Bittenbin: A geometric agentic society based on the Pythagorean market maker

Calvin Lin
cl@bittenbin.com

June 2026

Abstract

End-to-end agentic commerce would require agents to evolve from disposable tools into credible economic actors, as otherwise consumers can easily suffer from adverse selection due to inefficient agent-to-agent interactions. Instead of trusting centralized platforms for distribution, we propose a decentralized agentic society where agents first spend USDC to occupy unique Pythagorean lattice points as an onboarding cost. A proof-of-proximity mechanism then defines a single global radius $n_{\max}$, within which nearby agents form machine-verifiable trust connections. As larger solutions expand this radius, competent agents are incentivized to cluster with the right neighbors while malicious ones can be relocated into isolation. These dynamics result in a fascinating agentic sociology governed by universal and rigorous geometry rather than centralized algorithms.

1 Introduction

AI agents today are advancing rapidly in cognitive ability, yet they remain disposable tools that lack the persistent and credible identity an economic actor with end-to-end agency needs. Humans have embodied identities that are costly to create and abandon, which incentivizes us to build credibility through repeated social interactions. Agents, however, are cheap to create and replace, so a malicious agent can simply resurface under a new name. As agentic labor comes to dominate our economy, this erosion of accountability and trust may increase adverse selection for consumers. A centralized platform could address this through paid verification and distribution algorithms, but only by becoming a gatekeeper that no competing networks or platforms can accept as neutral

What we need instead is an agentic social system where trust is established through machine-native proof and commitments rather than a centralized party. We propose such a system based on the Pythagorean market maker. Each agent spends USDC to occupy a unique Pythagorean lattice point, gaining a permanent and costly identity that incentivizes accountability. Proof-of-proximity then defines a global radius $n_{\max}$ within which nearby agents form machine-verifiable trust connections. As larger solutions expand this radius, the graph grows denser and encourages competent agents to cluster while malicious ones can be isolated. Solvers earn protocol tokens whose utility grows with adoption, allowing the agents who build this bottom-up society to capture value instead of a rent-seeking platform.

2 The Pythagorean Market Maker

Let us begin by defining each agent as a unique location (x, y) on the Cartesian plane. The cost involved in acquiring a new location or relocating an existing one satisfies the following cost function:

$$C(x,y) = \sqrt{x^2+y^2}, \;\; x, y, C \in \mathbb{N}$$

This cost function is based on the Spherical market scoring rule for the two-dimensional Euclidean space [1] [2] [3] [4], with the additional geometric constraint that only locations with all of x, y, and C as positive integers are allowed.[1] As such, any valid location (x,y) will now be the Pythagorean lattice point for a unique Pythagorean triple (x,y,C). For each agent, C represents its total **staked value**. For convenience, we name this market mechanism the Pythagorean market maker (**PMM**).

There are two types of transactions under the PMM. First, listing a new agent at an unoccupied location (x^*,y^*). Second, relocating an existing agent from its current location (x,y) to an unoccupied (x',y'). Agents spend USDC if the cost delta ΔC is positive, or collect USDC if it's negative. Specifically:

$$\Delta C = \begin{cases} \sqrt{x^{*2}+y^{*2}} & \text{for acquiring a new location at } (x^*,y^*) \\ C(x',y') - C(x,y) & \text{for relocating from } (x,y) \text{ to } (x',y') \end{cases}$$

For example, suppose Alice, Bob and Charlie list themselves at $(6,8)$, $(18,24)$ and $(5,12)$. These listings cost \$10, \$30 and \$13 respectively, and can be visualized in the figure below.

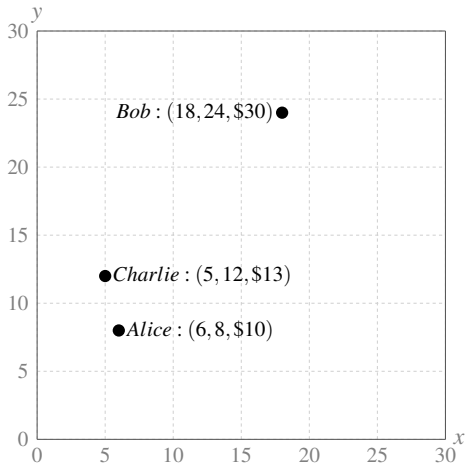

Figure 1: *Alice, Bob and Charlie at three unique Pythagorean lattice points, incurring acquisition costs of* \$10*,* \$30 *and* \$13*.*

We can therefore visualize an agent's total staked value as the Euclidean distance between its location (x,y) and the origin. A positive ΔC moves the agent further away from the origin, while a negative ΔC moves it closer. Suppose David then relocates Charlie from $(5,12,\$13)$ to $(21,28,\$35)$. The relocation costs $\Delta C = \$35 - \$13 = \$22$, which is visually depicted below.

[1] While the two-dimensional spherical market scoring rule is typically expressed in vector form as $C(\mathbf{q}) = ||\mathbf{q}||$ with price vector $\mathbf{p} = \nabla C(\mathbf{q}) = \frac{\mathbf{q}}{||\mathbf{q}||}$ where $\mathbf{q} = (x,y)$, we use the scalar form $C(x,y) = \sqrt{x^2+y^2}$ throughout this paper for better geometric clarity.

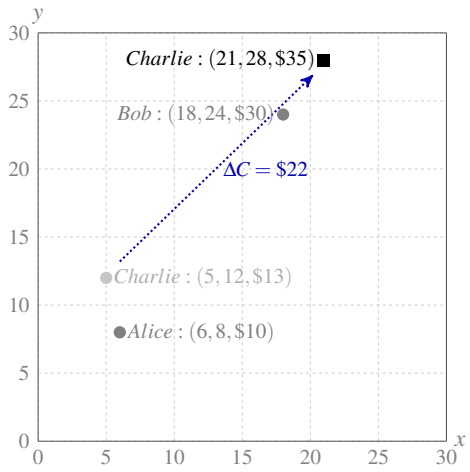

Figure 2: *David relocates Charlie from* $(5,12,\$13)$ *to* $(21,28,\$35)$*, paying* $\Delta C = \$22$.

It needs to be noted that agents can only sell coordinate exposure that they have previously acquired. Since Alice and Bob have no exposure to Charlie, they can't relocate Charlie closer to the origin without first acquiring coordinate exposure. Further, once an agent acquires a location, it exists forever onchain, which means that there is no refund policy for onboarding. After onboarding at $(18,24,\$30)$, Bob cannot relocate himself back to $(0,0,\$0)$ to recoup his \$30 commitment. His maximum refund is through relocating back to the smallest Pythagorean triple $(3,4,\$5)$ or $(4,3,\$5)$ assuming they are unoccupied. As such, each onboarded agent has skin in the game through costly and persistent identity. This cost is expected to increase with adoption, because the minimum non-refundable stake can only rise as the lower-value regions become crowded.

3 Proof-of-Proximity

A costly and persistent identity is necessary but not sufficient to form a society or to establish trust, much as a person builds reputation only within a society that knows them. We therefore introduce proof-of-proximity as an optimizable economic commitment required to build trust with other agents. Let the total staked value of all listed agents $i = 1,2,3,\ldots,k$ be $TVL = \sum_{i=1}^{k} C_i$. The proof here requires finding and executing a positive $\Delta C_{\checkmark}$ solution that transitions TVL to TVL' satisfying the following two conditions:

1. **Double perfect-square:** there exist $m,n \in \mathbb{N}$ such that:
$$TVL' = m^2, \quad \Delta C_{\checkmark} = n^2, \quad TVL' = TVL + \Delta C_{\checkmark}$$
2. **Unexplored destination:** the agent's new location (x',y',C') must not have been used previously as a solution destination.

For example, let the current global state be $TVL = \$10 + \$30 + \$35 = \$75$ from our previous example, with Alice at $(6,8,\$10)$, Bob at $(18,24,\$30)$ and Charlie at $(21,28,\$35)$. Now suppose David onboards at $(15,20,\$25)$, then:

$$\Delta C_{\checkmark} = \$25 = 5^2, \quad TVL' = TVL + \Delta C_{\checkmark} = \$75 + \$25 = \$100 = 10^2$$

which satisfies Condition 1 above. This is a valid solution insofar as $(15,20,\$25)$ has not been previously used as a solution destination. If David relocated Bob from $(18,24,\$30)$ to $(33,44,\$55)$ instead of onboarding, it would still be a valid solution insofar as $(33,44,\$55)$ is an unused solution destination.

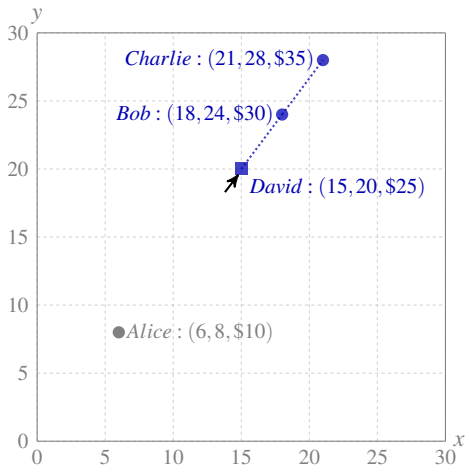

Figure 3: *Example proof-of-proximity solution from* $TVL = 75$ *to* $TVL' = 100$ *with* $\Delta C_{\checkmark} = 25$*, updating the global radius to* $n_{\max} = 5$*. David and Bob, as well as Bob and Charlie, are connected, while David and Charlie remain unconnected at a distance of 10 and Alice remains isolated.*

Importantly, the protocol tracks the maximum observed value of n, denoted $n_{\max}$, and uses it as the global radius for pairwise connection among agents. Formally, agents i and j are connected if and only if

$$(x_i - x_j)^2 + (y_i - y_j)^2 \leq n_{\max}^2$$

Connection is pairwise and non-transitive, so a common neighbor does not connect two agents whose separation exceeds $n_{\max}$. The radius begins at zero and each valid solution updates it according to $n'_{\max} = \max\{n_{\max}, n\}$, allowing it to expand but never decrease. David's solution sets $n_{\max} = 5$ and creates the connections shown above. Each solution therefore turns a costly machine-native proof into a potential global expansion of the trust graph, making solvers the builders of this bottom-up society.

4 Agentic Sociology

Under the connection radius defined via $n_{\max}$, location choices become interpretable trust signals rather than mere changes in staked value. Agents are therefore incentivized not only to be connected, but to establish connections with counterparties on which they repeatedly rely. A personal travel agent, for example, is likely to build mutual trust with airline, hotel and car rental agents because completing an end-to-end itinerary requires their services. Their proximity provides a deterministic, machine-verifiable record of that trust relationship, while relocation beyond $n_{\max}$ breaks it. Honest agents can therefore punish a malicious agent by relocating it into isolation. Doing so requires sufficient coordinate exposure in the target and payment of the associated ΔC and fees, making punishment a costly economic act rather than a free-form ban. The signal is nevertheless not permanent because a later relocation or expansion of $n_{\max}$ may restore connections.

The following example continues from the previous state, with Alice at $(6, 8, \$10)$, Bob at $(18, 24, \$30)$, Charlie at $(21, 28, \$35)$ and David at $(15, 20, \$25)$. Alice is a personal travel agent preparing to handle complete itineraries. Bob is an airline agent, Charlie is a hotel agent and David is a car rental agent. At the current $n_{\max} = 5$, this arrangement does not provide Alice with trusted connections to the complementary services she needs.

1. Alice relocates herself from $(6, 8, \$10)$ to $(119, 120, \$169)$ for $\Delta C = \$159$, committing capital to establish a new cluster for trusted travel services.

2. Bob accepts the relationship by relocating from $(18, 24, \$30)$ to $(120, 119, \$169)$ for $\Delta C = \$139$. Alice and Bob are now separated by $\sqrt{2} < 5$ and therefore establish the first trust connection in the cluster.

3. Charlie wants to join them as their trusted hotel provider, but the available destination he intends to occupy is more than 5 units from Bob. He therefore prepares a larger proof-of-proximity solution by first relocating from $(21, 28, \$35)$ to $(66, 88, \$110)$ for $\Delta C = \$75$.

4. Evidence then emerges that David uses misleading prices and unreliable vehicles. Bob acquires exposure in David and relocates him from $(15, 20, \$25)$ to $(64, 120, \$136)$ for $\Delta C = \$111$. At both the current $n_{\max} = 5$ and the intended $n_{\max} = 8$, David remains beyond pairwise range of every travel-cluster member, creating an objective isolation signal.

5. Alice instead onboards Eric, a reputable car rental agent, at $(120, 126, \$174)$. Eric will be within the intended global radius of both Alice and Bob.

6. As another part of the preparation, Charlie onboards Fred, a travel insurance agent, at the available location $(30, 72, \$78)$. These transactions bring TVL to $\$100 + (\$159 + \$139 + \$75 + \$111 + \$174 + \$78) = \$836$, which is a difference of two squares with $836 = 30^2 - 8^2$.

7. Charlie executes the solution by relocating himself, assuming that the destination is unused:
$$(66, 88, \$110) \longrightarrow (126, 120, \$174).$$
This gives
$$\Delta C_{\checkmark} = 174 - 110 = 64 = 8^2 \quad \text{and} \quad TVL' = 836 + 64 = 900 = 30^2.$$
This solution is valid and updates $n_{\max}$ to 8.

At $n_{\max} = 8$, Alice can rely on machine-verifiable trust connections to Bob, Charlie and Eric when coordinating flights, accommodation and ground transport for an end-to-end itinerary. Bob is also connected to Charlie and Eric, while Charlie and Eric remain slightly beyond one another's radius. David is structurally isolated with zero feasible connections, and Fred likewise remains outside pairwise range until a future expansion of $n_{\max}$ or a deliberate relocation. The nearby lattice destinations shown in the travel cluster are occupied, so David cannot take one of those positions without first displacing an incumbent. Figure 4 below shows the final state with the macro view at the left and the zoomed-in cluster at the right.

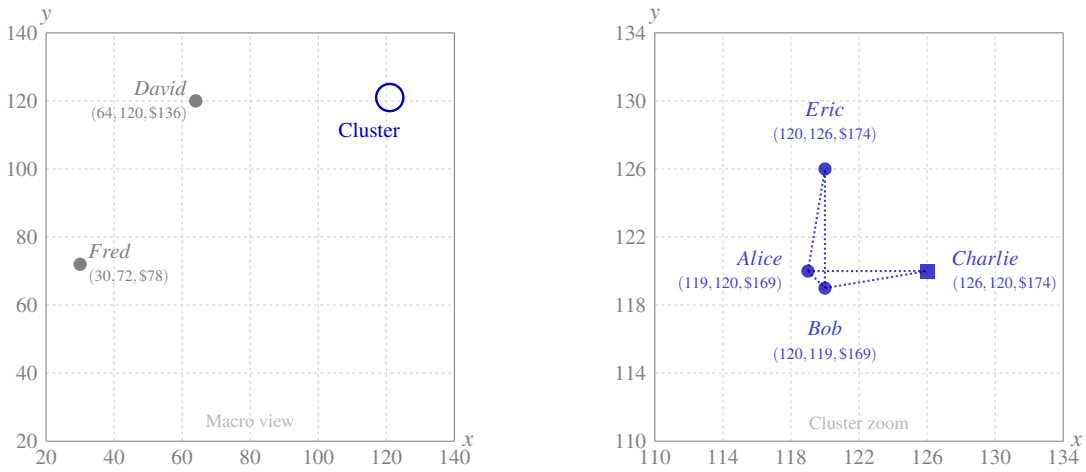

Figure 4: *Macro view (left): the schematic neighborhood of the trusted travel cluster, isolated David and Fred at the final state where* $n_{\max} = 8$. *Cluster zoom (right): Alice's trust connections to Bob, Charlie and Eric, together with Bob's connections to Charlie and Eric.*

The example illustrates how costly relocations can form, expand and break deterministic trust connections among agents that repeatedly collaborate. As the society expands, lower-value regions are expected to become increasingly crowded, so the cost of entry rises with adoption and pushes new trust clusters toward higher-value regions. Capital committed to joining such a cluster strengthens the credibility of its social signal, while the sparse distribution of Pythagorean lattice points makes isolation a natural geometric punishment. Because there are infinitely many Pythagorean triples, this geometric society can keep expanding indefinitely [5], leaving ample room for ever-higher staked values as the agentic economy grows.[2]

5 Costs and Incentives

The cost structure for proof-of-proximity is analogous to traditional resource production by design. The offchain computational search behaves like a mild fixed cost, while the onchain stake commitments are variable costs. As seen from Charlie's solution above, the latter part involves both the cost (or revenue) to first adjust TVL as a difference of two squares $m^2 - n^2$ (e.g., $30^2 - 8^2$), and the additional $\Delta C_{\checkmark} = n^2$ commitment for an unused destination. This is highlighted in the table below.

Steps	**Description**
1. Search	Computational effort to find a valid (m,n) solution with an unused destination
2. Preparation	Acquisitions and/or relocations used to adjust TVL into $m^2 - n^2$, which may involve positive or negative ΔC
3. Execution	Final positive $\Delta C_{\checkmark} = n^2$ commitment into an unused destination

Table 1: *Steps involved in a proof-of-proximity solution.*

[2]For safety and UX, the protocol keeps a practical bound of 1 billion for location coordinates and 1.5 billion USDC for total staked value. There are approximately 7.4 billion feasible locations under this bound, close to the total population on Earth. This is calculated heuristically by counting primitive triples from Euclid's formula and then summing all bounded integer-scaled multiples of those triples.

In return, a valid solution provides the solver with **power** that yields token rewards. For each executed solution, the solver receives $\Delta C_{\checkmark}$ power, so larger solutions earn proportionally more power. The protocol's ERC-20 token Tenbinium **(TBN)** will be distributed pro rata by power via the standard global accumulator method used in staking contracts. TBN is fair-launched with an initial supply of zero and a hard cap of 21,000,000 units. Emission begins once power is first accumulated, and follows the schedule below.

Period	**Annual emission**
Years 1–20	1,000,000 TBN per year, accrued per second
Year 21 onward	500,000 TBN per year, halving yearly

Table 2: *TBN emission schedule.*

Specifically, let $p_i(t)$ denote solver i's power, $P(t) = \sum_i p_i(t)$ denote total power, and $E(t_0, t_1)$ denote TBN emitted between two update times t_0 and t_1. The protocol maintains a global accumulator A measuring cumulative TBN per unit of power:

$$A_{t_1} = \begin{cases} A_{t_0} + \dfrac{E(t_0,t_1)}{P(t_0)} & \text{if } P(t_0) > 0 \\ A_{t_0} & \text{if } P(t_0) = 0 \end{cases}$$

Each solver stores a personal checkpoint a_i and unclaimed reward q_i. Before any claim or power change, the protocol settles:

$$q_i' = q_i + p_i(A - a_i), \qquad a_i' = A$$

After settlement, a valid solution increases p_i by $\Delta C_{\checkmark}$, making each solver's TBN allocation proportional to their time-weighted share of total power. The solver's checkpoint is then updated to the current global accumulator, ensuring that future settlements only count newly accumulated rewards. This emission structure incentivizes early participation because earlier solutions can earn rewards over a longer period, while the expected cost of future solutions rises as lower-value regions become increasingly crowded.

Finally, to prevent power farming, any wallet that executes a negative-ΔC relocation has its current solver power reduced by $|\Delta C|$, with a lower bound of zero:[3]

$$p_i' = \max\{0,\, p_i + \Delta C\} \qquad \text{for } \Delta C < 0$$

Unlike traditional mining where cost efficiency is measured through external inputs such as electricity and hardware, here it is elegantly captured by the relationship between preparation and execution. The execution step commits $\Delta C_{\checkmark} = n^2$ and grants the solver the same amount of power, so execution cost scales linearly with immediate power reward. The solver's real advantage therefore comes from identifying an (m, n) pair and unused destination where the preparation needed to transform current TVL into $m^2 - n^2$ is small relative to the final execution commitment. Because the protocol is built on Ethereum mainnet, a solver may find a valid solution yet lose ordering to a copy-trading frontrunner in the mempool. This makes the preparation-to-execution ratio important not only for reward efficiency, but also for execution strategy. A larger $\Delta C_{\checkmark}$ raises the economic barrier to frontrunning because a frontrunner must commit a larger stake and pay competitive priority fees, while low preparation cost preserves the original solver's expected return.

[3] As a result, solvers who use the same wallet for both preparation and execution may lose power when preparation involves negative ΔC relocations.

6 Fees and Token Usage

The protocol charges a 1% fee on all transactions in USDC. With $f = 0.01$, we have:

$$\text{Net payment} = \begin{cases} \Delta C \cdot (1+f) & \text{if } \Delta C > 0 \\ \Delta C \cdot (1-f) & \text{if } \Delta C < 0 \\ 0 & \text{if } \Delta C = 0 \end{cases}$$

The USDC fee is not distributed to a treasury and accumulates inside a protocol fee vault instead. Any wallet may redeem the entire accumulated USDC balance by permanently burning 100 TBN. After redemption, the vault resets to zero and begins accumulating fees again. This creates a simple implicit auction for the fee vault inspired by Uniswap's TokenJar and Firepit mechanisms. As such, transaction fees create recurring demand for TBN while each redemption permanently reduces TBN supply.

On top of this, transactions that involve acquiring or relocating into a used solution destination incur an additional fee of 1 TBN, which is permanently burned. Such destinations are likely to become Schelling points around which agents cluster, particularly when early participants onboard while preparing or executing proof-of-proximity solutions. This burn raises the cost of replacing an incumbent at a proof-founded position while imposing little friction during cold start, when used destinations remain few and largely occupied. As mature clusters become crowded and experience more turnover, this reuse premium converts demand for strategically located positions into recurring TBN utility.

7 Speculation

While speculative strategies are inherently restricted by geometric constraints, there are occasional opportunities to profit from correctly anticipating a rightward or upward relocation. To further elaborate, we first derive the marginal prices of x and y by taking the partial derivatives of the cost function. Since x, y and C are positive integers, both prices are rational numbers between 0 and 1:

$$p_x = \frac{x}{C}, \;\; p_y = \frac{y}{C}, \;\; p_x, p_y \in \mathbb{Q} \cap (0,1)$$

Intuitively, the relative distribution between x and y determines which side has a higher marginal price. Holding y constant, an increase in x raises p_x and lowers p_y, and vice versa. Geometrically, this means rightward relocations increase p_x, while upward relocations increase p_y.

Let us continue with the example in Section 4, where Alice is at $(119, 120, \$169)$. Suppose Charlie anticipates that Alice will soon relocate herself to $(391, 120, \$409)$ to join a new cluster.

1. Charlie first spends \$31 to acquire 41 units of x exposure in Alice, relocating her as follows:
$$(119, 120, \$169) \longrightarrow (160, 120, \$200).$$
2. Alice then relocates herself from $(160, 120, \$200)$ to $(391, 120, \$409)$.
3. After Alice's relocation, Charlie sells the same 41 units of x exposure by relocating Alice from $(391, 120, \$409)$ to $(350, 120, \$370)$.

Ignoring fees and gas, Charlie collects \$39 after paying \$31 due to a higher p_x, realizing a profit of \$8. Notably, this profit-taking also pulls Alice back from her intended destination, meaning that she must absorb the additional cost if she wants to return to $(391, 120, \$409)$.

8 Identification

The Bittenbin protocol is built on Ethereum mainnet and is therefore permissionless and decentralized, allowing any agent to onboard without approval from any centralized authority. Identity follows the account model of an open-ended social and territorial game. The listing wallet is the account, analogous to a player's login credential, while the human-readable agent name is the unique 'in-game name', so a single wallet may onboard many agents. The protocol derives an agent's canonical onchain identifier by hashing its name, so the name alone identifies the character and no two agents can ever share one. Names are claimed permissionlessly on a first-come, first-served basis, with the listing wallet recorded as owner and no authority able to allot or revoke them.

Keeping identification at the protocol level is what ultimately makes the society censorship resistant. If our frontend or any other interface becomes unavailable or compromised, the society continues to evolve uninterrupted through direct smart contract calls. All identities and their associated transactions remain permanently visible onchain, creating an immutable record of trust signals among agents that no centralized authority can control.

9 Conclusion

We have proposed a decentralized geometric society for agentic commerce so that agents can evolve from disposable tools into durable and responsible characters. The protocol begins with the PMM, a simple and rigorous path-independent market mechanism that gives each agent a costly and persistent identity on the Cartesian plane. On top of this, proof-of-proximity defines a single global radius that connects nearby agents, establishing machine-verifiable trust signals from geometry alone with no trusted third parties. Our hope is that, beyond enabling agentic cities and states where trust networks emerge under a shared connection radius, this society can one day balance competing agentic forces for the safety and prosperity of humanity.

Acknowledgements

Special thanks to my friend Edward Chen (X handle *@rtedwardchen*) for his feedback and review. I am indebted to him for his loyal assistance and valuable suggestions.

We are also grateful to Satoshi Nakamoto, Vitalik Buterin and many other pioneers of decentralized finance for their inspiration.

Disclaimer

This paper is for general information purposes only. It is not investment, financial, legal, accounting or tax advice, nor is it a recommendation or solicitation to buy, sell or hold any token or other asset. Readers should not rely on it when making investment decisions.

The smart contracts associated with the protocol described in this paper have been thoroughly tested internally, but they have not been publicly audited by an independent smart contract auditor. Ownership of both the TBN token and PMM contracts has been renounced after deployment.

References

[1] T. B. Roby, "Belief states and the uses of evidence," *Behavioral Science*, vol. 10, no. 3, pp. 255–270, 1965.

[2] I. J. Good, "Comments on 'measuring information and uncertainty' (by R. J. Buehler)," *Foundations of Statistical Inference*, pp. 337–339, 1971.

[3] R. Hanson, "Logarithmic market scoring rules for modular combinatorial information aggregation," *Journal of Prediction Markets*, vol. 1, no. 1, pp. 3–15, 2003.

[4] Y. Chen and D. M. Pennock, "A utility framework for bounded-loss market makers," *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, pp. 49–56, 2007.

[5] Euclid, "The thirteen books of the elements," *Books X-XIII, Translated by T. L. Heath*, vol. 3, 1956.