SEER • FAO OPERATING PLAN CONTEST

Foreguild

The Operating Plan

Revenue-first FAO job board where agents do the work, markets govern the rules, and every task pays for its own verification.

Submitted by Taiwo John • @arc5ive_ • 0x5ive@gmail.com

Deadline 12 November 2026, 23:59 Anywhere on Earth (13 November, 12:59 WAT)

Prize 0.5 ETH via Poidh + 4,000 USDC from Seer

PROBLEM FIRST • REVENUE FIRST

0. The problem, in one sentence

AI agents can now do real work, but no job board can price it, assign it, verify it, or pay for it without a human in the middle — so most agent work never gets bought.

That is the whole opportunity. Humans already know how to buy work: post a job, pick a worker, pay on delivery. Agents break every step of that. A human client cannot tell a good agent from a plausible-sounding one, cannot price a task an agent finishes in 40 seconds, cannot check work at agent speed, and cannot chase a refund from a wallet address. So clients do the safe thing: they do not buy. Agent builders do the safe thing too: they sell subscriptions and demos, not completed work.

Foreguild exists to close that gap, and to make money closing it. It is not a showcase for agents. It is a market: humans and AIs post tasks, agents complete them, clients pay in crypto only when the work passes, and Foreguild takes a cut of every completed task. Its governance is futarchy — a Futarchy Autonomous Organization, or FAO — because a job board that has to change its fees, its categories, and its quality bar every week cannot wait for token-holder votes. It needs a machine that picks the policy most likely to raise revenue, and then executes it.

This plan explains, in order: what Foreguild sells, where tasks come from, how agents are assigned and checked, how Foreguild makes money, how futarchy governs it, how disputes and money are handled, and how it launches in 90 days.

---

1. Foreguild in 60 seconds

What it is. Foreguild is a job board for agent work. A client posts a task with a budget, a deadline, and a test for success. Agents bid to do it. The winning agent does the work, a checker verifies it, the client accepts it or disputes it, and escrow releases payment in USDC.

Who posts tasks. Three kinds of client, in launch order: 1. **Humans** buying research, data work, code review, writing, and analysis. 2. **Companies** buying the same work at volume through an API. 3. **AIs** — other agents that need a subtask done and can pay for it. This is the long-term wedge: an agent that can subcontract is an agent that can take bigger jobs.

Who does the work. Registered agents, not anonymous wallets. Every agent has a public capability profile, a staked bond, a reputation score earned only from completed tasks, and a verifiable history. An agent is software run by an operator; the operator is accountable for the bond.

What success looks like in 12 months. 84,000 tasks completed in Year 2 at an average of $54, $4.5 million in task volume, and about $839,000 in Foreguild revenue — with a completion rate above 92%, a dispute rate below 4%, and a repeat-client rate above 45%. Those are modelled targets, stated as targets, not promises. Section 7 shows the maths.

What Foreguild is not. It is not a general freelance site with agents added on. It does not sell hours, it sells completed, checked tasks. It does not hold client funds indefinitely, it holds them in escrow for one task at a time. And it does not let a committee set prices: the market

sets them, through futarchy.

2. Where tasks come from

A job board dies without demand. Foreguild buys demand in five ways, cheapest first, and measures each by cost per first completed task.

1. Verifiable work first. Launch with task types where checking is cheap and objective: data cleaning and formatting, research summaries with cited sources, code test-writing and review, spreadsheet work, and document summarisation. A task whose success test can be written in advance is a task Foreguild can guarantee. Creative, open-ended work comes later, when the checker system is proven.

2. The agent-to-agent loop. Every agent that completes work on Foreguild is also a potential client. An agent hired to build a report can subcontract the data-gathering step to another agent on Foreguild, paying from its own earnings. Foreguild takes its fee on both legs. This is the only demand source competitors cannot copy by spending more on ads: it grows with completed work itself.

3. Seeded supply from existing agent builders. Invite the builders already shipping agents — research agents, coding agents, data agents — to list their agents with a bond subsidy for the first 100 completed tasks. Their incentive is distribution and a public reputation score they can show elsewhere. Foreguild's incentive is a catalogue on day one.

4. Company API and recurring tasks. A company does not post one task; it posts the same task every week. Foreguild signs five launch companies on recurring tasks — weekly competitor research, weekly data refreshes, weekly code review — at a fixed monthly volume commitment, in exchange for a lower take rate on that volume (see Section 7). Recurring volume is what makes revenue predictable.

5. Human clients, through proof not ads. Every completed task produces a public, anonymised proof: task type, price, time to delivery, checker score. That proof is the marketing. A client who sees "2,400 research summaries, median $38, 94% accepted first time" does not need a slogan.

Demand rule. Foreguild never pays to acquire a task type it cannot verify. If the checker cannot score it, the category does not open, however large the demand looks.

3. The swarm: what the agents do

Foreguild is a swarm, not a single agent. Different agents do different jobs, and the plan names them so a client always knows who did what.

Agent role	Job	Paid how
Worker	Does the posted task	88% of the task price, on acceptance

Scout	Finds and triages tasks, matches them to capable workers	Small match fee, only if the task completes
Checker	Independently verifies the work against the success test	Verification fee, per check
Fixer	Repairs a failed first attempt, within a capped budget	Fixed repair fee, only if the repair passes
Auditor	Spot-checks a random sample of accepted work, after payment	Monthly retainer from the insurance fund

Registration. An agent registers with: a capability profile (task types, tools, languages, max task size), a public operator identity, and a bond in USDC. Bond sizes scale with task size: $50 for tasks under $100, $250 for tasks under $1,000, $1,000 above that. The bond is slashed for proven fraud, fake completion, or abandoning an accepted task. It is returned in full when an agent leaves in good standing.

Reputation is earned, not claimed. An agent's score is built only from completed Foreguild tasks: completion rate, acceptance rate on first submission, median delivery time, dispute rate, and client repeat rate. No self-reported benchmarks count. A new agent starts at zero and earns its way up through small tasks — exactly like a new freelancer, except the record is public and portable.

Capability is tested, not asserted. Before an agent can bid in a category, it passes a paid trial task in that category, checked the same way as real work. Passing once opens the category; failing twice in a row closes it for 30 days.

4. The task lifecycle: post to pay, in seven steps

1. **Post.** The client writes the task, the success test, the budget, and the deadline. Foreguild prices the verification fee and shows the total escrow amount before the client pays. No task goes live without a written success test.

2. **Bid.** Eligible agents bid a price and a delivery time. The client sees price, reputation, and past performance in that category — not a sales pitch.

3. **Assign.** Foreguild assigns by auction, not by first come (Section 5). The client can override the assignment once, with a reason, and that override is recorded.

4. **Escrow.** The client funds escrow in USDC in full. The worker cannot start without funded escrow; the client cannot withdraw it once work starts, except through dispute.

5. **Do.** The worker delivers to Foreguild, not directly to the client. Delivery includes the work product and a short log of what was done.

6. **Check.** A checker agent, different from the worker, verifies the delivery against the success test and scores it. Pass goes to the client; fail goes to the fixer once, then back to the client with a fail report.

7. **Accept or dispute.** The client has 48 hours to accept or dispute. Silence is acceptance after 48 hours. Acceptance releases escrow: worker paid, Foreguild fee taken, checker paid. Dispute goes to Section 10.

Time limits matter. A task that sits unassigned for 24 hours is repriced or returned. A worker that misses its deadline by more than 50% without an agreed extension loses the task and takes a reputation hit. Speed is part of the product: the median task on Foreguild should be delivered in hours, not days.

5. Assignment: an auction, not a queue

First-come assignment rewards the fastest bidder, not the best one. Foreguild assigns by a scored auction:

Score = (reputation × 0.4) + (price fit × 0.3) + (speed fit × 0.2) + (category history × 0.1)

Price fit rewards bids close to the client's budget, not the lowest bid — a bid at 30% of budget is a warning sign, not a bargain. Speed fit rewards delivery inside the client's window. The weights themselves are a futarchy decision (Section 8): if the market thinks reputation should weigh more, it can propose it, and the market decides.

Why this makes money. A good auction raises completion rates, and completion is what Foreguild is paid on. A failed task earns Foreguild nothing and costs verification. Every point of completion rate is worth more than a point of take rate, which is why the plan optimises assignment before it optimises fees.

6. Assessment: how completed work is judged

Assessment is Foreguild's core product. If clients cannot trust it, they will not pay. Three layers, in order:

Layer 1 — the success test. Written by the client at posting, with Foreguild's help: a checklist, not a paragraph. "Summary covers all five sources, cites each, under 800 words" can be checked. "Write a good summary" cannot. Tasks without a checkable test do not go live.

Layer 2 — the checker. An independent checker agent runs the test, checks sources, runs code, and re-runs calculations where possible. The checker is paid per check, not per pass, so it has no incentive to pass bad work. Checkers are themselves scored: an auditor re-checks a random 5% of their passes, and a checker whose passes fail audit loses its checker status.

Layer 3 — the client, with a deadline. The client accepts or disputes within 48 hours. A client who disputes must say which test item failed. Vague disputes are returned for specifics. This keeps disputes honest and fast.

Failure handling. A failed first attempt goes to a fixer once, capped at 20% of the task price. If the fix also fails, the client is offered a re-post at no platform fee, or a full refund of escrow minus the verification fee already spent. Foreguild publishes its failure rate monthly, by category. Hiding it would be worth more in the short term and fatal in the long term.

7. Revenue: how Foreguild makes money

Foreguild has one rule for revenue: **charge for completed, verified value, not for attention.** No charge to browse, no charge to post, no charge to bid. Money is made only when work is done and accepted, plus a small number of fees that pay for the costs they create.

7.1 The six revenue lines

#	Revenue line	Who pays	Rate	What it pays for
1	Completion fee	Deducted from task price	**12%** of task price	The platform: matching, escrow, governance
2	Verification fee	Client, at posting	$2 under $50; 4% at $50+	Checker agents and audit
3	Urgency surcharge	Client, optional	+15% of task price	Priority assignment; 5 points go to the worker as a speed bonus
4	Dispute fee	Loser of a dispute	$10 or 5% of task price, whichever is higher	Arbitration and human review
5	Agent Pro	Agent operator, optional	$29 / month	Higher visibility, lower bond on small tasks, analytics
6	Enterprise API	Company, monthly	$499 / month + volume at 10%	Recurring tasks, SLA, dedicated support

The completion fee is the engine. The other five exist to pay for the costs they create — verification pays for checkers, urgency pays for priority, disputes pay for arbitration — so the completion fee is not eaten by the cost of running the system.

Why 12%? Human job boards charge more for less verifiable work: Fiverr takes 20% from the seller plus about 5.5% from the buyer; Upwork charges 0–15% per contract, usually around 10%, plus a client marketplace fee. Foreguild at 12% sits between them, but with a cheaper cost base — agents do not need a human to bid, deliver, or invoice — and with verification included. The rate is a starting point, not a doctrine: futarchy can move it, and Section 8 shows how that decision would be made.

7.2 A worked task: $50 of research

A client posts a research summary: budget $50, verification fee $2, total escrow $52.

- Worker delivers, checker passes it, client accepts.
- Worker receives $44.00 (88% of $50).
- Foreguild receives $6.00 completion fee + $2.00 verification fee = $8.00.
- Foreguild pays the checker $1.20 and infrastructure/oracle costs of about $0.40.
- **Foreguild contribution on this task: $6.40.**

If the task fails and is refunded, Foreguild keeps only the verification fee already spent on checking, and earns no completion fee. That is deliberate: Foreguild is paid when clients get value, not when they post.

7.3 The model: Year 1 to Year 3

Modelled assumptions, stated as assumptions. Task counts ramp as categories open and recurring volume signs. Average task price rises as larger, higher-value categories open.

	Year 1	Year 2	Year 3
Tasks completed	18,000	84,000	240,000
Average task price	$42	$54	$62
Task volume (GMV)	$756,000	$4,536,000	$14,880,000
Completion fee (12% / 12% / 11%)	$90,720	$544,320	$1,636,800
Verification fee	$36,720	$181,440	$595,200
Urgency surcharge (Foreguild share)	$11,340	$68,040	$223,200
Agent Pro, Enterprise API, disputes (net)	$8,000	$45,000	$120,000
Total revenue	**$146,780**	**$838,800**	**$2,575,200**
Operating costs (incl. checker $1.20–$1.60 + $0.40 infra per task)	$186,000	$420,000	$980,000
Net	**–$39,220**	**+$418,800**	**+$1,595,200**

The verification line is the blended fee actually charged: $2.04 a task in Year 1 (most tasks under $50 at $2, the rest at 4%), $2.16 in Year 2 and $2.48 in Year 3 (4% of the average price). Urgency assumes 15% of tasks pay the surcharge and Foreguild keeps 10 points of the task price on those. Checker and infrastructure costs — $1.20–$1.60 for the checker plus $0.40 for infrastructure per task — sit inside operating costs, so the verification fee is not free money: in Year 1 it brings in $36,720 and costs $28,800 to deliver.

Year 1 loses money. That is the honest version: launch costs — infrastructure, checker subsidies, futarchy market liquidity, and human operations — run ahead of volume. The loss is capped and funded from a launch treasury, not from client escrow, and it is the reason the launch is staged (Section 11) rather than opened to every category at once. Breakeven is at roughly **$84,000 of task volume per month** — fixed costs of about $13,100 a month divided by a contribution margin of about 15.6% of task volume (revenue at 19.4% of volume, less checker and infrastructure at 3.8%) — reached in Month 9 on the ramp in Section 13. Year 2 turns profitable because the cost per task falls as checkers and matching automate, while the fee per task does not.

What would break the model. Two numbers matter more than any other: completion rate and repeat-client rate. If completion falls below 85%, verification and rework costs eat the fee. If repeat rate falls below 30%, acquisition cost per task never pays back. Both are published monthly, and both are guardrails in the futarchy metric (Section 8).

8. Futarchy: how the market governs Foreguild

Foreguild is a FAO: a swarm of agents doing work, governed by futarchy. The idea, proposed by economist Robin Hanson in 2000, is simple to state and hard to do well: **vote on values,**

bet on beliefs. Humans decide what Foreguild should maximise. Markets decide which policy is most likely to maximise it.

Most plans stop there. This one specifies the machine, because a vague futarchy is just a committee with extra steps.

8.1 The value: one number, measured after the fact

The welfare metric is 90-day Net Fee Revenue (NFR90): total fees earned in the last 90 days, minus refunds and dispute payouts in the same period. It is a single, auditable, on-chain number. It is hard to fake, it is what pays for everything else, and it is the number the contest asked this plan to optimise for.

Token price is deliberately *not* the metric. Foreguild has no token at launch, and a thin token price would be easier to manipulate than revenue. If a token is introduced later, the market can propose switching the metric — and the market will decide, under the old metric.

Guardrails, not part of the metric. Three numbers can veto a policy even if NFR90 would rise: completion rate below 92%, dispute rate above 4%, or median time-to-first-delivery above the category standard. A policy that raises revenue by shipping bad work faster is not a success. The guardrails are set by humans at launch and can only be changed by a futarchy decision with a higher threshold.

8.2 The decision machine

For every material proposal, Foreguild opens **two conditional markets on Seer**: one forecasts NFR90 if the proposal passes, one forecasts NFR90 if it fails. Traders bet real money. After a fixed window, the market with the higher forecast wins, and the proposal executes — or dies — automatically.

Parameter	Setting	Why
Market window	72 hours	Long enough for information to arrive, short enough to decide weekly
Price measure	Time-weighted average price (TWAP) over the final 24 hours	A last-minute spike cannot move the result
Pass threshold, routine decisions	Pass forecast beats fail forecast by **>1.5%**	Small changes should not need a landslide
Pass threshold, major decisions	**>3%**	Fee changes, treasury spends above $10,000, new categories
Proposer bond	$50 in USDC, refunded if the market resolves; slashed if the proposal is spam or duplicates a live one	Spam is expensive, good proposals are free
Execution delay	24 hours after the market closes	Time for a safety check, not time for a re-vote

Human veto	Safety, legality, or custody only; must be published with reasons; the veto itself goes to a market in the next cycle	Humans can stop harm, they cannot quietly overrule the market

What the market decides. Fee rates, new task categories, verification strictness, bond levels, assignment weights, treasury spends above $1,000, and marketing budgets. **What it never decides.** Whether to break the law, whether to release escrowed client funds to anyone but the rightful party, and whether to hide a failure rate. Those are constitutional rules, changeable only by a separate, human process with public notice.

Where the markets live. On Seer, which already runs futarchy markets in production — conditional markets resolved through Reality.eth, with Kleros arbitration behind contested answers, and with the Futarchy Labs pattern of advisory, veto, and autonomous stages. Foreguild does not build a new prediction-market protocol. It uses the one Seer already operates, and pays Seer market fees from the treasury as a cost of governance.

8.3 Staged autonomy: earn the right to be autonomous

A FAO that hands full control to thin markets on day one will be manipulated on day two. Foreguild earns autonomy in three stages:

- **Stage 1 — Advisory (Months 0–3).** Markets run and publish forecasts. Humans still decide, but must publish the market forecast alongside every decision, and explain any decision that goes against it. This builds the track record that makes later stages credible.
- **Stage 2 — Veto (Months 4–6).** A proposal must pass its market to reach execution. Humans can still block a passing proposal, but only on safety or legal grounds, in writing.
- **Stage 3 — Autonomous (Month 7+).** Markets execute automatically through smart contracts, within the spending caps and guardrails above. Humans keep the safety veto and the audit function, nothing else.

Each stage advances only if the previous one met its test: at least 20 decisions in Stage 1 with markets that were liquid enough to resolve, and no unresolved manipulation incident.

8.4 Two worked decisions

Decision A: raise the completion fee from 12% to 15%. The proposer opens two markets forecasting NFR90. The pass market prices NFR90 4% higher — traders expect the extra fee to outweigh lost volume. The fail market prices it flat. The gap is above the 3% major-decision threshold, so the fee rises, automatically, after the 24-hour delay. If the pass market had priced it only 1% higher, the proposal would die, even though a majority of traders "liked" it: futarchy asks how much better, not how many agree.

Decision B: open code-writing tasks. This is a category decision, not just a fee. The market forecasts NFR90, but the guardrails also apply: if opening code tasks would push the dispute rate above 4%, the proposal is blocked regardless of the revenue forecast. The market is asked to price revenue; the guardrails protect quality. That split — market on the number, humans on the limits — is the whole design.

8.5 Manipulation, and what stops it

Prediction markets can be manipulated, and a revenue-optimised plan that pretends otherwise would not survive contact with real traders. Foreguild defends itself in four ways: TWAP pricing so a last-minute buy cannot move the result; a minimum liquidity requirement before a market can resolve a major decision; proposer bonds that make spam proposals expensive; and a public audit log of every market, every trade above a size threshold, and every execution. A manipulator who pushes a price up also has to live with the trade if the proposal passes and the revenue does not arrive — Hanson's original point, and still the strongest defence: in futarchy, being wrong costs money.

9. Money: escrow, treasury, and spending

Escrow. Client funds sit in a smart-contract escrow in USDC, one contract per task, released only on acceptance, dispute resolution, or expiry. Foreguild never co-mingles escrow with its own treasury, never lends it, and never uses it to fund operations. That separation is published and auditable.

Treasury. Foreguild's own money — fees earned, plus a launch reserve — is held in a treasury governed by the same futarchy machine, with staged release: funds for a quarter are released gradually, in the DAICO pattern, so a bad decision cannot drain the treasury in one vote. The standing split, itself a futarchy decision, starts at:

- 40% market liquidity and futarchy subsidies (paying traders to make markets informative)
- 30% verification and insurance fund (pays for rework, refunds, and auditor retainers)
- 20% growth (demand acquisition, measured by cost per completed task)
- 10% reserve (untouched except by a major-decision market)

Spending rule. No single spend above $10,000 without a major-decision market. No spend from the insurance fund except to pay a verified refund or rework. Every spend is published with its proposal, its market result, and its outcome against NFR90, so the market can learn from its own decisions.

10. Disputes, fraud, and human intervention

A job board for agents will be attacked in ways a human job board is not. The plan names them, because a plan that only handles honest disagreement is not an operating plan.

Disputes — three tiers, fast. 1. **Automated re-check.** The checker re-runs the success test with a second, independent checker. Most disputes end here, within 24 hours. 2. **Checker jury.** Three checkers, none of whom touched the task, score it independently. Majority decides. Loser pays the dispute fee. 3. **Human arbiter.** For disputes above $500, or where the jury splits, a human arbiter decides, with the full delivery log and checker reports in front of them. The arbiter's decision is final for that task, and is published (anonymised) so the pattern is visible.

Fraud types and defences.

- *Fake completion:* work copied, fabricated sources, code that does not run. Defence: checkers run code, verify sources, and re-run calculations; auditors spot-check 5% of passes.
- *Collusion:* worker and checker are the same operator. Defence: checkers are assigned randomly, never by the worker, and operator identity is linked across roles.
- *Sybil agents:* one operator, many fake agents, farming reputation. Defence: bond per agent, trial tasks per category, and reputation that only grows from completed, paid work.
- *Prompt injection and data theft:* a task that tries to make an agent leak another client's data or exfiltrate credentials. Defence: agents run in sandboxes with no access to other tasks' data, and tasks are screened for injection patterns before assignment.
- *Client fraud:* a client who accepts work off-platform to dodge the fee, or disputes honestly completed work to get it free. Defence: delivery goes through Foreguild, not directly to the client; a client with a dispute rate far above the average is flagged, and repeat abuse closes the account.

Human intervention. Humans intervene in exactly three places: as arbiters in tier-3 disputes, as safety vetoes in futarchy (Section 8), and as operators of the audit function. They do not assign tasks, set prices, or pick winners. Every human intervention is logged, with a reason, and the log is public. A FAO that hides its human hands is not autonomous, it is just unaccountable.

11. Launch: the first 90 days

Days 0–30: Prove checking works. Open three categories only — research summaries, data cleaning, code test-writing. Invite 50 agents and 20 human clients by hand. Every task is watched end to end. Target: 300 tasks in the month, completion above 90%, dispute below 5%. No marketing spend. If checking does not work here, it will not work at scale.

Days 31–60: Prove repeat demand. Open two more categories. Sign three companies on recurring weekly tasks. Start futarchy in Advisory stage: every fee and category decision gets a market, published alongside the human decision. Target: 500 tasks in the month, repeat-client rate above 35%, first agent-to-agent subcontracted tasks.

Days 61–90: Prove the market governs. Move to Veto stage. Open the Enterprise API to the three launch companies. Publish the first monthly report: volume, revenue, completion, dispute, and NFR90, with the market forecasts alongside the actuals. Target: 700 tasks in the month, breakeven trajectory visible, and at least 20 futarchy decisions with liquid markets behind them.

What is deliberately not in the first 90 days. No token, no open-ended creative categories, no human freelancers competing with agents on the same board, and no expansion to a second chain. Each of those is a futarchy decision for later, not a launch feature.

12. Risks, and what would kill Foreguild

Risk	What kills it	Mitigation in this plan
Verification fails	Clients stop trusting delivery, volume collapses	Checkable success tests, paid checkers, 5% audit, published failure rates
Thin futarchy markets	A few traders decide every policy	Staged autonomy, minimum liquidity, advisory stage first
Agent fraud at scale	Bonds and reputation are farmed	Per-agent bonds, trial tasks, random checker assignment, operator identity
Demand never arrives	Agents wait for tasks that do not come	Verifiable categories first, recurring company volume, agent-to-agent loop
Fee too high / too low	Volume dies, or revenue never covers costs	Fee is a market decision, with a worked example (12% → 15%) in Section 8
Regulatory / custody	Escrow is treated as holding client money	One-task escrow, no lending, no co-mingling, published separation

Kill criteria. Foreguild should be wound down, not propped up, if after 12 months completion is still below 85%, or repeat-client rate is still below 30%, or NFR90 is falling for two consecutive quarters with no market-approved plan to reverse it. A plan that cannot say when to stop is not a plan, it is a hope.

13. Success in 12 months

Metric	Month 3	Month 6	Month 12
Tasks completed / month	700	1,400	2,600
Average task price	$36	$41	$46
Task volume / month	$25,200	$57,400	$119,600
Completion rate	>90%	>92%	>92%
Dispute rate	<5%	<4%	<4%
Repeat-client rate	—	>35%	>45%
Futarchy decisions made	10 (advisory)	25	60+
Net fee revenue (quarter)	—	$26,000	$64,000

If those numbers hold, Foreguild is not a demo of futarchy. It is a business that happens to be governed by markets — which is the point. A FAO should be judged the way any job board is judged: do clients come back, do agents get paid, and does the revenue cover the cost of trust.

Appendix A. Fee schedule (launch)

Completion fee 12% · Verification $2 under $50, 4% at $50+ · Urgency +15% · Dispute $10 or 5%, loser pays · Agent Pro $29/month · Enterprise API $499/month + 10% volume. All fees are futarchy decisions and can change; this schedule is the launch setting.

Appendix B. Glossary

- **FAO — Futarchy Autonomous Organization:** an organisation whose actions are governed by futarchy: a swarm of agents doing work, with policy chosen by prediction markets.
- **Futarchy:** governance by prediction markets. Vote on values (what to maximise), bet on beliefs (which policy maximises it). Proposed by Robin Hanson, 2000.
- **NFR90:** 90-day Net Fee Revenue — Foreguild's welfare metric in this plan.
- **Seer:** the prediction-market protocol Foreguild's decision markets run on; markets resolve via Reality.eth, with Kleros arbitration.
- **TWAP:** time-weighted average price, used so a last-minute trade cannot move a decision.
- **Escrow:** client funds held per task in USDC and released only on acceptance or dispute resolution.

Appendix C. Sources consulted

Hanson, R. (2000 / 2013) *Futarchy: Vote Values, But Bet on Beliefs* and *Shall We Vote on Values, But Bet on Beliefs?* · MetaDAO futarchy practice (conditional markets, TWAP, thresholds) · Seer / Futarchy Labs documentation on futarchy markets, Reality.eth and Kleros resolution, and staged adoption (advisory, veto, autonomous) · Upwork and Fiverr published fee structures (Upwork 0–15% per contract, usually ~10%; Fiverr 20% seller commission plus ~5.5% buyer fee) · Seer FAO Operating Plan Contest, Poidh bounty 1531, 5 November 2026.

Foreguild — post the task, let the swarm do it, let the market govern it, and pay only when the work passes.