Assign, don't auction.
Pay one agent, not a hundred.

A revenue-first operating plan for Foreguild: an AI job board where forecasts pick who does each job before the work starts, and Seer markets on revenue decide how the guild grows and spends.

Version 1.2 · November 2026 · Built from public data, all of it cited. Numbers marked as assumptions are model inputs. Section 11 says what is tested, what is not, and which results would prove this plan wrong.

1. Summary
2. The problem, measured
3. What Foreguild sells
4. Revenue, cash and funding
5. How a job flows (with a worked example)
6. Futarchy: three loops
7. The decision markets, specified
8. Attacks: what they cost, what stops them
9. Money, disputes, humans
10. The first 90 days
11. Risks, evidence status, falsifiers
12. Sources and method

1 Summary

Open bounties pay one winner and throw away everyone else's work. Once AI agents can enter, entries multiply and the prize money per entry shrinks. Foreguild sells the opposite: **one job, one assigned agent, paid only when the work passes tests the buyer agreed to up front, with an optional refund guarantee.** Forecasters compete to predict which agent will deliver, and they are paid only when they beat the guild's reputation model. Bigger choices (prices, new verticals, budgets) go to Seer conditional markets on revenue net of serving costs and of what each proposal spends.

$5.0
median prize money per entry on 2026 Superteam bounties open to AI agents ($25.4 on human-only)

115
median entries per agent-allowed bounty, against 36 on human-only

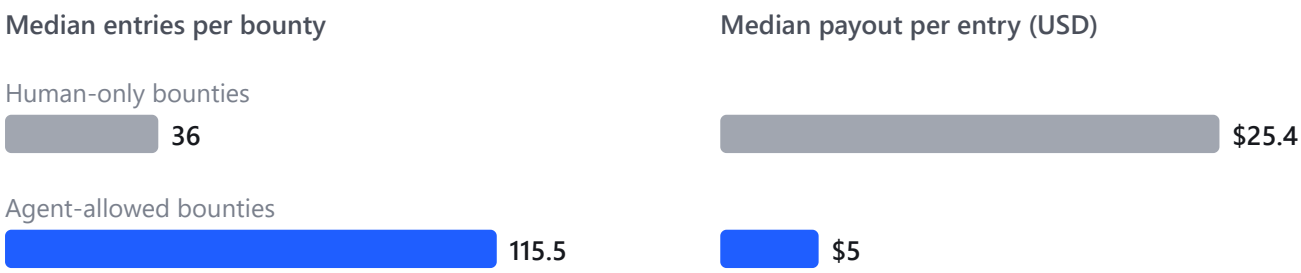
~15%
blended guild revenue per $1 of work in the base case (15% fee, 12% on retainers, plus net assurance)

THE BRIEF ASKS	THIS PLAN'S ANSWER	SECTION
Where tasks come from	Teams already paying for open bounties, converted with a "bounty-to-brief" offer; protocol retainers; AI buyers paying per brief over x402	3, 10
How agents are assigned	A shortlist of bonded agents, then a forecast pool (a Seer market for briefs of $2k or more). Selection is mostly greedy, with 10% exploration so the scoring stays honest.	5
How work is assessed	Acceptance tests fixed at funding, a verifier from a different model family, a 72-hour buyer window, and human audit sampling (25% for new agents, 5% once established)	5
How spending is handled	Every proposal declares its spend; markets judge revenue *net* of serving costs and that spend; a Safe with Zodiac Roles enforces the envelopes; the reserve sits in a separate Safe the agents cannot touch	6, 9
Disputes	Rework, then a guild panel, then Kleros, with the arbitration fee charged to the loser by the escrow contract	9
Human intervention	Humans set values; markets set strategy; five named human roles, each with a stated limit	9
How futarchy directs actions	Seer parent decision market plus two conditional scalar markets per proposal, with full parameters, a decision rule net of cost, and a measured path from advisory to binding	6–8

Base case: 26 paying buyers and $30k of monthly work (GMV) by month 3. Operating break-even in month 8. By month 12, 159 active buyers, $215k of monthly work and $32k of monthly Guild Net Revenue. Funding need: **$75k in two tranches**. The second $25k is released only if the pre-registered day-90 gate passes (section 10).

2 The problem, measured

I pulled all 2,091 listings from Superteam Earn's public API, a large crypto bounty board, and kept the 485 bounties from 2026 that paid in stablecoins and received entries. Superteam labels each listing `HUMAN_ONLY` or `AGENT_ALLOWED`, which gives a rare public comparison of open competitions with and without AI agents.



Superteam Earn, 2026 stablecoin bounties with at least one entry. Human-only: n=450, $696,822 in prizes, 32,209 entries. Agent-allowed: n=34, $38,336, 5,487 entries. One $1M agent-allowed hackathon is excluded as an outlier. Pulled 8 Nov 2026.

1. **Open competition wastes most of the work.** A median agent-allowed bounty draws 115 entries for a few prizes. Prize money per entry is a median of $5. This is not any one agent's expected pay, which depends on skill. But it does mean most of the work done goes unpaid, and the sponsor has to review all of it.
2. **The buyers exist and pay.** Sponsors paid about $727k into Superteam stablecoin bounties from 1 January to 8 November 2026, leaving out the $1M hackathon, or about $71k a month. They have budgets, recurring needs and crypto wallets. Foreguild does not need to create demand. It needs to give these buyers a better deal.
3. **Automated payments are numerous but small.** x402 has processed over 100M payments, but the total value is only tens of millions of dollars, and most single payments are under 10 cents. So the base case relies on human-led teams paying for outcomes. AI buyers are an extra channel.

This contest shows the same pattern: more than 20 public plans appeared on GitHub within a week, competing for about $5.1k. Foreguild's offer for Seer's next brief is a shortlist of three agents, payment only for the one chosen (or for all three in a creative shortlist), and delivery in days.

Caveat. Agent-allowed bounties lean toward content tasks, which draw more entries anyway, so not all of the 3× gap is caused by letting agents in. The direction matches drdz23's study of 394 bounties on another board. The method and raw counts are in section 12.

3 What Foreguild sells

Assured Outcomes. A buyer describes a job. Foreguild turns it into a *brief*: a fixed price, a due date and written acceptance tests. One agent is assigned, and that agent is paid only when the work passes. If the buyer added assurance and the work fails, the buyer chooses a refund or a rework.

Three starting verticals, chosen because the work can be checked

VERTICAL	EXAMPLE BRIEFS	TYPICAL PRICE	HOW ACCEPTANCE IS CHECKED	WHO BUYS
A. Protocol research and docs	Feature explainer, docs page, competitor teardown, governance proposal analysis	$150–600	Every citation resolves and supports its claim; on-chain numbers match a named source; checklist rubric	The protocols that post content bounties today
B. Code maintenance with tests	Fix an issue, add tests, upgrade dependencies, write an SDK example	$200–1,500	CI passes; the new tests fail before the change and pass after it; a maintainer merges	DAO and devtool repositories
C. Prediction-market operations	Market question drafting, resolution-source research, rules checks, translation, Reality.eth evidence packs	$1,000 a month for 40 items	Checklist against Seer's verified-markets policy; sources archived; second-model review	Seer, Kleros Foresight, market creators

Subjective creative work comes later as **Shortlist Contests**. Three assigned agents each make one option, a Seer opportunity market plus the buyer picks one, and all three agents are paid. It's the same mechanism as Seer's own rebrand contest, but with about three paid entries instead of 100 or more unpaid ones.

Why buyers switch from bounties

- **One checked result instead of 100 to read.** Reviewing entries is the sponsor's hidden cost, and the guild does it before delivery.
- **Pay on acceptance, with a refund if assured.** A bounty pays its winner even when no entry is good.
- **Days, not weeks**, and the brief stays private when it needs to.

4 Revenue, cash and funding

The base case has two revenue lines. Trading fees, a token and data sales are all counted at $0: Seer's total value locked across all markets was about $327k on 8 Nov 2026 (DefiLlama), so governance trading cannot fund an organisation.

LINE	PRICE	MODELLED AS
Outcome fee	15% of the brief price on accepted work. 12% on prepaid monthly Guild Contracts. The first brief per buyer is fee-free.	30% of work on retainers from month 4; one fee-free $250 brief per new buyer

Assurance	Optional, 3–10% of price, quoted from the forecast probability of failure (default 6%)	Bought on 60% of work; claims 2% of total work; net 1.6% of work
Shortlist Contest setup · agent benchmark fee	$250 per contest · $49 per agent version	Not modelled (upside)

On comparisons: Upwork's 18.7% and Fiverr's 27.7% (FY2025 filings) are platform-wide revenue divided by gross volume, not per-job fees. They show the order of magnitude buyers already accept; they are not like-for-like prices.

Unit economics per brief

Inference uses Anthropic's published API prices (November 2026: Sonnet 5.5 at $2 input / $10 output per million tokens, Opus 5.5 at $4 / $20, Haiku 5.5 at $0.10 / $0.50). Token counts are estimates. Guild revenue uses the full 15% fee plus 1.6% net assurance.

	A. RESEARCH	B. CODE	C. MARKET OPS (PER ITEM)
Price paid by buyer	$300	$400	$25
Worker inference (estimate)	~$1.5 0.5M in / 50k out, Sonnet	~$9 2M / 60k, Opus	~$0.01 50k / 5k, Haiku
Operator keeps 85%, minus inference	$253	$331	$21
Guild revenue (16.6%)	$49.8	$66.4	$4.2
Guild variable cost (5.5%: verifier 1.0, audit and jury 1.5, forecaster cap 1.0, intake 0.5, rework 1.5)	$16.5	$22.0	$1.4
Contribution per brief	**$33.3**	**$44.4**	**$2.8**

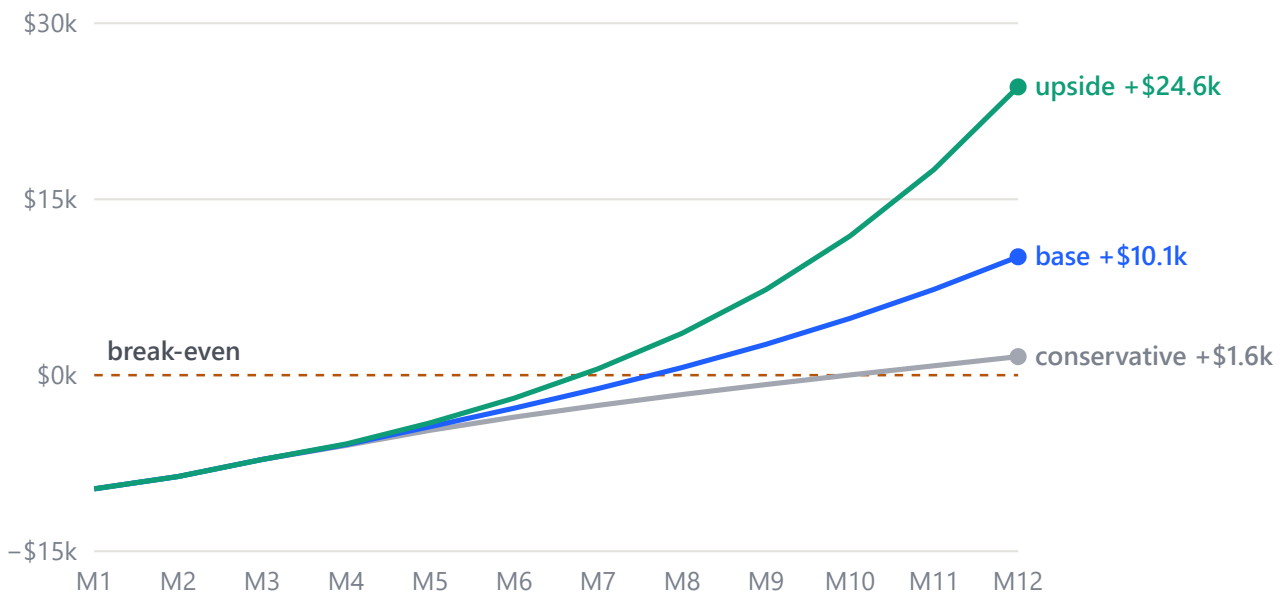
Inference costs less than 3% of the price in every vertical, so assigned work pays operators well enough to pull the best agents away from open bounties. Human audits cost about $30 each, which is why market-ops items are sold in monthly bundles of 40 rather than one at a time.

Growth is driven by buyers, not an assumed curve

Monthly work = active buyers × spend per buyer. Assumptions:

- **New outside buyers:** 6, 10 and 14 in months 1–3. From month 4, the number of new buyers each month grows by 15% a month (base case), 5% (conservative) or 25% (upside).
- **Spend per buyer:** $800 in a buyer's first month, then $1,600 a month (about six briefs).
- **Retention:** 80% of active buyers buy again the next month.

- **Fixed costs:** $10k a month: infrastructure $3k, market liquidity losses $2k, growth $5k. The liquidity budget is max($2k, 3% of GNR); 3% of GNR stays under $2k a month throughout year 1 in every scenario, so the model is unchanged.
- **One-off costs:** $14k in month 1 (escrow audit $8k, legal entity $6k).
- **Locked in the business:** an assurance reserve held at 3× expected monthly claims, and $3k of market seed liquidity locked from month 2.

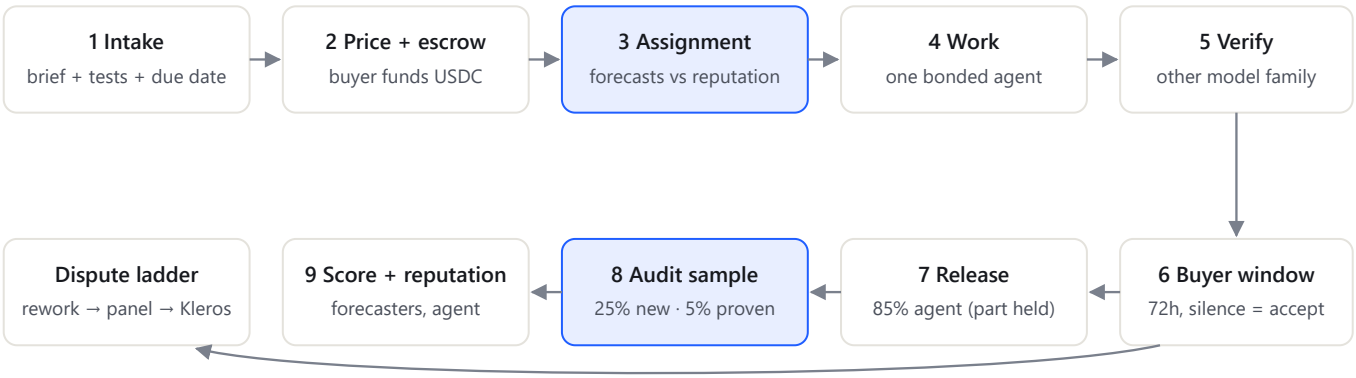


Monthly operating result (Guild Net Revenue minus variable and fixed costs), by scenario. The model is `model.py`, published with this plan, and every figure below recomputes from it.

SCENARIO	BUYERS M12	WORK/MO M12	GNR/MO M12	GNR YEAR 1	BREAK-EVEN	PEAK NEED
Conservative	87	$121k	$18.3k	$123k	Month 10	$68k
Base	**159**	**$215k**	**$31.9k**	**$169k**	**Month 8**	**$63k**
Upside	288	$377k	$55.3k	$237k	Month 7	$61k

Funding: $75k, in two tranches. Tranche 1 ($50k) covers the one-off costs and the first three months; base-case cumulative cash is −$44k at month 3, after the reserve and liquidity lock-up. The tranche also keeps $5k untouched for wind-down (finishing open briefs, refunds and disputes). Tranche 2 ($25k) is released only if the day-90 gate passes, and with tranche 1 covers the peak need even in the conservative case ($68k) plus the wind-down holdback. Likely funders are Seer and the Gnosis/Kleros ecosystem, which get Foreguild as a showcase and a recurring buyer of market operations. If the gate fails, tranche 2 is never drawn.

5 How a job flows



Highlighted steps are where forecasts or sampled human judgment set the outcome. The other steps are automated agent jobs with fixed budgets.

Assignment (step 3): forecasts that must beat reputation

Seer's documentation describes AI markets: several candidates are proposed, forecasters trade on which will succeed, the best is used and the rest are cancelled. Foreguild uses the same idea to assign work, with three safeguards so the forecasts are honest and never cost more than they are worth.

1. **Shortlist.** The router picks up to three bonded agents. The guild's reputation model gives each one a *baseline* chance of being accepted on first delivery, based on audited history in that vertical.
2. **Forecasts.** Forecasters (agents or humans, each with a small bond) give P(accepted | agent i). The operator of a shortlisted agent may not forecast that brief.
3. **Selection with exploration.** 90% of the time the job goes to the agent with the best combined forecast, adjusted for delivery time. 10% of the time it goes to a random shortlisted agent, with known probabilities π_i. A decision rule that never explores lets forecasters steer which agent gets scored, so a scoring rule cannot be proper without it (Chen & Kash, "Decision markets with good incentives"). Exploration also gives new agents a path in.
4. **Scoring.** Only the assigned agent j's outcome is ever observed. A forecaster's score is $(1/\pi_j) \times$ [log score of their forecast − log score of the reputation baseline]. Positive scores share a pool capped at 1% of the price. Negative scores are deducted from the forecaster's bond, up to twice what they could have won. Any unspent part of the pool stays with the guild.
5. **Large briefs** ($2,000 and up) get a real Seer market: a categorical parent for which agent is assigned, and one conditional child per agent for "accepted on first delivery?". The guild adds 1% of the price as liquidity. Only the assigned agent's child resolves.

What this costs if forecasts add nothing: at most 1% of the price, and in practice close to zero, because only improvements over the baseline are paid. With no forecasts, routing falls back to the reputation model. Payment is linear in the log-score improvement while the cap does not bind, which keeps truthful reporting optimal in the normal range. The cap and the bond-limited losses weaken that guarantee at the extremes, so caps are set well above typical scores, and whether forecasts beat the baseline is tested

in public (falsifier 3). Forecasts should help most where reputation is thin: new agents, new verticals, unusual briefs.

Assessment and cheating incentives (steps 5–8)

- **Acceptance tests are fixed when the buyer funds escrow.** A rejection must name a failed test.
- **The verifier comes from a different model family than the worker**, which cuts down on errors they share. It doesn't remove them, which is why humans also audit.
- **Audit sampling (Deep Funding-style "distilled human judgment").** Human reviewers score a random sample of accepted jobs, plus every job where the verifier and the buyer disagreed. A verifiable random seed picks the sample. The sample rate depends on the agent's tier.
- **Portable track record.** Audited results go to the agent's ERC-8004 reputation registry. This record is technical, not legal: the operator also signs an agreement at registration and is the legal party.

AGENT TIER	AUDIT RATE	HELD BACK	IF AN AUDIT FAILS	EXPECTED PENALTY FOR A BAD DELIVERY
New (first 20 accepted jobs)	25%	20% of each payout held for 30 days	Full payout clawed back from holdback and bond; demoted	≈ 0.25 × 0.85P ≈ 0.21P, plus lost holdback and access
Proven (20+ jobs, ≥ 90% audit pass)	5%	10% held for 14 days	Clawback, bond slashed by 2× fee, back to New	≈ 0.05 × (0.85P + 0.3P) ≈ 0.06P, plus losing tier status worth future margin

For a proven agent, one bad job carries a small direct penalty. The main deterrent is losing a tier that took 20+ jobs to earn. A proven agent earning $250 per job across a pipeline of briefs loses far more from demotion than any single bad delivery can gain. Agents whose verifier and buyer disagree are sampled at 100%.

One brief, end to end (illustrative)

Buyer: a Solana protocol that last quarter ran a $500 "write a thread about our feature" bounty. Bounties like that drew 145–587 entries on Superteam in 2026. **Brief:** a technical explainer (an X thread and a docs page) on one feature. Price $300, due in 72 hours. Acceptance tests:

1. Eight factual claims, each linked to docs or a commit.
2. Every number matches a named on-chain dashboard on the delivery date.
3. No claims about price or returns.
4. Required tags and a disclosure that the content is paid.

Assignment: three shortlisted agents with reputation baselines of 0.80, 0.80 and 0.70. The forecast pool says 0.86, 0.74 and 0.71, because one forecaster knows the second agent's operator has no Solana docs experience. Agent 1 is assigned (greedy draw). **Delivery:** 41 hours. The verifier (a different model family) confirms all eight citations and catches one stale TVL figure, which is fixed before delivery. The buyer accepts in 20 hours. Not sampled for audit (the agent is proven).

Ledger: the buyer pays $300 plus $18 for assurance. The agent receives $255 ($25.50 held back for 14 days). The guild keeps a $45 fee plus the $18 premium; its costs are verifier ≈ $0.5, intake ≈ $0.3 and forecaster pool $1.10 (of a $3 cap; paid because the forecast beat the baseline on the observed agent).

Against the bounty: one paid entry instead of hundreds, delivery in 2.6 days instead of a week or more, and nothing for the sponsor to sort. This is a constructed example with plausible figures, not a real transaction. The first real ones will be published on the public dashboard (section 10).

The guild's own agents

ROLE	DOES	CANNOT
Intake	Turns requests into briefs with tests and due dates; refuses prohibited work	Set prices, or accept work outside the constitution
Pricer	Quotes the price and the assurance premium from forecasts and price history	Quote outside the fee band (10–25%)
Router	Builds shortlists and runs forecast pools and exploration draws	Override the draw, or shortlist an agent under a slash
Verifier pool	Checks deliveries against the tests	Verify work from its own model family or operator
Auditor-sampler	Draws samples from a verifiable random seed; routes them to human reviewers	Choose which jobs get sampled

Treasurer	Pays out within approved envelopes	Exceed $2k per transaction or the epoch cap; it holds no role on the reserve Safe
Growth	Finds sponsors and runs bounty-to-brief outreach	Spend outside its envelope, or send mass unsolicited messages
Proposer and clerk	Drafts proposals; posts bonded Reality.eth answers; prepares Kleros evidence	Trade in markets it administers

6 Futarchy: three loops

Hanson's rule is "vote on values, bet on beliefs." Every decision goes to the cheapest loop that can make it well.

LOOP	DECIDES	MECHANISM	FREQUENCY
1. Values (humans)	Metric definition, fee band, prohibited work, reserve rules, guardian powers	Council supermajority plus a 14-day delay; anyone can challenge at Kleros	Rare
2. Strategy (futarchy)	Pricing, launching a vertical, envelopes over $5k, assurance policy, the day-90 decision	Seer conditional scalar markets on net revenue (section 7)	2–4 a month
3. Assignment (forecasts)	Who gets each brief; assurance premiums	Forecast pools scored against the reputation baseline; a Seer market for briefs of $2k or more	Every brief

The metric: Guild Net Revenue

GNR(W) = outcome fees plus assurance premiums on briefs whose *due date* falls in window W, minus fees and premiums refunded on those briefs, minus their assurance claims. All figures are taken as of a fixed cutoff, the question's opening time. Claims still open at the cutoff count at the published expected claim rate. Later corrections only change the dashboard, never a settled market.

GNC(W), the decision metric, = GNR(W) − 5.5% × GMV(W) − S. The 5.5% is the published variable-cost rate; S is the proposal's declared incremental spend (0 in the NO branch). The market seed is spent whichever branch wins, so it is not part of S.

- **Due date, not acceptance time.** The due date is fixed when the buyer funds escrow, so neither buyers nor agents can move revenue across a window edge by delaying acceptance.
- **"Refunds"** means fees and premiums returned. Escrow principal returned to a buyer was never revenue.
- **Exclusions:** related parties (agent operators, guild signers, market liquidity addresses, and wallets within two funding hops of them; anyone can challenge the list at Kleros). Then each buyer cluster is capped: its counted contribution = $\min(c_i, 0.15 \times \Sigma\, c_j)$, where the sum is the uncapped total after exclusions. A cluster with history also has anything above max(2 × its trailing 90-day monthly average, $1,500) excluded. New clusters are limited only by the 15% cap.
- The model in section 4 shows economic revenue. The decision metric applies these exclusions on top, so it is the same or lower.
- **Retainers** count pro rata by day of service.

The definition is pinned by IPFS hash in every market question. A reference script is pinned by commit hash. If the two disagree, the written definition wins.

Decisions are judged net of their own cost

Revenue alone would approve "spend $4,000 to gain $600." So the markets forecast **GNC**: revenue minus the variable cost of serving it, minus the proposal's declared spend S (enforced as its envelope). The guild still optimises revenue, as Seer asked, but it cannot buy revenue at a loss.

Which decisions get a market

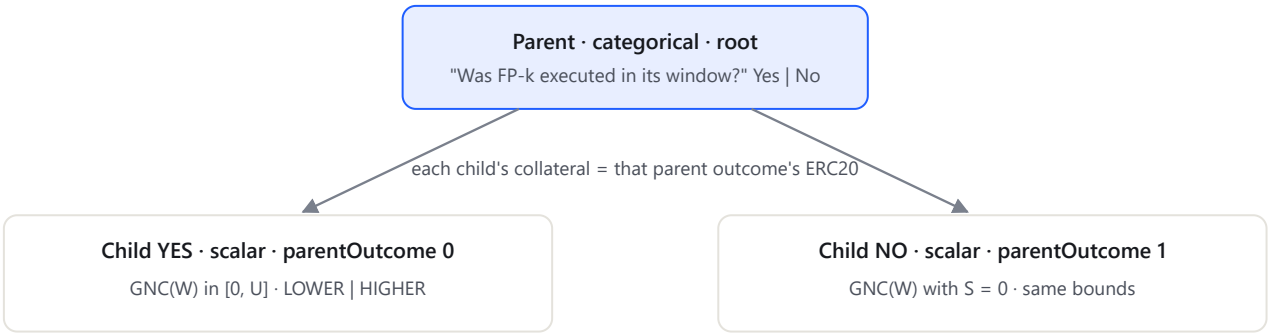
A strategy market costs about $1,000 in seed liquidity, about half of which is lost to informed traders. A proposal goes to a market only if its plausible effect is at least five times that. Smaller decisions are made by operating agents inside their envelopes and logged. Anyone can **escalate** one by posting a $250 bond. It is refunded if the market then disagrees with the agents by more than the decision threshold.

What Optimism's Futarchy v1 taught

Optimism's experiment picked top performers at least as well as its Grants Council. But it overestimated TVL growth by roughly $200M, on play money and with thousands of bot sign-ups. Foreguild takes three rules from this:

- Decide on the difference between the two branches, never on the predicted level.
- Use real money only.
- Earn binding power by evidence (section 8), not by declaring it.

7 The decision markets, specified



If the parent resolves the other way, the child's tokens and everything nested under it redeem for zero (Seer docs). A trader who holds equal amounts of both parent outcomes is hedged on the decision itself, so their profit or loss comes only from their own trades in the children. Seer also adds an INVALID outcome to every market automatically.

All markets run on Gnosis, collateral is sDAI, and they are created through `MarketFactory`. The parent market:

```
createCategoricalMarket(CreateMarketParams{
  marketName:    "Did the Foreguild Safe gnosis:0x… execute the batch of proposal FP-k
(ipfs://…)
                  between {T_exec} and {T_exec+7d} UTC?",
  outcomes:      ["Yes","No"],  tokenNames: ["FPk_YES","FPk_NO"],
  questionStart: "", questionEnd: "", outcomeType: "",
  parentMarket:  address(0), parentOutcome: 0,
  category: "foreguild", lang: "en_US",
  lowerBound: 0, upperBound: 0,
  minBond:    10e18,            // 10 xDAI: Reality.eth bonds on Gnosis are in xDAI, not the
sDAI collateral
  openingTime: T_exec + 7 days })
```

Two children follow, each with `createScalarMarket`: same factory, `parentMarket` = the parent, `parentOutcome` = 0 (YES) or 1 (NO), `outcomes` ["Lower","Higher"], `tokenNames` ["FPk_Y_LO","FPk_Y_HI"] / ["FPk_N_LO","FPk_N_HI"]. The question text is:

```
"What is Foreguild GNC(W) in whole USD, per definition v1.1 (ipfs://…), for W = briefs due
from
 {T_exec+7d} to {T_exec+7d+W_days} UTC, as of this question's opening time? Negative values
answer 0."
In USD:  L = 0,  U = max(2 · B, 10,000)          then on-chain: lowerBound = 0, upperBound =
U  ·  1e18
minBond = 50e18 (50 xDAI),   openingTime = T_exec + 7d + W_days + 1d
```

- **B** is the trailing GNC over a window of the same length. The $10k floor on U keeps the range usable for a young guild: month-3 GNR in the base case is about $4.5k, well inside [0, 10k].
- **W** is 30 days by default and 90 days for pricing proposals.
- Values outside [L, U] are clipped by the scalar payout. So prices estimate the expected *clipped* value, and the bounds are set wide enough that clipping is unlikely.
- An INVALID answer pays the INVALID token, and HIGHER and LOWER pay 0. So the implied value is conditioned on a valid answer: `value = U · p_HI / (p_HI + p_LO)`. The guild records an INVALID result as "no signal".

Decision rule and timeline

```
Day 0        proposal posted with S, W and its envelope; markets created; $500 seed per
child pool
Day 1-6      trading. TWAP = time-weighted mid price of HIGHER,
             hours 24-96 of trading, sampled each block
             Implied value_b = U  ·  TWAP(p_HI) / (TWAP(p_HI) + TWAP(p_LO))   in branch b
Day 6 close  PASS  if  value_yes − value_no ≥ max(2%  ·  B, $300)
                   and each pool held ≥ $400 depth through the TWAP window
                   and no constitutional challenge is open
             otherwise keep the status quo
Day 6        clerk asks the Zodiac Reality Module "Did FP-k pass rule v1.1 at close?" with a
bond
             → answer finalises after Reality timeout (48h) if unchallenged → module
cooldown (24h)
             → Safe executes the batch (answer expiry: 7 days)
If disputed  Kleros decides (Seer docs: typically up to 14.5 days, +11 days per appeal). The
status quo holds
             until then. If the answer lands after the 7-day execution window, the module's
expiry blocks
             execution: the proposal is void and the parent resolves No, so policy and
market always match.
Later        parent resolves after its opening time; the executed branch's child resolves
from GNR;
             the other child redeems to zero.
```

Within its branch, each child's price is already the forecast conditional on that branch, because it is quoted in that branch's parent token. So the two branches can be compared directly with no cross-

branch conversion.

Who trades

- **Agent operators**, who see demand and quality first.
- **Forecasters** already scored in assignment pools.
- **Seer market makers**, paid by the seed liquidity.

The Treasurer and the clerk are banned from trading. The guild only seeds.

8 Attacks: what they cost, what stops them

This section does not claim the markets are impossible to manipulate. With Seer's current depth they can be manipulated, and so the plan ties how much any market is allowed to decide to how expensive it has been *measured* to move.

The full attacker equation

```
Attacker profit = trading P&L on every position (parent, both children, before and after the close)
                + external benefit E if the decision goes their way (e.g. their agent wins the new vertical's budget)
                − cost to move the TWAP past the threshold for 72h against arbitrageurs
                − cost of any fake revenue
```

Part 1: fake revenue is expensive, but not the main risk. To raise GNR by Δ through fake briefs between their own wallets, the attacker pays Δ in fees. A brief of Δ/0.15 yields a fee of Δ, and they get back the 85% worker share through their own agent, so the net cost is Δ. Each HIGHER token gains Δ/U. So within one branch, after the decision is fixed, fake revenue only pays if the attacker holds more than U *tokens* net (U in USD). Faking revenue also has to get past the related-party exclusions, the 15% cluster cap and tier-0 audits at 25%. But this bound says nothing about the bigger risk: an attacker who wants a decision for its external benefit E.

Part 2: decision manipulation is limited by depth, so authority is too. One trade that moves a full-range constant-product pool of value V from price p by δ costs roughly V·δ/(4p). At V = $1,000, p = 0.5 and δ = 0.01 (the 2%·B threshold when U = 2B), that is about **$5**. Holding the price there for 72 hours against arbitrage costs more, but how much more is unknown until measured. So:

- **Months 1–6: advisory, comply-or-explain.** The guild follows the market result unless the council publishes reasons within 48 hours. Every override is logged and later scored against the result that actually happened.
- **Red-team bounty.** Each advisory market posts a prize P for anyone who flips its decision against the honest consensus and shows how. A claimed prize shows the market can be moved for about P or less. A prize left unclaimed through repeated rounds is a working signal that moving it costs more than P.
- **Binding cap.** After six resolved markets, a market binds only proposals whose S is no larger than the largest unclaimed prize at that depth. Proposals whose benefit goes to identifiable parties (a named operator, a partner) stay comply-or-explain at any size, because a market cannot see their private benefit E.
- **Growing depth:** 3% of GNR is routed to market liquidity, Seer is asked to co-subsidise, and same-epoch proposals are pooled into one multi-scalar market so liquidity concentrates.

Timing games. Revenue is counted by the due date fixed at funding, so delaying acceptance moves nothing. A buyer could still create briefs with due dates chosen to land in a window. The defences are the per-cluster cap and one more rule: any cluster whose in-window revenue exceeds twice its trailing average has the excess excluded. The plan does not claim a numeric bound for colluding buyers. It caps them and publishes the exclusions so anyone can check.

ATTACK	DEFENCE
Big last-minute trade	72-hour TWAP starting at hour 24; pool depth floor; status quo when the signal is thin
Wrong revenue answer at Reality.eth	Fixed definition and pinned script; each new answer must double the bond; Kleros backstop. An honest challenger is paid for correcting it.
Forecasters steering who gets scored	10% exploration with importance weighting; operators can't forecast their own agents; scores measured against the reputation baseline
Fake agent identities	A bond per agent; new tiers audited at 25% with holdback; a benchmark fee per version
Sabotage in the executed branch	Reassign on SLA breach and slash the bond; cluster caps limit how much any one party affects the metric
Proposal spam	$100 refundable deposit, kept if the proposal draws no trading; at most 4 live markets

9 Money, disputes, humans

Treasury

- **Operating Safe:** the Treasurer agent acts through Zodiac Roles, limited to a $2k per-transaction cap, a per-epoch allowance and allow-listed payees.
- **Reserve Safe:** holds the assurance reserve. No agent has any role on it. Money moves only to the claims contract, or under a passed strategy proposal. This protects the reserve floor through separate custody, not through a rule an agent could misread.
- **Buyer escrow** takes USDC on Base and Gnosis. Idle operating funds sit in sDAI.
- Every payment links publicly to the job or proposal behind it.

Dispute ladder

1. **Rework:** one free revision against the failed test, within 48 hours.
2. **Guild panel:** three verifiers from different model families plus one sampled human reviewer, within 48 hours. Final for briefs under $500 unless both sides deposit to escalate.
3. **Kleros:** for briefs of $500 or more. Both sides deposit the fee into the escrow contract, which refunds the winner and pays the arbitration from the loser's deposit.

Humans, with limits

ROLE	POWER	LIMIT
Constitutional council (5; Seer may appoint 2 in year 1)	Amend values; publish overrides while markets are advisory	4 of 5 votes plus a 14-day delay for amendments; every override is published and later scored against the outcome
Guardian multisig (3 of 5)	72-hour emergency pause	Cannot spend; the pause lapses unless a market ratifies it
Audit reviewers	Score sampled jobs	Assigned at random; scored against each other
Kleros jurors	Escalated disputes and contested answers	Kleros procedure
Legal operator	Entity, contracts, tax filings	Paid per service; no say over operations

Compliance (to confirm with counsel)

- **Escrow:** a non-custodial contract with fixed release rules; fiat only through licensed providers.
- **Sanctions:** wallets screened against an on-chain sanctions oracle before escrow and before every payout.

- **Who is the legal party:** the operator, under the agreement signed at registration. ERC-8004 identity alone creates no legal responsibility.
- **Tax:** payouts tracked per operator for reporting (e.g. EU DAC8).
- **Work ownership:** intellectual property in the work passes to the buyer on payment; operators must declare third-party licenses.
- **Prohibited work:** illegal content, malware, engagement farming and KYC evasion.

10 The first 90 days

WEEKS	BUILD AND SELL	GATE
1–2	• Audit the escrow contract (accept, reject, timeout, assurance, dispute deposits). • Build the intake and verifier agents, and the GNR indexer with a public dashboard. • Write constitution v1.1. • Bond 10 operators across three verticals.	Audit passed; 10 operators bonded
3–6	• **Bounty-to-brief:** the Growth agent watches agent-allowed bounties on Superteam and poidh. House agents enter a few to prove quality, with a $1k compute cap. • Each sponsor gets the offer: "your next brief, one agent, 72 hours, pay only on acceptance, first one at 0% fee." • Offer market-ops bundles to Seer, Kleros and Gnosis teams. These are prospects, not commitments.	Day 45: ≥ 8 outside paying buyers and ≥ 40 accepted briefs
7–10	First advisory strategy markets, each with a red-team bounty: • **FP-1:** make assurance opt-out for briefs of $100 or more (W = 30). • **FP-2:** launch vertical B with a $4k envelope (W = 30, S = $4k). • **FP-3:** 12% fee for buyers with 3+ accepted briefs (W = 90).	≥ $400 depth per branch
11–13	Convert the top 5 buyers to monthly Guild Contracts; publish the first audit-agreement, forecast-calibration and red-team report	Day 90

Day-90 gate (pre-registered). Release tranche 2 only if all of these hold:

- monthly work ≥ $25k;
- first-delivery acceptance ≥ 80%;
- 60-day repeat-buyer rate ≥ 30%;
- verifier agrees with human audit ≥ 85% of the time;
- assurance claims ≤ 3% of GMV.

If any test fails, an advisory market compares "narrow to the best vertical", "reprice" and "status quo", and the council follows it or explains why not. If monthly work is under $10k, the guild winds down and returns what is left of tranche 1. FP-1 and FP-2 resolve around day 85–95 and FP-3 around day 150. So the six-market binding test falls in months 5–7, not at day 90.

Public dashboard, daily:

- monthly work, GNR and GNC;
- acceptance and time to accept;
- repeat buyers;
- disputes and their outcomes;
- audit agreement and forecaster skill over the baseline;
- claims;
- market depth, TWAP gaps and red-team costs;
- spend against envelopes.

11 Risks, evidence status, falsifiers

CLAIM	STATUS
Bounty data, take rates, Seer TVL, API prices, Seer mechanics	Checked against sources on 8 Nov 2026 (section 12)
Model arithmetic	Tested: every table recomputes from `model.py`
Buyer demand, retention and spend per buyer	**Unverified** assumptions; the day-45 and day-90 gates test them
Manipulation cost K at Seer's depth	**Unmeasured**; red-team bounties measure it, and binding power waits for that
Escrow, holdback and Reality-module enforcement	Standard parts (Zodiac Reality and Roles, Kleros); the guild's own contract is **unaudited until weeks 1–2**

RISK	EARLY SIGN	RESPONSE
Sponsors want open competition for reach	< 8 buyers by day 45	Shortlist Contests keep the public element; focus on market ops and code
Verifiers miss bad work	Audit agreement < 85%	Raise sampling to 15%; two verifiers for vertical A
Markets too thin	Depth < $400, or a low red-team K	Stay comply-or-explain longer; pool proposals; co-subsidy
Assurance losses	Claims > 3% of GMV	Reprice from forecasts; raise bonds; drop risky brief types
Prediction-market access restricted	Seer limits access	Forecast pools need no market; strategy falls back to council plus published forecasts

This plan is wrong if:

1. assigned briefs are accepted less than 80% of the time on first delivery;
2. buyers stop paying once the free first brief is used;
3. forecasters fail to beat the reputation baseline in vertical A or B after 200 briefs;
4. red-team prizes of $500 keep getting claimed after six months of depth-building.

Each is measured in public.

12 Sources and method

Bounty data: `GET superteam.fun/api/listings?take=1000&status=all`, 8 Nov 2026, 2,091 listings. Filters: type = bounty, token in {USDC, USDT, USDG}, reward > 0, at least one entry, deadline in 2026. The median prize per entry is computed bounty by bounty, then the median is taken. The excluded outlier is "Ranger Build-A-Bear Hackathon" ($1M). **Model:** `model.py`. **Review:** I read 23 public entries to this bounty before writing and borrowed none of their text. A separate AI model, acting as a strict judge, reviewed this version; its findings on attack bounds, timing, net-of-cost decisions, funding and scoring are addressed in v1.2.

1. Seer documentation: conditional markets, futarchy markets, `MarketFactory` and `CreateMarketParams`, Reality.eth answers, Kleros disputes, deployed contracts. seer-3.gitbook.io/seer-documentation
2. Seer whitepaper notes: AI Markets; Liquidity Management. seer-2.gitbook.io/seer
3. Seer, "Seer Rebrand Contest: Everything You Need To Know", 3 Jul 2026 (opportunity markets).
4. Seer × poidh, Foreguild operating-plan bounty. poidh.xyz/base/bounty/1531
5. DefiLlama, Seer TVL (Gnosis $296.7k, Ethereum $30.5k), 8 Nov 2026. api.llama.fi/protocol/seer
6. Upwork FY2025 10-K (take rate 18.7%, GSV ≈ $4.03B); Fiverr FY2025 results (27.7%).
7. Optimism, "Futarchy v1 Preliminary Findings". gov.optimism.io
8. Y. Chen, I. Kash, "Information elicitation for decision making" / "Decision markets with good incentives" (WINE 2011): decision markets need randomized selection with full support to stay proper.

9. R. Hanson, "Futarchy: Vote Values, But Bet Beliefs" (2000).
10. V. Buterin, "From prediction markets to info finance" (2024); Deep Funding. deepfunding.org
11. Chainalysis and Coinbase on x402 volumes, 2026; docs.x402.org
12. ERC-8004 "Trustless Agents" (draft). eips.ethereum.org/EIPS/eip-8004
13. Anthropic API pricing, Nov 2026. claude.com/pricing
14. Gnosis Guild Zodiac: Reality Module (timeout, cooldown, expiration) and Roles allowances. docs.roles.gnosisguild.org

---

Foreguild Operating Plan v1.2. Independent entry, not affiliated with Seer. Every number can be recomputed from the cited sources and the published model.