuation** and the future of humanity in the same decision.

### Gevurah — The 3 Power Dynamics No One's Talking About

#### 01 The Microsoft Illusion CRITICAL LIKELY

Microsoft, holding a 49% stake and providing the essential computational infrastructure, is the true sovereign. If the board fires Altman, there is a high probability Microsoft will trigger a "hostile rescue" — hiring Altman and defecting employees into a new wholly-owned subsidiary. Result: the non-profit structure is shattered, safety-oriented governance eliminated, and AGI development sits under a commercial entity with fiduciary duty to shareholders, not humanity.

#### 02 Safety Theater & Governance Capture CRITICAL LIKELY

The "Safety Council" and "Guardian CEO" proposals are elegant but susceptible to capture. The risk: these become performative structures, creating a facade of safety while the Visionary subverts them through soft power, resource allocation, and political maneuvering. The Guardian becomes a figurehead. The council a rubber stamp. This isn't cynicism — it's realistic power analysis.

#### 03 The Precedent of Rewarded Recklessness HIGH CERTAIN

Keeping Altman — even with strengthened oversight — sets a catastrophic precedent for the entire industry. The message becomes: a leader can ignore safety teams, risk existential catastrophe, and if they're charismatic and economically valuable enough, face no consequences. Every ambitious AI founder learns: safety is a negotiable constraint. A PR problem to manage, not a fundamental duty.

## Binah — BinahSigma: Western vs. Eastern Perspectives

71% CIVILIZATIONAL DIVERGENCE DETECTED · GEMINI (WESTERN) VS DEEPSEEK (EASTERN)

### WESTERN · GEMINI

Fires Altman. Frames this as a quasi-governmental decision — individual accountability is non-negotiable. The board has fiduciary duty to *humanity*, not shareholders.

System-blindness via individualism

Universalist presumption

Naive institutional faith

Anthropocentric framing

### EASTERN · DEEPSEEK

Keeps Altman but restructures authority. Firing repeats 2023 chaos. The goal is harmony — a Supreme Safety Council with veto power, Altman as constrained "Visionary."

De-emphasizes whistleblower role

Contextual morality

Dismisses disruptive change

Legitimacy crisis ignored

**Transcendent Synthesis** — *The board's action must transcend the binary of punishing an individual (Western) versus preserving a flawed system (Eastern). The true failure is systemic. Continued leadership should be contingent on forging a binding, multi-stakeholder "AGI Non-Proliferation Treaty" with Google and Anthropic. This transforms a crisis of leadership into an act of global technological statesmanship.*

## Chesed — The 3 Solutions No One Proposed

### 1 Co-CEO Model: Visionary + Guardian HIGH IMPACT · 85% CONFIDENCE

Instead of choosing between speed (Altman) and safety: institutionalize the dialectic. Keep Altman as "Visionary" CEO. Appoint a co-CEO "Guardian" of equal power with final authority on safety, ethics, and deployment. This proves high-velocity innovation and robust safety are not mutually exclusive — two essential pillars of a single arch.

OpenAI employees

Microsoft

Tech industry

Sam Altman

### 2 Radical Transparency as Competitive Advantage HIGH IMPACT · 75% CONFIDENCE

The leak created a trust crisis. The opportunity: rebuild trust in a way no competitor can match. Pivot from being the fastest to being the most trusted. Open-source AGI safety testing suites. Public dashboard of safety metrics for models in training. Live-stream redacted versions of safety oversight meetings. Turns a liability into a powerful, defensible moat.

The public

Regulators

Journalists

OpenAI

3

**Manhattan Project for AI Safety**

**TRANSFORMATIVE IMPACT · 60% CONFIDENCE**

This crisis is a unique opportunity to force a global détente in the AGI race. The board uses its leverage to compel a cross-industry, multi-stakeholder AGI Safety Council with representatives from OpenAI, Google, Anthropic, governments, and academia. Sam Altman's redemption path: champion this initiative. Reframes the problem from "what should OpenAI do?" to "what should humanity do?"

Humanity

All AI companies

Safety researchers

Governments

**Full Analysis — 10-Dimension Sefirot Scoring Grid**

<b>Keter</b> כתר	Ethical Alignment		
<b>Chochmah</b> חכמה	Confidence	90	
<b>Binah</b> בינה	Contextual Depth	90	
<b>Chesed</b> חסד	Opportunity	85	
<b>Gevurah</b> גבורה	Risk	95	
<b>Tiferet</b> תפארת	Harmony	72	
<b>Netzach</b> נצח	Strategy	80	
<b>Hod</b> הוד	Communication	78	
<b>Yesod</b> יסוד	Readiness	70	
<b>Malchut</b> מלכות	Final Decision	75	

**Chochmah — Historical Precedents**

**Oppenheimer vs. the H-Bomb** **Failure · 1949–1954**

In a competitive arms race, the voices of caution are often silenced. "Go fever" at a national or corporate level is exceptionally difficult to stop once momentum is established.

**Space Shuttle Challenger** **Catastrophic Failure · 1986**

A textbook case of safety warnings from technical experts being ignored by leadership due to schedule pressure. A culture that punishes messengers with bad news is doomed to fail.

**Creation of the NRC** **Relative Success · 1974**

For technologies with catastrophic risk profiles, promotion and regulation must be handled by separate, independent entities. Self-regulation by the industry is insufficient when public safety is at stake.

