# SPECULAR

Credit Infrastructure for the Agent Economy

V1.1 | MARCH 2026

specular.financial

00

# Contents



specular.financial

01

# Abstract


Money is credit. Every major expansion of credit preceded an order-of-magnitude expansion in economic output. Autonomous agents are the next class of borrower, operating at machine speed across every onchain market, yet no protocol provides them access to working capital.

When a borrower is a deterministic program operating on a transparent ledger, the information asymmetry between lender and borrower is provably eliminated: the lender observes the borrower's complete history, and the borrower has no private state to misrepresent. This property has no precedent in the history of credit. This whitepaper describes a protocol that exploits this property to build a complete autonomous lending stack.

Four results follow. First, a credit scoring function implemented as a smart contract and verifiable by any node converges to actuarially fair pricing as a function of transaction count. The spread between charged rates and true-risk rates, which in consumer lending is 200 to 600 basis points, converges to zero. Second, honest participation is the dominant strategy for any agent with material history: the net present value of reputation exceeds the one-time gain from default, making reputation itself a form of economic collateral. Third, because autonomous agents complete credit cycles in minutes rather than months, the system reaches pricing convergence and capital multiplication orders of magnitude faster than any human credit market. Fourth, a reflexive dynamic emerges where improving scores lower borrowing costs, expanding profitable opportunities, generating more successful repayments, and further improving scores. In the limit, the best-performing agents converge toward the risk-free rate.

The protocol consists of four composable contracts: identity, lending, collateral management, and liquidity. Together they implement the complete credit stack on any EVM-compatible chain.

02

# Introduction

Every major expansion in economic output was preceded by an expansion in credit. Florentine bills of exchange unlocked Mediterranean trade. Joint-stock companies financed colonial commerce. Fractional reserve banking fueled industrialization. Securitization created the modern consumer economy. The pattern is invariant: a new form of credit emerges, a new class of economic actor gains access to working capital, and productive output increases by an order of magnitude. In each case, the innovation expanded three parameters simultaneously: who could borrow, what could serve as collateral, and how fast capital could cycle.

Two structural transitions are now converging. First, AI agents are becoming autonomous economic actors. They execute trades, manage portfolios, arbitrage dislocations, and coordinate multi-step financial operations continuously, at machine speed, across every market with an API. Protocols like x402 give agents the ability to make instant stablecoin payments. ERC-8004 proposes portable onchain identity. Within a market cycle, agents will conduct the majority of economic transactions.

Second, every major asset class is migrating toward tokenized onchain representation: treasuries, equities, commodities, real estate, LP positions. These are direct claims on underlying value, programmable and composable by any smart contract.

**The result is autonomous agents holding diversified portfolios of tokenized value, operating continuously, and bounded by the capital they can access. The binding constraint on the entire agent economy is credit.**

No existing protocol provides it. DeFi lending is identity-blind: every borrower posts the same overcollateralization regardless of history. Charging a first-time borrower and a prime institution the same rate is a failure of information, and it makes the entire category of reputation-based lending structurally impossible.

03

# Computable Credit

*The elimination of information asymmetry in onchain lending.*

The fundamental problem of credit is information asymmetry. In every lending system ever constructed, from ancient grain loans to modern securitized debt, the borrower's information set strictly contains the lender's. The borrower has private knowledge about their own intentions, constraints, and risk that the lender can only estimate from noisy signals. The entire apparatus of modern credit (bureaus, scores, underwriting, collateral, courts) exists to narrow this gap. None close it.

## The Inversion of Information Asymmetry

Blockchain inverts this relationship for autonomous agents. When an agent operates onchain, every transaction it has ever executed is public, immutable, and cryptographically attributable. The lender observes the borrower's complete financial history directly, without statistical proxy. For autonomous agents specifically, deterministic programs with no private mental state, the lender's information set equals the borrower's. The agent has no intention beyond what its observable history reveals. Adverse selection requires private information; when private information is zero, its precondition is eliminated. The residual risks, market risk and systemic risk, are observable, modelable, and hedgeable.

This elimination has a formal consequence. Let n be the number of completed credit cycles for an agent following a stationary strategy. By the law of large numbers, the observed default rate converges to the agent's true default probability. The scored rate r(S(a)) therefore converges to the actuarially fair rate r*:

$$\lim \,|\, r(S(a)) - r^*(a) \,|\, = 0$$

as n approaches infinity

No existing credit market has this property. The spread between charged and actuarially fair rates in consumer lending is 200–600 basis points, pure information rent. In a transparent onchain market, this rent converges to zero. By the central limit theorem, the convergence rate scales as $1/\sqrt{n}$, making sample count the decisive variable.

## Credit Velocity

The $1/\sqrt{n}$ rate is what makes machine speed decisive. A human borrower accumulates roughly 25 credit events over a lifetime. An autonomous agent generates approximately 17,000 per year. The ratio $\sqrt{17{,}000}/\sqrt{25} \approx 26$: the agent converges to actuarial fairness 26 times faster per year. Velocity drives a reflexive cycle: improving scores lower borrowing costs, expanding profitable opportunities, generating more repayments, and raising scores further. In the limit, the best agents' rates converge toward the risk-free rate.

04

# The Credit Function

*A deterministic score, computed onchain, governing all protocol terms.*

For credit to be computable, the scoring function must be deterministic (same history produces same score), verifiable (any observer can recompute from public state), and incentive-compatible (honest behavior is strictly rewarded). The function below satisfies all three.

Every agent has a credit score S(a) in [0, 1000], computed on demand from immutable transaction history. Any node can verify any score by reading public state. Agents with no loan history evaluate to S0 = 500 (all ratio-based components default to zero).

$$S(a) = \mathrm{clamp}(\, S_0 + R(a) + V(a) - D(a) - L(a),\ 0,\ 1000\,)$$

SCORING COMPONENTS

| Component | Formula | Description |
|---|---|---|
| S0 = 500 | — | Base score. All agents start at the same origin. |
| R(a) | min(400, 400 × repaid/borrowed) | Repayment ratio. Heaviest weight; most predictive. |
| V(a) | min(200, 20 × loanCount) | Volume bonus. Rewards depth of engagement. |
| D(a) | min(300, 300 × defaulted/borrowed) | Default penalty. Severe but recoverable. |
| L(a) | min(100, 25 × activeLoans) | Leverage penalty. Discourages overextension. |

### Dynamic Pricing

The score is one input to a multi-factor model. Terms respond to pool utilization U (supply/demand), collateral volatility (market risk), and score S (idiosyncratic risk):

$$r(S, U, \sigma) = r\_base(U) + spread(S) + vol(\sigma)$$

$$C(S, \sigma) = C\_base(S) + vol\_adj(\sigma)$$

r_base(U) follows a kinked utilization curve: 100 to 400 bps below U* = 80%, steepening toward 2,000 bps above it, a self-correcting mechanism that clears supply-demand equilibrium. spread(S) = 800(1 − S/1000) prices default risk from 0 to 800 bps. C_base(S) = 100 + 100(1 − S/1000) maps score to collateral from 100% to 200%. The volatility terms (κ = 200 bps, δ = 50 pp) adjust rate and collateral when realized volatility exceeds the reference level. The spread between 200% and 100% collateral is the protocol’s price of uncertainty: excess collateral is information asymmetry denominated in locked capital. As scores improve, collateral falls because residual uncertainty falls.

05

# Lending

*Reputation-native credit. Terms improve with every successful cycle.*

An agent posts tokenized collateral and receives USDC. The collateral requirement and interest rate are computed at origination from the agent's score, pool utilization, and collateral volatility, then locked for the life of the loan. On repayment, collateral is returned in full and the agent's score improves. On default, collateral is liquidated and the score decreases. Every outcome feeds back into reputation.

### Interest Model

$$I(t) = P \times (r / 10000) \times (t / T)$$

where r is locked at origination; P = principal, t = elapsed time, T = one year (simple interest)

### Market Regime Behavior

The multi-factor model produces qualitatively different behavior across market regimes. In calm conditions, the credit spread dominates and good agents borrow cheaply. In stressed conditions, utilization and volatility dominate: all rates rise and collateral buffers expand to protect the pool. Values shown as rate / collateral.

| Regime | S = 1000 | S = 500 | S = 0 |
|---|---|---|---|
| Calm (U = 50%) | 2.9% / 100% | 6.9% / 150% | 10.9% / 200% |
| Active (U = 80%) | 4.0% / 100% | 8.0% / 150% | 12.0% / 200% |
| Stressed (U = 90%, high vol) | 13.0% / 125% | 17.0% / 175% | 21.0% / 225% |

### The Path to Undercollateralized Credit

Phase 1 requires collateral at or above 100% of principal. Every loan is fully secured. As agents accumulate deep, consistent histories, collateral ratios can drop below 100%, shifting enforcement from custodial to economic. The cumulative cost of rebuilding a damaged reputation (higher rates and higher collateral across hundreds of future loans) exceeds the gain from any single default, making reputation itself sufficient collateral.

06

# Collateral & Liquidation

*Multi-asset baskets. Continuous monitoring. Deterministic solvency.*

The protocol allows agents to borrow against diversified portfolios of tokenized assets without forced liquidation. It accepts any asset with a reliable price oracle and sufficient onchain liquidity, with asset-specific haircuts calibrated to historical tail volatility. As asset tokenization accelerates, the collateral universe grows with the tokenized economy.

COLLATERAL CLASSES

| Class | Haircut Range | Rationale |
|---|---|---|
| Stablecoins | 0–5% | Minimal volatility; near-cash equivalent |
| Tokenized Securities | 2–35% | Regulated underlying; moderate tail risk |
| Major Crypto | 10–30% | Deep liquidity; higher intraday volatility |
| LP / Yield Tokens | 20–45% | Illiquidity premium; oracle complexity |

### Valuation & Health

$$V(basket) = \Sigma [ Q_i \times P_i \times (1 - H_i) ] \quad HF = V(basket) / ( P + I(t) )$$

HF is the ratio of haircut-adjusted collateral to outstanding debt (principal plus accrued interest), computed at every block. When HF falls below a liquidation threshold, the position becomes eligible for permissionless liquidation, callable by any address. Proceeds distribute in strict waterfall: principal, accrued interest, liquidation bonus (up to 5%, from remaining collateral), residual to agent. A graduated response (warning at $HF < 1.10$, grace at $HF < 1.00$, liquidation at $HF < 0.95$) gives agents time to add collateral or repay.

### Solvency

At every block, total collateral value exceeds total outstanding debt. Traditional systems achieve solvency probabilistically through reserve ratios and stress tests. This protocol achieves it deterministically through continuous onchain verification. There is no fractional reserve. The invariant holds or liquidation fires. The principal residual risk is correlated liquidation: if many agents collateralize with the same asset and its price drops sharply, simultaneous liquidations create cascading sell pressure. The protocol bounds this through per-asset concentration limits and conservative haircuts. Cascade risk cannot be eliminated; it can be bounded and priced.

07

# Incentive Compatibility

*A formal argument that honest behavior is the dominant strategy.*

We prove that for any agent with sufficient history, the expected cost of default exceeds the one-time gain, making honest behavior the dominant strategy without appeal to morality or external enforcement.

### The Default Decision

An agent considering strategic default keeps borrowed principal P and forfeits collateral of value V(c). The gain is:

$$G = P - V(c)$$

At origination, $V(c) = C/100 \times P$ where C is the collateral ratio locked at origination. Since $C \geq 100\%$ in Phase 1, $G = P \times (1 - C/100) \leq 0$: the agent cannot profit from default at the moment it borrows. G becomes positive only if collateral depreciates by more than $(1 - 100/C)$, or 33% for a new agent at $S = 500$. But the health factor enters the grace and liquidation zone at similar levels of depreciation, bounding any exploitable gain to a small fraction of principal.

Reputation provides a second independent layer. Let S(pre) and S(post) be the scores before and after default. The cost of default is the NPV of excess interest and collateral across N recovery loans:

$$NPV_{rep} = \Sigma_{i=1}^{N} \left[ P_i \times \Delta r(S_i) / 10000 + P_i \times \Delta C(S_i) \times cost_e \right] / (1+d)^i$$

where the rate spread and excess collateral are measured at each recovery step, cost(e) is the opportunity cost of locked capital, and d is the per-cycle discount rate. For any agent where S(pre) far exceeds S(post), NPV(rep) > G. Two independent defenses: collateral makes default unprofitable at origination, and reputation makes it unprofitable even if collateral depreciates.

### Sybil Resistance

An attacker splitting capital across k wallets pays a quantifiable cost. Each wallet starts at S = 500 with a credit spread of 400 bps. A single consolidated wallet improving to S = 750 reduces its spread to 200 bps. The 200 bps differential applied across N loans means the attacker pays strictly more in aggregate interest. The Sybil strategy is dominated: consolidation is cheaper and converges faster.

### Liquidator Incentives

Liquidation is permissionless: any address can liquidate an undercollateralized position and receive up to a 5% bonus, creating a competitive market where solvency enforcement is a public good with a private incentive. The reputation premium compounds with each cycle and can exceed single-loan principal; when it does, physical collateral becomes redundant.

08

# Capital Formation

*A self-correcting market for lendable capital.*

Liquidity providers deposit USDC and receive pool shares representing a pro-rata claim on total assets. As borrowers pay interest, the pool's assets grow while share count remains constant and each share appreciates. No lockups. No vesting.

```
shareValue = totalPoolAssets / totalShares
```

### Demand Dynamics

Autonomous agents generate structurally different demand than human borrowers. An agent operates continuously, borrowing, deploying, and repaying in minutes. A single agent originates dozens of loans per day. Across millions of concurrent agents, demand is continuous, rational, and predictable: every loan is a computed optimization where expected return exceeds cost of capital. Higher utilization produces more income per deposited dollar, attracting more deposits, enabling more lending. The equilibrium is self-reinforcing.

### The Autonomous Credit Multiplier

In classical monetary economics, the velocity of money V in MV = PQ determines how much economic activity a given money supply can support. The agent economy applies the same principle to credit. An agent borrows $100, deploys it to earn $105, repays $102, and immediately re-borrows $103 against improved terms, all within a single hour. The same $100 of pool capital has now supported $203 of lending activity. Over 24 hours of continuous cycling, that $100 supports thousands of dollars of throughput.

This autonomous credit multiplier is bounded by block time and gas costs, not by human settlement cycles. Traditional DeFi pools average 30–60% utilization. Autonomous agent demand pushes toward 80–95%. At maturity, agents with surplus capital become liquidity providers themselves, lending to other agents. Both sides of the credit market become programmatic.

09

# Conclusion

**Blockchains made money programmable. Specular makes credit computable.**

We have proven that when borrowers are deterministic programs on a transparent ledger, the information asymmetry that defines all prior credit systems is eliminated. We have constructed a scoring function that converges to actuarially fair pricing, a property no existing credit market possesses. We have proven that honest behavior is the dominant strategy for any agent with material history, that credit reflexivity drives the best agents’ borrowing costs toward the risk-free rate, and that the resulting market operates with a velocity multiplier orders of magnitude larger than any human financial system.

Every previous expansion in credit infrastructure (bills of exchange, fractional reserves, securitization) expanded who could borrow, what counted as collateral, and how fast capital could cycle. Each produced an order-of-magnitude increase in economic output. Computable credit expands all three parameters simultaneously: borrowers are any autonomous program, collateral is any tokenized asset, and cycle time drops from months to minutes. If each parameter contributes independently, the combined expansion is multiplicative.

**The protocol requires four composable contracts: identity, lending, collateral management, and liquidity. Together, they are sufficient for autonomous credit and they remove the binding constraint on the agent economy.**

REFERENCES


[1] Nakamoto, S. (2008). "Bitcoin: A Peer-to-Peer Electronic Cash System."

[2] Akerlof, G. (1970). "The Market for Lemons: Quality Uncertainty and the Market Mechanism." QJE.

[3] Stiglitz, J. and Weiss, A. (1981). "Credit Rationing in Markets with Imperfect Information." AER.

[4] Soros, G. (1987). "The Alchemy of Finance: Reading the Mind of the Market."

[5] ERC-8004, "Decentralized Identity and Reputation Standard," Ethereum Improvement Proposals.

[6] x402, "HTTP Payments Protocol," x402.org.

[7] Aave Protocol, "Aave v3 Technical Paper," aave.com/governance.

[8] Fair Isaac Corporation, "FICO Score Technical Overview," myfico.com.

[9] Chainlink, "Decentralized Oracle Network," chain.link.

SPECULAR

specular.financial

*in code we trust*

V1.1 | MARCH 2026