

Scientific Vote

For every major argument in society, a published vote of the 300 most significant researchers in that field — identified by a reproducible ranking, verified by cryptography rather than by us, and recorded where no one can quietly edit it later.

FIRST PANELS each 300 seats

Climate

Quantum Computing

Nutrition

OPERATING MODEL

No permanent operator; guardian sunsets on published conditions

FUNDING

ETH staked in a Rethink Finance vault; yield only, principal untouched

PRESENCE

.eth name resolving to content-addressed storage

IN THIS PAPER

1–3 The idea, the gap it fills, and why now 7 Governance and the path to no operator 4 The four steps, in detail 8–9 What could go wrong, and what this is not 5–6 Architecture and funding 10–11 Roadmap and open questions

soonami4good — social impact ventures under the soonami umbrella September 2026

1 In one page

Every major public argument — climate, AI, energy, nutrition, pandemic response — is conducted with claims about what "scientists say." Almost none of those claims are checkable. Scientific Vote makes them checkable.

For a given field, we identify the 300 most significant researchers using a meta-ranking built from multiple public sources, published with its criteria, its code and its input versions, so that anyone can re-run it and get the same 300. Those researchers claim their seat through a cryptographic proof that the project's own founders cannot forge or bypass. They then vote on precisely worded statements, and the result is published in full: every name, every answer, every confidence level, every dissent, every abstention.

Then we do the next field.

The project is built to outlive its founders and to have no permanent operator. The website lives on a .eth name pointing at content addressed storage. The treasury is ETH staked in a Rethink Finance vault, so operations are funded by yield rather than by donors with opinions about the answers. Governance runs on a four-month cycle, the soonami4good pattern.

Three things this paper refuses to hand-wave

The panel selects the answer. Cook et al. (2016) showed that the headline climate-consensus figure moves from roughly 73% to 100% purely by redefining who counts as an expert.^[1] A project that picks 300 people is choosing its own result unless the picking is pre registered, reproducible and contestable. Section 4 does that work.

Publicly ranking named scientists has been tried, and it went badly. Anderegg et al. (2010) ranked climate researchers by publication

record in PNAS and was widely attacked as a blacklist; a published comment accused it of using "publication metrics as a surrogate for expertise" and of assessing the worth of researchers rather than the validity of their science.^[2] Any project that publishes 300 names inherits that fight on day one. Section 9 says how we intend to survive it.

The scientists themselves draw a line we must respect. In Nature's 2021 survey of IPCC authors, 66% said they personally engage in advocacy — and 74% said the IPCC itself should not.^[3] Scientific Vote therefore

asks empirical questions by default. Where a normative question is asked, it is labelled and reported separately.

THE HONEST ONE-LINE VERSION

This does not settle scientific questions. Counting people never has. It measures what the field's most significant researchers currently think, publishes it in a form nobody can quietly retouch, and lets the public argument proceed with that on the table instead of in dispute.

2 The problem

Expert and public opinion diverge sharply, durably, and in a way that publishing more studies does not fix. Sapienza and Zingales measured the gap between US economists and the US public across nineteen questions: an average of 35 percentage points, and — the uncomfortable part — the gap was widest where the economists were most unanimous.^[4] Nothing suggests climate is different.

The evidence we have about what scientists actually think is thin in a specific way. It is not that the studies are bad; several are excellent. It is that every one of them is a snapshot, run once, by someone with a budget and a deadline:

SCIENTIFIC VOTE · SOONAMI4GOOD CONCEPT PAPER V0.1 1

STUDY YEAR SAMPLE RESULT / RESPONSE RATE

Doran & Zimmerman ^[5] 2009	3,146 of >10,200 Earth scientists	97% of actively-publishing climate scientists; ~31% response
Anderegg et al. ^[2] 2010	1,372 researchers, ranked by output	97–98% of the 200 most-published; drew "blacklist" accusations
Verheggen et al. ^[6] 2014	1,868 completed of 6,550 contacted	84%, rising to 90% above 10 climate papers; 29% response
Myers et al. ^[7] 2021	2,780 of 10,929 invited	98.7% among 153 confirmed experts; ~25% response
Nature / Tollefson ^[3] 2021	92 of 233 living IPCC AR6 authors	60% expect $\geq 3^\circ\text{C}$; 4% expect 1.5°C to hold; ~40% response
Wynes et al. ^[8] 2024	211 of 909 IPCC AR6-cycle authors	Median expected warming 2.7°C ; 23% response

None of these is a standing institution. None is repeatable by anyone other than its authors. None outlives its funding. And in the years between them, "the science says" is asserted daily and verified almost never.

Meanwhile there is a working model for exactly this, and it has been running for nearly two decades — in economics.

3 What exists, and what doesn't

The IGM Forum at Chicago Booth (now the Kent A. Clark Center) polls a standing panel of senior economists roughly weekly. Each item is a statement to agree or disagree with; each panellist gives a five-point response, a 1–10 confidence score, and an optional comment; every individual named answer is published.^[9] It is the closest existing thing to what this paper proposes, and it works.

FIELD STANDING PANEL CADENCE NOTES

Economics IGM / Kent Clark Center^[9] Weekly ~43 US, ~49 European, ~40 finance; named responses, confidence-weighted

Philosophy PhilPapers Survey^[10] ~Decadal ~2,000 respondents; includes a metasurvey predicting its own result

AI AI Impacts / Grace et al.^[11] ~3-yearly 2,778 of 18,459 contacted — a 15% response rate on cold email
Quarterly Named institutions, probability distributions, microdata published

Macro forecasting ECB Survey of Professional Forecasters^[12]

Climate science **Nothing** — One-off surveys and literature counts only

What changed in June 2026

Carbon Brief launched Project Cosmos: a database of over 1.8 million climate publications linked by roughly 40 million citation relationships, covering about 1.6 million authors, with a published "Cosmos 500" ranking of top authors, publications and institutions.^[13] It is built on OpenAlex, anchored on IPCC report references, and updated annually.

This is important for two reasons. First, it demonstrates that a large, defensible, climate-specific ranking is now buildable — the hard bibliometric work has been done at scale by a credible organisation. Second, Carbon Brief has deliberately not open-sourced the database, citing scraping risk, and Cosmos is a ranking, not a vote.^[13]

SCIENTIFIC VOTE · SOONAMI4GOOD CONCEPT PAPER V0.1 2

So the white space is precise: the ranking layer now exists; the voting layer does not. Scientific Vote's contribution is the vote, plus a ranking that is reproducible by anyone rather than proprietary to one publisher.

4 How it works

STEP 1 Build the meta-ranking

Combine several public rankings into one ordered list, publish the method and the code before publishing the names, and freeze the result with an on-chain commitment.

STEP 2 Let the 300 claim their seats — without us as gatekeeper

Each scientist proves control of an address at their own institution using a zero-knowledge proof of their mail server's signature. We never hold a secret that would let us claim a seat on someone's behalf.

STEP 3 Run the vote and publish everything

Precisely worded statements, five-point responses with confidence, ballots encrypted until close, results published both unweighted and confidence-weighted, with every named answer and comment.

STEP 4 Start the next field

The protocol, the code and the claiming mechanism carry over. Only the rankings and the institutional domain list are rebuilt.

Step 1 — The meta-ranking

Six candidate inputs, all public, with very different properties:

SOURCE LATEST COVERAGE DATA ACCESS CHARACTER

Cosmos 500^[13] Jun 2026 1.6M authors ranked Ranking public, DB Most climate-specific; ~50% US, 10% women closed

Clarivate HCR^[14] Nov 2025 6,868

researchers, 21 fields
Free to view, no bulk file

Includes integrity screening and retraction exclusions

Ioannidis top-2%^[15] Aug 2025 ~200,000 scholars Downloadable, CC-BY Six-component composite; self-citation variants shipped

OpenAlex^[16] Monthly Whole corpus CC0 bulk snapshot The only fully recomputable base

Research.com^[17] Jan 2026 9,674 environmental scientists

View only Discipline h-index threshold D≥30

IPCC author rosters^[18] AR7, 2025 664 experts (AR6: ~782)

Not bibliometric — peer nominated and balanced

Aggregation: not a rank-sum

Web database, scrapeable

The obvious method — rank each scientist in each list, add the ranks, sort — is the one Reuters used for its 2021 Hot List, and it is the weakest option available.^[19] Dwork et al. showed that no positional method, Borda included, can satisfy the Extended Condorcet Criterion, and that ECC-satisfying aggregators have a property we specifically need: if an entity manipulates fewer than half the input rankings, it gets pushed to the bottom

SCIENTIFIC VOTE · SOONAMI4GOOD CONCEPT PAPER V0.1 3

rather than up.^[20] Our inputs are citation-based and therefore gameable — self-citation, citation cartels and hyper-prolific publishing are all documented.^[21]

Proposed method: Reciprocal Rank Fusion (score = $\sum 1/(60 + \text{rank})$), which needs no tuning and no calibrated scores^[22] as the base, followed by local Kemenization to enforce the Condorcet property at $O(kn \log n)$.^[20] Both are decades-old, well-understood, and implementable in a few hundred lines.

The skew problem, stated before someone else states it

KNOWN DEFECT OF EVERY CITATION RANKING

Cosmos: roughly half of ranked authors US-based, 4% from the global south, 10% women, over 97% English-language — figures Carbon Brief publishes itself.^[13] Reuters 2021: 122 women out of 1,000, and five African scientists.^[19] By contrast, IPCC AR6 Working Group II was 41% women, and the AR7 author cohort is 46% women and 51% from developing countries or economies in transition.^[18]

A citation-only meta-ranking inherits and compounds this skew, and reviewers will say so — correctly. Two design consequences follow. First, the IPCC rosters enter as a co-equal input, not a footnote: they are the one large, credible, non-bibliometric roster of climate expertise, and they are dramatically more balanced. Second, the panel's own demographic composition is published alongside every result, so readers can discount it themselves rather than discovering it later.

Reproducibility as the substitute for authority

We cannot ask anyone to trust our judgement about who belongs in the top 300, and we should not want them to. Instead: the criteria are published before the list; the code is open; every input is pinned to a dated version; and the whole thing is committed to a Merkle root on-chain so the list cannot be edited after the fact without it being obvious. The target is not "trustless" — it is reproducible from stated criteria. Anyone can re-run it. Anyone who disagrees can publish a competing ranking using the same inputs and argue about the method in public, which is the correct venue.

Step 2 — Claiming a seat, without a gatekeeper

This is the step where the project either earns its credibility or loses it, and the research turned up something worth stating plainly: the requirement splits into two questions, and only one of them is a cryptography problem.

Minimise it by making it:

- Pre-registered — criteria before names ·
- Reproducible — open code, pinned inputs ·
- Committed — Merkle root, tamper-evident ·
- Contestable — the 300 can challenge it ·

T1 · WHO ARE THE 300?

Irreducibly an editorial judgement.
No cryptography removes it.

Declinable — anyone may opt out

T2 · WHICH PERSON HOLDS WHICH SEAT?

This is where cryptography actually helps — and where “cannot be hijacked” lives.

Trust root by mechanism:

· Emailed claim code → us. Hijackable. ·
ORCID OAuth → ORCID Inc. + our server. ·
DKIM-in-ZK → the university's mail server.
We never hold a secret. Proof is public. ·
eduGAIN → each institution's own IdP.

The trust split. No design eliminates T1; the honest move is to make it reproducible and contestable rather than to claim it away. T2 is genuinely solvable, and the choice of mechanism decides whether the project's own operators could seize seats.

SCIENTIFIC VOTE · SOONAMI4GOOD CONCEPT PAPER V0.1 4

The recommended mechanism: DKIM-in-ZK

Every mail server signs its outgoing mail with a private key whose public half is published in DNS. ZK Email verifies that signature inside a zero-knowledge proof, revealing only what you choose to reveal — and it can prove a specific address, not merely a domain.^[23] A scientist claims their seat by sending one email from their institutional address and generating a proof from it.

What makes this the right answer is narrow but decisive: the operator never learns a secret, and the proof is verifiable by anyone. Compare the alternatives honestly. With an emailed claim code, whoever sends the code knows the code — nothing stops the operator from claiming all 300 seats and telling each scientist their link is on the way. With ORCID, the OAuth token passes through our server, and ORCID Inc. can revoke or alter records; worse, an analysis of ORCID's 2024 public data file found that only about 0.9% of affiliations are institution-asserted — the rest are self-declared.^[24]

THE BLOCKER WE MUST TEST BEFORE COMMITTING

Microsoft 365 does not DKIM-sign custom domains unless an administrator configures it — otherwise mail is signed by the tenant's onmicrosoft.com domain. Google Workspace without custom DKIM signs as gapssmtp.com.^[25] A meaningful share of the 150-odd institutions in the panel may therefore be unprovable as themselves. Action: survey the DKIM signing domain, selector and key size for every institutional domain in the list before design freeze. Where coverage fails, fall back to eduGAIN federated identity — fifteen years old, run by the universities themselves, centralised per institution but distributed across them, which is a far better trust shape than one operator.^[26]

The surrounding protocol

Commit before opening. The Merkle root of the eligible list is published before claiming begins; names public, contact identifiers committed as salted hashes.

Challenge window. Claims are visible as they happen. Any of the 300 — and any observer — can contest a claim during a fixed window.

Bounded corrections. A correction may replace at most k leaves, may never push the count above 300, and is subject to a timelock long enough for third parties to diff the old and new trees. Contested corrections escalate to an external adjudicator such as Kleros or UMA — which replaces our discretion with someone else's, an improvement in contestability rather than an elimination of trust.^[27]

Wallets that a professor will actually use. A passkey-signed smart account: Touch ID or Windows Hello, no seed phrase, with ZK Email Recovery as the guardian — which reuses the same address they already proved.^[23]

AN HONEST LIMIT

Any recovery mechanism is also a delegation mechanism. We cannot both guarantee “no seat is ever lost to a forgotten key” and “only the named scientist ever voted.” A postdoc can be handed the laptop. The response is not to pretend otherwise: state the delegation policy, require a fresh signature per vote so standing delegation costs something, and publish the rules so readers can discount accordingly.

Step 3 — The vote

Most of the design here is borrowed, deliberately, from institutions that have already survived their own mistakes.

SCIENTIFIC VOTE · SOONAMI4GOOD CONCEPT PAPER V0.1 5

DESIGN CHOICE BORROWED FROM WHY

Statements, not questions; five-point scale plus "no opinion"

IGM^[9] Forces precision and makes disagreement legible

Self-reported confidence, 1–10 IGM^[9] Separates "the field is split" from "the field is unsure"

Standardised effect-size words: "substantial", "measurable" mode IGM^[9] Makes silence visible; blocks cherry-picking

Publish unweighted *and* confidence weighted, with non-respondents in the unweighted denominator

IGM, ECB SPF^{[9][12]} Dissent is the product, not the noise

Every named response and comment published

IGM, adopted 2022^[9] Vague wording was IGM's observed failure

Ballots encrypted until close Shutter shielded voting^[28] No running tally means no herding to the leader

Panel ratifies wording before a vote opens

scandal PhilPapers 2009^[10] Cheap, and it measures whether a

Optional metasurvey: predict the panel's own result

field knows itself

The empirical / normative line

IGM's 2019 MMT episode^[9] There, the wording itself became the

Scientific Vote asks, by default, questions of the form what is the case and what will happen — not what should be done. This is not squeamishness. It is the line the intended panel draws for itself: 74% of IPCC authors surveyed by Nature said the IPCC should not take an advocacy role, while 66% of the same people advocate personally.^[3] A panel that believes its institutional voice should stay empirical will not lend that voice to a policy ballot, and asking it to is the fastest way to lose the panel.

Where a normative question is genuinely worth asking, it is labelled as such, reported in a separate section, and never blended into an empirical headline number.

Step 4 — The next field

What carries over: the aggregation code, the claiming protocol, the ballot design, the publication format, the treasury and the governance process. What must be rebuilt each time: the input rankings for that field, the institutional domain list and its DKIM survey, and the panel's ratification of the first questions. The marginal cost of field number two should be a small fraction of field number one — which is the entire argument for building this as infrastructure rather than as a one-off climate survey.

Candidate criteria for choosing it: a live public argument, a field with usable public rankings, and a question set that is genuinely empirical.

5 Architecture

Presence — a .eth name over content-addressed storage

An ENS name stores a contenthash record pointing at content-addressed storage — IPFS, Arweave or Swarm — and gateways such as eth.limo resolve it over ordinary HTTPS.^[29] Two honest caveats belong in this paper rather than in a later post-mortem. Gateways are a real chokepoint: native browser support for ENS content resolution is weak and uneven, so in practice most readers arrive through a gateway operator who could be pressured. And somebody holds the key that can rewrite the contenthash — that key is the true control point of a "decentralised" website. Ours will be held by the DAO, with the ENS Name Wrapper fuses available as a

SCIENTIFIC VOTE · SOONAMI4GOOD CONCEPT PAPER V0.1 6

hardening step. The trade-off is stark and worth stating: burning CANNOT_SET_RESOLVER makes the site permanently unrewritable, including by us — and also permanently un-updatable. Fuses also reset at expiry, so renewal remains a live operational dependency.^[30]

For durability rather than availability, the canonical record of each vote is written to Arweave, whose one-time payment endowment model removes the recurring-payment single point of failure that IPFS pinning creates.^[31]

CORRECTION TO A COMMON ASSUMPTION

ENS's Namechain L2 was cancelled in February 2026; ENSv2 is deploying on Ethereum L1 only, on the reasoning that registration gas costs fell roughly 99% over the preceding year.^[32] Plans should not assume an ENS L2.

Treasury — a Rethink Finance vault

Rethink Finance is vault infrastructure for on-chain asset management rather than a fund. Each vault holds assets in a Safe smart account whose owner is an OpenZeppelin Governor contract, with a Zodiac Roles Modifier scoping what any executor may do.^[33] That shape matters here: it means the treasury's permitted actions are enumerated in code, and depositors can exit without the curator's cooperation.

soonami already runs on this: the Venture Staking vault takes WETH with a 32 ETH minimum and routes ETH yield into ventures.^[34] Scientific Vote uses the same pattern, which means the first version is a configuration exercise, not an engineering project.

DISCLOSURES WE SHOULD MAKE OURSELVES

Rethink is currently in closed beta with protocol fees switched off; the announced future fee is 0.15% of AUM annually plus 10% of curator fees.^[33] Its Fund and NAV contracts were audited by AuditOne in January–February 2024 — before the 2026 redesign — with one medium and three low findings recorded unresolved.^[35] And AuditOne is itself a soonami portfolio company, alongside Rethink.^[36] None of this is disqualifying. All of it is the sort of thing that reads far worse when someone else finds it first.

Chain and voting stack

Governance and treasury sit on Ethereum L1. This is now an affordable choice rather than a purist one: median mainnet transaction fees fell to roughly \$0.012 in Q1 2026 as blob capacity rose, and ENS Labs cancelled its own L2 on exactly this reasoning.^{[32][37]} L1 also avoids adding sequencer and upgrade-key trust to a project whose whole thesis is the absence of a privileged operator.

Voting uses Snapshot with a Merkle allowlist strategy for signalling — gasless for the scientist, which matters — and Safe with Zodiac optimistic execution for anything binding.^[38] Where anonymity is wanted, Semaphore lets a panellist prove membership of the 300 and vote once without revealing which member they are.^[39]

DEPENDENCY RISK, CURRENT AS OF THIS WRITING

The Ethereum Foundation's Privacy and Scaling Explorations group — which built both Semaphore and

MACI — was disbanded in June 2026, and MACI was archived on 19 August 2026.^[40] Do not build on MACI. Semaphore remains the correct primitive but now carries maintenance risk: vendor the contracts and circuits, pin versions, and budget to maintain them in-house. This argues generally for boring, widely-forked dependencies over clever ones.

6 Funding: yield, not donors

The point of funding operations from staking yield is not novelty. It is that a project funded by donors can be defunded by donors, and a panel whose results embarrass someone will eventually meet that pressure. An endowment that is never spent — only its yield — removes that lever permanently.

2.63%

Native ETH staking APR, 31 Aug 2026^[41]

~2.24%

Lido stETH net APR, same date, after the 10% protocol fee^[41]

~25 ETH

Annual yield on a 1,000 ETH endowment at ~2.5% net, principal untouched

Yields have compressed from over 4% in 2023, and this is structural rather than cyclical: Ethereum issuance scales roughly with the inverse square root of total ETH staked, and the staked share has passed 35%.^[41] MEV adds somewhere in the region of 10–30% on top of the base rate.^[42] Plan on the low number.

What ~25 ETH per year comfortably covers: ENS registration and renewal (a five-character name renews for about \$5 annually^[30]), Arweave permanence and IPFS pinning, gas for the governance cycle and vote execution, an annual rebuild of the meta-ranking, a security bounty pool, and bounties for the small amount of contract and circuit maintenance the project needs.

What it deliberately does not cover: salaried staff. That is the constraint that makes "no permanent operator" financially true rather than aspirational — if the project needs employees to survive, it has an operator by definition.

The honest risk: the budget is denominated in ETH and the yield is variable. A prolonged drawdown halves the operating budget in fiat terms while the work stays the same size. The mitigations are unglamorous — hold a fiat-equivalent buffer, keep the fixed cost base genuinely tiny, and size the endowment for the bear case rather than the bull.

The soonami4good commitment applies: over 95% of generated funds go to purpose, with full transparency of the flows on-chain.

7 Governance, and the path to no operator

Here is the position this paper takes, and it is deliberately less romantic than the brief: "fully decentralised" is an end state reached on published conditions, not a launch condition. The evidence from 2026 is not kind to the alternative.

WHAT HAPPENED DETAIL^[43]

Scroll DAO Fully suspended operations in January 2026 amid uncertainty over which proposals were even live

Jupiter Froze all governance votes and locked fund access, first to 2026, then to 2027 **Yuga Labs** Abandoned its DAO structure outright, citing inefficiency

Compound Drifted into what participants describe as a club of delegates

The sector Average turnout ~20%, often 10%; 1% of holders control 90% of voting rights; a 2025 Cornell analysis found 3–5 voters sufficient to sway most Compound and Uniswap proposals

Autonomy claimed on day one has a poor record. Autonomy reached on a schedule, with the powers being given up enumerated in advance, has a better one.

The proposed shape

A guardian multisig with strictly enumerated, monotonically decreasing powers. It can pause; it cannot direct a vote, alter a result, or move the endowment's principal. Every power it holds has a published condition under which it is surrendered.

Sunset tied to milestones, not dates: the first vote completed and published; the claiming protocol surviving one adversarial review; the ranking rebuilt reproducibly by a third party; treasury operations running for a full governance cycle without intervention.

Seven-day timelocks on everything, so that any action is visible before it is final.

Panel veto over guardian signers — the 300 can remove a guardian; the guardian cannot remove one of the 300.

A four-month governance cycle, carried over from soonami4good Education.

The legal wrapper

Some things genuinely cannot go on-chain: signing a contract, holding a domain registration, appearing in court, paying an auditor in fiat. The realistic options are a Swiss association (the structure Rethink itself uses^[33]), a Cayman foundation, or a Wyoming DUNA, which since July 2024 gives a decentralised association legal existence, the ability to contract and be sued, and limited liability between members.^[45] One trap worth

naming: a Wyoming DAO LLC auto-dissolves if no proposal is approved for a year — a live hazard for a project that deliberately votes infrequently.^[46] This paper does not choose between them; it flags that the choice needs qualified counsel, which is also an open issue in the sibling soonami4good Education project.

8 What could go wrong

Ordered by how likely they are to actually kill the project — which is not the order most people would

guess. SCIENTIFIC VOTE · SOONAMI4GOOD CONCEPT PAPER V0.1 9

RISK WHY IT IS REAL RESPONSE

Nobody votes This is the most likely failure. Expert survey response rates run 15–40%.^{[3][7][8][11]} In Optimism's RetroPGF 3 — a curated panel of 145 named people — the median participant spent **16 hours** and described it as far too much work.^[47]

"It's a blacklist" Anderregg et al. (2010) is the precedent and it was Our choice of 300 is therefore a substantive

ugly. The published critique: publication metrics used as a surrogate for expertise, and researchers assessed rather than their science.^[2]

The selection picks the answer Redefining the expert tier moves the climate consensus figure from ~73% to 100%.^[1]

claim, not a technicality.

Two-minute ballots. Low frequency at first. A standing committed panel rather than cold email. Publish the response rate every single time, prominently.

as non-participation, never as dissent. No one is labelled, ranked publicly by name, or scored.

Pre-registered criteria, open code, multiple co equal inputs including non-bibliometric ones, and the panel's composition published beside every result.

The list is an invitation, never a judgment. Anyone may decline and be removed, with no inference drawn. Non-participation is published

Question wording becomes the story
IGM's March 2019 MMT poll was publicly

called a dishonest representation of the position it purported to test.^[9]

Wording drafted in the open and ratified by the panel before the vote opens. Standardised effect-size language.

Harassment of participants The panel is named and public, and climate researchers are already targets. Publishing a list makes it easier, not harder.

audit of 55 ballots found no others.^[47]

Identity publication is opt-in at claim time. Pseudonymous participation available via Semaphore — provably one of the 300 without revealing which.^[39] No obligation to comment. A stated non-escalation policy.

Lobbying and capture A published vote of 300 named people is worth money to several parties. RetroPGF 3 recorded one enforced self-dealing violation among its 133 voters; a separate random

Encrypted ballots until close, so there is no running tally to lean on.^[28] Conflict-of-interest disclosure at claim time. Random ballot audits, as Optimism did.

Confusing "what will happen" with "what should be done"

74% of IPCC authors say their institution should not advocate.^[3] Blur the line and the panel walks.

Empirical by default. Normative items labelled, separated, and never folded into a headline figure.

Identity capture at claim time The Worldcoin precedent: credentials sold onward for as little as \$30 because the human at the enrolment desk was the weak point.^[48]

DKIM-in-ZK removes the enrolment desk entirely. Public claims plus a challenge window mean a wrong claim is visible to the 299 others.

Timelocks on everything, guardian pause during sunset, minimal privileged surface, and a treasury whose permitted actions are enumerated by the Roles Modifier.

Governance exploit \$8.5M lost by Term Finance in August 2026.^[44]

Vendor and pin everything. Prefer widely forked, boring dependencies. Budget for in house maintenance rather than assuming upstream survives.

Dependencies rot MACI archived, PSE disbanded, Allo Protocol in maintenance mode.^[40]

9 What Scientific Vote is not

It is not a replacement for peer review, and not a way of settling scientific questions. A majority among 300 people does not make a claim true, and the paper should say that in its own voice rather than wait to be told. What a vote measures is what the field's most significant researchers currently think — a narrower and more

SCIENTIFIC VOTE · SOONAMI4GOOD CONCEPT PAPER V0.1 10

modest thing than what is true, and still worth knowing precisely because it is currently asserted without evidence.

It is not a consensus-manufacturing machine. If the panel splits, the split is the result, and it is published exactly as prominently as agreement would be. A project that only produces headlines when scientists agree has selected on its outcome. There is also a real body of evidence that simply telling the public about a consensus shifts opinion less than advocates hope — Sapienza and Zingales found minimal movement when the public was informed of the economists' view,^[4] and the "gateway belief" literature is actively contested.^[49] This project is a record, not a persuasion campaign.

It is not a policy body, for the reasons in Section 4. And it is not a claim that the 300 are the only people worth

hearing. The ranking is a defensible, reproducible way to pick a panel of a workable size. It is not a statement about anybody's worth, and every design choice above is aimed at keeping that distinction visible.

10 Roadmap

PHASE WORK GATE TO PASS WHY THIS GATE

0 Feasibility: DKIM survey of ~150 institutional domains; ranking prototype on open inputs; approach Carbon Brief about Cosmos; scope the legal wrapper	control to the DAO DKIM coverage above an agreed threshold, or eduGAIN confirmed as a workable fallback	one empirical fact, and it is cheap to check
1 Publish criteria and code, then the 300; commit the Merkle root; open the challenge window	A third party reproduces the list independently	If it is not reproducible, it is just our opinion with extra steps
2 Claiming opens; wallets, recovery, support for non-crypto users	A pre-committed claim rate — we suggest 40%	Below that, the vote is not representative and should not be published as if it were
3 First vote: three to five empirical statements on climate, ratified by the panel	Result published in full, with response rate and composition	This is the actual product; everything before it is scaffolding
4 Sunset the guardian against the published conditions; ENS	One full governance cycle with no guardian intervention The whole anti-hijack design rests on this	Decentralisation demonstrated rather than announced
5 Second field Marginal cost materially below field one	community If it is not cheaper the second time, this is a survey, not infrastructure	

11 Open questions for the

Should the panel be 300 by citation rank, or 300 selected for balance?

IGM — the one comparable institution that works — picks by committee for geographic, generational, ideological and disciplinary balance, explicitly not by metric.^[9] That is real evidence against the founding premise of this paper, and it deserves a real answer rather than a defence.

What is the minimum claim rate at which a result may be published at all?

It must be fixed in advance, in public, or it will be negotiated downward after the fact.

SCIENTIFIC VOTE · SOONAMI4GOOD CONCEPT PAPER V0.1 11

Named votes or pseudonymous ones?

Named responses are what give IGM its authority. Pseudonymity is what protects a researcher from harassment. The panel should probably decide this, not us.

Who may propose questions?

Anyone, with the panel ratifying? The panel only? A rotating committee? This is the most attackable surface in the whole design.

Do we partner with Carbon Brief, or route around them?

Cosmos is the best climate ranking in existence and its data is deliberately closed.^[13] A published ranking can still be used as one input among several without access to the underlying database — but a partnership would be better for everyone.

What happens the first time a result is politically inconvenient?

Not a hypothetical. The answer needs to be structural — in the timelocks, the immutability of published records and the sunset conditions — rather than a matter of anyone's good character.

12 Sources

- [1] Cook et al., "Consensus on consensus," Environmental Research Letters
- [2] Anderegg, Prall, Harold & Schneider, PNAS 107(27):12107, 2010; and Bodenstein, PNAS 108(2):E1, 2011.
- [3] Tollefson, "Top climate scientists are sceptical that nations will rein in global warming," Nature 599:22–24, 1 Nov 2021.
- [4] Sapienza & Zingales, "Economic Experts versus Average Americans," American Economic Review P&P 103(3):636–642, 2013.
- [5] Doran & Zimmerman, "Examining the Scientific Consensus on Climate Change," Eos 90(3):22–23, 2009.
- [6] Verheggen et al., Environmental Science & Technology 48(15):8963–8971, 2014.
- [7] Myers et al., Environmental Research Letters 16(10):104030, 2021.
- [8] Wynes et al., Communications Earth & Environment 5:498, 2024.
- [9] Kent A. Clark Center for Global Markets (IGM Forum), kentclarkcenter.org — panel, methodology and critique history.
- [10] Bourget & Chalmers, "Philosophers on Philosophy: The 2020 PhilPapers Survey," Philosophers' Imprint 23(11), 2023.
- [11] Grace et al., "Thousands of AI Authors on the Future of AI," arXiv:2401.02843, 2024.
- [12] European Central Bank, Survey of Professional Forecasters, quarterly since 1999.
- [13] Carbon Brief, "Project Cosmos," launched 22 June 2026 —
- [25] Microsoft Learn, DKIM configuration for Microsoft 365; Google Workspace default
- [14] Clarivate, Highly Cited Researchers 2025, published 12 Nov 2025.
- [15] Ioannidis et al., PLOS Biology 18(10):e3000918, 2020; data v8, Elsevier Data Repository, Aug 2025.
- [16] OpenAlex — CC0 monthly snapshot; docs at help.openalex.org.
- [17] Research.com, "Best Scientists in Environmental Sciences," 5th ed., Jan 2026.
- [18] IPCC AR6 WGI/WGII/WGIII author pages; and IPCC, "AR7 authors selected," 18 Aug 2025.
- [19] Reuters Hot List methodology, April 2021; criticism in The Conversation, "Reuters' Hot List is geographically skewed," and Forbes, 15 Oct 2021.
- [20] Dwork, Kumar, Naor & Sivakumar, "Rank Aggregation Methods for the Web," WWW 2001.
- [21] Fire & Guestrin, GigaScience 8(6), 2019; Hicks et al., "The Leiden Manifesto," Nature 520:429–431, 2015.
- [22] Cormack, Clarke & Büttcher, "Reciprocal Rank Fusion," SIGIR 2009.
- [23] ZK Email — docs.zk.email; zkemail.nr; email-recovery module.
- [24] SCImago Lab, "ORCID Report III: Institutional Trust Markers," Feb 2026 — 0.89% of affiliations institution-asserted.
- [26] eduGAIN / GÉANT — interfederation of 11(4):048002, 2016. interactive.carbonbrief.org/cosmos/methodology. gapssmtp.com signing.
- national research and education identity federations.
- [27] Kleros Court and UMA Optimistic Oracle V3 — optimistic dispute resolution.
- [28] Shutter Network shielded voting on Snapshot — 881+ DAOs, 372,914+ encrypted votes.
- [29] ENS, "Decentralized websites"; eth.limo gateway documentation.
- [30] ENS pricing and Name Wrapper fuses — docs.ens.domains.
- [31] Arweave endowment model — ar.io.
- [32] The Block, "ENS Labs scraps Namechain L2, shifts ENSv2 fully to Ethereum mainnet," 6 Feb 2026.
- [33] Rethink Finance — rethink.finance and docs.rethink.finance.
- [34] soonami Venture Staking — soonami.io/venture-staking.
- [35] AuditOne, Rethink Finance audit report, Jan–Feb 2024.
- [36] soonami ventures portfolio — soonami.io/ventures.
- [37] Median L1 and L2 transaction fee data, Q1 2026, arXiv:2606.22206.
- [38] Snapshot documentation; Zodiac Reality / oSnap optimistic execution modules.
- [39] Semaphore Protocol — semaphore.pse.dev.
- SCIENTIFIC VOTE · SOONAMI4GOOD CONCEPT PAPER V0.1 12
- [40] MACI repository archived 19 Aug 2026; Ethereum Foundation restructuring and PSE dissolution, June 2026.
- [41] StakingRewards and Lido, figures as at 31 Aug 2026. Volatile — re-check before publication.
- [42] MEV-Boost contribution to validator returns, 2026 estimates.
- [43] ForkLog, "How DAOs are faring in 2026," incl. Cornell 2025 governance analysis.
- [44] The Block, "Term Finance loses an estimated \$8.5 million to governance exploit," 23 Aug 2026.
- [45] a16z crypto, "The DUNA for DAOs" — Wyoming SF50, effective 1 July 2024.
- [46] Wyoming DAO LLC Supplement, March 2021 — one-year dissolution provision.
- [47] Optimism, "RetroPGF 3 learnings and reflections"; and RetroPGF 3 conflict-of interest thread, gov.optimism.io.
- [48] The Block, "Worldcoin iris black market," 2023.
- [49] van der Linden et al., PLOS ONE, 2015 (gateway belief model); and the contested reanalyses in JCOM 16(05), 2017.

Figures dated 2026 — staking yields, fee levels, protocol status — were current at the end of August 2026 and several are volatile by nature. Items flagged during research as unverified have been either omitted or stated with their uncertainty. Before this paper is circulated externally, re-check: DKIM coverage across the institutional domain list; current staking yields; the licence terms of the Ioannidis dataset; and the maintenance status of Semaphore.

