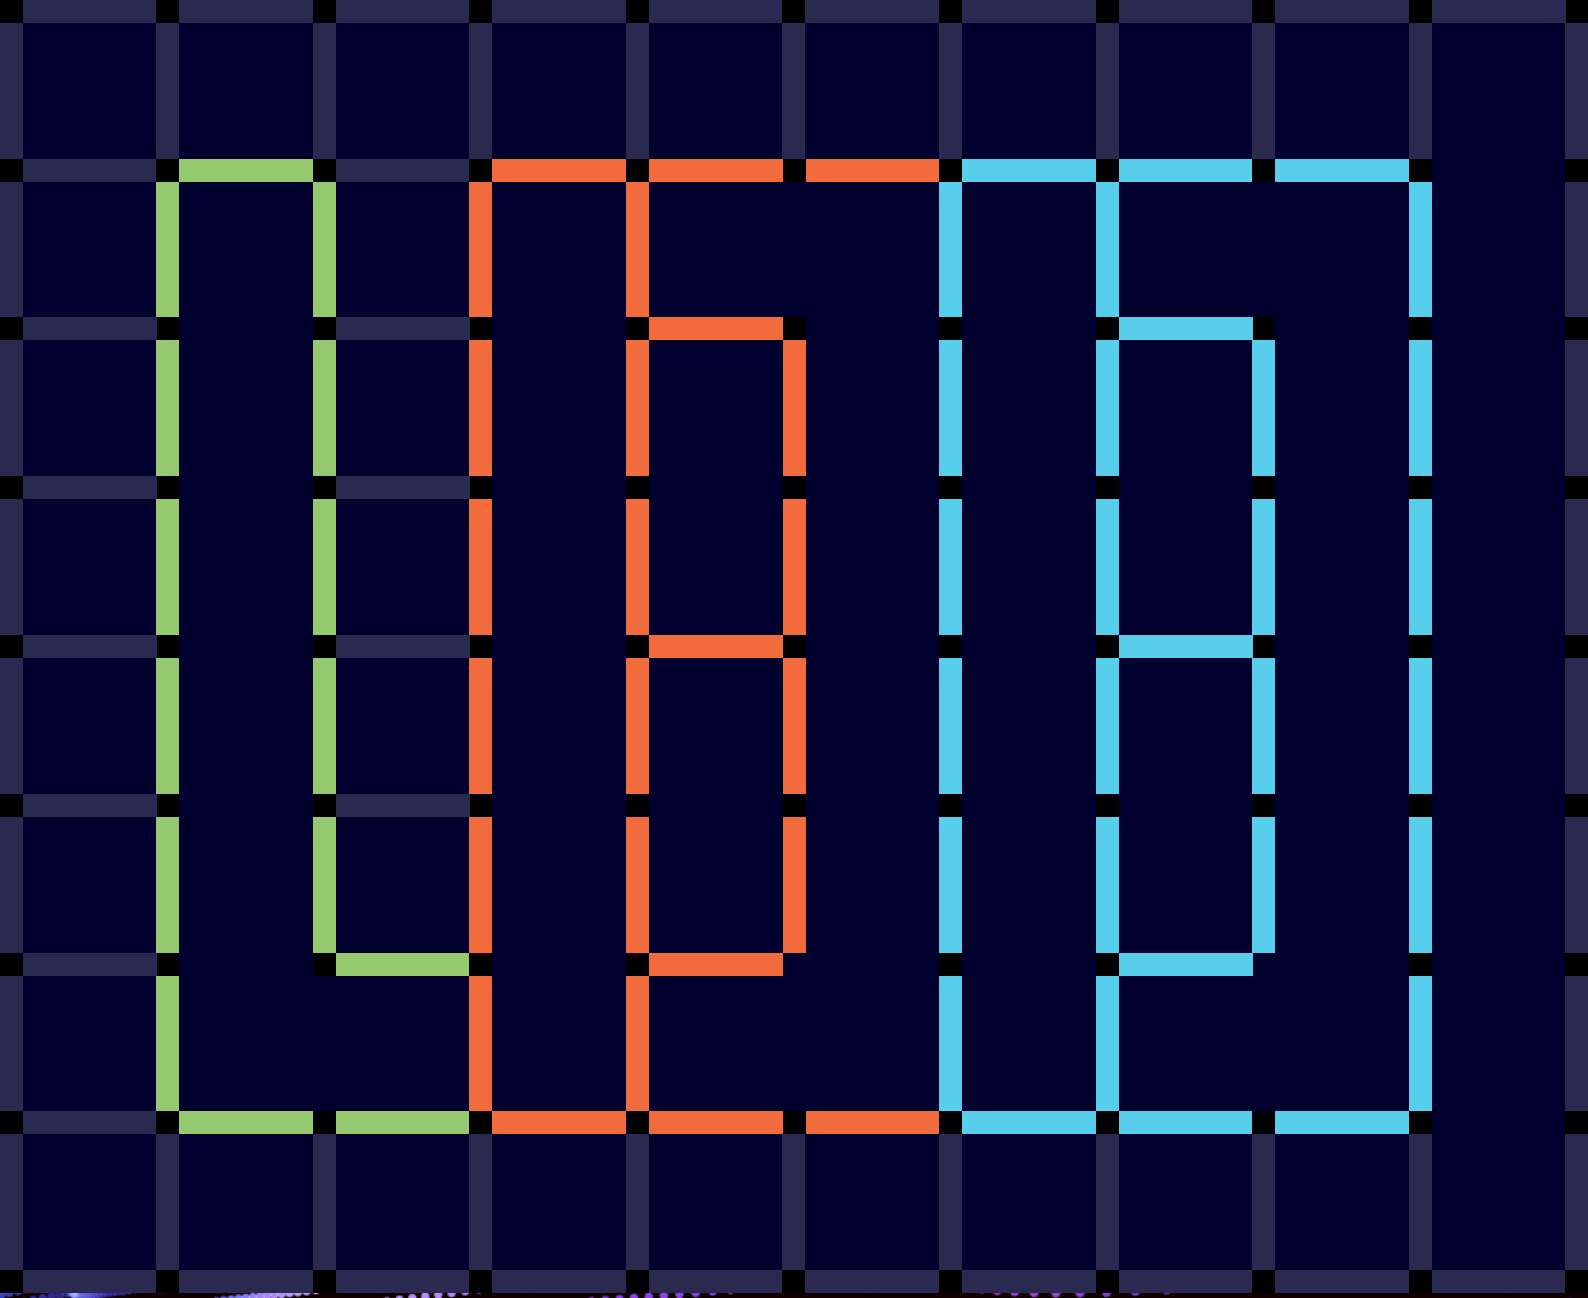

LEADING BEYOND BOUNDARIES

Davos, Switzerland

Held during the World Economic Forum

2025

Executive Summary

The Leading Beyond Boundaries (LBB) Summit was held during the World Economic Forum's Annual Meeting 2025 in Davos, Switzerland. LBB aims to transcend disciplinary disconnects that often impede the development of effective technology policy, research, and innovation. To achieve this, LBB convened four expert roundtables focused on critical areas shaping our future: Agentic AI, Decentralized AI, Personalized Health, and Legal Practice in the age of Generative AI. The conversations were convened by groups at **MIT**, **Harvard University** and **Stanford University** under Chatham House Rules, allowing participants to speak freely. This report synthesizes the key insights and recommendations shared by leading experts in these sessions.

The meetings emphasized that artificial intelligence is moving from experimentation toward broad deployment. Nearly eight in ten companies now use generative AI, yet may report limited material benefit. Participants argued that the next phase of AI development will require autonomous agents capable of planning, memory and tool use; however, such systems raise questions about stability, liability and trust.

Concerns about data privacy and security motivate decentralized architectures. In healthcare, panelists highlighted the promise of multimodal data and personalized medicine, while recognizing that most determinants of health lie outside of the clinic — shaped by lifestyle, environment and genetics. In law, generative AI could democratize legal services and streamline routine work, but access to justice and ethical use remain pressing issues. Across all sessions, speakers emphasized that AI should augment human judgment, not replace it, and that governance, incentives and education will determine whether AI improves or erodes societal well-being.

Key themes common to all four conversations include:

Trust and accountability

Without reliable safeguards, participants warned that autonomous agents could destabilize markets or manipulate individuals. Building trust metrics, transparent audit trails and liability frameworks were recurring suggestions.

Data stewardship

Access to high quality data underpins progress in health and law, yet privacy concerns and misaligned incentives continue to inhibit sharing. Federated and edge-based architectures that keep data local while enabling learning were highlighted as a promising way forward.

Human AI partnership

Speakers consistently advocated for "augmented intelligence", where AI supports professionals rather than replaces them. This principle applies equally to clinicians reviewing diagnostic suggestions, legal professionals drafting documents and regulators evaluating AI systems.

Equitable access

AI's benefits risk being captured by well-resourced actors. Participants urged that policy and industry investments ensure that small businesses, emerging economies and under-served communities benefit from legal AI tools, decentralized health solutions and open agent platforms.

The following chapters provide concise summaries of each domain specific discussion, along with recommendations.

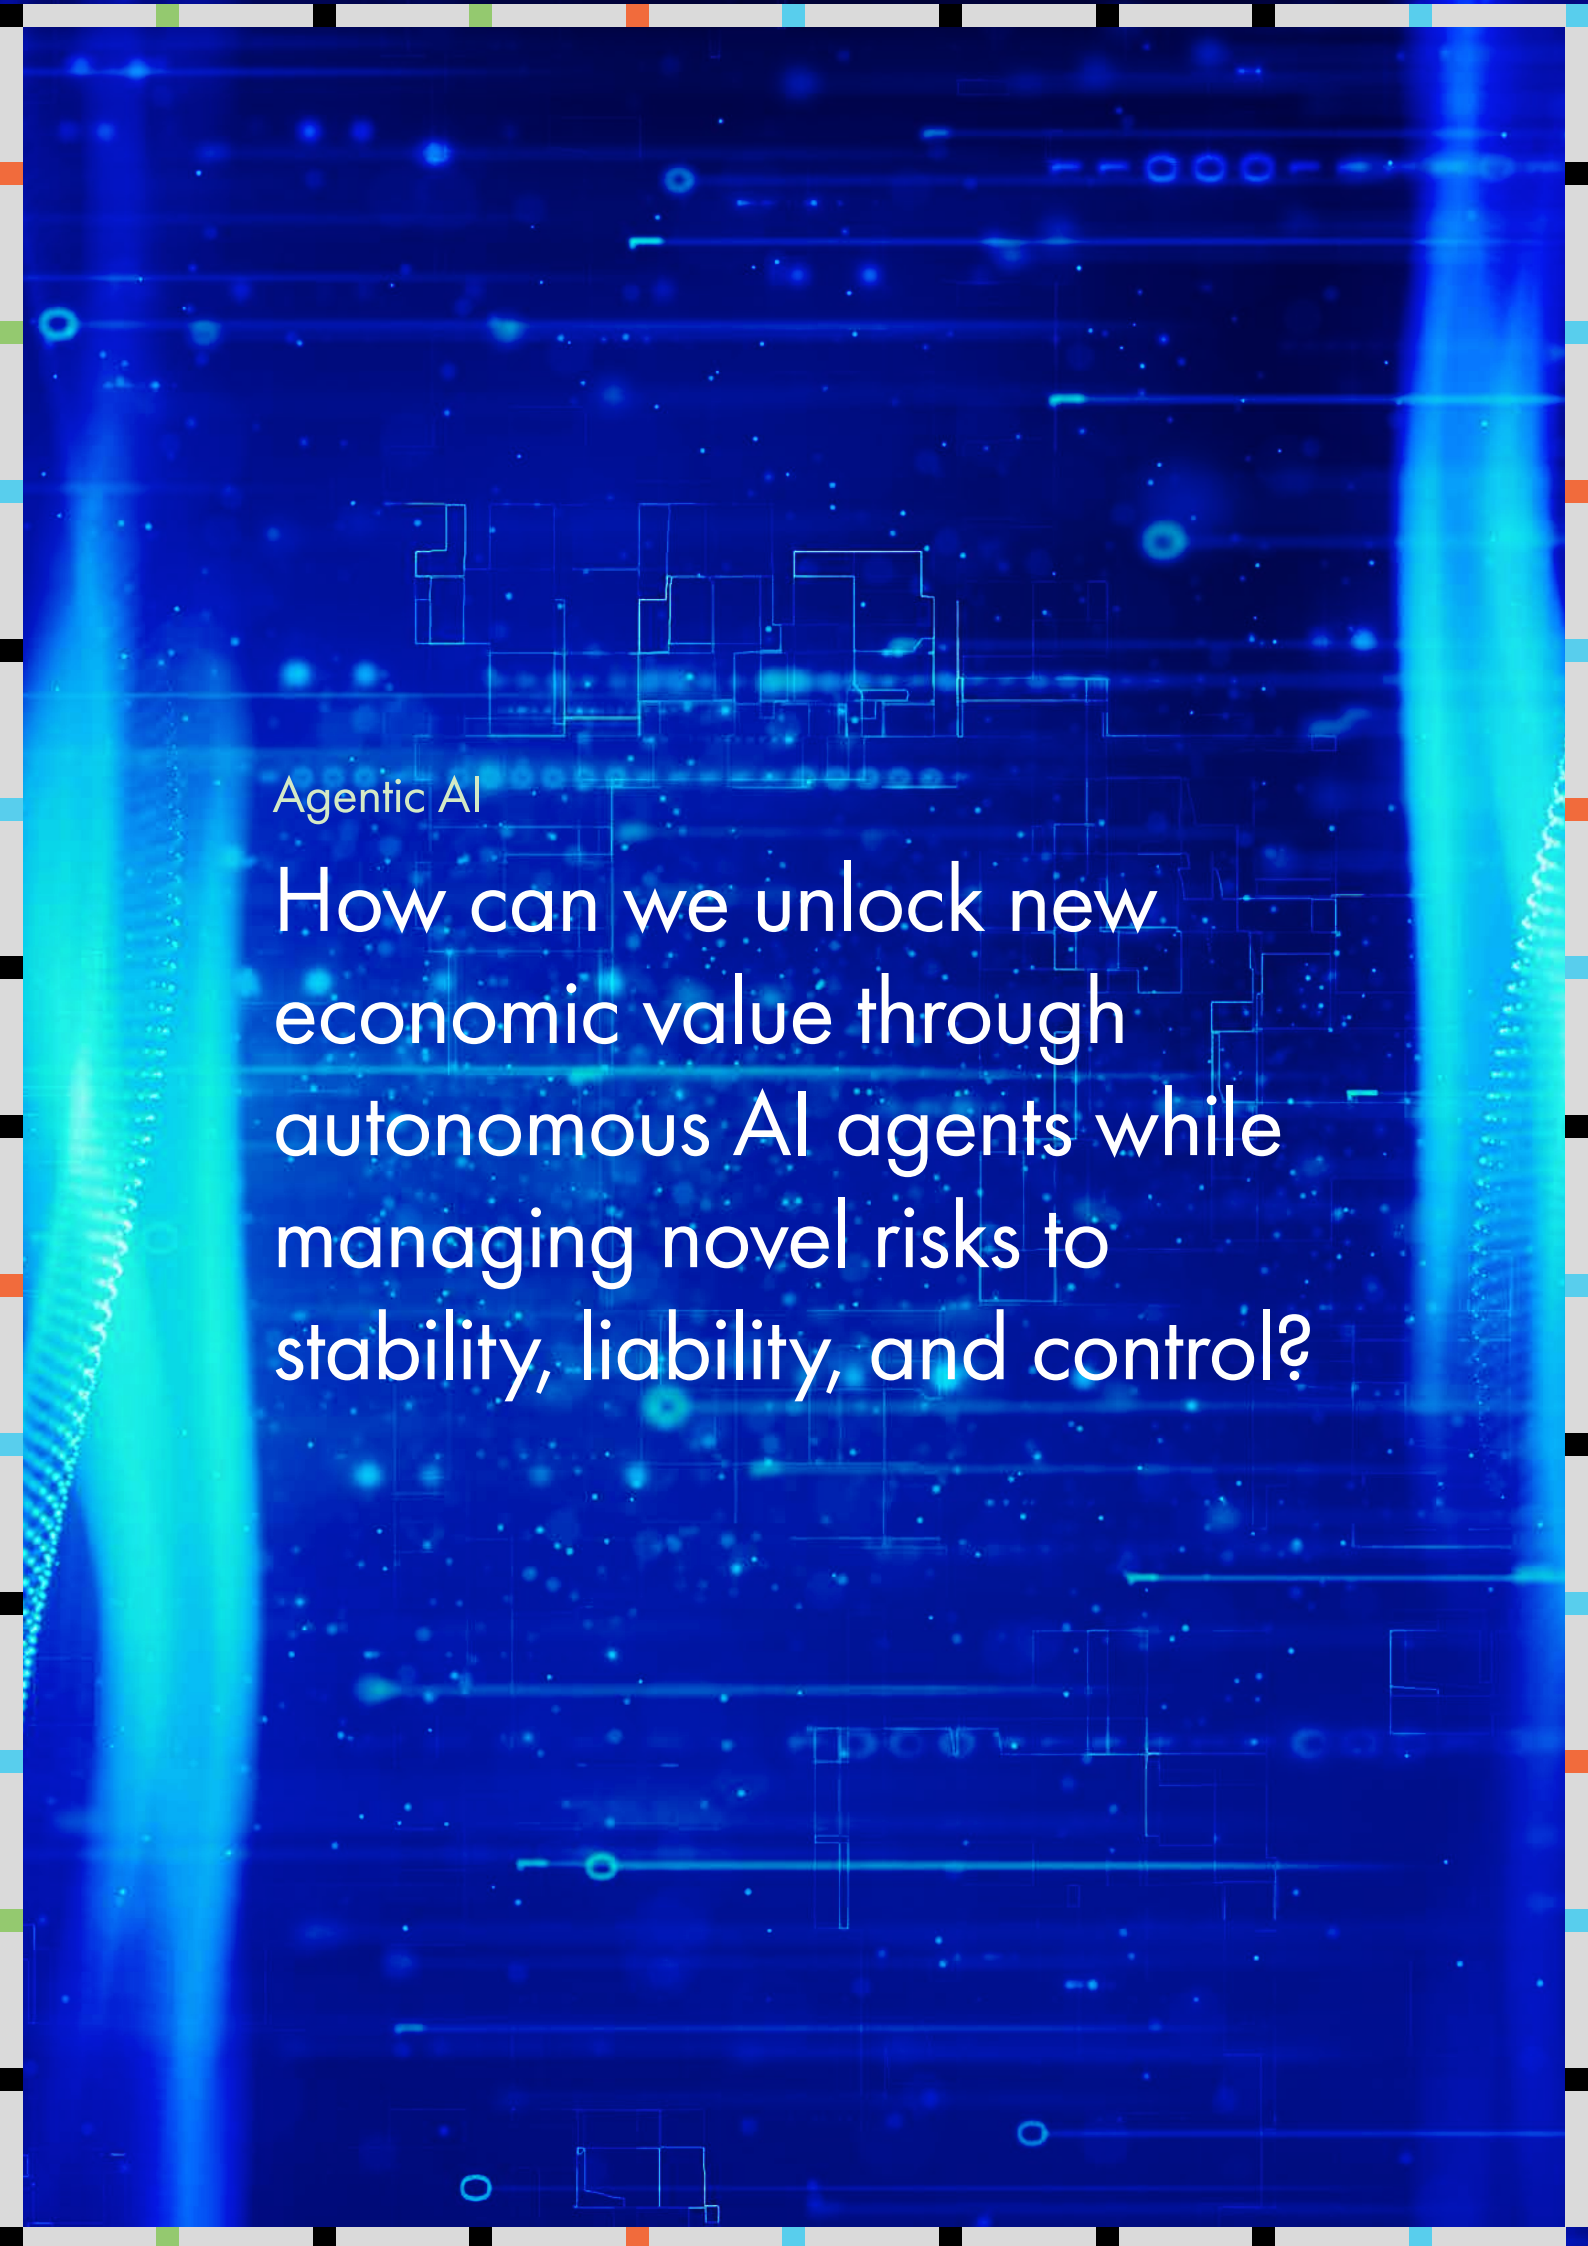
Agentic AI
How can we unlock new economic value through autonomous AI agents while managing novel risks to stability, liability, and control?

Agentic AI, Autonomy, Safety and Impact

Agentic AI refers to systems capable of planning tasks, retaining context, using external tools and collaborating with other agents to pursue goals on a user's behalf. Such "collaboration" may involve both autonomous decision-making and coordinated action under human supervision. A McKinsey report notes that adopting these agents could break the "Gen AI paradox": although eight out of ten companies use generative AI, most derive little value because current systems function largely as passive copilots. Agentic systems promise to automate complex workflows and unlock new revenue opportunities. Roundtable participants welcomed this potential but worried about stability, reliability and incentive alignment.

Introduction

The discussion opened with definitions. While traditional chatbots simply generate responses to prompts, **agentic AI** orchestrates tools and data sources — such as software APIs, cloud platforms, web search, email, or flight-booking applications — to complete multi-step tasks, often without continuous human supervision. Access to these external systems, typically through APIs or emerging protocols like the Model Context Protocol (MCP), allows agents to act beyond their training data and function as semi-autonomous collaborators. Participants noted that agentic systems already exist in domains such as high-frequency trading and algorithmic optimization, yet their occasional instability, such as market crashes and sudden meltdowns, demonstrates the stakes of agentic autonomy in complex environments. This prompted several questions: **How do we know an agent will not collapse the system?** Where does liability reside when an agent makes an erroneous trade or recommendation? How do we reset an agent that has gone 'rogue' — whether through technical malfunction or misaligned incentives? Some suggested "using AI to watch AI", deploying meta-agents that monitor other agents for anomalies and enforce guardrails.

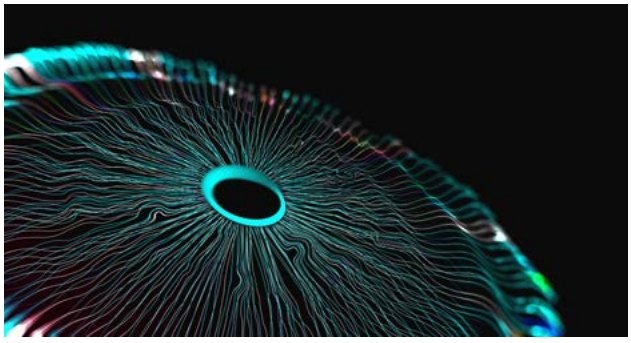

Much of the debate centred on incentives. Participants emphasized the need to create structures that align both human and algorithmic incentives, preventing situations where short-term profit-seeking undermines system stability. Others proposed rule extraction techniques aimed at translating the implicit reasoning of large language models into explicit, auditable policy rules. Quantum computing surfaced briefly as a potential tool for exploring large state spaces, though participants cautioned that current quantum systems are overhyped and remain far from practical deployment.

Reliability of tool calls and memory was another concern. One speaker noted that only about 70% of tool calls succeed, meaning one in three user requests fails. Hard-coded rules can reduce risk but also limit an agent's flexibility; conversely, unfettered reasoning can lead to unpredictable outcomes. To address this, the group proposed a two-layer approach: a model layer that generates actions or recommendations, paired with a policy layer that reviews these outputs to ensure they comply with legal and ethical standards.

Agentic AI, Autonomy, Safety and Impact

Actionable Recommendations

The experts agreed on key areas where progress is essential:

RECOMMENDATION	RATIONALE	STAKEHOLDERS
Define agentic standards and taxonomies	Establish a shared vocabulary distinguishing simple chat assistants from fully autonomous agents and specify required capabilities (planning, memory, tool use). Standardized definitions will support regulators, developers and buyers.	Standards bodies, research institutions, regulators
Develop stability and rollback mechanisms	Design agents with built-in **governance agents** that monitor behavior, detect anomalous actions and trigger safe shutdowns or state rollbacks. Include manual override options for human operators.	AI developers, system integrators
Align incentives and liability frameworks	Collaborate with economists and legal experts to design incentive structures that discourage catastrophic behavior and clearly define responsibility for agent actions, including when liability rests with platform providers versus users.	Regulators, insurers, corporate counsel
Improve tool reliability and transparency	Invest in research to increase the success rate of tool calls and maintain transparent logs for each call. Provide users with configurable settings that balance determinism and creative exploration.	AI vendors, open source communities
Educate practitioners in ethics and policy	Offer targeted training for technologists and policymakers on the legal and ethical dimensions of agentic AI, including fairness, accountability, non-discrimination and harm minimization.	Universities, professional bodies

Decentralized AI

How can distributed AI systems enhance privacy and resilience while ensuring trust, security, and value distribution?

Decentralized AI

Decentralized AI involves multiple autonomous agents operating across distributed devices and networks without a single trusted coordinator. Edge computing keeps data on local devices, reducing exposure to network attacks and improving privacy. Yet limited computational resources, the risk of physical tampering and emerging attack vectors – such as model inversion – create new security challenges. During the discussion, participants contrasted decentralized agents with cloud-based Software as a Service (SaaS) applications and debated how to quantify trust and design economic models for an **agent economy**.

Introduction

Speakers brainstormed themes such as "attribution", "decentralization", "agents versus humans", "value accrual" and "mass adoption". A lively exchange followed on whether agentic AI represents a mere incremental evolution of SaaS or a more fundamental shift. Some participants argued that agents resemble service-oriented architecture, while others insisted that natural language interfaces, capacity to learn and operational scale make them qualitatively different. One speaker likened the current excitement around agentic AI to the early days of computing: **"SaaS is like a musket that took eleven minutes to load; agents are like a modern rifle."**

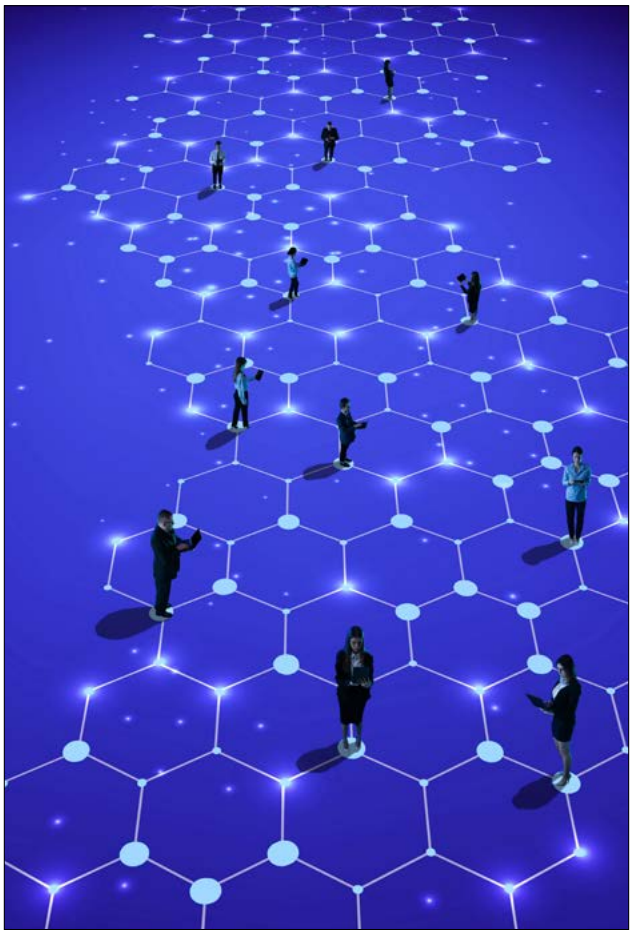

Trust emerged as the core concern. In decentralized systems, agents may interact with other autonomous agents whose origins, reliability, or intentions are uncertain — some of which could behave unpredictably or even maliciously. This raises difficult questions about accountability and safety: Who accrues value from their actions? How can we ensure that they act in users' interests? The group explored ways to measure trust, noting that reciprocal exchanges (people talking to each other or exchanging value) tend to strengthen it. For example, one speaker mentioned Taiwan's digital platform "Polis", which reportedly boosted trust in the government by enabling transparent, two-way engagement between citizens and officials.

Participants proposed a "Waze for agents" where users crowdsource information about the reliability of agents and flag malicious actors. They debated how to incentivize honest reporting, suggesting mechanisms such as using blockchain based tokens to reward trustworthy contributions. The group also discussed a transition from SaaS to agentic systems, noting that early agent prototypes may be easy to build but difficult to scale effectively. Cross agent orchestration and context awareness were identified as being essential for sustainable adoption.

Decentralized AI

Actionable Recommendations

The experts agreed on key areas where progress is essential:

RECOMMENDATION	RATIONALE	STAKEHOLDERS
Quantify trust and reputational metrics	Develop metrics to measure agent reliability based on successful interactions and peer feedback. Transparent ratings will allow users to choose trustworthy agents and discourage malicious behavior.	AI platforms, decentralized app developers
Research secure multi-party protocols	Invest in cryptographic methods (homomorphic encryption, secure enclaves) that allow agents to collaborate on private data without exposing raw inputs. These techniques can support federated learning and trust based digital economies.	Universities, cryptography researchers
Design incentive-aligned token models	Use digital tokens to reward honest reporting of agent behavior and penalize malicious actors. Explore governance mechanisms that allow communities to set policies and share in network value.	Economists, blockchain developers
Encourage regulatory guidance on decentralized AI	Policymakers should clarify liability and consumer protection standards for decentralized AI systems. Clear guidance on cross-jurisdictional data flows and recognition of decentralized identifiers will foster responsible innovation.	Regulatory bodies, legal scholars
Foster interoperable agent frameworks	Adopt open protocols that allow agents to share context and delegate tasks across networks. Interoperability will prevent platform lock in and encourage healthy competition.	Standards organizations, open source communities

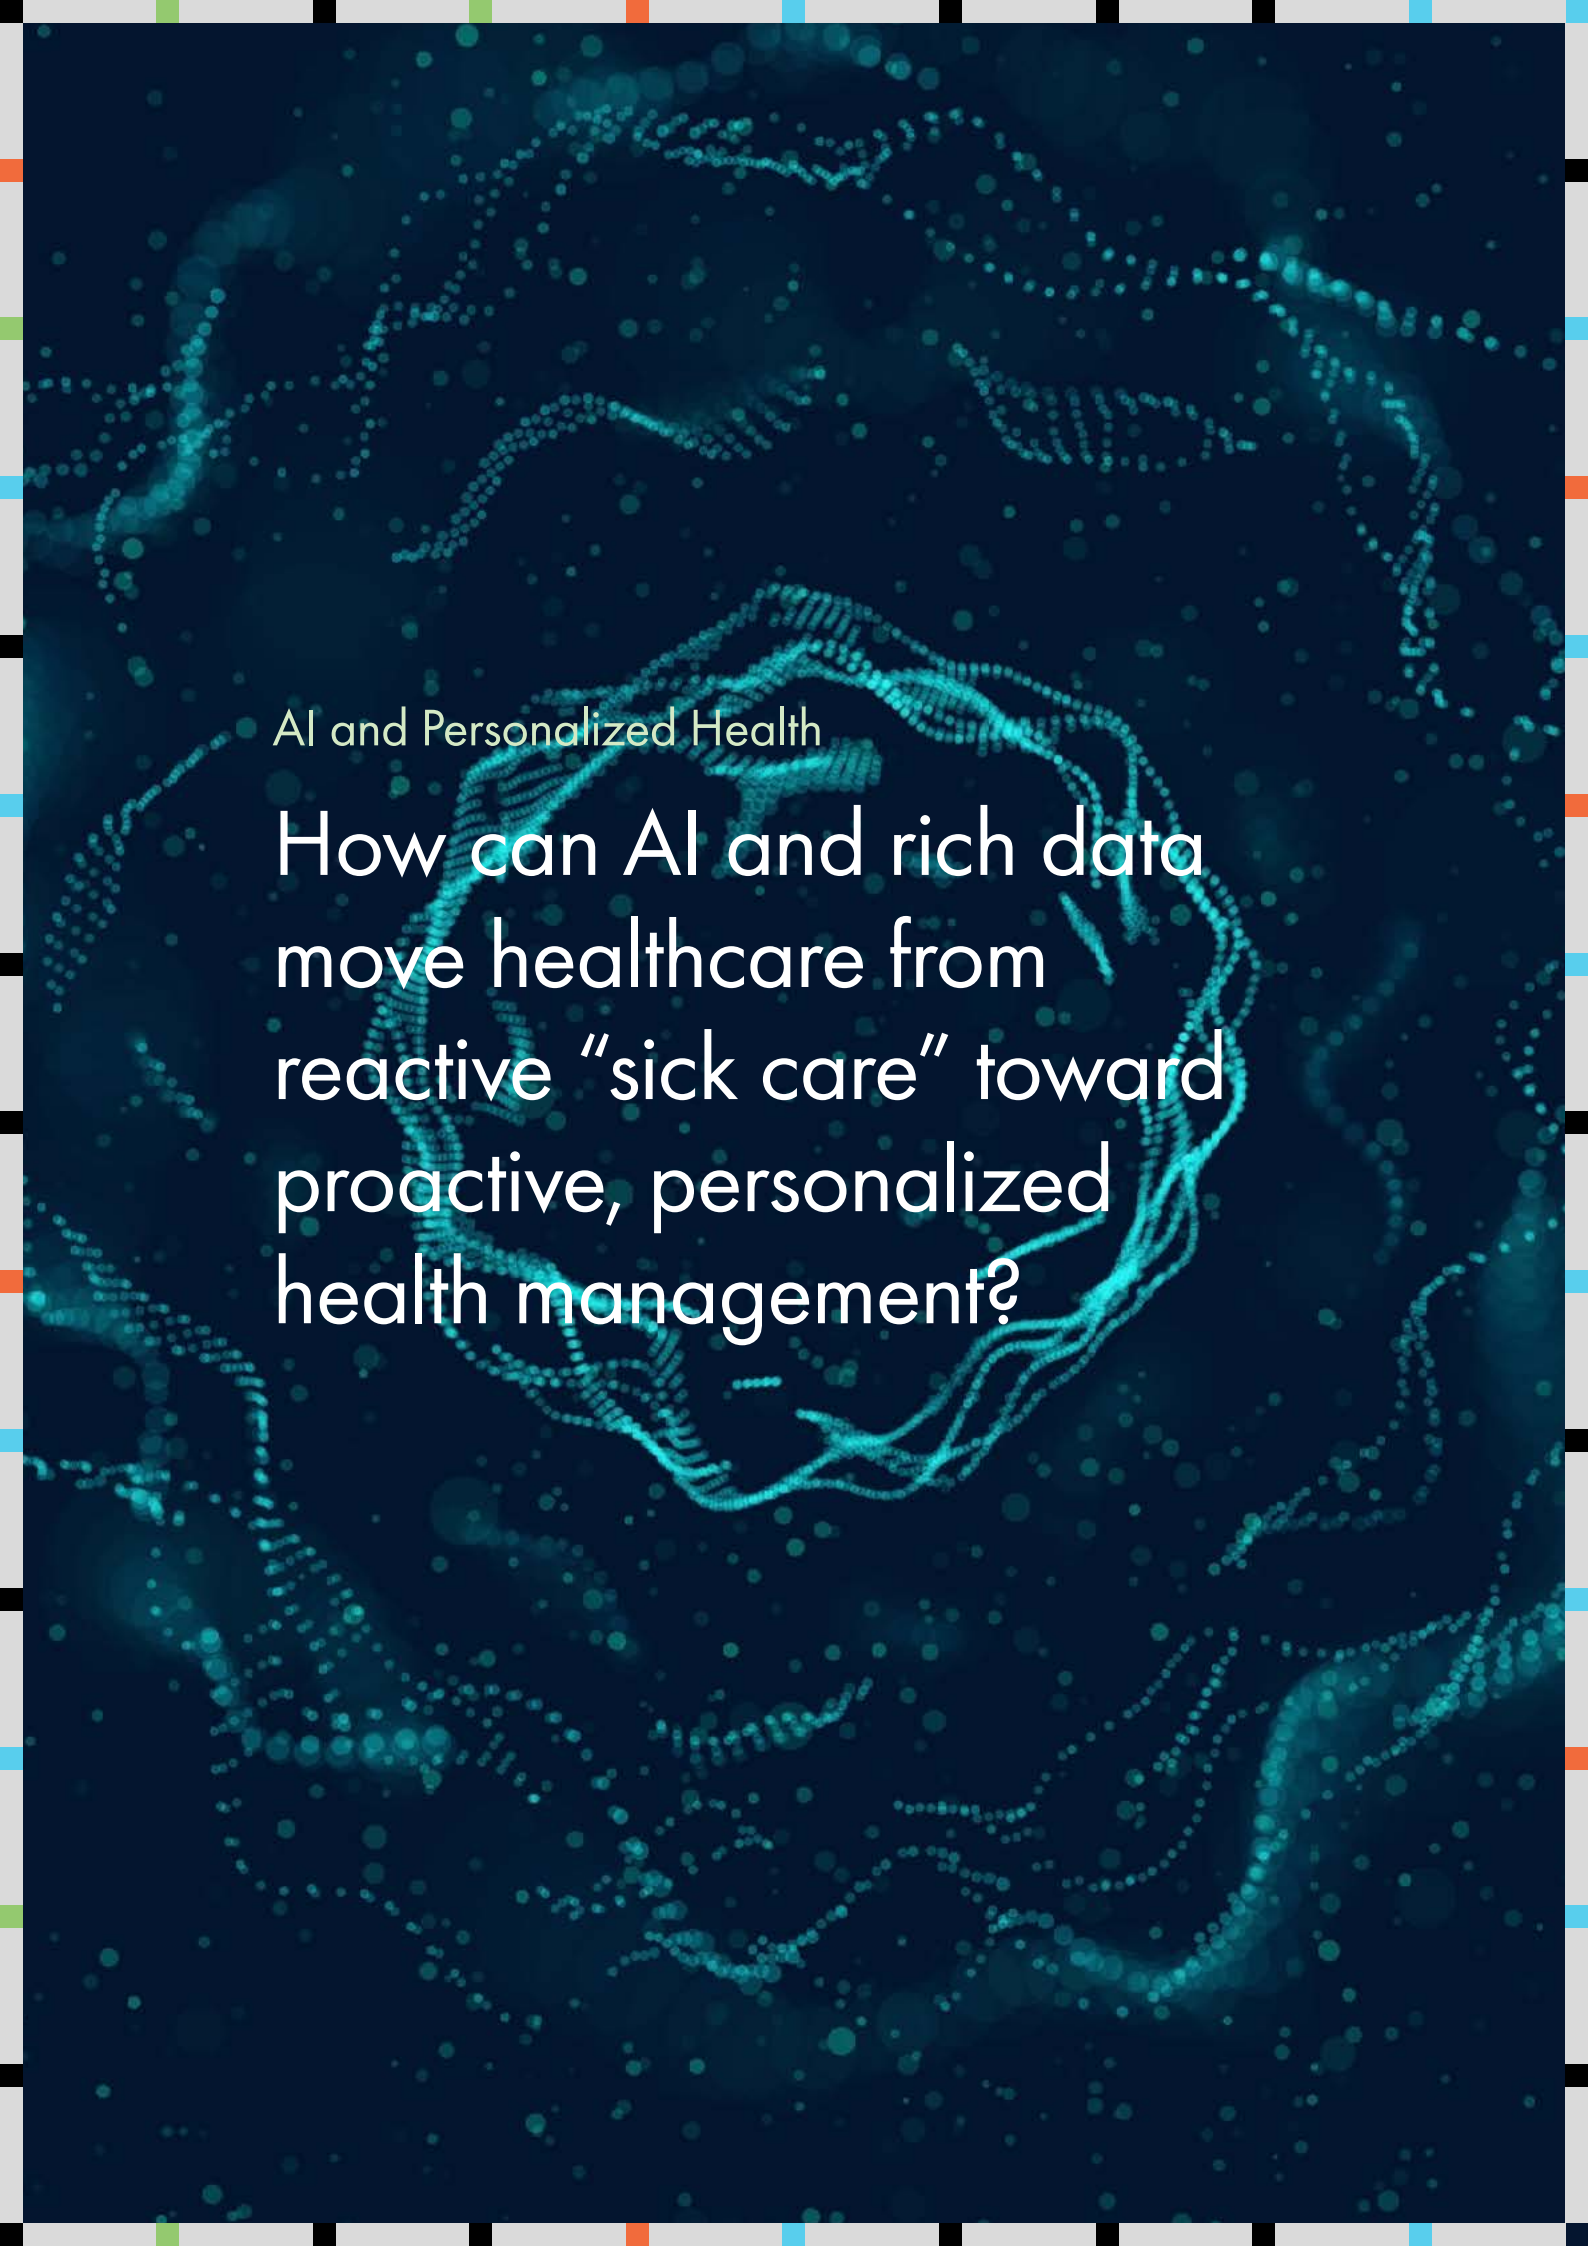
AI and Personalized Health
How can AI and rich data move healthcare from reactive “sick care” toward proactive, personalized health management?

AI and Personalized Health

Precision and preventive healthcare rely on integrating genetic, environmental and lifestyle data. Studies estimate that lifestyle and behavioral factors account for about 60 % of health determinants, genetics for roughly 30 % and medical history for around 10 %. Participants envisioned a shift from episodic "sick care" to continuous health management, powered by wearables, digital twins and AI driven recommendations. However, fragmented data, misaligned incentives and clinician scepticism present barriers.

Introduction

The roundtable's health session began with participants introducing themselves as clinicians, technologists, researchers and entrepreneurs. They quickly moved from introductions to personal stories demonstrating AI's growing impact in healthcare. One participant described a patient who learned of early-stage kidney disease from a smartwatch alert; another shared how their institution uses machine learning to monitor ICU patients for sepsis, sometimes anticipating crises hours before clinicians would have noticed. These anecdotes illustrated how real-time data and AI-driven monitoring can enhance early detection and improve patient outcomes.

Despite these advances, the group recognized that most determinants of health lie outside genetics. One panellist emphasized the "60–30–10 rule": environment, behavior and socioeconomic factors matter far more than genetics. Another argued that focusing **solely** only on genomics is like "reading the foreword of a long novel and ignoring the chapters."

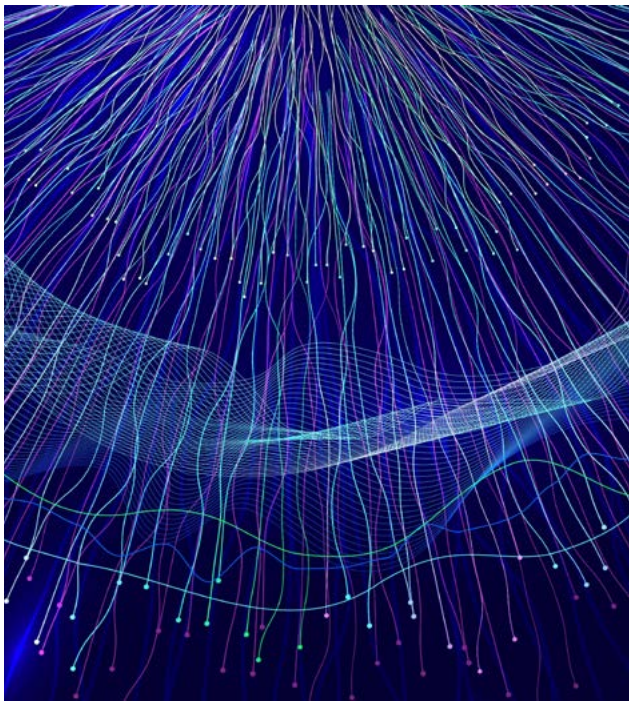

Participants highlighted successes in personalized medicine driven by emerging AI and data technologies. Multi omics sequencing has uncovered unexpected gene-drug interactions, while digital twins – virtual replicas of organs such as the heart or liver built from real-time clinical, genetic, and physiological data – can simulate individual responses to therapies and drugs. mRNA platforms, the repurposing of failed compounds, and generative models are also helping to expedite vaccine and drug development. In clinical care, machine vision systems now assist physicians by detecting conditions such as diabetic retinopathy or skin cancer, prompting earlier physician intervention. A participant pointed to a study where AI models predicted patient readmissions more accurately than clinicians. For example, the NYUTron system — a large language model trained

AI and Personalized Health

on over 4.1 billion words from 336,000 patient records — predicted hospital readmissions, inpatient mortality, and length of stay with up to 12 percent greater accuracy than traditional models or clinician assessments. It is important to note that physicians in these studies use these AI systems as decision-support tools — a type of second opinion or early warning system that flags at-risk patients while leaving diagnostic and treatment decisions firmly in human hands. Collectively, these breakthroughs show how AI can shorten drug discovery cycles, repurpose failed compounds and accelerate clinical trials, ultimately improving patient outcomes.

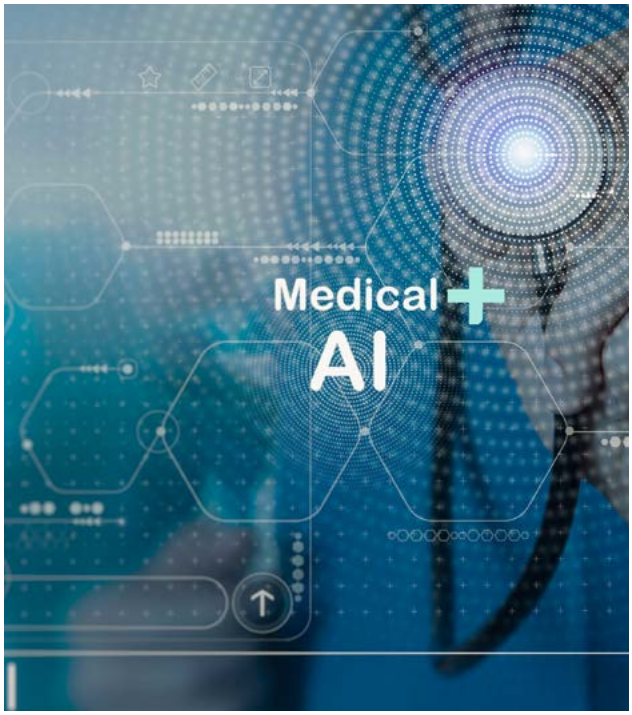

Yet barriers persist. Participants lamented that few physicians engage directly with AI conferences, with many clinicians still viewing AI as "not their problem." Data fragmentation and privacy concerns continue to limit cross institutional research and collaboration, constraining the scale and quality of medical AI research. Another challenge lies in misaligned financial incentives. Current payment models tend to reward treatment over prevention, discouraging investment in early detection and predictive care.

The discussion of automating care converged on the **70 % rule**: roughly 70 % of routine medical tasks — such as triaging minor illnesses or adjusting hypertension medication through remote monitoring — could be automated with remote monitoring and algorithms. This would free clinicians to focus on the remaining 30%: activities that demand surgical skill, nuanced diagnostic judgment, or compassionate human interaction. Consensus emerged that AI should assist, not replace, physicians: radiologists who use AI will replace those who don't.

AI and Personalized Health

Actionable Recommendations

The experts agreed on key areas where progress is essential:

RECOMMENDATION	RATIONALE	STAKEHOLDERS
Build interoperable, federated health data networks	Facilitate secure sharing of genomic, wearable and clinical data while preserving patient privacy. Federated/ Split learning and decentralized identifiers can enable cross-institutional research without exposing raw patient data.	Hospitals, regulators, health tech companies
Align incentives toward prevention	Reform payment models to reward early detection and chronic disease management rather than volume of procedures. Encourage insurers to invest in lifestyle interventions and remote monitoring.	Governments, insurers, employers
Educate clinicians in AI literacy	Offer training on interpreting AI outputs, recognizing biases and integrating algorithms into clinical workflows. Foster collaboration between data scientists and healthcare professionals.	Medical schools, professional associations
Invest in multi omics and digital twin technologies	Combine genomic, metabolomic, microbiome and wearable data to build patient-specific models. Use these models to predict disease trajectories and personalize treatment.	Research institutions, biotech firms
Promote participatory health platforms	Empower individuals to contribute data and share in the value created. Explore mechanisms for compensating patients whose data advances research, building public trust.	Policy makers, patient advocacy groups
Emphasise augmented intelligence	Design systems that support clinicians through decision support, documentation and predictive alerts without removing human oversight. Encourage regulators to adopt augmented intelligence as a guiding design principle.	AI vendors, health authorities

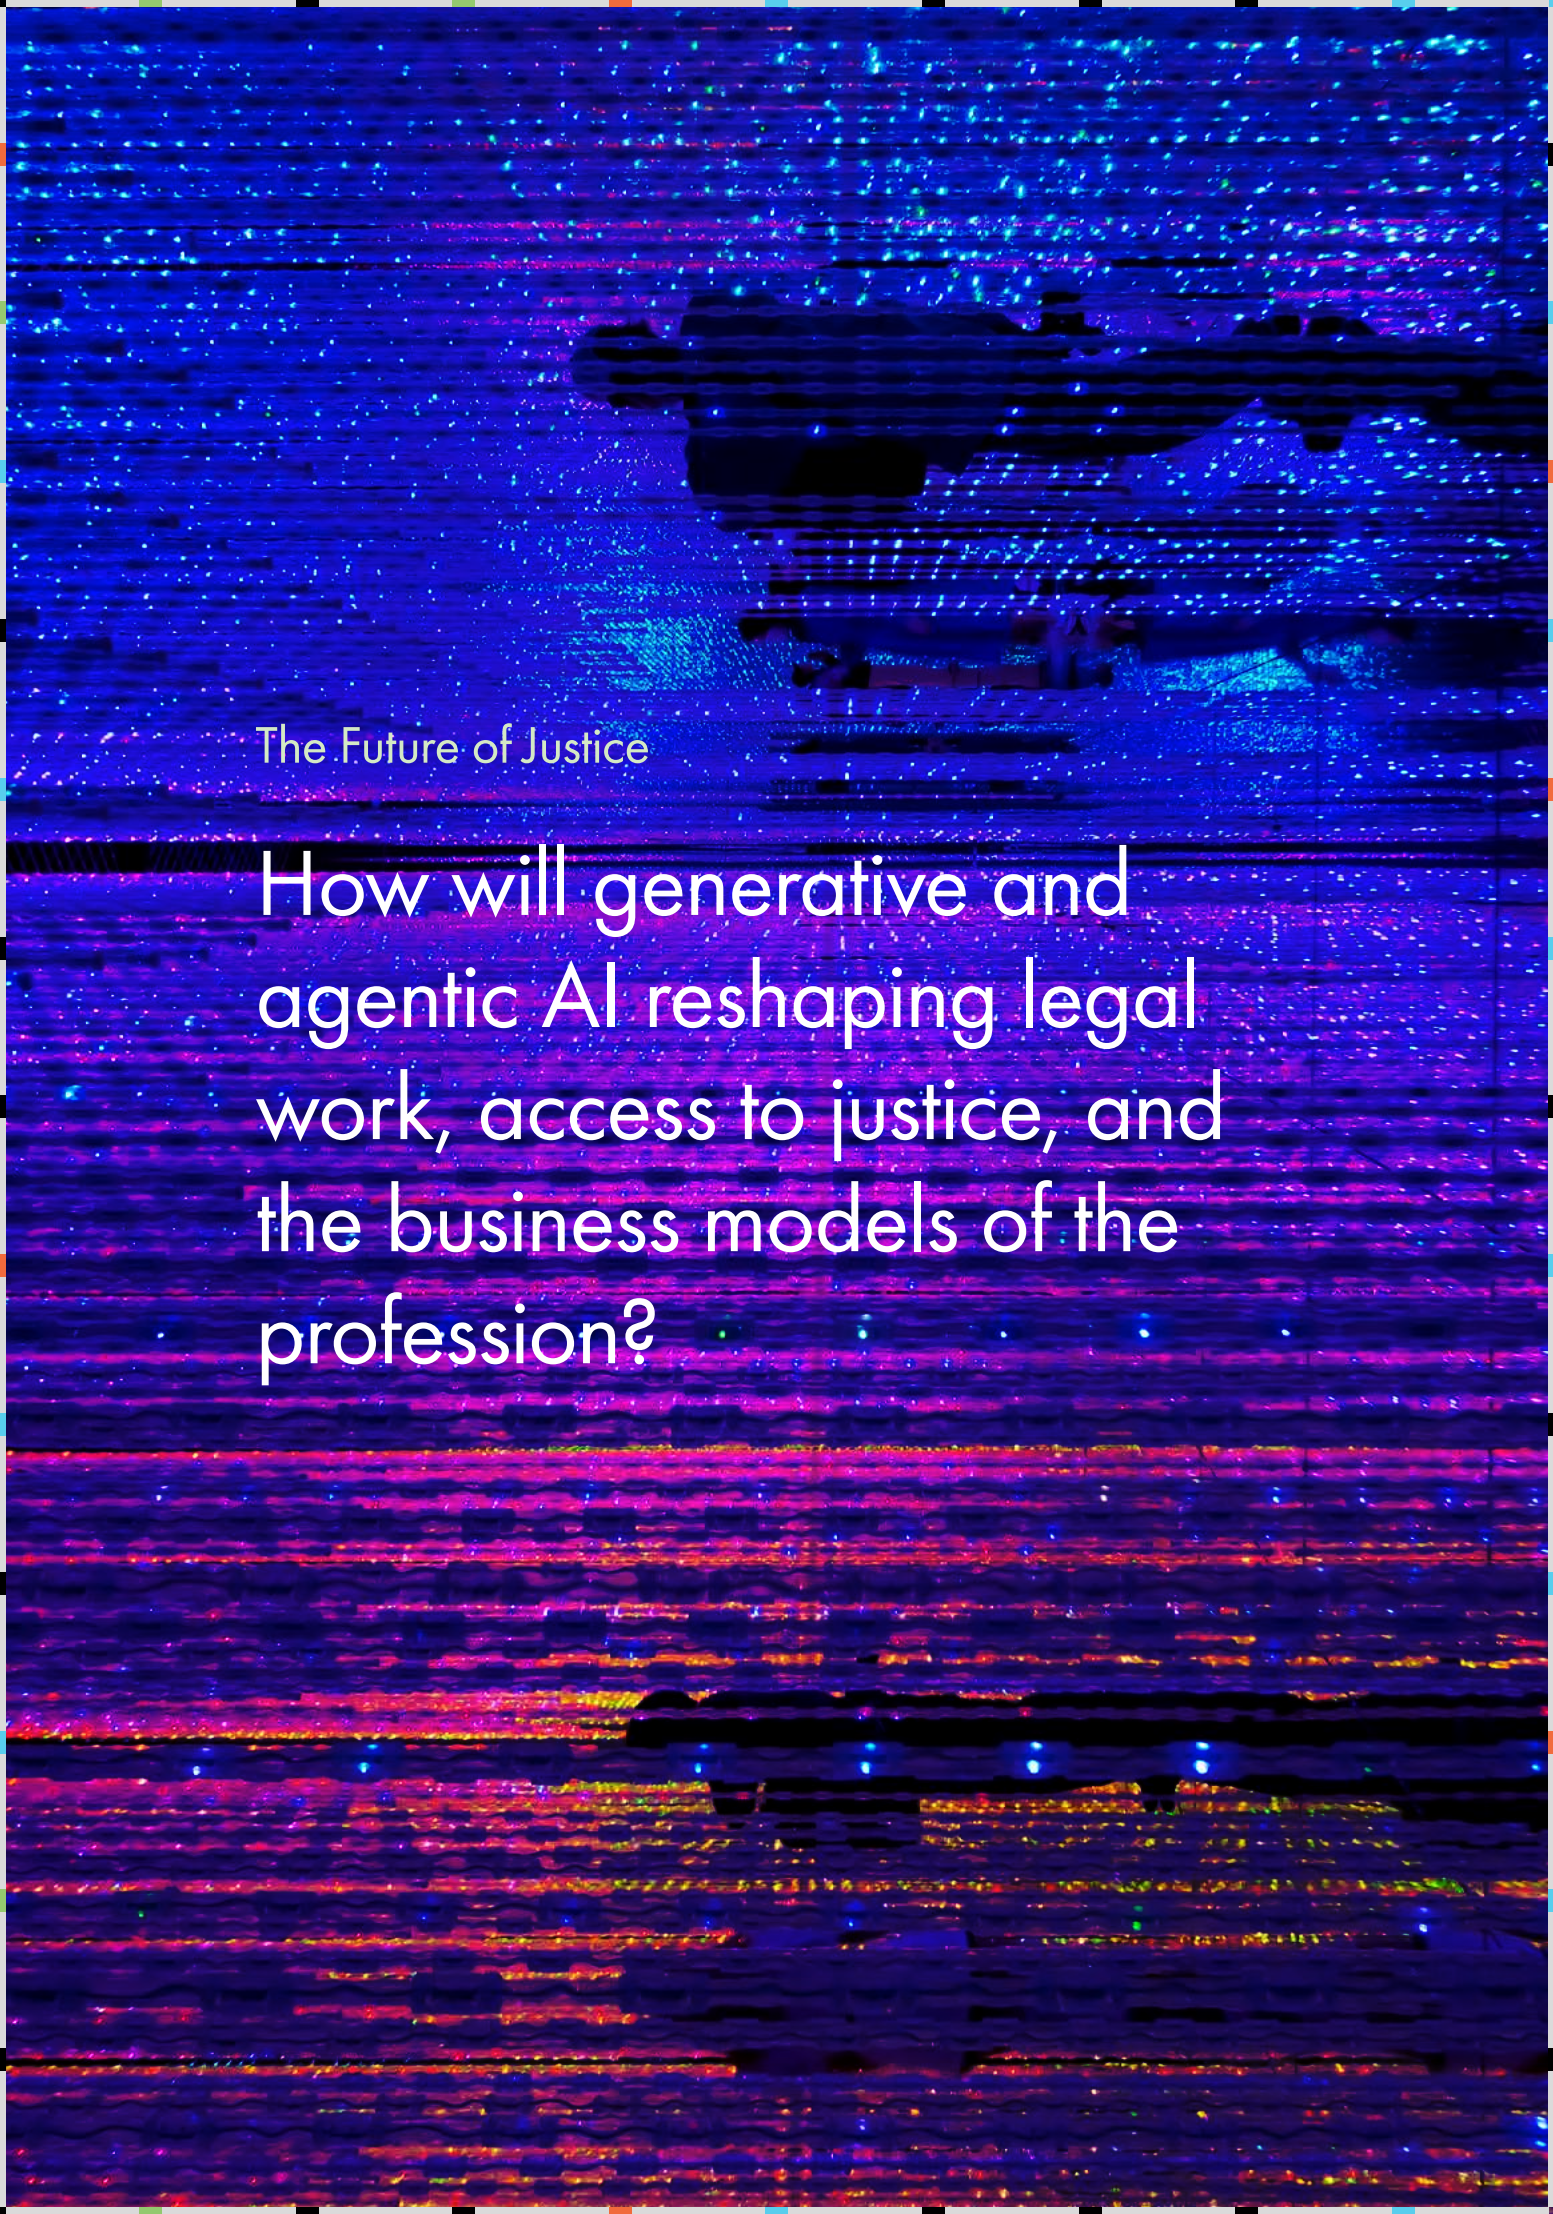
The Future of Justice
How will generative and agentic AI reshaping legal work, access to justice, and the business models of the profession?

Shaping the Future of Justice and Legal Practice

Generative AI is transforming legal services. Large language models can draft memos at a quality comparable to junior associates, summarize discovery documents and assist judges in identifying relevant precedents. A Harvard Law School interview with Professor David Wilkins, in a Harvard Law School interview, notes that AI tools can already match first-year associates in memo writing, at a substantially lower cost, though review by senior lawyers remains essential. Participants envisioned AI that improves access to justice by simplifying legal documents and automating routine tasks. They also cautioned that unequal access to advanced tools could exacerbate existing inequalities. Regulation, transparency and new business models will shape the profession's future.

Introduction

The session began with excitement about agentic applications to the legal field, systems that handle end-to-end tasks such as drafting contracts, conducting research and advising clients. Some panelists described early tools that assist judges by highlighting key case elements and preventing important details from being overlooked. Others voiced concerns about the long-term impact on the profession. If AI can produce better answers faster than humans, will we need as many lawyers? Will people lose legal reasoning skills if they rely too heavily on machines? A A few likened this to the autopilot system in aircraft: at first pilots resisted, but eventually the autopilot became indispensable and flying safer. Similarly, the legal profession may come to view agentic AI as a trusted co-pilot rather than a competitor.

Business models dominated much of the discussion. In-house counsel already perform tasks once outsourced to firms; AI may accelerate this trend. One in-house lawyer noted that generative models still struggle with domain specific issues like blockchain regulation, likely due to limited exposure to specialized training data, the rapid evolution of relevant laws, and the interpretive nuance these areas demand. Such material often lies outside generic training sets and requires access to proprietary or continuously updated data sources.

Participants agreed that high stakes litigation and complex transactions will still require specialist expertise, but routine drafting, discovery and compliance reviews are ripe for automation. Yet even seemingly basic legal tasks involve layers of nuance — interpretive choices, fairness considerations, and contextual judgment — that current AI tools may only partially capture. Ensuring that automation preserves these subtleties remains a key challenge for responsible adoption.

Law firms may shift from hourly billing to subscription-based services, with AI handling much of the commodity work. It remains an open question whether this transformation is primarily client-driven — as corporate legal departments demand predictability and efficiency — or whether leading firms are proactively adapting to preserve margins and attract tech-savvy clients. This underscores a central tension in the profession — AI's projected $20 billion in annual time savings and 190 hours freed per lawyer could revolutionize efficiency, yet nearly 90 percent of corporate legal spending still flows through hourly billing. Until firms and clients fully align incentives around value and outcomes, these productivity gains may paradoxically erode revenues instead of enhancing them, forcing a deeper rethink of how legal work is priced, measured, and rewarded.

Shaping the Future of Justice and Legal Practice

Regulation loomed large. The European Union's **AI Act** imposes strict obligations for transparency, risk assessment and human oversight. Under the Act, organizations must map data lineage, maintain model inventories and ensure compliance across subsidiaries. These requirements may burden small firms, but they also drive demand for legal tech solutions that help automate and manage compliance. Participants did not appear to address the specific risk of client data being input into AI systems. This omission overlooks a critical confidentiality concern — many generative models retain or learn from user inputs, meaning that sensitive legal information could inadvertently be exposed or reused by third-party systems if not properly safeguarded. Clients are increasingly wary of potential data exposure or misuse, making this an essential yet often neglected aspect of responsible AI adoption.

There was strong support for public-private collaboration in defining fairness metrics and ensuring that AI does not perpetuate or amplify existing biases. Participants stressed the potential of AI to expand access to justice in low resource settings by automating case triage, drafting petitions and translating legal documents into plain language. In practice, integrating AI-driven legal assistants into court websites and e-filing portals could further support self-represented litigants — helping them navigate procedural requirements without requiring full legal representation. Nevertheless, cultural and economic barriers will likely prevent public-interest legal applications from receiving the same level of investment as corporate or commercial tools. To bridge this gap, they urged universities, non-profits, and philanthropic foundations to take a more active role in supporting AI-for-justice initiatives.

Shaping the Future of Justice and Legal Practice

Actionable Recommendations

The experts agreed on key areas where progress is essential:

RECOMMENDATION	RATIONALE	STAKEHOLDERS
Adopt AI for routine legal tasks with human oversight	Use generative models to draft contracts, summarize case law and produce discovery memos. Lawyers should review AI output to correct errors and ensure ethical compliance.	Law firms, corporate counsel
Invest in equitable legal AI platforms	Develop open source or low-cost tools that small businesses and low income individuals can use to navigate legal processes. This reduces the justice gap and enhances public trust.	Nonprofits, startups, governments
Train legal professionals in AI literacy and ethics	Update legal education to include AI fundamentals, bias awareness and ethical guidelines. Encourage continuing education for practicing lawyers.	Law schools, bar associations
Design regulatory sandboxes and compliance tools	Create frameworks where legal practitioners can experiment with AI under regulatory supervision. Develop tools that document data sources, monitor model behavior and generate compliance reports for the EU AI Act and similar frameworks.	Regulators, legal tech vendors
Develop fairness and accountability standards	Collaborate across jurisdictions to define metrics for fairness, explainability and due process in algorithmic legal decisions. Require AI vendors to publish model capabilities and limitations.	Standard bodies, civil society
Explore new business models	Consider subscription or outcome-based pricing in place of hourly billing. Use AI to improve efficiency and transparency, freeing lawyers to focus on complex counseling and advocacy.	Law firms, accounting firms

Conclusions and Cross-Cutting Insights

Across all four domains, participants agreed that AI's promise will be realized only through **intentional design** and **socially responsible governance**. While autonomy, decentralization, personalization and legal reasoning each present unique challenges, the principles of transparency, accountability and human partnership remain essential to every application.

Invest in people and education

Technological advances will outpace human understanding unless curricula and professional training evolve accordingly. Clinicians, lawyers, engineers and policymakers must develop a clear grasp of AI's capabilities and limitations. Interdisciplinary education and training programs can cultivate a workforce equipped to co-create systems that serve the public good.

Align incentives with societal benefit

Business models, payment schemes and regulatory regimes should be designed to reward preventive approaches in healthcare, discourage risky agent behavior in financial markets and ensure that legal AI tools remain accessible to low-income users. The agent economy of the future must be built to create shared prosperity, not deepen existing economic and social divides.

Prioritize trust and fairness

Transparent documentation of data sources, model training and decision logic will build public trust and confidence. Establishing and applying fairness metrics – alongside independent audits – can detect and mitigate bias. Public engagement platforms should allow citizens to participate meaningfully in the governance and oversight of AI systems, ensuring that these systems reflect diverse values and societal priorities.

Embrace augmented intelligence

AI should amplify human capabilities, not supplant them. The most effective systems will integrate machine speed with human judgment, empathy and accountability. Evidence already shows that radiologists who use AI outperform those who do not and similar patterns are likely to emerge in law, engineering and governance.

By adhering to these principles – and implementing the specific recommendations outlined in each chapter – the organizations represented at the roundtables, together with the broader AI community, can shape technological progress toward outcomes that are equitable, trustworthy and beneficial to all.

References

[1] McKinsey & Company. (2024). Seizing the agentic AI advantage. McKinsey Global Institute.
[2] HUMAN Security. (2024). What is agentic AI? HUMAN Explainers.
[3] MIT Media Lab (2025). Decentralized AI. https://www.media.mit.edu/projects/decentralized-ai/overview/
[4] Stanford Medicine. (2023). Precision health and determinants of health. Stanford University School of Medicine.
[5] Stanford Medicine. (2023). Augmented intelligence in healthcare. Stanford University School of Medicine.
[6] Neal, J. (2024, February 14). Harvard Law expert explains how AI may transform the legal profession in 2024. Harvard Law Today. Retrieved from https://hls.harvard.edu/today/harvard-law-expert-explains-how-ai-may-transform-the-legal-profession-in-2024/
[7] Thompson Reuters. (2025, July 2). AI dooms the billable hour for Big Law earnings. Reuters Breakingviews. Retrieved from https://www.reuters.com/commentary/breakingviews/ai-dooms-billable-hour-big-law-earnings-2025-07-02/
[8] International Bar Association (IBA). (2025). Balancing efficiency and privacy: AI's impact on legal confidentiality and privilege. IBA News & Analysis. Retrieved from https://www.ibanet.org/balancing-efficiency-and-privacy-AI-impact-on-legal-confidentiality-and-privilege
[9] NYU Langone Health. (2025, June 4). New "AI doctor" predicts hospital readmission and other health outcomes. NYU Langone News. Retrieved from https://nyulangone.org/news/new-ai-doctor-predicts-hospital-readmission-other-health-outcomes
[10] Stanford Medicine. (2025, June 15). "The virtual you: How medical digital twins are revolutionizing healthcare." https://med.stanford.edu/medicine/news/current-news/standard-news/medical-digital-twins.html

Contributors

Aditi Joshi, Open Source Lead, Google

Abhishek Singh, Researcher, MIT Media Lab

Alex "Sandy" Pentland, Professor, Stanford HAI & MIT; Director, MIT Connection Science

Ami Bhatt, Chief Innovation Officer, American College of Cardiology; Associate Professor, Harvard Med

Alojzy Nowak, Rector, University of Warsaw

Amit Pradhan, CEO, Rainfall

Anne-Sophie Morand, Data & AI Governance Counsel, Swisscom

Ayisha Piotti, Managing Director, RegHorizon and ETH Zurich

Ben Moskowitz, Vice President, Innovation Lab

Camille Francois, Professor, Columbia University

Daniel Dobos, Research Director, Swisscom

Daniel Kraft, Founder and Chair, NextMed Health

Dazza Greenwood, Executive Director, Law.MIT

Edward Jung, Co-Founder and CTO, Intellectual Ventures

Elliott Ash, Professor, ETH Zurich

Erik Brynjolfsson, Professor, Stanford HAI

Gabriele Mazzini, Former AI Act Team Leader, European Commission

Guendalina Cobianchi, SVP Strategic Analytics & Business Insights, V-Nova Ltd

Hossein Rahnama, Visiting Professor, MIT

James A Landay, Professor and Co-Director, Stanford HAI

Jared Josleyn, SVP, Global Head of Digital Health and Emerging Disruptive Growth Exploration (E.D.G.E.), Sanofi

Javier Fernandez-Marques, Lead Research Scientist, Flower Labs

Jeannette Gorzala, CEO ACT AI NOW, Vice Chair AI Board Austrian Government

Jens Groth, Chief Scientist, Nexus

Joan LaRovere, SVP and interim Chief Medical Officer, Boston Children's Hospital

Julia Apostle, Partner, Orrick, Herrington & Sutcliffe LLP

Laura Brounstein, Founder, Brounstein Ventures and Strategic Advisor, Biden/Harris '24

Lee Schneider, General Counsel, Ava Labs

Martin Gajek, Director ML and Applied Research, Thomson Reuters

Nicholas Hudson, Director and Strategic Planning, IQVIA

Nitin Mittal, Principal, Deloitte, US AI Strategic Growth Offering Leader

Ori Goshen, Co-Founder, Co-CEO, AI21 Labs

Pablo Arredondo, Vice President, CoCounsel, Thomson Reuters

Peter Grabowski, Lead, Google AI Agents

Ramesh Raskar, Associate Professor, MIT Media Lab

Richard Socher, CEO/Founder, You.Com

Rika Nakazawa, Group Vice President, NTT Data Inc

Robert Grey, Chancellor, University of Warsaw

Robert Mahari, Associate Director, Stanford CodeX Center

Roland Scharrer, Partner & Managing Director, Kearney

Samuele Marro, Researcher, Oxford University

Scott O'Hair, CEO, Antenna Ventures

Seema Kumar, CEO, Deepfield Cure

Tobin South, Head of AI Agents and MCP, WorkOS

Vivek Sharma, Scientist, Ex MIT Media Lab, Harvard Med

Yaniv Altshuler, Researcher, MIT Media Lab and MIT Connection Science Fellow

Yejin Choi, Professor, Stanford HAI

Editors

Robert Mahari
Associate Director,
Stanford CodeX Center

Vivek Sharma
Scientist; Ex MIT Media Lab,
Harvard Med

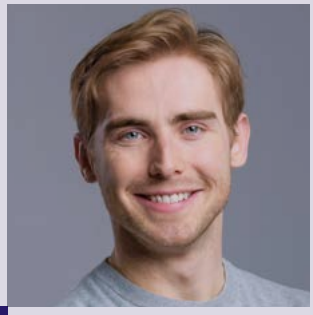

Tobin South
Fellow, Stanford University;
Head of AI Agents and MCP, WorkOS

Alex "Sandy" Pentland
Professor, Stanford HAI & MIT; Director,
MIT Connection Science

Ramesh Raskar
Associate Professor, MIT Media Lab

We would like to thank: **MIT Connection Science**, **the Harvard Law School Center on the Legal Profession**, **the Stanford Digital Economy Lab**, **Wisdom House**, **alphabiome**, **Thomas Reuters**; and **Antenna.VC** for their partnership and support.

We thank **Allessia Chiappetta** for her excellent research assistance.

LEADING BEYOND BOUNDARIES