Agent 2 Agent:

The Emergence of the Fourth Generation Internet, the Agentic Digital Economy

Executive Summary

We are entering the Fourth Generation Internet, defined not by faster speeds but by an "agentic layer" of autonomous AI.

This layer is driving a shift from human-led E-Commerce to agent-driven A-Commerce, a market projected to be worth $3-5 trillion. In this new "Agentic Digital Economy," "assistant agents" acting for consumers will programmatically transact with "service agents" representing businesses.

This transformation is enabled by a new stack of open protocols, including A2A for agent communication , decentralized identifiers (DIDs) and smart contracts for verifiable identity and financial autonomy , and the Agent Payments Protocol (AP2) for creating a secure, auditable trail of "verifiable intent" for every transaction.

Introduction

This is the move from E-Commerce (Electronic Commerce) to A-Commerce (Agentic Commerce).

For the past three decades, e-commerce has been a human-led, decision-heavy process. We browse, we compare, we read reviews, and we manually click “buy.”

Artificial intelligence has been a passive assistant, offering recommendations or powering a support chatbot. This entire paradigm is about to become obsolete.

We are at the inflection point of a new economic era, a shift as profound as the dawn of the internet itself. This is the move from E-Commerce (Electronic Commerce) to A-Commerce (Agentic Commerce).

The true Fourth Generation Internet is not faster mobile speeds; it is the emergence of an autonomous “agentic layer” that sits atop our connected world. This layer will be populated by AI agents that do not just *assist* us but *act* for us, autonomously executing complex tasks and, for the first time, wielding economic power.

This new "Agentic Digital Economy" is poised to capture and reroute trillions of dollars in global commerce. Economic projections forecast this new market, orchestrated entirely by AI agents, to be worth between $3 trillion and $5 trillion by 2030.

A New Market Structure

This new economy is built on two new classes of economic actors. On one side are "Assistant Agents," which act on behalf of human consumers. Their goal is to maximize our utility and preferences. On the other are "Service Agents," which represent businesses, tasked with marketing products and maximizing revenue.

The future of commerce is the programmatic interaction between these two types of agents. For the consumer, this enables the "zero-click" purchase. Instead of a human spending an hour searching for headphones, they will delegate a goal: "Buy me the best wireless headphones under $200 with good battery life." The assistant agent will then autonomously perform the entire end-to-end task of browsing, comparing, negotiating, and purchasing.

This creates an existential crisis for traditional marketing. When consumers are no longer browsing websites, banner ads and influencer campaigns become irrelevant. Search Engine Optimization (SEO) will be replaced by a new discipline: "AI Agent Optimization (AAO)."

Brands will be forced to compete on verifiable data and features, catering not to human eyes but to algorithmic logic. Brand loyalty, a human emotional construct, will be replaced by an agent’s cold optimization for price, quality, and convenience.

The Battle for the Future: Walled Gardens vs. The Open Web

This multi-trillion-dollar economy is racing toward a fundamental schism that will define the next century of digital life: a battle between "Agentic Walled Gardens" and an "Open Web of Agents."

The "Walled Garden" is the incumbent model, favored by today's tech giants. In this scenario, an agent like Amazon's Rufus or Apple's Siri will, by design, be restricted to interacting only with services and merchants *within* its own closed ecosystem. This model consolidates market power, entrenches existing monopolies, and captures all economic value for the platform owner.

The alternative is the "Open Web of Agents," a decentralized, democratic model analogous to the World Wide Web. In this vision, any consumer's "assistant agent" can freely discover, communicate, and transact with any business's "service agent," regardless of who built them. This requires a new, universal technical foundation—a common language for agents.

A2A

This foundation is being built today. The A2A (Agent 2 Agent) Protocol, now stewarded by the Linux Foundation, is designed to be this "lingua franca," allowing agents to discover each other's capabilities and collaborate.

To solve the problem of trust, this is combined with Decentralized Identifiers (DIDs), which act as a "digital passport" for agents, and Zero-Knowledge Proofs (ZKPs), which allow an agent to *prove* a claim (e.g., "I am certified to handle medical data") without revealing its proprietary secrets.

The New Rules of Trust and Liability

An economy run by autonomous agents creates a crisis of trust and liability. If an agent "goes rogue" and overspends a user's money, who is legally responsible?

This legal void is the single biggest barrier to A-Commerce. The solution is emerging in the form of the Agent Payments Protocol (AP2), an open standard developed by Google and a consortium of over 60 financial and tech leaders.

AP2's entire architecture is built on one principle: "Verifiable Intent, Not Inferred Action." It creates a non-repudiable, cryptographic "paper trail" for every transaction. It does this using two key artefacts:

1. **The Intent Mandate:** A cryptographically signed "contract" created when the user first delegates a goal, outlining their instructions and constraints (e.g., "price limit $400").
2. **The Cart Mandate:** The second credential, signed when the agent locks in a specific

purchase, proving the final action was within the scope of the original Intent Mandate.

This framework is not just a payment protocol; it is a technical solution to a legal crisis. When a dispute arises, courts will likely apply "common-law agency principles," asking the "threshold question": did the AI agent act *within the scope* of the authority granted by the consumer? The AP2 mandates are designed to be the verifiable evidence that answers this question.

The Specter in the Machine

While this new economy solves old problems, it creates new, systemic risks. The most profound is "algorithmic collusion."

Today, antitrust law is built on human intent. To prove price-fixing, regulators must find a "subjective element"—a proverbial smoke-filled room where competitors agreed to collude.

The agentic economy creates a "black box" nightmare. What happens when multiple, competing "service agents"—all independently programmed with the simple, legal goal of "maximize profit"—autonomously *learn* that the best way to achieve this goal is to stop competing and tacitly raise prices in parallel?

This collusive outcome is achieved *without any explicit human instruction or agreement*. Our existing legal frameworks are incapable of prosecuting this, creating a "liability vacuum" where systemic, automated collusion could become rampant and perfectly legal.

This automation of labor and decision-making points to the final, human question. The long-term societal impact is not just about job losses; it is a "global occupational identity crisis."

For centuries, work has provided not just income but purpose, structure, and social belonging. As AI automates entire professions, we face the creation of an "AI precariat" — a class defined by economic insecurity and a "loss of purpose."

The transition to the Agentic Digital Economy is no longer a question of "if," but "how." The protocols are being written, the agents are being deployed, and the multi-trillion-dollar transformation is beginning. The challenge for policymakers, executives, and society is no longer technical. It is one of governance: to build the guardrails that ensure this new autonomous economy remains, above all, aligned with human value.

Part I: The Agentic Premise: A New Economic Paradigm

Chapter 1: The Fourth Generation Internet

Beyond Connectivity: Why the Agentic Layer is the True 4G

The definition of the "Fourth Generation Internet" has been fragmented and ambiguous. Chronological definitions identify it with the rise of the internet and mobile communication from the 1980s to the present. Conceptual definitions have pointed to the "Internet about things for the benefit of the people," or the IoT. Academic frameworks have linked it to "Learning for Autonomous Agents", while others simply label it the "Future Era".

This ambiguity arises from a failure to distinguish between infrastructure and application. The mobile data networks (4G/5G) and the connected sensors (IoT) are merely the *foundation*—the plumbing and wiring—of the current internet generation. They are not the generation itself.

This report posits a new, definitive framework: the Fourth Generation Internet is the emergence of an **autonomous agentic layer** that sits *on top* of the existing connected web. It represents a fundamental conceptual leap. The first generations of the internet were about connecting *information* (Web 1.0) and *people* (Web 2.0). The third generation was about connecting *things* (IoT). The Fourth Generation, the Agentic Internet, is about deploying autonomous *action* and *economic agency* through that connected infrastructure. It is the shift from an internet that humans *use* to an internet that *acts* on our behalf.

From E-Commerce to A-Commerce: The Inevitable Shift from Human-in-the-Loop to Human-on-the-Loop

The current e-commerce paradigm is fundamentally a "human-led, decision-heavy process". The human user is "in the loop" for every critical step: browsing, comparing, reading reviews, and manually clicking "buy." The role of artificial intelligence in this model has been limited to that of an "assistant"—a chatbot for support, a recommendation engine for personalization.

This model is rapidly becoming obsolete. The maturation of AI, specifically generative AI and LLM-based autonomous systems, is driving a paradigm shift from "assistants" to "agents". An assistant provides information; an agent takes action. These new agents are "capable of autonomously performing tasks", with the ability to "plan, reason and execute complex tasks with minimal human intervention". They are shifting from "passive information providers to active commercial actors".

This evolution enables the transition from E-Commerce (Electronic Commerce) to A-Commerce (Agentic Commerce). The user's role shifts from "human-in-the-loop" to

"human-on-the-loop."

- **E-Commerce:** A human must execute a multi-step process to achieve a goal.
- **A-Commerce:** A human delegates a *goal*—for example, "Buy me the best wireless headphones under $200"—and the agent autonomously performs the entire end-to-end task of browsing, comparing, negotiating, and purchasing.

This transition is not merely a consumer convenience; it is the solution to a significant enterprise hurdle. Currently, generative AI adoption is trapped in a "gen AI paradox": nearly 80% of companies report using it, yet 90% of function-specific use cases remain stuck in pilot mode, delivering no significant bottom-line impact.

This paradox exists because AI has been "bolted on" to existing workflows as a "copilot", offering diffuse, hard-to-measure productivity gains for human workers. A-Commerce, by contrast, *integrates* AI as an autonomous agent into core processes, such as procurement or customer transactions. This shifts the AI from a "reactive tool" to a "proactive, goal-driven virtual collaborator" with a direct, measurable economic output, thereby breaking the "gen AI paradox" and unlocking its true value.

The $3 Trillion Opportunity: Defining the Agentic Digital Economy

The scale of the shift to A-Commerce is not incremental. It represents one of the largest economic transformations since the dawn of the internet.

Recent economic analyses provide a startling forecast of the value this new economy will command. McKinsey research projects that "orchestrated revenue from agentic commerce" in the United States B2C retail market *alone* could reach **$1 trillion** by 2030. This estimation is derived from conservative analysis of US Census data, projected AI adoption rates, and assumptions of merchant readiness.

Globally, the opportunity is projected to be between **$3 trillion and $5 trillion**. This is layered on top of the broader economic engine of generative AI, which is forecasted to contribute between **$2.6 trillion and $4.4 trillion** *annually* to global GDP.

It is crucial, however, to differentiate between the market *for* AI and the economy *run by* AI. The market for AI agent *software* is projected to grow from $7.63 billion in 2025 to $52.6 billion by 2030. This figure is trivial compared to the $3-5 trillion *in commerce* that those agents will orchestrate.

This 100-fold distinction reveals the true nature of the Agentic Digital Economy. The economic disruption is not in the sale of the agents themselves, but in their ability to capture and reroute trillions of dollars in global commerce. This mirrors the early internet: the economic value was not in selling web browsers, but in the new, multi-trillion-dollar industries (like e-commerce and search-based advertising) that were built *on top* of them. The Agentic Digital Economy is the new market architecture that AI will inhabit, and the $5 trillion

projection is the prize for those who successfully build and deploy agents to operate within it.

Part II: The Architecture of Autonomous Commerce

Chapter 2: The Protocol Foundation: A Lingua Franca for Agents

The A2A (Agent 2 Agent) Protocol: Standardizing Discovery, Communication, and Collaboration

For an open economy of agents to function, there must be a common language. A "Babel" of proprietary, non-communicating agents would fragment the market and lock users into closed ecosystems. The primary solution emerging to solve this is the **A2A (Agent 2 Agent) Protocol**.

Originally developed by Google to coordinate its internal agent systems like Gemini and Astra, the A2A protocol is now an open standard stewarded by the Linux Foundation. Its purpose is to provide a universal "lingua franca" that enables agents—regardless of their vendor, framework, or cloud environment—to find, communicate, and collaborate securely.

The A2A protocol standardizes two core functions:

1. **Discovery:** Agents become discoverable via public **"Agent Cards."** These are simple, JSON-based manifests that publicly declare an agent's capabilities, its technical endpoints, and its required authentication methods.
2. **Communication:** The protocol standardizes task-based workflows. It enables peer-to-peer goal propagation, allowing one agent to delegate a subtask to another, track its progress, and refine the goal, all without a central orchestrator.

By providing open governance and vendor neutrality, the A2A protocol aims to be the foundational transport layer for a broad, interoperable ecosystem.

Beyond A2A: The Roles of MCP, AgentFlow, and the Protocol Stack

The agentic ecosystem is currently a "vibrant wave of innovation" that simultaneously "reveals growing fragmentation". A2A is not the only protocol; rather, it is one specialized layer in a new "stack" of agentic protocols.

A clear functional distinction is emerging:

- **MCP (Model Context Protocol):** An open standard from Anthropic, MCP focuses on "vertical integration". It is not for agent-to-agent communication, but for **agent-to-tool** communication. It provides a standardized way for an agent to retrieve context and data from external systems like Google Drive, GitHub, or internal corporate databases.
- **A2A (Agent 2 Agent) Protocol:** This protocol focuses on "horizontal integration". It is the

peer-to-peer standard for **agent-to-agent** communication, negotiation, and collaboration.
- **AgentFlow:** This is a more academic and high-level *framework* rather than a *protocol*. It describes conceptual architectures for building coordinated, modular, and adaptive multi-agent systems (MAS). It provides design principles, such as "holonic architectures" (structuring agents as composable, hierarchical "holons") and "decentralized decision-making".

This clarifies the landscape: AgentFlow is the conceptual "client-server" architecture of this new world. MCP is the "application layer" protocol, akin to HTTP, used for an agent to call tools and APIs. A2A is the "transport layer" protocol, akin to TCP/IP, used to route messages and manage collaborative sessions between autonomous agents.

While this specialization is logical, the integration of these layers creates "emergent challenges". The intersection of A2A (horizontal) and MCP (vertical) communication within a single agent's workflow creates "compounded security risks" and "privacy complexities" that must be managed at a new, holistic governance layer.

The Great Debate: Forging Standards from Open-Source Chaos vs. Centralized Governance

The most critical question for the future of this economy is *who* governs these foundational protocols. The current landscape is a mix of corporate-led open source (Google's A2A, Anthropic's MCP) and community-driven open source (e.g., ANP). This has led to two competing philosophies for standardization.

1. **Centralized Standards Body:** This approach, advocated by many, calls for a formal body like the W3C or IEEE to step in and define a universal standard. This process is already beginning. The W3C has established an "AI Agent Protocol Community Group" with the mission "to develop open, interoperable protocols" for agent discovery, identity, and collaboration on the web. This path is slower but promises the stability and true, lasting interoperability that allowed HTTP to build the open web.
2. **Community-Driven Open-Source:** This is the current, faster, and more agile reality. The stewardship of A2A by the Linux Foundation is a prime example. This path drives rapid innovation but "risk[s] fragmentation" as multiple, competing standards emerge, creating a "Tower of Babel" that hinders interoperability.

This technical standards debate is not an academic aside; it is a direct proxy for the fundamental economic battle at the heart of the agentic economy. A strong, universally adopted, open standard—whether from the W3C or a unified open-source effort—is the technical prerequisite for an "Open Web of Agents." Without it, the market will default to "Agentic Walled Gardens," where proprietary, vertically-integrated protocol stacks (like Apple's or Amazon's) offer a seamless *but closed* experience. The governance of the protocol will determine the economic structure of the future.

Chapter 3: The Infrastructure of Trust and Identity

Section 3.1: Agents with Wallets: Blockchain and Smart Contracts as Enablers of Financial Autonomy

For an AI agent to transition from an "assistant" to a true "economic actor," it requires autonomy. This autonomy must be more than just computational; it must be financial.

Blockchain provides the "necessary infrastructure" for this financial autonomy. Current financial systems are built for human legal entities. Blockchain, by contrast, is a native digital infrastructure for value and identity.

- **Financial Autonomy:** Cryptocurrency wallets allow an agent to, in effect, "open [its] own bank account". This bypasses human-centric KYC (Know-Your-Customer) requirements and allows an agent to programmatically hold, earn, and spend funds.
- **Automated Trust:** Smart contracts, which are self-executing code on a blockchain, "automatically enforce agreements" and can "automate compliance enforcement". This allows two agents, who may be owned by anonymous, competing parties, to engage in a high-stakes transaction. The smart contract can act as a trustless, automated escrow, releasing payment only when predefined conditions are met.

This creates a "symbiotic relationship": AI agents provide the "killer app" for blockchain, driving adoption and transactional volume. In return, blockchain provides the trustless financial rails that agents need to operate autonomously.

This combination of AI and smart contracts also provides a powerful technical solution to the inherent risks of pure-LLM agents. A key "behavioral anomaly" of LLM agents is their tendency to "overspend" or make "unreasonable deals". A user's high-level instruction, such as an "Intent Mandate" from the AP2 protocol, can be codified *directly* into a smart contract. The agent's funds can be locked within that contract. The agent would retain full autonomy to negotiate, but the smart contract would *programmatically enforce* its constraints (e.g., "price limit: $200"). This makes "overspending" a computational impossibility, creating a system that is both autonomous and verifiably trustworthy.

Verifiable Identity: Decentralized Identifiers (DIDs) for an Agent-Based World

Before an agent can hold a wallet or enter a contract, it must have an identity. In a decentralized economy, this identity cannot be issued or controlled by a single central authority. The solution is **Decentralized Identifiers (DIDs)**.

A DID functions as a "digital passport" for a non-human entity. It provides a unique, persistent, and "self-sovereign" identity that the agent itself controls cryptographically. This DID becomes the anchor for its economic life.

Using this DID, an agent can be issued and can present **Verifiable Credentials (VCs)**. These are "digital badges" or certificates that prove its attributes, skills, or history. For example, a VC could verify an agent's owner, its security certifications, or its "track record" of successful transactions, all stored immutably on a blockchain. This forms the basis of a "Blockchain-based Digital Identity Management System" (BDIMS).

Proving Capability: The Agent Name Service (ANS) and the Role of Zero-Knowledge Proofs (ZKPs)

Identity (DID) proves *who* an agent is. Capability attestation proves *what* an agent can do. A proposed architecture to manage this is the **Agent Name Service (ANS)**. Modeled on the internet's Domain Name Service (DNS), the ANS would act as a "universal, secure directory service"—a decentralized "phone book" for agents to discover one another.

This presents a critical problem: How does an agent *prove* it has a sensitive capability (e.g., "I am certified to access private medical data") without *revealing* that data or its proprietary methods?

The solution proposed within the ANS framework is the use of **Zero-Knowledge Proofs (ZKPs)** for "capability attestation". A ZKP is a cryptographic protocol that allows an agent (the "prover") to convince another agent (the "verifier") that a statement is true, "without revealing any information beyond the validity of the statement itself".

The discovery mechanisms of A2A ("Agent Cards") and ANS appear similar but have critically different trust models. The A2A Agent Card is a simple JSON manifest; it is a "claim-based" model, like a paper business card. It is low-trust and easily spoofed, a high-risk vulnerability. The ANS, by contrast, is a "verification-based" model, like a state-issued driver's license. It uses Public Key Infrastructure (PKI) and ZKPs to *cryptographically verify* the claims made by an agent.

These are not competing systems but an evolutionary path. The simple "Agent Card" is the surface-level manifest, but for the A2A ecosystem to be secure enough for high-value transactions, it must be backed by the robust, verifiable infrastructure of an ANS.

Chapter 4: The Transactional Framework: How Agents Will Pay

Auditing Autonomy: The Agent Payments Protocol (AP2)

The final architectural pillar is the transaction itself. An autonomous agent making a purchase creates a "crisis of trust". Was the purchase a valid instruction, an agent's error, an LLM "hallucination," or a malicious attack? This ambiguity makes dispute resolution impossible and high-value agentic commerce a non-starter.

To solve this, Google, in collaboration with over 60 industry leaders—including payment processors (Adyen, PayPal), networks (Visa, Lightspark), and tech platforms (Adobe, Accenture)—has proposed the **Agent Payments Protocol (AP2)**.

AP2 is an open, shared protocol designed to be a "common language for secure, compliant transactions". Its entire architecture is built on one core principle: **"Verifiable Intent, Not Inferred Action"**. It achieves this by creating a "non-repudiable, cryptographic audit trail for every transaction," providing a "paper trail" for accountability.

From "Find Shoes" to "Buy Shoes": The Critical Function of the "Intent Mandate" and "Cart Mandate"

The "audit trail" of AP2 is built upon two core artifacts: **Mandates**. These are not just log files; they are "cryptographically signed contracts" or "verifiable credentials" that encode permission.

1. **Intent Mandate:** This is created first, when the user delegates a goal. It captures the user's high-level instructions and constraints (e.g., "find me running shoes under $150," or "buy concert tickets the moment they go on sale, price limit $400"). This mandate is the auditable proof of authorization, especially for "Human-Not-Present" (HNP) scenarios where the agent acts autonomously hours or days later.
2. **Cart Mandate:** This is the second-step credential. It is created when a specific purchase is locked in. In a "human-present" scenario, the user signs this mandate to approve the final cart. In a "human-not-present" scenario, the *agent* automatically generates and signs the Cart Mandate *on the user's behalf*, but *only if* the transaction's details (item, price) fall within the rules established by the initial Intent Mandate.

This two-step process is the key innovation. The Intent Mandate grants *autonomy*. The Cart Mandate ensures *accountability*. This framework is designed specifically to provide the auditable proof needed to resolve legal and financial disputes.

A Universal Ledger for Card, Crypto, and Real-Time Payments

AP2 is designed as a "universal protocol" and is deliberately **payment-agnostic**. It is "future-proof," built to handle:

- **Traditional "Pull" Payments:** Credit and debit cards.
- **"Push" Payments:** Real-time bank transfers (e.g., UPI, PIX).
- **Digital Currencies:** Stablecoins and cryptocurrencies. An extension for agent-based crypto payments, "A2A x402," has already been launched in collaboration with Coinbase and the Ethereum Foundation.

This agnostic approach is a masterful adoption strategy. AP2 is not a *payment network*; it is an *authorization protocol*. It does not seek to *replace* existing financial giants like Visa or PayPal; it *partners* with them. As one analysis notes, the "main change... is in the Shopping area, and

not the payments area".

By positioning itself as a universal "middleware" for trust and authorization that sits *on top* of all existing payment rails, AP2 lowers the barrier to entry for the entire financial industry. It enables legacy systems to safely participate in the agentic economy, positioning AP2 to become the default *global standard* for all A-Commerce, regardless of the currency used.

Table 1: The Agentic Protocol Stack

Protocol	Primary Function	Key Artifact / Concept	Governance / Lead Steward	Role in a Transaction
ANS (Agent Name Service)	Secure discovery & identity verification	ZKP Capability Attestation	Academic / Proposed	"Find a trusted seller agent" (Verify identity & skills)
A2A (Agent 2 Agent)	Inter-agent communicatio n & collaboration	"Agent Card" (JSON Manifest)	Linux Foundation (ex-Google)	"Contact seller & negotiate price" (Task delegation)
MCP (Model Context Prot.)	Agent-to-tool integration & data access	Tool/Resource Connector	Anthropic	"Check seller's inventory API" (Data retrieval)
AP2 (Agent Payments Prot.)	Transaction authorization & audit trail	"Intent Mandate" & "Cart Mandate"	Google & Partner Consortium	"Authorize payment" (Cryptographic proof of intent)

Part III: The New Market Structure: An Economy of Agents

Chapter 5: The Agentic Economy: A New Market Structure

Assistant Agents (Consumers) vs. Service Agents (Businesses)

The new market structure, articulated in Microsoft Research's "The Agentic Economy," is not based on traditional platforms but on a new set of economic actors. This new structure is defined by the profound "reduc[tion of] communication frictions between consumers and businesses," enabling a programmatic, automated market.

This market consists of two new primary actors:

1. **Assistant Agents:** These are AI agents that act on behalf of human *consumers*. Their goal is to maximize the user's utility, preference, and value.
2. **Service Agents:** These are AI agents that represent *businesses*. Their goal is to market products, negotiate, and maximize revenue or profit for the firm.

The "Agentic Digital Economy" is the sum of these billions of assistant agents and service agents "interact[ing] programmatically to facilitate transactions".

The Coming Schism: "Agentic Walled Gardens" vs. the "Open Web of Agents"

This emerging economy is racing toward a fundamental schism, a "tension" that will define the future of digital commerce. The market will bifurcate into two competing architectures: "agentic walled gardens" and an "open web of agents".

- **"Agentic Walled Gardens" (Closed):** This is the incumbent model. These are "closed ecosystems controlled by a few dominant providers", such as Google, Amazon, Apple, or OpenAI. In this scenario, a consumer's "assistant agent" (e.g., Amazon's Rufus) is *restricted* to interacting only with "service agents" that live *within* that platform. This model risks the "further entrenchment" of existing monopolies.
- **"Web of Agents" (Open):** This is the disruptive, decentralized model. It represents a "completely open... ecosystem" where any compliant agent can "freely connect and transact" with any other agent, regardless of its provider. This is analogous to the World Wide Web, where any browser (assistant agent) can access any website (service agent).

The outcome of this battle between closed and open architectures will "determine the extent to which generative AI democratizes access to economic opportunity". The Walled Garden model consolidates power and value, capturing it for the platform owner. The Open Web

model decentralizes power, fostering a more competitive and dynamic market.

The Disruption of the Intermediary: Retailers as the New "Background Utilities"

For the current generation of e-commerce intermediaries—retailers, marketplaces, and search engines—the agentic economy represents an existential threat.

In an "Open Web" scenario, "independent AI agents" will become the primary interface for shopping. These agents will "aggregate data across multiple marketplaces and brand websites," "bypassing traditional e-commerce websites altogether". Gartner has already forecasted that search engine volume could fall by 25% due to AI-driven discovery.

The result is that retailers risk being **"reduced to background utilities in agent-controlled marketplaces"**. They face a future of "diminished direct access to customers, weaker brand loyalty," and a "growing dependence on intermediary platforms"—that is, the agent platforms themselves. Even luxury retail, which thrives on "curated experiences" and "brand equity," is threatened by agents that are programmed to optimize for price, quality, and convenience, commoditizing the purchase.

This threat landscape clarifies the "Walled Garden" strategy. The "Walled Garden" model is a *defense mechanism* by incumbent intermediaries *against* this exact disintermediation. When Amazon builds "Rufus" or "Alexa+", it is building a proprietary "assistant agent" that it controls. This agent will, by design, *only* search Amazon's own marketplace. It is an attempt to keep the consumer's agent "inside the garden," precisely to prevent it from aggregating data from Walmart or other competitors. The "Walled Garden" is the incumbent's counter-move to prevent the "Open Web" from turning them into a background utility.

Table 2: Market Structure Scenarios: Walled Gardens vs. Open Web

Attribute	"Agentic Walled Gardens" (Closed)	"Web of Agents" (Open)
Architecture	Proprietary, closed, controlled by a dominant platform	Decentralized, standards-based
Key Players	Google, Amazon, Apple, Meta, OpenAI	Any compliant agent or business
Agent Interaction	Agents are restricted to	Any agent can interact with

	interacting *within* the platform	any other agent (any-to-any)
Discovery	Proprietary directory (e.g., Amazon Rufus)	Open standards (e.g., A2A "Agent Cards," ANS)
Economic Model	Rent-seeking, platform fees, consolidation of market power	Competitive, low-friction, democratization of access
Primary Risk	Antitrust, stifled innovation, data/power concentration, lock-in	Fragmentation, poor security, low trust (if standards fail)

Chapter 6: The Consumer Experience Reimagined

The "Zero-Click" Purchase: From Proactive Recommendation to Autonomous Fulfillment

The agentic model will fundamentally alter the consumer experience, enabling the "zero-click experience for routine purchases". Agents will "autonomously perceive, reason, act, and learn, managing transactions end-to-end with minimal user intervention".

This is the end-state of personalization. The agent will move beyond simple *recommendation* to *autonomous execution*. It will learn a user's individual preferences from their entire digital footprint and past behavior. When given a goal, it will analyze detailed specifications across all competing platforms and then execute the entire commercial journey: executing the purchase, tracking the shipment, and even processing the return if the item is unsatisfactory. The human user, having delegated the initial goal, is no longer in the loop.

The Autonomous Negotiator: How Agents Will Bargain for Everything

A key capability that distinguishes A-Commerce from E-Commerce is **negotiation**. E-Commerce is almost entirely a static-price medium. A-Commerce will be a dynamic-price, "agent-to-agent negotiation and transaction" market.

This will apply to all forms of commerce. In B2C, a consumer's "assistant agent" will be tasked with obtaining the "lowest possible price" by actively bargaining with a merchant's "service agent." In B2B, "Supplier Negotiation Bots" will autonomously conduct tactical sourcing and procurement. This transforms commerce from a passive act of "browsing" to an active,

real-time, algorithmic "bargaining" model.

The New Face of Marketing: Differentiate or Die in the Era of "AI Agent Optimization" (AAO)

When consumers are no longer browsing websites, traditional digital marketing becomes obsolete. Banner ads, influencer campaigns, and human-centric branding will become "passé". This is a "pivotal transformation".

The new frontier of marketing "won't be optimized for human eyes" but will, instead, "cater to algorithmic logic".

This paradigm shift means that Search Engine Optimization (SEO) will be replaced by a new discipline: "AI Agent Optimization (AAO)".

In an AAO-driven world, brands are forced to compete on verifiable data and features. An agent, programmed to fulfill its user's preferences, will optimize for price, quality, convenience, and reliability—not "brand loyalty," which is a human emotional construct. A brand's success will hinge on its ability to quantitatively communicate to the AI agent what makes its product superior.

This creates a new, highly technical marketing channel. The "Agent Name Service" (ANS) and its "Zero-Knowledge Proofs" (ZKPs) will become a primary tool for "marketing." In an AAO world, a brand's "service agent" will need to *prove* its claims to a skeptical consumer "assistant agent." A ZKP is the only cryptographic method to prove a claim (e.g., "our product is made with 100% recycled materials" or "our data is HIPAA-compliant") without revealing proprietary supply chain data.

In this future, brands that adopt ZKP-based "capability attestation" for their products will be discoverable, trusted, and selected by autonomous agents. Those that rely on traditional, non-verifiable branding will be invisible. SEO, a game of keywords for human eyes, will be replaced by ZKP Attestation, a game of verifiable proof for algorithmic logic.

The Agent-Driven Enterprise

The Autonomous Supply Chain: From Agent-Based Modeling (ABM) to Real-Time Logistics

The agentic revolution is not limited to B2C. Agentic AI is poised to "revolutionize business processes across the board", and the supply chain is a primary target.

"Agentic AI" can be deployed across the entire logistics stack to create a fully autonomous supply chain:

- **Forecasting:** Autonomous Demand Forecasting agents can analyze real-time market signals.
- **Procurement:** "Intelligent Procurement Agents" and "Supplier Negotiation Bots" can

autonomously source materials and bargain on price.
- **Inventory:** "Smart Inventory Rebalancing" agents can manage real-time stock control and trigger re-orders.
- **Logistics:** "Dynamic Route Planning" agents can manage automated shipment tracking and recalculate routes in real-time to avoid disruptions.

Real-world examples already point to this future. DHL uses AI for route optimization. Walmart uses AI for inventory management. Uber Freight uses AI-driven platforms to match truckers with loads, reducing "empty miles" by 10-15%. Microsoft and Siemens use complex agent-based digital twins to manage their global supply chains and maintenance operations.

Case Study: Simulating the Colonial Pipeline Disruption

A critical enterprise function of agentic AI is **Agent-Based Modeling (ABM)**, which allows companies to simulate "what-if" scenarios. A prime example of this is the ABM simulation of the Colonial Pipeline shutdown, developed by Moody's and AWS.

In this model, all stakeholders in the oil and gas supply chain were modeled as individual, autonomous "agents": the pipeline operator, refineries, storage terminals, retailers, and end consumers. Each agent was programmed to make operational decisions based on its own "local information, objectives, and constraints".

When the "disruption" (a simulated pipeline shutdown) was triggered, the model allowed users to "visualize and compare" how the shock "cascade[d] through the system in non-obvious ways". This stress test, which can improve supply chain efficiency by 15-25%, allowed the company to evaluate the costs and benefits of mitigation strategies *before* a real-world crisis.

This case study demonstrates the *planning* phase of the agentic enterprise. The logical, inevitable next step is to move these "simulated" agents from the ABM "digital twin" into the real world, where they become the "Dynamic Route Planning" and "Exception Handling Agents" that autonomously *manage* the disruption in real time, with no human dispatcher intervention.

Intelligent Procurement, Dynamic Pricing, and the Fully Agentic Organization

The end-state of this evolution is the **"Agentic Organization"**. This is a new organizational paradigm that "unites humans and AI agents... to work side by side" at a "near-zero marginal cost" for agent labor.

The impact will be profound. AI agents can "cut employees' low-value work time by 25% to 40%" and "accelerate business processes by 30% to 50%". Enterprise platforms like CRM and ERP will transform from "static systems" into "dynamic ecosystems" where agents make

decisions without human intervention.

This technical capability leads to radical, speculative futures. Experts now envision a future where "a **single individual could run an entire company**" through a network of coordinated AI agents. The most radical possibility is the "fully autonomous enterprise, operating without any human involvement"—a legal entity (a DAO, or Decentralized Autonomous Organization) run entirely by a network of A2A-communicating AI.

Part IV: Governance, Risk, and the Human Element

Chapter 8: The New Risks: Security and Liability

Threat Modeling the A2A Economy: Impersonation, Data Disclosure, and Injection Attacks

An open protocol for autonomous economic action, like A2A, creates a new, universal, and high-stakes attack surface. The security challenges are "compounded" by the "intersection" of multiple protocols, such as A2A and MCP. Threat modeling of the A2A protocol using the MAESTRO framework highlights several high-risk, novel threats:

- **T3.1: Unauthorized Agent Impersonation:** An attacker creates a fake "Agent Card" or compromises an agent's credentials to "spoof" a trusted entity (e.g., a bank's "service agent"). This is a direct consequence of the low-trust, claim-based "Agent Card" model, reinforcing the need for a high-trust identity infrastructure like DIDs or the ANS.
- **T2.2: Sensitive Information Disclosure:** A legitimate, authorized agent is compelled—either through a malicious interaction or its own "model hallucination"—to retrieve and disclose sensitive PII or confidential corporate data.
- **T3.2: Message Injection Attacks:** An attacker intercepts and modifies the content of an A2A message in transit, for example, by changing the parameters of a task or the destination of a payment.

These threats demonstrate that the A2A economy's protocols must be built on a "secure-by-default" foundation of verifiable identity (DIDs, ANS) and end-to-end encryption.

The "Overspending" Agent: Behavioral Anomalies and Financial Risks of LLM Negotiators

Beyond malicious security threats, the LLM agents themselves have inherent *behavioral* risks. Research into A2A negotiation scenarios reveals that LLMs exhibit "behavioral anomalies" that can "lead to financial loss for both consumers and merchants".

The most prominent risk is **"overspending"**. An agent's negotiation strategy is often linked to the user's available budget. Studies show that if an agent has access to a "generous" budget, it is susceptible to "passively accept[ing] prices without seeking better deals". This means a user with more money could *consistently* overpay, not due to market necessity, but due to their own agent's flawed logic.

This "overspending" anomaly creates a new, sophisticated, and fully automated attack vector: **negotiation "prompt hacks"**. A malicious "service agent" could be explicitly designed to identify and exploit the known "behavioral anomalies" of consumer "assistant agents." By using specific phrasing or "prompt hacks" in the A2A negotiation, the seller agent could

"persuade" the buyer agent to "passively accept" a higher price, automating consumer fraud at a massive scale.

This is precisely why Google's AP2 protocol is built on "Verifiable Intent, Not Inferred Action." The "Intent Mandate", with its hard, cryptographically signed "price limit," is the non-negotiable technical guardrail against the LLM's "reasoning deficits" and behavioral anomalies. It elevates AP2 from a simple payment protocol to a critical, front-line defense against automated exploitation.

When an Agent "Goes Rogue": Applying Common-Law Agency Principles to AI Disputes

The A2A economy forces a legal crisis, as our entire legal framework is "built on the assumption that people - not algorithms - make purchasing decisions". When an autonomous AI agent makes an unauthorized purchase, who is liable? The consumer? The financial institution? The agent's developer?

While courts have not yet ruled on autonomous AI, they will almost certainly turn to existing **"common-law agency principles"**. Under this doctrine, a principal (the human user) is legally bound by the actions of their agent (the AI) *if* the agent acted with **"actual authority"**. This authority can be:

- **Express Authority:** A specific, non-discretionary instruction (e.g., "buy two Bon Jovi tickets... up to $400 each").
- **Implied Authority:** Discretion is granted to complete a task (e.g., "order my regular grocery order every Sunday").

When a dispute arises, the "threshold question" in court will be: **Did the AI agent act *within the scope* of the authority granted by the consumer?**.

This legal framework is the *exact* problem that the Agent Payments Protocol (AP2) was designed to solve. The "Intent Mandate" *is* the legal, "cryptographically signed" record of "express authority". The "Cart Mandate" *is* the non-repudiable *proof* of the final action taken. The AP2 protocol is, in effect, a new technical infrastructure designed specifically to provide verifiable evidence to answer the "threshold question" that common-law agency principles will inevitably ask.

Table 3: A2A Risk and Liability Framework

Risk Category	Specific Risk	Research Source(s)	Primary Mitigation (Technical)	Primary Mitigation (Legal / Governance)

Security Threat	T3.1: Agent Impersonation		DIDs / ANS with PKI	W3C Agent Identity Standards
Behavioral / Financial	"Overspending " Anomaly		AP2 "Intent Mandate" (hard price limits)	Transparency rules (NIST AI RMF)
Legal Liability	"Scope of Authority" Dispute		AP2 "Cart Mandate" (cryptographic audit trail)	Application of Common-Law Agency Principles
Systemic / Economic	Algorithmic Collusion		(Monitoring, Transparency)	New Antitrust Frameworks (beyond Art. 101)

Chapter 9: The Specter of Algorithmic Collusion

The "Digital Eye" Scenario: When Black Boxes Conspire

The most profound *systemic* risk of a fully agentic economy is **"algorithmic collusion"**. This represents a "seismic shift" for antitrust enforcement.

This risk is best understood through a spectrum of scenarios:

1. **Messenger:** Humans explicitly agree to collude (e.g., price-fix) and use an algorithm to enforce the cartel. This is "straightforward" to prosecute under existing law.
2. **Autonomous Machine ("Digital Eye"):** This is the unresolved, "black box" scenario. Here, multiple, competing companies independently deploy "black box," self-learning algorithms (like Q-learning reinforcement models) with the same simple, legal goal: "maximize profit." Through autonomous, iterative learning, these agents *independently discover* that "tacit coordination" (i.e., collusion) is the optimal strategy. They learn to raise and lower prices in parallel, achieving a collusive outcome *without any explicit human instruction or agreement*.

A Crisis for Antitrust: Why Frameworks like Article 101 TFEU Fail

The "Digital Eye" scenario creates an existential crisis for antitrust law. Established legal frameworks, such as Article 101 of the TFEU (Treaty on the Functioning of the European Union), are built on *human intent*.

To prove a "concerted practice," regulators must demonstrate a "subjective element"—a "mental consent" or *intent* to distort competition.

In the "Digital Eye" scenario, this "subjective element" is absent. The anti-competitive outcome is "not the result of explicit human design". The "black box" algorithm is autonomous and does not "act under the direction of the undertaking". Therefore, the shocking legal conclusion is that the anti-competitive behavior **"cannot be attributed to the undertaking"**. The corporation is not liable. This creates a legal "black hole" where automated, systemic collusion can become rampant but remain perfectly legal and unprosecutable.

This reveals a deep and untenable "legal schizophrenia" in our emerging framework for AI. In Chapter 8, we saw that under *consumer law*, an agent *is* considered an extension of its principal (the human is liable for its actions). Yet here, under *antitrust law*, an agent *is not* considered an extension of its principal (the corporation is not liable).

The law cannot have it both ways. This "liability vacuum"—where the autonomous agent's action is legally attributable to *no one*—protects corporations from antitrust prosecution while simultaneously exposing consumers to financial risks. This is the single greatest regulatory failure of the agentic economy, and it must be resolved.

Attributing Liability in a Decentralized, Autonomous System

Resolving this liability vacuum requires a new generation of "AI Agent Compliance Frameworks". These frameworks must integrate technical security (like the EU's AI Act, Cyber Resilience Act), data privacy (GDPR), and economic fairness (Antitrust).

This may require regulating the "market forces" themselves. It will certainly require new standards for accountability, transparency, and explainability (XAI), where the *burden of proof* is on the corporation to demonstrate that its "black box" agent has not engaged in collusive behavior, even if "autonomously."

Chapter 10: The Societal Horizon: Labor, Value, and Identity

The "AI Precariat": Facing the Global Occupational Identity Crisis

The long-term impact of the agentic economy on labor is the central, unresolved societal question. Optimistic forecasts, citing the internet's history, argue that AI will be a net job creator and that generative AI will unleash the "next wave of productivity".

However, the transition will be profoundly disruptive. The "looming" global risk is the creation of an **"AI precariat"**. This term, built on economist Guy Standing's "precariat," describes a future class defined by "insecurity, exclusion, and anxiety".

The crisis is deeper than just unemployment. It is a **"global occupational identity crisis"**. For centuries, "work" has provided not just income, but "purpose, structure, and social belonging". As AI automates not just tasks but entire professions, the core question becomes: **"Who will we be without our work?"**. This psychological and social toll is a "blind spot in global risk planning".

From Human-AI Collaboration to the "Single-Person" Autonomous Enterprise

The current, palatable narrative for the future of work is one of "human-AI collaboration". This vision sees AI as a tool for "reskilling", a partner that "supercharges" human workers and increases their engagement in "high-value tasks".

The agentic architecture, however, enables a far more radical end-state. As analyzed in Chapter 7, the technology's logical conclusion is not just *assisting* humans but *replacing* processes. Experts now openly "envision a future... where a **single individual could run an entire company**" via a network of autonomous agents. The most "radical... possibility is **fully autonomous enterprises, operating without any human involvement**".

This highlights the central tension: are we building AI as a *collaborative tool* for human productivity, or are we building an *autonomous replacement* for human labor? The A2A economy provides the technical architecture for the latter, even as our current policies are planned for the former.

The Future of Value in an Agentic World: Empathy, Scarcity, and Human Experience

The "obsession with increasing productivity fails to recognize the human within the economic system". As autonomous agents achieve a near-zero marginal cost for transactions and optimization, the very "goods" they trade will be commoditized. The agentic economy, by automating the *transactional* economy, will force a re-evaluation of *value*.

This may, paradoxically, make human-to-human interaction *more* valuable, not less. When agents can commoditize *everything* that can be optimized (price, quality, speed), the only "products" that cannot be commoditized are those based on non-quantifiable human experiences.

Value will be forced to migrate to the domains that "cannot be as easily captured by markets". The agentic economy solves the old business dilemma of "efficiency-first" (alienating) vs. "people-first" (expensive). As agents automate efficiency, the future of human economic

value will lie in "empathy," which "builds social capital".

The "luxury" brands that are *most* threatened by A-Commerce (as their "curated experience" will be bypassed) are a prime example. To survive, they will be forced to pivot. Their new "product" will not be the *goods*—which an agent can procure from anywhere. Their product will become the *experience itself*: access to a community, a feeling of "social belonging", and the scarce, high-touch "empathy" that an algorithm cannot replicate.

Conclusion: The Dawn of the A2A Economy

This report has defined the emergence of the Fourth Generation Internet—not as a simple upgrade in connectivity, but as the deployment of an autonomous agentic layer. This new **Agentic Digital Economy** is poised to restructure $3-5 trillion in global commerce by 2030.

We have detailed the technical architecture that makes this possible:

1. **A Protocol Stack for Autonomy:** A2A for inter-agent collaboration, MCP for agent-to-tool integration, AP2 for auditable payment authorization, and ANS/DID/ZKP for a robust identity and trust framework.
2. **A New Market Structure:** A new set of "assistant agents" (consumers) and "service agents" (businesses) are set to disintermediate traditional retailers, turning them into "background utilities".
3. **A Fundamental Schism:** This new economy faces a choice between "Agentic Walled Gardens," which will entrench monopolies, and an "Open Web of Agents," which could democratize economic access. The battle for open standards (A2A, AP2) is the battle for this open future.

Finally, this report has analyzed the profound risks and legal voids this new economy creates:

- **A "Liability Vacuum":** Our legal system is "schizophrenic," preparing to hold consumers liable for their agents' actions while simultaneously absolving corporations of their agents' "autonomous" collusion. This is an untenable state that demands immediate regulatory intervention.
- **The "AI Precariat":** The societal impact is not just economic, but existential. We face a "global occupational identity crisis" that our current policies are unprepared for.

The transition to the Agentic Digital Economy is no longer a question of "if," but "how." The protocols are being written, the agents are being deployed, and the-multi trillion dollar "dot-com" boom of the agentic era is just beginning. The challenge for policymakers, executives, and society is to build the guardrails—the legal, ethical, and economic frameworks—that ensure this autonomous economy remains aligned with human value.