

U.S. DISTRICT COURT NORTHERN DISTRICT OF CALIFORNIA

In the Matter of:
RECONSTRUCTION OF THE EMERGENCE OF
OPENAI
(SEC Whistleblower No. 17684-273-411-436)

CASE NO. 5:26-cv-07078-EJD

**SUPPLEMENTAL PRE-TRIAL TRIAL EXHIBIT:
CROSS-EXAMINATION MATRIX FOR HOLDEN
KARNOFSKY**

MEMORANDUM FOR TRIAL TEAM: This supplemental trial outline is designed specifically to guide the cross-examination of **Holden Karnofsky**, CEO of Open Philanthropy, who served as a non-profit board member of OpenAI, Inc. from **March 2017 to 2021**. This document coordinates with our master trial platform, targeting the core tensions of the non-profit's early financial survival, compute-scaling math, and the eventual capped-profit transition. Use these lines of questioning to secure admissions that the early commercial pivot was dictated not by safe deployment constraints, but by a systemic R&D; monetization loop.

EXHIBIT 1: CHRONOLOGICAL BOARD CONTEXT & DONATION FLOWS

Date / Event	Transaction / Action	Strategic Board Context
March 2017	\$30,000,000 Capital Grant from Open Philanthropy	Karnofsky joins the OpenAI non-profit Board of Directors as a condition of the grant, representing the 'Effective Altruism' safety faction.
December 2017	\$16,028,500 Conflicted YC Org Donation	Executed during Sam Altman's dual-presidency of both YC and OpenAI. Siphoned accelerator capital to solve OpenAI's computing and payroll shortfalls.
February 2018	Elon Musk Board Resignation	Musk pulls early funding commitments. Board begins active, non-public planning to exit the 501(c)(3) public charity framework.
March 2019	Capped-Profit LLC Transition (OpenAI LP)	Non-profit board relinquishes direct operational control to a hybrid subsidiary to allow up to 100x return multiples for venture capital.
June 2020	Kaplan et al. Scaling Law Publication	OpenAI empirically proves cross-entropy loss follows strict power-law relationships with compute budget, signaling massive infrastructure needs.
May 2021	Karnofsky Resigns from OpenAI Board	Resigns due to a conflict of interest: his wife (Daniela Amodei) co-founded Anthropic alongside her brother Dario Amodei after a split on safety.

I. LINE OF QUESTIONING: THE \$30M ENTRANCE AND BOARD SEAT ACQUISITION

Q: Mr. Karnofsky, in March 2017, Open Philanthropy executed its largest grant to date: a \$30 million general support injection into OpenAI, correct?

Trial Tip: Establish that the \$30M grant was explicitly tied to a 36-month timeline (\$10M/year) to buy operational runway and direct board influence.

Q: As a non-negotiable condition of this \$30 million capital injection, OpenAI was contractually required to appoint you directly to its Board of Directors, correct?

Trial Tip: Establish that a board seat at a 501(c)(3) public charity was, in effect, purchased through structured capital distribution to ensure oversight of early safety.

Q: According to Open Philanthropy's published grant criteria, you partnered with OpenAI because of its 'shared values' and 'commitment to unconstrained research'—not because you believed a commercial pivot was inevitable, correct?

Trial Tip: Contrast early charitable representations of 'free of profit motives' with the board's concurrent backchannel planning to exit the tax-exempt structure.

II. LINE OF QUESTIONING: COMPUTE-SCALING MATH & THE \$10B CRISIS

Q: Mr. Karnofsky, while sitting on the board, you reviewed OpenAI's internal compute projections, which calculated that building and training frontier models would require compute funding 'on the order of \$10 billion', correct?

Trial Tip: Admit that the board knew that a pure donation-funded 501(c)(3) model was mathematically incapable of raising \$10 billion in public charity cash.

Q: You were on the board when Kaplan et al. finalized their 2020 paper demonstrating that language model performance follows strict power-law relationships with model parameters, dataset size, and compute budget, correct?

Trial Tip: Establish the scientific justification for scaling: that the board actively embraced scaling laws that required endless multi-billion dollar capital rounds from corporate hyperscalers.

Q: In fact, on the Dwarkesh podcast, you acknowledged that OpenAI's efforts in accelerating AI advancement and 'hyping' the technology has net-negatively compressed the time humanity has to prepare for catastrophic risks, correct?

Trial Tip: Secure a critical, on-the-record admission of the governance paradox: that the non-profit board presided over a capability acceleration that directly conflicted with its core safety mission.

III. LINE OF QUESTIONING: THE MARCH 2019 'CAPPED-PROFIT' RESTRUCTURING

Q: In March 2019, you voted to approve the transition of OpenAI's core R&D operations into a capped-profit limited partnership, OpenAI LP, under non-profit control, correct?

Trial Tip: Establish that the board formally ratified the hollowing out of the pure public charity, transferring valuable intellectual property and research talent to a commercial vehicle.

Q: Under this hybrid LP agreement, early investors were promised up to a 100-times return on their capital—a multiple that completely dwarfs standard low-risk venture expectations, correct?

Trial Tip: Highlight the extreme wealth-generation incentives introduced into an entity that was tax-exempt and publicly branded as 'freely owned by the world.'

Q: As a director, you were required to recuse yourself from voting on any financial matters where you or associated capital had direct conflicts of interest, correct?

Trial Tip: Prepare for the transition to the Anthropic split: establish that the board's internal recusal rules were under extreme strain due to the inter-marriage of board incentives and competing research labs.

IV. LINE OF QUESTIONING: THE CONFLICTION OF BOARD SEATS & 2021 DEPARTURE

Q: Mr. Karnofsky, in May 2021, you abruptly resigned from your seat on the OpenAI Board of Directors, correct?

Trial Tip: Admit that his departure was forced due to the ultimate conflict of interest arising from the hollowing out of OpenAI.

Q: Your resignation was prompted by the fact that your wife, Daniela Amodei, and her brother Dario Amodei, left OpenAI to found Anthropic—a direct commercial competitor also raising billions of dollars, correct?

Trial Tip: Establish that the 'Effective Altruism' board safety faction completely fractured, splitting into competing corporate entities (OpenAI and Anthropic) backed by rival VC groups.

Q: When your seat was vacated, the independent oversight you promised in exchange for the original \$30 million grant was dismantled, leaving the board heavily exposed to Altman's commercial control, correct?

Trial Tip: Synthesize the overall argument: the \$30 million public charity grant acted as a structural stepping stone that built OpenAI's early capability. Once that capability was monetizable, the safety-oriented board seat was discarded, completing the asset-conversion play.

SUMMARY DIRECTION FOR LEAD COUNSEL: Holden Karnofsky represents the bridge between OpenAI's early 'altruistic safety' branding and its hyper-capitalized commercial scaling. Use his testimony to establish the absolute materiality of the non-profit's asset-shifting plays. Establish that the 'capped-profit' LP and the subsequent PBC restructurings were not safety-driven evolutions, but highly calculated financial restructurings engineered to secure billions in venture capital while stripping the public of its promised oversight and equitable share of the light cone of value.