Tenbin: A decentralized reputation map via the Pythagorean market maker

Calvin Lin
cl@tenbin.finance

Edward Chen
rchen@tenbin.finance

July 14, 2025

Abstract

Digital platforms need reliable ways to distinguish credible content from misinformation, yet existing review and rating systems are either too cheaply manipulated and thus require centralized moderation or too speculative. We introduce a novel decentralized solution that maps reputation as unique Pythagorean lattice points. This geometric approach transforms stake-backed reputation measurement from a speculation game into meaningful territorial signaling, where proximity indicates similar reputation levels and monetary values. With real stakes involved, manipulations become costly and positions can be used to symbolize social status and connections. The mechanism operates as a permissionless and censorship-resistant public good that can withstand major currency crises, and the underlying mathematical principles can be applied to any opposing binary choices for measurements beyond reputation.

1 Introduction

Reputation is perhaps our oldest social technology — the invisible yet invaluable currency that enables trust between strangers and cooperation at scale that brought humanity to where it is today. Traditional societies often relied on repeated interactions within small communities to build trust, yet these methods don't usually apply to our global digital world where identities are fluid and communities are vast. Without reliable mechanisms to effectively measure and broadcast reputation, digital social platforms become fertile grounds for misinformation and fraud. When anyone can create content with equal apparent authority — whether they're a seasoned expert or a coordinated bot network — the signal-to-noise ratio plummets. The problem is likely to get worse as AI makes convincing but false content cheaper to produce than ever before.

Various systems have attempted to address this issue. Traditional rating and review mechanisms have helped to some extent, but they usually lack real economic stakes and thus can be cheaply manipulated. Centralized oversight is needed to reduce spams and manipulations, leaving reputation measurement in the hands of a single authority. Prediction markets like Polymarket do involve real stakes and are thus useful in measuring probability of verifiable events, but they become speculative when it comes to subjective choices that can't be objectively settled [1] [2]. In this paper, we propose a novel solution that turns this type of speculation game into a decentralized social territorial game. The solution is a new type of market maker that adds specific constraints to the two-dimensional Spherical market scoring rule [3] [4] [5] [6]. This market maker allows any entity's reputation to be mapped as exclusive positions via Cartesian coordinates. These positions are backed by real economic stakes, making them more credible and costly to manipulate, and their geometrical significance provides unique social and territorial value, encouraging transactions to be based on thoughtful commitments rather than pure speculation.

2 Position

Let us begin by defining each entity's **position** as (x,y) on the Cartesian plane. Here, x and y are respectively the amount of **Distrust** and **Trust** votes placed towards an entity, satisfying the following cost function:

$$C(x,y) = \sqrt{x^2 + y^2}, \;\; x, y \geq 0$$

This cost function represents the Spherical market scoring rule (SMSR) for the two-dimensional Euclidean space [5] [6], where marginal prices for buying and selling Distrust and Trust votes are continuous and monotonic. With two outcomes, these prices would sit on the unit circle instead of the unit sphere, and can be computed by taking the partial derivatives of the cost function[1]:

$$p_x = \frac{x}{\sqrt{x^2 + y^2}}, \;\; p_y = \frac{y}{\sqrt{x^2 + y^2}}, \;\; p_x, p_y \in [0,1]$$

Like other market scoring rules, the SMSR can be used as an automated market maker for binary options, with its convex and path-independent cost function ensuring smooth pricing and bounded losses for the market maker [5] [6]. The problem is that these mechanisms don't work too well for binary choices that can't be objectively settled (e.g., Distrust/Trust). From the voters' perspective, they would instead be buying votes at p_x or p_y only to sell it higher to someone else, turning the market into a speculation game.[2]

What we need is an alternative mechanism where voters and entities also benefit from non-monetary value from trading or receiving votes, in particular social value through status, connections, governance and commitments. To accomplish this, we created a specialized version of the two-dimensional SMSR by adding the following two spatial constraints:

1. **Scarcity:** Only positions where x, y, and C are positive integers are allowed. That is, instead of $x, y \geq 0$, we now have $x, y, C \in \mathbb{Z}_+$. Consequently, instead of $p_x, p_y \in [0,1]$, we now have $p_x, p_y \in \mathbb{Q} \cap (0,1)$.

2. **Exclusivity:** Each unique position (x,y) can only be occupied by at most one entity at any given time.

From the constraints above, we can conclude that any valid position (x,y) will now be the Pythagorean lattice point for a unique Pythagorean triple (x,y,C). For each entity, C represents their **market value** and also the total cost incurred by voters. Marginal prices p_x and p_y are now discrete and restricted to rational numbers in the

[1] While the two-outcome 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 geometrical clarity.

[2] The unique Bayesian-Nash Equilibrium would have all voters decline to buy votes at any positive price [1].

interval $(0,1)$. For convenience and ease of future references, we name this type of market maker as the Pythagorean market maker **(PMM)**.

3 Reputation

We define each entity's **reputation**, $R = {p_y}^2$, as a percentage score such that $R \in (0\%, 100\%)$. The closer R gets to 100%, the more reputable the entity becomes. However, R can never actually reach 0% or 100%, as p_y can never reach 0 or 1 due to our Scarcity constraint. This is intuitive from two perspectives. From a philosophical perspective, even the most reputable entities may not be completely honest all the time, and even the least trustworthy entities may not always tell lies. From a probability standpoint, the reputation score R works like an implied probability for binary outcomes under the two-dimensional SMSR[3], but since Trust and Distrust are our subjective perceptions rather than deterministic facts, the probability can never reach 0% or 100%. Hence, the mathematical bounds on any entity's reputation naturally acknowledge and reflect that the world is full of uncertainties, and that there are limits to human and machine judgement.

Geometrically, we can find entities with the same reputation by drawing a straight line from the origin through each position out to infinity. Positions that sit on the same line will always have the same reputation, and each of these lines is a unique **reputation line**. As an example, in Figure 1 below we can see that entities with position $(3,4)$ and $(6,8)$ will both be on the $R_1 = 64\%$ reputation line, whereas the entity with position $(8,6)$ is on the $R_2 = 36\%$ reputation line.

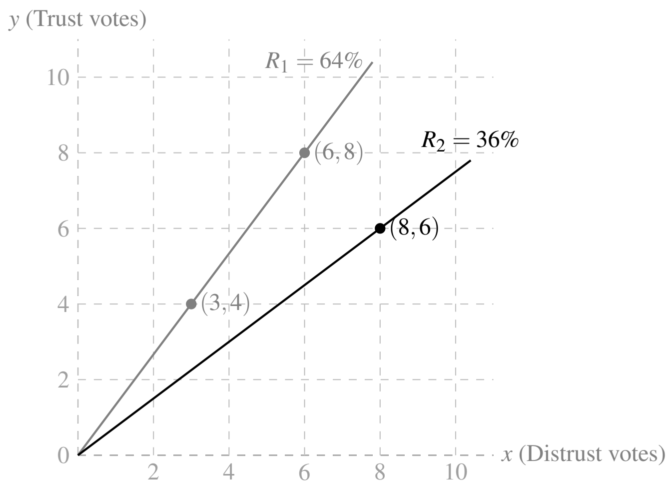

Figure 1: Reputation lines for positions $(3,4)$, $(6,8)$ and $(8,6)$.

Mathematically, it's obvious that any positions (kx, ky) where $k > 0$ such that $kx, ky \in \mathbb{Z}_+$ will have identical reputation. Hence, the Exclusivity constraint enables

[3] Recall that p_y is the marginal price of buying Trust votes. Under the context of binary options, its reciprocal $1/p_y = C/y$ is the marginal payoff of Trust votes, so neglecting fees and time value of money, this ratio of price to payout, ${p_y}^2$, gives you the implied probability of that the outcome will be 'Trust'.

voters to signal their intensity of commitment, through different values of k, upon entities on the same reputation line. For example, Alice might create Charlie the dishonest yoga teacher's market at $(24,7)$ with a market value of $25, and Bob the dishonest financial influencer's market at $(2400,700)$ with market value of $2,500. Under this scenario, despite both Bob and Charlie having same reputation of approximately 7.84%, the stakes that Alice committed for Bob is 100 times greater, signaling to society that Bob has greater reach and thus greater capacity for harm, or that Alice has substantially higher confidence in her assessment of Bob's untrustworthiness. This strength of commitment also applies to voters themselves as a way to broadcast social status. For instance, Alice could create her own market at $(700,2400)$ instead of $(7,24)$ to signal that her 92.16% reputation is more credible than positions of lower market value on her reputation line. From entities' perspective, this mathematical property somewhat reflects the sociological reality that while higher market value signals greater social status, it also amplifies social impact from both supporters and critics.

4 Transactions

The PMM supports two types of transactions – creating a new entity's market like what Alice did to Bob and Charlie at an unoccupied position (x^*,y^*), or moving an existing entity from position (x,y) to an unoccupied (x',y'). Voters spend USDC to create markets or buy votes, or collect USDC by selling votes. The cost calculation ΔC, ignoring fees, works as follows:

$$\Delta C = \begin{cases} \sqrt{x^{*2}+y^{*2}} & \text{for market creation} \\ C(x',y') - C(x,y) & \text{for buying/selling} \end{cases}$$

With a **one-percent transaction fee** ($f = 0.01$), the actual cost (or revenue if cost is negative) becomes:

$$\text{Cost} = \begin{cases} \Delta C \cdot (1+f) & \text{for buying/market creation } (\Delta C > 0) \\ \Delta C \cdot (1-f) & \text{for selling } (\Delta C < 0) \end{cases}$$

From voters' perspectives, transactions under this market maker are incentivized to be based on authentic social commitments, as opposed to pure speculation. While in many cases they can still make a handsome profit, that may no longer be the primary motivation as speculative strategies are inherently restrictive. With less speculation, voters face lower financial risk in general as they are more likely to break even by selling back their vote holdings if they change their mind about an entity, because it would typically require thoughtful commitments from other voters to move the entity away from that position. This creates a safer environment where honest voters can trust that most transactions are genuine and thus positions do usually reflect reality.

That said, it's important to note that market creators don't benefit from the 'refund policy' mentioned above, as market creation can't be undone. Once an entity's market is created, it exists forever onchain. If you create Alice's market at (7,24), you cannot sell back to (0,0) to recover your full investment. The most you can reduce her position to is the smallest Pythagorean lattice point (3,4) or (4,3), meaning market creators always bear some irreversible cost for bringing an entity into the system.

4.1 Manipulations

With real stakes involved, reputation data under the PMM is arguably more credible than traditional rating and review systems, especially for positions with higher market values where manipulations become increasingly costly. Consider our earlier example with Charlie at $(24,7)$ and Bob at $(2400,700)$, both with 7.84% reputation. Charlie could fake his reputation to 97.26% by spending just $120 to reach $(24,143)$, but Bob would need to spend $12,000 to achieve the same result at $(2400,14300)$ and a market value of $\$14,500$ — 100 times more expensive than Charlie. Likewise, imagine David, a highly reputable figure at $(24000,143000)$ with 97.26% reputation and $145,000 market value. An attacker who wants to maliciously move David to the 7.84% reputation line would need to spend at least $512,500 - $145,000 = $367,500, to move him to $(492000,143500)$.

4.2 Visualizations

We can visualize an entity's market value as the Euclidean distance between their position (x,y) and the origin, and their reputation change via the angle θ it makes with the x-axis, where $\theta \in (0,\frac{\pi}{2})$. A higher θ corresponds to an increase in reputation, and a simple way to achieve this is through a vertical movement from (x,y) to (x,y') where $y' > y$. This is shown in Figure 2 below, where both Charlie and Bob achieve identical angular changes from θ to θ', yet Bob's movement is far more costly as he needs to travel a greater distance to reach his new position at R'.

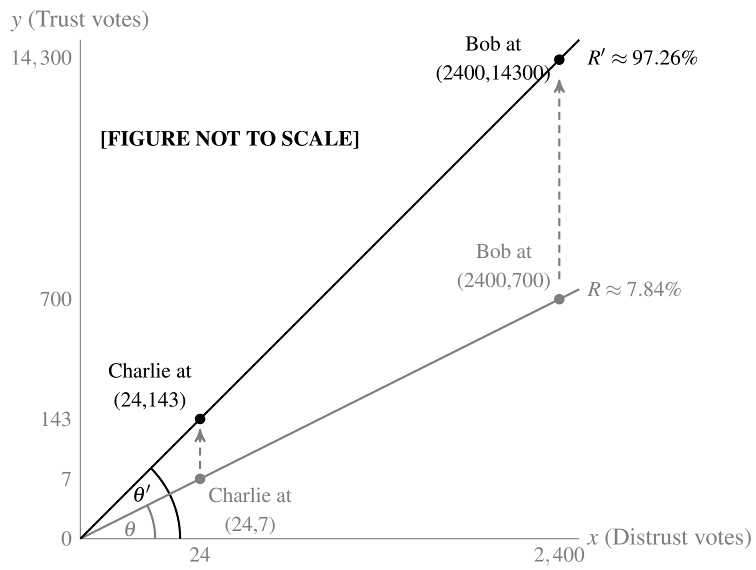

Figure 2: Reputation movements from approximately 7.84% to 97.26% for both Bob and Charlie.

Similarly, a lower θ indicates a drop in reputation. A simple way to achieve this is through a transaction that moves the target entity horizontally from (x,y) to (x',y) where $x' > x$. In the example of David who's positioned at $(24000,143000)$, the attacker could decrease his reputation by moving him to $(x',143000)$ instead of $(492000,143500)$ where $x' > 24000$, and the possible values of x' can be searched via the following steps:

1. Find all factor pairs (f_1, f_2) of 143000^2 where $f_1 \cdot f_2 = 143000^2$, $f_1 = C' - x'$, $f_2 = C' + x'$ and $f_1 < f_2$.[4]

2. Filter for $x' > 24000$ by computing $x' = (f_2 - f_1)/2$ for each factor pair and selecting those satisfying the constraint.

For our example, 143000^2 has a total of 151 factor pairs where $x' > 24000$. In other words, if an attacker or an opposing voter wants to decrease David's reputation without buying any Trust votes, there are 151 possible ways to do so, ranging from small commitments that slightly reduce David's reputation (e.g., $x' = 32175$), to expensive ones that achieve severe reputation damage, where x' can be as large as approximately 5.11 billion.[5] In Figure 3 below we demonstrate two examples, one with David's reputation reduced to $R' \approx 88.22\%$ at a cost of \$7,245, and the other with his reputation dropping below 10% to $R'' \approx 9.03\%$ at a cost of \$330,750.

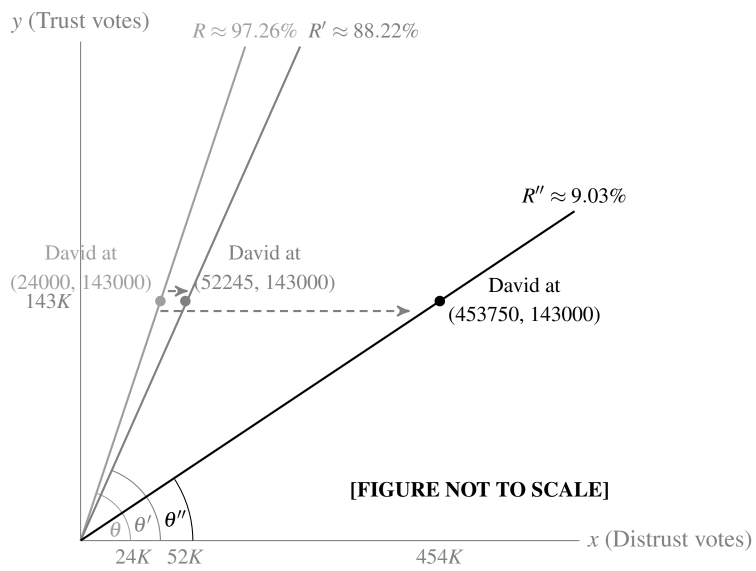

Figure 3: Reputation movements from approximately 97.26% to 88.22%, and 97.26% to 9.03% for David.

4.3 Monetary Profit

For profit-seeking voters, a relatively simple strategy is to 'be early' to either increases in the target entity's reputation with no changes to Distrust votes, or decreases in reputation with no changes to Trust votes. We can bring David back again to elaborate on the latter scenario using an example related to Figure 3 above.

1. Alice moves David from (24000, 143000) to (52245, 143000) by spending \$7,245 for 28245 Distrust votes.

[4]This can be done via a direct approach that checks all factors from 1 to 143000, or a prime factorization approach [7] that's substantially more efficient. For the latter, we would first compute $143000 = 2^3 \cdot 5^3 \cdot 11 \cdot 13$, then derive $143000^2 = 2^6 \cdot 5^6 \cdot 11^2 \cdot 13^2$ and generate all factors directly with minimal additional computation.

[5]Because 143000 is an even number, the largest Pythagorean triple can be written in the form of $(n^2 - 1, 2n, n^2 + 1)$, which in this case is $(5112249999, 143000, 5112250001)$ with $n = 71500$.

2. A wealthy Ethan thinks that David is a scammer, so he moves him to (453750, 143000) by spending $323,505 for 401505 Distrust votes.

3. Being a profit-maximizing voter, Alice sells 18315 Distrust votes, moving David from (453750, 143000) to (435435, 143000), netting a profit of $17,435 - $7,245 = $10,190 and now holds 9930 Distrust votes.

In this example, it's important to note that Alice might prefer David's reputation not to decline beyond $x > 453750$, where she could get 'stuck' as gaps between consecutive Pythagorean lattice points may exceed her portfolio of 28245 Distrust votes. As such, profit-maximizing voters anticipating substantial reputation swings may employ sophisticated hedging strategies, buying opposing vote types to reduce portfolio lock-up risk. These strategies are more complicated as they often require navigating through the irregular and complex patterns inherent in prime factorizations, and thus won't be discussed further in this paper.

5 Reputation Map

With the examples and illustrations provided above, one could view the reputation system as a dynamic **reputation map** where entities compete for positions that are both socially and economically valuable, in particular those that are close to the y-axis. From a social perspective, these positions are useful not only as proof of reputation and status, but also a gateway for connecting with other nearby entities. In particular, entities could signal their willingness to connect by positioning themselves close to others, a phenomenon that reflects the age-old principle that 'birds of a feather flock together.' Our system naturally facilitates this geometrically as spatial proximity indicates similar reputation and market values.

As such, we might see different social clusters form and evolve over time, and with continuous fiat inflation, lower market value regions are likely to become increasingly crowded, creating a continuously expanding reputation map under Euclidean geometry. Because there are infinitely many Pythagorean triples [8], the map can expand indefinitely, making the system resilient to major currency devaluation. In the hypothetical and unlikely case that the US dollar experiences hyperinflation, entities can simply reflect this by increasing their dollar-denominated stakes to preserve the real value of their position.

6 Decentralization

The Tenbin protocol is decentralized and thus permissionless, allowing anyone to create a reputation market for any entity without requiring approval from that entity or other centralized authorities. Specifically, each entity is uniquely identified via an arbitrary onchain key `platformID` at the protocol layer, and will be mapped to the corresponding ID of some centralized social platform (e.g., X, Instagram or TikTok) at the application layer. Market creators interacting with our frontend bear the responsibility of correctly entering the target entity's social platform ID as the `platformID`, as only those that correspond to existing accounts on that platform will be mapped and provided with context in our interface. Creators who misidentify targets risk creating markets that voters will ignore, and future users of that social platform whose IDs were mistakenly used as `platformID` can have their positions adjusted by themselves or other voters as needed.

This permissionless design contributes to the protocol's censorship resistance. If our frontend or any other interface becomes unavailable or compromised, the underlying mechanism continue operating uninterrupted through direct smart contract interactions. All transactions and positions remain permanently visible onchain, creating an immutable record of reputation assessments that no centralized authority can control. This transparency fosters a bottom-up ecosystem where other applications can build different mechanisms on top of this reputation layer, including for example incentive distribution, content curation and governance.

7 Future Works

Future works of the Tenbin protocol include both vertical integrations to enhance territorial incentives within the system, as well as horizontal expansions into decentralized platforms and measurements beyond reputation.

Vertical integrations may involve spatial features at the protocol level where each entities' position unlocks additional functionalities — for example, as a governance anchor for surrounding regions. Horizontal expansions may include native integration with decentralized social platforms (e.g., Farcaster), and applying the universal mathematical principles under the PMM to other pairs of orthogonal preference dimensions based on society's needs. These include Like/Dislike axes for individuals, companies, governments and countries to measure their popularity, and Buy/Sell for stocks, real-estates and other asset classes to weigh their value. Each of these pairs creates its own two-dimensional Euclidean space where different types of entities can be thoughtfully evaluated via stake-based data points that are hard to manipulate.

8 Conclusion

Digital platforms struggle with a basic problem: how to effectively determine what and who to trust online without a trusted party? Traditional rating systems can be manipulated at a small cost because they lack real stakes and thus require centralized enforcements. Prediction markets avoid this through economic backing, but often become speculation games when applied to subjective judgments. We've proposed a decentralized solution via the Pythagorean market maker that turns this type of speculation game into a territorial one where entities compete for meaningful positions on a reputation map. Economic stakes remain central as they help to mitigate manipulations, but geometric properties enable rich social signaling beyond simple betting.

This protocol operates as a permissionless and censorship-resistant public good, ensuring no single entity can control how society measures trust and reputation. The underlying mathematical principles extend well beyond reputation measurement, enabling thoughtful evaluation of any opposing binary preferences, and are also resilient against major currency crises.

References

[1] P. Milgrom and N. Stokey, "Information, trade and common knowledge," *Journal of Economic Theory*, vol. 26, no. 1, pp. 17–27, 1982.

[2] M. Brunnermeier and D. Abreu, "Bubbles and crashes," *Econometrica*, vol. 71, pp. 173–204, 2003.

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

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

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

[6] 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.

[7] Euclid, "The thirteen books of the elements," *Books III-IX, Translated by T. L. Heath*, vol. 2, 1956.

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