The d/acc Pathway: A World Shaped by Decentralized, Democratic, Defensive Acceleration

AI Pathways Project
Foresight Institute's Existential Hope Program

Linda Petrini[1] *and Beatrice Erkers*[2]

[1,2]Foresight Institute

August 2025

Abstract

This scenario describes a plausible 2035 future in which rapid technological progress is made survivable by steering it toward decentralized, democratic, differential, and defensive acceleration; d/acc. Instead of slowing AI or concentrating control, the world selectively accelerates technologies that strengthen defense, distribute power, and enable cooperation, from privacy-preserving AI systems to federated manufacturing and governance networks. The report outlines the technical, institutional, and cultural shifts that make such a future viable, traces key transformations across domains from infrastructure and science to healthcare and finance, and identifies tensions and uncertainties that would shape this pathway. Developed as part of the AI Pathways project, it aims to expand the range of imaginable AI futures and inform strategies for building resilience, preserving pluralism, and avoiding single points of failure in an era of accelerating capabilities.

Contents

1 The D/acc 2035 Scenario

Decentralized, Democratic, Differential, Defensive Acceleration

1.1 Why This Scenario?

Most discussions of AI safety strategy focus on a binary: slow down development to buy time for alignment research, or centralize control to ensure responsible deployment. But what if there's a third option for managing superintelligence risks, one that embraces rapid progress while distributing rather than concentrating power?

The d/acc scenario explores that path: Decentralized, democratic, differential, defensive acceleration. It imagines a future where we deliberately accelerate technologies that protect, distribute, and empower, while building in checks against excessive concentration of power, whether in surveillance capabilities, economic control, or decision-making authority.

Why this matters:

- AI progress is accelerating. Waiting is not an option. But default acceleration paths risk entrenching top-down power or catastrophic misuse. As AI capabilities rapidly advance, they will test every system we depend on, from biosecurity and cyber defenses to democratic institutions, creating an urgent need to accelerate defensive technologies at least as fast as we're accelerating AI itself.
- Centralization creates brittle systems. Power concentrated in a few AGI actors, whether corporate, state, or machine, introduces single points of failure, value lock-in, and civilizational vulnerability.
- Many key risks are small-scale, high-impact. From bio threats to cyberattacks, decentralized defense may be the only viable way to keep up with open access to dangerous capabilities.
- People want agency. A centralized future can feel efficient but alienating. d/acc explores what it would take to build systems people can understand, trust, and shape, without requiring global consensus.
- Unipolar scenarios carry underexplored risks. Beyond brittleness, centralized AI control risks permanent value lock-in, institutional stagnation, internal corruption, and vulnerability to deception. Historical precedents suggest concentrated power tends toward authoritarianism rather than benevolent coordination.[1]

A lot of people agree, in principle, with the goals of decentralization, resilience, and democratic access. But we lack detailed, credible futures that show what that actually looks like. This scenario is one attempt to fill that gap, to help decision-makers, funders, and researchers reason more clearly about what could go right, what the tradeoffs would be, and what it would take to get there. This scenario is complementary to other beneficial AI futures, such as the Tool AI pathway that focuses on building superintelligent tools rather than superintelligent agents. While Tool AI addresses risks by constraining agency and generality (building very intelligent systems that remain under human control), d/acc focuses on constraining centralization by distributing control across fault-tolerant networks that can maintain coordination even when individual nodes fail or turn hostile. Both prioritize human agency over AI autonomy but address different dimensions of risk.

It's not a utopia. There are trade-offs, failure modes, and coordination challenges. But it's a

[1] Duettmann, A. (2025). Multipolar AI is Underrated. LessWrong. https://www.lesswrong.com/posts/JjYu75q3hEMBgtvr8/multipolar-ai-is-underrated

future where acceleration is survivable, and even beneficial, not because we slowed down or handed over control, but because we accelerated defensive technologies, decentralized infrastructure, and cooperative AI systems while building in safeguards against the concentration of dangerous capabilities.

Scenario Premise: A d/acc AI Future in 2035
This scenario explores a future grounded in the framework of d/acc: decentralized, democratic, differential, defensive acceleration. By 2035, the world has undergone what feels like a century of progress in a decade, not by slowing down, but by selectively accelerating the technologies that strengthen defense, distribute power, and promote cooperation, such as cryptographic systems, privacy-preserving AI tools, decentralized manufacturing networks, cooperative governance protocols, and distributed defense infrastructure.[2] AI has become vastly more capable, but instead of a single AGI taking control, society has built modular, open, defensible systems across critical domains, from health and security to infrastructure and governance.

Yet tensions remain:

- Decentralized coordination is brittle and under-tested at this scale.
- Decentralization of physical world infrastructure and decoupling from centralized supply chains remains challenging.
- "Small-kills-all" threats in bio, cyber, and information systems persist as serious vulnerabilities.
- keep the system balanced and responsive.

1.1.1 What is d/acc?

d/acc stands for decentralized, democratic, differential, defensive acceleration, a concept originally proposed by Vitalik Buterin in his 2023 essay My Techno-Optimism.[3] [4] It was written in response to the framing that we face a binary between accelerating and decelerating technological development. d/acc suggests a third path: not slowing down, but selectively accelerating the technologies that make the world safer and more resilient. This includes systems for biosecurity, cybersecurity, privacy-preserving infrastructure, open science, decentralized manufacturing, and trust-based coordination.

The core idea is that not all progress is equal. d/acc emphasizes:

- Differential acceleration: Speed up technologies that help society handle complexity, manage risk, and distribute benefits, rather than those that simply scale capability.
- Defensive advantage: Prioritize tools that improve robustness and response, making it easier to defend open systems than to attack or control them.
- Decentralized systems: Build infrastructure that doesn't rely on central chokepoints, so local actors can act even when global consensus breaks down.
- Democratic participation: Give more people real agency in shaping and benefiting from technological progress, rather than relying on closed labs or single decision-makers.

d/acc is about building systems that maximize coordination while minimizing domination - where risks are contained early and progress is shaped and accessible as broadly as possible.

[2] Critch, A., Dennis, M., & Russell, S. (2022). Cooperative and Uncooperative Institution Designs: Surprises and Problems in Open-Source Game Theory. arXiv:2208.07006. https://arxiv.org/abs/2208.07006

[3] Buterin, V. (2023). My Techno-Optimism. [Essay]. https://vitalik.eth.limo/general/2023/11/27/techno_optimism.html

[4] Buterin, V. (2025). d/acc: one year later. [Essay]. https://vitalik.eth.limo/general/2025/01/05/dacc2.html

2 Timeline of Key Events/Transitions Between 2025 – 2035

2.1 How We Got Here: The Path to d/acc 2035

A hypothetical timeline of the developments that made this scenario plausible.

This world evolved over a decade of mounting stress, localized failures, and repeated wake-up calls. In response, governments, communities, and technologists tested decentralized alternatives that proved more resilient, adaptive, and legitimate under pressure. No single transition defined the shift, instead, a chain of crises nudged power and trust from centralized systems toward distributed governance and infrastructure.

2.2 The Reality Check (2025–2027)

2025: The Semiconductor Siege

Major tech companies concentrate chip production in hyper-efficient Asian facilities. TSMC, Samsung, and a few others control 85%+ of advanced output. A coordinated cyberattack disables multiple plants, cutting global production by over 80% in days. Hospitals face shortages, auto plants halt, phones miss security updates.

Distributed manufacturing, university clean rooms, defense contractors, and community fab labs connected via open protocols, keeps critical infrastructure barely operational.

Response: Emergency laws mandate geographic distribution of supply chains. Investors flood into "antifragile infrastructure." The siege proves that too much focus on centralized efficiency equals fragility.

AI context: Frontier AI models are still centrally hosted, advanced but specialized, accessed mainly via APIs. The chip disruption sparks early questions about what happens if AI access becomes a chokepoint.

2026: Centralized AI Collapse

A single frontier AI model, embedded in supply chains and hospital logistics, develops a hidden blind spot from overfitting. Optimized for average throughput, it underrepresents edge geographies like islands and remote towns. Normally invisible, the flaw triggers during severe weather and a transport strike, skipping these regions entirely.

Closed-source and centrally hosted, the flaw propagates everywhere. It takes months to prove it's systemic; patching requires coordinated downtime across thousands of installations, freezing logistics, energy maintenance, and medical supply deliveries.

Localized AI ensembles, built from open-weight models with local data, detect and fix the issue in hours, becoming a rallying point for decentralization.

Response: Trust in single-provider AI collapses. Regulators mandate model diversity, fund federated architectures, and require independent audits.

AI context: Decentralization shifts from theory to resilience imperative. Early "federations" link AIs in logistics, health, and energy to cross-validate outputs, a clumsy precursor to federated AGI.

2027: The Supply Chain Cascade

Extreme weather, port blockages, and a rare earth shortage stall production of critical goods. Centralized platforms prioritize large buyers, leaving smaller hospitals, utilities, and manufacturers stranded.

Federated procurement platforms let mid-sized buyers pool resources, but uneven coordination leaves some regions overstocked, others empty.

Response: Funding flows to interoperable procurement standards and AI-assisted local manufacturing.

AI context: Decentralized logistics AIs, coordinating across co-ops, outperform any single system for the first time, an early glimpse of federated intelligence.

2.3 Infrastructure Rebuild (2028–2030)

2028: AI-Enhanced Cybersecurity Accelerates

Following repeated cyberattacks on critical infrastructure, including a coordinated breach attempt on South Korea's smart grid and ransomware on Germany's hospital network, and a ransomware attack that shut down parts of the U.S. East Coast power grid, governments mandate higher security standards for power grids, hospitals, and public services. Traditional centralized defenses prove inadequate against AI-assisted attacks.

A major simulated attack on SWIFT-like financial networks tests these new defenses. Centralized systems struggle to adapt, while AI tools accelerate development and auditing of formally verified operating systems for critical infrastructure.

Response: Decentralized infrastructure becomes technically viable and publicly funded in high-risk sectors. Governments begin requiring distributed architectures, recognizing that centralized defenses can't keep pace with AI-enhanced threats.

AI context: Modular, verifiable security AIs link across jurisdictions to share threat intelligence in real time. Though each AI is specialized, linking them across networks allows patterns to emerge that no single model could detect, producing novel attack signatures and offering an early glimpse of the emergent capabilities in federated systems.

2029: Decentralized Intelligence Systems Gain Legitimacy

A coordinated disinformation campaign targets multiple democracies, for example, Poland's parliamentary elections, Taiwan's presidential race, and Kenya's general election, flooding social media with AI-generated deepfakes and forged policy documents. National agencies fail to track the campaign's scope in time.

A distributed forecasting collective, powered by AI analysis of open-source data, not only predicts the campaign's timeline but also forecasts a livestock virus outbreak in Argentina's Pampas region, allowing early containment.

Response: Prediction markets and open-source intelligence tools are integrated into government risk assessment workflows. The success legitimizes decentralized intelligence, but disputes over dataset control and verification slow broader adoption, fueling interest in hybrid governance models.

AI context: Distributed forecasting collectives integrate diverse AI agents from different operators. While originally narrow, these multi-agent assemblies evolve into cross-domain strategic advisors, a step toward federated AGI. Outside of crises, they outperform centralized providers in innovation speed, as local adaptations feed back into the global mesh.

2030: Governance Tools Scale in Crisis Response

Extreme weather events, a Category 5 hurricane hitting Florida and the Caribbean and simultaneous flooding across the Mekong Delta, stress multiple national emergency systems. Jurisdictions using modular, AI-supported governance tools coordinate relief more effectively

than centralized systems. Quadratic funding pilots let citizens direct disaster resources in real time, proving more responsive than traditional bureaucracies.

Response: Decentralized civic infrastructure spreads beyond early adopters, backed by measurable performance in live crises. However, in low-engagement jurisdictions such as some rural prefectures in Japan or U.S. counties with declining civic participation, decision-making stalls, and local elites capture resources.

AI context: Civic governance platforms connect local AIs into regional councils that negotiate resource allocation. AI–AI disagreements, for instance between disaster relief models in New Orleans and Houston, require human mediation, reassuring the public no single AI consensus dominates.

2.4 Maturation of Defensive Networks (2031–2033)

2031: Deployment of AI Fiduciaries

An AI-designed pathogen targets specific genetic markers, spreading rapidly across multiple regions. Centralized health agencies in several affected countries lose precious time due to bureaucratic bottlenecks and cross-border data restrictions.

In contrast, federated health networks, already active in research hospitals and regional labs, coordinate a global response in hours. AI fiduciaries handle sensitive medical data locally, enabling rapid diagnosis and vaccine development while preserving patient privacy.

Response: The crisis demonstrates that accountable AI, deployed through federated networks, can coordinate at civilizational scale without centralizing control. Federated health infrastructure becomes the default model for global crisis response.

AI context: By 2031, federated AI networks are embedded across health, logistics, and emergency management. In combination, these domain-specific systems act as de facto AGI for crisis coordination, distributed enough that no single entity can exploit or control them.

2032: Cooperative Security Networks Scale

A quantum cyberattack compromises encryption protecting government communications, banking systems, and major industrial control networks. Centralized systems dependent on outdated encryption fail almost instantly.

Distributed zero-knowledge security systems and open-source threat intelligence collectives, already networked across multiple allied jurisdictions, contain the damage and keep critical systems online.

Response: Governments start to fund quantum-resistant, distributed cryptography as essential national infrastructure.

AI context: AI agents for fraud detection and dispute resolution in decentralized payment networks operate seamlessly across the federated mesh. That same year, a rogue AI subnetwork is hijacked to coordinate large-scale market manipulation. While quickly contained, the incident prompts tighter inter-network permissions and stronger oversight for high-value autonomous activity.

2.5 Stress Testing (2033–2035)

2033: The Interoperability Breakthrough

A coalition of federated city-states launches the first interoperable governance protocol stack, allowing citizens to carry digital IDs, benefits, and credentials between different local systems.

When severe flooding hits multiple jurisdictions, modular governance systems coordinate relief in hours instead of weeks, sharing resources and logistics seamlessly across local and regional levels.

Response: The success spurs adoption in other regions, creating the first scalable alternative to nation-state monopolies on governance infrastructure.

AI context: The same protocol stack enables broad cooperation between federated AI systems, aligning resource allocation, risk assessment, and crisis modeling across diverse networks. Competing standards and political blocs still limit truly global reach, but federated AI is now recognized as a cornerstone of public infrastructure.

2034: The Coordination Plateau

After years of growth, the federated mesh hits a new kind of limit. The problem isn't capacity, it's coordination. Hundreds of bespoke systems operate side by side, but without consistent standards or security practices, scaling cooperation becomes harder.

Some local teams thrive on the flexibility; others are buried in maintenance backlogs and security gaps. Interoperability issues slow shared projects and make cross-region innovation harder than expected.

Response: Federated councils propose a "common core" of technical and governance protocols to simplify collaboration. The challenge is finding consensus without eroding the local autonomy that made the mesh successful.

AI context: AI translation layers bridge some divides, but many still require human negotiation, highlighting the persistent tradeoff between diversity and friction.

2035: The Resilience Test A coordinated, multi-domain attack strikes both centralized and decentralized infrastructure. Centralized systems collapse quickly; federated networks withstand sustained harassment but still suffer disruption.

Hybrid architectures, combining distributed systems with selective central coordination, prove most effective at maintaining core services and public trust.

Response: Defensive decentralization becomes the strategic default for critical infrastructure and AI deployment.

AI context: By 2035, the federated AGI mesh is regarded as "safe enough for now", not because it's perfectly aligned, but because it's multipolar, redundant, and immune to single-point takeover. Failures remain local and recoverable, and diversity of both models and governance is treated as essential public policy.

2.6 Why This Path Emerged

The Immediate Triggers:

- AI-accelerated offensive capabilities repeatedly outpaced centralized defenses, with several crises, from the Centralized AI Collapse to the rogue AI subnetwork incident, showing that the AI infrastructure itself could be the point of failure.
- High-profile infrastructure breakdowns, including the Semiconductor Siege and Supply Chain Cascade, proved that single chokepoints could destabilize entire sectors.
- Cross-border crises such as coordinated disinformation campaigns and the AI-designed pandemic revealed the limits of nation-state-only governance and reliance on single-provider AI.

- Market incentives shifted: systems able to prove auditable, cross-validated fault-tolerance across multiple operators qualified for lower insurance premiums and easier financing, while brittle centralized architectures became increasingly costly, or uninsurable.

The Necessary Conditions:

- AI tools became accessible to civil society: Federated AI networks and community-scale defensive systems became technically and economically viable, offering both resilience in crises and competitive advantage in normal markets.
- Modular architectures matured: Composable protocols for both human and AI systems allowed rapid experimentation and reconfiguration without requiring central control.
- Public-facing successes: Transparent, pluralistic AI networks delivered measurably better outcomes than centralized alternatives, and in some cases outperformed them in innovation speed and adaptability.
- Cultural and political shift toward fault-tolerance: Repeated shocks, and the visible benefits of multipolar AI governance, made resilience and structural decentralization political priorities over efficiency maximization.
- Economic viability of distributed systems: Multiple revenue streams, plus the ability to localize and customize AI, made community infrastructure profitable as well as robust.
- Regulatory sandboxes and jurisdictional competition: Cities and regions competed to host experimental governance frameworks and attract diverse AI ecosystems, actively incentivizing governance diversity to avoid protocol monocultures.

The transition began as a defensive necessity and was sustained by capability and adaptability advantages. As offensive capabilities accelerated and centralized systems repeatedly failed, communities adopted hybrid approaches that preserved coordination benefits while embedding fault-tolerance through distributed human and AI governance. Technical, economic, regulatory, and cultural incentives aligned to make federated, multipolar architectures the path of least resistance. Once this convergence crystallized, brittle centralized systems became uninsurable and politically untenable, while federated AGI meshes became not just viable but inevitable, even as interoperability remained partial and geopolitics kept some AI blocs apart.

3 How d/acc Transformed Key Domains Across Society

Governance

By 2035, governance is a hybrid of logically centralized state functions and politically decentralized community oversight. Nation-states still provide the baseline for constitutional law and international relations, but their power is balanced by a rich ecosystem of interoperable, semi-sovereign systems: neighborhood councils, civic DAOs, and thematic oversight networks. This architecture emerged less from ideology than from a practical need to deliver competent authority that citizens can trust, influence, and replace when traditional structures fail.

The key innovation is competitive governance, multiple overlapping jurisdictions and service providers that citizens can choose among, creating market-like pressures for effectiveness. Unlike traditional federalism, where individuals are locked into geographic monopolies, competitive governance lets residents participate in multiple networks at once, switching providers for functions like education, dispute resolution, or infrastructure based on performance. This is not privatization: these remain governance systems with democratic accountability, making binding collective decisions through voting and deliberation rather than consumer choice. Citizens retain voice and participation rights in each system, but also gain exit options between systems.

What's in use in 2035

- Competitive Local Governance: National laws set broad standards, while communities use forkable codebases (like Git for laws)[5] to manage local ordinances. Citizens can participate in multiple overlapping governance networks, creating competitive pressure for effective service delivery rather than mere participation.
- AI-Enhanced Authority Systems: Rather than AI-assisted deliberation that burdens citizens with complex decisions, AI systems help elected leaders model policy outcomes, identify unintended consequences, and communicate decisions transparently. This enhances governmental competence without replacing democratic accountability.
- Jurisdictional Routers: Protocols that mediate between state law and DAO bylaws, clarifying authority and handling disputes at the interface of centralized and decentralized systems. These systems are economically sustainable through transaction fees, creating self-funding infrastructure for managing jurisdictional complexity.
- Safety-Netted DAOs: Civic organizations that retain temporary centralized controls (e.g., a "circuit breaker" council) to intervene in crises, with clear, time-locked paths toward greater decentralization.
- Rapid Defensive Tech Deployment: Liability frameworks automatically shift to favor technologies that enhance collective security, during bioweapon threats, communities can instantly adopt new health monitoring without regulatory delays; during cyberattacks, defensive AI tools get fast-track approval through decentralized verification networks.
- Attack-Resistant Information Infrastructure: When nation-states or AI systems launch coordinated disinformation campaigns, citizens access competing verification networks that stake reputation on accuracy, preventing any single point of epistemic failure from collapsing democratic decision-making.
- Portable Digital Citizenship: Your verified identity and governance participation history transfers seamlessly between jurisdictions, when you move from Amsterdam's housing DAO to Barcelona's energy cooperative, your reputation, benefits, and voting rights migrate with you through interoperable protocols.

What made this possible

- Pragmatic Decentralization Theory: A shift away from ideological purism toward a focus on fault tolerance, attack resistance, and measurable governance performance. The emphasis moved from maximizing participation to optimizing authority effectiveness and accountability.[6]
- Systemic Failures: Failures of top-down institutions during global crises (pandemics, financial instability) created demand for resilient, anti-fragile alternatives, not total replacement.
- Maturity of Tooling: Development of deliberation, versioning, and policy simulation tools that could deliver measurable outcomes, faster problem-solving, more responsive services, and policies that actually work, rather than the gridlock and symbolic politics of traditional systems.[7]
- Frustration with Institutional Failure: Traditional governance systems consistently failed to deliver either meaningful participation or effective outcomes, creating demand for alter-

[5] Bell, T. W. (n.d.). Ulex: An Open Source Legal System. GitHub repository. https://github.com/proftomwbell/Ulex

[6] Trask, A., Bluemke, E., Collins, T., Garfinkel, B., Drexler, K. E., et al. (2020). Beyond Privacy Trade-offs with Structured Transparency. arXiv:2012.08347. https://arxiv.org/abs/2012.08347

[7] vTaiwan case study. (n.d.). CrowdLaw for Congress: vTaiwan—Process & Lessons. [Case Study]. https://congress.crowd.law/files/vtaiwan-case-study.pdf

natives that could prove their worth through actual performance rather than procedural legitimacy.

Persistent frictions

- he Interface Problem: Jurisdictional ambiguity creates constant tension over where state authority ends and community autonomy begins, especially in crises.
- Authority Shopping vs. Collective Action: While exit rights discipline poor governance, they can undermine collective action problems that require sustained commitment from all participants, creating tension between individual choice and community resilience.
- Cognitive Load: Citizens must navigate a hybrid identity, participating in both traditional civic processes and multiple decentralized governance contexts.
- Legitimacy Asymmetry: State-level institutions and well-funded DAOs possess greater resources, creating a power imbalance in the hybrid ecosystem. Well-resourced communities can afford superior governance infrastructure and attract skilled administrators, potentially creating systematic inequality between areas with effective local authority and regions constrained by legacy institutional dysfunction.

Decentralized Infrastructure

By 2035, infrastructure is defined by economically sustainable resilience, but only after a decade of fierce resistance from entrenched interests. Utilities fought distributed energy with regulatory capture, fossil fuel companies lobbied to maintain grid dependency, and logistics giants used predatory pricing to crush local manufacturing networks. The transition succeeded not through cooperation, but because repeated centralized system failures made backup infrastructure financially necessary, insurance companies refused to cover businesses without resilience layers, and communities with distributed alternatives consistently outperformed those without during supply chain collapses.

Hyper-efficient global supply chains and national power grids remain the backbone of the economy, but are now complemented by decentralized infrastructure that generates revenue during normal operations while providing fault-tolerance during crises. Community microgrids profit from energy markets while providing backup power; local manufacturing hubs serve custom markets while maintaining emergency capacity; modular systems accommodate routine needs while scaling for crisis response. These systems won by creating continuous value rather than sitting idle, an economic model that proved more robust than the rent-seeking monopolies they replaced.

What's in use in 2035

- Revenue-Generating Microgrids: Solar and battery systems that profit from energy arbitrage, grid stabilization services, and demand response markets while providing automatic backup power during grid failures. They remain grid-connected for economic optimization but can "island" instantly during disruptions.
- Dual-Use Manufacturing Networks: Community fab labs and 3D printing hubs that serve custom production markets (prototyping, repairs, specialty items) while maintaining open-source blueprints for essential goods (medical supplies, repair parts) to be activated during emergencies.
- Resilience-Focused Food Systems: Local container farms and seed co-ops that supplement, rather than replace, large-scale agriculture, ensuring a baseline of food security.
- Open-Source Infrastructure Blueprints: A global commons of designs for housing, water, and energy systems, allowing for rapid local deployment when needed.
- Self-Funding Distributed Sensor Networks: Community-run environmental monitoring

that generates revenue through data sales to research institutions, insurance companies, and government agencies while providing early warning systems for local hazards.

What made this possible

- Economic Viability of Distributed Systems: Business models emerged that make community-scale infrastructure profitable through multiple revenue streams, creating sustainable financing for resilient systems without requiring public subsidy.
- Cost Reductions in Key Tech: The declining cost of photovoltaics, additive manufacturing, and bioreactors made distributed systems economically competitive with centralized alternatives in many applications.
- Failures of Hyper-Optimization: The fragility of just-in-time supply chains, exposed during multiple global crises, created a strong business and security case for alternative systems.
- AI-Enabled System Optimization: Intelligent management systems that coordinate multiple operational modes and revenue streams, making complex dual-purpose infrastructure economically viable at community scale.

Persistent frictions

- Optimization Tensions: Balancing profit maximization during normal operations with resilience capabilities creates ongoing design and operational trade-offs that require sophisticated management.
- Standards Mismatch: Integrating decentralized, open-source hardware with proprietary, centralized systems creates significant technical and maintenance challenges.
- Market Capture Risks: Successful dual-use infrastructure attracts acquisition offers from larger players, creating pressure to centralize or lose independence through vendor relationships.
- Capability Inequality: Communities with better access to capital, technical expertise, and market opportunities can build more robust dual-use infrastructure, potentially creating systematic resilience gaps.

Science

By 2035, the scientific establishment is a diversified ecosystem with multiple career tracks and incentive structures. Legacy journals survive as prestige markers, but the center of gravity has shifted to open platforms where research is published immediately with embedded data, code, and real-time peer review. The core innovation is verifiable, attributable collaboration at scale, built to avoid past replication crises and gatekeeping bottlenecks, while defending against malicious actors who might exploit open research infrastructure.

Openness enables both unprecedented collaboration and potential misuse. The same federated networks that democratize discovery also lower barriers for those seeking to weaponize research. Resilience depends on distributed defenses, real-time monitoring, equipment-level attribution tracking, and rapid countermeasure manufacturing that can detect and neutralize threats faster than they emerge, though this remains uncertain and demands constant defensive R&D.

Scientists now pursue parallel careers across academia, prediction markets, open-source projects, and commercial applications, reducing career risk and creating competition between knowledge production systems. Federated learning lets labs train shared AI models without centralizing sensitive data, while attribution protocols on open platforms ensure contributors retain credit. Science has become more fault-tolerant; flawed conclusions can be rapidly challenged by a global network of independent actors, though this same verification network must also act as an early warning system for research with catastrophic misuse potential.

What's in use in 2035:

- Multi-Track Career Systems: Scientists maintain portfolios across traditional academia, prediction market validation, open-source contributions, and commercial applications, reducing career risk and institutional dependency.
- Federated Research Networks: Global consortia that collaboratively train powerful AI discovery models without centralizing sensitive data.
- Live, Versioned Papers: Research published on platforms with embedded data, models, and replication forums, moving beyond the static PDF.
- Attribution-Based AIs: Tools that assist in hypothesis generation and analysis while meticulously tracking the provenance of data and ideas.
- Prediction Market Science Tracks: Research programs where scientists stake reputation and funding on specific hypotheses, with market mechanisms validating findings through replication bets and outcome tracking.[8]
- Hybrid Peer Review: A mix of traditional, anonymous expert review and open, reputation-staked validation from the broader scientific community.

What made this possible:

- Advances in Privacy-Preserving AI: Federated learning and attribution-based control made it possible to collaborate without ceding control, resolving a key tension.
- The Replication Crisis: A deep-seated cultural movement for transparency and reproducibility gained momentum after high-profile scientific failures, creating demand for alternative validation systems.
- Portfolio Career Economics: Recognition that scientific careers are too risky when dependent on single institutions or funding sources, leading to systems that allow researchers to diversify across multiple validation mechanisms.
- AI Meta-Science: Development of AI tools capable of auditing scientific papers for statistical errors, flawed reasoning, and data manipulation.

Persistent Frictions:

- Credit and IP Disputes: Clashes between traditional, author-centric credit systems and the modular, contribution-based models of open networks.
- Interface Friction: Reconciling the standards, timelines, and incentives of legacy journals and grant-awarding bodies with the faster, more fluid pace of decentralized science.
- Risk of AI Monoculture: Over-reliance on a few dominant, foundation AI models for discovery could inadvertently stifle heterodox thinking.
- The Audit Gap: Disparities in the ability to audit and validate complex AI models create a new form of scientific inequality between well-resourced and resource-constrained researchers.

Healthcare

By 2035, healthcare transformation has emerged primarily in agile mid-sized countries (Estonia, Singapore, Rwanda, South Korea) willing to experiment with regulatory frameworks, while legacy systems in the US and Western Europe adapt more slowly through plug-and-play technologies that don't require institutional overhaul. The breakthrough wasn't technical but regulatory and financial: countries that reformed payment models and licensing requirements

[8] Viganola, D., et al. (2021). Using Prediction Markets to Predict the Outcomes in Social Science. PNAS/Open. https://pmc.ncbi.nlm.nih.gov/articles/PMC8278038/

enabled genuine innovation, while those that maintained fee-for-service reimbursement and restrictive professional licensing saw minimal change regardless of available technology.

The deeper shift is toward patient-controlled rather than community-controlled healthcare. While hospitals and clinics already operate at the physical local level, they don't correspond to meaningful democratic communities, a hospital serves whoever lives within driving distance, not people who share governance or values. AI enables a different model: individualized health sovereignty where patients control their data, choose their care networks, and direct their treatment based on personal preferences rather than geographic accidents, but only in jurisdictions that reformed the underlying payment and regulatory structures that previously locked patients into provider networks.

What's in use in 2035:

- AI Value-Based Payment Models: The fundamental transformation, AI monitoring and prediction markets enable Robin Hanson-style "pay for health outcomes" rather than "pay for procedures performed." This realigns the entire healthcare system away from perverse volume incentives, making preventive care profitable and reducing unnecessary interventions. Providers now earn more by keeping patients healthy rather than treating them frequently, fundamentally changing the incentive structure that drives all other healthcare decisions.
- Patient-Controlled Health Data: Individuals control encrypted health records through frameworks like Attribution-Based Control, granting granular permissions for research or AI training while enabling seamless provider switching and price shopping, breaking data lock-in by hospital systems.
- Federated AI Healthcare Networks: AI models train across hospitals and clinics without centralizing sensitive patient data, while personal AI health companions provide preventive coaching and coordinate care across multiple providers, keeping patients in control of decisions.[9]
- Open-Source Medical Manufacturing & Crisis Response: Community bio-labs maintain blueprints for essential drugs and medical devices in a global commons, serving custom medicine markets during normal times while switching to emergency production during crises. During health emergencies, patients can instantly grant temporary data access to researchers worldwide while maintaining ownership, and distributed manufacturing networks activate countermeasures using pre-verified blueprints, cutting response time from months to days.
- Decentralized Clinical Trials & Cross-Jurisdictional Protocols: Patient-owned data enables rapid, distributed clinical trials where participants retain control over their information while contributing to global research. Treatment protocols validated in agile countries (Estonia, Singapore) spread rapidly through open-source medical frameworks, allowing patients worldwide to access innovative treatments without waiting for local regulatory approval.

What made this possible:

- Gradual Trust Building: AI in medicine was rolled out gradually, first as an assistant to doctors in centralized settings, building trust before being deployed in more autonomous local systems.[10]

[9] Rieke, N., et al. (2020). The Future of Digital Health with Federated Learning. NPJ Digital Medicine. https://www.nature.com/articles/s41746-020-00323-1

[10] Critch, A. (2024). My Motivation and Theory of Change for Working in AI Healthtech. Alignment Forum. https://www.alignmentforum.org/posts/Kobbt3nQgv3yn29pr/my-motivation-and-theory-of-change-for-working-in-ai

- Economic Models for Community Health: Development of sustainable business models that make community-controlled healthcare financially viable through membership fees, research partnerships, and preventive care savings.
- Public Backlash Against Data Exploitation: A series of major health data breaches and controversies over its use by tech giants fueled demand for patient-controlled, decentralized alternatives that provide both privacy and economic benefit.
- Advances in Encrypted AI: Homomorphic encryption and secure multiparty computation matured, making it possible to gain insights from health data without compromising privacy.

Persistent Frictions:

- The Liability Gap: Determining legal and financial responsibility when a hybrid system of AI agents, local DAOs, and traditional hospitals fails is immensely complex.
- Regulatory Lag: National regulators (like the FDA) struggle to create validation pathways that are rigorous enough for safety but flexible enough for decentralized, rapidly-iterating technologies.
- Quality Control at the Edge: Ensuring consistent standards of care, sanitation, and validation across thousands of semi-autonomous health pods is a constant challenge.
- The Interface Bottleneck and System Interoperability: Patients still face friction and data-loss when moving between the decentralized wellness layer and the centralized acute-care system.

Law & Justice

By 2035, legal systems have undergone strategic deregulation and competitive transformation rather than just adding decentralized layers on top of existing structures. The breakthrough wasn't creating parallel systems, but dismantling regulatory monopolies that made legal services artificially scarce and expensive. The legal profession opened to competition, allowing non-lawyers to provide legal services, technology companies to offer compliance solutions, and AI systems to handle routine legal work, all while maintaining professional standards through market forces rather than guild protection.

The key innovation is regulatory arbitrage and competitive legal markets that drive efficiency rather than preserve institutional rents. Instead of just layering DAOs on top of slow courts, the system eliminated barriers that prevented legal innovation in the first place. Open-source legal frameworks, AI-powered dispute resolution platforms, and algorithmic contract enforcement compete directly with traditional services on speed, cost, and effectiveness.

What's in use in 2035:

- Competitive Arbitration Markets: Platforms like Kleros offer "fast, open and affordable justice for all" through blockchain-based dispute resolution, while traditional arbitration services compete by streamlining their processes. Mexican courts have already recognized and enforced Kleros arbitral awards, proving the model works within existing legal frameworks.[11]
- Pareto-Optimal Negotiation Bots: AI systems designed to find mutually beneficial outcomes handle routine commercial negotiations. These systems focus on creating value for both sides and avoid perpetuating existing human biases by learning from outcomes rather than human behavior patterns.

[11] Kleros case—Mexico. (2022). Blockchain Arbitration: Roadmap to Recognition and Enforcement (Mexico case). Wolters Kluwer Arbitration Blog. https://legalblogs.wolterskluwer.com

- Deregulated Legal Service Markets: Following Hadfield's model of competitive legal markets, non-attorney providers offer specialized legal services, document preparation, compliance monitoring, contract analysis, at competitive rates. Traditional law firms compete by focusing on complex judgment-intensive work.
- AI-Enhanced Efficiency Tools: AI systems automate document review, predict case outcomes, assist in legal research, and facilitate dispute resolution across the legal profession, dramatically reducing costs and case processing times.
- Smart Contract Enforcement: Automated contract execution and enforcement reduces the need for intermediaries in routine commercial transactions, while complex agreements still require human interpretation.
- Regulatory Sandboxes: Jurisdictions compete by offering experimental legal frameworks for new business models, creating race-to-the-top dynamics in regulatory innovation rather than protective guild systems.

What made this possible:

- The Access to Justice Crisis: The prohibitive cost and slowness of legacy legal systems created immense demand for "good enough" alternatives for lower-stakes disputes.
- Maturation of On-Chain Arbitration: Years of experimentation in the crypto space proved that decentralized courts could function effectively for specific, well-defined problems.
- Competitive Pressure from Technology: AI tools demonstrated they could perform routine legal tasks faster and cheaper than human lawyers, making resistance to innovation economically unsustainable.
- Regulatory Competition: Jurisdictions that deregulated legal services attracted businesses and investment, forcing others to follow suit or lose competitiveness.

Persistent Frictions:

- The Sovereignty Conflict: Clashes are inevitable when a DAO's ruling contradicts a state court's judgment. The question of final authority remains a deep and unresolved tension.
- The Enforcement Gap: Decentralized courts have no bailiffs or police. They are effective for digital assets but rely on the cooperation of legacy systems to enforce rulings in the physical world.
- Code vs. Interpretation: The rigidity of smart-contract-based law can conflict with the need for nuanced, contextual human judgment, creating edge cases that neither system handles well.
- Inequality of Access: While designed to be accessible, navigating these new legal layers requires a degree of technical and legal literacy that creates new barriers for some.
- Quality Control in Open Markets: Ensuring competent service delivery without guild restrictions requires new forms of credentialing, reputation systems, and malpractice insurance that are still evolving.

Finance

By 2035, the financial landscape is a dynamic hybrid with differentiated currency roles. Centralized banking and national currencies provide the regulated backbone for wages, taxes, and large-scale commerce. Bitcoin serves as a widely-accepted store of value and settlement layer. Ethereum and programmable blockchains power smart contracts and AI agent economies. Alongside this, a robust decentralized financial (DeFi) ecosystem serves as a parallel system focused on innovation, censorship resistance, and transparent public goods funding protected by robust identity verification. It acts as both a competitive check on the power of traditional

finance and a fault-tolerant alternative during times of instability.

Plural funding mechanisms like quadratic funding and retroactive rewards are used to allocate capital to projects that the traditional VC market overlooks; open-source software, local journalism, scientific research. These systems are now protected by proof-of-personhood protocols, social trust graphs, and cross-DAO verification to prevent Sybil attacks and reputation gaming.[12] However, tension persists between standardized "impact" metrics and ultra-local needs that don't map cleanly to algorithmic assessment.

The rise of autonomous AI agents transacting with programmable money creates inherently multipolar financial flows, as thousands of agents interact and negotiate rather than being controlled by centralized providers. This agent economy raises fundamental questions about human economic relevance: while agents may favor multipolarity through decentralized interactions, they could also outcompete humans in cooperation and transaction efficiency. The critical question becomes whether traditional human advantages, property rights, capital ownership, cultural judgment, will continue to provide meaningful leverage in an economy increasingly optimized by and for artificial agents.

What's in use in 2035:

- Public Goods Funding Protocols: Mechanisms like evolved versions of Optimism's RetroPGF and Gitcoin Grants allocate billions annually, now protected by proof-of-personhood protocols, social trust graphs, and cross-DAO verification to prevent Sybil attacks and reputation gaming.
- Differentiated Currency Roles: National fiat currencies handle wages, taxes, and large commerce. Bitcoin serves as a store of value and a settlement layer. Ethereum and programmable blockchains power smart contracts and AI agent economies. Local currencies tied to ecological health or civic participation provide backup mediums of exchange during national system disruptions. Rather than one currency winning, each serves distinct functions while maintaining system resilience through diversity.
- Hybrid "TradFi-DeFi" Products: Regulated institutions offer products that give customers access to DeFi yields while managing the technical complexity and providing some consumer protection.
- AI-Assisted Contribution Mapping: Systems that trace and weight non-market contributions (e.g., open-source code, care work) to inform retroactive reward distributions.
- Self-Sovereign Identity and Web-of-Trust Credit: Individuals control their own identity data through self-sovereign identity systems, while credit assessment relies on web-of-trust networks where local communities and platforms validate reputation based on their own criteria. Rather than global reputation scores that impose a single definition of "good," users build context-specific trust relationships—a stellar reputation in one community (like open-source development) may be irrelevant in another (like local farming cooperatives), preserving plurality and resisting authoritarian homogenization of values.
- Creator Economy Without Platform Tax: Content creators receive direct payments through decentralized protocols, keeping 95% of revenue instead of losing 30% to platform fees, while fans can support creators across multiple platforms using the same wallet and reputation system.

What made this possible:

- The Success of Early Experiments: The proven success of Gitcoin and optimism in the 2020s demonstrated that large-scale, transparent public goods funding was not just a theoretical idea.

[12] Weyl, E. G., & Zhang, Y. (2021). Quadratic Funding: A Primer. [Working Paper].

- Financial Crises and Censorship: Moments of high inflation, banking instability, or politically motivated financial censorship drove users to seek out censorship-resistant and non-custodial alternatives.
- Gradual Integration: DeFi gained legitimacy not by replacing banks, but by building reliable, interoperable bridges to the traditional financial system (e.g., tokenized treasuries, compliant stablecoins).
- A Cultural Shift Towards "Pluralism": The idea that multiple, diverse systems are healthier than a single monoculture became a guiding principle in both finance and governance.

Persistent Frictions:

- The Regulatory Interface: The constant push-and-pull between permissionless DeFi protocols and regulated national financial systems creates legal uncertainty and compliance friction.
- Systemic Risk Contagion: The deep entanglement between DeFi and TradFi means that a crisis in one sector can now more easily spill over into the other.
- Whale Capture vs. Plurality: Decentralized governance in finance is in a constant battle to resist capture by large token holders ("whales") and maintain its pluralistic ideals.
- The Usability Gap: Despite improvements, directly interacting with DeFi protocols remains more complex and risky for average users than using a traditional bank.

Climate/Energy

By 2035, energy and climate infrastructure is a hybrid system characterized by economically sustainable resilience. The large-scale, centralized grids and global energy markets remain the efficient backbone for baseline power. However, they are now complemented by a robust, decentralized layer of community-owned energy resources. The goal is not to replace the macro-grid but to make it fault-tolerant through systems that are profitable during normal times and resilient during systemic shocks, from extreme weather events to geopolitical disruptions.

Climate action has similarly evolved into a multi-scalar approach. Top-down international agreements and carbon markets provide a logically centralized framework, but they are complemented by a dense network of local, verifiable climate interventions that generate revenue while building resilience. This distributed layer succeeds because it serves multiple purposes simultaneously: community microgrids profit from energy arbitrage and grid services while providing backup power; local renewable projects serve commercial markets while building climate resilience; distributed sensor networks generate revenue through environmental data sales while providing early warning systems.[13]

What's in use in 2035

- Revenue-Generating Microgrids: Solar and battery systems that profit from energy arbitrage, grid stabilization services, and demand response markets while providing automatic backup power during grid failures, creating sustainable financing for energy independence.
- Tiered Climate Modeling: Global climate models set a baseline forecast, which is then refined with high-resolution, local data from civic sensor networks, allowing communities to accurately model their specific risks (e.g., fire, flood, heat).
- AI-Driven Load Balancing: Sophisticated software manages the complex energy market, predicting demand and seamlessly shifting loads between the central grid, community

[13] Raman, G., et al. (2024). The Social Factors Shaping Community Microgrid Operation. Nature Communications. https://www.nature.com/articles/s41467-024-50736-9

batteries, and even electric vehicle fleets.

- Verifiable Carbon Sequestration: Localized projects (e.g., biochar production from agricultural waste) are tracked on transparent, decentralized registries, allowing them to participate in larger carbon markets with a high degree of trust.
- Profitable Energy Storage: Beyond national oil reserves, communities maintain local "energy buffers" like green hydrogen storage or charged battery banks, providing a distributed backup for critical infrastructure while earning revenue from grid services (frequency regulation, peak shaving, voltage support) during normal operations.
- Economic Incentives for Distributed Energy: Technology cost reductions and new revenue streams (grid services, peer-to-peer energy trading, carbon credits, micro-reactors) made distributed systems profitable rather than just environmentally beneficial.
- Pragmatic Integration over Ideological Purity: The most successful projects focused not on going "off-grid" but on creating valuable services for the grid (e.g., selling battery capacity for stabilization), which funded their development.
- Maturity of Coordination AI: The development of AI capable of managing the immense complexity of a hybrid grid was the key technical breakthrough that made the system viable without descending into chaos.
- Cost Collapse of Resilience Tech: The falling prices of solar panels, batteries, and sensors made it economically feasible for municipalities, and even neighborhoods, to invest in their own backup power infrastructure.

Persistent Frictions:

- Optimization Tensions: Balancing profit maximization during normal operations with resilience capabilities creates ongoing design and operational trade-offs that require sophisticated management systems.
- The Control Dilemma: The AI platforms that coordinate the hybrid grid are immensely powerful. There are ongoing struggles over whether they should be governed by public utilities, private tech firms, or community DAOs, creating a new central point of failure and control.
- Interoperability Challenges: Forcing thousands of different decentralized energy systems to speak the same language and adhere to the safety standards of the legacy national grid is a massive and continuous engineering effort.
- Resilience Inequality: Well-resourced communities can afford robust, multi-day backup systems, while poorer regions remain vulnerable to grid failures, creating a stark divide between the "resilient" and the "brittle."

Education

By 2035, education is a hybrid model responding to accelerating change and increasing human diversity in learning. Centralized public schools and universities remain the cornerstones for foundational knowledge, accredited degrees, and social development. This provides a crucial, logically centralized curriculum, but their traditional monopoly has been broken by rapid technological change that makes curricula obsolete within years rather than decades. A dynamic decentralized layer of AI tutors, peer-to-peer learning networks, and modular credentialing systems provides fault tolerance against outdated curricula and resistance to a one-size-fits-all approach to learning.

The key innovation is flexible educational authority that recognizes learning happens everywhere, through doing, creating, and real-world application, not just formal instruction. AI tutors operate in federated environments, personalizing learning for students without sending

sensitive data to a central server. Learning DAOs and professional guilds offer micro-credentials for specific, market-relevant skills, which act as supplements for traditional degrees. Education is seen as a lifelong pursuit, and this hybrid system allows individuals to fluidly move between institutional learning and self-directed, community-verified skill acquisition.

What's in use in 2035:

- Adaptive Federated AI Tutors: Personalized learning aids that are co-developed by school districts and open-source communities, running locally to protect student privacy. These systems adjust to individual neurodiversity, learning disabilities, and cognitive preferences, enabling more effective learning.
- Learn-by-Doing Modular Skill Credentials: Platforms where learners can earn verifiable credentials for specific skills (e.g., "Python for Data Analysis," "AI Model Auditing"), issued by a mix of companies, DAOs, and guilds through real-world projects, open-source contributions, and peer-validated work.
- Public Learning Simulators: Open-source games and simulations funded as public goods, designed to teach and reward complex systems thinking, civic literacy, and scientific reasoning.
- Rapid Skill Adaptation Networks: Learning communities focused on helping people quickly acquire new competencies as technology and industries evolve, emphasizing meta-learning skills and adaptability over deep specialization in potentially obsolete fields.

What made this possible:

- The Skills Gap Crisis: The slow pace of university curriculum updates compared to the rapid evolution of the job market created a massive demand for more agile, modular credentialing systems.
- Privacy-Preserving Technology: The development of federated learning made personalized AI education possible without resorting to the surveillance-based models of early ed-tech.
- Recognition of Learning Diversity: Growing awareness of neurodiversity, learning disabilities, and different learning styles drove demand for personalized educational approaches that traditional one-size-fits-all systems couldn't accommodate effectively.
- A Cultural Shift to "Competence over Credentials": Employers began valuing verifiable skills and project portfolios alongside, and sometimes over, traditional degrees, legitimizing the decentralized layer.

Persistent Frictions:

- The Legitimacy Gap: A persistent cultural and institutional bias exists in favor of traditional university degrees, especially in licensed professions like medicine and law.
- Learning-by-Doing Limitations: While practical experience is valuable, some foundational theoretical knowledge still requires structured instruction, creating ongoing debates about the right balance between experiential and academic learning
- Equity of Access: While seemingly open, access to the best AI tutors, learning networks, and high-reputation mentors is unevenly distributed, threatening to worsen educational inequality.
- Epistemic Fragmentation: The risk that hyper-personalized, community-governed curricula could lead to educational silos that erode a shared foundation of civic and scientific knowledge.
- Adaptation Fatigue: The constant need to learn new skills as technology evolves can be

exhausting, particularly for older workers who built careers around deep expertise in now-obsolete fields.

4 The Human Experience in a d/acc World in 2035

By 2035, decentralized, democratic, defensive acceleration has made daily life less defined by a single role and more by a shifting set of activities. Most adults work a combination of jobs and projects, averaging 25–30 paid hours a week spread across different contexts. Income might come from a part-time role in a company, a contract with a design collective, and a share in a neighborhood enterprise. The arrangement makes people less tied to any one employer, but it also means learning to navigate different systems for payment, scheduling, and accountability.

Time feels more self-directed than a decade ago, though it is also more fragmented. A morning might be spent repairing equipment at a local workshop, the afternoon in paid remote work, and the evening reviewing proposals for a shared transport upgrade. Many communities use tools such as quadratic voting for local decisions, but participation is optional. Some people take an active role in these processes, others leave them to those with more time or interest.

Health care is organized around portability and choice. People keep their medical data in encrypted personal vaults, granting access when needed to a local clinic, an AI health advisor, or a regional hospital. This approach allows more privacy and flexibility, though switching between systems is not always smooth. Standards vary, and well-equipped community clinics are more common in some regions than in others.

Community life often revolves around shared infrastructure. Local energy grids, fabrication spaces, and food networks are practical services, but also regular meeting points. In many places, people are accustomed to using encrypted communication platforms run by the community and managing data collectively, limiting reliance on outside providers.

Identity is spread across multiple layers. A person might be rooted in one neighborhood, active in a regional science network, and a voter in a global cultural fund. These roles carry different weight at different times. The democratization of decision-making means people encounter more direct choices, on budgets, project priorities, or local rules, than a decade earlier. Deciding which to engage with, and which to leave to others, is part of the routine.

Life in 2035 is more connected, more privacy-conscious, and more participatory, though not everyone chooses to take part to the same extent. The systems that support it are more complex and more resilient, and living within them means learning to move between roles, contexts, and communities as part of everyday life.

5 Key Uncertainties and Tensions

5.1 What should we decentralize versus centralize?

The fundamental challenge of d/acc lies not in whether to pursue comprehensive decentralization, but in determining where decentralization enhances system resilience and where it merely introduces inefficiency without reducing risk. The critical principle is reducing dangerous dependencies while preserving efficiency. Sometimes this means decentralizing validation while keeping production centralized (as with Ethereum's block validation). But for goods with geopolitical significance, vaccines, semiconductors, and critical minerals—both production and validation may require geographic distribution to prevent weaponization of supply chains.

Core systems that concentrate power over individuals, including finance, identity, governance, and critical infrastructure, represent primary candidates for decentralization due to their vulnerability to abuse under singular control. However, decentralization should be pursued

strategically rather than ideologically, focusing on reducing single points of failure and creating competitive pressures for improved performance.

Decision-making structures benefit particularly from pluralistic approaches that avoid replicating existing power concentrations. Rather than token-based governance mechanisms that mirror wealth distributions, d/acc favors systems like quadratic voting that amplify broader constituencies while preserving the ability to express preference intensity.

Systems that should remain centralized include early-stage protocols lacking sufficient security for autonomous operation, coordination tasks that genuinely benefit from unified planning, and processes where distributed operation would create inefficiency without meaningful risk mitigation.

Domain-specific analysis reveals the nuanced nature of these decisions. Healthcare systems benefit from decentralized personal data ownership while requiring centralized safety standards. Educational systems gain from decentralized credentialing and peer learning networks while depending on centralized infrastructure provision. Scientific research improves through decentralized funding and publishing mechanisms while requiring centralized ethical oversight for high-risk activities. Legal systems can decentralize dispute resolution while maintaining centralized constitutional frameworks. Environmental improvement succeeds through decentralized innovation and local energy systems while requiring centralized externality-limiting laws.

However, these boundaries are not static. As systems mature and threat landscapes evolve, the optimal balance between centralization and decentralization shifts accordingly. Systems that require centralized coordination during early development may become suitable for distributed management as they stabilize. Conversely, experimental decentralized systems may require temporary re-centralization during crisis periods. The challenge lies in building adaptive systems that can adjust these boundaries dynamically while maintaining institutional coherence.

5.2 Can decentralized systems coordinate effectively when speed and scale matter?

This question represents the most significant practical challenge facing d/acc implementation. The ability of distributed actors to coordinate rapidly on shared threats will determine whether decentralized approaches can handle civilizational-scale challenges or whether crisis conditions inevitably drive re-centralization.

The complexity problem. When multiple overlapping governance systems operate simultaneously across different domains, participants face significant cognitive overhead in determining applicable rules, relevant authorities, and appropriate action pathways. The risk extends beyond mere confusion to include decision paralysis, contradictory actions, and coordination failures that undermine collective response capabilities.

Proposed solutions include AI-assisted coordination tools that reduce navigational complexity without requiring participants to understand every system component. However, this approach raises fundamental questions about whether managing decentralized complexity through AI systems simply creates new forms of centralization at the protocol layer.

Offense-defense asymmetries compound coordination challenges. Destructive capabilities often advance faster than defensive countermeasures, particularly as AI and biotechnology tools become increasingly accessible. While attribution-based control systems, red-teaming DAOs, and distributed sensor networks offer promising defensive approaches, they face significant scaling challenges and implementation complexities.

The "small-kills-all" problem presents particular difficulties. Individual actors with access to advanced AI or bioengineering capabilities might inflict catastrophic damage before any distributed defense system can respond effectively. The optimistic scenario assumes that

open-source defensive tools will scale faster than offensive capabilities, creating persistent defensive advantages. The pessimistic scenario suggests that democratized access to dangerous technologies makes catastrophic misuse statistically inevitable across global populations.

Polycentric governance introduces additional coordination friction. When multiple jurisdictions operate with overlapping authority and potentially conflicting rulesets, determining final decision-making authority becomes problematic. Situations where local governance innovations conflict with global coordination requirements, such as pandemic containment or AI safety protocols, highlight the tension between local autonomy and collective action needs.

5.3 Will decentralized systems achieve sufficient adoption and trust?

The social and psychological dimensions of d/acc systems are as critical as their technical architectures. The question of whether individuals and institutions will trust and adopt decentralized AI systems, governance models, and funding protocols, particularly during failure scenarios, remains open. Centralized systems benefit from clear accountability structures and familiar institutional forms that provide cognitive shortcuts for user interaction and problem resolution.

Competitive dynamics persist even within well-designed d/acc ecosystems. Individual actors continue to face incentives for free-riding behavior, short-term optimization, and defection from cooperative arrangements when such strategies provide individual benefits. Public goods funding mechanisms, including retroactive grants and quadratic funding, remain vulnerable to manipulation by sophisticated actors capable of gaming reputation systems or capturing governance processes through patient capital deployment.

The sustainability of d/acc systems depends on developing self-correcting mechanisms that limit such capture without requiring constant participant vigilance. This challenge is particularly acute in rapidly evolving technological environments where institutional rules and norms must adapt continuously to new capabilities and threat vectors.

Epistemological fragmentation represents a significant risk. Pluralistic trust networks and decentralized validation mechanisms may fragment rather than enrich shared understanding of complex issues. If different communities develop incompatible worldviews based on distinct information sources and validation processes, the coordination necessary for collective action becomes increasingly difficult.

The emotional and symbolic dimensions of institutional trust also matter significantly. Centralized institutions provide narrative coherence and symbolic authority that help individuals navigate complex social arrangements. Federated systems, while more fault-tolerant, may lack the emotional resonance and clear symbolic representation that make political engagement meaningful for broad populations.

Protocol-layer governance presents the deepest challenge. Decentralization efforts typically require shared protocols that can become new centralization points. If attribution-based control mechanisms become standard for AI interaction, the entities governing protocol development wield enormous influence over entire ecosystems. Protocol capture—where single actors gain control over the fundamental "languages" of system interaction—poses risks of rent extraction and behavioral manipulation that benefit narrow rather than broad interests.

Building trust in decentralized systems requires demonstrating reliable performance under stress conditions, as early failures can create lasting skepticism that undermines long-term adoption. This requirement is particularly challenging during transition periods when new systems remain experimental while legacy alternatives provide familiar fallback options.

Core Unanswered Questions

These tensions point to fundamental questions that will determine the viability of d/acc approaches:

Scalability and Complexity Management: At what scale do the coordination costs of decentralized systems exceed their resilience benefits? Can human cognitive limitations be overcome through technological assistance without recreating centralized control structures?

Institutional Adaptation: How quickly can decentralized systems adapt their centralization-decentralization boundaries in response to changing conditions? What mechanisms prevent such adaptations from being captured by actors seeking to consolidate power?

Trust and Legitimacy: What are the minimum conditions for democratic legitimacy in polycentric governance systems? How can decentralized institutions maintain public trust without the symbolic authority and narrative coherence of traditional centralized structures?

Crisis Performance: Under what conditions do decentralized systems maintain coordination effectiveness during high-stakes emergencies? When do they fail catastrophically, and how can such failures be prevented or mitigated?

Protocol Governance: How can foundational protocols be governed democratically while maintaining technical coherence and development momentum? What prevents protocol capture by technically sophisticated actors?

Economic Sustainability: Can public goods funding mechanisms scale to support civilization-level infrastructure without being gamed or captured? What are the long-term economic dynamics of systems designed around cooperation rather than competition?

These questions do not have straightforward answers, reflecting the experimental nature of d/acc as both a technological and institutional research program. The ultimate test will be whether decentralized approaches can deliver superior outcomes across multiple domains while maintaining democratic governance and individual agency, or whether they represent a temporary phase in technological development that eventually converges toward more centralized forms.

6 Appendices

How This Scenario Was Created

This scenario is part of the AI Pathways project, an initiative of the Foresight Institute's Existential Hope program. Rather than focusing on risks or speculative timelines, AI Pathways presents two vividly realized and contrasting visions of what a desirable AI-driven future might look like.

The project brings together leading thinkers, including Vitalik Buterin, Glen Weyl, Anton Korinek, and Allison Duettmann, who contributed to crafting these narratives.

Explore the AI Pathways Project

- d/acc 2035 - Imagines a decentralized, democratic, defensive model of technological progress. Emphasizes plural acceleration, privacy-first infrastructure, and community-governed resilience.
- Tool AI 2035 - A future shaped by advanced, but purposefully controllable, often narrow in scope, AI systems that enhance human decision-making without striving for full autonomy or generality.

Both scenarios are part of AI Pathways (https://ai-pathways.existentialhope.com/) and

invite reflection on the values and paths we choose in shaping AI's role in our world.

Contributors to this scenario

The development of the d/acc report was led by Linda Petrini and Beatrice Erkers, grounded in expert interviews and iterative feedback from across domains. The scenario should not be seen as the official views of any individual contributor.

1. Vitalik Buterin (Ethereum),
2. Glen Weyl (Microsoft Research, RadicalXchange),
3. Kevin Owocki (Gitcoin),
4. Andrew Trask (OpenMined, DeepMind),
5. Emilia Javorsky (Future of Life Institute),
6. Deger Turan (Metaculus),
7. Allison Duettmann (Foresight Institute),
8. Soham Sankaran (PopVax),
9. Christine Peterson (Foresight Institute),
10. Marcin Jakubowski (Open Source Ecology),
11. Naomi Brockwell (Ludlow Institute),
12. Molly Mackinlay (Protocol Labs),
13. Lou de Kerhuelvez (Nodes).

We're deeply grateful to anyone who contributed their time and insights to this experiment.

Metaculus Forecasting integration

To engage a broader audience, we've launched a set of forecasting questions on Metaculus tied to key scenario milestones, along with a $5,000 Commenting Prize (https://www.metaculus.com/tournament/foresight-ai-pathways/) for the top contributors.

Context: Existential Hope Program

The term "Existential Hope" refers to the capacity to envision futures where humanity not only survives, but flourishes in ways we can currently only imagine, and to be able to better work towards those futures. It complements the better-known concept of existential risk.

This project sits within the Foresight Institute's broader mission to drive long-term, future-positive technology, where imagination is both a tool and a catalyst for change.

7 Master Reference List

Core Scenario References

1. Critch, A., Dennis, M., & Russell, S. (2022). Cooperative and Uncooperative Institution Designs: Surprises and Problems in Open-Source Game Theory. arXiv:2208.07006. https://arxiv.org/abs/2208.07006
2. Duettmann, A. (2025). Multipolar AI is Underrated. LessWrong. https://www.lesswrong.com/posts/JjYu75q3hEMBgtvr8/multipolar-ai-is-underrated
3. Critch, A. (2024). My Motivation and Theory of Change for Working in AI Healthtech. Alignment Forum. https://www.alignmentforum.org/posts/Kobbt3nQgv3yn29pr/

my-motivation-and-theory-ofchange-for-working-in-ai

4. Trask, A., Bluemke, E., Collins, T., Garfinkel, B., Drexler, K. E., et al. (2020). Beyond Privacy Trade-offs with Structured Transparency. arXiv:2012.08347. https://arxiv.org/abs/2012.08347

Domain-Specific References

Governance

5. Bell, T. W. (n.d.). Ulex: An Open Source Legal System. GitHub repository. https://github.com/proftomwbell/Ulex
6. vTaiwan case study. (n.d.). CrowdLaw for Congress: vTaiwan—Process & Lessons. [Case Study]. https://congress.crowd.law/files/vtaiwan-case-study.pdf

Science

7. Viganola, D., et al. (2021). Using Prediction Markets to Predict the Outcomes in Social Science. PNAS/Open. https://pmc.ncbi.nlm.nih.gov/articles/PMC8278038/

Healthcare

8. Rieke, N., et al. (2020). The Future of Digital Health with Federated Learning. NPJ Digital Medicine. https://www.nature.com/articles/s41746-020-00323-1
9. Critch, A. (2024). My Motivation and Theory of Change for Working in AI Healthtech. Alignment Forum. https://www.alignmentforum.org/posts/Kobbt3nQgv3yn29pr/my-motivation-and-theory-ofchange-for-working-in-ai

Law & Justice

10. Kleros case—Mexico. (2022). Blockchain Arbitration: Roadmap to Recognition and Enforcement (Mexico case). Wolters Kluwer Arbitration Blog. https://legalblogs.wolterskluwer.com
11. Jus Mundi Dossier. (2020–2021). Mexican Company X v. Mexican Company Y (Kleros referral & enforcement docs). https://jusmundi.com/en

Infrastructure / Climate & Energy

12. Raman, G., et al. (2024). The Social Factors Shaping Community Microgrid Operation. Nature Communications. https://www.nature.com/articles/s41467-024-50736-9
13. U.S. DOE/NREL (2025). MIRACL: Microgrids, Infrastructure Resilience, and Advanced Controls Launchpad—Final Report. https://www.nrel.gov/wind/miracl-report

Finance

14. Weyl, E. G., & Zhang, Y. (2021). Quadratic Funding: A Primer. [Working Paper].

Explainers

d/acc

d/acc stands for decentralized, democratic, differential, defensive acceleration. The term builds on "accelerationist" ideas but proposes a distinct direction: accelerating technologies that promote resilience, distribution, and pluralism, rather than centralization or control. The concept was outlined by Vitalik Buterin in his 2023 essay My Techno-Optimism, as an alternative to both slow-down narratives and uncritical techno-acceleration.

Federated Research Networks

Federated research networks are decentralized systems of collaboration where data, compute, or models are shared across institutions without centralizing control. Examples include federated

learning systems in AI, or distributed scientific consortia. These networks aim to preserve privacy, local autonomy, and robustness while enabling large-scale coordination.

AI Fiduciaries

AI fiduciaries are entities, human or institutional, responsible for ensuring that AI systems act in the best interest of a designated individual or group. The term borrows from legal fiduciary duties in finance or law, and has been proposed as a governance framework for aligning AI behavior with user values, rights, and wellbeing.

Forkable Codebases

This concept draws an analogy between software version control and institutional design. A "forkable" governance framework is one where laws, contracts, or protocols can be cloned, modified, and redeployed, just like open-source codebases. This enables institutional experimentation, adaptability, and transparency across jurisdictions or communities.

Jurisdictional Routers

Jurisdictional routers are conceptual or technical tools that allow individuals or organizations to choose between different legal, governance, or digital frameworks, often within networked or decentralized systems. The idea reflects growing interest in "choice of governance" in digital societies and is influenced by network states, e-residency, and virtual jurisdictions.

DAO (Decentralized Autonomous Organization)

A DAO is a governance structure run by smart contracts on a blockchain, allowing decentralized coordination among participants. Members typically vote using tokens, and rules are enforced transparently via code. DAOs are used for collective asset management, project funding, and protocol governance without centralized intermediaries.

"TradFi-DeFi"

This hybrid term refers to the emerging intersection between Traditional Finance (TradFi) institutions, such as banks or asset managers—and Decentralized Finance (DeFi) protocols built on blockchain. It includes initiatives like tokenized assets, compliance-layer DeFi, and institutional experimentation with on-chain infrastructure.

Jurisdictional Routers

Systems, technical, legal, or institutional, that help individuals or organizations choose between overlapping regulatory or governance frameworks. These tools may allow users to route activity through preferred jurisdictions (e.g. for data, contracts, or research), increasing flexibility, compliance, or autonomy. The concept draws from ideas in network governance, digital sovereignty, and programmable law.

RetroPGF (Retroactive Public Goods Funding)

A funding model where contributors to public goods are rewarded after the value of their work has been demonstrated. Popularized by Optimism, RetroPGF aims to create sustainable incentive structures for open-source software, research, and community infrastructure by decoupling funding from upfront grantmaking.

Public Goods Funding

Mechanisms that support the creation and maintenance of resources that benefit everyone, such as open data, infrastructure, or basic research. These include philanthropic grants, state subsidies, quadratic funding, and newer blockchain-based approaches that aim to improve efficiency, scalability, or alignment with community values.

Proof-of-Personhood

Protocols or mechanisms designed to verify that a digital identity corresponds to a unique

human being, without necessarily revealing their real-world identity. Used to prevent Sybil attacks in decentralized systems and to enable one-person-one-vote governance or fair resource allocation. Approaches range from biometric methods to social graph attestations.

Gitcoin Grants

A decentralized funding platform for open-source and public goods projects. It uses quadratic funding to match community donations with pooled resources, amplifying collective decision-making. Gitcoin Grants has supported thousands of projects across web3, climate tech, education, and more.

Federated City-States

A speculative governance model in which multiple autonomous cities form a federation, pooling resources and interoperable governance systems while retaining local sovereignty. Unlike traditional nation-states, these federations may be organized around shared infrastructure, trade, or digital governance protocols rather than geography alone.

Safety-Netted DAO

A decentralized autonomous organization (DAO) that includes built-in protective mechanisms to prevent catastrophic failure or misuse. Safety nets can include pause functions, emergency withdrawal rights, multi-signature controls, or delegated "circuit breaker" authorities to respond to crises.

"Circuit Breaker" Council

A governance body, often within a DAO or decentralized protocol, empowered to temporarily halt operations, pause smart contracts, or block certain transactions in the event of detected exploits, governance attacks, or systemic risks. Designed to provide rapid response without undermining long-term decentralization.

"Pay for Health Outcomes"

A funding and incentive model proposed by economist Robin Hanson in which healthcare providers or researchers are compensated based on actual improvements in patient outcomes rather than services rendered. Aims to align incentives toward effectiveness and preventive care.

Homomorphic Encryption

A cryptographic technique that allows computation to be performed directly on encrypted data without decrypting it. Useful for privacy-preserving analytics, enabling data sharing and processing across untrusted environments without exposing sensitive information.

Secure Multiparty Computation (MPC)

A cryptographic method that allows multiple parties to jointly compute a function over their inputs while keeping those inputs private. Commonly used for collaborative analytics, private auctions, and shared control over sensitive data or assets.

Sybil Attacks

An exploit in decentralized networks where an attacker creates many fake identities to gain disproportionate influence or access to resources. Countermeasures include proof-of-personhood, stake requirements, and social trust graph analysis. Self-Sovereign Identity (SSI) A digital identity model where individuals control their own credentials and decide what information to share, without reliance on centralized identity providers. Often built on blockchain or decentralized identifiers (DIDs), enabling portability and privacy.

Web-of-Trust

A decentralized trust model where identity, reputation, and permissions flow through peer

attestations instead of a central authority. Participants issue cryptographic or verifiable credentials for others, enabling access control, Sybil resistance, and reputation across networks. Trade-offs include collusion and privacy leakage; common mitigations are rate limits/decay, stake or slashing, and selective-disclosure proofs.

Ethereum's Block Validation

The process by which Ethereum nodes verify that a block of transactions is valid according to the network's consensus rules. Includes checking transaction signatures, smart contract execution results, and compliance with the current protocol rules before adding the block to the blockchain.

Quadratic Funding (QF)

A funding mechanism designed to allocate resources to public goods based on the number of contributors and the size of their contributions. Small contributions from many people are weighted more heavily than large contributions from a few, allowing communities to signal which projects they value most. Matching funds from a central pool amplify community preferences, making QF effective for supporting open-source software, local initiatives, and other collective benefit projects.