



PROGRAM MATERIALS

Program #36209

October 13, 2026

AI Agents Unleashed: Understanding the Evolving Legal Risks of Autonomous AI and How to Stay Ahead of Them

Copyright ©2026 by

- **James Gatto, Esq. - Sheppard**

**All Rights Reserved.
Licensed to Celesq®, Inc.**

Celesq® AttorneysEd Center
www.celesq.com

5301 North Federal Highway, Suite 150, Boca Raton, FL 33487
Phone 561-241-1919



AI Agents Unleashed

Understanding the Evolving Legal Risks of Autonomous AI and How to Stay Ahead of Them

Presented by:



James G. Gatto
Partner | Washington, DC
jgatto@sheppard.com
+1.703.989.9288

Speaker Biography



James G. Gatto

J.D., Georgetown University Law Center, 1988

B.E., Electrical Engineering (Physics minor), Manhattan College, 1984

Former U.S. Patent Examiner

Bio

"I am passionate about the intersection of disruptive tech, law and business."

38 years of business-focused, legal advice on all aspects of intellectual property strategy, technology transactions, technology-related regulatory issues and litigation, especially ones driven by new business models and/or disruptive technology.

AI, Robotics and Quantum Team Leader

- **Artificial intelligence** (training, policies, IP, regulatory) 26 years
- **Quantum Technology and Robotics** (IP, regulatory, compliance) 15 years
- **Open Source** (audits, diligence, license issues, policies) 28 years
- **Blockchain** (blockchain games, crypto/NFTs, metaverses, digital art) 14 years

Some Recent AI Activity

- Adjunct Professor, George Mason Law School "Emerging AI Legal Issues"
- Adjunct Professor, Ole Miss Law School "Legal Issues with AI"
- Invited Speaker, Korean Copyright Office "AI and Open Source"
- Speaker, US Copyright Office Listening Session on AI Authorship
- Speaker, USPTO Listening Session on AI Inventorship Issues
- ABA-IPL AI/Machine Learning (AI/ML) Task Force
- Member, Artificial Intelligence Committee, International Technology Law Association

Learning Objectives



Develop an Understanding of AI Agents and Why Their Use is Surging




Learn About AI Agent Liability Issues



Identify and Learn How to Manage Legal and Regulatory Risks with AI Agents

Overview

- Why AI Agent Use is Surging
- What are AI agents and Agentic AI?
- Understanding and Managing Agent Risks
- Agentic AI Liability
- AI Agent Policies



Agentic AI v. Generative AI

Generative AI

User prompt



Content created

Agentic AI

User goal/task



Action taken

AI Agent

AI tool that:

- You give a specific goal/task
- It uses reasoning to develop a plan on how to accomplish it
- Uses data, tools and resources
- Can act autonomously or semi-autonomously to take action
- Can adapt over time




Agentic AI

- A system of agents
- Often with an orchestrator agent
- Typically give it a general goal
- It uses reasoning to determine the steps needed to accomplish the goal
- Often determines which agents to use to complete the steps
- “Orchestrates” the results from the various agents
- Takes action
- Can adapt over time



Why AI Agent Use is Surging

- **Marc Benioff (CEO, Salesforce)** - calls AI agents the new digital labor force that represent "the third wave" of AI (after predictive and generative AI) and that agents will work "longside human employees as a core part of the workforce
- **Dario Amodei (CEO, Anthropic)** - has spoken about AI agents that can go off and act autonomously... for hours or even days, suggesting agents could perform complex, multi-step knowledge work independently
- **Sam Altman (CEO, OpenAI)** - has described a future where AI agents join the workforce and materially change how work gets done, potentially performing the equivalent of a full day's work for a knowledge worker




Why AI Agent Use is Surging

- **Satya Nadella (CEO, Microsoft)**

- Agents are going to be the new apps... Agents will fundamentally replace all software
- reported a surge in AI adoption within 2 months of GA, with 40 million "Agent 365" registrations and 30 million paid Copilot seats, signifying AI agents outnumbering human AI seats

- **Jensen Huang (CEO, NVIDIA)** predicts there will be a trillion agents running 24/7 & compute demand will explode:


- The reason why we need so many more computers is that today we have a billion people using computers
- At any given point in time, only a hundred million people are using them at the same time around the Earth
- In the future, we will be surrounded by agents helping us do things and those trillion agents will be using computers all the time



Core Components of Agent Technology

Most AI agents include a set of components:

- **Model** that serves as the central reasoning and planning engine, or "brain," of the agent, processing instructions and generating contextually appropriate outputs
- **Instructions** - natural-language commands (such as a system prompt) that define the agent's role, capabilities, and behavioral constraints
- **Memory (context)** - information stored and made accessible to the model, whether short-term (within a session) or long-term (across sessions or users)
- **Planning and reasoning** - the model is trained to decompose a goal into a sequence of steps
- **Data/Tools/Resources** – what is available to enable the agent to interact with other systems, such as reading from or writing to databases, sending communications, or executing transactions
- **Protocols**, standardized methods for agents to communicate with tools (e.g., the Model Context Protocol) and with other agents (e.g., the Agent2Agent protocol) and to make payments (e.g., Agent Payments Protocol (AP2)), etc.



Agentic AI Characteristics


- **Autonomy and continuous operation.** Agents can initiate and complete multi-step tasks without step-by-step human instruction, and can operate 24 hours a day, which amplifies both their utility and their potential to cause harm at scale before any human notices
- **Delegation of judgment.** Traditional software executes deterministic logic written in advance by a programmer. An agent is instead given a goal and a set of tools and is left to determine, in real time, how to achieve that goal
- **Tool and system integration.** Agents can connect to and interact with various resources including external systems — search engines, databases, calendars, email clients, payment platforms, code execution environments — and use those tools to take real-world actions such as sending communications, executing transactions, or modifying records
- **Memory and context persistence.** Unlike a single-session chatbot, many agents retain memory across interactions, building a running context of prior tasks, conversations, and data processed, which raises data retention, privacy, and drift concerns not present in stateless tools
- **Interaction and unpredictability.** Agents interact with other agents, third-party systems, and unpredictable real-world inputs, and in multi-agent systems the overall behavior of the system may be emergent — not fully predictable from the behavior of any single agent examined in isolation
- **Opaque reasoning.** An agent's internal "chain of thought" or planning process is often not fully traceable or explainable after the fact, which complicates both real-time oversight and after-the-fact investigation or litigation



Action-Space and Autonomy




Action-Space: the range of actions the agent is permitted to take, the tools it can invoke, the data/data sources it can access, the systems it can access, etc.



Autonomy: the degree to which an agent can decide when and how to act toward a goal, as opposed to following a detailed, predetermined standard operating procedure and the degree of human involvement


Action-Space

- What actions will it be authorized to execute?
- What data and data sources will the agent have permissions to access?
 - Read only access or Read-Write?
- What other tools can it access (payments, accounts, code)?
- What systems - sandboxed environments only, internal systems, and/or external third-party systems?
- Can the agent communicate with other agents? If so, which ones?
- Which agent protocols does the agent use (e.g., MCP, A2A, ACP)?



Autonomy

- **Level 1 – User as Operator**
The human fully controls the agent's actions. The agent executes tasks only when explicitly instructed, with no independent decision-making
- **Level 2 – User as Collaborator**
The agent can suggest actions, provide recommendations, or perform limited tasks autonomously, but the human must confirm or approve each step
- **Level 3 – User as Consultant**
The agent can plan and execute tasks with some independence, but the human reviews and approves significant decisions. The agent may adapt based on context but remains under human oversight
- **Level 4 – User as Approver**
The agent can make most decisions and act autonomously, but the human retains the authority to approve major actions or override the agent when needed
- **Level 5 – User as Observer**
The agent operates fully autonomously, making and executing decisions without human intervention. The human's role is limited to monitoring outcomes and, if necessary, taking over



AI Agents Risk

- Elevated and distinct risks compared to traditional generative AI tools due to their:
 - Autonomy
 - Ability to plan and execute multi-step actions
 - Integration with tools/APIs/databases
 - Maintaining memory across interactions
 - Data privacy and compliance risks
 - Autonomous **adaptation** or multi-agent interactions - can cause agents to deviate from goals, create cascading failures, or operate unpredictably
- These capabilities enable powerful automation but also create new vulnerabilities in security, compliance, operations, and liability




Case Study – Agentic Commerce

Agentic AI Commerce

AI agents can autonomously initiate and execute transactions on consumers' behalf

- **Authentication** - Before a transaction is approved, it is important to authenticate the agent (e.g., that it is a valid agent, and that it is associated with a specific entity on whose behalf it is permitted to act)
- **Authorization** - Once authenticated, need to ensure the agent is authorized for the particular transaction (e.g., has the user empowered it to implement the type of transaction, the frequency, the dollar amount, etc.)
- **Consent** - users must review and consent to a merchant's terms of service and privacy policy; if an agent visits a merchant site that the user has never visited, and "clicks" I agree, is that binding consent to the terms by the user?
- **Fraud Risk** - Agentic systems introduce a new attack vector for fraud: the threat model shifts from stolen credit cards to malicious, stolen or manipulated agents
- **Privacy** - Agentic payment systems are also more likely to require broader data sharing across agents, wallets, issuers, merchants, and third-party AI providers


[Agentic AI Commerce: The Next Wave of Online Shopping and Retailer Risk](#)



Amazon v. Perplexity

- Amazon sued to prevent purchasing agents from buying things on Amazon despite user giving agent its account credentials and authorizing purchases
- Found that Perplexity's Comet browser agents access the platform **with the user's permission** but accesses the user's password-protected account, **without authorization** of the platform operator
- Court [issued](#) preliminary injunction: agents cannot access platforms' electronic systems **without permission from platform owners**
- 9th Circuit vacated (August 2026) - held Amazon was unlikely to succeed because it **could not show that Perplexity itself "accessed" Amazon's computers**; found it was the user, using Perplexity's AI agent as a tool, who accessed Amazon's computers to carry out shopping tasks


Amazon.com Services LLC v. Perplexity AI, Inc., No. 3:25-cv-09514 (N.D. Cal.).



Amazon v. Perplexity

Amazon amended complaint (September 21, 2026) with two new "access" theories:


- Agents are not merely software running locally at a user's direction — it claims Perplexity's own servers direct how Comet browser interacts with password-protected Amazon pages (what it clicks, where it navigates), calling Comet "Perplexity's instrument."
- On the iOS version, Perplexity's servers directly accessed password-protected areas of customers' Amazon accounts (rather than access occurring through the customer's own device)
- Perplexity tortiously interferes with contract by knowingly causing customers to violate Amazon's "Agent Terms":
 - Requiring AI agents to identify themselves
 - Allowing Amazon to restrict their access
 - Barring agents from continuing to access Amazon.com after being told to stay out
- This theory doesn't require Perplexity to be bound by the Agent Terms; the terms bind Amazon's customers and Perplexity knowingly induces those customers to breach them



Why Doesn't Amazon Want to Sell via Agents?

It fundamentally alters the economics of e-commerce by eliminating the human shopping journey!


- AI agents do not “view” ads - \$68 billion ad business (sponsored products and paid search ads)
- Loss of cross-selling and impulse buys
 - The human shopping experience on Amazon is engineered for discovery
 - Features like "Customers who bought this also bought...", personalized recommendations, flash deals, and subscription prompts rely on human impulse
 - AI agents operate algorithmically to fulfill exact requests, cutting out cross-selling, upsells, and brand discovery
- Security, Liability, and Fraud Risks
 - Amazon cites data security, identity verification, and contractual terms of use when blocking external agents;
 - Liability concerns re who is responsible if an agent purchases the wrong item, buys at an incorrect price point, or makes unauthorized account alterations



Amazon does not oppose all AI shopping

- It opposes **third-party AI agents it cannot monetize**
- Amazon wants users to shop through its own native AI assistants (where it controls the recommendation engine and paid placements) rather than allowing rival platforms to claim the customer relationship


Amazon [announces](#) Rufus, a new generative AI-powered conversational shopping experience



Agentic AI Liability Issues

Sample Legal and Liability Risks

- Agency Law
- Contractual Liability
- Tort and Negligence
- Multi-party Attribution



Agency Law

- Traditional agency law assumes a human (or a legal entity acting through humans) as principal and agent most jurisdictions do not recognize an AI system itself as a legal "agent" capable of holding rights or bearing obligations
- AI systems lack legal personhood — they cannot enter contracts or bear liability
- That does not mean agentic conduct falls into a legal void; rather, existing agency law is applied to determine whether, and to what extent, the humans and entities that deployed the AI agent are bound by, or liable for, what it does
- General principles of agency law - accountability rests with the humans/entities behind the agent
- Vicarious liability doctrine can make a principal or employer responsible for acts taken by an agent within the scope of its authorized function
- In the agentic AI context: what permissions, credentials, instructions, and system access did the company give the agent, and how was the agent's authority presented to users or counterparties?




Contract Formation Questions

- Whether contracts formed autonomously by an AI agent are binding on the deployer
- How offer-and-acceptance doctrine applies when an agent purports to make or accept an offer without contemporaneous human authorization
- Whether AI agents' lack of legal capacity creates ambiguity about the binding nature of AI-negotiated terms
- The risk that a consumer-facing or counterparty-facing agent creates unintended contractual obligations for the company
- Cross-border issues when different jurisdictions treat automated contracting differently



E-SIGN Act of 2000


- A federal statute that already anticipated this problem in a narrower context
- Defines an "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 at the time of the action or response" a definition broad enough to encompass many modern AI agents
- The statute provides that **a contract or record may not be denied legal effect solely because its formation involved the action of one or more electronic agents, so long as the action of the agent is legally attributable to the person to be bound**
- This confirms that agent-formed contracts can bind a principal, but leaves courts room to decide when that attribution is legally appropriate in any given case



Tort and Negligence - Product Liability


Where an AI agent's autonomous decision causes financial, physical, or reputational harm, several overlapping liability theories are already being tested against developers and deployers

- **Negligence** principles require identifying a duty of care, a breach, causation, and damages
- a company that deploys an agent without adequate testing, monitoring, or human oversight, or that grants an agent broader authority than its reliability warrants, may be found to have breached its duty of care even in the absence of a statute specifically addressing the conduct
- Because agentic systems act continuously and at scale, an unaddressed design flaw or an overly broad permission can multiply what might otherwise be a modest risk into large-scale harm before anyone intervenes




Tort, Negligence, and Product Liability

- **Product liability** theories — familiar from other AI-enhanced products such as chatbots, medical devices, and autonomous vehicles — are readily extrapolated to agentic systems and generally take three forms:
 - *Failure to warn* - where a developer or deployer fails to adequately disclose that AI is involved in a product's function, or fails to warn of AI-related risks not evident to a user at the time of use;
 - *Design-defect (or design negligence)* theories apply where a developer or deployer fails to make reasonable efforts to ensure an agent is safe for its intended use, including failing to anticipate and mitigate foreseeable misuse or failure modes;
 - *Manufacturing-defect analogs* apply to flaws introduced during training or deployment, such as corrupted training data or outdated models producing incorrect outputs




Causation/Foreseeability

- Causation is complicated in the agentic context because AI systems learn and evolve through use
- Something can "go wrong" at multiple points in the lifecycle - insufficient model design, low-quality or biased training data, or careless deployment and inadequate instruction - sometimes in combination, making it difficult to apportion liability among developer, deployer, and any intermediate integrator
- What counts as "foreseeable" is also a moving target - likely to expand as awareness of agentic AI risks grows, meaning conduct that seems reasonable today may be judged negligent in hindsight once the relevant risks become well known in the industry



Data Privacy and Confidentiality

- Agents that read email, access customer records, or process financial or health information necessarily ingest large volumes of personal and confidential data
- Need to ensure you are permitted to grant agents access to/use relevant data
- Consider:
 - PII, HIPPA, BIPA
 - Contractual obligations (e.g., third party data)



Multi-Party Liability Chains


- Several attribution models are emerging:
 - ***Developer Negligence***, where the developer is liable for failing to design, test, or warn adequately
 - ***Deployer/User Negligence***, where the entity that configured or directed the agent is liable for careless deployment or oversight
 - ***Control-and-Role Attribution***, liability tracks whichever party had the practical ability to prevent the harm
 - ***Joint or Layered Liability***, recognizing that in complex chains more than one party may bear responsibility simultaneously



Sample Laws

Core Liability Categories in California


Category	Applicable Statute / Legal Framework	How Liability Attaches
Enterprise Operations	Civ. Code § 1714.46 (AB 316)	Eliminates the defense that an AI agent acted as an independent, unforeseeable actor
Employment & Discrimination	Cal. Gov. Code § 12940 (FEHA)	Holds deployers vicariously liable for discriminatory actions or bias produced by AI agents
Consumer Protection	Bus. & Prof. Code § 17200 (UCL)	Holds companies liable for unauthorized contracts, misrepresentations, or deceptive trades made by their AI agents
Professional Conduct	Bus. & Prof. Code § 6068.1 (SB 574)	Bars non-delegation of core professional responsibilities to automated software



Designing/Deploying AI Agents to Mitigate Legal Risk


Bound the Risk Upfront Through Design

- **Define the action-space narrowly.** Limit the systems, tools, and data an agent can access to what is strictly necessary for its intended function, favoring internal, sandboxed, or read-only access over broad external, write-enabled access wherever the task allows
- **Match autonomy to risk tolerance.** Select the appropriate point on the human-involvement spectrum based on the detectability, reversibility and materiality of the agent's potential actions
- **Build in identity, credentialing, and least-privilege access controls** - give each agent its own verifiable identity and the minimum permissions necessary to perform its authorized function; recertify and review agent privileges periodically rather than treating initial provisioning as permanent
- **Use a risk-tiering methodology.** Classify each agent (low, medium, or high risk) based on its autonomy, the sensitivity of the data it can access, and the reversibility and materiality of the actions it can take, and apply governance overhead proportionate to that risk
- **Build In Governance and Human Oversight** - Ensure appropriate **human-in-the-loop** review; consider higher level to start with new agents




Test and Keep Testing

- **Test before deployment** - test agents for baseline safety and reliability **before** deployment; consider a gradual, staged rollout
- **Monitor continuously** – must ensure post-deployment monitoring as agents adapt and “drift” overtime



Contractual and Insurance - Best Practices


- **Negotiate agent-specific contract terms, not generic SaaS terms** - ensure that agreements for agent technology directly address autonomous decision-making, including:
 - Clear classification of the agent's function
 - Customer's oversight responsibility
 - Disclaimers requiring human verification of agent outputs before material business decisions.
 - Damage caps calibrated to the unpredictability of an adaptive system
 - Indemnification provisions that specifically address harm arising from the agent's autonomous conduct rather than only conventional software defects
- **Conduct AI-specific vendor due diligence.** Before onboarding any agent vendor, require a due diligence questionnaire covering the vendor's own risk assessment and regulatory classification of the system, training data practices, bias-mitigation measures, security posture, incident history, and data usage and retention practices, and review the vendor's actual technical documentation rather than relying solely on marketing claims about the system's capabilities
- **Confirm insurance coverage before scaling deployment.** Review commercial general liability, technology errors and omissions, and cyber insurance policies for AI-related exclusions, and engage with brokers and carriers proactively — particularly before deploying higher-risk or externally facing agents — rather than discovering a coverage gap only after a claim arises



Agentic AI Policies


AI Agent Development and Use Policies

- Many organizations already have an acceptable-use policy addressing generative AI tools such as chatbots; generally, does not address agentic AI (actions)
- Companies need to supplement their AI use policies to address AI agents




What AI Agent Policy Should Include

- Ensure responsible, secure, and legally compliant development, deployment, procurement, and use of AI agents
- Require a living inventory of every deployed agent — its owner, the data it can access, the systems and tools it can invoke, and its intended use case
- Classify each agent by risk tier based on its autonomy, the sensitivity of the data it touches, and the reversibility and materiality of the decisions it can make, with correspondingly stricter controls, review gates, and sign-off requirements as the tier increases
- Flag any agent affecting, employment, credit, housing, healthcare, or other legally sensitive decisions for mandatory legal and compliance review before deployment, given the state-law obligations discussed above



What AI Agent Policy Should Include


- **Authority Boundaries and Human Oversight**
 - Define in writing what each agent is and is not **authorized** to do, including any financial, contractual, or decisional limits
 - Specify **escalation thresholds** requiring human review or approval before an agent's action takes effect, particularly for actions that are irreversible, legally consequential, or above a defined materiality threshold; identify who has human oversight responsibility
 - Require a tested **kill switch or pause capability** with a clear internal process for who can invoke it and under what circumstances
 - Make clear that the person or team that authorized, deployed, or directed the agent remains responsible for its actions



What AI Agent Policy Should Include

- **Data Handling, Privacy, and Confidentiality Rules**


- Specify which categories of data (personal data, protected health information, trade secrets, privileged material, customer records) an agent may access, under what access controls, and on a least-privilege basis
- Build in the disclosures and opt-out mechanisms required by applicable privacy laws



What AI Agent Policy Should Include


- **Monitoring, Logging, and Audit Trails**

- Require comprehensive, tamper-resistant **logging** of every agent action, tool invocation, and data access, sufficient to reconstruct after the fact who or what triggered an action, what the agent decided, and why
- Implement **continuous monitoring** for anomalous agent behavior, decision drift, and policy violations, with defined response protocols for suspected malfunction, prompt injection, or unauthorized action
- Set a **retention schedule** for agent logs and audit records that satisfies both litigation-readiness needs and applicable regulatory retention requirements and treat those logs as potential litigation hold material from the outset



What AI Agent Policy Should Include

- **Marketing, Disclosure, and External Communications Review**
 - Require legal review of any external claim about what an AI agent can do, particularly claims about autonomy, accuracy, or the ability to replace human staff, given active regulatory enforcement against overstated AI agent capabilities (FTC/SEC)
 - Require disclosure to customers, employees, or job candidates when they are interacting with an AI agent rather than a human, particularly where applicable law or sector rules require it
- **Training, Enforcement, and Review Cycle**
 - Require role-appropriate training before employees are permitted to configure, supervise, or rely on an AI agent, with deeper technical and legal training for those who build, approve, or govern agents
- **Implement Technical Monitoring and Governance**





Questions?



James G. Gatto

*AI, Robotics and Quantum
Team Leader*

jgatto@sheppard.com

+1.703.989.9288

Let's [connect](#) on LinkedIn!

See articles on
Artificial Intelligence



Subscribe
to blog for updates





Thank you

Offices

Brussels • Century City • Chicago • Dallas • Houston • Los Angeles • London
New York • Orange County • San Diego (Del Mar) • San Diego (Downtown)
San Francisco • Seoul • Shanghai • Silicon Valley • Washington, DC