Harrison C. Margolin

An Analytical Comparison of the Ratio Method and Amodei's *Policy on the AI Exponential*

A comparison of Dario Amodei's *Policy on the AI Exponential* (June 2026) with *The Ratio Method* (52 St. Mary's L.J. 679, 2021).

June 2026 · Download PDF · GitHub

In June 2026, Anthropic CEO Dario Amodei published *Policy on the AI Exponential*, an essay calling for binding, enforceable regulation of frontier AI. Amodei's essay proposes regulating advanced models "like cars, airplanes, or drugs," modeled on agencies such as the Federal Aviation Administration. It argues that the industry should bear the harms it makes possible; and it reserves a more aggressive regulatory tier for future systems that resemble "weaponizable nuclear materials."

Several of these positions were shared, in method or reference, by *The Ratio Method*, published in the St. Mary's Law Journal in 2021. The Ratio Method is a procedure for selecting regulatory frameworks for emerging technologies by analogy to established industries.

The purpose of this comparison is to identify what their convergence may mean about the structure of the AI-governance problem, and how each overlap can indicate credibility or point to where more work is needed. The value of the comparison stems from the articles' independence.

Table of Comparison

Ex ante *is a Latin phrase meaning "before the event", commonly used in economics, finance, and law to describe analyses made in advance.*

Topic, concept, matter	**Policy on the AI Exponential** *(Amodei, 2026)*	**The Ratio Method** *(Margolin & Frazier, 2021)*	**Key Overlap / Analysis**
Technology Outpaces Policy	AI may progress from a "toy" to a "country of geniuses" within a single legislative cycle. Policymakers must act now while retaining capacity to escalate oversight.	Innovation outpaces regulation, leaving gaps before harms materialize. Match emerging industries to established liability regimes using comparative benefit-cost profiles.	Both address regulatory lag: Amodei through adaptive oversight, RM through existing liability doctrine.
Aviation & FAA Model	Like aircraft, cars, and drugs, frontier models provide substantial benefits but pose safety risks. Require pre-deployment testing and government authority to block or reverse unsafe deployments.	Apply FAA-style oversight to AI-enabled infrastructure: heightened duties of care, code review, operator vetting, safety standards, and emergency overrides.	Both invoke aviation safety, but Amodei targets frontier models; RM targets connected infrastructure.
AI & Cybersecurity	Mythos Preview demonstrates strategic cyber risks. Proposes mandatory third-party testing, model-weight security, red-teaming, penetration testing, and incident reporting.	Adaptive AI malware creates cascading risks analogous to oil spills. Apply Oil Pollution Act-style financial assurance, network-scaled liability, strict joint-and-several cleanup liability, and insurance.	Both prioritize cyber risks: Amodei emphasizes deployment safeguards; RM emphasizes liability and remediation.
Biological Weapons & Biosecurity	Bioweapon development is a mandatory evaluation area. Prioritizes testing biological uplift and strengthening biodefense as offensive capabilities improve.	CRISPR-engineered pathogens expose gaps in biosafety oversight, where lab standards are only advisory. Proposes internationally coordinated, capped strict liability with preparedness-based comparative negligence.	Both recognize exceptional biological risks; Amodei emphasizes prevention, RM liability and preparedness.
Labor Displacement	AI may cause persistent displacement, requiring wage insurance, retention credits, retraining, and potentially UBI funded by taxes on AI firms. Financial support alone cannot replace meaningful work.	Analogizes displacement to workers' compensation: verifiable harm, no-fault benefits, and profit-funded retraining. Firms retain productivity gains while protecting consumer demand.	Both favor funded worker support; Amodei proposes fiscal measures, RM producer-funded compensation.
Nuclear / Extreme-risk Tier	The most powerful AI systems may eventually warrant controls comparable to weapons-grade nuclear materials, though such measures should not be imposed prematurely.	Analogizes catastrophic risks to nuclear liability: absolute liability, mandatory coverage, damages caps, and shared responsibility where losses exceed insurable capacity.	Both invoke nuclear regulation for extreme risks; Amodei contemplates future controls, RM outlines liability mechanisms.
Autonomous / Self-evolving AI	Automated R&D and recursive self-improvement could accelerate other AI risks, warranting its own evaluation area.	Observes that adaptive AI malware may grow more capable through accumulated data, without addressing recursive model self-improvement.	Limited overlap: adaptive learning versus recursive frontier-model improvement.

Topic, concept, matter	**Policy on the AI Exponential** *(Amodei, 2026)*	**The Ratio Method** *(Margolin & Frazier, 2021)*	**Key Overlap / Analysis**
Innovation vs. Safety	Tailor regulation to specific risks, imposing safeguards on frontier models while accelerating beneficial downstream applications.	Society gives lighter liability to industries it values. Quantify that social value in advance, instead of leaving it to politics.	Both favor risk-sensitive regulation; RM formalizes the comparison quantitatively.
Ex Ante Review	Require pre-deployment testing above compute thresholds, backed by government authority to block or reverse deployment.	Draws on FAA airworthiness reviews, OPA financial-responsibility certificates, and biosafety oversight as preventive safeguards.	Both recognize preventive review; Amodei emphasizes deployment approval, RM integrates safeguards with liability.
Producer Cost Internalization	AI producers should bear disruption costs, including energy expenses and worker support funded through taxes on AI companies or capital gains.	Internalizes harms through cleanup liability, profit-funded retraining, and liability caps scaled to firm or network size.	Shared producer-pays principle; Amodei favors fiscal measures, RM tort-based incentives.
Quantitative Methods	Warns that rigid criteria may waste compliance resources or miss emerging risks; favors adaptable standards alongside compute thresholds.	Selects a liability regime algorithmically: compute Frequency and Average Damages, add Social Value and Litigation Costs, then match to a legacy industry.	RM offers quantitative screening; Amodei favors flexible standards. The approaches need not conflict.
Third-party Auditors	Proposes government evaluations or authorized private auditors subject to regulatory oversight ("regulatory markets").	Identifies analogous oversight mechanisms through FAA certifiers, biosafety registries, and financial guarantors.	Related oversight mechanisms, but only Amodei explicitly proposes frontier-model auditors.
Liability After Government Review	Focuses on certification and incident reporting without specifying liability for harms caused by certified models.	Evaluates negligence versus strict liability, claim aggregation, and appropriate compensation mechanisms by technology.	RM addresses the liability question Amodei leaves open, and could sit downstream of his deployment gate.
Civil Liberties / State Power	Warns against AI-enabled authoritarianism and surveillance; proposes a domestic ban on autonomous weapons, data-broker reform, a public right to AI counsel against the government, and checks on concentrated power.	Flags IoT-enabled mass surveillance and privacy concerns without developing broader institutional safeguards.	Primarily an Amodei theme; RM identifies surveillance risks without proposing comprehensive remedies.
Geopolitical Coalition	Proposes a democratic AI coalition coordinating compute infrastructure, manufacturing, risk standards, and defense.	Addresses international cooperation through biological, nuclear, and space-liability treaties.	Different objectives: geopolitical coordination versus cross-border liability.

Topic, concept, matter	Policy on the AI Exponential *(Amodei, 2026)*	The Ratio Method *(Margolin & Frazier, 2021)*	Key Overlap / Analysis
FDA / Biomedical Acceleration	Advocates validated AI-assisted research methods and more flexible regulatory pathways to prevent biomedical approval bottlenecks.	Analogizes CRISPR liability to drug/device regulation and vaccine-injury compensation programs.	Both reference biomedical regulation; Amodei addresses approval efficiency, RM liability for harm.

RM = The Ratio Method. Scroll sideways on narrow screens.

The Ratio Method in brief

The Ratio Method proposes that regulation can be solved by analogy. It characterizes a technology's harm profile using two primary ratios, **Frequency** (the share of the population affected by a harmful event) and **Average Damages** (the per-person cost of that harm relative to GDP per capita), supplemented by measures of the technology's **Social Value**, the **Type of Harm**, and the **Litigation Costs** of proving liability.

It then locates an existing, well-regulated **legacy industry** whose **benefit-cost** profile is comparable, and argues that the **standards of care**, **causation**, and **liability** governing that **legacy industry** are well suited to the emerging one.

The method is calibrated against historical outcomes in complex litigation, such as the Anthem and Target data-breach settlements, the Deepwater Horizon oil spill, the tobacco Master Settlement, and the workers' compensation system, among others, and is applied in the article to Artificial Intelligence, the Internet of Things (IoT), and CRISPR-Cas9 gene editing. The objective is stated in forward-looking terms: to help regulators identify, in advance, the liability regime whose incentives will produce "safe and steady industrial growth."

Amodei's Controlled Review vs. Ratio Method's Law-&-Economics

Both documents begin from the same premise: technology moves faster than policy can.

Amodei argues policy and legislation move very slowly while AI advances at a lightning pace. Margolin and Frazier frame it as the "pacing problem" and build their entire method

around situations "in which technological change outpaces legislative capacity." Both grapple explicitly with the dilemma of regulating under uncertainty.

Amodei wants ex ante regulatory gatekeeping: frontier models above a compute threshold get mandatory third-party testing, and their release should be blocked or reversed if they don't meet safety standards, administered by something modeled on the FAA. The control happens before deployment, through a government agency with the power to say no.

The Ratio Method operates from a lens of law-and-economics/tort. Their question is "how do we shape the landscape for technology to evolve, such that it grows safely." Its answer is to measure an emerging technology's harm profile (Frequency × Average Damages), match it to a legacy industry with a comparable cost-benefit profile, and borrow that industry's liability framework.

From this shared diagnosis, the two works select different mechanisms: Amodei proposes a pre-deployment gate: frontier models above a compute threshold would undergo mandatory third-party testing in four risk areas, cybersecurity, biological weapons, loss of control, and automated R&D, and the government would have the power to block a release that fails to meet safety standards, with evaluation performed by a government agency or authorized private auditors.

- **Amodei:** An administrative gatekeeper licenses or blocks the specific artifact (the model) before it reaches market. Direct permission-to-deploy.
- **Ratio Method:** A doctrinal regime (strict vs. negligence vs. no-fault) that shapes incentives, and determines which firms enter the market, and how safely they build.

The Ratio Method proposes to select, in advance, the liability regime whose anticipated incentives shape how an industry develops, which can include regulatory auditing by government agencies or independent parties, such as private insurers and underwriters.

Both frameworks are ex ante: each chooses the rules intended to steer the technology's development. The Ratio Method's mechanism is liability, but the policy choice it governs operates through anticipation, in the standard law-and-economics sense that the prospect of liability alters firms' behavior before they act.

Federal Aviation Act

Amodei's instinct is that AI should be regulated "like cars, airplanes, or drugs," with the FAA as his template. The Ratio Method reached for the same reference industries: aviation/FAA, pharmaceuticals/FDA, automobiles. But it uses the aviation framework for AI-driven IoT because, like aviation, IoT can control physical systems and threaten human safety; the drug/device and vaccine-injury regimes for CRISPR; and, notably, the oil-spill regime (OPA) for AI cybersecurity.

Amodei would put frontier AI broadly under an aviation-style pre-clearance model, which resembles the Ratio Method IoT framework. The Ratio Method viewed data breaches in AI-era as landing on oil-spill liability: certificates of financial responsibility, liability scaled to the size of the data "carrier," strict-but-capped damages.

Labor Displacement

Amodei proposes wage insurance, retention incentives, and potentially UBI financed by taxes on AI companies, while stressing that enduring job displacement is undesirable. The Ratio Method proposes a workers'-compensation-style no-fault scheme in which automating firms pay a share of their gains into retraining and support for displaced workers. Both authors land on a direct policy overlap: industry should help pre-fund the externalities it creates. Amodei's pledge to absorb datacenter-driven electricity price increases is a similar move in miniature.

The Extreme Tail – Nuclear Liability

At the catastrophic end, both documents converge again. Amodei says future AI may eventually look less like airplanes and more like weaponizable nuclear materials, a threat to humanity rather than just to public safety. Ratio Method, addressing a CRISPR-engineered synthetic pandemic, explicitly turns to nuclear plant liability (the Joint Protocol: strict liability, mandatory financial coverage, capped damages), because viral and radiological harms share that catastrophic, spreading character. They both treat nuclear governance as the reference point for "too dangerous for ordinary tort/regulatory tools."

Where they don't overlap

Amodei has substantial material with no analog in Ratio Method: a full section on civil liberties and state power, autonomous weapons, surveillance, the data-broker loophole, and an entire geopolitical strategy built around a democratic coalition, chip export controls, and AI as the new basis of national power. Ratio Method gestures at surveillance (IoT, the ACLU) and international bioweapons law, but they don't develop a theory of state power. Meanwhile, Ratio Method offers something Amodei lacks: a quantified methodology. Amodei states analogies at a high level, while Ratio Method offers a procedure with algorithmic inputs (Frequency, Average Damages, etc.) and worked examples (Anthem, Target, Deepwater Horizon). The authors even note, half in jest, that their automotive-industry calibration approximates to the Golden Ratio.

Why the comparison matters

The two documents share a diagnosis and come to similar conclusions. The reason the comparison merits attention is what the pattern reveals about the structure of the problem.

First, the convergence is independent. A thesis reasoning from possibility in 2021 and an AI-lab executive reasoning from observed model capabilities in 2026 arrive at several of the same reference points – aviation, pharmaceuticals, automotive, nuclear – and same principles of harm-tiered classification, producer cost internalization, and a reserved extreme-risk tail. Independent convergence of this kind is evidence that the authors are tracking structural features of the AI regulatory problem, and raises the credibility of the overlap since they have no common source.

Next, the comparison isolates a step in Amodei's argument. Amodei suggests regulating frontier AI "like airplanes," but does not specify why aviation is the correct referent. The Ratio Method provides a selection procedure for that step. Read together, the two works redirect the central question in AI governance toward which existing regime fits which risk. The comparison demonstrates that a disciplined, replicable method for that question may exist.

Additionally, Amodei's pre-deployment review approach determines what enters the market, while the Ratio Method's liability approach determines what framework governs a deployed system. The pieces are therefore complementary: Amodei's proposal does not specify what liability controls over a certified model; the Ratio Method is designed to resolve that question and structure incentives around it.

The comparison also highlights a consequential question: whether ex ante rules in AI regulation can be quantitatively determined, as Ratio Method suggests, or must remain adaptable, as Amodei proposes. This framing replaces a difference in approach with a specific, unresolved question: whether rules that precede harm can be structured or must be left deliberately open.

Finally, the Ratio Method's legacy-doctrine thesis transfers further than AI-exceptionalism may predict, but reaches its limit at two of Amodei's named risks, "loss of control" and recursive self-improvement. No prior industry has produced an artifact capable of pursuing its own objectives or improving its own capabilities. The comparison, therefore, establishes how much of AI governance can be adapted from existing regulation, and the specific areas it cannot, which points to where genuinely novel regulation is actually required.

Sources

Dario Amodei, *Policy on the AI Exponential* (June 2026), darioamodei.com/post/policy-on-the-ai-exponential

Harrison C. Margolin & Grant H. Frazier, *The Ratio Method: Addressing Complex Tort Liability in the Fourth Industrial Revolution*, 52 St. Mary's L.J. 679 (2021). Article page · SSRN

© 2026 Harrison C. Margolin · See also The Ratio Method report: Precedents