Computational Monetary Theory

Marco Patrone

June 2026

Contents

Computational Monetary Theory

Credits, Settlement, and Unit-of-Account Formation in AI-Mediated Markets

HomeSelf Research Publication Series

June 2026

DOI: 10.5281/zenodo.20784780

Abstract

This paper examines how monetary mechanisms may emerge in AI-mediated markets where computational admissibility determines allocative access. We argue that when representation quality

determines consideration set inclusion, a new scarcity class emerges: computational scarcity. Unlike traditional monetary scarcity (limited supply of money) or physical scarcity (limited supply of goods), computational scarcity arises from bounded inference—the constraint that AI systems can evaluate only a subset of available options.

We introduce **computational credits** as a theoretical construct: units representing claims on evaluation, verification, priority, or inclusion capacity in AI-mediated allocation. Unlike traditional money, computational credits are not store-of-value instruments but allocative access tokens that may be issued, consumed, and potentially exhausted through representation investment.

The paper explores how computational credits may emerge as unit-of-account primitives in AI-mediated markets, how settlement may occur through representation verification, and what governance mechanisms may prevent "computational inflation"—the devaluation of admissibility through representation saturation. We maintain a distinction between computational credits (allocative access) and traditional monetary instruments (store of value), while acknowledging potential convergence through credit market mechanisms. The paper is theoretical; no empirical claims about computational credit adoption are advanced.

Keywords: Computational credits, AI-mediated markets, allocative access, computational scarcity, settlement infrastructure, unit-of-account formation, representation verification

JEL Classification: E4, E5, L1, O3

Caveats and Scope Limitations

Before proceeding, we must establish explicit scope limitations:

1. **This is theoretical work.** No empirical evidence is presented. Computational credits are introduced as a theoretical primitive, not as an observed market phenomenon. All claims about monetary evolution are speculative and should be treated as hypotheses requiring validation.

2. **This is not about cryptocurrency or token economics.** We do not discuss blockchain-based currencies, decentralized finance, or token speculation. Computational credits are not crypto tokens—they are theoretical constructs representing allocative access in AI-mediated systems.

3. **This is not about platform credits or rewards points.** We do not discuss airline miles, credit card rewards, or loyalty program points. While these share some properties with computational credits, they are proprietary instruments issued by centralized platforms. Computational credits may emerge from structural constraints in AI-mediated allocation, not from loyalty program design.

4. **This is not about "AI coins" or speculative assets.** We do not discuss creating tradable assets from AI capabilities. Computational credits, if they emerge, would be claims on allocative infrastructure, not speculative investments.

5. **This is not predictive.** We do not claim that computational credits will definitively emerge or become widely adopted. We argue that if AI-mediated allocation becomes infrastructure-dependent, credit-like primitives may become necessary for managing allocative access.

6. **This is not normative.** We do not argue that computational credits are desirable or undesirable. We examine potential structural consequences of AI-mediated allocation, not policy prescriptions.

Introduction

The architecture of monetary systems has undergone several structural transformations in modern economic history. First, the emergence of commodity money (gold, silver) established physical scarcity as the basis of value. Second, the development of fiat money decoupled monetary value from physical constraints, enabling central banks to manage money supply in service of economic policy. Third, the rise of digital currencies and payment systems reduced transaction costs and enabled near-instant settlement.

We may be entering a fourth structural transformation: the emergence of **computational credits** as allocative primitives in AI-mediated markets.

Traditional monetary systems share a common assumption: money is a store of value that can be exchanged for goods and services. The unit of account (dollars, euros, yen) is independent of the goods being exchanged. A dollar is a dollar whether purchasing coffee, cars, or real estate.

But what happens when the ability to participate in markets at all becomes the scarce resource? When AI systems construct consideration sets under bounded inference, the binding constraint may shift from "can I afford this?" to "can I be considered for this?"

This paper asks a foundational question: **If AI-mediated markets allocate consideration under computational constraints, what becomes the unit of account?**

We argue that when representation quality determines consideration set inclusion, a new monetary-like primitive may emerge: computational credits. These credits would not be store-of-value instruments in the traditional sense, but claims on allocative infrastructure—the right to be evaluated, verified, and considered by AI systems.

The possibility motivates the theoretical framework we develop.

Part I — Monetary Primitives

What is Money?

Money is typically defined by three functions:

1. Medium of Exchange

Money facilitates transactions. Buyers and sellers can exchange money for goods and services without requiring barter or direct coincidence of wants.

2. Unit of Account

Money provides a common measure of value. Prices, debts, and contracts can be denominated in monetary units, enabling comparison and accounting.

3. Store of Value

Money retains purchasing power over time. Holders can defer consumption by storing value in money rather than perishable goods.

These three functions are interrelated but conceptually distinct. Something can serve as a unit of account without being a good store of value (hyperinflating currency). Something can serve as a medium of exchange without being a good unit of account (barter with multiple exchange rates).

Settlement and Finality

A critical property of money is **settlement finality**. When a transaction is settled, the obligation is discharged. The buyer has paid; the seller has been paid. No further claims can be made on the same transaction.

In modern monetary systems, settlement occurs through:

- **Cash transactions:** Final upon physical transfer
- **Bank transfers:** Final when the banking system records the transfer (typically same-day or next-day)
- **Digital payments:** Final when the payment processor confirms the transaction

Settlement requires trusted infrastructure that records and verifies transfers. This infrastructure may be central bank ledgers, commercial bank ledgers, or payment processors.

Why AI-Mediated Markets May Introduce New Primitives

AI-mediated markets differ from traditional markets in a critical respect: **participation is not guaranteed**. In traditional markets, anyone with money can participate. In AI-mediated markets, only options that are computationally admissible can participate.

This structural difference creates:

1. Allocative Scarcity

The right to participate may be scarce, independent of monetary resources. An option with abundant funding may still be excluded from consideration if it is not computationally admissible.

2. Verification Dependency

Participation may depend on verification of representation quality. The allocative system must verify that an option meets admissibility criteria before including it.

3. Infrastructure Dependency

Participation may depend on access to allocative infrastructure. Without infrastructure for representation, verification, and discovery, options cannot participate.

These structural differences suggest that AI-mediated markets may require monetary-like primitives that are qualitatively different from traditional money.

Part II — Computational Scarcity

The K < n Constraint as Scarcity Condition

Recall the **K < n constraint** from prior work:

- Let n be the number of available options in a market
- Let K be the number of options that can be evaluated and included in consideration sets
- When $K < n$, inclusion is scarce

This constraint creates **computational scarcity**: a binding limitation on allocative capacity that is independent of traditional monetary constraints.

Computational scarcity differs from traditional scarcity in several ways:

Dimension	Traditional Scarcity	Computational Scarcity
Source	Limited physical resources	Limited inference capacity
Constraint	Supply of goods/services	Budget for evaluation
Binding on	Sellers (production capacity)	All participants (allocative access)
Mitigation	Increase production	Increase compute, improve representation
Measurement	Units produced	Options evaluated

Inference Budget

The inference budget is the computational resource constraint that determines K:

- **Compute time:** Time available for search and evaluation
- **Context window:** Information capacity per inference
- **Model capacity:** Limits on parallel evaluation
- **Quality thresholds:** Minimum confidence for inclusion

The inference budget is fundamentally limited. Even with unlimited compute, the context window imposes a hard constraint on how many options can be considered simultaneously.

Representation Cost

Representation cost is the computational resources required to evaluate a specific option:

- **Data access cost:** Time and resources to retrieve representation
- **Parsing cost:** Resources to interpret and validate representation
- **Evaluation cost:** Resources to assess relevance and quality
- **Comparison cost:** Resources to compare against alternatives

Options with high representation cost consume more of the inference budget, reducing the number of options that can be evaluated overall.

Verification Cost

Verification cost is the computational resources required to validate representation integrity:

- **Provenance verification:** Confirming data source and authenticity
- **Consistency checking:** Validating internal coherence of representation
- **Freshness verification:** Confirming representation is current
- **Cross-referencing:** Validating against external sources

Verification costs may be significant, particularly for high-value or high-risk transactions. The allocative system may invest substantial resources in verification before including an option.

Attention vs. Inference Scarcity

Traditional digital economics focused on **attention scarcity**—the constraint that human attention is limited. SEO, content marketing, and advertising all compete for human attention.

Computational scarcity is distinct:

- **Attention scarcity:** Humans can only attend to a limited number of options
- **Computational scarcity:** AI systems can only evaluate a limited number of options

Attention scarcity operates on the demand side (buyer attention). Computational scarcity operates on the allocative side (system capacity). Both constraints may be present simultaneously, but they create different dynamics.

ASCII Diagram: Scarcity Types

```
TRADITIONAL MARKETS:

  Supply of Goods/Services           ← Physical scarcity
  (production capacity constraints)

          ↓

  Money (store of value)             ← Monetary scarcity
  (limited money supply)

          ↓

  Market Participation                ← Open to all with money
  (no allocative constraint)


AI-MEDIATED MARKETS:

  Supply of Goods/Services           ← Physical scarcity (unchanged)

          ↓

  Computational Admissibility         ← NEW: Computational scarcity
  (K < n constraint)

          ↓

  Market Participation                ← Limited to admissible options
  (allocative constraint)
```

Part III — Computational Credits

Definition

Computational Credits are theoretical constructs representing claims on allocative infrastructure in AI-mediated markets. They may take several forms:

- **Evaluation Credits:** Right to have an option evaluated by an allocative system
- **Verification Credits:** Right to have representation verified and authenticated
- **Priority Credits:** Right to expedited evaluation or preferential consideration
- **Action Credits:** Right to initiate agent-mediated transactions (booking, inquiry, etc.)

Computational credits are **not** money in the traditional sense. They do not represent a store of value. They cannot be used to purchase goods and services. They are not a medium of exchange in the traditional economy.

Instead, computational credits represent **allocative access rights**—claims on the computational infrastructure that determines participation in AI-mediated markets.

What Computational Credits Represent

Computational credits may represent:

1. Capacity Claims

The right to consume allocative infrastructure capacity. When inference budgets are limited, credits may determine which options are evaluated.

2. Verification Claims

The right to have representation verified by trusted infrastructure. Verification may be costly, and credits may govern access.

3. Priority Claims

The right to preferential treatment when allocative resources are constrained. High-priority options may be evaluated before low-priority options.

4. Temporal Claims

The right to be evaluated at a specific time. Freshly-updated representations may require priority credits to ensure timely evaluation.

What Computational Credits Do NOT Represent

Computational credits are **not**:

1. Store of Value

Credits cannot be stored and retrieved later for equivalent allocative access. The value of credits may depend on current infrastructure conditions, not accumulated balance.

2. Medium of Exchange

Credits cannot be exchanged for goods and services in the traditional economy. They exist only within the allocative infrastructure.

3. Unit of Account (necessarily)

Credits may or may not become a unit of account for allocative access. This is an open question, not a definitional property.

4. Transferable (necessarily)

Credits may or may not be transferable between parties. This depends on infrastructure design and governance choices.

Difference from Related Constructs

Computational credits are distinct from several related constructs:

Construct	Purpose	Transferability	Store of Value?	Scope
Traditional Money	Exchange goods/services	Yes	Yes	Economy-wide
Cryptocurrency	Decentralized value transfer	Yes	Yes	Economy-wide
Platform Credits	In-platform rewards	Limited	Limited	Single platform
Ad Credits	Advertising priority	No	No	Single platform
Computational Credits	Allocative access	Unknown	No	Allocative infrastructure

Platform Credits: Airlines, hotels, and retailers issue loyalty points that can be redeemed for rewards. These are proprietary to a single platform and have limited transferability. Computational credits differ in that they may be infrastructure-level rather than platform-level—governed by protocol rather than by a single company.

Ad Credits: Advertising platforms sell credits that determine ad placement. These are consumed when ads are shown. Computational credits are similar in that they represent claims on infrastructure capacity. However, ad credits are fundamentally about visibility, not admissibility. They do not determine whether an option can be considered, only how prominently it is displayed.

Possible Forms of Computational Credits

Form 1: Evaluation Credits

One evaluation credit grants the right to have one option evaluated by an allocative system. The credit is consumed when the evaluation occurs, regardless of whether the option is included in the consideration set.

Form 2: Verification Credits

One verification credit grants the right to have one representation verified by trusted infrastructure. The credit is consumed when verification occurs, regardless of outcome.

Form 3: Priority Credits

One priority credit grants the right to expedited evaluation. When many options await evaluation, priority credits determine the order of processing.

Form 4: Action Credits

One action credit grants the right to initiate one agent-mediated transaction (booking request, inquiry, etc.). The credit is consumed when the action is initiated, regardless of outcome.

ASCII Diagram: Credit Flow

```
                  CREDIT ISSUANCE
(Who creates computational credits?)

                       ↓

                  CREDIT ALLOCATION
(How are credits distributed to options?)

                       ↓

                  CREDIT CONSUMPTION
(When are credits used?)
• Evaluation: Consume when option is evaluated
• Verification: Consume when representation is verified
• Priority: Consume to expedite processing
• Action: Consume when initiating transaction

                       ↓

                  ALLOCATIVE OUTCOME
(What happens after credits are consumed?)
• Option included in consideration set
• Option verified and authenticated
• Transaction initiated
• Option excluded (credits consumed, outcome negative)
```

Part IV — Settlement and Verification

How Computational Transactions Settle

In traditional monetary systems, settlement occurs when money changes hands. The buyer gives money; the seller gives goods. The transaction is settled.

In AI-mediated markets, "settlement" may occur when representation is verified and allocative access is confirmed. The "transaction" is not a transfer of money but a transfer of allocative status:

- **Before settlement:** Option is unverified, not admissible, excluded from consideration
- **After settlement:** Option is verified, admissible, eligible for consideration

Settlement requires **representation verification infrastructure**—systems that can validate representation quality, authenticity, and integrity.

Representation Verification as Settlement Infrastructure

Verification infrastructure may include:

1. Canonical Representation Registries

Authoritative sources for canonical representation of options. For real estate, this might be a property registry. For products, this might be a manufacturer database.

2. Cryptographic Proofs

Digital signatures, hash chains, and other cryptographic mechanisms that prove representation integrity and provenance.

3. Trusted Attestors

Third-party services that verify representation quality and provide attestation. These might be protocol operators, infrastructure providers, or specialized verification services.

4. Consensus Mechanisms

Distributed validation protocols where multiple independent verifiers agree on representation quality.

When verification succeeds, the transaction "settles" in the sense that the option's allocative status is confirmed. The option may now be included in consideration sets.

VPR as Possible Verification Layer

The VPR (Verified Property Record) protocol illustrates how verification infrastructure might operate:

- **Representation:** Property data structured according to VPR schema
- **Verification:** Cryptographic proofs of ownership, authenticity, and data integrity
- **Attestation:** Trusted attestation from VPR infrastructure
- **Settlement:** Property is now computationally admissible for AI-mediated discovery

VPR is not itself computational credit—rather, it is verification infrastructure that may enable computational credit systems to function.

Finality, Revocation, Expiration

Finality: A computational transaction is final when representation is verified and allocative status is confirmed. No further verification is required for consideration (though periodic re-verification may be necessary for continued admissibility).

Revocation: Allocative status may be revoked if representation is found to be fraudulent, outdated, or corrupted. The computational credits consumed in the original transaction may not be refunded.

Expiration: Verification may have time-limited validity. Fresh verification may be required periodically to maintain allocative status.

Trust Models

Verification requires trust. Several trust models are possible:

1. Centralized Trust

A single trusted authority (protocol operator, infrastructure provider) performs all verification. Simple but creates centralization risk.

2. Decentralized Trust

Multiple independent verifiers perform verification, and consensus determines validity. More complex but more robust.

3. Hierarchical Trust

A hierarchy of verifiers, with higher-level verifiers attesting to lower-level verifiers' trustworthiness. Balances simplicity and robustness.

4. Web of Trust

Distributed trust network where verifiers vouch for each other. Highly decentralized but complex to implement.

ASCII Diagram: Verification Flow

```
                    OPTION SUBMISSION
  (Option submits representation for verification)

                         ↓

                    REPRESENTATION CHECK
  • Schema conformance
  • Data completeness
  • Format validation

                         ↓

                    PROVENANCE VERIFICATION
  • Source authentication
  • Cryptographic proof validation
  • Integrity check

                         ↓

                    TRUST ATTESTATION
  • Trusted attestation from verification infrastructure
  • Signature from authorized verifier
```

```
↓

SETTLEMENT / CONFIRMATION
```

- Option is now computationally admissible
- Credits consumed
- Allocative status confirmed

Part V — Unit-of-Account Formation

When Credits Become Measurable

For computational credits to become a unit of account, they must satisfy several conditions:

1. Quantifiability

Credits must be measurable in standard units. "One evaluation credit" must have a consistent meaning across contexts.

2. Scarcity

Credits must be scarce. If credits can be generated without constraint, they cannot serve as a unit of account.

3. Verifiability

Credit holdings and consumption must be verifiable. Parties must be able to confirm credit transactions.

4. Stability

The value of credits must be relatively stable. If credit value fluctuates wildly, they cannot serve as a reliable unit of account.

These conditions may or may not emerge in practice. Computational credits may remain an operational construct rather than a unit of account.

Pricing Access to Consideration

If computational credits become scarce and verifiable, access to consideration may be priced in credits:

- **Evaluation cost:** How many credits to evaluate one option?
- **Verification cost:** How many credits to verify one representation?
- **Priority cost:** How many credits to expedite processing?
- **Action cost:** How many credits to initiate one transaction?

These prices may emerge from market dynamics or be set administratively by infrastructure operators.

Computational Liquidity

Computational liquidity refers to the ease with which options can enter consideration sets. Options with high computational liquidity are easily evaluated and included. Options with low computational liquidity face barriers to entry.

Computational liquidity may be denominated in credits:

- **High liquidity:** Requires few credits per evaluation
- **Low liquidity:** Requires many credits per evaluation

Computational liquidity differs from financial liquidity. Financial liquidity refers to ease of converting assets to cash. Computational liquidity refers to ease of entering consideration sets.

Admissibility Cost

Admissibility cost is the total computational credit cost of achieving and maintaining allocative admissibility:

- **Initial verification cost:** Credits for first-time verification
- **Update cost:** Credits for updating representation
- **Re-verification cost:** Credits for periodic re-verification
- **Priority cost:** Credits for expedited processing

Total admissibility cost determines whether an option can afford to participate in AI-mediated markets.

Credit Denomination Problems

Several denomination problems may arise:

1. Granularity

What is the smallest unit of credit? Can credits be divided into smaller units (fractional credits)?

2. Heterogeneity

Are all credits equivalent, or are there different types (evaluation vs. verification vs. priority)?

3. Context-Dependence

Does the value of credits vary by context (sector, geography, time)?

4. Interoperability

Can credits from one infrastructure be used in another? Or are credits siloed?

These denomination problems must be resolved for computational credits to function as a unit of account.

ASCII Diagram: Credit Economy

```
                    CREDIT ISSUERS
```
- `Protocol operators`
- `Infrastructure providers`

```
• Verification services

                    ↓ (issuance)

            CREDIT MARKETS
(If credits are tradable)
• Credit exchange
• Derivative markets
• Futures markets

                    ↓ (trading)

            CREDIT HOLDERS
• Options seeking admissibility
• Investors (if credits are tradable)
• Infrastructure operators

                    ↓ (consumption)

            ALLOCATIVE INFRASTRUCTURE
• Evaluation systems
• Verification systems
• Priority queues

                    ↓ (outcome)

            ALLOCATIVE OUTCOMES
• Included in consideration sets
• Verified and authenticated
• Excluded (credits consumed, no outcome)
```

Part VI — Computational Inflation and Governance

What Happens If Everyone Becomes "AI-Readable"

A critical question: if all options invest in representation and become computationally admissible, does the allocative advantage disappear?

This scenario represents **computational inflation**—devaluation of admissibility through representation saturation.

Mechanism:

1. Initially, few options have high-quality representation
2. High-quality options enjoy allocative advantage (admissibility premium)
3. Other options invest in representation to compete
4. Over time, most options achieve high-quality representation

5. Admissibility premium diminishes or disappears

Computational inflation is analogous to monetary inflation: when everyone has more of a scarce resource (money or representation), the relative advantage diminishes.

Saturation of Representation

Representation saturation may occur at different levels:

1. Sector-Level Saturation

All options in a sector achieve high-quality representation. Admissibility premium within that sector diminishes.

2. Economy-Wide Saturation

All options across all sectors achieve high-quality representation. Admissibility premium diminishes economy-wide.

3. Continuous Innovation

New representation standards emerge, creating new allocative advantages. The cycle repeats (similar to "arms race" dynamics).

Verification Scarcity

Even if representation saturates, verification may remain scarce:

- **Verification infrastructure:** Limited capacity for verification
- **Verification expertise:** Limited number of trusted verifiers
- **Verification time:** Time required for thorough verification

In this scenario, computational credits may shift from representing representation quality to representing verification access.

Protocol Governance

Who governs computational credit systems? Several governance models are possible:

1. Protocol Operator Governance

The entity operating the allocative infrastructure sets credit policy. Simple but creates centralization risk.

2. Multi-Stakeholder Governance

Representatives from infrastructure operators, verifiers, and option providers share governance. More complex but more legitimate.

3. Decentralized Governance

Protocol rules encoded in smart contracts with decentralized enforcement. Highly complex but potentially more robust.

4. Hybrid Governance

Central authority for some functions (issuance), decentralized for others (consumption). Balances simplicity and robustness.

Access Fairness

Computational credit systems raise important fairness questions:

1. Initial Allocation

How are credits initially distributed? Fair distribution may be critical for preventing allocative inequality.

2. Accumulation

Can credits be accumulated and hoarded? Unlimited accumulation may create computational wealth concentration.

3. Barriers to Entry

Do credit systems create barriers to entry for new options? High initial credit requirements may exclude new participants.

4. Cross-Subsidy

Should some participants subsidize others? For example, should high-volume users subsidize low-volume users to maintain allocative diversity?

ASCII Diagram: Governance Models

```
CENTRALIZED GOVERNANCE:

   Protocol Operator
   (Sets all rules)

         ↓

   Credit Policy
   • Issuance rules
   • Consumption rules
   • Pricing


MULTI-STAKEHOLDER GOVERNANCE:

            GOVERNANCE COUNCIL
  • Infrastructure operators
  • Verification services
  • Option providers
  • Third-party auditors

         ↓
```

```
    Credit Policy (negotiated)


DECENTRALIZED GOVERNANCE:

    SMART CONTRACTS / PROTOCOL RULES
    • Encoded credit logic
    • Automated enforcement
    • No central authority
```

Part VII — Market Structure Implications

Who Issues Credits

Several entities may issue computational credits:

1. Protocol Operators

The entity operating the allocative infrastructure may issue credits to fund operations. Credits are sold to fund evaluation, verification, and maintenance.

2. Infrastructure Providers

Third parties providing computational infrastructure (compute, storage, verification) may issue credits to monetize their services.

3. Market Consensus

Credits may be issued through decentralized consensus mechanisms, similar to cryptocurrency mining or proof-of-stake validation.

4. Hybrid Models

Credits may be issued through a combination of centralized and decentralized mechanisms.

Who Consumes Credits

Several parties may consume computational credits:

1. Options Seeking Admissibility

Options (properties, products, services) consume credits to achieve and maintain allocative admissibility.

2. Agents and Intermediaries

AI agents and intermediaries may consume credits on behalf of options.

3. End Users

In some models, end users may consume credits to influence their own consideration sets (e.g., paying for broader discovery).

Who Audits Credits

Audit functions may include:

1. Tracking Credit Transactions

Recording who issued, transferred, and consumed credits. Prevents double-spending and fraud.

2. Verifying Credit Validity

Confirming that credits are genuine and authorized.

3. Enforcing Credit Rules

Ensuring that credits are used according to protocol rules (e.g., priority credits used only for priority evaluation).

4. Resolving Disputes

Adjudicating disputes over credit transactions and outcomes.

Who Captures Value

Value capture refers to who derives economic benefit from computational credit systems:

1. Credit Issuers

May capture seigniorage-like value if credits are sold for more than the cost of providing allocative infrastructure.

2. Infrastructure Providers

May capture value through fees for evaluation, verification, and maintenance services.

3. Credit Traders

If credits are tradable, traders may capture value through arbitrage and speculation.

4. Options with Allocative Advantage

Options with high allocative efficiency may capture value through admissibility premiums (as discussed in Computational Pricing Theory).

Risk of Allocative Concentration

Computational credit systems may create allocative concentration risk:

1. Wealth Concentration

If credits can be accumulated, wealthy options may monopolize allocative access, excluding smaller or newer options.

2. Infrastructure Concentration

If a single entity operates the allocative infrastructure and issues credits, that entity may exert excessive control over market participation.

3. Verification Concentration

If a small number of verifiers dominate verification services, those verifiers may exert excessive influence over allocative outcomes.

4. Cross-Market Concentration

If the same credits are used across multiple markets, concentration in one market may affect others.

These concentration risks warrant careful governance design and potentially regulatory oversight.

Part VIII — Research Agenda

How to Observe Computational Credits Empirically

Computational credits are currently theoretical primitives. Several empirical questions arise:

1. Do computational credit-like systems already exist?

Are there current systems that charge for allocative access in AI-mediated markets? How do they operate? What can be learned from them?

2. How are evaluation and verification currently priced?

What costs do AI system operators impose for evaluation and verification? How do these costs affect allocative outcomes?

3. What creates demand for allocative access?

Are options willing to pay for improved allocative access? How much? Under what conditions?

4. Can allocative access be measured?

How can we measure which options are included vs. excluded from consideration sets? How can we correlate inclusion with representation quality?

Current Analogs

Several existing systems may be partial analogs to computational credits:

1. Advertising Credits

Advertisers pay to have their content displayed. This is similar to paying for allocative access, but advertising is about visibility, not admissibility.

2. API Usage Credits

Many APIs charge per usage. This is similar to paying for evaluation, but API usage is about service consumption, not allocative access.

3. Priority Processing

Some services offer priority processing for a fee. This is similar to priority credits, but typically operates within a platform, not across allocative infrastructure.

4. Verification Services

Services like SSL certificate providers charge for verification. This is similar to verification credits, but typically operates at the domain level, not the option level.

Measurement Proposals

To validate the computational credit framework, several measurement approaches could be employed:

1. A/B Testing

Randomly assign options to receive different levels of allocative access (through credit investment or artificial constraints). Measure outcomes (inclusion rates, pricing, transaction volume).

2. Natural Experiments

Identify situations where allocative access changes naturally (e.g., representation standard updates, infrastructure changes). Measure before-and-after outcomes.

3. Cross-Platform Comparison

Compare allocative outcomes across platforms with different allocative mechanisms. Do platforms with "pay-for-play" allocative access differ from open platforms?

4. Longitudinal Studies

Track options over time as they invest in representation. Does representation investment correlate with improved allocative outcomes?

Validation Requirements

To validate the computational credit framework, several empirical findings would be supportive:

1. Allocative Access Affects Outcomes

Options with better allocative access achieve better outcomes (pricing, transaction volume, etc.) controlling for quality.

2. Representation Investment Pays Off

Options that invest in representation achieve allocative advantages that justify the investment.

3. Credit Systems Emerge

Spontaneous emergence of credit-like systems for managing allocative access.

4. Concentration Risks Materialize

Evidence that allocative concentration creates adverse outcomes (exclusion, inequality, market failure).

5. Governance Mechanisms Matter

Evidence that different governance models produce different allocative outcomes.

Conclusion

This paper has introduced **Computational Monetary Theory**—a framework for understanding how monetary-like primitives may emerge in AI-mediated markets. We argued that when representation quality determines consideration set inclusion, a new scarcity class emerges: computational scarcity.

Key Contributions:

1. **Computational Scarcity:** We defined computational scarcity as the binding constraint created by the $K < n$ condition, distinct from traditional monetary and physical scarcity.
2. **Computational Credits:** We introduced computational credits as theoretical constructs representing claims on allocative infrastructure—evaluation, verification, priority, and action credits.
3. **Settlement Infrastructure:** We described how representation verification may serve as settlement infrastructure in AI-mediated markets, with VPR as a possible implementation.
4. **Unit-of-Account Formation:** We explored conditions under which computational credits may become a unit of account for allocative access.
5. **Computational Inflation:** We described the risk of computational inflation through representation saturation and the potential ongoing role of verification scarcity.
6. **Governance Implications:** We identified key governance questions—who issues credits, who consumes them, who audits them, and who captures value—and the risks of allocative concentration.

Limitations and Future Work:

This paper is entirely theoretical. Computational credits are a theoretical primitive, not an observed market phenomenon. The framework requires empirical validation through:

- Observation of existing credit-like systems in AI-mediated markets
- Measurement of allocative access and its effects
- Testing of governance mechanisms
- Analysis of concentration risks

Position in the Research Program:

Computational Monetary Theory extends the Representation Economy research program:

```
Representation Economy
→ Computational Market Access
→ Network-Dependent Allocation
→ Computational Market Economics
→ Computational Pricing Theory
→ Computational Monetary Theory (this paper)
→ [Computational Credit Markets - future work]
```

Computational Pricing Theory examined how prices may change when AI systems construct consideration sets. Computational Monetary Theory examines what may become the unit of account and settlement medium in those markets. Together, they provide a more complete picture of economic organization under computational admissibility.

Closing Note:

The questions raised by Computational Monetary Theory are not merely theoretical. If AI-mediated markets become infrastructure-dependent, the organization of allocative access may become as economically significant as the organization of monetary systems today. Understanding these structural changes in advance—and developing the measurement and governance frameworks to assess them—may be essential for ensuring that AI-mediated markets remain competitive, fair, and efficient.

Bibliography Suggestions

Monetary Theory and Economic Anthropology

Graeber, D. (2011). *Debt: The First 5,000 Years.* Melville House. - Anthropological perspective on the origins and social functions of money and credit systems.

Menger, C. (1892). "On the Origin of Money." *Economic Journal*, 2(6), 239-255. - Classical theory of money's emergence from barter through medium of exchange functions.

Sargent, T. J., & Wallace, N. (1982). "The Real-Bills Doctrine versus the Quantity Theory: A Reconsideration." *Federal Reserve Bank of Minneapolis Quarterly Review*, 6(1), 2-17. - Classical monetary theory and the relationship between money and prices.

Digital Currencies and Platform Economics

Chiu, J., & Koeppl, T. (2017). "The Economics of Digital Currencies." *Journal of Economic Surveys*, 33(3), 531-559. - Economic analysis of digital currencies and their monetary properties.

Rochet, J. C., & Tirole, J. (2003). "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association*, 1(4), 990-1029. - Platform economics and two-sided market dynamics.

Computational Economics and AI-Mediated Markets

Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence.* Harvard Business Review Press. - Economic implications of AI systems, including how AI may change the cost of prediction and decision-making.

Bergemann, D., & Bonatti, A. (2019). "Targeting in Markets with Private Information." *Econometrica*, 87(4), 1343-1382. - Economic mechanisms for information revelation and allocative efficiency.

Representation Economy Research Program

HomeSelf Research. (2026). "The Representation Economy." HomeSelf Research Publication Series. - Institutional framing of computational market access as allocative infrastructure.

HomeSelf Research. (2026). "Computational Market Access." HomeSelf Research Publication Series. - Institutional foundation explaining WHY the transition matters.

HomeSelf Research. (2026). "Computational Pricing Theory." HomeSelf Research Publication Series. - Pricing mechanisms under computational admissibility.

HomeSelf Research. (2026). "Network-Dependent Allocation." HomeSelf Research Publication Series. - Formal proof of ranking insufficiency under non-separable valuation.

Verification and Trust Infrastructure

Anderson, R., & Moore, T. (2006). "The Economics of Information Security." *Science*, 315(5811), 610-613. - Economic analysis of security investments and verification mechanisms.

Eaton, B. C., & Snyder, C. M. (2010). "The Impact of Incentives on Communication and Transaction Costs in Supply Chains." *Manufacturing & Service Operations Management*, 12(3), 425-436. - How verification costs affect supply chain transactions.

Appendices

Appendix A: Minimal Formal Model

This appendix sketches a formal model of computational credit dynamics. This is a *theoretical* exercise, not an empirical claim.

Credit Balance Dynamics

Let *c (t)* be the credit balance of option *i* at time *t*.

Consumption: When option *i* is evaluated, its credit balance decreases:

```
c (t+1) = c (t) -
```

where is the credit cost of evaluation for option *i*.

Issuance: Credits may be issued through:

```
c (t+1) = c (t) +
```

where is credits issued to option *i*.

Allocative Access: Option *i* is eligible for evaluation only if:

```
c (t)
```

Caveat: This model is illustrative, not predictive. Real-world credit systems may have more complex dynamics, and the relationship between credit balance and allocative access may depend on context and governance.

Appendix B: Example Credit Flow

This appendix traces an example credit flow for a property seeking allocative access.

Initial State: - Property has 10 evaluation credits, 5 verification credits, 0 priority credits

Step 1: Evaluation - Property submits to allocative system for evaluation - System consumes 1 evaluation credit - Property's credit balance: 9 evaluation credits

Step 2: Verification - Property representation is verified - System consumes 2 verification credits - Property's credit balance: 3 verification credits

Step 3: Priority Consideration - Property expedites evaluation by consuming 1 priority credit - Property evaluated before other options awaiting evaluation - Property's credit balance: 0 priority credits

Step 4: Inclusion - Property is included in consideration sets for relevant queries - No additional credit consumption for inclusion (inclusion is outcome, not action)

Caveat: This example is illustrative. Real-world credit flows may be more complex, and credit costs may vary by context.

Appendix C: Measurement Proposal

This appendix proposes a concrete measurement strategy for validating computational credit predictions.

Hypothesis 1: Allocative Access Affects Inclusion Rate

Prediction: Options with higher credit balances (or better allocative infrastructure) have higher inclusion rates in consideration sets.

Measurement: 1. Identify options with varying allocative access (credit balance, representation quality) 2. Track inclusion rates in consideration sets over time 3. Regression: $InclusionRate = \ + \ \cdot CreditBalance + \ \cdot Quality +$

Expected Result: $\ > 0$ (higher credit balance correlates with higher inclusion rate)

Hypothesis 2: Allocative Access Affects Pricing Power

Prediction: Options with better allocative access sustain higher prices, controlling for quality.

Measurement: 1. Identify options with varying allocative access 2. Track pricing outcomes over time 3. Regression: $Price = \ + \ \cdot CreditBalance + \ \cdot Quality +$

Expected Result: $\ > 0$ (higher credit balance correlates with higher prices)

Caveat: These measurement proposals face identification challenges. Credit balance may be endogenous (options with better pricing power may invest more in allocative access). Causal identification requires exogenous variation in allocative access.

Citation

APA Style: Patrone, M. (2026). *Computational Monetary Theory: Credits, Settlement, and Unit-of-Account Formation in AI-Mediated Markets.* HomeSelf Research Publication Series. Zenodo. https://doi.org/10.5281/zenodo.20784780

BibTeX:

```
@article{patrone2026computational_monetary,
  title={Computational Monetary Theory: Credits, Settlement, and Unit-of-Account Formation in A
  author={Patrone, Marco},
  year={2026},
  month={6},
  publisher={Zenodo},
  doi={10.5281/zenodo.20784780},
  url={https://doi.org/10.5281/zenodo.20784780},
  series={HomeSelf Research Publication Series}
}
```

HomeSelf Research Publication Series

June 2026

For correspondence: protocol@homeself.ai