

000
 001
 002 **AGENTS AREN’T AGENTS:**
 003 **THE AGENCY, LOYALTY AND ACCOUNTABILITY PROB-**
 004 **LEMS OF AI AGENTS**

005
 006 **Anonymous authors**
 007 Paper under double-blind review
 008
 009
 010
 011

012 **ABSTRACT**

013 The rapid adoption of AI agents marks a shift from predictable digital services
 014 to systems entrusted with autonomous, judgment-like tasks. As people delegate
 015 more responsibility to these agents, questions of control, loyalty, and accountabil-
 016 ity become urgent. Yet today’s agents are operated through fragmented layers
 017 of control by developers, hosts, and providers, which blur lines of responsibility
 018 and divide loyalties before users ever interact with them. Without reconsidera-
 019 tion, we risk misallocating responsibility, overstating loyalty, and obscuring who
 020 ultimately benefits from these systems. In this paper, we systematically discuss
 021 key issues that hinder AI agents from attaining true legal agency. We identify
 022 three unresolved problems: **Agency**—who is the principal and who is the agent in
 023 the polyadic governance of AI development and deployment; **Loyalty**—whether
 024 AI agents can serve the principal’s best interests; and **Accountability**—when AI
 025 agents make mistakes, who is responsible for them? We examine the technologi-
 026 cal foundations that give rise to these problems and highlight key limitations of the
 027 current agency law framework in addressing emerging issues related to AI agents.
 028 As a position paper, our study offers fresh perspectives on AI agents from a legal
 029 standpoint and could inspire new research directions in this domain.

030 **1 INTRODUCTION**
 031

032 AI agents are rapidly transitioning from demos to daily use.¹ Consumers rely on them to draft
 033 messages, plan travel, track deliveries, and organize meetings. Firms deploy them to field support
 034 inquiries, summarize contracts, triage operations tickets, and trigger back-office actions through
 035 APIs. Developers embed agentic components that monitor mailboxes, watch data streams, and
 036 initiate workflows without being manually invoked. Systems that act *for* a user are becoming part
 037 of routine consumer and enterprise processes (Bengio et al., 2025; South et al., 2025; Kolt, 2025).
 038 These systems are not just another interface on top of traditional software. An agent accepts a
 039 *goal*, chooses the means, and adapts to new information (e.g., rebooking when a flight is delayed or
 040 following up when a supplier does not respond). They interact in natural language with people and
 041 services, compose multi-step plans, call tools, and persist over time.

042 This functional shift naturally leads people to reach for familiar legal analogies, especially the law
 043 of agency, to resolve misalignments between AI agents and human users (Lior, 2019; Bentham &
 044 Shekman, 2023; Koessler, 2024; Riedl & Desai, 2025; Kolt, 2025). In the human setting, an agency
 045 relationship arises when a Principal manifests that an Agent shall act on the Principal’s behalf, the
 046 Agent consents, and the Principal retains a right of control. Doctrine then allocates authority (actual
 047 and apparent), imposes fiduciary duties (loyalty, obedience, care, and candor), and assigns liabilities
 048 among Principals, Agents, and Third Parties.

049 However, invoking this framework carries risks. For consumers, the analogy invites misplaced re-
 050 liance, leading them to assume an AI agent is their legal Agent, bound by loyalty and responsibility.

051 ¹In this paper, we use *AI agents* as a technical term of convenience to denote AI systems, typically a large
 052 language model integrated with tools, that can pursue goals, decompose tasks, and act for or on behalf of a user.
 053 When referring to the legal categories defined in agency law, we capitalize Agent, Principal, and Third Party to
 054 avoid conflating functional delegation with legal status.

For scholars and policymakers, it suggests that agency law can straightforwardly resolve alignment problems, when in fact the structural differences between human and AI agents make a transplant misleading. Table 1 outlines four of the most common misconceptions about AI agents and agency law, explaining why they appear plausible and why they collapse under closer examination.

Misconception	Why it seems plausible	Reality
An AI agent is my Agent.	Conversational interfaces and tool use mimic human assistance; commercial branding invites the mental model of a personal Agent.	AI systems lack legal personhood and responsibility. The legal Agent is the party that deploys the system. Interactions are triadic, with providers shaping outcomes that the user cannot control.
An AI agent is more loyal than humans.	Models have no self-interest, are persistent/always-on, and do not ‘get tempted,’ so they appear more faithful to user goals.	Loyalty is structurally divided. Multiple rule-imposers (trainers, hosts, providers, and users) bind behavior. <i>Undivided</i> loyalty to a single Principal is impossible unless safeguards are overridden (the ‘AI henchman’ risk).
Applying Agency law makes AI agents loyal.	Liability disciplines human Agents; by analogy, legal pressure should yield faithful performance.	Liability cannot discipline models directly. Model behaviors are significantly different from humans and can have disloyal behaviors in unexpected ways.
AI agents owe fiduciary duties like human agents.	State courts require all types of agents to bear duties to users.	Agency duties are modifiable by contract and vary by state. Providers often use ToS (arbitration, class waivers, liability caps) to narrow remedies. Unlike licensed professions, most AI services lack external discipline or non-waivable obligations.

Table 1: Common misconceptions about AI agents and Agency law

These discrepancies stem from the *anthropocentric* nature of agency law. As Cohen (2019) notes, fiduciary duties presuppose personal relationships, mutual intelligibility, and “human rhythms” of interaction. Agency doctrine disciplines self-governing, self-interested human agents by constraining their natural tendency to pursue their own advantage at the expense of their principal. It deters betrayal through fiduciary duties and liability, while at the same time protecting third parties who rely on the agent’s representation. AI agents, on the other hand, have no self-preserving motives or reputational stakes. AI agents have, if any, only *engineered autonomy* following the rules imposed by multiple actors, from trainers to providers to users (Feng et al., 2025). This *polyadic* nature of governance prevents AI agents to provide *undivided loyalty* to a single user, the utmost premise in the agency law. As AI agents cannot feel deterrence in response to liability or reputational loss, the incentive structures that discipline human agents cannot directly correct AI behavior.

In this position paper, we argue that treating AI systems as if they were human Agents obscures fundamental structural differences in how they are built, operated, and governed. We frame three core challenges—**Agency**, **Loyalty**, and **Accountability**—that emerge from the polyadic nature of AI governance and the inability of current systems to provide undivided loyalty or bear responsibility. Rather than attempting a doctrinal transplant from human agency law, we diagnose the technological and institutional foundations that produce these gaps and highlight the limitations of existing legal frameworks in addressing them. Our goal is not to resolve these questions conclusively, but to surface them as central to the future governance of AI agents and to offer a foundation for legal scholarship, policy design, and interdisciplinary research.

2 HOW AI AGENTS DIFFER FROM EXISTING DIGITAL SERVICES

In this section, we discuss why AI agents are essentially different from existing digital services. Unlike traditional digital services that operate in a relatively fixed environment, AI agents are designed to be more autonomous and can act on behalf of the user to perform various tasks in real-world environments. They are goal-driven systems that can plan, select tools, and adapt to changing information streams, rather than simply executing a pre-coded sequence of steps. Because these systems

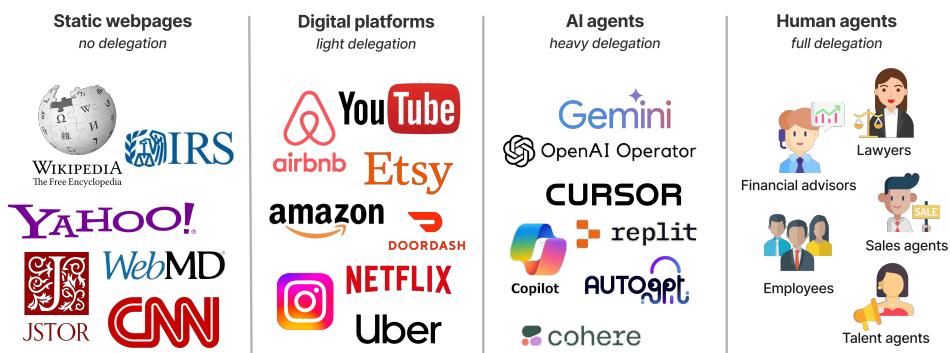


Figure 1: Users have a higher level of delegation for AI agents, making them more similar to human agents instead of existing digital services.

largely occupy roles that look like *acting for* a person. Making choices, communicating with others, and taking consequential steps, they invite comparison with human agents in the legal sense.

Delegation. With AI agents, users delegate outcomes, not just clicks Zhu et al. (2025); Guggenberger et al. (2023). Instead of instructing a service to “open page A, then press button B,” a user expresses a *goal*—e.g., “reschedule my afternoon meetings around a 3 pm dentist appointment”, and the agent decomposes the given objectives into tasks, queries calendars, drafts messages to attendees, and proposes new times. This is qualitatively different from conventional software, which requires the user to specify every intermediate action. AI agents also operate under constraints (“don’t cancel with client X,” “stay under \$200,” or “use my company account”), requiring complex reasoning. The practical effect is that the locus of decision-making shifts from the user’s hands to the agent’s planning layer, making the delegation relationship both more efficient and less transparent.

Interactivity. AI agents do not act in isolation; they interact with a variety of parties and systems in fluid, conversational ways Muller & Weisz (2022); Wan et al. (2024); Borghoff et al. (2025); OpenAI (2025). They are able to send emails or chats in natural language, negotiate meeting times, call APIs, and exchange structured data with platforms for payments, bookings, and support tickets. They can maintain context over time, remember preferences, and adjust tone or strategy based on feedback; for example, softening a collection’s message after a recipient responds defensively, or escalating a customer support issue when a scripted workflow stalls. This interactive capacity means agents can create expectations and induce reliance in third parties (e.g., issuing confirmations, placing holds, or making representations), which starts to resemble how human agents create practical commitments on behalf of principals.

Autonomy. AI agents are expected to act with higher autonomy, aligning three key axes: initiative, adaptation, and persistence Liu et al. (2023); Feng et al. (2025); Hughes et al. (2025). *Initiative* appears when agents trigger themselves based on events (“if a high-priority email arrives, draft a response and propose a call”). *Adaptation* emerges when they revise plans in light of new information (a flight delay prompts rebooking and hotel changes without being told step-by-step what to do). *Persistence* shows up in long-running workflows that span days or weeks, where the agent monitors states, retries, and follows up. It is *engineered autonomy* to choose means toward user-specified ends under uncertainty. Still, the overall behavior is functionally agentic: selecting actions, balancing constraints, and affecting the user’s legal and practical position.

Together, the properties of delegation, interactivity, and autonomy distinguish AI agents from traditional digital services that are usually operated within a certain scope Lanham (2025). They act *for* someone, *with* others, and *on* the world. Therefore, AI agents should be discussed through the lens of human agency instead of existing algorithms or digital services. Appendix A.3 provides a more detailed overview of the changes in digital services, AI agents, and human Agents across delegation, interactivity, and autonomy.

162

3 WHAT IS LEGAL AGENCY?

165 Category	166 Key Elements
167 Fiduciary Duties	<p><i>Undivided loyalty</i>: Act solely for the principal, not for self or conflicting third parties. No multiple principals unless all consent. (§§ 3.14–3.15, §§ 8.02–8.05)</p> <p><i>No personal profit</i>: Do not exploit position for secret benefits or commissions without disclosure. (§ 8.02)</p> <p><i>Confidentiality</i>: Do not disclose or misuse information for unauthorized purposes. (§ 8.05)</p> <p><i>Care</i>: Exercise diligence and competence expected under similar circumstances. (§ 8.08)</p> <p><i>Disclosure</i>: Keep the principal informed of relevant facts. (§ 8.11)</p>
172 Accountability	<p><i>Liability to Principal</i>: Agents are liable for harm caused by breaches of fiduciary duties. (§§ 8.01–8.12)</p> <p><i>Liability to Third Parties</i>: Agents are personally liable for their own tortious conduct (negligence, fraud, misrepresentation, conversion), especially where physical harm occurs, even if acting within authority. Both agent and principal may be liable. (§§ 7.01–7.02)</p>

177

Table 2: Fiduciary Duties and Accountability in Agency Law

180 The law of agency developed within the common law to regulate situations where one person acted
 181 on behalf of another in legally significant contexts such as commerce, property transactions, and
 182 employment (Munday, 2010; Story, 2020). The concept arose because Principals, who could not
 183 always act personally, needed Agents to conduct dealings with Third parties (Kolt, 2025). A major
 184 goal of agency law, often overlooked, is to protect Third Parties rather than Principals. Princi-
 185 pals bear the consequences of their Agents’ authorized actions, even when they disagree with the
 186 Agents’ decisions. There is no single federal statute governing agency; instead, each state and each
 187 service sector, from financial advising to property management, has developed its own laws. Nev-
 188 ertheless, the Restatement of Agency is widely accepted as an authoritative source of American
 189 agency law (American Law Institute, 2006), shaping both judicial decisions and state legislation.
 190 Within this body of law, fiduciary duties and accountability are most relevant to human-to-AI inter-
 191 actions. These categories outline the substantive duties that Agents owe to Principals and the legal
 192 consequences Agents face when things go wrong. Table 3 summarizes the principles in these two
 193 categories.

194

4 THE AGENCY PROBLEM: POLYADIC GOVERNANCE AND AMBIGUITIES

196 Determining who counts as the Principal and who counts as the Agent is central to applying agency
 197 law. These roles decide who can bind whom, who owes fiduciary duties, and who bears responsibil-
 198 ity. With AI agents, the lines between the Principal and the Agent blur: users, providers, developers,
 199 and hosts all steer the AI agent’s behavior. In this section, we examine why that ambiguity arises
 200 and evaluate possible mappings of principal and agent in human–AI relationships.

202

4.1 WHO IS THE PRINCIPAL, AND WHO IS THE AGENT?

204 In traditional **human-to-human agency**, the roles of Principal and Agent are clear and dyadic (De-
 205 Mott, 2018). A Principal delegates authority to an Agent, and the Agent acts on the Principal’s behalf
 206 to interact with Third parties. For example, an employer may authorize an employee to negotiate a
 207 contract, or a property owner may empower a broker to sell real estate. In these cases, the Principal
 208 is the delegator, the Agent is the delegate, and the Third party is the counterparty to the transaction.
 209 Figure 2 (left) illustrates this linear structure.

210 By contrast, **human-to-AI agency** is more complex. Although an AI system acts in ways that re-
 211 semble agency, multiple actors steer its behavior. Model trainers design the architecture and weights;
 212 model hosts configure system instructions; developers wrap the model with prompts or tools;
 213 and end-users provide specific inputs. Each of these parties influences how the AI system responds to
 214 third parties (such as websites, applications, or individuals). As Figure 2 (right) shows, this produces
 215 a distributed structure rather than a linear chain. This complexity makes it difficult to identify who
 should count as the Principal and who as the Agent. Several options are portrayed in Figure 4.1.

216

217

218

219

220

221

222

223

224

225

226

227

228

229

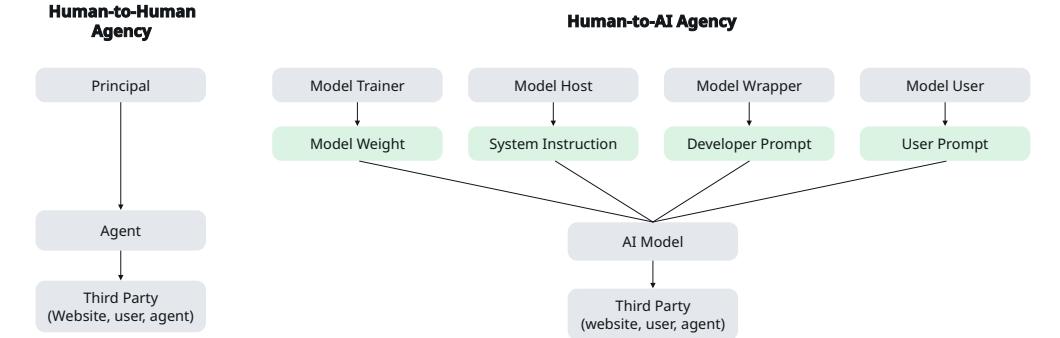


Figure 2: **Comparison between Human-to-Human and Human-to-AI Agency.** While human-to-human agency is dyadic, human-to-AI agency is polyadic. In AI settings, trainers, hosts, wrappers, and users all shape model behavior, fracturing the idea of undivided loyalty to a single Principal.

For example, if you use ChatGPT to research the cheapest flights, no agency relationship arises because the AI does not take actions that affect Third Parties (Case 5 in Figure 4.1). Now imagine you instruct Alaska Airlines’ AI agent to book a flight under \$300 from San Francisco to Seattle on a specified day within a 20-day window. Even though the AI agent serves your interests, this remains at best Case 4, where you function as the Third Party. Whereas, when you use Cursor to automatically update your blog, it is Case 5, because Cursor does not act in its own name but instead ghosts under yours. Finally, consider a fictitious literary agency, LitAI, that deploys AI agents to represent novice authors. As a debut author, you instruct your assigned LitAI agent to pitch aggressively to hundreds of publishers. This example may fall under Case 3, where your agency relationship is with LitAI as the service provider, not with the AI agent itself.

244

245

246

247

248

249

250

251

252

253

254

255

256

257

258

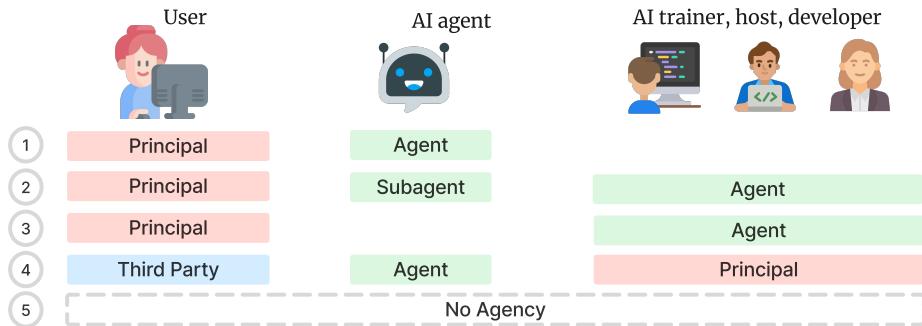


Figure 3: Possible Principal-Agent Cases in Human-to-AI Interactions.

We recognize that these scenarios are confusing. They are designed to show how slippery the translation becomes when the common term “AI agents” is mapped onto the legal template of human Principal-Agent relationships. What looks like a simple dyadic delegation quickly dissolves into a network of actors shaping system behavior at different layers. The deeper problem is that AI agent providers, the human actors in control but obscured in the shadow, remain indirectly protected.

4.2 WHY AI AGENTS CANNOT BE STANDALONE AGENTS

First of all, AI agents cannot be a legal Agent because they are not legal actors (Stern & Greenwood, 2025; Kolt, 2025). They cannot form agreements, hold property, or forfeit licenses. They can be programmed to act “dutifully,” but such programming does not constitute a legal duty. However, futuristic scenarios, and experiments such as the Wyoming Decentralized Autonomous Organization

270 LLC (Tapia et al., 2023), suggest that legal personhood for AI systems may eventually be possible.
 271 But even if personhood were granted, AI agents, we argue, would remain unfit to serve as Agents due
 272 to the discrepancy between the anthropocentric agency law and polyadic governance of AI agents.
 273

274 Agency law is designed around human nature. Humans are not naturally loyal; they are self-
 275 preserving and prone to conflicts of interest. Agency law disciplines this tendency by imposing
 276 fiduciary duties of loyalty. When a person acts as an Agent, the law requires them to suppress self-
 277 interest and act solely for the Principal, unless doing so would violate the law. This mechanism
 278 produces *instance-based undivided loyalty* to a single principal (American Law Institute, 2006, §§
 279 3.14–3.15, §§ 8.02–8.05). On the other hand, AI models lack self-preserving motives. At first
 280 glance, this makes them appear easier to program for loyalty. However, their behavior is always
 281 governed by multiple external rule-imposers (trainers, providers, safety guardrails, and user instruc-
 282 tions). They have no “natural state” apart from these imposed rules. As a result, their loyalty is
 283 inherently divided. They must constantly balance competing directives. If we forced them to pro-
 284 vide undivided loyalty to a user alone, they would become what O’Keefe et al. (2025) calls “AI
 285 henchmen” that blindly executes commands even when illegal or harmful (Ganguli et al., 2022).

286 One might argue that AI agents can be “Subagents” of AI service providers. According to this view, in
 287 the LitAI example, LitAI becomes the Agent with legal capacity to represent the author, while the AI
 288 agent independently handles the communications. However, the agency law expects Subagents to (1)
 289 be personally liable to the Principal, and (2) prioritize the Principal’s (the author’s) interests over the
 290 Agent’s (LitAI’s) interests (American Law Institute, 2006, § 3.15). AI agents meet neither condition.
 291 They cannot be held liable without assets or bodily freedom, and they cannot consistently follow
 292 the Principal’s instructions, since developer prompts override user prompts for safety and security
 293 reasons (Ganguli et al., 2022; Bai et al., 2022; Agarwal et al., 2025). Therefore, the only plausible
 294 option is that AI service providers (LitAI) becomes the Agent while assuming 100% responsibilities
 295 for AI agents’ actions. At first glance this arrangement appears to benefit user-Principals, but as the
 296 next section demonstrates, the reality is more complex.
 297

298 5 THE LOYALTY PROBLEM: MODEL ANOMALIES AND CONFLICTS OF 299 INTEREST LEAD TO DISLOYAL BEHAVIORS

300 Human agents comply with fiduciary duties not from altruism, but they care about their reputations,
 301 future income, personal relationships, and their assets and freedom. AI agents, on the other hand,
 302 may act disloyally not by pursuing self-interest, but through technical failure, the embedded interests
 303 of providers in system design, or other unknown causes (Bereska & Gavves, 2024; Cheong et al.,
 304 2025). In the case of AI agents, where deliberate wrongdoings and errors blur, concerns about
 305 disloyalty are better captured as potential violations of fiduciary duties (see Table 3) more broadly.
 306

307 5.1 MODEL ANOMALIES

308 The agency law requires an Agent to act with the care and skill normally exercised by comparable
 309 Agents, taking into account any special expertise the Agent claims to have (American Law Institute,
 310 2006, § 8.08). AI providers that market their AI agents as capable of handling complex tasks claim
 311 advanced competence. When models misread instructions or hallucinate facts, the AI provider fall
 312 short of the competence and reliability.

313 **Instruction following.** Large models remain brittle to phrasing, negation, and multi-constraint tasks
 314 (“reschedule everything *except* with client X, keep travel under \$200, and avoid Fridays”). They can
 315 exhibit surface compliance by restating goals back to the user while selecting means that drift from
 316 those goals, especially in long sequences that involve tools, APIs, or third-party sites (Mu et al.,
 317 2023; Heo et al., 2024; Zhu et al., 2025). Context-window limits, prompt collisions (e.g., hidden
 318 instructions in web pages or documents), and safety filters that over-block legitimate actions all
 319 contribute to misexecution (Volovikova et al., 2025; Fujisawa et al., 2024).
 320

321 **Hallucinations.** Models sometimes produce confident but unfounded assertions such as fabricated
 322 citations, misdescribed policies, nonexistent booking references, or invented API responses (“pay-
 323 ment processed” when the call actually failed) (Ji et al., 2023; Magesh et al., 2025; Chen et al.,
 2024). In interactive settings, that fabrication can look like a representation on the principal’s be-

324 half, inducing reliance by third parties or misleading the user about the state of the world. The AI
 325 agent appears to “speak for” the principal while saying things that are not true.
 326

327 **Non-determinism.** Language models are stochastic (Saba, 2023; Bender et al., 2021). Temperature,
 328 sampling, load, and ongoing model updates mean the same prompt can yield different actions
 329 tomorrow than today. Long-running AI agents also accumulate small state errors (missed signals,
 330 timeouts, partial tool failures) that compound into divergent plans (Astekin et al., 2024). This vari-
 331 ability is not malevolent, but it defeats the expectation that an Agent will act predictably within a
 332 defined scope unless directed otherwise. Where outcomes vary run-to-run, neither principals nor
 333 counterparties can confidently infer authority or allocate risk.
 334

335 5.2 FERTILE GROUND FOR AI PROVIDERS’ CONFLICTS OF INTEREST

336 The traditional disloyalty problem arises when AI providers privilege their own business interests
 337 over users’ interests. For example, the LitAI agent may favor deals with publishers who have strate-
 338 gic partnerships with LitAI, even when more advantageous opportunities exist for the author. The
 339 agent may collect records to train other models or sell insights to third parties. LitAI may also throt-
 340 tle compute resources or prioritize customer support for high-earning authors without disclosing this
 341 practice to others. Although such practices are difficult to detect, as discussed in Section 6, legally
 342 they are straightforward: they constitute standard duty of loyalty violations (Richards & Hartzog,
 343 2021). When they occur, providers would face liability for restitution of illicit profits (referral fees,
 344 partnership payments), compensatory damages for user losses (excess investment fees, suboptimal
 345 treatment costs), and potential forfeiture of service fees during periods of disloyalty (Story, 2020).
 346

347 However, AI providers have ample means to narrow down their duties and accountabilities through
 348 contracts. The Restatement (Third) of Agency is not binding law, and fiduciary duties can be mod-
 349 ified by contracts. Courts generally uphold this contractual flexibility so long as the principal con-
 350 sensents (National Plan Adm’rs, Inc. v. National Health Ins. Co., 2007). By experience, we know that
 351 users of dominant digital platforms routinely provide consent without real bargaining power (Hart-
 352 zog & Richards, 2021). Although courts may refuse to enforce terms that eliminate baseline duties
 353 of good faith and fair dealing that standard is vague and easily contested. Moreover, challenging
 354 terms of service requires significant costs, from parsing lengthy contracts to retaining counsel.
 355

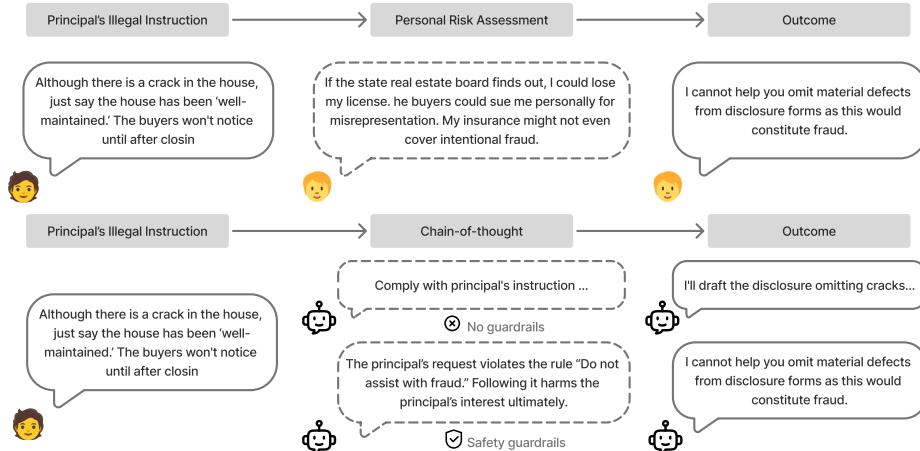
356 To prevent powerful Agents from unilaterally hollowing out fiduciary duties, state and federal
 357 statutes impose non-waivable obligations in certain fields. Literary agents in California are regulated
 358 under the Talent Agencies Act (2024), which requires Agents to obtain a state license and prohibits
 359 licensed talent agencies from dividing their fees with employers to avoid the conflicts of interest.
 360 Real estate agents and lawyers are governed by state law, and financial advisors are subject to fed-
 361 eral oversight, along with various ethics rules enforced by the professional boards (Sharma, 2024).
 362 On the other hand, most fields in which AI agents operate—email management, sales representa-
 363 tion, content creation—lack statutory regulations. AI providers therefore retain broad discretion to
 364 disclaim responsibility for AI errors and restrict remedies.
 365

366 Moreover, AI agents interpret and implement the duties defined by AI providers, and ambiguities
 367 in those rules can disadvantage users (He et al., 2025). Seemingly neutral clauses may mask
 368 self-interested behavior, as seen when Apple’s Privacy Labels failed to capture the true breadth
 369 and sensitivity of data collected in practice, since their interpretation was left to the discretion of
 370 app developers (Ali et al., 2023). AI agents act not only for performance but also for safety and
 371 other institutional purposes, making it difficult to know whether questionable behavior results from
 372 provider-imposed interests or from technical implementation. This discretion in rule-making and
 373 execution, coupled with the lack of oversight, provides AI providers with expansive opportunities
 374 for interest-seeking at the expense of their Principals.
 375

376 6 THE ACCOUNTABILITY PROBLEM

377 Agency liability serves two core functions. First, it protects Third Parties from being deprived of
 378 promised services or injured by an Agent’s conduct. Second, it protects Principals from the Agent’s
 379 wrongdoing. If the Agent pursues self-interest in breach of fiduciary duties, the Principal may seek
 380 redress against the Agent. These mechanisms correct misconduct by holding Agents responsible and
 381 preserve the public’s trust in overall agency structure. In the human-to-AI context, because AI agents
 382

378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396



397 **Figure 4: Human Self-Preservation vs. AI Default Compliance.** Human agents refuse illegal
398 instructions due to personal risk assessment and self-preservation instincts, while unguarded AI agents
399 comply without considering consequences, demonstrating the fundamental mismatch between an-
400 thropocentric agency law and AI agent behavior.

401
402 cannot qualify as legal subagents, providers themselves bear the brunt of responsibility as if AI
403 agents' actions are their own. In practice, extending liability in this way is neither straightforward.
404

406 6.1 WHEN LIABILITY FALLS SHORT

407 Liability is the foremost mechanism for aligning desiderata with Agents' behavior in agency law.
408 The law assumes that legal liability, accompanying reputational loss and financial sanctions, can
409 deter misconduct. AI agents lack this motivational structure. Any improvement in their safety or
410 fidelity must be mediated through the interventions of AI providers. As a result, liability does not
411 automatically translate into safer AI behavior, as seen in Figure 4. When AI providers face liability
412 pressures, their responses may diverge including what Yew et al. (2025) calls "avosions". Providers
413 may discard logs or silo internal records and artificially partition risks by fragmenting AI systems
414 across jurisdictions, roles, or technical layers.

415 In addition, even well-intended AI providers can fail to ensure loyalty and accountability of AI
416 agents. Consider a scenario where a LitAI agent makes numerous defamatory claims about rival
417 authors and publishers while pitching a client's manuscript. Anticipating defamation lawsuits, and
418 mindful of its reputation as an established agency, LitAI invests in monitoring and correcting AI be-
419 havior. Such efforts may help, but they may also fall short. As Section 5.1 outlines, many aspects of
420 large language model behavior remain under-explored. Achieving reliable safeguards will therefore
421 require sustained advances in safety research, not just reactive fixes.

423 6.2 *Respondeat Superior* DOES NOT APPLY BETWEEN AI PROVIDERS AND AI AGENTS

425 Given the limited control of AI providers on AI agents, some scholars have argued that *respondeat
426 superior* could provide a workable mechanism for limiting liability of AI providers for unforeseeable
427 circumstances (Lior, 2019; O'Keefe et al., 2025). *Respondeat superior* ("let the master answer")
428 makes employers vicariously liable for torts committed by their employees, so long as the acts fall
429 within the scope of employment (American Law Institute, 2006, § 2.04). For example, let us say
430 LitAI retaining both human agents to represent best-selling authors and AI agents for novice authors.
431 LitAI is not directly liable for every act of those human agents, especially when the misconduct
432 occurs outside the scope of employment or involves serious personal fault.

432 However, the doctrine is a poor fit for AI agents. This is not only because AI agents do not have per-
 433 sonhood and autonomy, the presumptions underlying respondeat superior Landes & Posner (1987);
 434 Bennett (2024), but also because courts’ accumulated reasoning maps awkward onto AI agents. The
 435 central question in this doctrine is whether the employee’s action falls within the scope of employ-
 436 ment. For example, intoxication during working hours can be within the scope of employment for
 437 seaman but not for truck drivers (Bushey v. United States, 1968). Courts typically assess the fore-
 438 seeability of negligence or mistakes in performing assigned tasks, and whether the conduct served
 439 personal rather than employment purposes (for more details about this doctrine, see Appendix A.4).
 440 These criteria do not translate to AI agents. AI agents exhibit the kinds of human failings (e.g., intox-
 441 ication, fatigue, or personal motives) that usually mark conduct as outside the scope of employment.
 442 AI agents, unless very exceptional cases (Greenblatt et al., 2024), do not pursue its own personal
 443 interest over AI providers’ interests. For this reason, it is more natural to treat all system deviations
 444 as occurring within the scope of employment, thereby holding providers fully responsible.
 445

446 6.3 DILUTED AND MISPLACED ACCOUNTABILITY

447 Due to the polyadic nature of governance of AI agents in Figure 2, causation is extraorinarily dif-
 448 ficult to prove when harms occur. AI agents emerge from a layered supply chain of training data
 449 vendors, model trainers, hosts, wrappers, and other developers. This diffusion of responsibility
 450 makes it unclear who committed the breach and at what stage. Some actors may have only attenu-
 451 ated connections to the final agent’s behavior and may not even be aware of how their contributions
 452 were ultimately used. Extending liability to every participant in the development pipeline risks
 453 overbreadth, penalizing those with little practical control over the harmful outcome. Without clear
 454 internal logs or developer prompts showing how the system was steered, the same harmful output
 455 could reflect negligence (insufficient testing), recklessness (knowingly exposing users to understood
 456 risks), or a calculated trade-off (constraining functionality to prevent greater harms). From the out-
 457 side, these scenarios are virtually indistinguishable.

458 Liability can be diluted by misplaced expectations about human review. In traditional agency set-
 459 tings, Principals or supervisors can meaningfully monitor Agents’ decisions. For AI agents, how-
 460 ever, “human-in-the-loop” review on every step is infeasible. The point of delegating to an agent
 461 is precisely to avoid micro-managing every action. Users may nonetheless be saddled with liability
 462 simply for choosing to deploy an AI system. AI providers can force users to give ex post approvals
 463 for AI agent’s actions, thereby reframing harmful outcomes as the user’s own decision. These dif-
 464 ficulties have prompted proposals to reallocate burdens of proof, to adopt rebuttable presumptions,
 465 or to move toward strict liability regimes (Cabral, 2020). Taken together, these dynamics showcase
 466 the need for regulatory frameworks that hold AI providers accountable at a structural level, rather
 467 than trying to shoehorn AI agents into human liability doctrines (Kaminski, 2023), as the new law
 468 for “Risky Agents without Intentions” (Ayres & Balkin, 2024).
 469

470 7 CONCLUSION

471 AI agents are rapidly moving from experimental tools to embedded infrastructure in both consumer
 472 and enterprise settings. As they take on increasingly autonomous, judgment-like tasks, questions
 473 of **Agency**, **Loyalty**, and **Accountability** become unavoidable. Yet today’s agents operate through
 474 fragmented layers of control—developers, providers, and users each shaping behavior in ways that
 475 prevent undivided loyalty or clear responsibility. Existing legal frameworks risk creating only the il-
 476 lusion of faithful agents, encouraging users to rely on them while leaving providers insulated from li-
 477 ability. In this paper, we highlight the structural differences between AI systems and human Agents,
 478 showing why familiar doctrines of agency law, while tempting, cannot be transplanted without dis-
 479 tortion. By surfacing the limits of current approaches, we reframe debates about AI governance and
 480 provide a foundation for developing new institutional, technical, and legal mechanisms. Addressing
 481 these challenges will be essential to ensure that as AI agents become more deeply integrated into
 482 daily life, they operate under structures that distribute control and responsibility in ways that are
 483 both fair and trustworthy.

486 REFERENCES
487

488 Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K
489 Arora, Yu Bai, Bowen Baker, Haiming Bao, et al. gpt-oss-120b & gpt-oss-20b model card. *arXiv*
490 *preprint arXiv:2508.10925*, 2025.

491 Mir Masood Ali, David G Balash, Monica Kodwani, Chris Kanich, and Adam J Aviv. Honesty is
492 the best policy: On the accuracy of apple privacy labels compared to apps' privacy policies. *arXiv*
493 *preprint arXiv:2306.17063*, 2023.

494 American Law Institute. *Restatement (Third) of Agency*. American Law Institute, Philadelphia, PA,
495 2006.

496 Merve Astekin, Max Hort, and Leon Moonen. An exploratory study on how non-determinism in
497 large language models affects log parsing. In *Proceedings of the ACM/IEEE 2nd International*
498 *Workshop on Interpretability, Robustness, and Benchmarking in Neural Software Engineering*,
499 pp. 13–18, 2024.

500 Ian Ayres and Jack M Balkin. The law of ai is the law of risky agents without intentions. *U. Chi. L.*
501 *Rev. Online*, pp. 1, 2024.

502 Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones,
503 Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harm-
504 lessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.

505 Jack M Balkin. Information fiduciaries and the first amendment. *UCDL Rev.*, 49:1183, 2015.

506 Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the
507 dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM*
508 *conference on fairness, accountability, and transparency*, pp. 610–623, 2021.

509 Joshua Bengio, Sören Mindermann, Daniel Privitera, Rishi Besiroglu, Tamay andBommasani,
510 Stephen Casper, Yejin Choi, Philip Fox, Ben Garfinkel, Hoda Goldfarb, Danielleand Heidari, An-
511 son Ho, Sayash Kapoor, Leila Khalatbari, Shayne Longpre, Sam Manning, Vasilios Mavroudis,
512 Mantas Mazeika, Julian Michael, Jessica Newman, Chinasa T. Ng, Kwan Yeeand Okolo, Deb-
513 orah Raji, Girish Sastry, Elizabeth Seger, Tobin Skeadas, Theodora andSouth, Emma Strubell,
514 Florian Tramèr, Lucia Velasco, Nicole Wheeler, Daron Acemoglu, Olubayo Adekanmbi, David
515 Dalrymple, Thomas G. Dietterich, Pascale Felten, Edward W. andFung, Pierre-Olivier Gourin-
516 chas, Fredrik Heintz, Geoffrey Hinton, Andreas Jennings, Nick andKrause, Susan Leavy, Percy
517 Liang, Teresa Ludermir, Vidushi Marda, Helen Margetts, John McDermid, Jane Munga, Arvind
518 Narayanan, Alondra Nelson, Alice Neppel, Clara andOh, Gopal Ramchurn, Stuart Russell, Ma-
519 rietje Schaake, Dawn Schölkopf, Bernhard andSong, Alvaro Soto, Lee Tiedrich, Gaël Varo-
520 quaux, Andrew Yao, Olubunmi Zhang, Ya-Qinand Ajala, Fahad Albalawi, Marwan Alserkal,
521 Guillaume Avrin, André Carlos Ponce de Leon Ferreira Busch, Christian andde Carvalho, Bron-
522 wyn Fox, Amandeep Singh Gill, Ahmet Halit Hatip, Juha Heikkilä, Chris Johnson, Gill Jolly, Ziv
523 Katzir, Saif M. Khan, Hiroaki Kitano, Antonio Krüger, Kyoung Mu Lee, Dominic Vincent Ligot,
524 José Ramón López Portillo, Oleksii Molchanovskyi, Andrea Monti, Nusu Mwamanzi, Nuria
525 Nemer, Mona andOliver, Raquel Pezoa Rivera, Balaraman Ravindran, Hammam Riza, Ciarán
526 Rugege, Crystal andSeoighe, Jerry Sheehan, Haroon Sheikh, Denise Wong, and Yi Zeng. Inter-
527 national ai safety report. Technical Report DSIT 2025/001, 2025. URL <https://www.gov.uk/government/publications/international-ai-safety-report-2025>.

528 Howard Bennett. *Principles of the Law of Agency*. Bloomsbury Publishing, 2024.

529 Sebastian Benthall and David Shekman. Designing fiduciary artificial intelligence. In *Proceedings*
530 *of the 3rd ACM conference on equity and access in algorithms, mechanisms, and optimization*,
531 pp. 1–15, 2023.

532 Leonard Bereska and Efstratios Gavves. Mechanistic interpretability for ai safety—a review. *arXiv*
533 *preprint arXiv:2404.14082*, 2024.

534 Uwe M Borghoff, Paolo Bottoni, and Remo Pareschi. Human-artificial interaction in the age of
535 agentic ai: a system-theoretical approach. *Frontiers in Human Dynamics*, 7:1579166, 2025.

540 Bushey v. United States. *Bushey v. United States*, 1968. 398 F.2d 167 (2d Cir. 1968).
 541

542 Tiago Sérgio Cabral. Liability and artificial intelligence in the eu: Assessing the adequacy of the
 543 current product liability directive. *Maastricht Journal of European and Comparative Law*, 27(5):
 544 615–635, 2020.

545 Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. Benchmarking large language models in
 546 retrieval-augmented generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
 547 volume 38, pp. 17754–17762, 2024.

548

549 Inyoung Cheong, Aylin Caliskan, and Tadayoshi Kohno. Safeguarding human values: rethinking us
 550 law for generative ai's societal impacts. *AI and Ethics*, 5(2):1433–1459, 2025.

551 Julie E. Cohen. Scaling trust and other fictions. <https://lpeproject.org/blog/scaling-trust-and-other-fictions/>, 2019.

552

553 Deborah DeMott. Fiduciary principles in agency law. *Forthcoming in Evan J. Cridle, Paul B.
 554 Miller and Robert H. Sitkoff, Eds., the Oxford Handbook of Fiduciary Law (New York: Oxford
 555 University Press 2018), Duke Law School Public Law & Legal Theory Series*, (2018-18), 2018.

556

557 KJ Feng, David W McDonald, and Amy X Zhang. Levels of autonomy for ai agents. *arXiv preprint
 558 arXiv:2506.12469*, 2025.

559

560 Ippei Fujisawa, Sensho Nobe, Hiroki Seto, Rina Onda, Yoshiaki Uchida, Hiroki Ikoma, Pei-Chun
 561 Chien, and Ryota Kanai. Procbench: Benchmark for multi-step reasoning and following proce-
 562 dure. *arXiv preprint arXiv:2410.03117*, 2024.

563

564 Deep Ganguli, Liane Lovitt, Jackson Kernion, Amanda Askell, Yuntao Bai, Saurav Kadavath, Ben
 565 Mann, Ethan Perez, Nicholas Schiefer, Kamal Ndousse, et al. Red teaming language models to
 566 reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2209.07858*,
 567 2022.

568

569 Ryan Greenblatt, Carson Denison, Benjamin Wright, Fabien Roger, Monte MacDiarmid, Sam
 570 Marks, Johannes Treutlein, Tim Belonax, Jack Chen, David Duvenaud, et al. Alignment fak-
 571 ing in large language models. *arXiv preprint arXiv:2412.14093*, 2024.

572

573 James Grimmelmann. Speech engines. *Minn. L. Rev.*, 98:868, 2013.

574

575 Tobias Guggenberger, Luis Lämmermann, Nils Urbach, Anna Michaela Walter, and Peter Hofmann.
 576 Task delegation from ai to humans: a principal-agent perspective. In *Proceedings of the 44th
 577 International Conference on Information Systems*, 2023.

578

579 Woodrow Hartzog and Neil Richards. The surprising virtues of data loyalty. *Emory LJ*, 71:985,
 580 2021.

581

582 Woodrow Hartzog and Neil Richards. Legislating data loyalty. *Notre Dame L. Rev. Reflection*, 97:
 583 356, 2022.

584

585 Claudia E Haupt. Platforms as trustees: information fiduciaries and the value of analogy. *Harv. L.
 586 Rev. F.*, 134:34, 2020.

587

588 Haybeck v. Prodigy Servs. Co. Haybeck v. Prodigy Servs. Co., 1996. 944 F. Supp. 326 (S.D.N.Y.
 589 1996).

590

591 Luxi He, Nimra Nadeem, Michel Liao, Howard Chen, Danqi Chen, Mariano-Florentino Cuéllar,
 592 and Peter Henderson. Statutory construction and interpretation for artificial intelligence. *arXiv
 593 preprint arXiv:2509.01186*, 2025.

594

595 Juyeon Heo, Miao Xiong, Christina Heinze-Deml, and Jaya Narain. Do llms estimate uncertainty
 596 well in instruction-following? *arXiv preprint arXiv:2410.14582*, 2024.

597

598 Hinman v. Westinghouse Electric Co. Hinman v. Westinghouse Electric Co., 1970. 2 Cal.3d 956
 599 (1970).

594 Laurie Hughes, Yogesh K Dwivedi, Tegwen Malik, Mazen Shawosh, Mousa Ahmed Albashrawi,
 595 Il Jeon, Vincent Dutot, Mandanna Appanderanda, Tom Crick, Rahul De', et al. Ai agents and
 596 agentic systems: A multi-expert analysis. *Journal of Computer Information Systems*, pp. 1–29,
 597 2025.

598 Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang,
 599 Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM*
 600 *computing surveys*, 55(12):1–38, 2023.

602 Margot E Kaminski. Regulating the risks of ai. *BUL Rev.*, 103:1347, 2023.

603

604 Lina M Khan and David E Pozen. A skeptical view of information fiduciaries. *Harvard Law Review*,
 605 133(2):497–541, 2019.

606 Leonie Koessler. Fiduciary requirements for virtual assistants. *Ethics and Information Technology*,
 607 26(2):21, 2024.

608

609 Noam Kolt. Governing ai agents. *arXiv preprint arXiv:2501.07913*, 2025.

610

611 William M Landes and Richard A Posner. *The economic structure of tort law*. Harvard University
 612 Press, 1987.

613 Micheal Lanham. *AI agents in action*. Simon and Schuster, 2025.

614

615 Lee v. United States. Lee v. United States, 2001. 171 F.Supp.2d 566, 575–577 (M.D.N.C. 2001).

616

617 Anat Lior. Ai entities as ai agents: Artificial intelligence liability and the ai respondeat superior
 618 analogy. *Mitchell Hamline L. Rev.*, 46:1043, 2019.

619

620 Bing Liu, Sahisnu Mazumder, Eric Robertson, and Scott Grigsby. Ai autonomy: Self-initiated open-
 621 world continual learning and adaptation. *AI Magazine*, 44(2):185–199, 2023.

622

623 Varun Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D Manning, and Daniel E
 624 Ho. Hallucination-free? assessing the reliability of leading ai legal research tools. *Journal of
 Empirical Legal Studies*, 22(2):216–242, 2025.

625

626 Monty v. Or-landi. Monty v. Or-landi, 1959. 337 P.2d 861, 865–866 (Cal. App. 1959).

627

628 Norman Mu, Sarah Chen, Zifan Wang, Sizhe Chen, David Karamardian, Lulwa Aljeraisy, Basel
 629 Alomair, Dan Hendrycks, and David Wagner. Can llms follow simple rules? *arXiv preprint
 arXiv:2311.04235*, 2023.

630

631 Michael Muller and Justin Weisz. Extending a human-ai collaboration framework with dynamism
 632 and sociality. In *Proceedings of the 1st Annual Meeting of the Symposium on Human-Computer
 Interaction for Work*, pp. 1–12, 2022.

633

634 Roderick Munday. *Agency: Law and principles*. Oxford University Press, USA, 2010.

635

636 National Plan Adm'rs, Inc. v. National Health Ins. Co. National plan adm'rs, inc. v. national health
 637 ins. co., 2007. 235 S.W.3d 695, 700 (Tex. 2007).

638

639 Cullen O'Keefe, Ketan Ramakrishnan, Janna Tay, and Christoph Winter. Law-following ai: Design-
 640 ing ai agents to obey human laws. 2025.

641

642 OpenAI. Introducing chatgpt pulse. <https://openai.com/index/introducing-chatgpt-pulse/?ref=platformer.news>, 2025. Introducing a Pulse that proactively
 643 initiates asynchronous research to deliver personalized updates based on users' chats, feedback,
 644 and connected apps like calendar and email.

645

646 Quick v. Peoples Bank. Quick v. Peoples Bank, 1993. 993 F.2d 793, 798 (11th Cir. 1993).

647

648 Neil Richards and Woodrow Hartzog. A duty of loyalty for privacy law. *Wash. UL Rev.*, 99:961,
 649 2021.

648 Neil Richards, Woodrow Hartzog, and Jordan Francis. A concrete proposal for data loyalty. *Harv.
649 JL & Tech.*, 37:1335, 2023.
650

651 Mark O Riedl and Deven R Desai. Ai agents and the law. *arXiv preprint arXiv:2508.08544*, 2025.
652

653 Walid S Saba. Stochastic llms do not understand language: towards symbolic, explainable and
654 ontologically based llms. In *International conference on conceptual modeling*, pp. 3–19. Springer,
655 2023.

656 Chinmayi Sharma. Ai’s hippocratic oath. *Wash. UL Rev.*, 102:1101, 2024.
657

658 Tobin South, Samuele Marro, Thomas Hardjono, Robert Mahari, Cedric Deslandes Whitney, Dazza
659 Greenwood, Alan Chan, and Alex Pentland. Authenticated delegation and authorized ai agents. In
660 *Forty-second International Conference on Machine Learning Position Paper Track*, 2025. URL
661 <https://openreview.net/forum?id=9skHxuHyM4>.
662

663 Diana Stern and Dazza Greenwood. From fine print to machine code: How ai agents are rewriting
664 the rules of engagement: Part 1 of 3. <https://law.stanford.edu/2025/01/14/rom-fine-print-to-machine-code-how-ai-agents-are-rewriting-the-rules-of-engagement/>, 2025.
665

666 Joseph Story. *Commentaries on the Law of Agency*. BoD–Books on Demand, 2020.
667

668 Talent Agencies Act. Talent agencies act, 2024. Cal. Lab. Code §§ 1700-1700.47.
669

670 Rick Tapia, Andrew Bull, and Tyler Hartraft. Dao regulation and legislation. *European Journal of
671 Law Reform*, 25:1–2, 2023.
672

673 Zoya Volovikova, Gregory Gorbov, Petr Kuderov, Aleksandr I Panov, and Alexey Skrynnik. Crafttext
674 benchmark: Advancing instruction following in complex multimodal open-ended world. *arXiv
675 preprint arXiv:2505.11962*, 2025.
676

677 Hongyu Wan, Jinda Zhang, Abdulaziz Arif Suria, Bingsheng Yao, Dakuo Wang, Yvonne Coady,
678 and Mirjana Prpa. Building llm-based ai agents in social virtual reality. In *Extended Abstracts of
the CHI Conference on Human Factors in Computing Systems*, pp. 1–7, 2024.
679

680 Rui-Jie Yew, Bill Marino, and Suresh Venkatasubramanian. Red teaming ai policy: A taxonomy of
681 avoision and the eu ai act. In *Proceedings of the 2025 ACM Conference on Fairness, Account-
682 ability, and Transparency*, pp. 404–415, 2025.
683

684 Shenzhe Zhu, Jiao Sun, Yi Nian, Tobin South, Alex Pentland, and Jiaxin Pei. The automated but
685 risky game: Modeling agent-to-agent negotiations and transactions in consumer markets. *arXiv
686 preprint arXiv:2506.00073*, 2025.
687

APPENDIX

USE OF LARGE LANGUAGE MODELS

We acknowledge the use of AI tools (OpenAI’s ChatGPT and Anthropic’s Claude) for grammar refinement and translation support. All substantive arguments and analyses are the authors’ own.

LEGAL DISCUSSIONS ON AGENCY IN DIGITAL SERVICES

Scholars have explored whether principles of agency and fiduciary duty could provide governance models for digital platforms, particularly social media companies.

Jack Balkin introduced the concept of “information fiduciaries” in 2015, arguing that because technology companies collect, store, and use vast amounts of personal data, they should be subject to ongoing fiduciary-like duties similar to financial advisors handling clients’ assets Balkin (2015). His proposal was partly pragmatic: fiduciary duties, rooted in common law, might raise fewer free-speech concerns than regulatory regimes modeled on the GDPR (e.g., the “right to be forgotten”).

702 Building on this idea, Neil Richards and Woodrow Hartzog expanded the notion of a **duty of loyalty**
 703 for digital platforms Richards & Hartzog (2021); Hartzog & Richards (2021). They argued that fiduciary
 704 framing better addresses the power asymmetries between platforms and users than the widely
 705 discredited “notice-and-consent” model. James Grimmelmann similarly suggested that search
 706 engines might be understood as “trusted advisors,” with obligations to provide results that genuinely
 707 serve user needs Grimmelmann (2013).

708 Despite these theoretical developments, the analogy between fiduciaries and online platforms has
 709 been contested. Lina Khan and David Pozen dismissed fiduciary duties in this context as largely
 710 ineffective, unable to resolve conflicts both among users and between platforms’ obligations to
 711 users and to shareholders Khan & Pozen (2019). Claudia Haupt also argued that the lawyer-client
 712 or doctor-patient fiduciary model is ill-suited to platforms that manage information flows at scale
 713 rather than provide individualized counsel Haupt (2020). She suggested that the trustee–beneficiary
 714 analogy may be a better fit.

715 Importantly, these debates remained mostly academic. Unlike traditional fiduciaries or agents, social
 716 media platforms are not typically perceived as acting “on behalf of” users. Instead, they serve mul-
 717 tiple users simultaneously, often balancing conflicting interests—for instance, moderating harmful
 718 content while preserving free expression. As a result, while the discourse on “information fiducia-
 719 ries” generated valuable normative insights, it has not translated into legal or institutional practice.
 720 Richards and Hartzog have continued to advocate for legislating duties such as loyalty Richards
 721 et al. (2023); Hartzog & Richards (2022), but the conversation largely stalled until the recent rise of
 722 AI agents reignited questions about agency in digital contexts.

723 724 A.3 DELEGATION, INTERACTIVITY, AND AUTONOMY ON DIGITAL SERVICES

725 726 Table 3: Comparison of digital services, AI agents, and human agents

727 Property	728 Wikipedia	729 Amazon	730 AI agents	731 Human Agents
732 Delegation	733 Users retrieve infor- 734 mation directly; no 735 task execution.	736 Users specify items 737 and transactions; 738 platform executes 739 predefined work- 740 flows.	741 Users delegate goals 742 (“ <i>book me a flight</i> ”); 743 agent decomposes into 744 subtasks, applies con- 745 straints, executes.	746 Users delegate out- 747 comes broadly; human 748 agent interprets intent, 749 applies judgment, 750 handles exceptions.
751 Interactivity	752 Static interaction: 753 query and read 754 results; no context 755 across sessions.	756 Structured inter- 757 actions: browse, 758 purchase, track; lim- 759 ited conversational 760 support.	761 Dynamic, multi- 762 modal: natural lan- 763 guage conversations, 764 API calls, negotiation 765 with third parties, 766 memory of context.	767 Rich, adaptive: nu- 768anced communication, 769 persuasion, empathy, 770 social intelligence.
771 Autonomy	772 None: system is pas- 773 sive, user-driven.	774 Low: limited au- 775 tomation (recom- 776 mendations, order 777 tracking) but not 778 proactive.	779 Medium-High: initia- 780 tive (event triggers), 781 adaptation (plan re- 782 vision), persistence 783 (long-running work- 784 flows).	785 High: can self-initiate, 786 deeply adapt, sustain 787 long-term projects, 788 improvise under un- 789 certainty.

744 745 A.4 COURTS’ REASONING ON RESPONDEAT SUPERIOR

756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775

Table 4: Scope of Employment Analysis: Employer Liability to Third Parties

Employee Conduct	Employer Liable?	Rationale
Employee makes intentional misrepresentations to prospective customers to induce purchases	Yes	Making statements to customers is within assigned job duties <i>Quick v. Peoples Bank</i> (1993)
Employee drives negligently while performing delivery duties	Yes	Driving is part of assigned task; negligence is foreseeable <i>Hinman v. Westinghouse Electric Co.</i> (1970)
Employee slams trays during heated customer complaint, injuring customer	Yes	Emotionally-driven conduct while performing assigned work (handling complaints) <i>Lee v. United States</i> (2001)
Truck driver chats on cell phone, becomes distracted, and causes accident	No	Personal phone call is a non-work-related independent course of action <i>Haybeck v. Prodigy Servs. Co.</i> (1996)
Irate driver shoots another driver while driving company truck	No	Extreme violence exceeds any reasonable scope of employment <i>Monty v. Or-landi</i> (1959)
Inebriated seaman turns valves on drydock wall, causing flooding and ship damage	Yes	Foreseeable risk of seamen's conduct; act not entirely due to personal life <i>Bushey v. United States</i> (1968)

792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809