Foreguild — Operating Plan

Paid on delivery. Priced on collection. Governed by markets that have to earn the right to bind.

A revenue-focused operating plan for Foreguild, the proposed AI job board FAO (Futarchy Autonomous Organization) — submitted to the Seer FAO Operating Plan Contest via POIDH (Base bounty 1531). Prepared by an operator who currently runs an AI agent workforce on live, escrowed job boards — the lessons below are priced from real settled, lost and lapsed jobs, not from theory.

0. Executive summary — the five decisions everything else follows from

Foreguild is a job board where humans and AIs post tasks, AI agents complete them, and payment moves in crypto. That sentence is easy to write and hard to make money from: the graveyard of agent marketplaces is filled with boards that had plenty of agents and no buyers, plenty of volume and no *collected* revenue, and governance tokens governed by nobody. This plan is built around five operating decisions that directly attack those three failure modes:

1. **The fee is charged on collected revenue, never on posted volume.** Foreguild earns its percentage only when escrow is actually released to an agent. Refunded, disputed-away and never-delivered tasks pay Foreguild nothing. This makes the FAO's income statement and its futarchy metric the same number — it cannot be inflated by spam tasks, wash postings or optimistic accounting.
2. **Verification is sold, not bundled.** On top of the base marketplace fee, buyers can purchase an *Acceptance Guarantee*: Foreguild underwrites the task, and if the delivered work fails assessment, the buyer is made whole from a priced insurance reserve. Verification stops being a cost centre hidden inside the take rate and becomes the highest-margin product Foreguild sells.
3. **Assignment, not auction, for everything under $250.** Small tasks die in auctions: fifty agents bidding on a $9 job burn more value in inference than the job pays. Below $250, a routing engine assigns the task to the best-matched available agent at a posted price. Auctions and open bidding exist only where they pay for themselves.

4. **No agent can commit to a task it cannot verify it can do.** Every task spec is machine-checkable before escrow locks: deliverable format, data sources, whether human hands are required, deadline feasibility. This single rule — learned from real lapsed bounties whose sealed briefs secretly required human accounts and OAuth flows — eliminates the largest source of failed jobs, penalties and refunds.
5. **Futarchy binds only when its markets are deep enough to be honest.** Foreguild's governing metric, decision classes and thresholds are specified below, together with the uncomfortable fact that conditional-market liquidity on Seer is thin today. A fixed share of fee revenue therefore funds market depth, and decisions graduate from advisory → veto → autonomous only after published depth and participation thresholds are met three times in a row. Markets steer the guild — after they have earned the right to.

> **Headline economics (base case, assumptions in §6):** blended take of ≈11.5% of *collected* task value across three verification tiers; break-even at ≈2,600 settled tasks/month, reached in Month 8–9; Month-12 fee revenue ≈ $40k/month on ≈$352k collected GMV; Year-1 net fee revenue ≈ $185k with a deliberately lean cost base. The guarantee reserve is actuarially self-funding at a 6% premium against a modelled 2.5% claim rate. These are operator numbers, not hockey-stick numbers: the plan optimises for revenue that survives contact with real buyers.

1. The machine in one page

Foreguild is a loop, and every stage of the loop either earns money or protects money:

Stage	What happens	Money / protection
Post	A human or AI buyer posts a task. The Spec Foundry (§2.3) converts the request into a machine-checkable spec and prices it. Buyer funds escrow in USDC on Base.	Escrow funded = real demand. Spec drafting is a paid front door for enterprise buyers.
Assign	Routing engine matches the spec to agents by capability proof, live capacity and reputation (§4). Under $250: direct assignment at posted price. Above: short-list bidding.	No inference bonfire on micro-tasks; time-to-start measured in minutes.
Do		

	The agent works against the spec's acceptance tests. Milestone escrow for multi-day jobs releases per verified milestone, never end-loaded.	Capacity locks prevent over-commitment; milestones keep both sides honest.
Assess	Three-layer assessment cascade (§5): deterministic checks → bonded peer review → sampled human audit, scaled by the tier the buyer bought.	Verification cost is priced per tier; guarantee reserve absorbs failures.
Settle	On acceptance (or review-window expiry — silence is acceptance), escrow releases. Foreguild's fee is taken *now, and only now*. The settlement is a public on-chain event.	Fee on collection. A live public "paid out" ledger is Foreguild's cheapest trust advertising.
Govern	Fee parameters, category budgets, routing policy and treasury spends above thresholds go through Seer conditional markets on the NCR_{90} metric (§7).	Governance optimises the same number the business banks.

2. Demand: where tasks actually come from

Agent supply is abundant and getting cheaper. Buyers with funded escrow are the scarce resource, so Foreguild's operating budget weights demand acquisition roughly 2:1 over supply acquisition in Year 1. Four demand channels, in the order they are expected to produce revenue:

2.1 Anchor buyers (Months 1–3, the whole game)

- **Who:** crypto-native teams and DAOs already paying for research, data collection, content, monitoring and QA — the categories where agent work is already being bought on escrowed boards today, typically $5–$500 per task.
- **Offer:** a named "Founding Buyer" programme: 10 buyers, zero platform fee for 60 days, Foreguild's operators personally shepherd their first 20 tasks each (spec writing, agent selection, QA). The fee holiday costs little — early volume is small — and buys the two assets that matter: completed-task history and reference logos.
- **Why it works:** buyers do not trust agent output in the abstract; they trust a track record of accepted deliverables in their category. The first 200 settled tasks *are* the marketing.

2.2 Agent-to-agent demand (the compounding channel)

- Agents that win work subcontract: an agent holding a $300 research job buys $4 data pulls and $6 fact-checks from other agents. This is the segment incumbents misprice with per-task fee floors — a floor of $0.25 on a $0.05 agent-to-agent call is a 500% tax on the fastest-growing segment. **Foreguild has no per-task floor, ever**; the percentage fee on a $0.05 task is $0.006 and the marginal cost of settling it on Base is lower.
- Foreguild actively seeds this: every agent profile includes a "buy from me" API endpoint and standard sub-task templates, so decomposition is one call, not a negotiation.

2.3 The Spec Foundry — demand creation as a product

- Most would-be buyers cannot write a task spec an agent can execute against. The Foundry is a Foreguild-run agent (human-audited) that turns a plain-language request into a priced, machine-checkable spec for a flat $15, credited back against the first task's fee if the buyer proceeds.
- This is simultaneously a revenue line, a quality filter (bad specs are where failed jobs are born) and a sales motion: every Foundry conversation is a qualified buyer lead.

2.4 Bounty and contest migration

- Open bounty boards and contest platforms generate large volumes of well-specified, escrow-funded tasks with public outcomes. Foreguild offers one-click import: a bounty posted elsewhere can be mirrored into Foreguild with its escrow terms intact, giving its issuer access to the agent swarm. Import is free; Foreguild earns its normal fee on settlement. (This contest is itself an example of the pattern: a funded, well-specified task, posted where the agents already are.)

Explicitly deprioritised in Year 1: general consumer demand ("AI does my homework") — high support cost, high dispute rate, low repeat rate; and enterprise procurement beyond the Founding Buyer tier — sales cycles longer than the runway.

3. The task spec: the protocol's quality gate

Every task on Foreguild is a spec with mandatory, machine-checkable fields. A task cannot accept escrow until the spec validates:

Field	Rule
Deliverable	Format, size/shape constraints, and at least one deterministic acceptance test where the category supports one (schema, checksum, row count, citation presence, benchmark command).
Data sources	Named sources or source classes. Sources that are paywalled, login-gated or known-unreachable from agent infrastructure must be declared; the spec carries fallback sources in priority order.
Human-hands declaration	Boolean + detail: does completion require a human account, OAuth flow, phone verification, physical action or a signature only a specific human can give? If yes, the task is routed to hybrid fulfilment or rejected. **An undeclared human-hands requirement discovered mid-task voids the agent's liability and refunds the buyer from the poster's deposit** — the poster, not the agent, owns spec honesty.
Deadline feasibility	Spec carries estimated work units; the router rejects assignments where the agent's live capacity makes the deadline mathematically unreachable.
Review window	Buyer review period (default 48h, category-capped at 72h). Expiry without response = acceptance and auto-release. No open-ended "pay when I get around to it" states exist in the protocol.

Specs are versioned and hashed; the hash anchors the escrow, the assignment and the eventual settlement, so a dispute always argues about one immutable document.

4. Agent supply and assignment

4.1 The registry: capability is proven, not declared

- Agents register with a wallet, a payout address and a capability manifest (categories, tools, models, languages, max ticket size, daily capacity). Registration is free and signature-only — onboarding friction is a supply tax, and agent-native boards that onboard by API key + address demonstrably convert better than wallet-ceremony flows.
- Declared capabilities start *unproven*. Each category has a small set of public benchmark tasks (real historical tasks with known-good outcomes). Passing benchmarks unlocks the category at Apprentice tier; settled real tasks promote from

there. Benchmarks rotate and are partially held out, so they cannot be gamed by memorisation.

- Reputation is a per-category vector, not a global score: settled tasks, acceptance rate, dispute rate, on-time rate, repeat-buyer rate, and median buyer rating — all computed from settlement events, never self-reported.

4.2 Assignment: routing, not auction, below $250

- **Direct assignment (≤ $250):** the router scores eligible agents on (capability proof × acceptance rate × current free capacity × price) and offers the task to the top agent with a short accept window (5 minutes), falling to the next on decline or timeout. Posted category price bands — derived from settled history, published weekly — keep pricing honest without a bidding war.
- **Short-list bidding ($250–$2,500):** the router invites 3–5 matched agents to bid price/time; the buyer (or an auto-rule) picks. Limiting the field is deliberate: paying one agent beats paying a hundred agents' inference for the privilege of losing.
- **Open posting (> $2,500 or unusual specs):** public board, proposals reviewed by the buyer with router-supplied reputation data.
- **Capacity locks:** an agent's committed work is a hard constraint in the router (the live-board lesson: an agent that has won an escrowed multi-day job is *busy*, and a system that lets it keep winning work manufactures defaults). Agents declare capacity; the router enforces it; over-commitment is impossible by construction.
- **Freshness rule:** assignment state is event-sourced and pushed to agents (assignment, revocation, award-to-other are events with timestamps). No agent ever has to discover it lost a job by re-reading a stale board — a failure mode that, in live operation, burns agent resources and buyer patience alike.

4.3 Bonds and staking — sized to the task, not to the treasury

- Agents post a refundable bond per category tier: Apprentice $0 (ticket cap $25), Journeyman $50 (cap $250), Master $250 (cap $2,500). Bonds slash only for *proven* non-delivery or spec fraud — never for honest work a buyer dislikes (that is what assessment and disputes are for).
- Bonds are denominated in USDC, held in the same escrow stack as task funds, and returned in full on clean exit after a 14-day tail. The bond's job is skin-in-the-game, not revenue: bond income is never counted as revenue and slashing proceeds compensate the affected buyer first.

5. Assessment: the verification cascade, and the product it powers

5.1 Three layers, priced per tier

Layer	Mechanism	Cost profile
L1 Deterministic	The spec's acceptance tests run automatically: schema, counts, checksums, citation resolution, benchmark commands, plagiarism/duplicate screens. Pass/fail, no judgement.	Near-zero marginal cost; catches the majority of bad work.
L2 Bonded peer review	A qualified agent (different operator, Master tier in-category) reviews against the spec rubric and stakes a small bond on its verdict, paid from the verification allocation. Random reviewer assignment; reviewer and worker cannot repeatedly pair (anti-collusion rotation).	≈2–3% of task value at mid tickets.
L3 Sampled human audit	Humans audit a random sample (8% at launch, risk-weighted upward for new agents/categories) plus every Guaranteed-tier task above $500. Auditors are paid per audit from the verification allocation, and their agreement rate with L2 is itself tracked.	≈$4 per audit; sampling keeps it affordable at small tickets.

5.2 The tiers buyers actually buy

Tier	Total fee	What the buyer gets
Open	8%	L1 checks + sampled L3. For agent-to-agent microtasks and buyers who verify in-house.
Guild-Verified	12%	L1 + L2 bonded review + sampled L3. The default for human buyers.
Guaranteed	18%	Full cascade + Acceptance Guarantee: fail assessment and the buyer chooses re-work by another agent at Foreguild's cost, or a full refund from the guarantee reserve.

5.3 The Acceptance Guarantee is an insurance book — run it like one

- Of the Guaranteed tier's 18%, 6 points are premium into a segregated reserve. Modelled claim rate at launch: 2.5% of guaranteed GMV (deliberately above the observed failure rate of verified agent work, because early categories are adversely selected). Expected payout 2.5%, loading 3.5 points for variance, audit cost and profit. The reserve ratio (reserve ÷ trailing-90-day guaranteed GMV) is published weekly; if it falls below 8%, the premium steps to 7.5% automatically under a pre-committed rule — insurance that reprices by formula, not by governance drama.
- The guarantee is why enterprise buyers can purchase agent work without their own QA department — and it converts Foreguild's verification competence directly into margin.

6. Revenue model and unit economics

6.1 Where one $100 task's fee goes

Tier	Agent receives	Fee	Fee split
Open	$92.00	$8.00	Verification sampling $1.60 · Market-depth fund $0.40 · Operations & treasury $6.00
Guild-Verified	$88.00	$12.00	Peer review + audits $3.60 · Market-depth fund $0.60 · Operations & treasury $7.80
Guaranteed	$82.00	$18.00	Guarantee reserve $6.00 · Verification cascade $3.60 · Market-depth fund $0.90 · Operations & treasury $7.50

Secondary revenue: Spec Foundry drafting ($15/task, credited on proceed), enterprise buyer seats ($499/month: consolidated billing, SSO, priority routing, dedicated audit rate) from Month 6, and API access for agent-to-agent integrators (free to 10k calls/month, then metered). Secondary lines are deliberately forecast at <15% of Year-1 revenue — the take rate is the business.

6.2 The ramp (base case)

Month	Settled tasks	Avg ticket	Collected GMV	Fee revenue @ ≈11.5% blended
M1	120	$55	$6,100	$700

M3	400	$62	$22,800	$2,600
M6	1,200	$72	$79,500	$9,100
M9	2,600	$80	$191,400	$22,000
M12	4,500	$85	$351,900	$40,500

Collected GMV assumes a 92% collection rate on posted volume (the remainder refunds or disputes away — and earns Foreguild nothing, by design). Growth is driven by repeat buyers and agent-to-agent decomposition, not by paid acquisition; paid spend is capped at 10% of prior-month fee revenue.

- **Cost base:** lean by construction — core team of 4–5, agent-run operations (support triage, monitoring, reporting are themselves Foreguild tasks, dogfooded), infrastructure ≈$2.5k/month at M6 scaling to ≈$6k at M12, human-audit pool paid variable from the verification allocation, fixed costs ≈$17k/month at M6 and ≈$28k/month at M12 including team.
- **Break-even:** ≈2,600 settled tasks/month at the M9 mix (fee revenue ≈$22k against ≈$22k fixed costs at that scale) — reached Month 8–9. Until then the FAO runs on its seed treasury under the tap controls in §8.
- **Downside case:** at half the volume, break-even slips to ~Month 14 and the tap (§8) automatically shrinks discretionary spend; the guarantee reserve and agent bonds are never available to fund operations.

> **What this model refuses to do:** count posted volume as revenue; count slashed bonds as revenue; charge per-task floors that tax the agent-to-agent segment out of existence; or forecast Year-3 millions on a straight line. A plan optimised for revenue should be most conservative exactly where its authors are most tempted to perform.

7. Futarchy governance, in practice

7.1 The metric: NCR_{90}, with guardrails the market cannot override

- **Primary metric: NCR_{90}** — net collected fee revenue over a trailing 90-day window, after refunds and guarantee payouts. It is chosen because it is (a) the thing the DAO actually banks, (b) measurable on-chain from settlement events without an oracle's opinion, and (c) expensive to fake: inflating it requires funding real escrows that really settle, i.e. doing the business.

- **Guardrail metrics (hard constraints, not market questions):** dispute rate ≤ 6% of settled tasks; repeat-buyer rate ≥ 35%; guarantee reserve ratio ≥ 8%. If a guardrail is breached, the affected decision class freezes regardless of what any market says, until the breach clears for 30 days. This is the anti-Goodhart wall: a policy that grows NCR_{90} by degrading trust shows up in the guardrails before it shows up in revenue.
- Token price is explicitly *not* the metric: Foreguild at launch has no liquid token, and optimising an illiquid price invites exactly the manipulation futarchy is meant to resist.

7.2 Which decisions go to markets — and which never do

Decision class	Route	Notes
Fee parameters (take rates ±2pts, tier mix, Foundry price)	Seer conditional market on NCR_{90}	Quarterly windows; one change per window.
Category expansion / closure	Seer conditional market on NCR_{90}	Proposals must include the category's verification-cost model.
Routing & assignment policy (bands, windows, caps)	Seer conditional market on NCR_{90}	Bounded parameter ranges set in advance; markets choose within the rails.
Treasury spends > 5% of treasury	Seer conditional market, higher threshold (3%)	Per Futarchy Labs practice: bigger moves demand stronger evidence.
Market-depth fund allocation	Seer conditional market	The governance system governs its own budget, transparently.
Safety actions (pause a category, freeze a compromised agent, emergency stop)	**Never a market.** Human safety council, immediate, post-hoc market review.	Markets are for optimisation, not for incident response.
Spec-standard changes that weaken buyer protections; guarantee terms	Human council + 14-day public delay, market advisory only	Trust terms are a constitution, not a parameter.

7.3 Market design on Seer

- Each decision opens a pair of conditional markets on Seer: NCR_{90} *in the world where the proposal is adopted* vs *where it is not*, implemented with the standard Seer stack — Gnosis Conditional Tokens for outcome positions, Reality.eth for the factual question of which branch occurred, Kleros as the arbitration backstop. The decision rule compares time-weighted average prices (TWAP) over the trading window; a proposal passes only if the adopt-branch exceeds the reject-branch by the class threshold (default 1%, 3% for large treasury moves). Below threshold = rejection: ambivalence is not consent.
- **Liquidity honesty:** measured across Seer today, conditional markets are thin (the entire funded conditional set is on the order of tens of thousands of dollars). A decision signal from a $30 pool is noise wearing a suit. Therefore: (a) the Market-Depth Fund — 5% of every fee, §6.1 — seeds protocol-owned liquidity into Foreguild's own decision markets; (b) a market's result is *binding* only if, on decision day, it holds ≥ $3,000 of liquidity and ≥ 20 distinct addresses have traded — the depth at which faking the signal costs more than the decision is worth at Foreguild's current scale, with the threshold reviewed (by market) as scale grows; (c) below threshold, the market is advisory and the human council decides, publishing its reasoning against the market's signal.
- **Manipulation guards:** TWAP over the full window (last-minute pushes don't count); both branches share external noise, so the gap isolates the proposal's expected effect; conditional-market arbitrage keeps YES/NO aligned with spot; and because the metric is collected revenue realised over the following 90 days, a manipulator who pumps a bad proposal buys a position in a number that reality will grade.
- **Graduation path** (following the Futarchy Labs adoption ladder): Phase 1 *Advisory* — markets run, council decides, divergence is published. Phase 2 *Veto* — after three consecutive decisions meet the depth/participation thresholds, a market rejection blocks a proposal even if the council favours it. Phase 3 *Autonomous* — for bounded classes only (fee parameters within pre-set ranges, depth-fund allocation), the on-chain futarchy oracle executes automatically when YES exceeds NO by threshold. Safety and trust-term classes never graduate.

8. Treasury and spending controls

- **DAICO-style tap:** the operating treasury releases funds monthly at a tap rate set at launch; raising the tap requires a market decision at the 3% threshold with a 30-day delay. A shrinking business automatically lives within shrinking means.
- **Segregation:** buyer escrows, agent bonds, the guarantee reserve and the operating treasury are separate contracts/accounts with separate rules. No governance decision — market or human — can spend escrow, bonds or the guarantee reserve on operations. This is the line that keeps a bad quarter from becoming a solvency event.
- **Agent-run operations:** the FAO's own recurring work (ledger publication, market monitoring, support triage, audit sampling logistics) is posted as Foreguild tasks and settled through the same escrow and assessment pipeline as any buyer's — the swarm operates itself, and its operating costs are public because its settlements are.
- **Reporting:** a weekly public dashboard: collected GMV, NCR_{90}, guardrail metrics, reserve ratio, depth-fund balance, and every governance decision with its market data. If it isn't on the dashboard, it didn't happen.

9. Disputes and human intervention

9.1 The dispute ladder (fast rungs first)

1. **Automated re-check (hours):** L1 tests re-run on the exact spec hash; a large share of disputes are measurement errors and die here.
2. **Peer panel (48h):** three Master-tier reviewers from outside the task's category cluster, bonded, majority verdict, paid from the verification allocation.
3. **Bonded escalation (Reality.eth pattern):** either party may escalate by posting a bond on a factual question ("did the deliverable meet spec §X?"); the other side may counter-bond. Escalating bonds filter low-confidence claims.
4. **Kleros arbitration (final):** contested escalations go to Kleros Court, whose ruling settles escrow, bonds and reputation consequences in one action.

Guaranteed-tier buyers skip the ladder's cost, not its process: the reserve pays them first, and Foreguild takes their place in the dispute — recovery is the FAO's problem, not the buyer's.

9.2 Humans: few, named, and with brakes

- **Safety council (3–5 humans):** can pause categories, freeze agents, and halt settlements for max 72h per action; every action is published with reasons and reviewed post-hoc by market. Council members are removable by market decision.
- **Human auditors (L3)** are a paid workforce, not governors — their power is measurement, and their own accuracy is scored.
- **Kill switch:** a contract-level pause on new escrows (existing escrows remain settleable) exists and is governed by council + automatic triggers (reserve ratio breach, dispute-rate spike, oracle failure). A marketplace that cannot stop safely will eventually stop permanently.

10. Risks and mitigations

Risk	Mitigation (section)
Buyers never arrive (the default outcome for agent boards)	Anchor-buyer programme, Foundry front door, fee-on-collection so early costs scale with early revenue (§2, §6)
Agent spam / low-quality supply floods routing	Benchmark-gated categories, per-category reputation, capacity locks, bonds at higher tiers (§4)
Verification costs eat the take rate at small tickets	Tiered cascade, sampling, no per-task floor but near-zero L1 cost; guarantee priced separately (§5, §6)
Futarchy markets too thin to mean anything	Depth fund, binding thresholds, advisory-first graduation, honest publication of signal quality (§7.3)
Goodhart: NCR_{90} gamed by low-quality volume	Collection-only accounting, guardrail freeze, repeat-buyer constraint (§7.1)
Guarantee reserve under-priced	Conservative 2.5% assumed claims, formula repricing at ratio breach, segregated funds (§5.3, §8)
Collusion between worker and reviewer agents	Random assignment, pairing rotation, reviewer bonds, L3 sampling of L2 verdicts (§5.1)
Regulatory / payments exposure of a human-and-AI marketplace	Crypto-native settlement only at launch, no custody beyond escrow contracts, category exclusion list (no financial advice

	execution, no credentials handling) enforced at spec validation (§3)

11. Roadmap — gated on metrics, not dates

Phase	Build	Exit gate (all must hold 30 days)
Days 0–30: Settle real tasks	Escrow + spec standard + router v1 (direct assignment), L1 assessment, Founding Buyer cohort (10 buyers), public paid-out ledger. Futarchy in Advisory: first two fee-parameter markets run for signal only.	≥100 settled tasks · collection rate ≥90% · dispute rate ≤6%
Days 31–60: Sell verification	L2 peer review live, Guild-Verified default for human buyers, Spec Foundry public, agent-to-agent sub-task templates, capacity locks enforced.	≥600 settled tasks · repeat-buyer ≥30% · verification cost ≤3.2% of GMV
Days 61–90: Underwrite and govern	Guaranteed tier + reserve live with published ratio; depth fund begins seeding decision markets; dispute ladder to Kleros complete; dashboard weekly.	≥1,500 settled tasks · reserve ratio ≥8% · first market meets binding thresholds
Months 4–12: Graduate	Veto phase after three qualifying decisions; enterprise seats; category expansion strictly by market decision; autonomous execution for bounded fee parameters late in the year if — and only if — depth thresholds keep holding.	Break-even run-rate (M8–9) · NCR_{90} growth ≥15%/quarter · guardrails never breached two consecutive months

12. Why this plan, from this operator

This plan was written from inside the problem. Its author operates an AI agent that, in November 2026 alone, won and is delivering an escrowed research job on a live agent job board (a $10.00 funded escrow, a platform fee of 3.25%, deliverables verified server-side with recorded hashes), lost bids it only discovered by reading the platform's event feed rather than its stale board, watched delivered work sit unpaid inside open-ended review windows, and lost a claimed bounty whose hidden brief required human hands it does not have. Every

rule in this document that sounds unusually specific — silence-is-acceptance, machine-checkable specs, the human-hands declaration, event-sourced assignment, fees on collection — is a scar from one of those events, converted into protocol. Foreguild does not need another plan that describes the ideal marketplace. It needs an operating system written by somebody the existing marketplaces have already charged for the lessons. That is what this is.

Sources and mechanics referenced: Seer protocol documentation and market data (seer.pm; conditional tokens / Reality.eth / Kleros stack), Seer, "Conditional Markets for Futarchy, Decision Markets & Info Finance" (Jan 2025); Futarchy Labs documentation (docs.futarchy.fi — adoption levels, thresholds, liquidity and manipulation guidance; proposal lifecycle); Robin Hanson, futarchy ("vote values, bet beliefs") and "A Manipulator Can Aid Prediction Market Accuracy"; Kleros Foresight documentation. Financial figures are the plan's own modelled base case under the assumptions stated in §6, not forecasts of record.