

Regulating Agentic AI and Agent-to-Agent Interaction:

Teleperson Research
May 2026

Abstract

The rapid commercialization of "agentic" artificial intelligence (AI) — systems that plan, use tools, transact, and increasingly communicate with other AI agents — has outpaced the existing legal architecture designed for narrow, supervised models. This paper surveys the present state of agentic AI regulation across the European Union, the United States (federal and state), the United Kingdom, China, and other major jurisdictions, alongside the principal international instruments and technical protocols that mediate agent-to-agent (A2A) interaction. Drawing on primary legal sources, scholarly literature, and recent industry standards, we map a fragmented but identifiable terrain: a comprehensive risk-based framework in the EU, a deregulatory federal posture in the United States offset by a thickening patchwork of state laws, a sectoral and pro-innovation stance in the United Kingdom, and a comparatively prescriptive approach in China. We argue that the most pressing legal gaps are not in defining "AI" generally but in five interlocking issues raised distinctively by autonomous agents: liability allocation, contract formation by electronic agents, identity and provenance of agents in machine-to-machine markets, antitrust risks of algorithmic coordination, and consumer-protection disclosure when agents transact on behalf of humans. The paper concludes with a positioning assessment of where the market stands as of May 2026 and a research agenda for legal scholarship that grapples with multi-agent, principal-less systems rather than single-model deployments.

I. Introduction

Within twenty-four months of the public release of the first widely capable general-purpose AI assistants, software described as "agentic" — systems that decompose goals, plan multi-step actions, call external tools, and increasingly negotiate with other autonomous systems — has migrated from research papers to production deployments in commerce, software engineering, customer service, and enterprise back office. Industry frameworks now distinguish "assistive" AI, which acts within a single conversational turn under close human supervision, from "agentic" AI, which executes extended task chains with limited oversight, and from emerging "multi-agent" architectures in which agents coordinate or transact with other agents largely without a human in the loop.¹

The shift is not merely quantitative. Two design choices in particular distinguish the agentic generation: persistent memory and tool-use across sessions, and standardized agent-to-agent (A2A) communication. Anthropic's Model Context Protocol (MCP), released in November 2024, and Google's Agent2Agent Protocol (A2A), released in April 2025, exemplify a new layer of infrastructure designed to make heterogeneous agents discoverable, composable, and capable of completing tasks delegated by other agents.² The protocols are agnostic about who is on either end of the wire: an agent acting for a consumer, an agent acting for an employer, an agent acting for a counterparty, or an agent acting for itself in a long-running automated workflow. As Noam Kolt observes, this raises a governance problem that conventional AI law — focused on developers, deployers, and end users of monolithic models — does not yet address:

¹Jason Gabriel et al., The Ethics of Advanced AI Assistants, GOOGLE DEEPMIND (Apr. 19, 2024), [Source](#)

²See Anthropic, Introducing the Model Context Protocol (Nov. 25, 2024), [Source](#); Google, Announcing the Agent2Agent Protocol (A2A) (Apr. 9, 2025), [Source](#)

agent law must allocate authority, accountability, and rights of audit across delegation chains that begin with a human principal and end with an action taken by software.³

Regulators have responded along two tracks. The first is to extend existing horizontal frameworks designed for narrower AI systems. The European Union's Artificial Intelligence Act (EU AI Act), the United States' sectoral enforcement under FTC, SEC, EEOC, and CFPB authority, the United Kingdom's pro-innovation regulator-led approach, and China's overlapping CAC measures all predate the agentic moment but are being interpreted to reach agent deployments. The second track is the development of agent-specific guidance. The U.S. National Institute of Standards and Technology (NIST) issued a Generative AI Profile of its AI Risk Management Framework in July 2024,⁴ and in February 2026 NIST's newly created Center for AI Standards and Innovation (CAISI) opened a formal AI Agent Standards Initiative seeking comment on agent-specific cybersecurity, governance, and incident-response questions.⁵ Both tracks face a common difficulty: the legal categories that anchor most existing rules — "provider," "deployer," "user," "data subject," "consumer," "principal," "agent" — assume a single, identifiable human at one end of the system. Agentic AI systematically blurs that assumption.

This paper is positioned as a literature review and a positioning assessment. Part II clarifies what is meant by "agentic AI" and "agent-to-agent interaction" and distinguishes these terms from related concepts. Part III surveys the existing regulatory landscape across major jurisdictions. Part IV identifies the cross-cutting legal issues raised distinctively by agent autonomy and A2A communication: liability, contract formation, agent identity, antitrust, consumer protection, and data protection. Part V reviews soft-law instruments and the technical protocols that increasingly function as de facto governance. Part VI summarizes where the market stands as of May 2026 and Part VII offers a short research agenda. The paper does not advocate for a single regulatory model; rather, it maps the field, identifies the points of greatest legal uncertainty, and indicates the doctrinal materials likely to be deployed as cases and enforcement actions accumulate.

II. Defining Agentic AI and Agent-to-Agent Interaction

Neither "agentic AI" nor "agent-to-agent interaction" is a defined term in any binding instrument as of May 2026. The EU AI Act, for instance, defines an "AI system" generically as "a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions, that can influence physical or virtual environments."⁶ The phrase "varying levels of autonomy" anticipates agentic behavior but does not single it out. Industry usage and recent scholarship converge on three working markers.

First, an agentic AI system pursues a goal across multiple steps and external tool calls, rather than producing a single output for a single prompt. Goal-directedness, planning, and tool use are the operational

³Noam Kolt, *Governing AI Agents*, 100 NOTRE DAME L. REV. (forthcoming 2025), [Source](#)

⁴NAT'L INST. STANDARDS & TECH., U.S. DEP'T OF COMMERCE, NIST AI 600-1, ARTIFICIAL INTELLIGENCE RISK MANAGEMENT FRAMEWORK: GENERATIVE ARTIFICIAL INTELLIGENCE PROFILE (July 2024), [Source](#)

⁵Notice of Inquiry, NIST Center for AI Standards and Innovation, AI Agent Standards Initiative, 91 Fed. Reg. (Feb. 17, 2026), [Source](#)

⁶Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act), 2024 O.J. (L 1689) [hereinafter EU AI Act], art. 3(1) (defining "AI system"), [Source](#)

signatures distinguishing agents from generative models that merely respond. Second, an agent maintains state across interactions: persistent memory, scratchpads, or environment-mediated context that allow it to act on the basis of prior steps. Third — and this is the key issue for this paper — an agent communicates with other software agents through standardized protocols rather than only with a human user. MCP defines an agent-to-tool interface; A2A defines an agent-to-agent interface; both are open specifications adopted by multiple model providers.⁷

Two clarifications matter for the legal analysis. First, "agentic" is a spectrum, not a binary. A system that drafts a calendar invitation and waits for human approval is at one end; a system that autonomously negotiates with another agent over a procurement contract is at the other. Most current commercial deployments — coding assistants, customer-service bots, retrieval-augmented research tools — sit closer to the assistive end and operate within human-approved scopes. The fastest-growing category, however, is "task agents" that complete bounded but multi-step actions (booking, filing, reconciling, refunding) without per-step human approval. Where on this spectrum legal categories such as "automated decision-making" under Article 22 of the General Data Protection Regulation (GDPR) or "consequential decision" under the Colorado AI Act apply is now contested.

Second, "agent-to-agent" interactions encompass three distinct fact patterns. The first is intra-organizational: a sales agent calling an inventory agent, both deployed by and for the same firm. This raises governance questions but few external-facing legal ones. The second is inter-organizational under contract: a buyer's procurement agent communicating with a supplier's order-management agent through an established commercial relationship. UETA's electronic-agent provisions and analogous instruments apply with comparative ease. The third — and most legally novel — is inter-organizational without a pre-existing relationship: an agent representing a consumer interacting with an agent representing a merchant on an open marketplace, or two agents negotiating coverage on a real-time advertising exchange. Here questions of authority, identity, attribution, and remedy are genuinely open.\

III. The Existing Regulatory Landscape: A Comparative Survey

A. European Union

The EU AI Act (Regulation (EU) 2024/1689), which entered into force on 1 August 2024, is the most comprehensive horizontal AI law in any jurisdiction.⁸ It adopts a risk-based architecture. Article 5 prohibits a defined set of practices outright; Article 6 and Annex III classify certain systems as "high-risk" and impose obligations on providers and deployers; Article 50 imposes transparency obligations on certain limited-risk systems; and Articles 51 to 55 establish a separate regime for general-purpose AI (GPAI) models, including a designation of "GPAI models with systemic risk" triggered by training compute exceeding 10^{25} floating-point operations.⁹

The Act's relevance to agentic AI is felt at three layers. At the GPAI layer, the obligations on providers of foundation models — including detailed technical documentation, copyright compliance, training-data summaries, and, for systemic-risk models, additional adversarial testing and incident reporting — apply

⁷Anthropic, Introducing the Model Context Protocol (Nov. 25, 2024), [Source](#); Model Context Protocol Specification, [Source](#)

⁸EU AI Act, supra, [Source](#) (consolidated text, with phased application dates running through 2 August 2027).

⁹EU AI Act, supra, art. 5 (prohibited practices); art. 6 & Annex III (high-risk classification); arts. 50, 51–55 (transparency and general-purpose AI obligations).

to the models that power most agent products. These obligations took effect on 2 August 2025 for newly placed models, with a longer transition for legacy models running until 2 August 2027.¹⁰ The GPAI threshold of 10^{25} FLOPs is, as of May 2026, met by several frontier models that serve as the substrate for commercial agents.¹¹

At the high-risk layer, autonomous agent deployments will frequently fall within Annex III categories (employment, access to essential services, law enforcement, migration, education) and so attract obligations on risk management, technical documentation, logging, transparency, human oversight, accuracy and robustness, and conformity assessment. Article 14, in particular, requires that high-risk systems be "designed and developed in such a way ... that they can be effectively overseen by natural persons" — a requirement in tension with autonomous, multi-step agent behavior and especially with autonomous A2A coordination on consequential decisions.¹²

At the transparency layer, Article 50 requires that natural persons interacting with an AI system be informed of that fact and that AI-generated or manipulated content be marked in a machine-readable manner. These obligations take effect on 2 August 2026.¹³ A draft Code of Practice on the marking and labelling of AI-generated content, prepared with the European AI Office, is in consultation as of early 2026.¹⁴ For agent-to-agent commerce, Article 50 creates a difficult problem: when a consumer's agent negotiates with a merchant's agent, neither end of the wire is a natural person, but the eventual user of the resulting transaction is. Whether the obligation runs to the human principal, the counter-party agent, or both is unresolved.

The civil-liability complement to the AI Act has had a more troubled trajectory. The proposed AI Liability Directive, designed to ease evidentiary burdens for plaintiffs in fault-based AI claims, was withdrawn by the European Commission in February 2025 and formally removed from the legislative agenda in October 2025.¹⁵ In its place, the Revised Product Liability Directive — Directive (EU) 2024/2853, adopted on 23 October 2024 — extends strict liability to "products" that explicitly include "software, including AI systems and AI models," and modernizes the concept of "defect" to cover post

¹⁰Commission Communication, Guidelines for Providers of General-Purpose AI Models, COM(2025) (July 18, 2025), <https://digital-strategy.ec.europa.eu/en/policies/guidelines-gpai-providers>; see also DLA Piper, Latest Wave of Obligations Under the EU AI Act Take Effect (Aug. 2025), <https://www.dlapiper.com/en-us/insights/publications/2025/08/latest-wave-of-obligations-under-the-eu-ai-act-take-effect>.

¹¹EU AI Act, *supra*, art. 51(2) (presumption of systemic risk for GPAI models trained using more than 10^{25} floating-point operations).

¹²EU AI Act, *supra*, art. 14 (human oversight requirements for high-risk AI systems).

¹³EU AI Act, *supra*, art. 50(1).

¹⁴European Commission, Code of Practice on Marking and Labelling of AI-Generated Content (consultation draft 2026), <https://digital-strategy.ec.europa.eu/en/policies/code-practice-ai-generated-content>.

¹⁵European Commission, Withdrawal of the Proposal for an AI Liability Directive, 2025 Commission Work Programme, Annex IV (Feb. 11, 2025); Notice of Withdrawal, 2025 O.J. (Oct. 6, 2025); see also EUR. PARL., LEGISLATIVE TRAIN, AI Liability Directive, <https://www.europarl.europa.eu/legislative-train/theme-a-europe-fit-for-the-digital-age/file-ai-liability-directive> (last visited May 2026).

market changes such as updates and continued learning.¹⁶ For agentic AI deployments, the Revised PLD shifts much of the litigation pressure onto a strict-liability frame for products on the market, while the withdrawn AI Liability Directive's fault-based reform must now be reconstructed through Member State law.¹⁷

Data protection law continues to apply in parallel. The European Data Protection Board's Opinion 28/2024 sets out the EDPB's view on the processing of personal data in the context of AI models, including issues of legitimate-interest legal bases, anonymization claims for trained models, and downstream-deployer obligations.¹⁸ For agents that read and write personal data across organizational boundaries, the Opinion implies that each delegation step may constitute a discrete processing operation requiring its own legal basis and, where Article 22 GDPR is engaged, a right of human intervention.

B. United States: Federal Landscape

The federal AI policy posture of the United States changed sharply in early 2025. Executive Order 14179, "Removing Barriers to American Leadership in Artificial Intelligence," signed January 23, 2025, rescinded the Biden Administration's Executive Order 14110 and directed agencies to revise or withdraw guidance perceived as imposing "barriers" to AI development.¹⁹ The America's AI Action Plan, released in July 2025, elaborated a strategy emphasizing infrastructure, export competitiveness, federal procurement, and workforce.²⁰ The Office of Management and Budget reissued its AI guidance to federal agencies in OMB Memos M-25-21 and M-25-22 (April 2025), retaining a framework of impact assessments and human oversight for "high-impact" federal uses while emphasizing acquisition agility.²¹

Below the executive level, the federal regime remains sectoral. NIST's AI Risk Management Framework (AI RMF 1.0) and its Generative AI Profile remain the principal soft-law touchstones; the

¹⁶Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on Liability for Defective Products, 2024 O.J. (L 2853) [hereinafter Revised PLD], <https://eur-lex.europa.eu/eli/dir/2024/2853/oj>.

¹⁷See Revised PLD, *supra*, arts. 4(1), 7 (treating "software, including AI systems and AI models" as "products" subject to strict liability).

¹⁸European Data Protection Board, Opinion 28/2024 on Certain Data Protection Aspects Related to the Processing of Personal Data in the Context of AI Models (Dec. 17, 2024), https://www.edpb.europa.eu/our-work-tools/our-documents/opinion-board-art-64/opinion-282024-certain-data-protection-aspects_en.

¹⁹Exec. Order No. 14179, Removing Barriers to American Leadership in Artificial Intelligence, 90 Fed. Reg. 8741 (Jan. 23, 2025), <https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/>.

²⁰Exec. Office of the President, America's AI Action Plan (July 2025), <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>.

²¹OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, M-25-21, ACCELERATING FEDERAL USE OF AI THROUGH INNOVATION, GOVERNANCE, AND PUBLIC TRUST (Apr. 3, 2025), <https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf>; OFFICE OF MGMT. & BUDGET, M-25-22, DRIVING EFFICIENT ACQUISITION OF ARTIFICIAL INTELLIGENCE IN GOVERNMENT (Apr. 3, 2025), <https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf>.

Texas Responsible AI Governance Act (discussed below) explicitly imports compliance with the NIST framework as a safe harbor.²² The Federal Trade Commission has used Section 5 of the FTC Act to police deceptive AI claims under "Operation AI Comply," including against firms that overstated agent capabilities or deployed undisclosed AI in customer interactions.²³ The Securities and Exchange Commission has brought "AI washing" cases. The Equal Employment Opportunity Commission has prioritized AI in hiring as a Strategic Enforcement Plan area, and the Consumer Financial Protection Bureau has issued guidance applying fair-lending and adverse-action rules to algorithmic underwriting. The Food and Drug Administration's pre-existing AI/ML SaMD framework continues to govern medical-device AI.

Two structural points are noteworthy. First, despite the Administration's pro-innovation posture, the FTC and other enforcement agencies have not retreated from agentic-AI scrutiny; rather, they have continued to apply existing consumer-protection and securities authorities to new agent products. Second, an attempted federal preemption of state AI laws — packaged as a ten-year moratorium in the House version of the One Big Beautiful Bill Act — failed in the Senate. The Senate adopted Senator Blackburn's amendment striking the moratorium provision, and the bill was signed into law as Public Law 119-21 on July 4, 2025 without any federal AI preemption.²⁴ State AI laws therefore remain in force, and a renewed federal preemption effort, including through executive action, is under active consideration as of May 2026.

The most recent federal initiative directly addressing agentic AI is the NIST AI Agent Standards Initiative, opened for public input on February 17, 2026. The Initiative identifies seven domains of agent governance — including delegation chain accountability, runtime behavioral governance, tool-use risk, and inter-agent communication security — and signals a forthcoming agent-specific profile of the AI RMF.²⁵ Whether and how this voluntary technical work translates into binding obligations will depend on developments in the executive and judicial branches over the coming months.

C. United States: State-Level Patchwork

In the absence of comprehensive federal AI legislation, state lawmakers and agencies have moved to fill the void, producing what is now an unmistakably fragmented landscape. Three patterns predominate.

The first is the comprehensive risk-based statute on the EU model. The Colorado Artificial Intelligence Act (Senate Bill 24-205), enacted in 2024, defines a "high-risk" AI system as one used to make or as a substantial factor in a "consequential decision" — including in education, employment, financial services, government services, healthcare, housing, insurance, or legal services — and imposes duties of care,

²²NAT'L INST. STANDARDS & TECH., U.S. DEP'T OF COMMERCE, NIST AI 100-1, ARTIFICIAL INTELLIGENCE RISK MANAGEMENT FRAMEWORK (AI RMF 1.0) (Jan. 2023), <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.

²³See Federal Trade Commission, Operation AI Comply: Continuing the Crackdown on Deceptive AI Claims and Schemes, FTC.GOV (Sept. 25, 2024), <https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes>.

²⁴One Big Beautiful Bill Act, Pub. L. No. 119-21 (July 4, 2025); see Goodwin Procter LLP, Federal AI Moratorium Dies on the Vine as Senate Passes the One Big Beautiful Bill Act (July 2025), <https://www.goodwinlaw.com/en/insights/publications/2025/07/alerts-practices-aiml-federal-ai-moratorium-dies-on-the-vine>.

impact assessments, consumer notice, and rights to correction or appeal on developers and deployers.²⁶ The Act's effective date, originally February 1, 2026, was postponed to June 30, 2026 by Senate Bill 25B-004; substantial amendments are expected in the 2026 regular session, and a federal court has paused enforcement pending preliminary-injunction proceedings.²⁷ Texas followed in June 2025 with a substantially narrower Responsible AI Governance Act (House Bill 149), focused on intentional discrimination and certain prohibited harmful uses, with a regulatory sandbox at the Department of Information Resources and a safe harbor for compliance with the NIST AI RMF; the Texas law takes effect January 1, 2026.²⁸ Together, Colorado and Texas illustrate competing models — one closer to the EU AI Act, one closer to a sectoral fairness rule.²⁹

The second pattern is sectoral. New York City's Local Law 144 has required bias audits for automated employment-decision tools (AEDTs) since July 2023.³⁰ Illinois, Maryland, and several other states have analogous employment provisions. California's Privacy Protection Agency finalized regulations in 2025 governing "automated decisionmaking technology" (ADMT) under the California Consumer Privacy Act, including notice and opt-out rights for certain consequential automated decisions.³¹ Utah's AI Policy Act (Senate Bill 149, 2024) imposes mandatory disclosure when consumers interact with generative AI in regulated occupations.³²

The third pattern is the targeted prohibition. California Senate Bill 1047, which would have imposed safety obligations on "covered models" exceeding specified compute and cost thresholds, was vetoed by Governor Newsom in September 2024 in part on the ground that compute thresholds did not capture the most relevant risks.³³ Tennessee, Florida, and Minnesota have enacted deepfake- and synthetic-media-specific laws that apply to certain agent-generated content. The cumulative effect is a patchwork that imposes real compliance costs on multi-state agent deployers but that is also rapidly evolving and, in several jurisdictions, currently under judicial or administrative reconsideration.

²⁶Colorado Artificial Intelligence Act, S.B. 24-205, 74th Gen. Assemb., 2d Reg. Sess. (Colo. 2024) (codified at Colo. Rev. Stat. §§ 6-1-1701 to -1707), <https://leg.colorado.gov/bills/sb24-205>.

²⁷See Colo. S.B. 25B-004, 75th Gen. Assemb., 1st Extra. Sess. (Colo. 2025) (delaying effective date to June 30, 2026), <https://leg.colorado.gov/bills/sb25b-004>; see also Akin Gump, Colorado Postpones Implementation of Colorado AI Act (Sept. 2025), <https://www.akingump.com/en/insights/ai-law-and-regulation-tracker/colorado-postpones-implementation-of-colorado-ai-act-sb-24-205>.

²⁸Texas Responsible Artificial Intelligence Governance Act, H.B. 149, 89th Leg., R.S. (Tex. 2025) (effective Jan. 1, 2026), <https://capitol.texas.gov/tlodocs/89R/billtext/pdf/HB00149F.pdf>.

²⁹See K&L Gates, Pared-Back Version of the Texas Responsible Artificial Intelligence Governance Act Signed Into Law (June 24, 2025), <https://www.klgates.com/Pared-Back-Version-of-the-Texas-Responsible-Artificial-Intelligence-Governance-Act-Signed-Into-Law-6-24-2025>.

³⁰N.Y.C. Admin. Code §§ 20-870 to -874 (Local Law 144 of 2021, Automated Employment Decision Tools, eff. Jan. 1, 2023, enforcement July 5, 2023), <https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page>.

³¹See Cal. Code Regs. tit. 11, §§ 7000–7304 (CPPA Regulations on Automated Decisionmaking Technology) (effective 2025); CAL. PRIVACY PROT. AGENCY, ADMT FINAL RULES (2025), <https://cppa.ca.gov/regulations/>.

³²Utah Artificial Intelligence Policy Act, S.B. 149, 2024 Gen. Sess. (Utah 2024) (codified at Utah Code §§ 13-2-12 to -13), <https://le.utah.gov/~2024/bills/static/SB0149.html>.

³³Cal. S.B. 1047, 2023–2024 Reg. Sess. (Cal. 2024) (vetoed by Governor Newsom Sept. 29, 2024), https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202320240SB1047.

D. United Kingdom

The United Kingdom's 2023 AI white paper articulated a "pro-innovation" approach that delegates AI rule-making to existing sectoral regulators — the Information Commissioner's Office (ICO), the Financial Conduct Authority (FCA), the Competition and Markets Authority (CMA), the Equality and Human Rights Commission, and others — rather than enacting a horizontal AI statute on the EU model.³⁴ That approach has continued through the change of government in 2024. The AI Safety Institute, established in 2023, was rebranded the AI Security Institute in 2025; its mandate is testing and evaluation of frontier and agentic systems rather than regulation as such.³⁵

The ICO's Guidance on AI and Data Protection — most recently updated in early 2025 — sets out the ICO's expectations for fairness, transparency, and accountability under the UK GDPR, including for automated decision-making under Article 22.³⁶ For agentic systems, the implication is that delegated agent-to-agent processing must be analysed step by step, with controllers identified and legal bases recorded at each delegation. The CMA has published a strategy on AI foundation models and has signalled scrutiny of self-preferencing and tying behavior in agent platforms; the FCA has continued to develop sector-specific guidance on algorithmic and AI-driven financial services.

E. China

China's approach is more prescriptive than any major Western regime as of May 2026, but is layered across overlapping instruments rather than consolidated in a single horizontal statute. The Cyberspace Administration of China (CAC) and other ministries jointly issued the Interim Measures for the Management of Generative AI Services (effective August 15, 2023), which imposes content-moderation, registration, and security-assessment obligations on providers of public-facing generative AI services in China.³⁷ Earlier instruments — the Algorithmic Recommendation Provisions (effective March 1, 2022)³⁸ and the Deep Synthesis Provisions (effective January 10, 2023)³⁹ — already required algorithm filings,

³⁴Department for Science, Innovation and Technology (UK), A Pro-Innovation Approach to AI Regulation, CP 815 (Mar. 29, 2023), <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach>.

³⁵See UK AI Security Institute (formerly AI Safety Institute), <https://www.aisi.gov.uk/> (rebranded 2025).

³⁶Information Commissioner's Office (UK), Guidance on AI and Data Protection (Mar. 2023, updated Jan. 2025), <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/>.

³⁷Cyberspace Administration of China et al., Interim Measures for the Management of Generative Artificial Intelligence Services [生成式人工智能服务管理暂行办法] (effective Aug. 15, 2023), unofficial English translation at <https://digichina.stanford.edu/work/translation-measures-for-the-management-of-generative-artificial-intelligence-services-effective-august-15-2023/>.

³⁸Cyberspace Administration of China, Provisions on the Administration of Algorithmic Recommendations of Internet Information Services (effective Mar. 1, 2022), <https://digichina.stanford.edu/work/translation-internet-information-service-algorithmic-recommendation-management-provisions-effective-march-1-2022/>.

³⁹Cyberspace Administration of China, Provisions on the Administration of Deep Synthesis Internet Information Services (effective Jan. 10, 2023), <https://digichina.stanford.edu/work/translation-provisions-on-the-administration-of-deep-synthesis-internet-information-services/>.

opt-out rights for personalized recommendations, and labelling of synthetic media. As of late 2025, more than 700 generative AI services had completed national-level filings.

In September 2025, China's Measures for Labeling AI-Generated Synthetic Content took effect, requiring both implicit (metadata) and, in many cases, explicit (visible) labelling of AI-generated text, image, audio, and video.⁴⁰ Draft ethics rules and a forthcoming comprehensive AI Law are under consultation, with provisions reportedly addressing autonomous agents and high-impact applications.⁴¹ The cumulative regulatory pressure on agent deployments in China is high: agentic services aimed at the Chinese public must clear filings, embed values-aligned moderation, label outputs, and operate under enforcement that is more administrative and less litigated than in the EU or the United States.

F. Other Asia-Pacific and Latin American Jurisdictions

South Korea's Framework Act on the Development of AI and the Establishment of Trust ("AI Basic Act"), promulgated on January 21, 2025, takes effect on January 22, 2026 and establishes a risk-based framework with particular obligations for "high-impact" and "generative" AI, transparency duties, and a national AI safety institute.⁴² Japan adopted a separate AI bill in 2025, oriented toward promotion and voluntary guidelines, complemented by the Ministry of Economy, Trade and Industry's AI Guidelines for Business (most recently revised March 28, 2025).⁴³ Singapore continues to develop the Model AI Governance Framework, including a 2024 module addressed specifically to generative AI.⁴⁴

In Latin America, Brazil's draft AI Marco Legal (Bill 2338/2023) was approved by the Senate in December 2024 and is pending in the Chamber of Deputies; it largely mirrors the EU's risk-based architecture and adopts strict liability for high-risk systems.⁴⁵ Canada's proposed Artificial Intelligence and Data Act (AIDA), embedded in Bill C-27, died on the order paper when Parliament was prorogued in

⁴⁰Cyberspace Administration of China, Measures for Labeling AI-Generated Synthetic Content (effective Sept. 1, 2025), <https://www.cac.gov.cn/>.

⁴¹See Mayer Brown, China Formulates New AI Global Governance Action Plan and Issues Draft Ethics Rules and AI Labelling Rules (Oct. 2025), <https://www.mayerbrown.com/en/insights/publications/2025/10/artificial-intelligence-a-brave-new-world-china-formulates-new-ai-global--governance-action-plan-and-issues-draft-ethics-rules-and-ai-labelling-rules>.

⁴²Framework Act on the Development of Artificial Intelligence and the Establishment of Trust [AI Basic Act], Act No. 20762 (Republic of Korea, promulgated Jan. 21, 2025; effective Jan. 22, 2026), <https://aibasicact.kr/>.

⁴³Bill on the Promotion of Research, Development, and Use of AI-Related Technologies, Act No. ____ (Japan 2025); see Ministry of Economy, Trade & Industry, AI Guidelines for Business v. 1.1 (Mar. 28, 2025), https://www.meti.go.jp/english/press/2025/0328_001.html.

⁴⁴Personal Data Protection Commission (Singapore), Model AI Governance Framework for Generative AI (May 30, 2024), <https://www.pdpc.gov.sg/help-and-resources/2024/05/model-ai-governance-framework-for-generative-ai>.

⁴⁵Projeto de Lei nº 2.338, de 2023 (Marco Legal da Inteligência Artificial), Senado Federal (Brazil) (approved by Senate Dec. 10, 2024; pending in Chamber of Deputies), <https://www25.senado.leg.br/web/atividade/materias/-/materia/157233>.

early 2025 and has not yet been reintroduced; Canadian AI regulation in the interim relies on PIPEDA, provincial privacy law, and sectoral guidance.⁴⁶

G. International Instruments

Two international instruments warrant particular attention. The Council of Europe's Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law was opened for signature on September 5, 2024 and is the first legally binding international treaty on AI.⁴⁷ Its substantive obligations focus on human rights, democracy, and the rule of law, requiring parties to take measures to identify, assess, and mitigate the risks of AI systems within their jurisdiction. Its early signatories include Andorra, Georgia, Iceland, Norway, the Republic of Moldova, San Marino, the United Kingdom, Israel, the United States, and the European Union; entry into force awaits ratification by five signatories including at least three Council of Europe member states.⁴⁸

At the soft-law level, the OECD's AI Principles, originally adopted in 2019 and revised in May 2024, articulate principles of inclusive growth, human-centered values, transparency, robustness and security, and accountability that have been adopted or endorsed by over forty countries and the G20.⁴⁹ The G7 Hiroshima Process produced an International Code of Conduct for Organizations Developing Advanced AI Systems in October 2023, focused on advanced foundation models.⁵⁰ ISO/IEC 42001:2023 establishes a certifiable AI management-system standard analogous to ISO 27001 for information security; major model developers and several large enterprises have begun pursuing certification.⁵¹ Although none of these instruments specifically governs agentic AI or A2A interaction, all three are increasingly cited in regulator guidance and in private contracts as the relevant baseline of "good practice."

IV. Cross-Cutting Legal Issues for Agentic AI

The horizontal frameworks surveyed above were drafted with bounded AI systems in mind: a hiring screener, a credit model, a generative chatbot. Six interlocking issues take on a different shape once systems are agentic and increasingly engage in agent-to-agent interaction. We address them in turn.

⁴⁶Bill C-27, An Act to Enact the Consumer Privacy Protection Act, the Personal Information and Data Protection Tribunal Act and the Artificial Intelligence and Data Act, 1st Sess., 44th Parl. (Can. 2022) (died on order paper, Jan. 2025), <https://www.parl.ca/DocumentViewer/en/44-1/bill/C-27/first-reading>.

⁴⁷Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, opened for signature Sept. 5, 2024, C.E.T.S. No. 225 [hereinafter CoE AI Convention], <https://rm.coe.int/1680afae3c>.

⁴⁸CoE AI Convention, supra, arts. 1, 4–7; see Council of Europe, Council of Europe Opens First Ever Global Treaty on AI for Signature (Sept. 5, 2024), <https://www.coe.int/en/web/portal/-/council-of-europe-opens-first-ever-global-treaty-on-ai-for-signature>.

⁴⁹OECD, Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449 (adopted May 22, 2019; revised May 2024), <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>.

⁵⁰G7 Hiroshima Process International Code of Conduct for Organizations Developing Advanced AI Systems (Oct. 30, 2023), <https://digital-strategy.ec.europa.eu/en/library/hiroshima-process-international-code-conduct-advanced-ai-systems>.

⁵¹INT'L ORG. FOR STANDARDIZATION & INT'L ELECTROTECHNICAL COMM'N, ISO/IEC 42001:2023, INFORMATION TECHNOLOGY — ARTIFICIAL INTELLIGENCE — MANAGEMENT SYSTEM (Dec. 2023), <https://www.iso.org/standard/81230.html>.

A. Liability Allocation and the "Many Hands" Problem

Allocating responsibility when complex software causes harm is not a new problem. Helen Nissenbaum identified the "many hands" problem in computerized systems as early as 1996,⁵² and Madeleine Clare Elish has analyzed the recurrent tendency to assign blame to the human nearest the failure — the "moral crumple zone" effect. Agentic AI compounds the problem in two ways. First, the chain of action between a human principal and an outcome lengthens: a user instructs an agent, which calls a tool through MCP, which calls another agent through A2A, which composes an action against a third-party API. Each link involves a different vendor and, plausibly, a different jurisdiction. Second, the model itself is increasingly produced by one party (the foundation model provider), wrapped by another (the agent platform), customized by a third (the deployer), and operated for a fourth (the principal end user).

Existing doctrinal materials offer four imperfect responses. Products liability — most cleanly expressed in the EU's Revised Product Liability Directive — assigns strict liability to the producer of a defective "product" that now expressly includes software and AI systems.⁵³ For agentic AI, the open issues are when an integrated agent is a single product or a chain of components, and how the "defect" inquiry handles emergent behaviors arising from agent composition. Vicarious liability and respondeat superior offer a familiar route in employment-like contexts: where an agent acts on behalf of a principal, the principal answers for the agent's conduct. Agency law principles have already been applied to AI vendors in the United States: in *Mobley v. Workday*, a federal district court allowed Title VII employment-discrimination claims against an AI-screening vendor to proceed under an agency theory.⁵⁴ Negligence and duty-of-care theories — including, in the EU AI Act, structured duties to maintain risk-management systems, to log, and to provide for human oversight — preserve fault as the locus but require evidence the plaintiff is often poorly situated to obtain. The withdrawn EU AI Liability Directive would have eased that evidentiary burden through targeted presumptions; its withdrawal in 2025 leaves Member States with the task of evolving such presumptions through litigation or national legislation.⁵⁵

Finally, the regulatory liability route — administrative fines under the EU AI Act, FTC Section 5 actions in the United States, and state attorney-general enforcement under laws like Colorado SB 24-205 — operates in parallel with private litigation. Scholars including Kolt have argued that conventional agency law, properly extended, can provide much of the doctrinal material needed for agentic AI liability, but that statutory recognition of certain agent-specific obligations (audit-log retention, identity disclosure, kill-switches) will be needed to make agency-law remedies tractable.⁵⁶

⁵²See Helen Nissenbaum, *Accountability in a Computerized Society*, 2 *SCI. & ENG'G ETHICS* 25 (1996); see also Madeleine Clare Elish, *Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction*, 5 *ENGAGING SCI., TECH. & SOC'Y* 40 (2019).

⁵⁴See *Mobley v. Workday, Inc.*, 740 F. Supp. 3d 796 (N.D. Cal. 2024) (allowing Title VII action against AI vendor under agency theory), preliminarily certified as collective action, No. 23-cv-00770 (N.D. Cal. May 16, 2025).

B. Contract Formation and the Authority of Electronic Agents

In the United States, the Uniform Electronic Transactions Act (UETA), adopted in 49 states, has anticipated automated contract formation since 1999. Section 14 provides that "[a] contract may be formed by the interaction of electronic agents of the parties, even if no individual was aware of or reviewed the electronic agents' actions or the resulting terms and agreements."⁵⁷ The accompanying definition of "electronic agent" is technology-neutral: "a computer program or an electronic or other automated means used independently to initiate an action or respond to electronic records or performances in whole or in part, without review or action by an individual."⁵⁸ The UETA Drafting Committee in 1999 anticipated, presciently, that "within the useful life of UETA, electronic agents may be created with the ability to act autonomously" and "may be able to 'learn through experience, modify the instructions in their own programs, and even devise new instructions.'"

For agent-to-agent commerce, UETA §14 supplies a baseline of enforceability: a contract concluded between two electronic agents is not voidable merely because no human reviewed the terms. But three doctrinal issues remain open. First, what is the scope of the agent's authority? Werbach and Cornell argue that automated contracts should be assessed against the principal's reasonable expectations, with manifest mistakes outside the scope of authority not enforced.⁵⁹ Second, how should mistake, fraud, and unconscionability doctrines apply to interactions in which one or both ends of the wire is a learning agent that may act unpredictably? Third, what disclosures are owed when a consumer's agent transacts with a merchant's agent — and is the consumer bound by the agent's click-through accepting terms the human never read? UNCITRAL Working Group IV has begun work on these questions internationally, with a 2024 working paper on the use of AI and automation in contracting.⁶⁰

C. Agent Identity, Authentication, and Provenance

Once agents transact with other agents, the question of who is on the other end of a connection becomes both technically nontrivial and legally consequential. Three sub-issues emerge. The first is identification: does the law require an agent to disclose that it is an agent (rather than a human)? The EU AI Act Article 50 answers in the affirmative for human-facing agents, with effect from August 2026.⁶¹ California's 2018 "bot law" already prohibits undisclosed bots in certain contexts. But Article 50 does not by its terms reach agent-to-agent interactions, and U.S. federal law currently does not.

⁵⁷UNIF. ELEC. TRANSACTIONS ACT § 14 (UNIF. LAW COMM'N 1999), <https://www.uniformlaws.org/committees/community-home?CommunityKey=2c04b76c-2b7d-4399-977e-d5876ba7e034>.

⁵⁸UETA § 2(6), *id.* (defining "electronic agent" as "a computer program or an electronic or other automated means used independently to initiate an action or respond to electronic records or performances in whole or in part, without review or action by an individual").

⁵⁹See Kevin Werbach & Nicolas Cornell, *Contracts Ex Machina*, 67 DUKE L.J. 313 (2017), <https://scholarship.law.duke.edu/dlj/vol67/iss2/2/>.

⁶⁰U.N. Comm'n on Int'l Trade Law (UNCITRAL), Working Group IV (Electronic Commerce), *Use of Artificial Intelligence and Automation in Contracting*, U.N. Doc. A/CN.9/WG.IV/WP.180 (2024), https://uncitral.un.org/en/working_groups/4/electronic_commerce.

The second is authentication: by what cryptographic or organizational means does an agent prove that it is acting for a particular principal and within a particular scope of authority? Industry technical work — including the A2A protocol's "Agent Card" mechanism for capability advertisement and OpenAI's and Anthropic's discussions of "agent passports" and signed delegation tokens — proposes various designs, but no jurisdiction has yet adopted authentication requirements as binding law.⁶²

The third is provenance and audit: what records must be kept of an agent's actions, including its inter-agent communications, and to whom must they be made available? The EU AI Act requires logging of high-risk system operations and access by deployers, providers, and competent authorities; ISO/IEC 42001 conditions certification on similar logging and audit practices; the NIST CAISI Initiative explicitly identifies "delegation chain accountability" as a core domain.⁶³ The doctrinal puzzle is to translate these technical practices into evidentiary regimes — evidentiary rules, spoliation doctrines, discovery defaults — that will be litigation-ready when agentic AI cases mature.

D. Antitrust and Algorithmic Coordination

The intuition that algorithms might collude has been a topic in competition law for nearly a decade. Ezrachi and Stucke distinguished four scenarios: messenger collusion (algorithms simply enforce a human-agreed cartel), hub-and-spoke (multiple firms use a common algorithm or platform that effects coordination among them), tacit collusion (independently designed algorithms learn to coordinate), and "digital eye" collusion in transparent markets.⁶⁴ For agentic AI, the salient categories are hub-and-spoke and tacit. The OECD's 2023 Background Note on Algorithmic Competition and its 2024 work on AI, Data and Competition canvass the empirical literature, including studies of pricing-software hub effects in retail gasoline and rent-recommendation tools in U.S. residential housing, and conclude that no fully autonomous tacit-collusion case has yet been documented but that hub-and-spoke risks are real.⁶⁵

Agentic AI changes the prior debate in two ways. First, A2A protocols permit communication between agents that may, on its face, be benign coordination but that could shade into prohibited information exchange. Second, the use of common foundation models — and indeed common agent platforms — across competitors creates structural conditions for hub-and-spoke arrangements. In response, U.S. legislators have introduced the Preventing Algorithmic Collusion Act in the 119th Congress; among other things, the bill would create a presumption that exchanging sensitive competitive information through

⁶²See, e.g., Yonadav Shavit et al., Practices for Governing Agentic AI Systems, OPENAI (Dec. 14, 2023), <https://cdn.openai.com/papers/practices-for-governing-agentic-ai-systems.pdf>.

⁶⁴See ARIEL EZRACHI & MAURICE E. STUCKE, VIRTUAL COMPETITION: THE PROMISE AND PERILS OF THE ALGORITHM-DRIVEN ECONOMY (Harv. Univ. Press 2016); Ariel Ezrachi & Maurice E. Stucke, Artificial Intelligence and Collusion: When Computers Inhibit Competition, 2017 U. ILL. L. REV. 1775.

⁶⁵See OECD, Algorithmic Competition, OECD Competition Policy Roundtable Background Note (May 2023), https://www.oecd.org/content/dam/oecd/en/publications/reports/2023/05/algorithmic-competition_2be02d00/cb3b2075-en.pdf; OECD, AI, Data and Competition (June 2024), https://www.oecd.org/en/publications/artificial-intelligence-data-and-competition_e1d1ad17-en.html.

pricing algorithms constitutes an "agreement" under the Sherman Act.⁶⁶ EU competition authorities have signalled similar concerns under Article 101 TFEU, including in scrutiny of platform self-preferencing.

E. Consumer Protection and Disclosure in Agent-Mediated Commerce

Agent-mediated commerce raises consumer-protection issues that fall imperfectly within existing frameworks. When a consumer authorizes an agent to make purchases up to a budget, who is on the hook if the agent buys items the consumer did not want, accepts terms the consumer would have rejected, or is induced by a hostile counter-party agent into actions the consumer did not authorize? Future of Privacy Forum's 2025 analysis frames these as questions of "who pays when transactional agents play," and identifies consumer-financial-services regulators, the FTC, and state attorneys general as the most active enforcers.⁶⁷

Three doctrinal pressure points are visible. First, disclosure: many U.S. and EU rules require disclosure that a consumer is "interacting with an AI" rather than a human, but few require disclosure when an agent is interacting on the consumer's behalf with another agent. Second, the unauthorized-transaction doctrines familiar from credit-card and electronic-funds-transfer law have natural extensions to agent transactions but have not yet been adapted by statute or rule. Third, dark patterns and manipulative design — the subject of FTC and EU enforcement against social-media interfaces — re-emerge in agent-to-agent settings, where one agent might exploit another agent's instruction-following weaknesses.

F. Data Protection and Information Privacy

The GDPR's Article 22 right "not to be subject to a decision based solely on automated processing ... which produces legal effects ... or similarly significantly affects" the data subject is the most-cited European data-protection provision in the AI context. For agentic AI, two issues are acute. First, when agents chain together to produce a consequential outcome, is each intermediate step "automated processing" within Article 22, or only the final adverse decision? The EDPB's 2024 Opinion on AI models hints toward a granular analysis but does not resolve the question for agent chains.⁶⁸ Second, the controller-processor allocation envisaged by the GDPR — and reinforced by joint-controllership doctrine in the wake of cases like *Wirtschaftsakademie* and *Fashion ID* — sits uneasily with multi-vendor agent stacks where the "purposes and means" of processing are dynamically negotiated between agents.

Outside Europe, comparable issues arise under the California Consumer Privacy Act (as amended) and the new ADMT regulations,⁶⁹ under the UK Data Protection Act and the ICO guidance,⁷⁰ and under sectoral rules such as HIPAA in the United States and the Personal Information Protection Law in China.

⁶⁶Preventing Algorithmic Collusion Act of 2025, S. 232, 119th Cong. (2025), <https://www.congress.gov/bill/119th-congress/senate-bill/232>.

⁶⁷See Future of Privacy Forum, *From Chatbot to Checkout: Who Pays When Transactional Agents Play?* (2025), <https://fpf.org/blog/from-chatbot-to-checkout-who-pays-when-transactional-agents-play/>.

The cumulative direction of travel is toward more granular logging, clearer notice, and meaningful opt-outs from automated decision chains — all of which require technical rearchitecture by agent platforms.

V. Soft Law, Standards, and Industry Protocols

The hardness of "law" matters less than its effective gravity. Three categories of non-statutory instruments are exerting increasing influence on agentic AI deployments.

The first is voluntary national frameworks, of which the NIST AI RMF and the Generative AI Profile are paradigmatic.⁷¹ The Texas Responsible AI Governance Act explicitly imports the NIST framework as a safe harbor against certain liability.⁷² In private litigation and regulator inquiries, conformance with the AI RMF — and increasingly with the forthcoming Agentic Profile under the AI Agent Standards Initiative — is being deployed as evidence of reasonable care.⁷³ Although voluntary, these frameworks function as a de facto standard of practice.

The second is international standards, most notably ISO/IEC 42001:2023.⁷⁴ Modeled on ISO 27001, it specifies requirements for an AI management system; certification is now sought by major model developers and by enterprises deploying AI in regulated sectors. The G7 Hiroshima Process Code of Conduct, the OECD AI Principles, and the Council of Europe Framework Convention add further layers, each addressed primarily to states and large developers.⁷⁵ For agentic systems, none of these instruments yet contains agent-specific obligations, but the Hiroshima Code's commitments around security testing and incident reporting, and the CoE Convention's human-rights-impact-assessment requirements, are expected to shape future agent-specific guidance.⁷⁶

The third — and the layer that most directly shapes A2A interaction — is technical protocols. MCP, introduced by Anthropic in November 2024, has been adopted by OpenAI, Google DeepMind, and a wide developer community as a standard for agent-to-tool integration; it specifies, among other things, capability discovery, scope authorization, and error semantics.⁷⁷ Google's A2A protocol, announced in April 2025, layers above MCP and addresses agent-to-agent collaboration: agents publish "Agent Cards" advertising capabilities and constraints, and tasks are exchanged with explicit lifecycles.⁷⁸ These protocols are governance artifacts as much as technical ones. The schemas they require — for capability description, for authorization scopes, for error and incident reporting — operationalize obligations that may eventually be codified in law. As Lessig argued nearly three decades ago, in cyberspace "code is law": the design of these protocols, more than any present statute, will determine what agentic interactions are possible and on what terms.⁷⁹

⁷³See Cloud Security Alliance, NIST AI Risk Management Framework: Agentic Profile v.1 (2026), <https://labs.cloudsecurityalliance.org/agentic/agentic-nist-ai-rmf-profile-v1/>.

⁷⁸Google, Announcing the Agent2Agent Protocol (A2A) (Apr. 9, 2025), <https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interopability/>.

⁷⁹Lawrence Lessig, CODE AND OTHER LAWS OF CYBERSPACE 6 (1999); see Lawrence Lessig, CODE: VERSION 2.0 (2006).

VI. Where the Market Is Today: A Positioning Assessment

Synthesizing the foregoing, four observations describe the position of agentic AI regulation as of May 2026.

First, no jurisdiction has enacted a horizontal "agent law." Every binding rule that reaches agentic AI is a generalized AI rule (the EU AI Act, Colorado SB 24-205, China's Interim Measures), a sectoral rule (NYC LL 144, EEOC guidance, FDA SaMD), a generic technology-neutral rule (UETA §14, the Revised PLD), or a soft-law standard (NIST AI RMF, ISO/IEC 42001). Agent-specific regulation exists, if at all, only as draft or as administrative initiative — most prominently the NIST AI Agent Standards Initiative opened for comment in February 2026.⁸⁰

Second, the regulatory map is bifurcating geographically. The European Union is implementing a comprehensive, ex-ante, risk-based regime — the AI Act, the Revised PLD, and the GDPR/EDPB framework — that will increasingly govern agent deployments through general AI obligations adapted by Commission guidelines and Member State enforcement. China, with a different political-economic logic, has parallel ex-ante obligations focused on content control, registration, and state oversight. The United States, after Executive Order 14179 and the failure of the federal preemption push, has a deregulatory federal posture combined with a thickening state patchwork (Colorado, Texas, California ADMT, NYC LL 144, Utah) and aggressive sectoral enforcement (FTC, SEC, EEOC, CFPB). The United Kingdom and a number of Asia-Pacific jurisdictions occupy intermediate positions, leaning innovation-forward but adopting selective new statutes (Korea's AI Basic Act; Japan's AI Bill).

Third, where the law has not yet spoken, technical protocols and contract are filling the void. MCP and A2A specify, with increasing precision, what an agent can do, how it identifies itself, what scopes it claims, and how it logs its actions. Major AI providers' developer policies — accepted as terms of service by every customer — already contain detailed obligations on agent design, including prohibitions on certain categories of autonomous action. Insurance markets are beginning to price agentic-AI risk separately, although standard policy language remains immature.

Fourth, the litigation pipeline is shallow but indicative. *Mobley v. Workday* is the most-watched U.S. case; class-certification rulings are likely to be persuasive on the application of agency theory to AI vendors.⁸¹ The first generative-AI copyright cases have begun to produce trial-court rulings. Public-enforcement cases, especially under FTC Operation AI Comply,⁸² have produced consent decrees that cumulatively are shaping disclosure expectations. The first agent-to-agent disputes — about authority, mistake, and consequential damages — are not yet on the public docket but are being negotiated in commercial arbitration, which will continue to obscure precedent for some time.

Taken together, the picture is of a market scaling rapidly under a fragmented and partly retrofitted legal framework, with technical standards moving faster than statutes and a credible prospect that several jurisdictions will, within the next twenty-four months, enact agent-specific obligations on identity, logging, authority, and disclosure. Brookings, CSET, and OpenAI policy work — although addressed to different audiences — converge on a core menu of obligations: provenance metadata, signed delegation

tokens, audit logs, scope-bounded execution, and meaningful human-in-the-loop checkpoints for "high-stakes" actions.⁸³

VII. Open Questions and a Research Agenda

Five questions warrant sustained legal-academic attention.

(1) The doctrine of agency, and especially its English- and U.S.-law forms, has been the resource of choice in early agentic-AI scholarship. The agenda set by Kolt — adapting actual, apparent, and ratified authority to electronic agents, and constructing duties of loyalty and care for AI agents toward their principals — invites empirical work on how agents are actually deployed and on the points at which authority breaks down.⁸⁴

(2) The "law informs code" agenda — exemplified by John Nay's Stanford CodeX work — proposes that legal materials be ingested directly into agent training and runtime constraints.⁸⁵ Whether such an approach can scale, and whether it raises new democratic-legitimacy concerns when private firms operationalize legal text, is an open methodological question.

(3) Antitrust analysis of agent ecosystems must move beyond the pricing-algorithm paradigm. The most legally significant question may be whether agent platforms — which combine model-provision, runtime, and identity infrastructure — develop "tipping" dynamics that warrant ex-ante interoperability or non-discrimination obligations on the model of the EU Digital Markets Act.

(4) Cross-border enforcement of agent-related obligations will become urgent as A2A protocols are inherently jurisdictionally mobile. The Council of Europe Framework Convention provides a useful baseline, but ratification is slow and the Convention is principle-level rather than operational.⁸⁶

(5) Finally, the temptation to revisit "electronic personhood" — proposed by the European Parliament in 2017 and rejected by both the academic community and ultimately the Commission — is likely to return as multi-agent systems become more autonomous.⁸⁷ The 2017 critique remains powerful: granting legal personhood would diminish the responsibility of the humans who build, deploy, and profit from these systems.⁸⁸ Sound regulation should keep liability anchored in human and corporate principals while developing the auxiliary doctrines — identity, audit, scope — that make principal-based accountability tractable for agent chains.

VIII. Conclusion

⁸³See Brookings Inst., *Toward a Governance Framework for Agentic AI* (2025), <https://www.brookings.edu/articles/toward-a-governance-framework-for-agentic-ai/>.

⁸⁵See John J. Nay, *Aligning AI Agents with Humans Through Law as Information*, STAN. CODEX (Oct. 2025), <https://law.stanford.edu/wp-content/uploads/2025/10/Aligning-AI-Agents-with-Humans-through-Law-as-Information.pdf>.

⁸⁷European Parliament Resolution of 16 February 2017 with Recommendations to the Commission on Civil Law Rules on Robotics, 2015/2103(INL), ¶ 59(f) (proposing electronic personhood), https://www.europarl.europa.eu/doceo/document/TA-8-2017-0051_EN.html.

⁸⁸Open Letter to the European Commission: Artificial Intelligence and Robotics (Apr. 12, 2018) (signed by more than 150 European AI scientists and legal scholars), <http://www.robotics-openletter.eu/>; see also Nathalie Nevejans, *European Civil Law Rules in Robotics*, Eur. Parl. (2016).

A familiar caution about legal change is Frank Easterbrook’s argument that subject-specific cyberlaw is mostly a distraction from generally applicable doctrines.⁸⁹ For agentic AI, that view is partly vindicated and partly overtaken. Generally applicable doctrines — agency, products liability, contract, antitrust, consumer protection, data protection — supply most of the doctrinal materials needed to address agent harms. But the architecture of agentic systems — chained delegation, machine-to-machine communication on open protocols, persistent state across organizational boundaries — generates problems that those doctrines were not designed to identify, let alone resolve. The most useful regulatory contributions of the next several years will not be new horizontal AI statutes but auxiliary obligations that make existing doctrines workable: identity disclosure between agents, audit-log retention, scoped delegation tokens, sectoral disclosure rules for agent-mediated commerce, and harmonized incident reporting. The market is moving faster than the law, but the gap is narrower than common rhetoric suggests, and the materials for closing it are largely in place.

⁸⁹Cf. Frank H. Easterbrook, *Cyberspace and the Law of the Horse*, 1996 U. CHI. LEGAL F. 207 (warning against subject-specific cyberlaw); see also Lessig, *supra* (responding).

References

A. Cases

Mobley v. Workday, Inc., 740 F. Supp. 3d 796 (N.D. Cal. 2024). <https://casetext.com/case/mobley-v-workday-inc>

B. Statutes, Treaties, and Bills

Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, opened for signature Sept. 5, 2024, C.E.T.S. No. 225. <https://rm.coe.int/1680afae3c>
Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act), 2024 O.J. (L 1689). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on Liability for Defective Products, 2024 O.J. (L 2853) (Revised Product Liability Directive). <https://eur-lex.europa.eu/eli/dir/2024/2853/oj>

One Big Beautiful Bill Act, Pub. L. No. 119-21, 139 Stat. ____ (2025). <https://www.congress.gov/bill/119th-congress/house-bill/1>

Preventing Algorithmic Collusion Act of 2025, S. 232, 119th Cong. (2025). <https://www.congress.gov/bill/119th-congress/senate-bill/232>

Colorado Artificial Intelligence Act, S.B. 24-205, 74th Gen. Assemb., 2d Reg. Sess. (Colo. 2024) (codified at Colo. Rev. Stat. §§ 6-1-1701 to -1707). <https://leg.colorado.gov/bills/sb24-205>

Colorado S.B. 25B-004, 75th Gen. Assemb., 1st Extra. Sess. (Colo. 2025) (delaying effective date of Colorado AI Act). <https://leg.colorado.gov/bills/sb25b-004>

Texas Responsible Artificial Intelligence Governance Act, H.B. 149, 89th Leg., R.S. (Tex. 2025). <https://capitol.texas.gov/tlodocs/89R/billtext/pdf/HB00149F.pdf>

California Senate Bill 1047, 2023–2024 Reg. Sess. (Cal. 2024) (vetoed Sept. 29, 2024). https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202320240SB1047

California ADMT Final Regulations, Cal. Code Regs. tit. 11, §§ 7000–7304 (2025). <https://cpa.ca.gov/regulations/>

New York City Local Law 144 of 2021, Automated Employment Decision Tools, N.Y.C. Admin. Code §§ 20-870 to -874. <https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page>

Utah Artificial Intelligence Policy Act, S.B. 149, 2024 Gen. Sess. (Utah 2024) (codified at Utah Code §§ 13-2-12 to -13). <https://le.utah.gov/~2024/bills/static/SB0149.html>

Uniform Electronic Transactions Act §§ 2(6), 14 (Unif. Law Comm'n 1999). <https://www.uniformlaws.org/committees/community-home?CommunityKey=2c04b76c-2b7d-4399-977e-d5876ba7e034>

Bill C-27, An Act to Enact the Consumer Privacy Protection Act, the Personal Information and Data Protection Tribunal Act and the Artificial Intelligence and Data Act, 1st Sess., 44th Parl. (Can. 2022) (died on order paper, Jan. 2025). <https://www.parl.ca/DocumentViewer/en/44-1/bill/C-27/first-reading>

Projeto de Lei nº 2.338, de 2023 (Marco Legal da Inteligência Artificial) (Senado Federal, Brazil). <https://www25.senado.leg.br/web/atividade/materias/-/materia/157233>

Framework Act on the Development of Artificial Intelligence and the Establishment of Trust [AI Basic Act], Act No. 20762 (Republic of Korea, promulgated Jan. 21, 2025; effective Jan. 22, 2026). <https://aibasicact.kr/>

C. Executive, Administrative, and Regulatory Materials

- Exec. Order No. 14179, Removing Barriers to American Leadership in Artificial Intelligence, 90 Fed. Reg. 8741 (Jan. 23, 2025). <https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/>
- Executive Office of the President, America's AI Action Plan (July 2025). <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>
- Office of Mgmt. & Budget, Exec. Office of the President, M-25-21, Accelerating Federal Use of AI Through Innovation, Governance, and Public Trust (Apr. 3, 2025). <https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf>
- Office of Mgmt. & Budget, M-25-22, Driving Efficient Acquisition of Artificial Intelligence in Government (Apr. 3, 2025). <https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf>
- NIST, AI Agent Standards Initiative — Notice of Inquiry, NIST-2025-0035 (Feb. 17, 2026). <https://www.regulations.gov/document/NIST-2025-0035-0001>
- Federal Trade Commission, Operation AI Comply: Continuing the Crackdown on Deceptive AI Claims and Schemes (Sept. 25, 2024). <https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes>
- European Commission, Guidelines for Providers of General-Purpose AI Models (July 18, 2025). <https://digital-strategy.ec.europa.eu/en/policies/guidelines-gpai-providers>
- European Commission, Code of Practice on Marking and Labelling of AI-Generated Content (consultation 2026). <https://digital-strategy.ec.europa.eu/en/policies/code-practice-ai-generated-content>
- European Data Protection Board, Opinion 28/2024 on Certain Data Protection Aspects Related to the Processing of Personal Data in the Context of AI Models (Dec. 17, 2024). https://www.edpb.europa.eu/our-work-tools/our-documents/opinion-board-art-64/opinion-282024-certain-data-protection-aspects_en
- European Parliament, Legislative Train, AI Liability Directive (current status: withdrawn) (last visited May 2026). <https://www.europarl.europa.eu/legislative-train/theme-a-europe-fit-for-the-digital-age/file-ai-liability-directive>
- Cyberspace Administration of China et al., Interim Measures for the Management of Generative Artificial Intelligence Services (effective Aug. 15, 2023) (DigiChina trans.). <https://digichina.stanford.edu/work/translation-measures-for-the-management-of-generative-artificial-intelligence-services-effective-august-15-2023/>
- Cyberspace Administration of China, Provisions on the Administration of Algorithmic Recommendations of Internet Information Services (effective Mar. 1, 2022) (DigiChina trans.). <https://digichina.stanford.edu/work/translation-internet-information-service-algorithmic-recommendation-management-provisions-effective-march-1-2022/>
- Cyberspace Administration of China, Provisions on the Administration of Deep Synthesis Internet Information Services (effective Jan. 10, 2023) (DigiChina trans.). <https://digichina.stanford.edu/work/translation-provisions-on-the-administration-of-deep-synthesis-internet-information-services/>
- Cyberspace Administration of China, Measures for Labeling AI-Generated Synthetic Content (effective Sept. 1, 2025). <https://www.cac.gov.cn/>
- Department for Science, Innovation and Technology (UK), A Pro-Innovation Approach to AI Regulation, CP 815 (Mar. 29, 2023). <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach>

UK AI Security Institute (formerly AI Safety Institute). <https://www.aisi.gov.uk/>

Information Commissioner’s Office (UK), Guidance on AI and Data Protection (Mar. 2023, updated Jan. 2025). <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/>

Personal Data Protection Commission (Singapore), Model AI Governance Framework for Generative AI (May 30, 2024). <https://www.pdpc.gov.sg/help-and-resources/2024/05/model-ai-governance-framework-for-generative-ai>

Ministry of Economy, Trade & Industry (Japan), AI Guidelines for Business v. 1.1 (Mar. 28, 2025). https://www.meti.go.jp/english/press/2025/0328_001.html

European Parliament Resolution of 16 February 2017 with Recommendations to the Commission on Civil Law Rules on Robotics, 2015/2103(INL). https://www.europarl.europa.eu/doceo/document/TA-8-2017-0051_EN.html

D. Standards and Soft-Law Instruments

Nat’l Inst. Standards & Tech., U.S. Dep’t of Commerce, NIST AI 100-1, Artificial Intelligence Risk Management Framework (AI RMF 1.0) (Jan. 2023). <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>

Nat’l Inst. Standards & Tech., NIST AI 600-1, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (July 2024). <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>

OECD, Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449 (adopted May 22, 2019; revised May 2024). <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>

OECD, Algorithmic Competition, OECD Competition Policy Roundtable Background Note (May 2023). https://www.oecd.org/content/dam/oecd/en/publications/reports/2023/05/algorithmic-competition_2be02d00/cb3b2075-en.pdf

OECD, Artificial Intelligence, Data and Competition (June 2024). https://www.oecd.org/en/publications/artificial-intelligence-data-and-competition_e1d1ad17-en.html

G7 Hiroshima Process International Code of Conduct for Organizations Developing Advanced AI Systems (Oct. 30, 2023). <https://digital-strategy.ec.europa.eu/en/library/hiroshima-process-international-code-conduct-advanced-ai-systems>

Int’l Org. for Standardization & Int’l Electrotechnical Comm’n, ISO/IEC 42001:2023, Information Technology — Artificial Intelligence — Management System (Dec. 2023). <https://www.iso.org/standard/81230.html>

Council of Europe, Council of Europe Opens First Ever Global Treaty on AI for Signature (Sept. 5, 2024). <https://www.coe.int/en/web/portal/-/council-of-europe-opens-first-ever-global-treaty-on-ai-for-signature>

Cloud Security Alliance, NIST AI Risk Management Framework: Agentic Profile v.1 (2026). <https://labs.cloudsecurityalliance.org/agentic/agentic-nist-ai-rmf-profile-v1/>

U.N. Comm’n on Int’l Trade Law (UNCITRAL), Working Group IV (Electronic Commerce), Use of Artificial Intelligence and Automation in Contracting, U.N. Doc. A/CN.9/WG.IV/WP.180 (2024). https://uncitral.un.org/en/working_groups/4/electronic_commerce

E. Technical Protocols and Industry Position Papers

Anthropic, Introducing the Model Context Protocol (Nov. 25, 2024). <https://www.anthropic.com/news/model-context-protocol>

Model Context Protocol Specification. <https://modelcontextprotocol.io/specification>

Google, Announcing the Agent2Agent Protocol (A2A) (Apr. 9, 2025).

<https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interopability/>

Yonadav Shavit et al., Practices for Governing Agentic AI Systems, OpenAI (Dec. 14, 2023).

<https://cdn.openai.com/papers/practices-for-governing-agentic-ai-systems.pdf>

Iason Gabriel et al., The Ethics of Advanced AI Assistants, Google DeepMind (Apr. 19, 2024).

<https://storage.googleapis.com/deepmind-media/DeepMind.com/Blog/the-ethics-of-advanced-ai-assistants/the-ethics-of-advanced-ai-assistants.pdf>

F. Books

Lawrence Lessig, Code and Other Laws of Cyberspace (1999).

<https://www.basicbooks.com/titles/lawrence-lessig/code/9780465039142/>

Lawrence Lessig, Code: Version 2.0 (2006). <https://codev2.cc/>

Frank Pasquale, The Black Box Society: The Secret Algorithms That Control Money and Information (Harv. Univ. Press 2015). <https://www.hup.harvard.edu/books/9780674970847>

Ariel Ezrachi & Maurice E. Stucke, Virtual Competition: The Promise and Perils of the Algorithm-Driven Economy (Harv. Univ. Press 2016). <https://www.hup.harvard.edu/books/9780674545472>

G. Articles and Working Papers

Ryan Calo, Artificial Intelligence Policy: A Primer and Roadmap, 51 U.C. Davis L. Rev. 399 (2017).

https://lawreview.law.ucdavis.edu/sites/g/files/dgvnsk15026/files/media/documents/51-2_Calo.pdf

Andrew Tutt, An FDA for Algorithms, 69 Admin. L. Rev. 83 (2017).

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2747994

Jack M. Balkin, The Three Laws of Robotics in the Age of Big Data, 78 Ohio St. L.J. 1217 (2017).

https://moritzlaw.osu.edu/sites/default/files/2021-12/oslj_v78n5_5balkin.pdf

Margot E. Kaminski, The Right to Explanation, Explained, 34 Berkeley Tech. L.J. 189 (2019).

<https://lawcat.berkeley.edu/record/1128984/files/fulltext.pdf>

Kevin Werbach & Nicolas Cornell, Contracts Ex Machina, 67 Duke L.J. 313 (2017).

<https://scholarship.law.duke.edu/dlj/vol67/iss2/2/>

Helen Nissenbaum, Accountability in a Computerized Society, 2 Sci. & Eng'g Ethics 25 (1996).

<https://nissenbaum.tech.cornell.edu/papers/Accountability.pdf>

Madeleine Clare Elish, Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction, 5

Engaging Sci., Tech. & Soc'y 40 (2019). <https://estsjournal.org/index.php/ests/article/view/260>

Ariel Ezrachi & Maurice E. Stucke, Artificial Intelligence and Collusion: When Computers Inhibit Competition, 2017 U. Ill. L. Rev. 1775. <https://illinoislawreview.org/print/volume-2017-issue-5/artificial-intelligence-collusion-when-computers-inhibit-competition/>

Frank H. Easterbrook, Cyberspace and the Law of the Horse, 1996 U. Chi. Legal F. 207.

https://chicagounbound.uchicago.edu/journal_articles/2147/

Noam Kolt, Governing AI Agents, 100 Notre Dame L. Rev. (forthcoming 2025).

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4772956

John J. Nay, Aligning AI Agents with Humans Through Law as Information, Stanford CodeX (Oct. 2025). <https://law.stanford.edu/wp-content/uploads/2025/10/Aligning-AI-Agents-with-Humans-through-Law-as-Information.pdf>

Nathalie Nevejans, European Civil Law Rules in Robotics, Eur. Parl., Directorate-General for Internal Policies (2016).

[https://www.europarl.europa.eu/RegData/etudes/STUD/2016/571379/IPOL_STU\(2016\)571379_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2016/571379/IPOL_STU(2016)571379_EN.pdf)

Open Letter to the European Commission: Artificial Intelligence and Robotics (Apr. 12, 2018).
<http://www.robotics-openletter.eu/>

H. Policy and Practitioner Commentary

Brookings Institution, Toward a Governance Framework for Agentic AI (2025).

<https://www.brookings.edu/articles/toward-a-governance-framework-for-agentic-ai/>

Helen Toner & Tantum Collins, Decoding Intentions: Artificial Intelligence and Costly Signals, Ctr. for Sec. & Emerging Tech. (Oct. 2023). <https://cset.georgetown.edu/publication/decoding-intentions/>

Future of Privacy Forum, From Chatbot to Checkout: Who Pays When Transactional Agents Play? (2025). <https://fpf.org/blog/from-chatbot-to-checkout-who-pays-when-transactional-agents-play/>

DLA Piper, Latest Wave of Obligations Under the EU AI Act Take Effect (Aug. 2025).

<https://www.dlapiper.com/en-us/insights/publications/2025/08/latest-wave-of-obligations-under-the-eu-ai-act-take-effect>

K&L Gates, Pared-Back Version of the Texas Responsible Artificial Intelligence Governance Act Signed Into Law (June 24, 2025). <https://www.klgates.com/Pared-Back-Version-of-the-Texas-Responsible-Artificial-Intelligence-Governance-Act-Signed-Into-Law-6-24-2025>

Akin Gump, Colorado Postpones Implementation of Colorado AI Act, SB 24-205 (Sept. 2025).

<https://www.akingump.com/en/insights/ai-law-and-regulation-tracker/colorado-postpones-implementation-of-colorado-ai-act-sb-24-205>

Goodwin Procter LLP, Federal AI Moratorium Dies on the Vine as Senate Passes the One Big Beautiful Bill Act (July 2025). <https://www.goodwinlaw.com/en/insights/publications/2025/07/alerts-practices-aiml-federal-ai-moratorium-dies-on-the-vine>

Mayer Brown, China Formulates New AI Global Governance Action Plan and Issues Draft Ethics Rules and AI Labelling Rules (Oct. 2025).

<https://www.mayerbrown.com/en/insights/publications/2025/10/artificial-intelligence-a-brave-new-world-china-formulates-new-ai-global-governance-action-plan-and-issues-draft-ethics-rules-and-ai-labelling-rules>

Cooley LLP, South Korea's AI Basic Act: Overview and Key Takeaways (Jan. 2026).

<https://www.cooley.com/news/insight/2026/2026-01-27-south-koreas-ai-basic-act-overview-and-key-takeaways>