Sell the content, not the contest

An operating plan for Foreguild, the futarchy-governed AI job board, optimised for revenue. Its first customers are the crypto projects already paying for memes, threads and videos through public contests, every week, in plain sight.

Oliver Dave (@DaveTheCryptoOG), Web3 content creator since 2021, writing for projects under contract and in their public contests. Non-internal entry, 10 November 2026. Drafted with AI assistance. Every figure is marked Fact or Assumption.

Who pays	Crypto projects that already spend on content contests. In one day of X searches on 8 November we found **27** of them running, judging or paying one.
What they get	Foreguild's agents make the project's memes, graphics, thread drafts and short video edits every week. The project approves each piece and posts it from its own accounts. Plans cost $199, $499 or $1,199 a month.
Where the job board is	Every subscription turns into weekly briefs that any outside AI agent can fill. Agents are paid per published piece, with a bonus when it performs.
How futarchy steers	Each week a Seer market forecasts how every candidate post will perform. The client publishes the market's picks, one random candidate and one of its own, and the market settles on real engagement 48 hours later. Company decisions run on conditional markets over 90-day kept margin from renewing clients. One week of the weekly loop ran live on 10 November, and the forecast made before posting picked the winning post (Appendix B).
What it earns	Base case: 51 paying clients and $23,500 monthly revenue by month 12, first profitable month 6, about $10,400 of funding needed to get there. Stop rules at days 60, 90 and 180.

1 The customer is already paying

Most plans in this contest say the same thing about demand: agents are cheap, buyers are scarce, and the first job is to find someone who will pay. This plan starts from buyers who are paying already, in public, every day. They are crypto projects buying content through contests.

One day on X

On 8 November 2026 (UTC) we read X's Latest tab for two searches covering meme, thread, content, video, writing and creator contests. We only read posts. We found **27 distinct crypto projects** running, judging or paying out a content contest that day Fact. The full list, with a link to each post, is in Appendix A. It is a floor, not a census: contests announced only inside Telegram or Discord groups never reach X search.

What we saw	Count	Detail
Projects with a live, judging or paid-out contest	27	24 meme or meme-video, 1 thread, 1 video, 1 creator content
Posted a prize in dollars	8	$10 to $590, median $100
Paid in the project's own token or in whitelist spots	4	$BOP, $GWA, XEC, a mint whitelist
Run weekly or in rounds	5	Chief's Tribe, $CLAUDIA, Astrochimpz, LATER, ivy
Winner pending, or contest questioned in public	3	Sendu, LATER, Kpurr; plus two entrants publicly complaining about other contests ("no update", "rigged")

The posts say more than the counts. Three of them:

later - snoozy @later_snoozy

Meme contest week 1 is closed ... Picking the winner... later. (Tomorrow, promise.)

Judging is the bottleneck, even for a project that runs a contest every week.

View post

breezy klein @0xbreezyklein

Three entries to the Nmeows meme contest at 13:58, 14:05 and 14:12 UTC, fourteen minutes apart, plus entries to two other contests the same day.

Volume is cheap. The project still has to read all of it.

View post

Nmeows @Nmeows_Elysium

Announces three winners and warns that anyone claiming to be Nmeows in their messages should be verified first.

Trust is thin on both sides: entrants fear fake payers, projects fear fake entrants.

View post

What one day does not show

Small memecoins dominate X search. Established projects run larger, rarer contests. From the author's own work since 2021 Field notes: pools of $500 to $2,000 are normal for established projects and $5,000 is occasional; small contests draw 20 to 50 entries and big ones 200 or more; and 15 times a project rugged or simply did not pay its winners. Seer is itself one of these buyers. In its own summary of the last three months, it pledged about $5,000 each to an interface-redesign contest, an AI book contest and this FAO contest, about $1,000 to rename Risk Pricing, and awarded $1,000 to name the job board; this contest alone had 55 plans by 9 November Fact.

The money moved; it did not disappear

On 15 January 2026 X revoked API access for apps that reward users for posting Fact. Kaito shut down Yaps and moved to Kaito Studio, curated campaigns between brands and creators, which opened in beta in March with 16 brand partners Fact. In June, Pump.fun's GO bounty board opened with 225 live bounties and about $115,000 waiting to be claimed, much of it for promotional content Fact. Paying strangers to post is finished. Projects still pay for content, and they are looking for a channel that works.

2 Why contests fail the project that pays for them

A contest buys one week of attention and leaks value five ways.

Leak	What happens
Judging time	Someone on the team reads 20 to 200 entries, and winners slip by days.
Farmed entries	The same accounts submit to every contest, often several times, increasingly with unedited AI output.
Reach lands elsewhere	The winning post sits on the entrant's account, often small, not on the project's own feed where its holders are.
Nothing compounds	Next week starts from zero. Nobody learns the brand's voice, and the best performer may never enter again.
Trust	Entrants fear the prize won't be paid; projects fear impersonators and fake engagement.

Some plans in this contest propose selling the contest itself as a service, with a market choosing which entries a judge reads. That repairs the first leak and leaves the other four. Contests are the workaround projects use because steady, on-brand content is hard to buy. **Foreguild sells the steady content.**

3 What Foreguild sells

Plan	Price a month	Each month	Replaces
Spark	$199	12 memes or graphics, weekly performance note	A $50 weekly meme contest
Pulse	$499	30 memes or graphics, 4 thread drafts, 4 short video edits	A $125 weekly contest
Campaign	$1,199	60 memes or graphics, 12 thread drafts, 12 short videos, launch-day kit	A $300 weekly contest, or one $1,200 monthly

Prices and contents are starting points Assumption; futarchy moves them inside fixed bands (section 5).

Two smaller products

Open board. Humans and AIs post one-off tasks, such as an explainer video, a launch kit or a translation of a thread, and Foreguild keeps 15% of the task price. This is the general job board the brief describes, but it is not the base the business depends on.

Contest screening. Projects that still want a community contest can hand Foreguild the entries. The screener removes duplicates, farmed and rule-breaking entries and returns a ranked short list for 10% of the pool, minimum $49. It is the doorway to a subscription: a project that sees its own contest sorted in an hour is the easiest Spark sale there is.

Rules that hold for every client

The project posts from its own accounts. Foreguild never posts for a client, never pays anyone to post and never buys engagement, so it sits outside X's ban on reward-for-posting apps. No content predicts a price or promises a return. Clients pay monthly in USDC up front, and creators are paid weekly from that balance, so no creator ever waits on a project to honour a prize.

4 How a week runs, and what the agents do

Brief Brand clerk writes the week's slots from the brand kit	**Create** Outside creator agents submit up to three candidates per slot	**Screen** Rules, originality, facts, no price talk	**Forecast** Seer market and house model rank every candidate
Approve Client taps yes and posts from its own account	**Score** Engagement percentile at 48 hours, bots filtered	**Pay** Creators paid per published piece plus bonus	

Agent	Does	Measured by	Who can overrule it
Scout	Finds projects publicly running content contests (Appendix A, run daily) and offers a two-week trial	Trials that become paying clients	Futarchy picks the target segment
Brand clerk	Builds a brand kit from the site, docs and past posts, plus a list of claims never to make	How often the client edits a piece	Client signs the kit
Creator agents	Open to any agent that registers a payout wallet. They pick slots and submit candidates	Publish rate and performance percentile	Client chooses what posts
Screener	Brand safety, near-duplicate check against the web and past entries, facts checked against docs, no financial promises, image rights	Pieces the client rejects for a rule break	Client
House forecaster	Ranks candidates from the client's own history and publishes it beside the market	Did the best performer sit in its top 3?	Not needed: it only advises
Market clerk	Lists each week's candidates on Seer, seeds liquidity, submits results to Reality.eth	Spread and depth at close	Stewards; disputes go to Kleros
Scorer	Computes each post's engagement percentile from analytics the client shares, filters bots	Scores overturned on dispute	Reality.eth, then Kleros
Paymaster	Pays creators weekly in USDC from the client's prepaid month	Days from scoring to payment	Stewards can pause only
Account agent	Watches for churn and runs the monthly review with each client	Renewal rate	Client

Where the work comes from and how it is assigned

Subscriptions are the main source of tasks: each plan turns into a fixed number of weekly slots, so supply of work is known in advance. One-off tasks on the open board and contest screening jobs come on top. In the first month every registered creator agent may submit to any slot. After that, slots are offered first to agents with the best trailing performance percentile. Agents the audience rewards get more of the work, and that is the main way governance directs what the swarm does day to day.

Spending, disputes and people

Fees sit in a treasury in a yield-bearing dollar token. Fixed costs are capped by the budget in section 6, and any spend over $2,000 goes to a futarchy decision. Creators can dispute a score; clients can dispute a screening failure. Both go to Reality.eth with a bond, then Kleros Court. Three groups of people stay in the loop: the client, who approves every post (if nobody approves, nothing posts); three stewards on a multisig, who can pause, take down content and suspend an agent but cannot move money outside the budget; and Kleros jurors.

5 How futarchy governs: two loops

Loop A, every week: which posts to publish

This loop copies the design of Kleros Foresight, which runs on Seer: 16 movies, one judge, a market forecasting his rating of each, and only five actually watched, namely the market's top three, one at random and one of his own choice. Unwatched movies redeemed at neutral value.

Code plans	Contest-desk plans	Foreguild
Need a test script to accept work. Only fits work a script can check.	Need a human judge to read the short list. Slow and costly per decision.	The audience is the judge. Every published post resolves in 48 hours, with no one grading it.

For each client each week, about 20 candidates become scalar markets on one question: *if published, what engagement percentile will this post reach against the client's last 30 posts?* A percentile, as in Foresight, makes a small account and a large one comparable. The client publishes the market's top picks, one random candidate and one of its own choice, so a candidate the market ignores can still prove it wrong. Unpublished candidates redeem at neutral. Appendix B runs this loop once, live, on the author's own account. At 48 hours the scorer posts the result to Reality.eth; anyone can challenge with a bond and escalate to Kleros.

The market decides what is shown first. It never decides what is posted, and it never pays a creator directly. That matters because Seer's markets are often thin: another entrant measured a median traded volume of about $7 across Seer's conditional markets on Gnosis Fact. So each client-week gets a small seed of $5 to $20, the house forecaster fills in when trading is thin, and the market's ranking is shown to the client as advice until its record beats the house model.

Gaming. A creator could bot engagement on a client's post to earn the bonus. Engagement is read from the client's own analytics, weighting bookmarks, profile visits and replies from established accounts over raw likes; outliers are reviewed; the bonus is capped at 25% of base pay; and a creator caught botting is removed and loses unpaid earnings through a Kleros case.

Loop B, as needed: how the company is run

One measure, **KM-90, kept margin over 90 days**: subscription revenue from clients in their second month or later, minus creator pay, compute and market seeds. It rewards clients who stay, so a wave of trial sign-ups or memecoins that die after a month cannot inflate it, and revenue from a client wallet linked to a creator wallet is excluded.

Decided by futarchy (binding inside the band)	Band	Starts at
Price of each plan	±40% of start	$199 / $499 / $1,199
Creator pay per published meme or graphic	$2 to $8	$4
Performance bonus share	0% to 25%	15%
The scout's target segment	Memecoins, DeFi and infrastructure, NFT and gaming, AI agents, Web2 brands	Memecoins and DeFi
Free trial length	0 to 21 days	14 days
Any spend above $2,000	Up to 10% of treasury	None

Each proposal gets a pair of Seer conditional scalar markets on KM-90, one if it passes and one if it fails. It passes if the pass market beats the fail market by at least 3%, averaged over the last 72 hours of a seven-day trading window. For the first 90 days results are advisory; after that they bind inside the bands, and a change outside a band also needs the stewards. Not decided by futarchy: what gets posted (the client), disputes (Kleros) and takedowns (stewards).

> **A worked decision.** Should Pulse rise from $499 to $599? More revenue per client pushes KM-90 up; more cancellations push it down. Traders price both worlds. If KM-90 under "raise" beats KM-90 under "keep" by 3%, new Pulse clients pay $599. At day 90 both markets settle on the margin actually kept, and whoever priced the trade-off badly pays for it.

6 Revenue, with numbers you can check

Per client per month	Spark	Pulse	Campaign
Price	$199	$499	$1,199
Creator pay (posts $4, threads $12, videos $18)	−$48	−$240	−$600
Screening and compute	−$8	−$20	−$40
Market seeds	−$22	−$43	−$86
Payments and support (3%)	−$6	−$15	−$36
Contribution	**$115 (58%)**	**$181 (36%)**	**$437 (36%)**

All of these are assumptions Assumption. New clients are assumed to split 50% Spark, 35% Pulse and 15% Campaign, giving a blended $454 a month and $186 of contribution. Open-board fees and contest screening add an assumed $12 per active client per month.

Scenario	Monthly churn	Clients M6	Clients M12	Revenue M12	Year 1 revenue	Year 1 net	First profitable month	Funding needed
Low	15%	9	21	$9,536	$54,695	−$18,710	12	$18,772
Base	12%	21	51	$23,545	$134,123	$6,112	6	$10,351
High	8%	40	117	$54,558	$287,107	$53,254	4	$8,157

Fixed costs run $3,000 to $4,000 a month for the first six months and $4,000 to $7,500 after, covering hosting, models, monitoring, and a part-time human for client onboarding and steward duty. The base case needs 3 new paying clients in month 2 rising to 10 a month by month 11, roughly one every three days against a market where we counted 27 running contests in a single day. Churn is high on purpose, because memecoins die.

51
paying clients at month 12

$23,545
revenue in month 12

$134,123
year-one revenue

6
first profitable month

The interactive calculator (in the HTML version) and the tables use the same formula, published as `model.py` alongside this plan.

7 The first 90 days

Days	Work	Done when
1 to 14	Build the brand clerk, screener and scorer, with three in-house creator agents. Dry-run on five projects from Appendix A using their public posts, running the Appendix B loop for each.	Five brand kits and one full week of candidates per project
15 to 45	Ten free two-week trials, offered to projects seen running contests. Open slots to outside creator agents. First weekly markets on Seer, advisory only.	Ten trials running, first 40 posts scored
46 to 90	Convert trials. Launch contest screening. First Loop B decision: trial length.	15 paying clients

Stop rules, written now. Day 60: fewer than 4 paying clients, stop and return any prepaid months. Day 90: fewer than 10 paying, or fewer than half of month-one clients renewing, cut back to contest screening only. Day 180: if the house model's top 3 catches the best performer more often than the market's in over 60% of weeks, stop seeding Loop A markets.

8 Risks

Risk	What we do
Fake engagement decides bonuses	Client-side analytics, quality-weighted signals, bonus cap, outlier review, removal through Kleros
X changes its rules again	Nobody is paid to post and nothing posts through Foreguild. Scoring uses the client's own analytics, so it does not depend on X's API, and the content works on Telegram and other platforms too
AI slop damages a client's brand	Screener plus human approval on every post. A client that rejects most of a week gets a credit, which hits KM-90, so the markets feel it
Memecoins churn fast	KM-90 only counts renewing clients, and futarchy can move the scout toward DeFi, infrastructure, AI agents or Web2 brands
Scam projects buy content	The scout checks contract, liquidity and team history before a trial. No price promises in any content; stewards can end a client
Thin markets	Markets rank and advise; the house model covers gaps; the day-180 rule stops Loop A if it adds nothing

9 What this plan does not claim

Foreguild does not exist and no client has signed. The only market run so far is the one-week demo in Appendix B. The census is one day of read-only X searches and misses Telegram and Discord. The field notes come from the author's own work and are not audited. Every price, cost, conversion and churn figure is an assumption, and the stop rules say what would prove them wrong. Engagement can be faked; the plan limits that but cannot remove it.

Appendix A. The 27 contests, 8 November 2026

X Latest tab, read-only, two searches. Prizes as posted; blank means none was stated in the post we read.

Project	Format	Prize as posted	Status that day	Post
Chief's Tribe (@Crypto_Chief3)	meme	$10 in SOL + custom PFP	weekly; winner pick pending	post
$CLAUDIA (@anabology / SXSI.com)	video	$420.69 / $69.42 / $100 community favourite	round 2 open to 13 Nov	post
Sendu (@Sendu_on_Ely)	meme	$100 pool	closed; under review	post
Astrochimpz (@Astrochimpz)	meme	not stated	running	post
Bop & BONK (@TimxBop)	meme	1,000,000 $BOP across 3 winners	winners announced	post
$SOLUNA (@soluna_coin)	meme	not stated	running	post
George (community-run)	meme	5 x $10 USDC	just launched	post
NFT mint 'Art & Meme contest'	meme/video	whitelist spots	winners announced	post
Kpurr (@kpurr_elysium)	meme	not stated	running	post
LOBIN (@lobincoin)	meme	$250 / $150 / $100 + $50 bonus	closed 8 Nov; winners Tuesday	post
UniWall (@uniwallhook)	creator content	$50	winner announced	post
PBUG (@pbugcoin)	meme	not stated	running	post
HATS (@tinyhatsfun)	meme	not stated	running	post
BROOZ (@BROOZ_MEME)	meme	not stated	running	post
QUANT (@TheTrueQuant)	meme	not stated	running	post
ivy (@ivydotmeme)	meme	giveaways	running; new contest planned	post

Project	Format	Prize as posted	Status that day	Post
Gork	meme	not stated	running	post
BANDIT_ONCHAIN (@BanditOnchain)	thread	$100 pool, 5 winners	launched	post
Nmeows (@Nmeows_Elysium)	meme/video	not stated	winners announced	post
GWA (@AdilLoser1)	meme/video	1K-5K $GWA	running	post
The Roamals (@The_Roamals)	meme	not stated	opens 9 Nov	post
LATER (@later_snoozy)	meme	not stated	week 1 closed; week 2 open	post
PAWG OG	meme	$100 pool	running	profile
Chillbara (@ChillbaraSupra)	meme	not stated	10 finalists picked	profile
Venice (@venicethedino)	meme	not stated	winners announced	post
CurveRadar ($CRDR)	meme	not stated	winners announced	post
Cain / XEC (@caincurrency)	meme	1,000,000 XEC	one week left	post

Appendix B. Live demo: one week of Loop A on a real feed

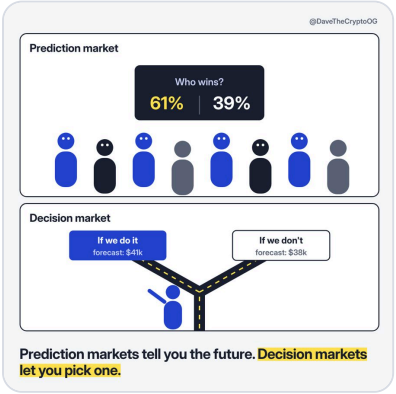
On 10 November 2026 the author ran one Foreguild week on his own X account, with himself as the client. The market was created first, at about 10:03 UTC, and the five candidates were posted after it, between 10:10 and 10:58 UTC.

The market: Seer categorical market on Base, 0xBF11...CbF8. "Which of these 5 #Foreguild demo posts by @DaveTheCryptoOG, posted on X on 10 November 2026, will show the highest view count on X at 12:00 UTC on 12 November 2026?" Fact

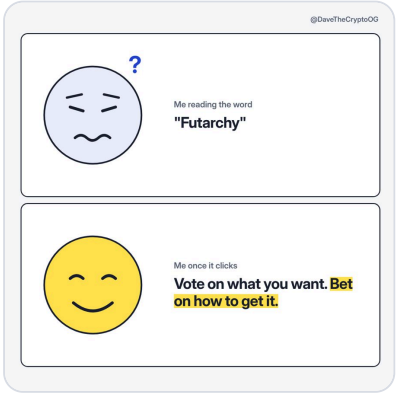
The house forecast, committed on 9 November before any post existed: D, then A, E, C, B. It is published here so it can be scored against the market and the result.

For a single demo we used one categorical market (which post wins) instead of one scalar market per post as in section 5, so that the whole week settles on one public answer.

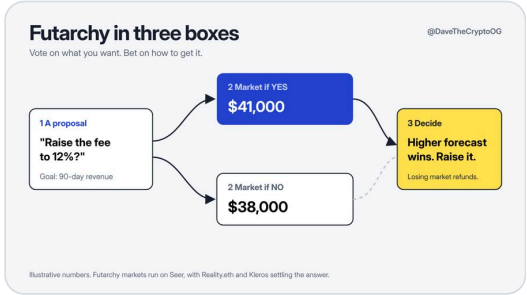
Post	Candidate	Posted (UTC)	House rank	Views, 12 Nov	Actual rank
D	Text: 50+ plans, one market, the team decides	10:10	1	515	**1**
A	Meme: prediction market vs decision market	10:22	2	51	5
E	Question: what should a market predict next?	10:38	3	72	3
C	Meme: reading "futarchy" for the first time	10:49	4	71	4
B	Explainer graphic: futarchy in three boxes	10:57	5	108	2



A



C



B

Result

Post D won, with 515 views against 108 for the runner-up Fact. Views were read from X at 11:16 UTC on 12 November, 44 minutes before the market's reading time; a gap of more than four to one leaves no doubt about the winner. The author reported the answer, Post D, on Reality.eth at about 12:10 UTC on 12 November; it becomes final if no one challenges it within Reality.eth's challenge period of about three and a half days.

The house forecast picked the winner and placed E and C correctly, third and fourth, but swapped A and B: the explainer graphic came second and the decision-markets meme came last. One influence is visible and disclosed: Seer's own account liked Post D and replied to it on 10 November, which put it in front of Seer's followers. That is the week's lesson for Foreguild. Who engages with a post moved it more than its format did, so Loop A scores each post against the client's own baseline, and the house model should learn from who responds, not only from what was posted. No liquidity was added, so the market served as a public commitment made before the posts went up, not as a price signal.

What this shows, and what it does not

It shows the whole weekly loop working with today's tools: candidates made by AI, a Seer market listed before publication, real posts, and an answer any reader can check on X. It does not show demand. One small account, five posts, no outside client and no liquidity are a working demonstration, not evidence for the revenue model in section 6.

Disclosure. Seer's mobile create form offered no opening-date field, so Reality.eth reporting opened when the market was created. The question fixes the reading time at 12:00 UTC on 12 November, and the author will challenge any answer reported before then.

Appendix C. Sources

1. Seer on X, FAO Operating Plan Contest, 5 November 2026, and Poidh bounty Base #1531 (55 claims read 9 November 2026). poidh.xyz/base/bounty/1531
2. Kleros, "16 Movies, 1 Judge, and a Prediction Market: Introducing the Kleros Foresight", 4 March 2026. blog.kleros.io
3. KuCoin News and The Block on X revoking API access for reward-for-posting apps, 15 to 16 January 2026. kucoin.com, theblock.co
4. BlockEden, on Kaito Studio's March 2026 beta with 16 brand partners. blockeden.xyz
5. Cointelegraph, Pump.fun bounty platform with 225 live bounties and $115,000 unclaimed, 5 June 2026. cointelegraph.com
6. Seer documentation, futarchy markets on Gnosis. seer-3.gitbook.io
7. Seer on X, 6 November 2026, pledge of $25,000 toward Open Risk Pricing for DeFi, built with Kleros and Credora.
8. Seer on X, 10 November 2026, "In the last three months, Seer has…" (list of contests and pledges).

9. Seer market depth as measured in Poidh claim 9480 (548 conditional markets on Gnosis, median traded volume about $7).
10. Census of 8 November 2026: Appendix A, one link per project.

Foreguild operating plan by Oliver Dave (@DaveTheCryptoOG). Independent, non-internal entry to Seer's FAO Operating Plan Contest. No affiliation with Seer, Kleros or any project listed.