

AI Disclosures Project

Human+AI Markets

Architectures and Mechanisms for Distributed Value Creation

[Bellagio Convening](#) • Italy • 27 April – 1 May 2026

Participant Bios



Tadas Antanavicius

Co-Founder, PulseMCP

Tadas is a career software engineer and founder, currently serving as a member of the MCP Steering Group where he co-maintains the MCP Registry and MCP Contributor open source community. He writes a semi-regular newsletter about MCP and agentic engineering via his startup, PulseMCP, and spends much of his time in consulting engagements with software engineering teams working to adopt coding agents in their day to day work. He is passionate about exploring MCP's promise for facilitating incentive-aligned productivity gains in an AI-first era, and empowering people with productivity tooling that helps them get more done, while staying ahead of having their work be replaced by ever-improving AI capabilities.



Geoffrey Bilder


Head of Technology Strategy, MIT Press

Geoffrey Bilder is Head of Technical Innovation and Strategic Partnerships at MIT Press. His career spans three decades of building the technical infrastructure that underpins how scholarly knowledge is identified, verified, attributed, and exchanged across distributed publishing networks. He co-founded Brown University's Scholarly Technology Group in 1993, providing the academic community with technology consulting in support of research and scholarly communication. He subsequently served as head of IT R&D at Monitor Group, the global management consulting firm, before moving into scholarly publishing as CTO of Ingenta from 2002 to 2005, and later as a publishing technology consultant before joining Crossref.

At Crossref, the nonprofit that maintains the DOI infrastructure underlying global scholarly publishing, Bilder led the technical development and launch of several services that are now foundational to how research is identified and verified. These include Similarity Check, the plagiarism screening service used by thousands of journals; Crossmark, which allows readers to see whether a published work has been corrected or retracted; ORCID, the persistent researcher identifier system now used by millions of academics and integrated into the publishing workflows of most major journals and funders; and the Open Funder Registry, which standardizes how research funding acknowledgments are recorded and linked. Each of these systems addresses a version of the same underlying problem: distributed networks of knowledge producers need shared identifiers and protocols to attribute, verify, and link their work reliably. Crossref's metadata now covers over 180 million scholarly records and handles over two billion DOI resolutions per month.

At MIT Press, Bilder is working through those same questions in the context of AI. The scholarly publishing infrastructure he helped build was designed for a world in which humans navigated research through citations, DOIs, and structured metadata. AI systems consume, synthesize, and recirculate that material in ways that often bypass those attribution and provenance mechanisms. Much of the debate around attribution has focused, understandably, on compensation. But there is an epistemic issue underneath it: if people cannot identify trustworthy content, they will be less likely to rely on it, and less likely to pay for it. The challenge is therefore not only to protect the economic value of knowledge work,

but to preserve the signals that make knowledge credible, inspectable, and worth supporting in the first place. His presence at the convening reflects both his technical depth in building open scholarly infrastructure and his practical familiarity with what it takes to get competing institutions to adopt shared standards over time.



Alissa Cooper

Executive Director, Knight-Georgetown Institute (KGI), Georgetown University

Alissa Cooper is the inaugural Executive Director of the Knight-Georgetown Institute (KGI), a center based at Georgetown University dedicated to connecting independent research with technology policy and design. KGI is designed to provide practical resources that policymakers, journalists, and private and public sector leaders can use to tackle tech policy issues in real time. Alissa leads and oversees the organization's work across platform governance, competition policy, and artificial intelligence.

Alissa is a recognized leader in the development of global Internet standards, policy, and governance. She has served in a variety of roles in the tech industry, including Chair of the Internet Engineering Task Force (IETF), the world's premier Internet standards organization. She led the IETF through significant transitions related to Internet security and encryption, network performance, privacy, and real-time voice and video.

Prior to joining KGI, Alissa spent a decade at Cisco Systems in senior engineering and executive roles, including Vice President of Technology Standards and Vice President and Chief Technology Officer for Technology Policy. She played a central part in shaping Cisco's approach to technology standardization, public policy, and privacy across the company's networking, security, and collaboration businesses. Alissa was the first woman in Cisco history to be promoted to Fellow, the company's highest engineering distinction.

Prior to joining Cisco, Alissa served as the Chief Computer Scientist at the Center for Democracy and Technology, where she was a leading public interest advocate and technologist focused on privacy and

net neutrality. Alissa holds a D.Phil from the Oxford Internet Institute and M.S. and B.S. degrees in computer science from Stanford University. She currently serves on the board of The Tor Project and the advisory board of the Princeton Center for Information Technology Policy (CITP).



Jacques Crémer

Professor of Economics, Toulouse School of Economics

Jacques Crémer received his undergraduate degree from the Ecole Polytechnique, followed by an SM in Management and a PhD in Economics from MIT, completed in 1977. He held appointments at the University of Pennsylvania and Virginia Tech before joining the Toulouse School of Economics in 1991 as a CNRS Research Professor. At TSE he has served as Director of the Institut d'Economie Industrielle, Scientific Director of the school, and as the founding director of its Digital Center. He is a Fellow of the Econometric Society, a Fellow of the European Economic Association, and has been a member of the Econometric Society's Council. His early work produced foundational contributions to planning theory, auction theory, and the economics of organizations, and he is among the economists who helped establish the modern theory of the firm as a problem of information and incentives rather than simply one of ownership.

His research on digital markets began in earnest in the late 2000s and has developed into one of the most sustained bodies of work on platform competition and incumbency dynamics in the field. His most influential contributions in this area concern switching costs and network effects, and in particular the conditions under which an established platform retains dominance even when a better or cheaper alternative exists. A series of papers with Gary Biglaiser and others analyzes what they call incumbency advantage: the structural fact that users of a dominant platform have incentives to wait for others to migrate before moving themselves, which means no migration takes place even when it would be collectively beneficial. A key counterintuitive result from this work is that more frequent migration opportunities can actually increase incumbency advantage, because each individual opportunity to

migrate raises the bar for the next. This line of research has direct implications for data portability mandates, interoperability requirements, and the design of remedies in platform competition cases.

From April 2018 to March 2019, Crémer served as a Special Adviser to European Commissioner for Competition Margrethe Vestager, during which time he co-authored with Yves-Alexandre de Montjoye and Heike Schweitzer the report "Competition Policy for the Digital Era," published by the European Commission in April 2019. The report became one of the most cited policy documents in the history of EU digital regulation. It argued that the fundamental goals of competition law remain valid in the digital era but that established tools and methodologies need to be updated to account for the speed of market tipping, the competitive role of data, the self-reinforcing dynamics of digital ecosystems, and the inadequacy of narrow consumer welfare standards for capturing harms that unfold over time. Its recommendations fed directly into the design of the Digital Markets Act, and Crémer has continued to contribute to debates around the DMA's implementation, enforcement, and the role of economic analysis in proceedings against designated gatekeepers.



Paul Farrow

Principal Product Manager, Microsoft

Paul Farrow holds a PhD in Neuroscience from the University of Bristol and has spent close to a decade working at the intersection of digital advertising, content economics, and privacy. He came to Microsoft via AppNexus/Xandr, and he serves on the boards of Prebid, IAB Europe, and the Network Advertising Initiative, as well as co-leading Prebid's LLM and Publisher Monetization Task Force for 2026, which is developing scalable monetization frameworks for an AI-driven web.

At Microsoft Advertising, Paul led the privacy-preserving advertising strategy, including the development and rollout of the Ad Selection API. This API enables personalized targeting through differential privacy, k-anonymity, and trusted execution environments, delivering relevance without third-party cookies or cross-site identifiers. His current focus includes the AI Economy and the Publisher Content Marketplace

(PCM), which is helping to create the infrastructure for the AI answer engines and content monetization of the future.



Lucky Gunasekara

Cofounder, Miso.ai

Lucky Gunasekara is the cofounder of Miso.ai, where he works on grounded Answers, private publisher-specific LLM systems, and publisher-controlled AI infrastructure. Earlier in his career, he cofounded medic.org while at Stanford, spent two years working with HIV and maternal and child health clinics across East Africa, and built epidemicQ, an early multilingual OSINT and outbreak-intelligence system. That work was an early training ground in how to aggregate, ground, and attribute insights from news and public information at scale, research he later pursued further into OSINT-based market forecasting as a Visiting Researcher at Intel.

The work that led to Miso.ai began during his time as a Technologist in Residence at Cornell Tech, where he worked on privacy-first personalization, retrieval, and early answer systems. At Miso.ai, he, and his cofounder Andy Hsieh, have worked closely with O'Reilly Media to develop O'Reilly Answers, an early proof point for citation-backed Answers, grounded retrieval, and author-attribution and royalty logic. Today, Miso.ai supports more than 5 million AI Answer conversations each month across 50+ partnerships with major news, research, and media organizations. A consistent theme in this partnership-driven work has been building AI systems that are private to the publisher or institution itself: retrieval-first, grounded in the organization's own corpus, tuned to its own language and policies, and designed so prompts and usage do not simply become training fuel for shared public systems.

More recently, his work has focused on the standards, telemetry, and market infrastructure needed for a more permissioned AI web. He leads Project Sentinel, which tracks machine access behavior across roughly 12,000 publisher and news sites and has identified more than 2,000 bots and scrapers; Miso's recent findings suggest publishers explicitly target only about 22 bots on average, that roughly half of

domains still do not block even obvious bots, and that clear disallow rules are still often ignored in practice. Alongside that research, he is helping build Studio, Miso's MCP- and workflow-oriented tooling for publishers, and Bonsai, an agentic research network designed around transparent, paid, permissioned use of news, research, and knowledge sources. Across that work, his focus is on giving knowledge producers more agency: clearer ways to control how their work is used, better attribution and telemetry, and market structures where access, compensation, and demand are visible instead of hidden inside scraping.



Alex Hancock

Software Engineer, Block

Alex Hancock is a software engineer at Block, the fintech company behind Square, Cash App, and Afterpay, where he works on goose, Block's open-source AI agent framework. Goose is a local-first agent that combines language models with extensible tools and MCP-based integrations, designed to provide a structured and reliable environment for building and executing agentic workflows. It was released in early 2025 and quickly became one of the more widely used open-source agent frameworks in the developer community, used by thousands of engineers internally at Block and by a growing global user base outside the company.

Hancock was among the first engineers pulled into the effort to iterate on MCP. In February 2025, MCP's creators David Soria Parra and Justin Spahr-Summers reached out to the PulseMCP and goose teams to help build a centralized server registry. Hancock was involved in the early stages of this project. More recently he has taken over the development and maintenance of the official Rust SDK for MCP. Hancock has represented Goose and Block at MCP Developers Summits and is involved in the technical governance of MCP through the AAIF.

In addition to his work on MCP, Hancock is also working on the Agent Client Protocol as another open standard in the AI agent space and is contributing in the Transports Working Group of that protocol.

In December 2025, Block contributed Goose to the Linux Foundation's newly formed Agentic AI Foundation, a vendor-neutral consortium co-founded by Anthropic, Block, and OpenAI with platinum membership from AWS, Google, Microsoft, Bloomberg, and Cloudflare. The AAIF's three founding projects are MCP, Goose, and OpenAI's AGENTS.md, collectively intended to serve as the open infrastructure layer for the agentic AI era, analogous in ambition to what the W3C and IETF provided for the web. Under AAIF governance, Goose retains its open-source license and transitions to community stewardship, with Block engineers continuing as active contributors and maintainers.



Ola Hungerford

Principal Engineer, Nordstrom

Ola Hungerford is a Principal Engineer at Nordstrom, where she leads AI enablement initiatives with a focus on platform engineering for MCP and AI agents. Her technical background spans nearly two decades and covers a wide range of domains: front-end web development, recommender systems, cashless payments, business intelligence, enterprise resource planning, and scoring software for professional bowling tournaments. At Nordstrom she has spent the better part of a decade on the personalization team, working on how to make product discovery feel tailored without crossing into territory that feels intrusive, a problem she approaches as requiring both technical and human judgment in roughly equal measure.

Her involvement with the Model Context Protocol began early in the project's public life. She maintains the MCP Inspector, the developer tool used to test and debug MCP server implementations, and the Reference Server repository, which serves as the canonical educational example for developers building their own servers. She also moderates the MCP Contributors Discord server and its working groups, which have grown to include over 2,900 members. In November 2025, when the MCP core team published a retrospective on the protocol's first year, Hungerford was quoted on the coordinating function of the protocol itself, noting that the value of MCP is less the technical specification and more the shared reason it gives people to coordinate and talk about the same thing. She has spoken at the MCP Dev Summit North America, and co-authored posts on the official MCP blog covering server

instructions and tool annotations. In March 2026 she co-authored a detailed technical analysis of tool annotations as a risk vocabulary for agentic workflows, alongside contributors from GitHub and AWS.

Outside of her engineering work, Hungerford is a vocalist, songwriter, and studio co-founder. She records music with her husband, composer and producer Johnny Dexter Goss, under the names Cock and Swan, and has performed at KEXP, Bumbershoot, and Decibel Festival. She has also contributed to recording sessions for other artists as part of Johnny's recording and production studio, Dandelion Gold.



Nicole Immorlica

Professor of Computer Science, Yale University; Microsoft Research

Nicole Immorlica received her PhD from MIT in 2005, followed by postdocs at Microsoft and CWI and a professorship at Northwestern. In 2012, she moved to Microsoft Research New England, where she is a member of the Economics and Computation Group. In 2025 she additionally joined Yale as a Professor of Computer Science. She is an ACM Fellow and a Society for the Advancement of Economic Theory Fellow, and has received the Sloan Fellowship, the Microsoft Faculty Fellowship, and the NSF CAREER Award.

Convinced that algorithmic tools could reshape how markets work in practice, Immorlica was among the first cohort of computer scientists to take economics seriously as a research domain. Her most cited work spans large-scale auction design, which now underpin much of the revenue architecture of major technology platforms; social network dynamics, including segregation models and the role of referral networks in inequality; and matching market theory, where her 2005 paper "Marriage, Honesty, and Stability" received the 2023 SIGecom Test of Time Award for its analysis of when participants in large random matching markets have strong incentives to misreport preferences.

More recently Immorlica has turned toward studying generative AI as an economic phenomenon. A 2024 paper with Brendan Lucier and Aleksandrs Slivkins proposes a framework in which AI is modeled not as a tool that reduces costs for human agents, but as an economic agent itself, with potentially different information and preferences from the user it serves. In this framing, the equilibria that emerge when

users consult AI before acting in a game can differ qualitatively from equilibria without AI, raising questions about how market design needs to change when the agents making decisions are partially autonomous. She views theory as indispensable for exactly this kind of forward-looking work, noting that unlike data, which can only tell you what has happened, theoretical models allow you to reason about possible futures that do not yet exist.



Mallory Knodel

Founder & Executive Director, Social Web Foundation

Mallory Knodel is the founder and Executive Director of the Social Web Foundation, a nonprofit launched in 2024 to expand the fediverse, improve the ActivityPub protocol and its user experience, inform policymakers about federated social networks, and educate the public about participating in them. Before SWF she spent four years as Chief Technology Officer of the Center for Democracy and Technology in Washington, where she directed technical work on encryption, censorship circumvention, content moderation infrastructure, and cybersecurity policy, and represented CDT in U.S. and international standards and policy venues. She was the Head of Digital at ARTICLE19 in London and the technical specialist for the Association for Progressive Communications in Johannesburg, after an early career as an Indymedia activist.

Her standards work has been continuous and substantive. She chairs the Human Rights Protocol Considerations research group at the Internet Research Task Force, chairs the ATP working group at the IETF, has served on the Internet Architecture Board, chaired the Public Interest Technology Group, and also remains an active contributor at the W3C, IEEE, and the United Nations on technical policy questions. She is a US delegate to the ITU. She is co-author of *How the Internet Really Works*, a widely translated introduction to protocols, privacy, and internet governance aimed at non-specialists.

Her nonprofit work has attracted nearly ten million dollars in funding from the Ford Foundation, Open Society Foundations, Omidyar Network, MacArthur Foundation, Project Liberty, the US State

Department, SIDA (Sweden), Open Tech Fund, Sovereign Tech Agency, Meta, Apple, Google, Ghost, Mastodon, Bluesky, and more.

She advises the Freedom Online Coalition governments and the Open Technology Fund, teaches internet governance and technology policy, and publishes regularly on technology standards and human rights. Her through-line is a question: what would it mean to design technology that serves people and accelerates social movements?



Mike Linksvayer

VP of Developer Policy, GitHub

Mike Linksvayer is VP of Developer Policy at GitHub, where he leads the company's engagement with policymakers, regulators, and standards bodies on issues affecting how software gets built and shared. He has spent over two decades working at the intersection of open licensing, intellectual commons, and technology policy. Before GitHub, he joined Creative Commons as CTO in 2003, became VP in 2007, and spent nearly a decade there, during which he co-authored ccREL (the Creative Commons Rights Expression Language): the machine-readable standard for expressing licensing metadata on the web. In 2000, he co-founded Bitzi, the file-metadata company that developed magnet links, the identifier format still used in BitTorrent and elsewhere. He is also a Fellow at Open Forum for AI and has previously served as a volunteer director of Software Freedom Conservancy, chaired the Open Definition Advisory Council, contributed to The Wealth of the Commons essay collection, and written for The Data Journalism Handbook and The Art of Community.



Markus Mobius

Senior Principal Researcher, Microsoft Research

Markus Mobius is Senior Principal Researcher at Microsoft Research New England, a Research Associate at the NBER, and affiliated faculty at the University of Michigan School of Information. He earned his PhD in Economics from MIT in 2000, an M.Phil. in Economics and B.A. in Mathematics from Oxford, and was previously Associate Professor of Economics at Harvard.

His earlier work helped establish the modern economics of social networks. "Trust and Social Collateral" (Karlan, Mobius, Rosenblat, Szeidl, QJE 2009) formalized how network ties function as enforcement mechanisms for informal contracts — a framework that has shaped a decade of work on informal credit, labor referrals, and risk-sharing, including his "Consumption Risk-sharing in Social Networks" (AER 2014) and "Naive Learning with Uninformed Agents" (Banerjee, Breza, Chandrasekhar, Mobius, AER 2021). A separate empirical line, including "Why Beauty Matters" (Mobius & Rosenblat, AER 2006) and his work on motivated beliefs and self-confidence, demonstrated how non-merit signals propagate through labor and information markets.

For roughly the past decade, Markus has used Microsoft's browsing and search data to study the political economy of digital news. "Measuring the News and its Impact on Democracy" (Mobius, Rothschild & Watts, PNAS 2021) is one of the most-cited empirical contributions to the misinformation debate. His earlier work with Susan Athey on Google News and local content and his collaborations with David Rothschild and Duncan Watts have shaped how platform economists think about aggregation, news production incentives, and the structure of attention markets.

Markus is a co-author of "The Agentic Economy" (Rothschild, Mobius, Hofman, Dillon, Goldstein, Immorlica, Jaffe, Lucier, Slivkins & Vogel, 2025), a brief that argues the architecture of agentic communication is the decisive design choice for whether generative AI redistributes economic opportunity or recentralizes it. The paper distinguishes between *unscripted* interactions — agent-to-agent exchanges where the protocol allows open-ended natural-language negotiation rather

than pre-defined API call patterns — and *unrestricted* interactions – exchanges that aren't fenced in by a platform's market rules (i.e., your assistant agent can transact with *any* service agent, not just ones inside a walled garden).



Mark Nottingham

Member, Internet Architecture Board; Standards Lead, Cloudflare

Mark Nottingham is a standards engineer based in Melbourne, Australia who has spent over twenty-five years working on the core protocols of the internet. He has authored or substantially contributed to more than thirty IETF RFCs and W3C Recommendations, covering HTTP semantics and caching, HTTP/2, QUIC, web security, and URL standards. He has chaired the IETF HTTP Working Group since 2007, previously chaired the QUIC Working Group, and has served on the W3C Board of Directors, the Internet Architecture Board, and the W3C Technical Architecture Group. He currently works as Standards Lead at Cloudflare and holds a Graduate Diploma in Communications Law from Melbourne Law School. His career has moved across Akamai, Fastly, Yahoo, and various standards bodies, consistently focused on the question of how protocol design shapes the distribution of power on the internet.

In 2023 Nottingham published RFC 9518, a formal IETF document on centralization, decentralization, and internet standards, which grew out of work he began during his time on the Internet Architecture Board. The document examines how protocols can either resist or inadvertently enable centralization, and reaches a careful conclusion: standards bodies have a limited but real role to play, and should not be expected to solve what are ultimately market and political problems on their own. This body of thinking directly informs his current work. In 2025 he took on the chair role of the newly formed IETF AI Preferences (AIPREF) Working Group, which is developing standardized ways for content creators and publishers to express machine-readable preferences about how their work is collected and used in AI training pipelines. The group's position is that the current patchwork of robots.txt conventions is not working, that AI crawlers regularly ignore or route around them, and that a common vocabulary and

attachment mechanism embedded at the protocol level is needed to give publishers meaningful recourse.

Alongside the AIPREF work, Nottingham co-authored a draft IETF specification for paid web crawling, setting out technical requirements for a protocol-level system in which AI crawlers would compensate publishers at the point of access. The draft treats payment as an engineering problem requiring the same kind of careful requirements analysis that the IETF applies to any other protocol feature, including security considerations and interoperability constraints. His broader interest is in how legal regulation and technical standards interact, and whether the internet standards community can help preserve the decentralized character of the web as AI systems increasingly become a primary means of accessing content.



Tim O'Reilly

CEO, O'Reilly Media; Founder and co-director, AI Disclosures Project

Tim O'Reilly is the founder, CEO, and chairman of O'Reilly Media, the technical publishing and learning company he has run for over four decades. During that time he coined the terms "open source" and "Web 2.0," organized the first conferences around many of the technologies that now underpin the modern internet, and developed a reputation for identifying structural shifts in technology markets before they become legible to the broader industry. His books, conferences, and essays have shaped how generations of software engineers and technology leaders think about platforms, standards, and the relationship between technical architecture and economic power.

In 2024 O'Reilly co-founded the AI Disclosures Project with economist Ilan Strauss, where he serves as principal investigator. The project sits at the intersection of corporate transparency, market design, and AI governance. Its work includes producing disclosure frameworks for AI companies, analyzing the competitive dynamics that AI is generating across industries, and making the case that existing

regulatory tools, including SEC disclosure requirements, can be adapted to surface information about how AI systems are trained and deployed. A 2025 paper from the project drew attention for applying membership inference methods to show that GPT-4o had likely been trained on paywalled O'Reilly Media books without authorization.

O'Reilly's current argument, developed across a series of essays, policy briefs, and presentations, is that disclosure frameworks should be understood less as warning labels and more as networking protocols. His contention is that the open web succeeded because shared protocols allowed independent actors to innovate without asking permission from a central platform, and that the same logic applies to AI markets now. As major AI providers race to lock users in through proprietary context stores and metered API access, O'Reilly argues that open protocols and auditable telemetry are the mechanism by which a more competitive and permissionless AI ecosystem becomes possible. He has presented this argument at venues including Oxford University, the Wharton School, and to government policy teams in the U.S. and Europe.



Serkan Piantino


Independent and Venture Partner at Floating Point Advisors; former VP Product, Reddit

Serkan Piantino studied computer science at Carnegie Mellon before joining Facebook in 2007, where he spent nearly a decade working on some of the platform's most consequential products. Shortly after starting, he developed EdgeRank, the algorithm that determined what content users saw in News Feed, a system Wired later described as the most important invention in the history of the social web. He subsequently led the development of Facebook Timeline, which went from hackathon concept to shipped product in under six months with a team of the company's top engineers, and played a central role in building Facebook Messenger, which also originated at a hackathon and became one of the platform's most widely used features.

In 2012 Piantino moved to New York to build out Facebook's engineering presence there, selecting the Wanamaker Building in Astor Place as the location for an office designed by Frank Gehry. He grew that office from 100 people in a midtown sales space to over 2,000 engineers. During this period he also

co-founded Facebook AI Research in 2013, one of the earliest dedicated industrial AI research labs, at a time when deep learning was still largely an academic pursuit. He served on Mayor Bloomberg's Council on Technology and Innovation and on the board of Tech:NYC, and has been a board member of the Economic Security Project since 2016.

After leaving Facebook he founded Spell in 2016, a machine learning platform designed to make it easier for engineering teams to run large-scale ML experiments without specialized hardware infrastructure. Spell was named a Cool Vendor by Gartner and ranked among the hottest ML startups of 2021 by CRN. Reddit acquired Spell in 2022, bringing Piantino on as VP of Product to lead the company's AI, ML, search, and feed teams, where he focused on closing the gap to state-of-the-art and turning research prototypes into broad, self-serve capabilities across the platform. He is now a Venture Partner at Floating Point Advisors, a fund focused on bringing technology to complex, legacy industries, alongside other venture partners including Evan Moore, co-founder of DoorDash, and physicians and operators from healthcare and insurance.



Sruly Rosenblat

LLM Researcher, AI Disclosures Project

Sruly Rosenblat is a researcher for the AI Disclosures Project. He holds a Computer Science degree from Hunter College. He has worked on four research papers covering topics ranging from membership inference attacks to LLM citation analysis, forthcoming in *Ethics & AI*, *Data & Policy*, and other leading interdisciplinary journals. His work examining LLM pre-training data has been covered by TechCrunch, Fast Company, and The Register.

He is also published in O'Reilly Radar, Towards Data Science, and Asimov's Addendum. These articles often explore specific gaps in the AI ecosystem, covering topics such as the lack of standardization in chat templates, how thin and concentrated AI markets create systemic risks, memory architectures, and how limited controls available to developers that build on top of hosted LLMs increase user risk.



Ido Salomon

Co-Creator & Maintainer, MCP Apps/UI

Ido Salomon is the creator of MCP-UI and AgentCraft, and the co-creator of GitMCP and MCP Apps. He serves on the MCP Steering Committee as a maintainer of MCP Apps. Through his work as a software architect, Ido has helped shape how humans, agents, and the web interact.

Previously, Ido was chief architect and an inventor of the enterprise browser at Talon Cyber Security. Building on that foundation, he pioneered browser-native AI both independently and within Palo Alto Networks. In parallel, he became a core contributor to the MCP ecosystem through open-source work, starting with GitMCP — one of the early MCP servers that made web content accessible to agents.

MCP-UI pioneered universal agentic UI and the emerging agentic web paradigm, forming the foundation for MCP Apps, built alongside Anthropic and OpenAI. MCP Apps is now the official MCP extension for interactive applications, supported by hosts including ChatGPT, Claude, and Copilot. Following GitMCP joining [monday.com](https://www.monday.com), Ido led AI and MCP initiatives in the CEO's office, helping drive the transition from traditional SaaS interfaces toward agent-native products.

Today, Ido focuses on the future of human-agent interaction at scale. As agents become first-class citizens, the foundations built now will determine how much control people retain over their connection to the world, and whether the gap between the technically fluent and everyone else widens or narrows. Through his work on the agentic web, he is building the protocols and mechanisms that preserve user

agency while increasing agent autonomy. Through AgentCraft, he is exploring how AI itself can level the playing field, turning agent orchestration into a capability available to everyone.



Ilan Strauss

Co-Director, AI Disclosures Project

Ilan Strauss is Co-Director of the AI Disclosures Project (housed at Code for Science and Society), with Tim O'Reilly (founder) since 2024. He holds a PhD in economics from the New School for Social Research and an Msc in Development Economics from SOAS (University of London). He has taught macroeconomics at NYU and Rice University. His academic work spans digital platform economics, Bayesian panel estimation, corporate investment behavior, and international economics.

Ilan helped establish the National Minimum Wage Research Initiative in South Africa which used research and advocacy to pave the way for the first national minimum wage, legislated on May 29, 2018.

Before the AI Disclosures Project, Strauss was a senior research associate at University College London's Institute for Innovation and Public Purpose (IIPP), where he led digital economy research with Mariana Mazzucato and Tim O'Reilly on a multi-year project funded by the Omidyar Network on what Tim O'Reilly coined "algorithmic attention rents". He remains an Honorary Senior Fellow at UCL IIPP and is a Visiting Associate Professor at the University of Johannesburg. He has received funding from the Economic Security Project to investigate Big Tech's acquisition of technological capabilities (published in Applied Economic Letters).

He is currently working on a book exploring why AI's employment impacts are conditional on macroeconomic distributional relationships. He is interested in understanding further how coding architectures shape market structures.



Liad Yosef

Co-founder & CTO, Era Labs; Co-creator, MCP Apps

Liad Yosef is the co-founder and CTO of Era Labs, the co-creator of MCP Apps and GitMCP, and the co-builder of the MCP-UI project. He is an AI lead and software architect with 20+ years of experience across web development, user interfaces, and human-AI interactions. Previously the architect of Shopify's agentic commerce solution, Liad now works at the frontier of the Human-Agentic Web, defining the transition from human-operated interfaces to agent-driven systems.

At Era Labs, Liad is building the interaction and trust layers for a world where AI agents, not browsers, mediate how people access and engage with digital services. His work advances a clear and controversial thesis: that traditional websites and standalone applications are rapidly becoming obsolete, giving way to dynamic, agent-native experiences delivered directly within conversational and autonomous workflows. As these interfaces fade, the future of the digital economy will depend on seamless, trustworthy, and expressive experiences.

Liad is a core contributor to the Model Context Protocol (MCP) ecosystem, serving as a maintainer in its steering committee and co-leading the MCP Apps working group alongside collaborators from OpenAI and Anthropic. He is the co-creator of MCP Apps: the emerging standard for interactive agentic interfaces, now adopted by leading companies including OpenAI, Anthropic, Microsoft, Cursor, and others.

Prior to that, Liad co-built and maintained MCP-UI, which pioneered the vision of interactive experiences inside agentic flows, and paved the way to the current paradigm of the agentic user interfaces and the new agentic web. He is also the co-creator of the widely-used project GitMCP, which contributed to the shift from static APIs toward live, agent-accessible environments.

Prior to founding Era Labs, Liad led AI and MCP infrastructure initiatives within the CEO offices of Shopify and monday.com, following the acquisition of his previous project. During his tenure at Shopify, he invented and built the agentic storefronts solution from the ground up, pioneering how commerce functions in an agent-first world.

Liad holds degrees in mathematics, physics, computer science, and behavioral economics, shaping a perspective that bridges deep technical systems with human behavior and decision-making.

Alongside his technical work, Liad is a published poet and an analog astronaut for the European Space Agency (ESA). This intersection of human experience and technical rigor informs his work on how systems built for human navigation evolve into systems designed for machine interpretation without losing trust, identity, or meaning.

Person 20 (MIT Scholar)

Bio not shared