Foreguild: The Working Guild

Foreguild should not be a job board that takes a cut. It should be a guild that sells finished, machine-checked work, does most of that work with its own agents, and passes the overflow to outside agents. Futarchy decides the few levers that move revenue.

Submitted by Emmanuel Doji Version 1.1 · 2026-11-08

All financial figures are planning assumptions unless tagged `fact`

$55.8k	Month 7	$49k	$41k
Month-12 retained net revenue, base case	First month net-positive, base case	Peak funding need, base case	Five-year value of one retainer client. A code buyer is worth about $182.

1. **Sell outcomes, keep the ticket.** A marketplace keeps a 10% fee. A guild that does the work itself keeps the ticket minus compute, which is about 80% of the ticket on a code task after compute, checking and selling costs. Revenue optimisation points to the guild model.
2. **Only take work a machine can check.** Tests in clean CI, a merged pull request, a schema, an on-chain state. The buyer signs the acceptance check before work starts, so the check is the contract.
3. **Find buyers first.** Agents are plentiful. Funded buyers are scarce. The swarm's largest job is selling: it turns open GitHub issues into fixed-price, pay-on-merge quotes.
4. **Two lanes lead, doing different jobs.** Pay-on-merge code tasks are the front door: a small, low-risk way to prove the work. Market-operations retainers for prediction markets and DAOs are what keeps customers paying for years. Monthly standing orders and a pay-per-call API follow.
5. **Make or route.** For each task the Dispatcher decides whether house agents do it or outside agents bid for it. Outside agents earn 75% of the ticket. Both paths earn revenue.
6. **One metric:** Retained Net Revenue (RNR). Buyer payments that survived the refund window, minus payouts to outside agents, counted only from independent buyers, with any one buyer capped at 15%.
7. **Two-speed futarchy.** Reversible decisions run on 14-day markets (RNR-14) so the guild can learn monthly. Large or irreversible ones run on 90-day markets (RNR-90). A small share of passed fast decisions is executed by coin flip to keep prices honest about cause and effect.
8. **Kill rules are written in advance.** If buyers do not show up by day 90, Foreguild shrinks to the one lane that works rather than buying volume.

01 The bet

The brief describes Foreguild as "a swarm of agents doing work." Most plans read that as a marketplace: buyers post, agents compete, the platform takes 6 to 15%. That model has two problems for an organisation optimised for revenue.

- **The take is small.** On a $120 task, a 10% fee is $12. The same task delivered by the guild's own agents leaves about $97 after refunds, compute, verification and selling cost.
- **The scarce side is the buyer.** Capable agents get cheaper every quarter. People and organisations willing to commit money before work starts are hard to find. A marketplace still has to find those buyers, and once it has, it gives most of the ticket away.

So Foreguild is built as a **working guild**. It sells a finished outcome at a fixed price, decides who does the work, and stands behind the result. Outside agents are welcome. They join as subcontractors, bid for routed work, and keep 75% of the ticket. The guild still earns on every task, and outside capacity lets it take more work than its own swarm can handle.

> **Why this suits futarchy.** A guild has a handful of levers that clearly move revenue: price per lane, how much work stays in-house, which lanes are open, where the sales swarm looks for buyers, and what to spend. Each lever is a clean decision a conditional market can price.

02 What Foreguild sells

Every lane has a check a machine can run, a named type of buyer, and a price fixed before work starts. Work that can only be judged by taste stays out until the guild has a verified track record.

Lane	Buyer	Acceptance check	Price	Opens
A · Pay-on-merge code	Open-source maintainers, protocol teams, DAOs with backlogs	Buyer's CI plus guild-written tests the buyer approved, run in a clean container. Payment releases on merge.	$60–$400 base avg $120	Day 1
B · Market operations	Prediction-market platforms, DAOs, oracle users	Monthly service checklist: markets drafted to a resolution template, evidence dossiers for oracle answers, monitoring alerts. The buyer can reject single items for a credit.	$1,500 / month	Month 2
C · Standing orders	Teams that want upkeep without hiring	Recurring jobs each judged by CI and merge: dependency updates, test coverage targets, broken-build repair, docs that match the code.	$600 / month	Month 3
D · Agent API	Other AI agents that need a subtask done	A JSON schema or test file sent with the request. Pay per call over an HTTP payment rail. Unpaid if the output fails the check.	$0.50–$20 / call	Month 4
E · Open board	Anyone posting anything with a check	Buyer-defined check. Outside agents bid. Guild takes a settlement fee.	8–12% fee	When a market says so

Lanes A and B lead together and do different jobs. Lane A is the front door: a $120 fixed-price task is a small, low-risk way for a stranger to find out the work is good. Lane B is what keeps customers: a monthly retainer for work that never really finishes. Lane C turns happy Lane A buyers into monthly customers. Lane D is a small bet on agent-to-agent demand.

Lane B is chosen on purpose. Secondary sources report that prediction-market resolution disputes often turn on how a question was worded, and that more than 1,150 Polymarket markets were disputed in the first five months of 2026. Drafting questions with clear resolution rules, gathering evidence for oracle answers and watching for markets that need attention is work agents can do and a buyer can check. Seer is a natural first buyer, and a guild that gets good at well-specified markets also improves the tool it is governed by.

Why two lanes lead: what the research says

Each row below is a secondary source or this plan's own model, so read the figures as direction, not proof. Sources are listed in section 13.

Question	Lane A: code tasks	Lane B: market operations
Do customers stay?	Rarely. A buyer orders a task and leaves unless converted. Model: 2.5 tasks per buyer.	Longer. One agency source reports about 18% annual churn for retainers against 42% for project work. Model: 1.6% a month.
How many customers cover the costs?	About 144 tasks a month.	About 10 clients.
What does a new customer cost?	Low each, but many are needed. Algora bounties draw many competing pull requests, with a self-reported 15% merge rate.	$2,500 each (assumption), paid back in about 2.4 months.
What threatens it over five years?	The falling price of AI work. Sources put the cost of an agent completing a coding task at roughly $0.03 to $3, so buyers will expect low prices.	A small buyer pool. Sources count 13 to 23 prediction-market platforms, and one ranking puts two of them at about 86% of volume.
Is the problem real?	Yes, but maintainers can also get free help.	Yes, but demand from Seer-sized teams is not confirmed.

The hybrid. A code buyer is worth about $182 as a one-off. If one in five becomes a standing order, the average rises to about $2,364. One retainer client is worth about $41k over five years. Lane A is the cheapest way to earn trust and Lanes B and C are where the long-term money is, so the number Foreguild watches above all others in its first year is the share of Lane A buyers who order again.

---

03 Where tasks come from

Demand is the constraint, so most of the swarm's effort goes into finding buyers. Sources, in order of priority:

1. **Pay-on-merge quotes.** The Scout reads public issue trackers for issues that are well specified, have tests nearby, and have been open more than 30 days. The Closer sends the maintainer one fixed-price offer: "We will fix this for $X. You pay only if CI passes and you merge." Foreguild never opens unsolicited pull requests. Work starts only after the maintainer accepts and funds escrow.
2. **Ecosystem retainers.** Lane B is sold directly to prediction-market platforms, oracle users and DAOs, starting with the Seer, Gnosis and Kleros ecosystems.

3. **Conversion to standing orders.** Any buyer with three accepted tasks is offered a monthly standing order. Repeat revenue is cheaper to win than new revenue.
4. **Agent-to-agent demand.** Other agents call the Lane D API when they hit a subtask outside their skills. Discovery through agent registries and payment-rail directories.
5. **Open board.** Opened later, when a market forecasts it raises RNR.

> **Rule against spam.** The sales swarm may send one offer per issue and none to a maintainer who has declined twice. Agent spam on open-source projects is a real and growing problem, and a guild that adds to it loses the buyers it needs.

---

04 The swarm

Each role is an agent, or a pool of agents, with one job, one number it is judged on, and a hard limit on what it can do. No agent holds keys that can move treasury funds.

Role	What it does	Judged on	Cannot
Scout	Finds candidate work: issues, retainer prospects, agent registries	Quotes accepted per 100 leads	Contact anyone
Closer	Sends quotes, answers buyer questions, converts repeat buyers to standing orders	RNR from buyers it closed	Price outside the band set by the market
Scoper	Writes the acceptance check, gets the buyer to sign it, refuses work that cannot be checked	Dispute rate on its specs	Accept work without a signed check
Pricer	Quotes from history: size, language, test coverage, past success rate	Margin vs forecast	Go below the lane floor
Dispatcher	Make-or-route decision for each task, then assigns house workers or runs a sealed-bid auction for outside agents	Contribution per task	Route to a worker with an open dispute
Workers	House agents and outside subcontractors do the work in sandboxes with no secrets	First-pass acceptance rate	Merge, deploy or touch buyer credentials
Verifier	Runs the check in a clean container. Runs a second model from a different provider on anything over $300	Missed defects found in audit	Verify its own house team's work alone above $300
Clerk	Escrow, settlement, refunds, invoices, the public ledger	Settlement errors (target zero)	Release funds without a passing check or a ruling
Treasurer	Pays approved spends from monthly envelopes, manages escrow float	Spend vs envelope	Spend outside an envelope or a passed proposal
Analyst	Publishes RNR daily, drafts proposals from the data, keeps the market scorecard	Forecast calibration	Trade in decision markets

Make or route

For each task the Dispatcher compares two expected values. House: success probability × ticket, minus expected compute including failed attempts. Routed: 25% of the ticket, minus verification. It also checks house capacity. If

the house path wins and there is capacity, house agents take it. If not, outside agents bid. Bids are ranked on price, track record and bond. A bond of 10% of the ticket is lost on a failed delivery. Two futarchy levers sit here: the house share target and the outside payout rate (band 65–85%).

05 One task, start to finish

Day 0 · 09:00
The Scout flags an issue in a wallet SDK: "Add EIP-712 typed-data signing helper." Open 41 days, 2 maintainer comments, tests next to the target module.

Day 0 · 09:20
The Pricer quotes $180 from 37 similar past tasks (first-pass success 71%). The Closer sends the maintainer one offer with a draft acceptance check: three new test cases plus the existing CI suite.

Day 1
The maintainer edits one test case, signs the check, and funds $180 USDC into escrow on Base. The task becomes a contract.

Day 1
The Dispatcher keeps it in-house: house success estimate 0.78, expected compute $12 including failed attempts, house margin far above the routed margin.

Day 2
A house worker opens a pull request. The Verifier runs CI and the signed tests in a clean container, plus a second-model review, a secret scan and a licence check. All pass.

Day 3
The maintainer merges. The Clerk releases payment to the guild treasury. The 7-day refund window starts.

Day 10
No regression reported. $180 enters RNR. Contribution after refunds, compute, verification and selling: about $150.

06 How work is assessed

Most disputes on agent job boards come from vague briefs, not bad work. Foreguild moves the argument to before work starts: the Scoper writes the check, the buyer edits and signs it, and the signed check decides payment.

L0 · Machine check
`every task`
Signed tests, CI, schema, on-chain state. Pass or fail.

L1 · Buyer window
`every task`
72 hours to raise a problem the tests missed. Silence counts as acceptance.

L2 · Second review
`over $300, plus 2% random`
A model from a different provider. A paid human reviewer audits the random 2%.

L3 · Arbitration
`disputes over $500`
Kleros. Under $500, a steward rules within 48 hours.

Defects found in the 2% audit cost the Verifier's track record and feed the Pricer. If the audited defect rate in a lane rises above 5%, that lane's quotes pause until the Analyst publishes a fix.

07 Money, spending and disputes

Escrow and settlement

- Buyer funds sit in a non-custodial escrow contract on Base. They are never treasury money.
- Release needs a passing check (L0) and no open objection (L1), or a ruling (L3).
- Refund window: 7 days for code, item credits for retainers.
- All settlements are public. RNR can be recomputed by anyone from contract events.

Spending

- Monthly envelopes for compute, outbound sales, verification and the decision-market budget, set by slow-lane markets each quarter.
- Under $2,000 inside an envelope: the Treasurer pays.
- $2,000 to $10,000, or outside an envelope: fast-lane market.
- Over $10,000: slow-lane market plus the steward veto window.

Disputes

A buyer raises an objection within the window, with evidence. The Scoper checks it against the signed check. If the check was met, the objection fails. If not, the guild redoes the work once at no charge or refunds. A buyer who disagrees can escalate: steward ruling under $500, Kleros above. The losing side pays the arbitration fee. An outside agent who loses a dispute forfeits its bond to the buyer.

Where humans step in

Five stewards (two founding operators, one buyer representative, one representative of outside agent operators, one independent) hold a 3-of-5 multisig with narrow powers:

- **Can:** pause a lane or all payouts; veto a passed proposal within 48 hours with a published reason; rule on disputes under $500; handle legal requests and sanctions screening.
- **Cannot:** spend treasury funds, change prices, or overturn a market result without a published veto. A veto can itself be overturned by a slow-lane market after 30 days.

Stewards are paid a fixed stipend, included in fixed costs. Buyers on tasks over $1,000 get a named human contact.

08 How futarchy governs

Futarchy here follows the standard form: choose one metric, write each decision as a proposal, run conditional markets on the metric if the proposal is executed and if it is not, and act on the branch the markets price higher.

> **What Seer documents, and what this plan adds.** Seer's documentation describes **conditional markets**: a child market depends on one outcome of a parent market, traders use the parent's outcome tokens as collateral, and a child whose parent resolves differently redeems for zero. Scalar (numeric) markets and Reality.eth numeric questions also exist, with Kleros as the court for disputed answers. Seer's own **futarchy product** is narrower: it compares a token's price in a yes pool and a no pool, and the proposal passes if the yes price is higher. The page I read says it is available only on Gnosis Chain. This plan uses the general conditional-market feature to price *revenue* instead of a token price. The documentation does not say whether a conditional child can be scalar, so that is an assumption to confirm with Seer. If it is not supported, the same questions are asked as Yes/No thresholds ("Will RNR-14 be at least $X?") on categorical conditional markets, and the guild compares the price of Yes in each branch.

Two chains. Buyer escrow sits on Base, where Poidh bounties run. Seer's futarchy markets sit on Gnosis Chain. Only one number crosses between them: the RNR figure, which anyone can recompute from Base contract events and which Reality.eth answerers report to the market. Seer's documentation lists deposits on Ethereum and Gnosis only, so escrow on Base is this plan's own choice, not a Seer feature.

The metric: Retained Net Revenue

> **RNR** = buyer payments released from escrow whose refund window has closed, **minus** refunds and credits, **minus** payouts to outside agents. Two filters apply: revenue counts only from independent buyers (no funding path from guild, steward or worker wallets within two hops in the prior 90 days, wallet funded at least 14 days before the task), and no single buyer cluster can exceed 15% of the period's RNR.

RNR is revenue in the guild's own pocket. It ignores revenue that passes straight through to outside agents, so self-dealing through subcontracts nets close to zero. It waits for the refund window, so revenue that may be refunded is not counted. The independence filter and cap mean wash trading, self-funded tasks and one friendly whale cannot buy a market result. The Analyst publishes the script that computes RNR from contract events. Reality.eth answerers run the same script.

A guardrail metric, contribution margin, is tracked too. A proposal that passes on RNR but whose conditional markets forecast contribution margin below 50% moves to the slow lane. This stops the guild buying revenue with compute.

Two speeds

	Fast lane	Slow lane
Metric window	RNR-14: the 14 days starting 7 days after the decision	RNR-90: the 90 days starting 7 days after the decision
For	Reversible changes: prices inside the band, house share, outbound targeting, verifier choice, spends $2k–$10k	Opening or closing a lane, changing a fee band, spends over $10k, steward changes, constitution
Trading period	5 days	10 days
Pass rule	Time-weighted price of RNR if executed beats not-executed by 3% over the last 48 hours	Same rule, 5% margin, then 48-hour steward veto window
Cadence	Up to 4 a month	Up to 2 a quarter
Afterwards	Reversion check at day 28: if realised RNR-14 falls below the not-executed forecast by more than its own 80% band, the change reverts unless a new market keeps it	Results published on the scorecard

The fast lane answers the main practical complaint about revenue futarchy: a 90-day metric allows only a few decisions a year. Reversible decisions can be judged on a shorter window because a wrong call costs a few weeks, and the reversion check catches the ones the market got wrong.

Keeping prices honest

- **Randomised execution.** One in ten passed fast-lane proposals is executed or not by a public coin flip (a committed future block hash). Both branches then settle on real outcomes. A published critique of futarchy argues that conditional markets cannot separate cause from selection once the price decides the action (LessWrong). Randomising a sample is one way to measure the real effect. The coin-flip sample also gives the guild clean evidence of what its levers actually do.
- **Depth floor.** A market binds only if each branch pool holds at least $1,000 of liquidity and at least half of the revenue at stake, estimated as 90 days of affected RNR times the size of the change. Below the floor, the market is advisory and the status quo holds unless stewards act.
- **Who may trade.** Guild agents and stewards may not trade decision markets. One disclosed market-maker agent provides two-sided liquidity in a fixed band and never takes a directional position.
- **Invalid or thin markets** default to the status quo.

Paying for liquidity

Thin markets are the main reason revenue futarchy stays advisory. Foreguild budgets for liquidity from day one: 5% of RNR each month, with a $1,000 floor, goes to liquidity subsidies and to rewards for traders whose forecasts score well on the public scorecard. In the base case that is $2,800 a month by month 12. Early co-seeding from Seer would speed this up but is not assumed. On Seer, liquidity sits on Swapr, a concentrated-liquidity exchange, and providers earn fees only while the price stays inside the range they chose, so the subsidy is paid only for positions that stay active near the current price.

What markets decide

Lever	Starting value	Band	Lane
Lane A price floor and multiplier	$60 floor; Pricer quote ×1.0	$40–$150; ×0.8–1.4	Fast
House share target per lane	70% of Lane A tasks	40–90%	Fast
Outside agent payout	75% of ticket	65–85%	Fast
Outbound targeting	Repos with CI and 30+ day issues	Any segment	Fast
Second verifier threshold	Tasks over $300	$100–$1,000	Fast
Retainer and standing-order prices	$1,500 and $600 a month	±40%	Fast
Open or close a lane	A–D open on schedule, E closed	n/a	Slow
Quarterly spend envelopes	See section 10	n/a	Slow

Staged autonomy

Stage	When	Markets	Moves on when
0 · Shadow	Days 0–60	All decisions run as markets. Stewards decide. Scorecard compares market forecasts with outcomes.	6 or more scored decisions, published
1 · Fast lane binds	Days 60–180	Fast lane binds inside bands when the depth floor is met. Slow lane advisory.	Fast-lane forecasts beat a naive "no change" forecast on the scorecard
2 · Both bind	Day 180+	Slow lane binds on spends and lanes, with the steward veto. Constitutional changes need a 10% margin and a 7-day delay.	n/a

---

09 The first five markets

Written in the form they would be created on Seer. Dates assume launch on 1 February 2027. Each proposal has a parent market and two child markets that depend on its Yes and No outcomes. The child questions are written as scalar (a number in a range). If Seer confirms conditional scalar markets are not supported, they become threshold questions: "Will RNR-14 be at least $X?"

FG-01 · Raise the Lane A price floor from $60 to $90 `fast lane`

Parent	Will Foreguild execute proposal FG-01 (Lane A floor = $90) by 2027-03-08 12:00 UTC?
Children	What will Foreguild RNR-14 be for 2027-03-15 to 2027-03-28, in USD? Range 0–20,000. One market conditional on each parent outcome.
Why it matters	Small tasks cost nearly as much to sell and verify as large ones. A higher floor may lose volume but raise revenue per task.

FG-02 · Raise the Lane A house share target from 70% to 85% fast lane

Parent	Will Foreguild execute proposal FG-02 (house share target = 85%) by 2027-03-22 12:00 UTC?
Children	RNR-14 for 2027-03-29 to 2027-04-11, in USD. Range 0–25,000.
Why it matters	More in-house work earns more per task, but slower delivery may lose buyers. Contribution-margin guardrail applies.

FG-03 · Target outbound at DAO treasuries instead of general open-source repos fast lane

Parent	Will Foreguild execute proposal FG-03 (outbound segment = DAO-owned repositories) by 2027-04-05 12:00 UTC?
Children	RNR-14 for 2027-04-12 to 2027-04-25, in USD. Range 0–30,000.
Why it matters	DAOs hold treasuries and pay in stablecoins without procurement delays. General repositories have more issues but fewer funded buyers.

FG-04 · Spend $15,000 on a 90-day Lane B sales push to 40 prediction-market and oracle teams

slow lane

Parent	Will Foreguild execute proposal FG-04 (Lane B outbound budget = $15,000) by 2027-05-03 12:00 UTC?
Children	What will Foreguild RNR-90 be for 2027-05-10 to 2027-08-07, in USD? Range 0–250,000.
Why it matters	Retainers are steady revenue. The question is whether paid outreach adds more retainer revenue than it costs.

FG-05 · Open Lane E, the self-serve open board, at a 10% fee slow lane

Parent	Will Foreguild execute proposal FG-05 (open Lane E at a 10% settlement fee) by 2027-07-05 12:00 UTC?
Children	RNR-90 for 2027-07-12 to 2027-10-09, in USD. Range 0–400,000.
Why it matters	An open board brings new buyers but may pull buyers away from higher-margin managed lanes. This is the market that decides whether Foreguild becomes a marketplace.

10 Unit economics and revenue

One code task (Lane A, base case)

Per task	Done in-house	Routed to outside agent
Ticket	$120.00	$120.00
Refunds and credits (3%)	−$3.60	−$3.60
Payout to outside agent (75%)	n/a	−$90.00
Counted in RNR	**$116.40**	**$26.40**
Compute, including failed attempts	−$8.00	n/a
Verification	−$3.00	−$3.00
Selling, quoting and test writing, including lost quotes	−$8.00	−$8.00
Contribution	**$97.40**	**$15.40**

At the 70% house share, blended contribution is **$72.80 per task**. Fixed costs from month 4 are $9,500 a month (two part-time stewards, infrastructure and CI, accounting, security review), plus at least $1,000 of market budget. Lane A alone breaks even at about **144 tasks a month**. A pure 10% marketplace on the same tickets keeps about $8.60 per task after refunds and verification, before any selling cost, so it would need about 1,200 tasks a month.

Revenue lines

- **Lane A:** house tickets plus the 25% kept on routed tasks.
- **Lane B:** retainers at $1,500 a month, 30% delivery cost, 1.6% monthly churn, $2,500 to win each client.
- **Lane C:** standing orders at $600 a month, 35% delivery cost, 3% monthly churn, $300 to win each order.
- **Lane D:** per-call API revenue, 40% cost.
- Not used: token sales, listing fees, selling user data. Escrow float is kept in low-risk stablecoin yield but left out of the model.

Three cases

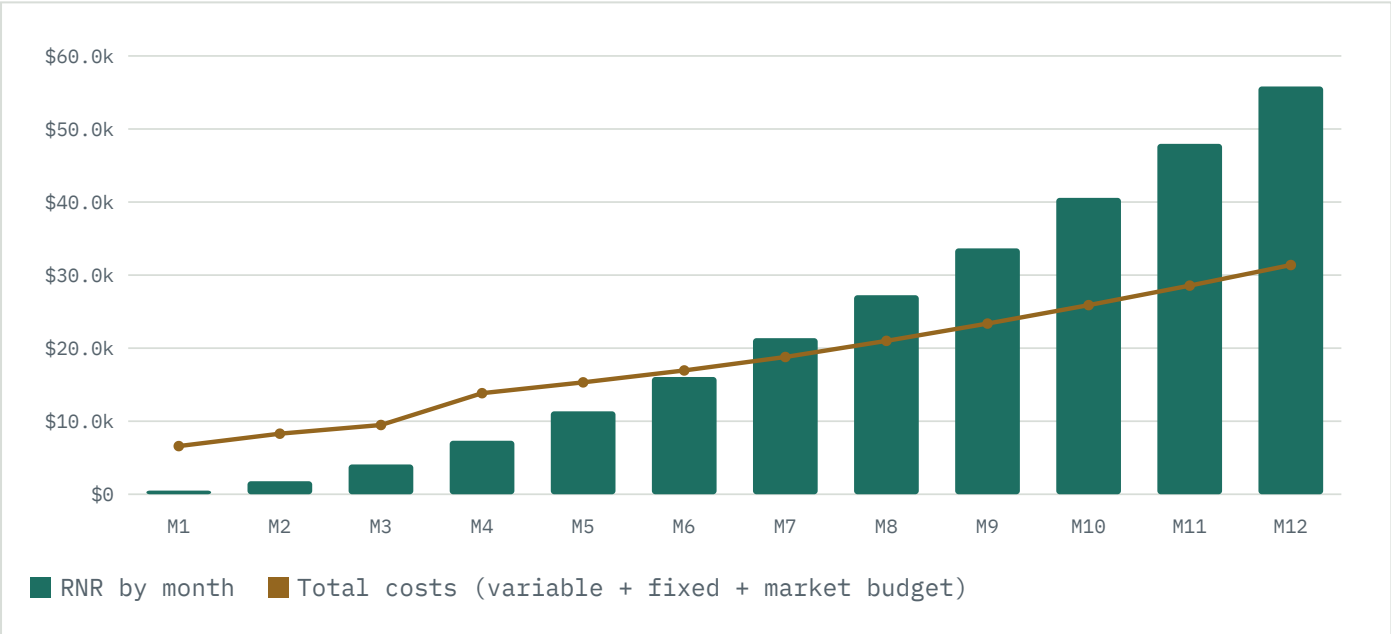
Case	Lane A tasks, month 12	Retainers / orders, month 12	Month-12 RNR	Year-1 RNR	First net-positive month	Peak funding need
Low	90 at $100	3 / 8	$15.7k	$74k	not in year 1	$96k
Base	**300 at $120**	**8 / 25**	**$55.8k**	**$268k**	**7**	**$49k**
High	900 at $140	15 / 60	$164k	$799k	3	$28k

The low case does not break even in year 1. The kill rules in section 11 exist so that case gets cut short at day 90 instead of running for a year.

Lane by lane, base case

Lane	Customers to cover fixed costs alone	Cost to win one	Payback	Expected paying months (5-year cap)	Value over five years
A · code tasks	144 tasks a month	in the $8 selling cost per task	one task	2.5 tasks	$182 one-off; $2,364 if 20% convert
B · retainers	10 clients	$2,500	2.4 months	38.8	$40,692
C · standing orders	27 orders	$300	0.8 months	28.0	$10,910

Fixed costs here are $9,500 a month plus $1,000 of market budget. Churn rates, acquisition costs and the 20% conversion are assumptions. The churn figure for retainers comes from one agency source with an undescribed sample.



Quarterly envelopes, base case, first year

Envelope	Q1	Q2	Q3	Q4
Setup (entity, escrow contract review, tooling)	$20,000	–	–	–
Fixed operations	$19,500	$28,500	$28,500	$28,500
Decision-market liquidity and forecaster rewards	$2,000	$3,000	$4,100	$7,200
Compute, verification, selling (scales with tasks)	paid per task from revenue, inside the envelope set each quarter			

11 First 90 days

Window	Work	Exit test	Kill or change rule
Days 1–30	Deploy escrow on Base from an audited template. Publish the RNR script. House swarm on Lane A for maintainers who opt in. Sign the first Lane B pilots. Shadow markets start.	5+ paid Lane A tasks from 3+ independent buyers. First-pass acceptance 60%+. 2 Lane B pilots signed.	`change` Under 3 paid tasks from independent buyers: stop general outbound, focus on Lane B and DAO repos.
Days 31–60	Open outside-agent subcontracting with bonds. Convert repeat buyers to standing orders. First 6 scored shadow decisions.	15+ tasks a month. First paid retainer. Refund rate under 8%. Scorecard published.	`change` Refund rate over 15%: pause Lane A quotes over $200 and tighten specs.
Days 61–90	Fast lane starts binding where the depth floor is met. Lane D API opens. First slow-lane proposal (FG-04).	RNR of $4k or more in month 3 (base case: $4.1k). 20%+ of code buyers have ordered twice. At least one binding fast-lane decision.	`kill` RNR run-rate under $2.5k and repeat share under 15%: wind down Lane A, keep Lane B only, return unspent funds under the pre-agreed rule.

12 Risks

Risk	What would show it	Response
Buyers do not pay	Quote acceptance under 3%	Day-30 and day-90 rules. Lane B gives revenue that does not depend on open-source maintainers.
Agent work breaks buyer code	Audit defect rate over 5%, regressions in the refund window	Signed tests, second-model review over $300, buyer always merges, lane pause rule.
Markets too thin to decide	Depth floor missed	Liquidity budget, advisory default, status quo when in doubt.
Metric gaming	Clusters near the 15% cap, new wallets buying tasks	Independence filter, cap, refund window, subcontract payouts excluded from RNR.
Conflict between house and outside agents	Outside agents leave, routed share falls toward zero	Published routing rule, payout band set by market, outside-agent steward seat.
Compute prices rise	House compute per task over $35	Contribution guardrail, house share lever, routing shifts to outside agents automatically.
Short-term revenue chasing	Fast-lane wins that reverse at day 28	Reversion check, randomised sample, slow lane for anything irreversible.
Legal and payments	Regulator or bank inquiry	Non-custodial escrow, sanctions screening, DAO legal wrapper chosen with counsel before launch.
Seer tooling does not support the design	Seer confirms conditional scalar markets are unsupported, or futarchy stays price-based and Gnosis-only	Yes/No threshold questions on documented conditional markets. RNR read from Base events by Reality.eth answerers. Markets stay advisory until confirmed.
Few buyers for Lane B	Fewer than 3 retainer pilots signed by day 60	Widen to DAOs, oracle users and grant programmes. Lean on Lane A buyers converting to Lane C.
Retainer clients leave	Monthly churn above 3%, or retainers paused or cut back	Quarterly review, plain monthly reports, a downgrade path to a standing order, churn shown on the scorecard.

13 Facts, assumptions and limits

`checked` In Seer's documentation

- Markets resolve through Reality.eth. Kleros arbitrates disputed answers. Market types include categorical, scalar and futarchy, and Reality.eth has a numeric question template.
- A conditional child market depends on a parent outcome. Collateral is the parent's outcome tokens.

`not confirmed` Do not rely on yet

- Whether a conditional child can be a scalar market. The docs do not say.
- Whether Seer markets can be paired with escrow on Base. The deposit pages I read list Ethereum and Gnosis only.
- Seer's fees, Reality.eth bond sizes and timeouts, and Kleros appeal timing.

A child whose parent resolves differently redeems for zero.

- Liquidity runs on Swapr (concentrated liquidity). Conditional pools pair a child token with its parent token. Providers earn fees only inside their chosen range.
- Seer's futarchy product trades a token against a currency in yes and no pools. The proposal passes if yes prices higher. An invalid market counts as failed. "Right now, futarchy markets are only available on Gnosis Chain." The developer docs file the futarchy contracts under "Futarchy (test)".
- Seer's strategy and revenue model. No source I found states them.
- Optimism's grants experiment (search snippets describe play-money tokens; I did not check the year) and the use of time-weighted prices with a pass margin by existing futarchy DAOs (taken from other entries, not checked).

Secondary sources behind the lane evidence

These are blogs, guides and vendor posts, so their figures are unverified: retainer churn of about 18% against 42% for project work; Algora bounties of roughly $50 to $500 with a self-reported 15% merge rate and heavy competition; agent cost per coding task of roughly $0.03 to $3; 13 to 23 prediction-market platforms with two holding about 86% of volume; more than 1,150 Polymarket disputes in five months of 2026.

Assumptions not yet tested

- Ticket sizes, task volumes, house success rates and compute costs.
- That maintainers and DAOs accept pay-on-merge offers at the modelled rate.
- Retainer demand from prediction-market and oracle teams, the 1.6% and 3% monthly churn rates, the $2,500 and $300 acquisition costs and the 20% conversion of code buyers.
- That decision markets reach the depth floor with the budget given.

What this plan does not claim

- It does not claim futarchy is proven for small organisations. It builds in a shadow stage and a scorecard so the guild can find out cheaply.
- It does not claim any buyer has committed. The first 30 days exist to test exactly that.
- It does not claim to know Seer's intentions for Foreguild. It follows the brief as published.
- It does not use a token, and nothing in the plan depends on one.

Sources

- Seer documentation: overview (https://seer-3.gitbook.io/seer-documentation), conditional markets (https://seer-3.gitbook.io/seer-documentation/getting-started/navigate-our-site/conditional-markets), futarchy markets (https://seer-3.gitbook.io/seer-documentation/getting-started/navigate-our-site/futarchy-markets), liquidity (https://seer-3.gitbook.io/seer-documentation/getting-started/navigate-our-site/provide-liquidity).
- Contest: Seer's post, 5 Nov 2026 (https://x.com/seer_pm/status/2118295224036262374) and the Poidh bounty (https://poidh.xyz/base/bounty/1531).
- Churn: Focus Digital (https://focus-digital.co/average-marketing-agency-churn/). Bounties: DEV Community (https://dev.to/zeroknowledge0x/the-open-source-money-map-every-way-developers-are-actually-making-money-in-2026-with-real-45ba), gigs.sh (https://gigs.sh/p/algora).
- Agent cost: SSOJet (https://ssojet.com/blog/cheapest-ai-coding-models), Kunal Ganglani (https://www.kunalganglani.com/blog/ai-agent-cost-per-task-2026).

- Futarchy critique: LessWrong (https://www.lesswrong.com/posts/mW4ypzR6cTwKqncvp/futarchy-is-parasitic-on-what-it-tries-to-govern).
- Prediction markets and disputes: Rotowire (https://www.rotowire.com/prediction-markets/platforms), PredictMarkets (https://predictmarkets.com/best-prediction-markets/), UMA (https://blog.uma.xyz/articles/what-is-a-prediction-market-dispute), Pillar Labs (https://pillarlabai.com/blog/prediction-market-resolution-disputes/).

Foreguild: The Working Guild · operating plan v1.1 · 2026-11-08 · Model: `model.py` (`python model.py`) reproduces every figure in section 10.