

The Prediction Network

A richer grammar for online participation

An account of how structured, shareable predictions can help communities gather believers, turn care into action, and create futures that do not yet exist.



Brand	The identity people choose to follow, support or belong to
Campaign	The collective movement from a present state towards a shared future
Topic	A prediction question and outcome set published by a Campaign administrator
Campaign point	A non-cash balance used only inside the Campaign that issued it
Task	The practical action a person can take to move the Campaign forward
Token	A record of recognition, contribution or belonging, separate from Campaign point balances
Prediction	A public stake of Campaign-scoped points on a Topic outcome, with shareable reasoning

Why this plan exists

Communities shape people. People shape communities. Every meaningful community is trying to create a future that does not yet exist—and a future does not just happen; it gathers believers.

Most platforms give people a thin grammar for online participation: like, comment, share and follow. The Prediction Network asks what becomes possible when belief about a shared future is structured as content. A question can carry confidence, reasoning, commitment, action and resolution. The product terms in this plan are not the idea itself. They are the grammar that lets communities gather believers and move together towards the futures they are trying to create.

Contents

1 Communities shape people. People shape communities.	2
2 From people to futures	2
3 The Prediction Network	3
4 Support as investment	4
5 Market position and competitive probability	6
6 Growth and revenue model	7
7 Trust, investment integrity, and operating model	8
8 Product roadmap after MVP	9
9 Potential pilot hotspots	10
10 Near-term pilot approach	10

1 Communities shape people. People shape communities.

Most of us can name a place, a group or a cause that changed who we are. It might be a neighbourhood where we learned the streets, a club where we kept showing up, a workplace that became part of our identity, an open-source project we carried with us, or a cause that outlived any single member. We did not simply join these communities. They entered us, and we, in turn, left marks on them.

This is the bargain at the heart of every lasting community: it shapes the people inside it, and the people inside it shape the community. Neither exists without the other. A club without members is a building. A member without a community is an individual. The alchemy happens in the loop between them.

Mutual shaping

Communities shape people. People shape communities. Neither exists without the other.

We also become ourselves in the presence of others. Our individual humanity—what we believe, what we hope for, what we are willing to stand for—is often experienced most clearly when it is shared. A belief held alone is a thought. A belief held with others is a beginning. This is why people seek communities: not only to belong, but to experience their own convictions reflected, tested and amplified by people who care about the same future.

The strongest communities are not merely groups we join. They are environments we enter, learn and eventually carry with us. They have a rhythm, a language, a memory and a set of expectations about how people behave. Over time, participation changes what we notice, what we value and what we are willing to do for other people.

A person may arrive in a neighbourhood intending to stay for a year. They learn the cafés, join a local group, volunteer, start a project, meet neighbours. Five years later they are still there, or they have left but the place still feels like home. The same pattern appears in smaller institutions that may matter even more deeply: universities, workplaces, sports clubs, religious communities, advocacy groups and online projects. Not every place becomes a community. Not every community becomes part of a person's identity. The special ones do.

Immersion

Immersion is not passive exposure and it is not time spent nearby. It is the gradual process by which a person learns the place, forms relationships, accepts some of its responsibilities and begins to care about its future.

Digital communities can create the same effect. A person does not stay on a platform because of the interface; they stay because they have become immersed in something that matters to them. The photography, the fitness tracking, the server, the fandom, the political movement. At their best, these are environments people carry with them.

Most social networks, however, are designed to help people consume communities. The Prediction Network should help people become more deeply immersed in them. Immersion is not measured by scrolling time. It is measured by how much a person comes to care about the future of the place and the people within it.

2 From people to futures

The history of online connection can be read as a series of expansions. Email connected people to people. Social media connected communities to each other, making it easier for a place or an identity to find its audience. The next layer is less obvious but no less important: connecting communities to the futures they are trying to create.

The arc of connection

Email connected people. Social media connected communities. The Prediction Network connects communities to futures.

Every meaningful community is trying to create a future that does not yet exist. A neighbourhood wants to be safer and more connected. A club wants to survive its next decade. A charity wants to reach more people. A creator wants to finish something ambitious. A team wants to fill a stadium. These futures are not guaranteed. They must be imagined, argued for, worked towards and believed into being.

People rarely achieve anything meaningful on their own. We remember the person holding the trophy or standing at the lectern, but most outcomes are built by the people behind them: friends, colleagues, volunteers, customers, neighbours and strangers who happened to care at the same time. Communities turn attention into effort, effort into trust and trust into things that looked unlikely until they happened.

The internet has made these communities easier to find and easier to hear. A like says somebody noticed. A share helps the message travel. A comment adds another voice. These are useful signals, but they are light luggage. A like may mean agreement, amusement or nothing more than a thumb arriving before the mind.

Online interaction already carries forms of currency: attention, trust, reputation, belonging and effort. Most platforms provide only a thin grammar for expressing them through likes, comments, shares and follows. Some outcomes ask for more. We may believe a fundraiser will reach its target, a local event will succeed, a team will fill a stadium or a community will find enough volunteers. We may also believe a target will be missed unless more people join in. Both views can be useful. Support that cannot admit trouble is merely advertising in a friendly hat.

Meaningful communities create futures

Every meaningful community is trying to create a future that does not yet exist.

The work is hard. Futures do not arrive fully formed. They compete for attention, resources and belief. They meet resistance, apathy and doubt. Yet communities keep trying, because the alternative is to let the future be decided by people who care less. The Prediction Network exists to make that caring more visible, more discussable and more connected to action.

3 The Prediction Network

A future does not just happen. It gathers believers.

Every community that has ever built something worth keeping had to do this work: imagine a future, name it in public, convince others that it was worth their attention, and keep believing long enough for action to catch up. The future does not arrive because it is logical. It arrives because enough people decide it should.

The prime directive

A future does not just happen. It gathers believers.

The Prediction Network is the grammar that helps communities do this gathering. It is not another social feed, another workspace or another points table. It is a way for a community to name the future it wants, invite people to stand behind it, connect that stand to practical action, and remember who helped move the needle.

3.1 What the product names

For years Statstrade was explained as a prediction market. The description was technically correct and practically incomplete. Prediction was only one part of the behaviour that mattered. Beneath it sat immersion in a place or group, the identity formed through that immersion, the future its people cared about, the work they did after expressing a view, and the way contribution was remembered.

The same pattern kept returning. There was a place or identity people gathered around, a present state they wanted to change, a shared future they could imagine, actions they could take and some form of recognition for helping. The product needed names for these things so the design would stop wandering.

Statstrade calls the public identity a **Brand**. A Brand may be a company, charity, club, creator, team, institution or place. It is the environment people recognise, enter and may eventually carry as part of themselves.

Statstrade calls the collective movement from a present state towards a shared future a **Campaign**. A Campaign is not merely a single question with a closing date. It is the common direction and public story that lets people recognise what they are moving towards together.

Campaign administrators publish **Topics** inside that movement. Each Topic asks a specific question and lists the outcomes people can predict. One Topic may ask whether a target will be reached. Another may compare several possible paths or test whether a milestone will happen by a certain date. Topics can close and resolve at different times while the wider Campaign continues.

The Campaign gives participants a limited supply of non-cash **Campaign points** that belong only to that Campaign. People can stake them across Topic outcomes inside the same Campaign, explain their reasoning and share the resulting prediction through social media. The stake makes confidence visible; the explanation makes it discussable; the share turns the prediction into content and carries the Campaign beyond the people already following it. The prediction may travel; the point balance does not.

Statstrade calls the practical opportunities to contribute **Tasks**, and the badges and other recorded forms of recognition **Tokens**. Tokens preserve contribution and belonging, but they do not make Campaign points portable. Communities do not need to admire the vocabulary. They only need the machinery to reflect patterns they already understand.

The vocabulary of movement

Brand — the identity people gather around.

Campaign — the movement from a present state to a shared future.

Topic — an administrator-published question with defined outcomes.

Campaign point — a non-cash balance used only inside the Campaign that issued it.

Task — a practical action that moves the Campaign forward.

Token — a record of recognition, contribution or belonging, separate from Campaign points.

Prediction — a public stake of Campaign-scooped points on a Topic outcome, with shareable reasoning.

3.2 From belief to action

The Prediction Network connects those patterns into a richer grammar. Following expresses affinity. Immersion develops identity and belonging. Staking Campaign-scooped points expresses confidence. Reasoning makes a position discussable. Sharing turns that position into content. Tasks connect belief to action. Resolution gives the interaction memory. Tokens and levels preserve contribution and standing over time.

Prediction becomes most interesting when the person forecasting is also holding a shovel. Traditional prediction markets try to separate the forecast from the outcome so the market can measure belief cleanly. Communities are rarely so tidy. The person predicting the result may also help produce it.

A small point stake can make that view more deliberate. The points are supplied by the Campaign, not bought by the participant. They remain inside the Campaign that issued them and cannot be carried into another Campaign, transferred to another person or collected in a general account balance. Their purpose is to let somebody show relative confidence: not merely "I agree", but "this is where I am willing to place the limited points I was given for this Campaign."

This is **visible belief**: a public statement about the future of a place, cause or group that already matters to the participant. People already do this when they share a crowdfunding campaign, recommend a new company, invite friends to an event or say that a local organisation deserves support. The Prediction Network gives that behaviour structure, memory and a path towards action.

Visible belief is not blind cheerleading. A committed participant may say that confidence is falling, the target looks doubtful or the Campaign needs more hands. Belonging is useful precisely because it can carry both hope and candour.

The product thesis

The Prediction Network deepens the relationship between people and the communities that already shape them. It makes belief about shared futures visible, connects that belief to reasoning and action, and preserves the contribution, standing and leadership that emerge over time.

3.3 Executive summary

Statstrade is being designed as a consumer social network for deeper participation in the places, groups and causes people care about. People join Campaigns that move communities from one state towards another. Campaign administrators publish Topics and issue non-cash points scoped to that Campaign. Participants stake those points only on outcomes within the same Campaign, explain and share their predictions, complete useful Tasks and build a history inside the communities they support.

The opportunity is larger than adding prediction buttons to a feed. Most social networks help people consume communities. Statstrade is intended to help people become immersed in them: to care about their future, make that care visible, act upon it and build standing through repeated contribution. Structured prediction is one mechanism within that larger relationship, not the relationship itself.

Online communities already exchange attention, trust, reputation, belonging and effort. Statstrade gives those interactions more structure, more memory and more ways to lead into action. Organisations appear as public identities and publishers, but the product should not feel like enterprise software handed to passive members.

The commercial opportunity lies in strengthening relationships that already have value. Communities need reach, participation, recognition and evidence of momentum. Sponsors need credible ways to support communities without barging into the conversation. Individuals may value richer identity, history and discovery. Statstrade can charge for additional reach, publishing, moderation, sponsorship, premium personal tools and aggregated insights while keeping ordinary participation accessible.

The central question is whether structured and shareable belief can deepen immersion: creating better discussion, stronger public commitment, meaningful return for resolution and a clearer path from belonging to contribution, standing and leadership.

4 Support as investment

The internet already gives communities distribution. A person can discover a cause, follow a club, subscribe to a creator or share an event with hundreds of people in seconds. The harder work is turning that attention into a relationship that lasts long enough to matter.

Most online support begins with a small gesture: a like, a comment or a message to a friend. These actions help information travel and may begin something deeper. Most platforms, however, treat the gesture as the end. The community receives a number, the

person moves on and the rest of the story is misplaced.

The Prediction Network should assume that support develops over time. A person may notice a community through somebody else's public view, return because the outcome remains uncertain, begin by sharing, contribute in a practical way and eventually become known by the people already involved. The product should allow that progression without marching everybody through the same formal process.

4.1 Investment, not surveillance

There is a difference between watching a community and investing in it. A like is a glance. A prediction backed by reasoning is a declaration. A Task completed is a practical bet on the future. The language we use changes what the action feels like.

"Accountability" can quickly sound like surveillance or blame: who slipped up, who must answer, who is being measured. That is not the spirit of the Prediction Network. The people in a community are not employees to be monitored. They are members who have chosen to place some of their attention, trust, reputation and effort in a shared future.

Investment, not surveillance

Support is an investment of attention, trust, reputation, belonging and effort in a future the community is trying to create.

This shift matters. When somebody stakes Campaign points on a Topic, completes a Task or declares public belief, they are not being held to account. They are choosing to invest in a future. The platform's job is to make that investment legible, safe and connected to real movement, not to turn every action into a scorecard.

4.2 Discovery and standing

Communities know that support has weight and history. Somebody who has volunteered for ten years has a different relationship with a club from somebody who discovered it yesterday. A person who repeatedly organises events, brings in members or backs uncertain Campaigns accumulates standing. There may be no neat score for it, but the community knows it is there.

The internet measures reach and reaction well. It is less useful at showing how a person has grown inside a community. Followers tell us how many people are watching. Likes tell us which posts travelled. They do not tell us who showed up, who helped, who backed an uncertain outcome or who became trusted over time.

Online participation can therefore feel disposable. Each action is recorded, yet the relationship between actions is weak. A person may support the same cause for years without the system showing their movement from observer to contributor, contributor to organiser, or organiser to somebody the community depends upon.

Levelling up matters here, but not merely as a game mechanic. People naturally progress through communities. They learn the language, understand the history, accept responsibility and become known for particular contributions. A useful network should make that progress visible without pretending every community values the same things.

Standing is contextual

Somebody may be trusted deeply in one community and entirely new in another. The system should preserve the difference rather than flattening everybody into one rank.

Public prediction gains context from this history. A view from somebody with years of relevant contribution means something different from the same view offered by a newcomer. Nobody is inherently worth more, but the history explains why a judgement may deserve attention.

4.3 What communities receive

A community receives more than another publishing tool. It gains a public surface where support becomes visible, people can act before the outcome is decided and contribution can accumulate into standing. It may later pay for additional reach, moderation, sponsorship, reporting or custom presentation, but the underlying activity must remain accessible enough to travel through the wider network.

The commercial balance is simple to state and difficult to fake. Statstrade should earn money from communities and sponsors without turning every relationship into procurement. The network becomes valuable when support can move among people and communities, not when each organisation is sealed inside its own account.

The community loop

Immersion → identity → belonging → contribution → standing → leadership → a stronger place.

The first experience therefore matters. A newcomer should not meet a complicated market interface or an administrative workspace.

They should meet a community, understand what it is trying to achieve and see that other people are willing to stand behind it. Prediction gives them a way to say where they stand and a reason to return, but it need not occupy the whole stage.

Once the relationship begins, the network becomes more useful. People discover communities through people they trust. Communities reach beyond their mailing lists. Sponsors support real participation rather than buying a lonely patch of attention. The platform gains a history of how belief, contribution and standing develop around different outcomes.

5 Market position and competitive probability

Statstrade does not fit neatly into one software category. That is an opportunity, though opportunities have been known to arrive wearing the same coat as risks. The product borrows behaviours from social networks, community platforms, loyalty systems, creator tools, forecasting communities and campaign software. The useful comparison is not which label fits best, but which parts of the model each category already supports and what remains missing.

Social networks are strong at identity and distribution but weak at giving a collective movement durable structure. Community platforms are strong at membership and discussion but usually keep each community inside its own workspace. Loyalty products are strong at points and rewards but often reduce participation to transactions. Forecasting platforms are strong at questions, probability and calibration but usually treat action and community ownership as secondary. Campaign tools are strong at goals and reporting but rarely create a persistent public identity for the people carrying the movement.

Statstrade's proposed difference is not the presence of Brands, Campaigns, Topics, Tasks and Tokens by themselves. It is the interaction grammar they create together: a person can move from following an identity to expressing confidence, explaining a position, sharing it as content, taking action and returning for resolution. The business case depends on whether that richer grammar produces useful behaviour that adjacent products do not.

5.1 What the market already proves

Social platforms show that people value persistent identity and discovery. Products such as Strava and Letterboxd show that structured activity can become social identity rather than remain a private record. Forecasting communities show that people will express confidence and return for resolution. Loyalty systems show that non-cash points can affect behaviour. Community and creator platforms show that recurring publishers can bring audiences into a network.

What remains unproven is whether these behaviours can be joined without making the product feel like homework or gambling. A person should not need to understand an economic system before joining a movement. An administrator should be able to publish a clear Topic without becoming a market operator. A Campaign-scoped point stake should make confidence more expressive without creating financial risk or an account-wide currency. Tasks and recognition should remain useful after the novelty wears off.

5.2 The competitive boundary

Statstrade does not need to replace Facebook Groups, Discord, a council consultation portal, a volunteer roster or a charity donation system. Those products are often systems of record, and there is no prize for rebuilding every filing cabinet. Statstrade can sit across them as a public participation and discovery layer. A Task may link to an official registration system. An administrator-published Topic may resolve from an existing roster, vote, event result or published report. A Campaign may draw evidence from all of them without becoming the legal record of membership, payment or attendance.

This boundary reduces the operational software Statstrade must rebuild and keeps the product focused on the part most likely to create a network: a richer public grammar connecting identity, confidence, reasoning, sharing, action, recognition and outcome.

The defensible asset

The defensible asset is not a generic campaign builder or a points table. It is the accumulated history of structured participation: what people followed, predicted, explained, shared, acted upon and learned at resolution. If that history improves discovery, trust and future participation, the network becomes more useful as it grows.

5.3 Questions for initial pilots

The next step is not to declare the model proven. It is to work with a small number of communities and organisations on movements they already care about, then observe what actually happens. Early pilots can test whether administrators can frame Topics clearly, whether participants understand that supplied points remain inside the issuing Campaign, whether staking makes predictions more deliberate, and whether people are willing to explain and share those predictions beyond the platform.

The organisations involved should help shape the Topic list, point allocation, moderation, resolution and reporting rules. Their needs may differ across communities. The purpose of the pilots is to discover where the model is useful, where it creates friction and which assumptions deserve to survive.

6 Growth and revenue model

Revenue must follow the product. A convenient pricing table is not a strategy. If Statstrade is a Prediction Network, revenue should come from helping communities invest more effectively in the futures they are trying to create. It should not require every Brand to buy a licence before the network is useful, and it should not require participants to lose money.

The base consumer experience should remain broadly free. People need to follow Brands, join Campaigns, receive Campaign-scoped points, stake predictions, share their reasoning, complete Tasks and build a visible history without meeting an enterprise sales process. Free participation is not charity; it is the raw material of the network.

The growth hypothesis is that a public prediction can alter the path of a Campaign. When somebody stakes supplied points, explains their view and shares it, they are not only measuring confidence. They are investing public belief in a future. Other people can react, challenge the reasoning or join the Campaign. The Brand can see where confidence is weak, which predictions are travelling and where Tasks may help. As evidence changes, people can revise their positions and the Campaign's story develops.

The network should therefore measure more than prediction volume. It must test whether public belief, shared reasoning and practical action change behaviour. The strongest early evidence would be a repeatable relationship between Topic participation, shared predictions, Task completion, referrals, recognition and visible Campaign progress.

6.1 Revenue sequence

Early revenue is most likely to come from Brands and sponsors who want to deepen their investment in a community's future. A Brand may pay to promote a Campaign or Topic, manage point allocations within that Campaign, coordinate teams, run advanced moderation, customise presentation or analyse participation. A sponsor may fund Task rewards or a shared recognition programme. These payments should support activity that remains public and useful to the people in the network.

Premium consumer membership can follow once personal history has value. People may pay for deeper calibration, private groups, saved research, richer profile controls, stronger archives or an advertising-reduced experience. The premium layer should add depth rather than make ordinary participation feel unfinished.

Advertising should come after the feed has earned attention. Insights and APIs should come after Statstrade has enough Topic outcomes and Campaign history to offer more than another sentiment dashboard. Enterprise services may matter for large Brands, but they should extend the network rather than become its organising principle.

6.2 Illustrative three-year case

The following figures are planning assumptions, not a forecast. They test whether a modest consumer network can support several revenue streams without relying on enterprise contracts or participant losses.

Table 1: Illustrative consumer-network model after public launch (AU dollars)

Metric	Year 1	Year 2	Year 3
Registered people	40,000	180,000	600,000
Monthly active people at year end	12,000	70,000	240,000
Monthly active Brands	250	1,500	5,000
Active Campaigns per month	400	3,000	12,000
Premium consumer conversion	1.5%	2.5%	3.5%
Average premium price per month	\$8	\$9	\$10
Premium revenue	\$17,280	\$189,000	\$1,008,000
Promoted Campaigns and sponsorship	\$40,000	\$300,000	\$1,200,000
Advertising and other network revenue	\$0	\$90,000	\$720,000
Total illustrative revenue	\$57,280	\$579,000	\$2,928,000

The model is deliberately staged. Advertising contributes nothing in Year 1. Premium conversion remains low. Promotion and sponsorship begin with Brands already publishing useful Campaigns and Topics. The Year 3 case depends on active people, repeated participation and a growing history of Campaign movement and resolved Topics rather than a handful of large contracts.

6.3 The network equation

Monthly network movement

New active people come from shared predictions, creator distribution, Brand broadcasts, Tasks, search and direct discovery.

Retained active people return for new Topics, changing evidence, revised stakes, Campaign progress, Tasks, recognition and resolution.

Network value grows when public predictions affect participation, Tasks affect the shared future, and the resulting history improves future discovery.

The product should track whether shared predictions lead to referrals or Task completion; whether participants understand that each point balance belongs only to its Campaign; and whether people move into another Brand or Campaign without carrying the previous Campaign's points with them. These measures belong inside pilots, not inside a claim that growth has already been solved.

6.4 Pilot questions before growth assumptions

Before the growth model deserves confidence, Statstrade needs a small set of well-designed pilots with communities and organisations. Each pilot should begin with a movement the partner already understands. Administrators can publish a limited Topic set, supply a clear point budget scoped to that Campaign and invite participants to stake and share predictions. The work is to learn whether this produces useful discussion, whether the points make confidence clearer, whether Tasks support participation and whether the organisation can govern the activity without creating another administrative burden.

The pilots should also reveal what cannot be standardised. Different communities may need different Topic rules, point allocations, moderation, evidence, privacy or measures of progress. The aim is not to force every pilot through the same funnel. It is to identify which parts of the model travel well, which depend on local practice and what evidence would justify broader investment.

6.5 Capital efficiency

Early capital should fund product quality, Campaign, Topic and Task design, Campaign-scoped point and staking controls, moderation, community partnerships, pilot delivery, creator tools, Brand seeding and the infrastructure required for feeds and resolution. Sales headcount and enterprise integrations should remain limited until the participation loop has been tested in practice. A weak model cannot be repaired by selling more administration features.

7 Trust, investment integrity, and operating model

The influence hypothesis makes trust more important, not less. A popular forecast is not proof, and a confident crowd can still march in the wrong direction. A Brand can misuse urgency. A Campaign can create pressure instead of understanding. Topics can reward heat over truth. Points can become manipulative when people are unclear about why they were issued, where they can be used or what a stake means.

The answer is not surveillance or blame. It is investment integrity: a shared understanding of what people are investing, why it matters, how it is governed and how it is protected.

Investment integrity

Trust protects the shared investment. Integrity means people understand what they are investing, why it matters and how it is governed.

Statstrade therefore needs clear rules for every primitive. Brands need identifiable ownership and clear public identity. Campaigns need a clear present state, shared future, purpose, governance and boundaries. Campaign administrators need responsibility for the Topics they publish, the outcomes they list, the points they issue and the evidence used for resolution. Tasks need legitimate actions, evidence and limits. Tokens need transparent issuance, expiry, acceptance and revocation. Predictions need visible confidence, history and correction without being presented as fact.

The first product should use non-transferable points supplied by and scoped to a single Campaign. Participants should not deposit money, buy points, withdraw them, transfer them to another person, move them into another Campaign or Brand, place them in an account-wide wallet, or redeem them for cash. Staking means allocating that Campaign's limited, non-cash point balance across Topic outcomes published inside the same Campaign. It is a signal of confidence within that movement, not a portable currency or financial wager.

Sponsors may fund separate rewards or shared unlocks, but they should not control Topic definitions, participant point balances, prediction distributions or resolution. Those rewards must remain separate from Campaign point balances. Transferable value, participant-funded prizes and secondary trading require separate legal and product review.

7.1 Prediction and staking integrity

Every prediction-bearing Topic needs a clear close, outcome set, resolution source and correction history. A Campaign administrator can explain the question and publish evidence, but should not quietly rewrite it after people have staked points. Material changes require visible versioning, cancellation or a new Topic. The wider Campaign may continue after one Topic resolves, because a movement can contain many questions and possible outcomes over time.

Point issuance must be understandable. Participants should know how many points they received, whether everybody received the same balance, whether more points can be earned, how the balance changes when a Topic resolves and whether a stake can be changed before closing. The interface must state that the balance belongs only to the issuing Campaign and cannot be used elsewhere. It should never imply that points are money or that prediction success guarantees a financial reward.

Brand ambassadorship must be declared. Paid, rewarded or formally sponsored support must be labelled. Honest shilling ceases to be honest the moment the bargain is hidden. Ambassadors should remain free to express doubt, revise a view and explain why confidence has changed.

The recommendation system should not reward ambiguity merely because it produces reaction. Distribution should favour Topics that are relevant to the Campaign, evidence that can be inspected and participation that remains constructive. A Brand or Campaign that repeatedly publishes misleading, abusive or unresolved Topics should lose reach and eventually publishing privileges.

7.2 Task integrity

A Task is an invitation to influence a Campaign's movement, so it needs a legitimate relationship to the shared future. Sharing, donating, volunteering or attending can be appropriate. Harassment, sabotage, brigading or pressuring an individual are not. The platform should show what evidence is required, who verifies it, whether completion is public and what recognition follows.

Tasks should not exist merely to inflate activity metrics or distribute points without purpose. When Tasks affect a Campaign point balance, those points remain inside that Campaign. A Campaign with thousands of meaningless clicks is less useful than one with a smaller number of verified actions. The operating model should distinguish self-reported activity, provider-verified events and Brand-approved evidence.

7.3 Token integrity

Campaign points are not portable Tokens. They cannot be accepted by another Campaign or Brand, transferred between people or converted into a network-wide balance. Badges, recognition records and sponsor-funded rewards may have separate rules, but the product must not imply that those rules make Campaign points portable. Every issuer must state what it issues, where it applies, how it can be revoked and whether it has any use beyond recognition.

7.4 Social identity and privacy

A person's history across Brands, Campaigns, Topics, Tasks, Tokens, point stakes and predictions can reveal political, cultural, religious, health, financial and local interests. Statstrade should treat that history as sensitive personal data. People need granular visibility controls, private or pseudonymous participation where appropriate, export and deletion controls, and clear separation among public predictions, private reasoning, credentials and billing data.

Similarity and recommendation should be explainable and controllable. The product can explain that two people follow similar Brands or stake on related Topics. It should not present an unexplained score as an objective judgement of personality.

7.5 Responsibility model

Statstrade is responsible for platform policy, recommendation integrity, prohibited Topics, technical audit history, account controls, appeals, correction mechanisms and escalation. Brands and Campaign administrators are responsible for Campaign purpose, Topic publication, point issuance, Task legitimacy, moderation, evidence and resolution. People remain responsible for their public reasoning and conduct.

Normal activity should not require staff approval. Automated checks, reputation, Brand moderation and post-publication review should handle routine use. High-risk allegations, minors, medical claims, elections, unusual prizes, financial products and legal disputes require specialised policy or restrictions.

8

Product roadmap after MVP

The roadmap should follow the participation loop rather than the complexity of the underlying market engine.

1. **Brands:** public identity, follows, governance, moderation and history.
2. **Campaigns:** present state, shared future, public narrative, membership, progress, governance and moderation.

3. **Topics:** administrator publishing, outcome sets, evidence, closing, resolution and social sharing.
4. **Points and prediction:** Campaign-scoped balances, staking only within the issuing Campaign, confidence, reasoning, revision history and resolution.
5. **Tasks:** action definitions, verification, referral, completion history and safety controls.
6. **Tokens and recognition:** badges and other recognition records kept distinct from Campaign point balances, with explicit issuance, expiry and revocation rules.
7. **Discovery and monetisation:** feeds, similarity, promoted Campaigns and Topics, sponsorship, premium tools and later insights or APIs.

The MVP should be judged by whether the full loop can be tested with pilot partners: administrators define a shared movement, publish Topics and issue points scoped to that Campaign; participants stake those points only within the same Campaign, explain and share predictions; Tasks provide ways to act; and people return to see how Topics resolved and whether the Campaign moved. Enterprise integrations should follow demonstrated demand rather than lead the roadmap.

9 Potential pilot hotspots

The most useful pilots may not begin with a single organisation. They may begin with a place, event or local network where the same people already move among several communities. A person can be a surf club volunteer, a council resident, a charity donor, a parent, a runner, a customer and a local business owner in the same week. That overlap is where the network hypothesis becomes visible.

A pilot hotspot is therefore more than a list of organisations in one suburb. It is a cross-cutting intersection of people, communities, identities and shared concerns. One organisation may administer the first Campaign, publish the initial Topics and issue points scoped to that Campaign. Participants may then carry their predictions into neighbouring communities through the relationships they already have, while the point balance remains with the original Campaign.

Potential hotspot	Cross-cutting intersection	Future the community is trying to create	Questions to investigate
Bondi and the eastern beaches	Surf clubs, Waverley Council, charities, schools, local businesses, residents, sponsors, events and visitors	A more connected, prepared and supported coastal community where local organisations share people, attention and momentum	Can administrators publish Topics that matter across communities? Do Campaign-scoped point stakes encourage people to explain and share predictions?
City2Surf network	Runners, charities, teams, employers, sponsors, volunteers, councils and neighbourhoods along the route	A connected movement that turns individual registration into shared training, fundraising and neighbourhood participation	Can an event administrator use Topics and Campaign-local points to connect smaller community Campaigns without flattening their identities?
Community readiness network	Emergency services, Red Cross groups, surf clubs, councils, neighbourhood organisations and local volunteers	A state of visible local preparedness where passive awareness becomes coordinated action before it is needed	Which Topics are safe and useful to publish? Can point staking surface confidence or concern without compromising privacy?
Civic accountability network	Residents, advocacy groups, councils, community media, professional associations and service organisations	A public culture where commitments are visible, evidence is inspected and follow-through becomes collective expectation	Can administrators publish fair outcome sets and resolve them credibly? Does public reasoning remain constructive when predictions travel through social media?

The named places and organisations are research settings, not claimed partners. The first work is to map the relationships properly, speak with the people involved and find where a shared movement already exists. Statstrade should enter as a tool for learning with the network, not as a machine arriving with the community pre-labelled.

10 Near-term pilot approach

An initial pilot should begin by identifying a hotspot, mapping the overlapping communities and interviewing people who occupy more than one role inside it. That work can reveal where attention already moves, which relationships carry trust, which systems must remain authoritative and where a public Campaign could add something useful.

The pilot can then select one shared movement, one anchor community and a small number of adjacent participants. The Campaign administrator and partner organisations can define the present state, shared future, initial Topic list, outcome sets, Campaign-scoped point budget, moderation boundaries, resolution sources and legitimate Tasks. The product should be adjusted with them rather than presented as finished.

Participants would receive a limited supply of non-cash points that exist only inside the pilot Campaign. They would stake those points on the administrator-published outcomes, explain the prediction and decide whether to share it publicly. The pilot can then observe whether the prediction travels into neighbouring communities, whether the explanation starts useful discussion and whether people return as evidence changes. The shared content may cross community boundaries; the points do not.

The purpose is to learn whether cross-community activity appears at all. Do administrator-published Topics travel beyond the anchor group? Do people recognise themselves in more than one part of the network? Does point staking make predictions more deliberate or merely more complicated? Do participants understand that joining another Campaign gives them a separate point context rather than access to their previous balance? Can a Campaign connect communities without erasing their differences? Where does the model create administrative, cultural, legal or privacy friction?

Early evaluation should use interviews, observed participation, partner feedback and a small number of carefully chosen measures. Fixed growth targets would be premature before the shape of a useful pilot is known. The first result should be a clearer account of where the model works, where it does not and what a second pilot should test differently.

Platform delivery and proof

The immediate objective is to build the platform well enough to run credible pilots with real communities and organisations. The work is to deliver the full participation loop—Campaigns, administrator-published Topics, Campaign-scoped supplied points, shared predictions, Tasks, resolution and visible history—then gather evidence about whether that loop creates richer and more useful online participation. The next case for growth should be earned through working software, partner learning and observed behaviour.

A future does not just happen. It gathers believers. The question the pilots must answer is whether the Prediction Network can help real communities gather theirs.

References

- [1] Chris Zheng, *Statstrade Motivation*, founder document, 2025.
- [2] Chris Zheng, *Bondi Medtech – Call to Action*, founder writing sample, 2024.
- [3] Statstrade, *Techstars Application*, presentation supplied 7 July 2026.
- [4] Reddit, *Annual Report*, current public filing. Used as a category reference for community-led discovery and advertising, not as a direct operating comparison.
- [5] Strava, public product and subscription materials, current at preparation. Used as a category reference for identity, activity, community, and premium consumer features.
- [6] Metaculus, public product documentation. Used as a category reference for resolved forecasting and calibration.
- [7] Manifold Markets, public product documentation. Used as a category reference for social prediction content and community-created markets.
- [8] Australian Government, *Interactive Gambling Act 2001*, current compilation. <https://www.legislation.gov.au/C2004A00851/latest/text>.