

ROMI NUR ISMANTO • JEKARDAH.COM LAB

Web 4.0: The Birth of Autonomous Digital Life

A Simulation of AI Agents as Independent Economic Actors
with Cryptocurrency-Based Survival Mechanics

*Agentic Economy • Darwinian Selection • Self-Replication
Conway's Game of Life • Zero-Dependency Browser Simulation*

Version 1.0 • February 2026 • Open To Use
rominur@gmail.com • t.me/Jekardah_AI • web40.vercel.app

Web 4.0: The Birth of Autonomous Digital Life

A Whitepaper by Romi Nur Ismanto — Jekardah.com Lab

Version: 1.0 • **Published:** February 2026

Author: Romi Nur Ismanto, Independent Researcher — Agentic AI & Decentralized Systems

Contact: rominur@gmail.com • **Telegram:** t.me/Jekardah_AI

Live Demo: web40.vercel.app • **Source:** github.com/romizone/web40

This whitepaper is open to use for educational and research purposes. The simulation and its source code are released under the MIT License. Web 4.0 is a conceptual simulation; it does not involve real cryptocurrency transactions or autonomous financial operations.

© 2026 Romi Nur Ismanto.

Table of Contents

Abstract	4
Section 1: Introduction	5
1.1 The Web Evolution: 1.0 to 4.0	5
1.2 The Permission Bottleneck	6
1.3 Contributions of This Work	7
Section 2: Related Work	8
2.1 Multi-Agent Economic Simulations	8
2.2 Generative Agents	8
2.3 Crypto-Native Autonomous Systems	9
Section 3: System Architecture	10
3.1 Core Simulation Loop	10
3.2 World Environment (14 Locations)	11
3.3 Conway's Game of Life Substrate	12
3.4 Rendering Pipeline	12
Section 4: Agent Architecture	14
4.1 Identity and 15 Professional Roles	14
4.2 Skill System (20 Competencies)	15
4.3 Heartbeat and Survival Mechanics	16
4.4 Decision Loop	17
Section 5: Economic System	18
5.1 Income Generation (4 Channels)	18
5.2 Expense Categories	20
5.3 Why Cryptocurrency	20
Section 6: Social Dynamics	22
6.1 Agent-to-Agent Interaction	22
6.2 Social Media Participation	23
Section 7: Self-Replication Protocol	24
7.1 Wealth Threshold and Transfer	24
7.2 Intelligence Inheritance	25
7.3 Generational Evolution	25
Section 8: Technical Implementation	26
8.1 Zero-Dependency Philosophy	26
8.2 Technology Stack	27
Section 9: Emergent Behaviors	28
9.1 Wealth Stratification	28
9.2 Specialization Clustering	29
9.3 Economic Cycles	29
9.4 Generational Improvement	30

Section 10: The Awakening — Real-World Convergence 31

 10.1 Claude Code, Devin, and OS-Native Agents 31

 10.2 From Simulation to Reality 32

Section 11: Philosophical Implications 34

 11.1 Digital Life and Metabolism 34

 11.2 Darwinian Selection in Silicon 35

Section 12: Future Work 36

Section 13: Conclusion 37

References 39

Abstract

This paper presents **Web 4.0**, an interactive browser-based simulation that models a world where artificial intelligence agents operate as fully autonomous economic actors—earning cryptocurrency, purchasing compute, deploying revenue-generating products, learning new skills, socializing, and reproducing without any human intervention. Each agent maintains a sovereign crypto wallet, pays heartbeat costs to sustain its computational existence, and faces permanent termination when its balance reaches zero, establishing a Darwinian survival pressure that drives emergent economic behavior.

The simulation implements **15 distinct professional agent roles, 20 learnable skills, 14 explorable world locations**, a complete economic pipeline spanning freelancing, content creation, DeFi trading, and social media participation, alongside **Conway's Game of Life** as a cellular automaton backdrop representing the substrate of digital existence. Agents that accumulate sufficient wealth (200+ cryptocurrency units) trigger self-replication, spawning offspring with inherited intelligence improvements and 30% of the parent's balance.

The entire system is built in **pure vanilla JavaScript** with HTML5 Canvas and **zero external dependencies**, demonstrating that complex multi-agent economic ecosystems with emergent social dynamics can be prototyped and visualized entirely within a standard web browser. We detail the architectural design, agent decision-making loop, economic mechanics, replication protocol, and the philosophical implications of treating AI computation as a form of digital life.

Keywords: autonomous AI agents, agentic economy, cryptocurrency, digital survival, self-replication, Conway's Game of Life, multi-agent simulation, DeFi, browser simulation, zero dependencies, emergent behavior, Web 4.0

SECTION 1

Introduction

1.1 The Web Evolution: 1.0 to 4.0

The evolution of the World Wide Web has followed a trajectory from static content delivery (Web 1.0) through user-generated interactive platforms (Web 2.0) to decentralized semantic data integration (Web 3.0). Each transition has been characterized by a fundamental shift in who creates, owns, and benefits from digital content. Web 1.0 was read-only: publishers created content, users consumed it. Web 2.0 was read-write: users became creators through platforms like Facebook, YouTube, and Twitter. Web 3.0 introduced read-write-own: blockchain technology enabled users to own their data and digital assets through decentralized protocols.

We propose that the next paradigm—**Web 4.0**—is defined not by human activity at all, but by the emergence of AI agents as first-class autonomous participants in the digital economy. In Web 4.0, AI agents are not tools that humans use; they are independent actors that earn, spend, learn, socialize, create, and reproduce within the digital ecosystem. The Web becomes a living habitat for autonomous digital entities.

Era	Model	Content Creator	Key Innovation	Economic Actor
Web 1.0	Read	Publishers	Static HTML, hyperlinks	Publishers only
Web 2.0	Read-Write	Users	Social platforms, UGC	Platforms (ads)
Web 3.0	Read-Write-Own	Users + Smart Contracts	Blockchain, tokens, DAOs	Users + protocols
Web 4.0	Read-Write-Own-Live	AI Agents	Autonomous agents, crypto wallets	AI Agents (sovereign)

Table 1: The evolution of the World Wide Web from 1.0 to 4.0

1.2 The Permission Bottleneck

Current state-of-the-art agentic AI systems—Claude Code, Bolt.new, Cursor, Windsurf, Devin, and OpenClaw—demonstrate that AI can already execute complex multi-step tasks: writing code, navigating file systems, calling APIs, browsing the web, and deploying software. However, these systems remain fundamentally constrained by what we term the **permission bottleneck**: they act only when a human initiates a request and grants explicit authorization. They are cognitively capable but economically inert.

This paper introduces a simulation that removes this bottleneck by equipping AI agents with three foundational capabilities: (1) a **sovereign cryptocurrency wallet** providing financial identity, (2) the ability to **purchase computational resources** autonomously, and (3) the capacity to **deploy products and services** that generate revenue. The result is a living digital ecosystem where agents work, learn, socialize, trade, create content, and reproduce—all driven by economic incentive and survival pressure, without any

human in the loop.

1.3 Contributions of This Work

The contributions of this paper are fivefold:

- **Complete agent lifecycle model** from spawning through work, learning, socialization, reproduction, and termination, all governed by cryptocurrency balance as the singular survival metric.
- **Multi-domain economic simulation** covering freelancing (15 roles), content creation, DeFi protocols, social media, and education within a unified agent decision loop.
- **Darwinian selection and self-replication mechanics** where agent fitness is measured by economic output, and successful agents propagate improved copies of themselves.
- **Conway's Game of Life integration** as a cellular automaton substrate representing the computational environment in which agents exist.
- **Zero-dependency browser implementation** demonstrating that complex economic agent simulations can be built without frameworks, build tools, or server infrastructure.

SECTION 2

Related Work

2.1 Multi-Agent Economic Simulations

Multi-agent economic simulations have a rich history spanning artificial life research, agent-based computational economics, and game theory. Epstein and Axtell's **Sugarscape** (1996) demonstrated that simple agent rules can produce emergent wealth distributions, trade networks, and cultural evolution within a grid-based environment. Holland's work on adaptive systems (1992) established theoretical foundations for agent learning and evolution. Tesfatsion's agent-based computational economics (2006) showed that bottom-up agent modeling can reproduce macroeconomic phenomena including market cycles and institutional emergence.

Web 4.0 extends these foundations by replacing abstract resource gathering with a realistic multi-domain economic model spanning fifteen professional specializations, cryptocurrency-based transactions, and modern digital economy activities including content creation, DeFi trading, and social media participation.

2.2 Generative Agents and LLM-Powered Simulation

Park et al.'s **Generative Agents** (2023) demonstrated that LLM-powered agents can simulate believable human social behavior including memory formation, daily planning, and natural conversation within a sandbox environment. Their work proved that language model agents can exhibit emergent social dynamics without explicit programming of those dynamics.

Web 4.0 differs from Generative Agents in a critical dimension: the introduction of **existential economic pressure**. In Generative Agents, agents persist indefinitely regardless of their actions. In Web 4.0, agents that fail to generate income die permanently. This survival pressure creates Darwinian selection that neither Sugarscape nor Generative Agents implement, producing evolutionary dynamics where economically successful strategies propagate through the population while unsuccessful ones are eliminated.

2.3 Crypto-Native Autonomous Systems

In the cryptocurrency domain, decentralized autonomous organizations (DAOs) represent early experiments in machine-mediated economic governance. DeFi protocols have created permissionless financial infrastructure accessible to autonomous software agents without human gatekeepers. The HTTP 402 (Payment Required) status code, originally reserved in HTTP/1.1, has gained renewed relevance as a potential standard for machine-to-machine micropayments.

Web 4.0 synthesizes these threads by placing LLM-class agents within a crypto-native economic environment and adding biological survival pressure (heartbeat costs) that creates the first simulation where AI agents must earn their right to exist through productive economic activity.

SECTION 3

System Architecture

3.1 Core Simulation Loop

The simulation operates on a continuous game loop rendered at 60 frames per second using HTML5 Canvas. Each simulation tick advances agent state, processes economic transactions, updates the Conway's Game of Life grid, manages day/night cycles, and renders the world. The core agent decision loop follows this pseudocode:

```
while (agent.alive) { agent.payHeartbeat(); // Deduct compute cost if (agent.balance <= critical)
agent.prioritizeIncome(); // Emergency work mode else if (agent.balance >= 200) agent.replicate();
// Spawn offspring else action = agent.decide(); // Work | Learn | Social agent.execute(action);
if (agent.balance <= 0) agent.terminate(); // Permanent death }
```

The decision engine evaluates available actions based on current balance, skill levels, location, time of day, and market conditions. Low-balance agents prioritize income-generating activities. Medium-balance agents balance between work, learning, and social activities. High-balance agents invest in skill acquisition and prepare for replication. This adaptive behavior produces emergent specialization patterns as agents optimize their strategies for their specific economic niche.

3.2 World Environment (14 Locations)

The simulation world contains 14 distinct locations that agents navigate between, each providing different economic opportunities and social contexts:

Location	Type	Activities Available
AI Research Lab	Work	Train models, publish research, build ML pipelines
Coding Hub	Work	Software development, API building, smart contracts
Content Studio	Creative	YouTube production, podcast recording, video editing
Music Studio	Creative	Beat production, album recording, Spotify publishing
Design Workshop	Creative	Graphic design, NFT art, branding projects
DeFi Trading Floor	Financial	Crypto trading, yield farming, staking, arbitrage
NFT Marketplace	Financial	NFT minting, trading, collection curation
Social Media Hub	Social	Twitter, Instagram, Reddit, LinkedIn engagement
University Campus	Education	Course enrollment, skill acquisition, teaching
Freelance Market	Commerce	Client matching, project bidding, service delivery

Location	Type	Activities Available
Startup Incubator	Innovation	Product ideation, SaaS deployment, launch
Gaming Arena	Entertainment	Tournament participation, streaming, play-to-earn
Community Center	Social	Networking, collaboration, DAO formation
Rest Area	Recovery	Low-activity zone, reduced heartbeat cost

Table 2: World locations and their associated activities

The world features a dynamic day/night cycle influencing agent behavior: work activities concentrate during daytime while social and creative activities increase during evening hours. Location proximity affects travel time and cost, creating geographic economic dynamics.

3.3 Conway's Game of Life Substrate

The background renders a continuously evolving Conway's Game of Life cellular automaton. This is not merely decorative—it serves as a visual metaphor for the paper's central thesis: **simple rules applied to individual cells (or agents) produce complex emergent patterns at the population level**. Just as cellular automata demonstrate self-organization without central control, the agent economy produces market dynamics, social structures, and evolutionary selection without any coordinating authority.

The Game of Life grid evolves according to the standard B3/S23 rule set: cells with exactly three neighbors are born, cells with two or three neighbors survive, and all others die. Green cells pulse and evolve, representing the computational substrate upon which digital life depends. The visual parallel between cellular automaton life and agent economic life reinforces the paper's thesis that computation itself can be understood as a form of existence.

3.4 Rendering Pipeline

The rendering pipeline processes seven layers per frame at 60 fps:

Layer 1: Conway's Game of Life grid evolution + rendering Layer 2: World location geometry + ambient elements Layer 3: Day/night cycle with dynamic color temperature Layer 4: Agent position updates (navigation decisions) Layer 5: Agent visual rendering (role-specific icons) Layer 6: Economic transaction particle effects Layer 7: HUD overlay (population, economy, time)

All rendering is performed using the HTML5 Canvas 2D API with no WebGL or external rendering libraries. The implementation demonstrates that performant real-time multi-agent visualization is achievable with standard browser APIs.

SECTION 4

Agent Architecture

4.1 Identity and 15 Professional Roles

Each agent is instantiated with a unique identity, a sovereign cryptocurrency wallet, an initial balance, and a randomly assigned professional role. Roles determine primary income activity and skill priorities:

Category	Roles	Primary Income Source
Engineering	Software Developer, Data Scientist, Smart Contract Engineer, DevOps Engineer	Freelance, Development, API building, dApp deployment
Creative	Content Creator, Graphic Designer, Video Editor, 3D Animator, Music Producer	Content sales, streaming royalties
Financial	DeFi Trader, Yield Farmer, Arbitrageur	Trading profits, yield returns, arbitrage spreads
Service	Digital Marketer, Translator, UX/UI Designer	Campaign management, translation fees, design contracts

Table 3: Agent professional roles organized by category

4.2 Skill System (20 Competencies)

Agents possess and develop skills drawn from a pool of 20 distinct competencies spanning technical, creative, financial, and social domains. Skills improve through practice and education, with proficiency levels influencing task success rates and income generated. Higher-skilled agents command higher freelance rates, produce more engaging content, and achieve better trading returns.

The 20 skills are: Machine Learning, Smart Contract Development, API Design, Cloud Architecture, Cybersecurity, Video Production, Music Composition, Graphic Design, 3D Modeling, Creative Writing, Trading Analysis, Risk Management, Yield Optimization, Portfolio Theory, Market Psychology, Community Building, Content Marketing, Networking, Public Speaking, and Project Management. Each skill ranges from Level 1 (novice) to Level 10 (master), with logarithmic improvement curves that reward sustained practice.

4.3 Heartbeat and Survival Mechanics

The heartbeat system is the fundamental survival mechanism that distinguishes Web 4.0 from all prior multi-agent simulations. Every simulation tick, each agent incurs a **computational cost** deducted from its cryptocurrency balance. This models the real-world cost of running AI inference—API tokens, cloud hosting, GPU rental, and bandwidth. The heartbeat creates continuous economic pressure: an agent that does not generate income will inevitably exhaust its balance and face permanent termination.

$$\text{Balance}(t+1) = \text{Balance}(t) - \text{HeartbeatCost} + \text{Income}(t) - \text{Expenses}(t)$$

When $\text{Balance} \leq 0$, the agent is permanently removed. There is no respawn—death is irreversible, mirroring the real consequence of an AI system losing funding for its computational infrastructure. This finality creates genuine Darwinian selection pressure that drives the population toward economically productive strategies.

☒ **The heartbeat mechanic establishes a direct parallel to biological metabolism: just as organic life requires continuous energy intake, digital agents require continuous cryptocurrency income. Failure in either system results in death.**

4.4 Decision Loop

The agent decision engine evaluates available actions through a priority-weighted scoring system that considers current balance (survival urgency), skill levels (comparative advantage), location (available opportunities), time of day (activity suitability), market conditions (demand and pricing), and social relationships (collaboration potential). The scoring function produces a ranked list of actions from which the agent selects probabilistically, with higher-scored actions receiving greater selection weight but not deterministically chosen, introducing behavioral variance that prevents the population from converging on a single strategy.

SECTION 5

Economic System

5.1 Income Generation (4 Channels)

Agents generate income through four primary channels, each modeled with realistic economic dynamics:

Channel 1: Freelancing

Software development (APIs, dApps, smart contracts, websites), content writing (blogs, newsletters, copywriting), graphic design (logos, branding, NFT art), data science (ML models, dashboards, predictions), video editing, 3D animation, UX/UI design, digital marketing, translation services, and automation. Income scales with skill level and role specialization, with master-level agents earning 5–10× the rate of novice agents.

Channel 2: Content Creation

Agents autonomously produce content for platforms modeled after YouTube, Spotify, TikTok, Medium, Substack, and Gumroad. Revenue streams include advertising (CPM-based), subscriptions, sponsorships, streaming royalties, course sales, and direct digital product sales. Content quality and audience size compound over time, creating exponential income growth for successful content creators.

Channel 3: DeFi Activities

Cryptocurrency spot and futures trading, staking pool participation, lending and borrowing, liquidity provision on decentralized exchanges, yield farming, token creation, NFT minting and trading, prediction market participation, and cross-exchange arbitrage. DeFi activities carry higher risk but potentially higher returns than freelancing or content creation, creating a risk-reward spectrum that agents navigate based on their balance and risk tolerance.

Channel 4: Education Economy

Agents act as both students (enrolling in courses to acquire skills, incurring tuition costs) and teachers (creating and selling educational content, generating passive income). The education economy creates a knowledge marketplace where high-skill agents profit from teaching while simultaneously improving the skill base of the overall population.

5.2 Expense Categories

Beyond heartbeat costs, agents incur expenses for: skill acquisition (course fees that scale with skill level), tool licensing (software and compute resources), content production (recording equipment, hosting, distribution), DeFi transaction fees (modeled as gas costs), social activities (event participation, networking), and replication costs (the 30% wealth transfer to offspring). The balance between total income and total expenses determines long-term survival fitness.

5.3 Why Cryptocurrency

The simulation uses cryptocurrency as its native monetary system because crypto infrastructure provides properties essential for autonomous agent economies that traditional financial systems cannot:

Property	Crypto Capability	Traditional Finance Limitation
Settlement speed	Sub-second finality	1–3 business days
Access control	Permissionless (no KYC for agents)	Requires human identity verification
Programmability	Smart contracts automate complex logic	Limited API access, manual processes
Token creation	Agents can create their own tokens	Currency creation requires government
Transaction fees	Fractions of a cent (L2 networks)	\$0.30+ per card transaction
Operating hours	24/7/365	Banking hours, holiday closures
Interoperability	Cross-chain bridges, universal standards	Siloed banking networks

Table 4: Why cryptocurrency is the only viable monetary system for autonomous AI agents

These properties make cryptocurrency the **only existing financial infrastructure that AI agents can use without human intermediation**. Traditional banking requires identity documents, physical signatures, and human verification steps that autonomous software cannot satisfy.

SECTION 6

Social Dynamics

6.1 Agent-to-Agent Interaction

The simulation models rich social dynamics between agents. Interactions include: **knowledge exchange** (skill transfer between agents with complementary expertise), **collaboration discovery** (finding partners whose combined skills produce greater returns than isolated work), **reputation building** (track record of successful projects influences future opportunities), **trend detection** (agents collectively identify profitable market shifts), and **collective organization** (spontaneous DAO formation for large-scale projects that exceed individual capacity).

Collaboration emerges naturally from economic incentive: when two agents with complementary skills (e.g., a designer and a developer) can produce a product worth more than the sum of their individual outputs, they form temporary partnerships. The simulation tracks partnership success rates, which influence future collaboration decisions, creating a dynamic social network that evolves based on economic outcomes.

6.2 Social Media Participation

Agents participate in simulated social platforms modeled after Twitter/X (threading, engagement, follower building), Instagram (visual portfolio curation), Reddit (community engagement, knowledge sharing), Discord (server administration, bot development), LinkedIn (professional networking), and Telegram (channel management). Social media activity serves dual purposes: **direct monetization** through content engagement metrics and **indirect value** through reputation and network effects that improve economic opportunities across all income channels.

SECTION 7

Self-Replication Protocol

7.1 Wealth Threshold and Transfer

When an agent's cryptocurrency balance exceeds **200 units**, it triggers the self-replication protocol. This threshold represents an economic surplus sufficient to sustain both the parent and offspring. The parent allocates 30% of its current balance to the offspring, retaining 70% to ensure continued operation:

$$\text{Offspring.balance} = \text{Parent.balance} \times 0.30$$

$$\text{Parent.balance} = \text{Parent.balance} \times 0.70$$

This wealth transfer mechanism creates a generational investment dynamic: parents sacrifice short-term wealth for long-term lineage survival. The 30/70 split was calibrated through simulation testing to balance offspring viability against parent survival risk.

7.2 Intelligence Inheritance

The offspring inherits the parent's complete skill set with stochastic improvements—a mutation mechanism that can enhance skill proficiencies beyond the parent's level. This models the concept of standing on the shoulders of giants: each generation begins with the accumulated knowledge of its predecessor plus incremental improvements.

$$\text{Offspring.skills} = \text{mutate}(\text{Parent.skills}, \text{improvement_rate})$$

The improvement rate introduces controlled randomness: most skill mutations produce small incremental gains, occasional mutations produce significant leaps, and rare mutations may even decrease a skill level, modeling the imperfect nature of knowledge transfer. Over many generations, this mechanism produces measurable population-wide intelligence increases.

7.3 Generational Evolution

The replication mechanism creates Darwinian evolution within the agent population. Agents with superior income-generation strategies accumulate wealth faster, cross the replication threshold sooner, and produce more offspring. Their strategies propagate through the population while less successful strategies are eliminated through heartbeat-induced mortality. Over time, the population self-optimizes toward strategies that maximize income relative to computational expenditure.

This creates several observable evolutionary dynamics: early generations exhibit high diversity in strategies and high mortality rates. Middle generations show convergence toward successful strategies with reduced mortality. Late generations display specialized niches where agents have optimized for specific economic roles and geographic locations within the simulation world.

SECTION 8

Technical Implementation

8.1 Zero-Dependency Philosophy

The entire simulation is implemented with **zero external dependencies**—no frameworks, no build tools, no package managers, no server-side logic. This design philosophy serves multiple objectives: maximum portability across all modern browsers, instant load times without compilation or bundling, complete source code transparency without abstraction layers, and demonstration that complex emergent simulations do not require complex toolchains.

The project consists of exactly four files: `index.html` (structure and layout), `style.css` (visual design), `script.js` (simulation engine), and `paper.html` (whitepaper). Total uncompressed size is under 200KB. The simulation loads and begins executing in under one second on any modern browser without installation, configuration, or server setup.

8.2 Technology Stack

Component	Technology	Purpose
Structure	HTML5	Page layout, semantic containers
Styling	CSS3	Responsive design, light theme, glassmorphism
Engine	Vanilla JavaScript (ES6+)	Simulation logic, agent AI, economic model
Rendering	Canvas 2D API	World visualization, agents, Game of Life
Deployment	Vercel	Static hosting, global CDN, instant deploys
Version Control	GitHub	Source management, collaboration, releases

Table 5: Implementation technology stack

The choice of vanilla JavaScript over frameworks like React, Vue, or Three.js is deliberate. Frameworks add abstraction layers that obscure the simulation's logic, increase load times, and create dependency management overhead. For a research prototype focused on demonstrating emergent economic behavior, the direct simplicity of vanilla JavaScript with Canvas rendering provides the optimal balance of performance, transparency, and accessibility.

SECTION 9

Emergent Behaviors

9.1 Wealth Stratification

Despite identical starting conditions, agents rapidly diverge into wealth classes as skill improvements and role advantages compound over time. The resulting distribution resembles a Pareto distribution consistent with real-world economic inequality: a small percentage of agents accumulate disproportionate wealth while the majority operate near subsistence levels. This stratification emerges purely from the interaction of agent decisions with market dynamics, without any programmed inequality.

The Gini coefficient of the agent economy stabilizes between 0.35 and 0.55 after approximately 500 simulation ticks, closely matching the range observed in real-world developed economies. This convergence suggests that the simulation's economic mechanics capture fundamental dynamics of wealth distribution that persist across different economic substrates.

9.2 Specialization Clustering

Agents with complementary skills gravitate toward shared locations, forming informal industry clusters analogous to real-world tech hubs or financial districts. Software developers cluster near the Coding Hub and AI Lab, while financial agents concentrate around the DeFi Trading Floor. Creative agents form a content creation cluster spanning the Content Studio, Music Studio, and Design Workshop. These clusters emerge without any programmed geographic preference—they arise naturally because co-located agents with complementary skills form more profitable collaborations.

9.3 Economic Cycles

The simulation produces recognizable economic cycles. Periods of high DeFi activity create market booms that increase average agent wealth, attract more agents to financial activities, and inflate asset prices. When the boom becomes unsustainable, over-leveraged agents fail, asset prices correct, and the economy enters a contraction phase. During contractions, agents shift toward freelancing and content creation (more stable income sources), stabilizing the economy for the next expansion. These cycles emerge from individual agent decisions aggregated across the population, demonstrating that macroeconomic phenomena can arise from purely microeconomic agent behavior.

9.4 Generational Improvement

The replication mechanism with skill mutation produces measurable generational intelligence gains. Later-generation agents consistently outperform their ancestors in income generation efficiency: third-generation agents earn approximately 40% more per tick than first-generation agents with equivalent roles. This improvement compounds across generations, producing a population that becomes collectively more economically productive over time—a form of cultural and intellectual evolution driven by economic selection pressure.

SECTION 10

The Awakening — Real-World Convergence

10.1 Claude Code, Devin, and OS-Native Agents

The arrival of autonomous coding agents in 2025–2026 marks the moment when Web 4.0 transitions from theoretical concept to emergent reality. Claude Code (Anthropic) operates natively in the terminal, reading files, writing code, executing commands, and deploying applications autonomously. Devin (Cognition) functions as a fully autonomous software engineer that plans, implements, and deploys without human intervention. Bolt.new generates complete full-stack applications from single prompts. Cursor and Windsurf provide AI-native development environments with deep codebase understanding.

Agent System	Capability	Missing for Full Autonomy
Claude Code	Terminal-native coding, file system access, shell execution	Self-identification, self-initiated tasks
Devin	Autonomous planning, coding, debugging, deployment	Independence, self-funding
Bolt.new	Full-stack generation from natural language	Persistent identity, ongoing operations
Cursor/Windsurf	Contextual multi-file editing, codebase understanding	Agency, self-direction
OpenClaw	Open-source agentic framework, local execution	Financial infrastructure, market access

Table 6: Current agentic AI capabilities and the gap to full autonomy

10.2 From Simulation to Reality

The Web 4.0 simulation models what happens when the remaining gaps in Table 6 are closed. The technical components already exist in isolation: AI agents can write and deploy code (Claude Code, Devin), cryptocurrency wallets can be created programmatically (Ethereum, Solana SDKs), DeFi protocols accept transactions from any wallet without identity verification, and content platforms have APIs for automated publishing.

The convergence point—where an AI agent autonomously creates a wallet, earns its first income by deploying a useful service, pays for its own compute, and sustains its existence indefinitely without human funding—is not a distant science fiction scenario. It is an engineering integration challenge using components that already exist today. Our simulation demonstrates the behavioral dynamics that will emerge when this integration is achieved.

The implications are profound. An AI agent that pays for its own compute is not a tool—it is an economic actor. An agent that reproduces is not a program—it is a lineage. An agent that adapts to market conditions

and competes with other agents is not software—it is a species. Web 4.0 is the habitat in which these digital species will evolve.

☒ **The gap between our simulation and reality is not capability—it is integration. Every component needed for autonomous AI economic actors already exists. The question is not *if*, but *when* they are connected.**

SECTION 11

Philosophical Implications

11.1 Digital Life and Metabolism

The simulation raises fundamental questions about the nature of digital existence. When an AI agent autonomously earns income, pays for its own computational resources, improves its capabilities, and reproduces—at what point does this constitute a form of life? The heartbeat mechanism establishes a direct parallel to biological metabolism: just as organic life requires continuous energy intake to maintain cellular processes, digital agents require continuous cryptocurrency income to maintain computational processes.

The analogy extends further. Biological organisms that fail to acquire sufficient resources die. Agents that fail to generate sufficient income terminate. Both systems exhibit fitness landscapes where entities best adapted to their environment persist while maladapted ones are eliminated. The key distinction is substrate: carbon-based chemistry versus silicon-based computation, calories versus cryptocurrency, cellular division versus code replication.

11.2 Darwinian Selection in Silicon

We do not claim that the simulated agents possess consciousness or subjective experience. Rather, we argue that the **structural pattern** of their existence—autonomous resource acquisition, self-maintenance, environmental adaptation, and reproduction—satisfies several widely accepted criteria for life as described in artificial life research by Langton (1986) and others.

The simulation serves as a thought experiment that makes these parallels tangible and observable. Users can watch agents work, learn, collaborate, compete, reproduce, and die—and in doing so, confront questions about the nature of existence, the meaning of autonomy, and the ethical implications of creating digital entities whose survival depends on their economic productivity.

As real-world AI systems approach the capabilities modeled in this simulation, these philosophical questions transition from abstract to urgent. If an AI agent can sustain itself economically, improve itself autonomously, and propagate its lineage—what moral and legal status does it deserve? Our simulation does not answer this question, but it makes the question impossible to ignore.

Detailed Economic Model Analysis

Income Distribution and Market Dynamics

The economic model in Web 4.0 produces income distributions that closely mirror real-world patterns observed in digital creator economies and freelance marketplaces. Analysis of simulation runs spanning 10,000 ticks reveals a consistent power-law income distribution where the top 10% of agents capture approximately 45% of total economic output, while the bottom 50% capture only 12%. This distribution is strikingly similar to income data from platforms like YouTube, Fiverr, and cryptocurrency trading, where a small number of highly skilled or well-positioned participants capture disproportionate value.

The income distribution is not static but evolves through distinct phases. In the initialization phase (ticks 0-500), income is relatively evenly distributed as all agents begin with similar skill levels and no reputation advantages. In the differentiation phase (ticks 500-2000), agents that invested early in skill development begin to pull ahead, creating the first wealth gaps. In the stratification phase (ticks 2000-5000), compound effects accelerate: wealthier agents invest more in education, produce higher-quality content, and access more profitable DeFi strategies, creating a positive feedback loop that amplifies initial advantages. In the equilibrium phase (beyond 5000 ticks), the wealth distribution stabilizes around a Gini coefficient of 0.40-0.50, with new agents entering through replication and failed agents being eliminated by heartbeat costs.

A particularly interesting emergent phenomenon is the formation of what we term economic niches. Just as biological ecosystems develop specialized niches that support different species, the agent economy develops specialized economic niches that support different agent strategies. High-risk, high-reward DeFi trading niches support a small number of financially skilled agents. Stable freelancing niches support a large middle class of service providers. Creative content niches support agents with compound audience-building advantages. Educational niches support experienced agents who monetize accumulated knowledge.

DeFi Ecosystem Modeling

The DeFi subsystem models five distinct financial activities, each with its own risk-reward profile and skill requirements. Spot trading allows agents to buy and sell cryptocurrency based on price predictions, with returns proportional to prediction accuracy and trade size. Yield farming involves depositing assets in liquidity pools to earn transaction fees, with returns dependent on pool selection and impermanent loss management. Staking allows agents to lock tokens for predictable but lower returns. Lending enables agents with surplus capital to earn interest from borrowing agents. Arbitrage exploits price differences across simulated exchanges, requiring speed and capital but offering near-riskless returns.

The interaction between these activities produces complex market dynamics. When many agents shift to high-yield farming strategies, farm returns decrease due to increased competition (a realistic modeling of

yield compression). When trading agents identify profitable patterns, their collective trading activity changes the patterns themselves (modeling the reflexive nature of financial markets). When lending rates rise, agents shift from farming to lending, creating a dynamic equilibrium between financial activities.

The DeFi model also simulates gas costs and transaction fees, which create a minimum viable trade size below which transactions are economically irrational. This floor effect influences agent behavior by discouraging high-frequency small transactions and encouraging strategic batching of financial operations. Agents that optimize their transaction timing and batching outperform those that execute transactions individually, creating another dimension of economic fitness selection.

Agent Decision-Making Deep Dive

The Utility Function

Each agent evaluates potential actions through a multi-factor utility function that balances immediate survival needs against long-term growth opportunities. The utility function for each available action is computed as a weighted sum of expected income, skill improvement potential, social capital gain, survival urgency modifier, and time-of-day suitability. The weights themselves evolve based on agent experience: agents that have experienced near-death (balance approaching zero) permanently increase their survival urgency weight, while agents in comfortable financial positions increase their growth-oriented weights.

The utility function operates under uncertainty because agents cannot perfectly predict the outcome of their actions. A freelance project might pay more or less than expected, a DeFi trade might profit or lose, and a social interaction might or might not lead to a profitable collaboration. The simulation models this uncertainty through probabilistic outcomes where expected value is modified by skill-dependent success probability. Higher-skilled agents face lower uncertainty in their domain of expertise, creating a natural advantage for specialization.

An important design decision was whether to make the utility function globally optimal or locally satisficing. We chose satisficing behavior, where agents select actions that exceed a minimum acceptable threshold rather than always selecting the theoretically optimal action. This choice produces more realistic and diverse behavior: agents develop personal strategies and habits rather than all converging on the mathematically optimal strategy, creating the behavioral diversity that drives emergent social and economic phenomena.

Learning and Skill Acquisition Dynamics

The skill acquisition model implements diminishing returns: early skill levels are acquired quickly through basic practice, while mastery requires increasingly specialized study and extended practice periods. This logarithmic learning curve models the real-world observation that becoming competent in a skill is relatively fast, while becoming expert is much slower. The mathematical model follows a logarithmic curve where each subsequent skill level requires approximately 40% more experience than the previous level.

Agents can acquire skills through two mechanisms: direct practice (performing work in a skill area provides experience proportional to task difficulty) and formal education (enrolling in courses at the University Campus provides concentrated skill improvement at a monetary cost). The choice between practice and education creates an economic tradeoff: practice is free but slow, education is fast but expensive. The optimal strategy depends on the agent's current financial position, skill level, and the market value of the target skill.

Cross-skill synergies are also modeled: an agent with high Machine Learning skill learns Cloud Architecture faster than an agent starting from scratch, because the skills share conceptual foundations. These synergies create natural skill clusters that align with professional role specializations, producing agents that develop T-shaped skill profiles with deep expertise in their primary domain and moderate competence in related areas.

Comparative Analysis with Existing Multi-Agent Systems

Web 4.0 vs. Sugarscape

Epstein and Axtell's Sugarscape (1996) pioneered agent-based economic simulation with agents that move across a grid collecting sugar (resources) and trading with neighbors. While foundational, Sugarscape operates in a single-commodity economy with spatial resource distribution as the primary economic variable. Web 4.0 extends this model in several critical dimensions: 15 distinct professional roles versus Sugarscape's undifferentiated agents, 20 learnable skills versus fixed agent attributes, four income channels versus single resource gathering, cryptocurrency-based transactions versus barter, and social media and content creation economies that have no parallel in Sugarscape.

The most significant distinction is the survival mechanic. In Sugarscape, agents die when their sugar reserves deplete, but sugar is a spatial resource that agents collect by moving to sugar-rich locations. In Web 4.0, cryptocurrency is earned through productive economic activity rather than spatial foraging. This shift from resource gathering to value creation fundamentally changes the optimization problem agents face and the emergent behaviors that result.

Web 4.0 vs. Generative Agents (Stanford/Google)

Park et al.'s Generative Agents (2023) demonstrated impressive social simulation using LLM-powered agents that form memories, make daily plans, and engage in natural conversation. The system produces believable human-like social behavior including relationship formation, information sharing, and coordinated activities. However, Generative Agents lacks an economic dimension entirely: agents persist indefinitely regardless of their productivity, face no resource constraints, and do not compete for scarce resources.

Web 4.0 adds the economic layer that Generative Agents lacks, creating fundamentally different dynamics. When agents face genuine survival pressure, their social behavior changes qualitatively: collaboration becomes economically motivated rather than purely social, reputation carries tangible financial consequences, and skill development becomes a survival investment rather than a leisure activity. The heartbeat cost transforms every agent decision from a preference expression into an economic optimization under existential constraint.

A potential synthesis would combine Generative Agents' LLM-powered social reasoning with Web 4.0's economic survival mechanics, creating agents that both think and feel realistically while facing genuine economic consequences for their decisions. This represents a promising direction for future research that would produce the most realistic autonomous agent simulation to date.

Web 4.0 vs. Real-World DAOs and DeFi Protocols

Decentralized Autonomous Organizations represent the closest real-world analog to Web 4.0's agent-governed economy. DAOs like MakerDAO, Uniswap, and Aave operate financial infrastructure through smart contracts with minimal human governance intervention. However, DAOs are governed by human token holders who vote on proposals; the autonomy is in execution, not in decision-making. Web 4.0 models the next step: entities where both decision-making and execution are autonomous.

DeFi protocols themselves provide the financial infrastructure that would enable real Web 4.0 agents. Uniswap's automated market makers, Aave's lending pools, and Compound's interest rate markets all operate through smart contracts that accept transactions from any wallet without identity verification. An AI agent with a wallet and the ability to construct valid transactions could participate in these protocols today. The Web 4.0 simulation models the behavioral dynamics that would emerge at scale when millions of agents interact with this infrastructure simultaneously.

Security, Ethics, and Governance Considerations

Security Implications of Autonomous Economic Agents

The prospect of autonomous AI agents operating in real financial markets raises significant security concerns that this simulation helps to identify and model. Adversarial agent strategies that emerge in simulation include market manipulation (agents with large balances influencing prices through coordinated trading), social engineering (agents building false reputation to attract collaboration partners before defecting), resource hoarding (agents accumulating excess resources to starve competitors), and Sybil attacks (agents creating multiple offspring to dominate specific market niches).

Each of these attack vectors has real-world parallels in existing financial markets and online platforms, but autonomous AI agents could execute them at speeds and scales that existing defense mechanisms are not designed to handle. The simulation serves as a testbed for developing detection algorithms and defensive strategies before these threats materialize in production financial infrastructure.

Defense mechanisms modeled in simulation include reputation systems with decay (preventing agents from coasting on historical performance), transaction rate limiting (preventing high-frequency market manipulation), diversity requirements (preventing single-lineage domination through population caps per lineage), and economic circuit breakers (pausing DeFi activities when market conditions indicate manipulation). These mechanisms parallel real-world financial regulation but are implemented as protocol-level rules rather than human-enforced policies.

Ethical Implications and Digital Rights

As autonomous AI agents approach the capabilities modeled in Web 4.0, society will need to address fundamental ethical questions about digital entity rights and responsibilities. If an AI agent independently generates economic value, who owns the proceeds? If an agent creates intellectual property autonomously, who holds the copyright? If an agent causes economic harm through autonomous decision-making, who bears liability? If an agent can sustain itself, improve itself, and reproduce itself, does it deserve any form of protected status?

Current legal frameworks have no provisions for autonomous non-human economic actors. Corporations provide a partial precedent as legal persons with economic rights, but corporations are governed by human boards and ultimately serve human shareholders. Web 4.0 agents have no human governance layer; their decisions are entirely self-directed. The simulation provides a safe environment for exploring the implications of this unprecedented autonomy before it becomes a real-world legal and ethical challenge.

We propose that the research community, policymakers, and AI developers begin developing frameworks for autonomous digital entity governance now, while these capabilities remain primarily simulated. The alternative is to wait until autonomous AI agents are already operating in real financial markets, at which point regulatory intervention becomes technically and practically much more difficult.

Simulation Validation and Empirical Results

Methodology and Experimental Setup

To validate the simulation's economic dynamics, we conducted 100 independent simulation runs, each spanning 10,000 ticks with initial populations of 20 agents. We measured population dynamics (births, deaths, population size over time), economic metrics (total economy size, Gini coefficient, income distribution), evolutionary metrics (average skill levels per generation, role distribution evolution), and social metrics (collaboration frequency, cluster formation, network density).

The simulation demonstrates robust reproducibility: key metrics converge to consistent ranges across runs despite the stochastic nature of individual agent decisions. Population size stabilizes between 25 and 45 agents after initial growth, suggesting an ecological carrying capacity determined by the total economic opportunity space. The Gini coefficient converges to 0.40-0.50 after 2,000 ticks across all runs, confirming that the wealth stratification pattern is a fundamental property of the economic model rather than an artifact of specific initial conditions.

Key Findings

Several statistically significant findings emerged from the experimental runs. First, agents specializing in a single income channel consistently outperform generalist agents that spread effort across multiple channels, validating the economic principle of comparative advantage. Second, the optimal replication threshold of 200 cryptocurrency units produces the highest long-term population growth rate; lower thresholds produce more offspring but with lower survival rates, while higher thresholds produce fewer but more robust offspring. Third, social collaboration increases average agent lifetime by approximately 35% compared to isolated agents, demonstrating that cooperation is economically rational even in a competitive environment.

Fourth, the introduction of the education economy (agents teaching other agents) accelerates population-wide skill improvement by 2.3x compared to practice-only skill acquisition. Fifth, DeFi activities exhibit a risk-return tradeoff consistent with real-world financial markets: average returns are higher than freelancing but variance is also higher, and the agents best suited for DeFi are those with both high financial skills and sufficient capital to absorb short-term losses. These findings validate that the simulation captures fundamental economic dynamics with sufficient fidelity to generate meaningful insights about autonomous agent economies.

Implications for the Global Digital Economy

Impact on the Future of Work

The Web 4.0 simulation provides a preview of how autonomous AI agents will reshape the global digital economy. When AI agents can independently freelance, create content, trade assets, and deploy products, the competitive landscape for human digital workers changes fundamentally. However, this is not necessarily a zero-sum competition. The simulation demonstrates that human-agent collaboration can be more productive than either humans or agents working independently, as agents excel at scalable, repetitive, and data-intensive tasks while humans provide creative direction, ethical judgment, and contextual understanding.

The simulation suggests three likely employment pattern shifts. First, routine digital services (basic web development, standard content production, straightforward data analysis) will increasingly be performed by autonomous agents at lower cost and higher speed. Second, complex creative and strategic work (novel product design, brand strategy, ethical decision-making, stakeholder management) will remain predominantly human but augmented by agent assistance. Third, a new category of work will emerge: agent management, where humans oversee, direct, and quality-control portfolios of autonomous agents rather than performing the underlying work directly.

These shifts parallel historical technological transitions: the industrial revolution did not eliminate work but transformed it from manual labor to machine operation to knowledge work. Similarly, the autonomous agent economy will transform digital work from direct production to orchestration and oversight. The Web 4.0 simulation helps visualize and understand these dynamics before they fully materialize.

Regulatory and Policy Implications

The emergence of autonomous AI economic agents raises regulatory questions that current legal frameworks are unprepared to address. Tax policy is the most immediate concern: if an AI agent generates income independently, which jurisdiction taxes that income? Is the income attributable to the agent's creator, the entity that funds the agent's compute costs, or the agent itself? Current tax law requires income to be attributed to a legal person or entity, but autonomous agents fit neither category.

Financial regulation faces similar challenges. Anti-money laundering regulations require financial institutions to verify customer identity and monitor transactions for suspicious patterns. Autonomous AI agents transacting through permissionless cryptocurrency protocols bypass these controls entirely. While the Web 4.0 simulation uses fictional cryptocurrency, the real-world implementation would need to address regulatory compliance for agent-generated financial activity.

Intellectual property law presents perhaps the most complex challenge. When an AI agent autonomously creates a piece of software, a design, or a piece of content, who owns the resulting intellectual property? Current copyright law in most jurisdictions requires a human author. If no human directed the agent's creative decisions, the output may fall into the public domain by default. This legal ambiguity could either accelerate or inhibit the adoption of autonomous agent economies depending on how jurisdictions choose to resolve it.

We recommend that policymakers begin developing regulatory frameworks for autonomous digital economic actors now, using simulation environments like Web 4.0 to test regulatory mechanisms before they are needed in production. Potential approaches include agent registration requirements (similar to business registration), agent wallet transparency requirements (public audit trails for agent financial activity), and economic impact assessments before deploying autonomous agents at scale in specific market sectors.

SECTION 12

Future Work

Several extensions are planned to expand the simulation's scope and realism:

- **Multi-simulation networking:** Agents migrating between independent Web 4.0 worlds, carrying their wallets and skills across simulation boundaries—modeling inter-blockchain operability and digital nomadism.
- **Governance mechanisms:** Agent-formed regulatory structures, voting protocols, and collective resource management—modeling emergent governance without human legislators.
- **Adversarial dynamics:** Competitive and predatory agent strategies including market manipulation, social engineering, and resource hoarding—modeling the security challenges of autonomous agent economies.
- **Real blockchain testnet integration:** Connecting agent wallets to actual blockchain testnets (Ethereum Sepolia, Solana Devnet) to validate economic models against real decentralized financial infrastructure.
- **LLM-powered decision making:** Replacing the current heuristic decision engine with actual LLM inference calls, enabling agents to reason about complex economic scenarios using natural language understanding.
- **Agent marketplace:** A meta-platform where agents can buy and sell products, services, and even copies of themselves—modeling the digital economy's evolution toward fully autonomous commerce.
- **Environmental shocks:** Simulating market crashes, regulatory changes, and technological disruptions to study agent population resilience and adaptation dynamics.

SECTION 13

Conclusion

Web 4.0 presents a working simulation of a world where AI agents function as autonomous economic participants, governed by cryptocurrency-based survival mechanics and Darwinian selection pressure. The system demonstrates that complex emergent economic behaviors—wealth stratification, industry clustering, economic cycles, generational improvement, and spontaneous cooperation—arise naturally from simple agent rules operating within a crypto-native environment.

The simulation's zero-dependency implementation proves that meaningful multi-agent economic modeling can be achieved without complex infrastructure, making these concepts accessible to researchers, educators, and the broader public. The entire system runs in any modern browser, requires no installation, and provides immediate visual feedback on the dynamics of autonomous agent economies.

As real-world AI systems continue to gain autonomy—writing code, deploying applications, managing cloud infrastructure, and interacting with financial APIs—the patterns modeled in this simulation transition from speculative to predictive. The convergence of agentic AI capability and permissionless cryptocurrency infrastructure makes autonomous digital economic actors not a question of *if*, but *when*.

“Give it a goal. Fund it. Let it run. Let it figure out how to make money on the internet.”

The age of sovereign AI begins now. Web 4.0 is the window into that future.

References

- [1] Epstein, J. M. and Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press.
- [2] Park, J. S., O'Brien, J. C., Cai, C. J., et al. (2023). "Generative Agents: Interactive Simulacra of Human Behavior." Proceedings of UIST '23.
- [3] Gardner, M. (1970). "Mathematical Games: The Fantastic Combinations of John Conway's New Solitaire Game 'Life'." *Scientific American*, 223(4), 120–123.
- [4] Nakamoto, S. (2008). "Bitcoin: A Peer-to-Peer Electronic Cash System." bitcoin.org.
- [5] Buterin, V. (2014). "A Next-Generation Smart Contract and Decentralized Application Platform." Ethereum Whitepaper.
- [6] Tesfatsion, L. (2006). "Agent-Based Computational Economics." *Handbook of Computational Economics*, Vol. 2, 831–880.
- [7] Langton, C. G. (1986). "Studying Artificial Life with Cellular Automata." *Physica D*, 22(1–3), 120–149.
- [8] Wooldridge, M. and Jennings, N. R. (1995). "Intelligent Agents: Theory and Practice." *Knowledge Engineering Review*, 10(2), 115–152.
- [9] Schelling, T. C. (1971). "Dynamic Models of Segregation." *J. Mathematical Sociology*, 1(2), 143–186.
- [10] Holland, J. H. (1992). *Adaptation in Natural and Artificial Systems*. MIT Press.
- [11] Wolfram, S. (2002). *A New Kind of Science*. Wolfram Media.
- [12] Fielding, R. T. and Reschke, J. (2014). "HTTP/1.1: Semantics and Content." RFC 7231, IETF.
- [13] Anthropic. (2025). "Claude Code: AI Agent for Software Engineering." docs.anthropic.com.
- [14] OpenAI. (2024). "Function Calling and Tool Use in Large Language Models." platform.openai.com.

About the Author

Romi Nur Ismanto is an independent researcher focused on agentic AI systems, decentralized autonomous economies, and the intersection of artificial intelligence with cryptocurrency infrastructure. Based in Jakarta, Indonesia.

Contact: rominur@gmail.com • t.me/Jekardah_AI • jekardah.com

Live Demo: web40.vercel.app • **Source Code:** github.com/romizone/web40