Less Than One Thought in a Billion is About Breakthrough-Level Cosmology

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

December 28, 2025

This paper proposes a deliberately provocative, audit-friendly claim: when we restrict attention to breakthrough-level, new-to-world, field-shaping cosmology, the fraction of total human cognition devoted to such work is plausibly less than 1 part per billion (ppb). The point is not to diminish cosmology, nor to glorify it as a priesthood, but to clarify how extreme the selection filters become when we combine (i) a narrow domain (cosmology), (ii) a stringent novelty threshold (field-shaping breakthroughs rather than incremental progress), and (iii) a population-wide denominator (all human cognition across the planet). We treat "fraction of thoughts" as a proxy for fraction of cognitive time, because thoughts are not countable and time is. Using order-of-magnitude reasoning, we build a simple funnel model: global population -> cosmology-capable workforce -> active cosmology contributors -> breakthrough-producing subset -> breakthrough-mode time share. Even generous parameter choices push the final fraction below ppb, and stricter definitions push it far lower. We close by proposing empirical anchors and falsifiable bounds, and by reflecting on what such rarity implies for epistemic humility, funding narratives, and public understanding of "frontier knowledge."

Keywords: cosmology, breakthrough cognition, order-of-magnitude reasoning, attention economics, scientific creativity, research workforce, cognitive time budget, novelty thresholds, heavy-tailed innovation, epistemic humility, bibliometrics, science of science.

A collaboration with GPT-5.2-Thinking. 54 pages.

Outline

1. The Claim and Why It Sounds Wrong

1.1 The headline claim: sub-ppb cognition share

The headline claim is intentionally sharp: if we restrict attention to "breakthrough-level / new-to-world / field-shaping cosmology," then the fraction of total human cognition devoted to that target may be less than 1 part per billion (1e-9). The claim sounds absurd because it collides with two everyday intuitions that feel incompatible: first, cosmology is culturally loud (it appears constantly in documentaries, headlines, and "big questions" discourse); second, breakthroughs feel psychologically huge (they reorganize narratives, not merely datasets). How could something that seems so prominent and so consequential occupy such a vanishingly small share of what humanity is thinking?

The core move is a denominator shift. We are not estimating the fraction of scientists doing cosmology, nor the fraction of academic work hours spent on cosmology. We are dividing by the cognitive life of the entire species: billions of people, most awake hours, most daily concerns, most nontechnical attention, most entertainment, most labor, most caregiving, most routine. Under that denominator, even large global industries become small fractions; under this denominator, "breakthrough-level cosmology" becomes smaller still because it adds two more filters: a narrow domain boundary and a stringent novelty threshold.

Because "thoughts" are not countable in any clean way, the paper uses a proxy: fraction_of_thoughts ~= fraction_of_cognitive_time. This is not a metaphysical claim about the nature of thought. It is a practical measurement stance: time budgets are at least definable and, in principle, measurable. The sub-ppb claim is therefore best read as a time-share claim: if you sum across humanity the time spent in genuine breakthrough-mode cosmological cognition, and divide by total human awake cognition time, you plausibly land below 1e-9.

1.2 Why "about cosmology" is a stricter filter than it looks

Most people hear "about cosmology" and imagine any mental contact with the cosmos: looking at stars, wondering about origins, watching a space video, feeling awe, repeating a popular phrase about black holes, or having a late-night metaphysical tremor. But those are not the same as cognition that is about cosmology in the disciplinary sense. The stricter filter used here is: thoughts that participate in cosmology as a knowledge-making practice rather than cosmology as a cultural aesthetic.

That difference matters because cosmology is not merely "thinking about the universe." Cosmology is a constraint-dense activity: it is thinking under specific empirical, mathematical, and methodological commitments. It lives inside instrument pipelines, inference methods, priors and systematics, calibration choices, consistency checks, internal debates about tensions, and continuous negotiation between theory, simulation, and observation. In this sense, much of what the public calls "cosmology" is better described as cosmology-flavored narrative. Narrative can be sincere and meaningful; it can even be philosophically deep. But it is not typically cosmology as a field, because it does not run on the same error model.

This is why the "about cosmology" filter is severe. It excludes not only the majority of human cognition, but even the majority of "space-related" cognition. It excludes most astronomy hobbyism, most sci-fi, most popular science consumption, and much of what scientists do when they are teaching, administrating, or writing grants. It narrows the candidate pool toward those engaging the actual moving boundary where cosmological claims are made, contested, refined, and sometimes overturned.

1.3 Why "breakthrough-level" changes everything

Even within professional cosmology, most cognition is not breakthrough cognition. The field advances through an ecology of increments: debugging pipelines, checking systematics, refining likelihoods, comparing models, running simulations, extending measurements, improving instruments, reconciling calibration discrepancies, writing methods sections, reproducing results, responding to referee reports, and maintaining collaborative infrastructure. This work is not a lesser form of thought; it is the work that makes thought honest. But it is not what we mean by "breakthrough-level / field-shaping."

The breakthrough filter is qualitatively different from the "doing the field" filter. It demands novelty that is new-to-world and field-shaping rather than locally clever, personally new, or publishable-but-incremental. A breakthrough, in the strict sense used here, changes what future competent practitioners must take into account. It alters default models, standard interpretations, canonical tensions, or the set of admissible explanations. It does not merely add another paper; it changes what counts as a good question.

This makes the distribution heavy-tailed. In a heavy-tailed regime, most participants will have near-zero probability of producing a field-shaping shift in any given year, and a small number of people or teams will account for a large share of the field's direction changes. Moreover, "breakthrough-mode" time is often a small fraction even for the most capable: the deeper the constraints, the more time is consumed by verification, calibration, and disciplined doubt. If cosmology has any defining temperament, it is the price it pays for claims: you do not get to say "the universe is like this" without paying the tax of error analysis.

So the breakthrough filter drives the time-share down twice: fewer people qualify, and even among those people, the fraction of time that is genuinely in breakthrough-mode is small. It is entirely coherent, under this strict meaning, that the final share lands below 1e-9 without implying anything mystical about the field.

1.4 What the claim is not saying: worth, dignity, or "who matters"

A claim about fractions can easily be misheard as a claim about value. This paper rejects that slide. The denominator is not a moral ranking. It is a scaling choice. Saying "sub-ppb of cognition is breakthrough cosmology" is not the same as saying "cosmology matters more than ordinary life" or "ordinary life is trivial." On the contrary: the very reason the fraction is tiny is that human life is full. The species is busy surviving, loving, working, raising children, tending elders, cooking meals, repairing roofs, grieving losses, laughing, arguing, praying, and enduring. A civilization in which everyone is thinking about cosmology all day would be pathological or imaginary.

Nor is the claim saying that only credentialed academics "count" as thinkers. If a noncredentialed person truly produces a new-to-world, field-shaping cosmological insight that survives contact with the field's constraint structure, then by definition their cognition belongs in the numerator. The claim is about how rare that numerator is, not about who is allowed into it.

Finally, the claim is not a call to mystify cosmology as a priesthood or to treat cosmologists as oracles. Rarity can produce unhealthy social effects: the temptation to defer, to mythologize, to outsource judgment, to replace method with charisma. The paper takes the opposite stance: if the numerator is tiny, the duty of transparency grows. Small numerators invite storytelling; disciplined fields must answer with audit. The point of the claim is to locate the work on the scale of humanity without turning scale into sanctification.

In short: the claim sounds wrong because we confuse cultural volume with cognitive share, and because we confuse consequence with frequency. Once the filters are made explicit, "sub-ppb" stops being an insult or a cultic boast and becomes a plausible scaling fact about a narrow, constraint-heavy, breakthrough-defined frontier.

2. Definitions and Units: What Counts as a "Thought" Here

2.1 Thoughts are uncountable; time is countable

If we try to literally count "thoughts," we immediately hit a measurement wall. A "thought" does not have a stable physical boundary. Is a thought one sentence? One intention? One image? One associative chain? A ten-minute reverie? A sudden insight that recurs for weeks? Even within one person, the granularity changes with fatigue, emotion, task, and attention. Across a population, the concept becomes even less well-defined. People vary in inner speech, imagery, attentional fragmentation, and the degree to which cognition is explicit versus tacit. Any attempt to count thoughts would therefore smuggle in arbitrary segmentation rules that would dominate the result.

Time, by contrast, is countable in a practical sense. We can define windows of activity. We can talk about awake hours per day and work hours per year. We can estimate the fraction of a day spent in a particular kind of cognitive mode. Time is not perfect, but it is tractable. The paper therefore adopts a deliberately prosaic stance: when we say "fraction of thoughts," we really mean "fraction of total cognitive time." This converts a metaphysical question ("what is a thought?") into an accounting question ("how much time is spent in mode X across humanity?").

This move does not reduce thought to time in a philosophical sense. It is an operational choice for scaling arguments. We need a denominator that is coherent, and time provides one.

2.2 Time-share proxy: fraction_of_thoughts ~= fraction_of_cognitive_time
We introduce the proxy explicitly in CopySafe ASCII Math:

fraction_of_thoughts ~= fraction_of_cognitive_time

The symbol "~=" is not claiming equality; it is acknowledging that we are using time as a surrogate measure. The assumption behind the surrogate is modest: if you sample cognition at random moments across humanity, the probability you land in "breakthrough cosmology cognition" is roughly proportional to the share of time such cognition occupies.

This is a sampling interpretation rather than a metaphysical one. Imagine a hypothetical "random moment sampler" that selects a random awake moment from a random person and asks: what is the cognitive content at that moment? If we can approximate the proportion of sampled moments that belong to the target category, we can talk coherently about a global fraction. That fraction is naturally a time-share.

There are known distortions. Some cognitive states are more "dense" than others: a 10-second flash of insight may carry more novelty than an hour of routine activity. Conversely, an hour of careful derivation might contain fewer "distinct thought tokens" than an hour of daydreaming. But the point of this paper is order-of-magnitude scaling. For that purpose, time-share is the least assumption-heavy metric. It yields a conservative framework: it forces us to count the laborious parts of science that are cognitively real (verification, calibration, error modeling) rather than romanticizing breakthroughs as disembodied lightning bolts.

A secondary reason this proxy is useful: time budgets can be triangulated by multiple independent methods (surveys, work logs, publication cadence, meeting time, teaching load), whereas thought counts cannot.

2.3 Awake-time normalization and why it is conservative
We define the denominator as total human awake time rather than total clock time. This matters because sleeping cognition is both harder to interpret (dream content, fragmented

awareness) and less relevant to the kinds of disciplined, constraint-heavy cognition under discussion. The simplest normalizing choice is:

total_awake_time ~= N * awake_hours_per_day * 365

Where N is world population. Using awake time is conservative in a specific sense: it increases the denominator relative to using "work time" or "educated adult time" or "internet-active time," which would make the target fraction larger. We intentionally choose the broadest plausible denominator because the claim is population-wide. If the paper's conclusion holds even under the largest denominator, it becomes more robust.

Awake-time normalization is also conservative because it implicitly credits the entire species with cognitive reality. It does not narrow the denominator to "people who have the leisure to think abstractly" or "people with formal education." It includes caregivers, laborers, children (to the extent they have awake cognition), the sick, and the elderly. Again: this is not virtue signaling; it is a scaling discipline. If you want to claim something about humanity-wide cognition shares, you must not quietly exclude most humans from the denominator.

A further conservative feature is that "awake time" still overestimates "available reflective time." Much awake time is constrained by tasks, stress, pain, and urgency. Using all awake time as eligible cognition makes the denominator larger than a denominator restricted to "discretionary reflective time." That pushes the estimated fraction downward, which is appropriate when making a claim of extreme rarity: we want to avoid inflating the numerator by shrinking the denominator.

2.4 "About cosmology" versus "uses cosmology-adjacent tools"

To prevent category drift, we distinguish two levels of inclusion.

(1) about_cosmology: cognition whose object is a cosmological claim, question, method, dataset, inference, or theoretical structure. It is cosmology as the topic.

(2) cosmology_adjacent: cognition that uses tools or methods that also appear in cosmology, but whose object is not cosmology. Examples include statistics, Bayesian inference, signal processing, HPC, numerical simulation, machine learning, detector calibration, time-series analysis, and visualization. These tools are cosmology-adjacent, but using them does not mean one is thinking about cosmology.

This distinction is crucial because modern science shares toolchains. A large fraction of contemporary technical cognition uses the same computational and statistical substrate. If we blurred the categories, we could mistakenly count huge amounts of non-cosmology work as cosmology cognition. That would destroy the meaning of the claim.

Operationally, a thought is "about cosmology" if changing the cosmological context would change what the cognition is trying to accomplish. If you can substitute "financial time series" or "genomics" or "computer vision" with minimal change, then the cognition is not about cosmology; it is about the shared substrate. Cosmology-adjacent tools can be necessary conditions for cosmology, but they are not sufficient conditions for cosmology-aboutness.

This is also where public discourse can mislead. Popular cosmology content often features the aesthetic shell of cosmology: dark matter, black holes, the Big Bang, "the universe is expanding." But it rarely engages the constraint structure that makes cosmology cosmology: systematics, selection effects, instrumental transfer functions, likelihood construction, parameter degeneracies, priors, cross-survey consistency, and the hard work of arguing with your own pipeline.

So the "about cosmology" filter intentionally points toward cognition that is entangled with cosmology's constraints, not merely its vocabulary.

2.5 Breakthrough / new-to-world / field-shaping: operational criteria

"Breakthrough" is the most loaded term in the paper, so we specify operational criteria rather than relying on intuition. We do not require that a breakthrough be immediately recognized, but we do require that it has stable downstream consequences.

A candidate cosmology cognition episode counts as breakthrough-level if it satisfies most of the following criteria over a meaningful time horizon (often years):

(1) new_to_world: the core claim, method, or conceptual reframing is not merely a rephrasing or routine recombination of known components. It introduces a genuinely new element into the field's option space.

(2) constraint_surviving: it survives contact with the field's error model. In cosmology, that means it is not a fragile artifact of an unexamined systematic, a miscalibration, or a degenerate fit presented as certainty.

(3) field_shaping: it changes what competent practitioners must consider. This can happen in multiple ways:

- it becomes a new standard method or pipeline component
- it resolves or reframes a major tension or anomaly
- it establishes a new measurement regime or observable
- it forces revision of default priors or baseline models
- it opens a new program that attracts sustained follow-on work

(4) irreducibility_of_impact: its influence is not merely sociological hype. There is a durable technical residue: later work depends on it, tests against it, or uses it as a baseline.

(5) counterfactual dependence: imagine removing this idea from the world. Would the field, within a decade, look meaningfully different? If the answer is yes, that is evidence of field-shaping status.

These criteria allow a spectrum rather than a binary. The paper's "sub-ppb" thesis is about the extreme tail: cognition episodes that meet the strongest version of these criteria, not the broad class of "interesting ideas."

The criteria also explain why the breakthrough filter is so severe: cosmology is a domain where correctness is expensive. Field-shaping claims must be robust across datasets, instruments, teams, and analysis choices. The higher the evidential bar, the more rare the cognition episodes that clear it.

Finally, we emphasize a discipline of humility: a thought can feel like a breakthrough to the thinker and still fail these criteria. The paper is explicitly not counting "breakthrough feelings." It is counting what remains after the field has done its adversarial work. That adversarial work is part of cosmology's dignity and part of why the numerator in our sub-ppb claim collapses so strongly.

3. The Funnel Model: From Humanity to Breakthrough Cosmology

3.1 Stage A: total population and total awake-hours

Stage A is the denominator discipline. We are not comparing cosmology to "other science," or to "educated people," or to "people online." We are comparing against the cognitive life of the species.

We define:

- N = world population
- h_awake = average awake hours per person per day (a coarse global mean)
- T_awake = total human awake-hours per year

In ASCII math:

T_awake ~= N * h_awake * 365

This is intentionally blunt. The goal is not precision to three significant digits; it is a stable denominator that does not secretly exclude most humans. Using awake time is also conservative in the sense that it makes the denominator large. If the numerator remains tiny even under this “maximal” denominator, the rarity claim becomes more robust.

A useful mental picture: if you imagine sampling a random awake moment from a random human, Stage A is the sample space of that random draw. Everything that follows is an increasingly strict set of filters on what that moment could be “about.”

3.2 Stage B: cosmology-capable population (training and context)

Stage B is the first major “collapse.” Most humans are not cosmology-capable in the operational sense used in this paper. “Capable” here does not mean intelligent, curious, or spiritually attuned to wonder. It means: able to produce or critically evaluate cosmology claims inside the field’s constraint structure.

Cosmology-capability typically requires:

- advanced mathematical literacy (linear algebra, calculus, probability, statistics)
- physics fluency (classical mechanics, EM, QM, GR at least conceptually; often more)
- comfort with inference under uncertainty (likelihoods, priors, degeneracies, systematics)
- some exposure to the actual objects of practice (datasets, simulations, pipelines, calibration)
- cultural context: knowing what the field currently treats as live questions vs settled scaffolding

This stage matters because it separates “cosmology as awe” from “cosmology as knowledge production.” Many people think about the universe; far fewer can think cosmologically in the field’s technical dialect.

We can name the Stage B population size:

- K = number of cosmology-capable individuals worldwide

Then the Stage B share is:

K / N

We do not need to pick a precise K to see the main point: K is far smaller than N, and this collapse happens before we even talk about “active research,” much less “breakthroughs.”

3.3 Stage C: active contributors (publishing, analyzing, building instruments)

Stage C narrows from "capable" to "currently engaged in the field as a producing system." This includes people who are:

- publishing cosmology research (theory, observation, data analysis, simulation)
- building or operating instruments and surveys where cosmological inference is an explicit aim
- developing core methods used in cosmological pipelines (calibration, inference frameworks, map-making, systematics control)
- maintaining crucial shared infrastructure (software, catalogs, pipelines) when that work is directly cosmology-facing

We call this population:

- C = number of active cosmology contributors

Stage C is where cultural prominence and cognitive share diverge sharply. Cosmology can be highly visible in public imagination while C remains small, because visibility scales with media reproduction, not with the number of people doing frontier work.

A subtle point: even within this active group, much cognition is not "about cosmology" in the strict sense at each moment. A large fraction of work time is meetings, debugging, documentation, writing, administration, teaching, and maintenance. Those activities can be necessary conditions for cosmology, but they are not always cosmology-aboutness at the content level. That reality is handled later via the breakthrough-mode time share g.

3.4 Stage D: breakthrough-capable subset (track record, positioning, resources)

Stage D addresses a second collapse that many accounts ignore: not all active contributors are equally positioned to produce field-shaping breakthroughs, even in principle.

Breakthrough capability is not just raw intelligence. It is an alignment of multiple constraints:

- track record and deep familiarity with the field's failure modes (especially systematics traps)
- proximity to high-leverage data, instruments, surveys, or simulations
- access to compute, collaboration networks, and the right technical stack
- enough uninterrupted time to perform deep work (the rarest resource)
- the social positioning to pursue risky questions without immediate career penalty

- the temperament for cosmology's adversarial burden: rigorous self-attack before public claim

We define:

- B = number of breakthrough-capable cosmology contributors (a subset of C)

So:

$B <= C <= K <= N$

Stage D is where the "heavy-tail" structure begins to dominate. In heavy-tailed innovation, a relatively small subset of contributors accounts for most of the "direction changes." This is not a moral statement; it is an empirical regularity across many creative domains. Cosmology amplifies it because the evidential bar is high and the toolchains are expensive.

A key consequence: even if C is larger than people intuit, B can remain small, because breakthrough capability depends on positional and infrastructural constraints, not only on personal qualities.

3.5 Stage E: breakthrough-mode time share (g) within that subset

Stage E introduces the final and often decisive collapse: even among the breakthrough-capable subset, only a small fraction of cognitive work time is truly "breakthrough-mode" under the strict definition.

We define:

- g = fraction of a contributor's cognitive work that is in true breakthrough-mode

Breakthrough-mode is not "working on a cosmology project." It is the narrow cognitive regime where a person is actually generating, stress-testing, and integrating candidate ideas that could plausibly be new-to-world and field-shaping. That regime is small because cosmology forces long overhead:

- verification and re-verification
- error modeling and robustness checks
- pipeline discipline (calibration, selection effects, instrument modeling)
- reconciling with prior literature and known constraints
- conversion of insight into testable form (predictions, diagnostics, falsifiers)
- collaboration coordination and reproducibility burden

Even for exceptional people, most days are not "breakthrough-mode." Many days are necessary maintenance that makes breakthroughs real rather than delusional.

There is also an important distinction between:

- attempted_breakthrough_mode: time spent pursuing ideas that feel potentially field-shaping
- true_breakthrough_mode: time spent on ideas that, after adversarial contact with the field, actually qualify as new-to-world and field-shaping

The paper's claim concerns the second category. Operationally, you can treat g as already "net of self-illusion," meaning it counts only the time that survives the discipline of constraints. This pushes g downward relative to naive intuition, and again that downward push is not cynical; it is exactly the cost of honesty in a high-uncertainty domain.

If desired, one can split g into two parameters:

- g_attempt = time share spent in attempted breakthrough-mode
- s = survival fraction of attempted breakthrough-mode that later qualifies as true breakthrough-mode

Then:

g_true = g_attempt * s

But for the funnel's headline scaling, it is enough to work with a single effective g (implicitly g_true).

3.6 The core formula in ASCII math

The minimal model given is:

share ~= (C / N) * g

This is the population-wide cognition share of true breakthrough-mode cosmology, under the time-share proxy.

However, because Stage D matters, the sharper version is:

share ~= (B / N) * g

Where B is the breakthrough-capable subset and g is the breakthrough-mode time share within that subset.

If you want to keep the user-provided symbols while honoring the funnel, you can write:

Let:

- N = world population
- C = number of active cosmology contributors
- f = (B / C) = fraction of contributors who are breakthrough-capable
- g = fraction of contributor cognitive work in true breakthrough-mode

Then:

share ~= (C / N) * f * g

This expression makes the funnel explicit: Stage C contributes (C/N), Stage D contributes f, Stage E contributes g.

A concrete sanity-check (purely illustrative, not a claim of measured truth):

- N ~= 8e9
- B ~= 2e3 (a few thousand breakthrough-capable cosmology contributors worldwide)
- g ~= 1e-3 (0.1% of their work cognition is true breakthrough-mode)

Then:

share ~= (2e3 / 8e9) * 1e-3
share ~= (2.5e-7) * 1e-3
share ~= 2.5e-10

That is 0.25 ppb, i.e., comfortably below 1 ppb.

If you tighten either parameter even slightly (B smaller, or g smaller under stricter “true” criteria), the share drops further. If you loosen them substantially (B larger, g larger), you can push above 1 ppb, but you must do so by making strong claims about how many people are truly positioned for field-shaping work and how much of their cognitive time is genuinely in that mode.

The funnel’s purpose is not to force one number. Its purpose is to show why the rarity regime is structurally stable: the product of several small fractions is extremely small, and cosmology’s constraint burden makes the final factor g smaller than most people expect.

4. Order-of-Magnitude Estimates

4.1 Bounding C: what counts as "active cosmology contributor"

The variable C is doing heavy lifting, so we bound it carefully. "Active cosmology contributor" is narrower than "astronomer" and narrower than "astrophysicist." It is also narrower than "person who publishes something with the word cosmology in it." The operational meaning in this paper is:

An active cosmology contributor is someone who is currently spending a material share of their professional cognitive work on producing, validating, or enabling cosmological inference at the research frontier.

That definition includes several roles:

- **Cosmology theory**: model building, analytic derivations, consistency checks, predictions, model comparison frameworks.
- **Observational cosmology**: survey analysis, map-making, parameter inference, systematics control, cross-survey reconciliation.
- **Simulation and numerical cosmology**: large-scale structure simulations, synthetic sky generation, emulator building, forward models, uncertainty quantification.
- **Instrumentation and pipelines (cosmology-facing)**: detectors, calibration, transfer functions, end-to-end pipelines explicitly constructed for cosmological observables.
- **Methods that are cosmology-specific**: inference frameworks whose constraints and systematics are cosmology-shaped rather than generic statistics.

It excludes:

- **General astronomy** (stellar, exoplanets, galactic structure) unless it is explicitly cosmology-facing in a frontier sense.
- **Generic ML/statistics/HPC** work unless the object of inference is cosmological and the work lives inside cosmology's constraint structure.
- **Public outreach** and popularization (important, but not frontier contribution in the numerator sense).
- **Education and teaching** unless it is directly tied to producing frontier results (again: important, but different category).

With that boundary, C is plausibly on the order of **a few thousand to perhaps around ten thousand** worldwide at any given time. A tighter plausible band (for the purpose of order-of-magnitude reasoning) might be:

- **C_low ~= 1,000** (restricting to direct frontier analysis + core theory + cosmology instrument leads)
- **C_mid ~= 3,000 to 7,000** (including broader survey teams, simulation groups, pipeline engineers)
- **C_high ~= 10,000 to 20,000** (a generous upper bound that risks counting a lot of cosmology-adjacent work)

The paper does not need the exact value. The key point is structural: even the high end of C is tiny relative to N ~= 8e9. That alone pushes the share into the ppm-to-sub-ppm range before we apply the breakthrough filter.

4.2 Bounding g: why breakthrough-mode time is tiny even for experts

The parameter g is where intuition usually fails. People imagine breakthroughs as the central activity of frontier scientists. But in a constraint-heavy field like cosmology, most cognitive time is spent paying the "truth tax":

- verifying that the signal is not an artifact
- modeling systematics and calibration
- stress testing assumptions, priors, and degeneracies
- reproducing results across pipelines and teams
- writing and communicating in forms that survive adversarial review
- coordinating in large collaborations where the unit of work is often not an individual mind

Even for a highly capable person, the fraction of time spent in what we are calling **true breakthrough-mode** is small. It is not "working hard" and it is not "thinking about cosmology." It is a specific regime: generating or integrating a candidate idea that (i) is plausibly new-to-world and (ii) has a credible path to changing what the field must account for.

A useful decomposition is:

g_true = g_attempt * s

Where:

- g_attempt = fraction of work cognition spent in *attempted* breakthrough-mode (high-risk ideation, deep reframing attempts, new method proposals, radical consistency checks)
- s = survival fraction of attempted breakthrough-mode that later qualifies as *true* breakthrough-level under the field's constraint structure

In many domains, g_attempt could be a few percent for some people. But s can be extremely small when the evidential bar is high. Cosmology is famous for "beautiful ideas" that do not survive contact with data, systematics, or other constraints. The numerator in this paper is explicitly **net of that attrition**.

So for g_true, plausible order-of-magnitude bounds might be:

- **g_low ~= 1e-4** (0.01% of work cognition)
- **g_mid ~= 3e-4 to 1e-3** (0.03% to 0.1%)
- **g_high ~= 3e-3** (0.3%) as a generous upper bound under strict "true breakthrough" criteria

These values can feel offensively small until you convert them into time. If a researcher works ~2,000 hours/year, then:

- g = 1e-3 corresponds to ~2 hours/year of true breakthrough-mode time
- g = 3e-3 corresponds to ~6 hours/year
- g = 1e-4 corresponds to ~12 minutes/year

Those numbers are not saying people rarely have insights. They are saying the subset of insights that end up being **field-shaping breakthroughs** is extremely sparse. Cosmology does not merely reward creative production; it demands robust survival across an adversarial constraint environment.

4.3 A worked example yielding sub-ppb

We now produce a worked example using only the funnel math. Use:

- N = 8e9
- C = 6,000 active cosmology contributors
- f = B / C = 0.25 breakthrough-capable share (so B = 1,500)
- g = 1e-3 true breakthrough-mode time share within B

Then the population-wide share is:

share ~= (B / N) * g
share ~= (1500 / 8e9) * 1e-3
share ~= (1.875e-7) * 1e-3
share ~= 1.875e-10

That is **0.1875 ppb**, comfortably below 1 ppb.

A more "generous-to-the-numerator" variant:

- N = 8e9
- B = 3,000 (larger breakthrough-capable pool)
- g = 2e-3 (0.2% true breakthrough-mode time)

share ~= (3000 / 8e9) * 2e-3
share ~= 3.75e-7 * 2e-3
share ~= 7.5e-10

That is **0.75 ppb**, still below 1 ppb.

These are not claims of measured truth; they are demonstrations that sub-ppb is not exotic. It emerges naturally once you take the strict filters seriously and then use a species-wide denominator.

4.4 What parameter values would be required to exceed 1 ppb

To exceed 1 ppb, we need:

share > 1e-9

Using the B-form:

(B / N) * g > 1e-9
B * g > N * 1e-9

If N ~= 8e9, then:

B * g > 8

This is the key threshold condition: the product of "breakthrough-capable headcount" and "true breakthrough-mode time share" must exceed about eight (in units where g is a fraction).

This lets us see what it would take:

- If **B = 1,000**, then g must exceed 8 / 1000 = 0.008 = **0.8%**
- If **B = 2,000**, then g must exceed 8 / 2000 = 0.004 = **0.4%**
- If **B = 5,000**, then g must exceed 8 / 5000 = 0.0016 = **0.16%**
- If **B = 10,000**, then g must exceed 8 / 10000 = 0.0008 = **0.08%**

So exceeding 1 ppb is not impossible, but it requires at least one of the following to be true:

1. **B is surprisingly large** (thousands upon thousands of people truly positioned for field-shaping cosmology), or
2. **g is surprisingly large** (a large fraction of those people's cognitive work is genuinely in true breakthrough-mode), or
3. **the definition of "breakthrough" is loosened** (raising g by counting many more events as "field-shaping"), or
4. **the denominator is quietly narrowed** (not N, but some smaller population such as "working adults" or "educated adults" or "people online").

The paper's sub-ppb thesis is most defensible when: (i) "breakthrough" is strict, (ii) N is truly species-wide, and (iii) g is interpreted as net-of-survival after adversarial constraint checking.

4.5 Heavy tails: why averages hide the real structure

Even the worked examples above may mislead, because they speak in averages. Breakthrough production is almost certainly **non-uniform** across people and across time.

Two properties matter:

- **Heavy-tailed contribution**: a small number of individuals and teams account for a disproportionate share of major directional shifts. Most contributors produce essential incremental work but do not generate field-shaping changes.
- **Bursty time structure**: even for a "breakthrough generator," breakthrough-mode time is not evenly spread. It comes in bursts: a few intense weeks or months embedded in long periods of infrastructure, verification, and communication.

Averages therefore hide what is actually happening. The world is not:

"Everyone does 0.1% breakthrough-mode time every year."

It is more like:

"Most people do ~0% true breakthrough-mode time in most years; a small minority do nonzero amounts; and a tiny minority occasionally produce bursts that matter."

This is important for interpretation because it changes what the ppb claim means. It does not mean "there are hardly any cosmology thoughts." It means "if you sample randomly from humanity's cognitive stream, the chance you land inside the rare bursts where field-shaping cosmology is actually being generated is extremely small."

You can formalize this as a mixture model:

- With probability (1 - p), a random draw lands in non-breakthrough cognition.
- With probability p, it lands in a breakthrough-generating regime (and even then, only some of those survive as true breakthroughs).

The ppb-level share is then a statement about p, not a statement about the importance of what happens inside the rare regime when it occurs.

This heavy-tail picture also warns against moralizing the statistic. In heavy-tailed systems, a tiny fraction of activity can have outsized consequences. The fact that a process occupies sub-ppb of cognition does not imply it is negligible. It implies it is rare. Cosmology is a prime example of a rare cognition stream with disproportionate cultural and philosophical impact, precisely because it speaks to boundary conditions for almost everything else we think we know.

In short: order-of-magnitude estimates push us toward sub-ppb under strict definitions, and the heavy-tail structure explains why that can be simultaneously true and culturally world-shaping.

5. Sensitivity Analysis: Where the Claim Could Fail

5.1 If C is much larger than assumed (broadening the field boundary)

The sub-ppb claim depends on the numerator being small under a strict definition of "active cosmology contributor." The most direct way to break the claim is to make C much larger by broadening the boundary of what counts as "cosmology."

There are at least three ways this broadening could happen:

(1) **Include most of astronomy/astrophysics.**
If we treat any astronomy research as cosmology-adjacent enough to count, then C can jump by an order of magnitude. But this move is definitional sleight-of-hand. Much of astronomy does not aim at cosmological parameters, early-universe physics, large-scale structure inference, or

the interpretive framework of Lambda-CDM alternatives. Counting all astronomy would mean the category is no longer "about cosmology," but "about space."

(2) **Include generic methods work.**
Modern cosmology uses generic tools: Bayesian inference, large-scale computation, ML, statistical calibration, signal processing. If we count everyone building such tools as a cosmology contributor, C becomes enormous. But then the category has become "cosmology-adjacent tools," not "cosmology."

(3) **Include public discourse and contemplation.**
If "about cosmology" includes watching a cosmology documentary, reading a pop-sci article, or having wonder-filled reflections about origins, then C effectively becomes "anyone who thinks about the universe," which could plausibly be a large fraction of humanity. Under that definition, the sub-ppb claim becomes false by construction. But it becomes false because the numerator has changed categories: we are no longer counting field-shaped knowledge production; we are counting cosmological sentiment and narrative.

So yes: the claim can fail if C is defined so broadly that it loses its disciplinary meaning. That failure mode is important to name explicitly because it distinguishes "myths about the cosmos" from "cosmology as a constraint-heavy research practice." The paper's headline is about the latter.

5.2 If g is much larger than assumed (broadening "breakthrough")

The second main way to break the sub-ppb conclusion is to inflate g, the breakthrough-mode time share. This can happen for two reasons: genuine empirical reality (cosmologists spend more time in breakthrough-mode than we estimate) or definitional broadening (we count many more cognitive episodes as "breakthrough").

There are several "broadening moves" that can make g jump:

(1) **Count attempted breakthroughs rather than true breakthroughs.**
If g counts "time spent trying to make breakthroughs," g can be a few percent for some people. But the paper's claim is about the strict tail: ideas that survive adversarial constraint checking and actually become field-shaping. In cosmology, the survival rate of attempted breakthroughs is likely small. If you replace g_{true} with $g_{attempt}$, you will drive the estimated share upward, potentially above ppb.

(2) **Count incremental novelty as breakthrough.**
If any genuinely new method tweak, parameter constraint improvement, or novel dataset integration counts as "breakthrough-level," then g inflates dramatically. But that is exactly the boundary the paper is trying to enforce: "breakthrough-level / field-shaping" is not "publishable novelty." It is "alters what the field must account for."

(3) **Count toolchain innovation as cosmology breakthrough even when not cosmology-shaped.**
If someone invents a generic ML method that later proves useful in cosmology, do we count their ideation as “breakthrough cosmology”? Under the paper’s strict framing, no: the cognition was not about cosmology; it was about a general method. If we count it anyway, g inflates. But again: the category drifts from “about cosmology” to “influences cosmology eventually.”

Even under strict definitions, there is one real-world scenario that could make g larger than expected: a future in which AI and automation collapse overhead (calibration, checks, pipeline maintenance) so dramatically that human researchers spend a much higher fraction of their time in generative regime. That is a live possibility, but it is precisely something the later sections address: AI may increase g, C, or both. The paper’s present claim is about the current structure of the field (or near-current), not a guaranteed invariant.

5.3 If "thought" should not be time-weighted (alternate metrics)
The proxy fraction_of_thoughts ~= fraction_of_cognitive_time is defensible, but it is not the only possible metric. If a critic rejects time-weighting, the claim might fail depending on what alternate weighting they choose.

Here are three alternate metrics and how they would change things:

(1) **Information-weighted cognition.**
Instead of counting minutes, count “bits of novelty” or “bits of model update” per unit time. A short insight can have huge informational consequence compared to hours of routine work. Under information weighting, breakthrough cognition would receive a heavier weight, pushing the fraction upward.

However, this move has a cost: “bits of novelty” is even harder to measure than time, and it creates incentives to romanticize unvalidated novelty. In cosmology, a high-information-looking insight can still be wrong. If you information-weight raw ideation, you risk overweighting elegant errors. If you information-weight only validated breakthroughs, you are back to time-like accounting because validation requires time and field uptake.

(2) **Outcome-weighted cognition.**
Weight thoughts by downstream impact: citations, paradigm shifts, educational influence, technological spinoffs, cultural presence. Under outcome weighting, cosmology breakthroughs could occupy a larger share than their time footprint, because their cultural and conceptual impact is disproportionate.

But outcome-weighting changes the question. The paper’s claim is explicitly about cognition share, not impact share. It is entirely consistent that cognition share is sub-ppb while cultural impact share is far larger.

(3) **Attention-weighted cognition within connected media.**
If you restrict the denominator to internet attention or media consumption, cosmology may become a larger fraction. But that is again a denominator switch. The sub-ppb claim is about all human cognition, not the cognition of a subset (online, educated, or leisure-rich).

So yes: if “thought” is not time-weighted, the headline number can move. But most alternate metrics either (i) become unmeasurable, (ii) change the question to impact rather than cognition, or (iii) sneak in a smaller denominator.

5.4 If the breakthrough process is distributed outside the formal workforce
A legitimate objection is that the funnel model may be too institutional: it may assume breakthroughs emerge primarily from credentialed cosmologists embedded in universities, labs, and survey collaborations. What if breakthrough cosmology is substantially distributed outside the formal workforce?

There are several forms of “outsider contribution”:

(1) **Independent theorists and cross-discipline physicists.**
Some may not label themselves cosmologists but could contribute field-shaping ideas. This enlarges C and possibly B.

(2) **Method innovators in adjacent fields.**
Statisticians, applied mathematicians, computer scientists, and ML researchers could invent techniques that unlock a cosmology breakthrough when applied properly.

(3) **Citizen science and distributed analysis.**
In some domains, distributed volunteers have made genuine contributions (classification, anomaly spotting). Cosmology has some citizen-science footprints, though breakthrough-level inference is harder because it is toolchain- and systematic-heavy.

If you count these distributed sources, does the claim fail? Usually not, for a simple reason: even if you expand the pool, you do not automatically expand true breakthrough-mode time share. Outsider ideas face a harsh translation cost: connecting the idea to cosmology’s constraint structure, gaining access to relevant data, controlling systematics, and convincing the field through reproducible pathways. That translation is expensive and often must be done in partnership with the formal ecosystem.

So distributed contribution is real, but its effect on the numerator is limited by the same constraint that governs insiders: cosmology breakthroughs require survival across adversarial checks. Outsiders can contribute, but the number of outsiders who both generate and successfully integrate field-shaping cosmology is likely still small relative to humanity-wide denominators.

5.5 Robustness: why the claim survives most reasonable perturbations

After naming failure modes, we can now state why the sub-ppb conclusion is robust under most reasonable perturbations.

The robustness comes from multiplicative smallness. The share is:

share ~= (B / N) * g

To exceed 1e-9 with N on the order of 1e10, you need B * g on the order of 10. That is a surprisingly strong requirement under strict definitions.

If you perturb parameters modestly:

- Double B (more breakthrough-capable people)
- Double g (more breakthrough-mode time)
- Increase C by a factor of 3 (broaden boundary a bit)

The share scales linearly with each change, but you still need large combined multipliers to jump an order of magnitude. In practice, you only cross 1 ppb if you do one of these:

- broaden “cosmology” until it becomes “space-related thinking”
- broaden “breakthrough” until it becomes “interesting novelty”
- abandon the species-wide denominator
- count attempted breakthroughs as if they were successful field-shaping breakthroughs
- assume a future AI-mediated regime where overhead collapses and g rises dramatically

If you hold steady on the strict meanings, the conclusion is stable because:

- N is enormous and does not move much.
- B is limited by training, positioning, and resources.
- g is limited by cosmology’s truth tax and by the scarcity of ideas that survive adversarial constraint checking.

This is the core robustness claim: you can argue about the exact number, but under strict definitions, you are arguing within a regime that is naturally ultra-small. Sub-ppb is not a fragile point estimate; it is a stable basin of plausibility produced by (i) a species-wide denominator and (ii) a highly selective numerator defined by both domain and novelty.

In other words: the headline can be wrong in the third decimal place of ppb, but it is unlikely to be wrong in the direction of becoming "common." Under strict definitions, breakthrough cosmology is structurally rare. That rarity is not a judgment. It is a scaling fact.

6. Alternative Proxies and Cross-Checks

6.1 Bibliometric proxy: share of top-0.1% cited cosmology work

One way to cross-check the "sub-ppb cognition share" thesis is to use bibliometrics as a proxy for field-shaping events. The underlying intuition is simple: if a contribution truly changes what the field must account for, it tends to become disproportionately referenced, reused, debated, taught, and built upon. Citation counts are not a moral measure, and they are not pure epistemic truth, but they are a visible trace of uptake.

An operational proxy could be:

- Define a cosmology corpus (journals + arXiv categories + key survey papers).
- Identify the top 0.1% by citations (or by field-normalized citation impact).
- Treat that set as a rough approximation to "field-shaping" contributions.
- Estimate the annual production rate of such top-tail papers and the number of distinct authors responsible.

Then ask: how many people, how many author-hours, and how many years of training are concentrated in the production of that tail? Even without precise numbers, the method clarifies structure. The top 0.1% tail is, by construction, extremely sparse. If cosmology publishes on the order of thousands of papers per year, the top 0.1% tail is on the order of single-digit papers per year. That is already consistent with the rarity thesis: field-shaping items are not a continuous stream.

Caveats matter here:

- Citations lag; breakthroughs can be recognized late.
- Some breakthroughs are "silent infrastructure" (methods or calibrations) that may be under-cited relative to their importance.
- Citations inflate around fashionable topics and collaborations.
- A paper can be heavily cited because it is wrong but influential as a foil.

Even with these caveats, bibliometrics has value as a cross-check because it forces us to confront the heavy-tail regime quantitatively rather than rhetorically. The proxy does not prove the ppb claim, but it can validate that "field-shaping outputs" live in an extreme tail, and that extreme tails naturally correspond to ultra-small cognition shares under species-wide denominators.

6.2 Prize/recognition proxy: how often the field truly "turns"

Another cross-check is to ask: how often does cosmology undergo a widely recognized turn, where the field's baseline story, baseline parameters, or baseline observables change meaningfully?

We can operationalize "turn" in a few ways:

- A new observational regime becomes standard (e.g., a survey era that resets constraints).
- A major tension becomes a durable, organizing anomaly (or is resolved).
- A new method becomes canonical (a pipeline or inference approach everyone must use or respond to).
- A new theoretical framing becomes unavoidable (even if contested).

Prizes, keynote invitations, major review articles, high-profile missions, and canonical textbooks are all weak signals of turning points. Recognition is not truth, but it is a marker that the community has collectively decided something mattered enough to incorporate into its shared memory and training.

If you view "turns" as rare—occurring on the scale of years or decades rather than months—then the time density of breakthrough-level events is necessarily low. This aligns with the sub-ppb thesis even before quantifying headcounts: if the field truly turns only occasionally, then the total "breakthrough-mode" cognition producing those turns must be concentrated into rare bursts.

This proxy also exposes a subtlety: a turn does not always map to a single person or paper. It can be a distributed phase transition of many incremental contributions. That does not weaken the rarity claim; it strengthens it. Distributed turns still require a special alignment of data, methods, and interpretation, and they do not happen continuously. A phase transition can be distributed and still rare.

6.3 Preprint/author network proxy: size and activity of the core graph

Cosmology has a visible preprint culture, which makes network analysis feasible. You can build a co-authorship graph:

- nodes = authors

- edges = co-authorship on cosmology-tagged papers
- weights = frequency or recency of collaboration

Then define a "core" by network centrality measures (degree, betweenness, eigenvector centrality) and by sustained activity (e.g., at least X cosmology preprints in the last Y years).

This proxy gives two useful measurements:

(1) **Core size**: the number of authors who are repeatedly active in the frontier production stream. This is a practical estimate for C.
(2) **Core structure**: whether the field is dominated by a few massive collaborations (dense clusters) and how many bridging nodes connect subfields.

Breakthrough-level activity, under strict criteria, is likely concentrated in:

- the highly central core (because it has access to data and collaboration infrastructure)
- or in rare bridging nodes (because they connect methods/ideas across subfields)
- or in peripheral nodes that occasionally inject new ideas that survive (rare but possible)

Network structure therefore provides a cross-check for the heavy-tail story. If the cosmology author graph has a relatively small, highly connected core and a large periphery of low-activity nodes, then "random moment sampling" across humanity will almost never land inside the cognitive stream of core actors, much less inside their rare breakthrough-mode moments.

This proxy also highlights the denominator mismatch: even if the cosmology author graph contains tens of thousands of nodes, most nodes will be low-activity or transient (students, short-term contributors), and only a smaller subset sustains frontier engagement over long periods.

6.4 Instrumentation proxy: telescope teams, survey pipelines, compute budgets
Cosmology breakthroughs are not purely mental events; they are often enabled by instrumentation, surveys, and compute. A fourth cross-check is therefore infrastructural:

- count major cosmology-facing instruments and surveys (CMB experiments, large-scale structure surveys, gravitational lensing campaigns, etc.)
- estimate team sizes and the distribution of roles (instrumentation, data management, pipeline, analysis, theory integration)
- estimate compute budgets and pipeline complexity (as a proxy for the overhead burden)

This proxy is valuable because it ties the breakthrough process to the physical economy of the field. If breakthroughs require rare, expensive, and coordinated infrastructures, then the number of people embedded in those infrastructures is bounded. The pipeline itself imposes constraints on who is even in a position to do "breakthrough-level cosmology" in the strict sense: you need access to the data, to the pipeline, to the systematic models, and to the collaboration's interpretive regime.

Instrumentation also clarifies why g is small. In instrument-heavy science, a huge fraction of cognitive work is spent building and verifying the apparatus that makes any inference trustworthy. Breakthrough-mode time is sparse because the system demands long periods of calibration, validation, and reconciliation.

Compute budgets add another angle: if cosmological inference and simulation are compute-intensive, then the number of actors who can run end-to-end experiments is limited. That limitation constrains B (breakthrough-capable subset) even if C (active contributors) is larger.

In short: infrastructure provides a natural cap on how large C and B can realistically be at any given time, reinforcing the plausibility of ultra-small population-wide cognition shares.

6.5 Comparing proxies and reconciling contradictions

These proxies will not always agree, and that is a feature, not a flaw. Each proxy measures a different shadow of the same underlying phenomenon.

- Bibliometrics measures **uptake** (what the field cites and builds on).
- Recognition measures **collective memory and narrative** (what the field commemorates).
- Network structure measures **participation topology** (who is repeatedly active and connected).
- Instrumentation measures **physical and organizational constraints** (who can access the frontier and what overhead costs exist).

Contradictions arise for predictable reasons:

- A contribution can be field-shaping yet under-cited (quiet infrastructure, calibration methods).
- A contribution can be highly cited yet not field-shaping (popular review papers, fashionable but ultimately transient topics).
- Recognition can reflect sociology and timing as much as epistemic substance.
- Networks can inflate headcounts via transient contributors while hiding the small number of people who actually steer inference choices.

- Instrumentation can concentrate power in collaborations, making breakthroughs appear "collective" rather than traceable to individuals.

The reconciliation strategy is therefore not to pick a single proxy and declare victory. It is to ask whether all proxies converge on two structural claims:

(1) The process is **heavy-tailed and bursty**, not uniform.
(2) The set of actors and moments that plausibly qualify as "breakthrough-level cosmology" is **small enough** that, under a humanity-wide denominator, the time-share plausibly falls into the sub-ppb regime.

If bibliometrics shows an extreme tail, if recognition shows rare turns, if networks show a small core, and if instrumentation shows expensive constraints and large overhead, then the ppb-level conclusion is not an artifact of one assumption. It is the intersection of independent, imperfect measurements that all point to the same geometry: cosmology breakthroughs are rare events produced by a small, constraint-burdened frontier community embedded in expensive infrastructures and adversarial verification norms.

Where proxies genuinely diverge, the divergence itself becomes informative. For example:

- If bibliometrics suggests many "breakthrough" candidates per year but recognition suggests few turns, then we may be overcounting breakthroughs and should tighten criteria.
- If networks suggest a very large active graph but instrumentation suggests only a small fraction can touch cosmology-facing data products, then C should be tightened to data-facing contributors.
- If recognition suggests frequent "turns" but systematics debates show instability, then the "turns" may be narrative churn rather than genuine field-shaping change.

The broader methodological point is that the paper's headline is not a single-number prophecy. It is a claim about regime: under strict definitions and species-wide denominators, multiple independent cross-checks make it plausible that breakthrough-level cosmology occupies less than one part per billion of human cognitive time.

7. Why Breakthrough Cognition is Spiky

7.1 Constraint accumulation: why breakthroughs require long preparation

Breakthrough cognition is spiky because cosmology is not a sandbox where you can freely invent. It is a constraint-saturated environment. The relevant constraints are not only physical laws (GR, QFT, thermodynamics) but also observational realities: selection effects, instrumental transfer functions, calibration drift, foreground contamination, survey geometry, covariances, and inference degeneracies. Every mature cosmological claim lives at the intersection of many constraints, and the constraints accumulate over time.

This creates a structural asymmetry between *having an idea* and *having an idea that can survive*. The space of conceivable cosmological stories is vast; the space of stories that can be consistently embedded in the existing constraint lattice is tiny. Most "new" ideas are either:

- already implicitly ruled out by some known constraint, or
- consistent only in a narrow region that collapses under better error modeling, or
- untestable without new observables, or
- equivalent to an existing idea under reparameterization.

Therefore breakthroughs often require years of preparation in which the thinker internalizes the constraint lattice deeply enough that their creative moves land in the thin viable manifold rather than in the vast forbidden space. The preparation phase is not visible as breakthrough cognition. It looks like reading, re-deriving, debugging, writing methods, calibrating, and arguing with past papers. But it is exactly what makes later "insight" possible.

This is one reason the time share g is small: much of the work that enables breakthroughs is not itself breakthrough-mode. It is constraint accumulation: acquiring the ability to "feel" which moves are even legal. Spikiness is not just psychological; it is the emergent rhythm of a mind learning a constraint landscape.

7.2 Toolchain dependence: data, code, pipelines, collaborations

Cosmology breakthroughs are rarely detachable from tools. The modern field is pipeline-shaped. A breakthrough might be a conceptual reframing, but it must ultimately hook into data: CMB maps, galaxy catalogs, lensing shear, supernova distances, BAO signatures, large-scale structure statistics, and their systematic models. That hooking is done through code.

Code and pipelines introduce a second constraint lattice: the toolchain lattice. The result is that breakthrough cognition becomes spiky for logistical reasons:

- access to the right data products may arrive in discrete releases
- pipeline maturity crosses thresholds (a bug fixed, a calibration updated, a systematic modeled)
- compute access comes in bursts (allocations, cluster availability, queue cycles)
- collaboration decisions gate what is even analyzed (what metrics are prioritized, what sanity checks are mandatory)

- time is fragmented by releases, deadlines, and coordination overhead

Under toolchain dependence, insight is often waiting for a moment when the machinery can actually express it. A person can have a correct conceptual move in mind, but until the pipeline allows a decisive test or a robust measurement, the move remains speculative. When the toolchain catches up—new data, new calibration, new simulation run—the insight can crystallize quickly. That looks like a spike, but it is really a phase transition triggered by infrastructure.

This toolchain structure also creates another form of spikiness: breakthroughs often emerge not from “pure thought” but from **contact points** between components of the toolchain that were previously treated separately: theory meets pipeline, simulation meets catalog systematics, inference meets new foreground models, a cross-check reveals a mismatch, a mismatch forces reinterpretation. These contact points happen in bursts when projects align, which is why the cognitive signal looks intermittent.

7.3 Social selection: funding, hiring, prestige gradients, and path dependence

Breakthrough cognition is also spiky because the social economy of science shapes where attention can go. Cosmology is expensive. Telescopes, surveys, detectors, and compute budgets impose a funding regime. Funding regimes impose a career regime. Career regimes impose risk tolerance.

In practice, this generates selection effects:

- **Funding selects questions.** Some questions are “safe” and legible to reviewers; others are risky and easy to reject. This biases the idea space explored.
- **Hiring selects trajectories.** People adapt to what gets jobs and grants; path dependence accumulates.
- **Prestige gradients concentrate resources.** High-prestige institutions and collaborations can run larger projects and attract top candidates, amplifying inequality of opportunity (not of dignity).
- **Narrative legibility matters.** Ideas that can be told cleanly get attention; ideas that require long technical translation can be ignored, even if correct.

These social forces tend to reduce the continuous exploration of the idea space and instead create “windows” where certain topics become fundable, publishable, and career-sustainable. During such windows, many people can work intensely on a question, increasing the chance of a turn. Outside the window, the same question might be neglected. That produces spikiness at the field level: bursts of attention followed by long quiet.

This is not a cynical claim. It is simply to say that cosmology is not only epistemic; it is also institutional. The structure that makes the field possible also shapes its rhythm. Breakthrough cognition, being rare and high-risk, is the first thing to be squeezed when institutions optimize for steady output.

7.4 Cognitive ecology: deep work, error bars, and the cost of being wrong

Cosmology is unusually punitive of confident error. A wrong claim can be publicly costly, but more importantly it can waste the time of collaborators, distort subsequent work, and pollute the literature with a seductive false trail. The field therefore builds habits of disciplined doubt: error bars, systematics budgets, robustness checks, blinded analyses, preregistered choices in some contexts, and “show me the residuals” cultures.

That culture interacts with cognition in a way that naturally creates spikes:

- Deep work is required for genuine reframing, but deep work is scarce in collaboration-heavy life.
- The cognitive load of error modeling is high; it slows ideation but increases truth.
- The need to verify before claiming means that many insights remain private or tentative for long periods, then appear suddenly when the verification threshold is crossed.
- Attention fragmentation (meetings, emails, coordination) chops time into slices, and deep insights tend to require uninterrupted stretches.

In this ecology, “breakthrough-mode” is not a default state. It is a fragile state that must be cultivated: protected time, quiet mental bandwidth, and a tool environment that allows rapid testing without premature publication. When those conditions are present, cognition can enter a high-yield regime and produce a spike. When those conditions are absent, the same mind may spend months doing important but non-breakthrough work.

The cost-of-being-wrong also forces an internal gating mechanism: a cosmologist with mature judgment will often suppress their own generative impulses until they can attach them to a credible error model. That self-gating reduces the visible frequency of “breakthrough thoughts,” but it increases the survival rate of those that do surface. Spikiness is therefore partly ethical: restraint delays output until it is responsibly formed.

7.5 The "rare alignment" view: right person, right tools, right moment

All of the above can be unified into a single picture: breakthroughs are spikes because they require a rare alignment of conditions. Not just “a genius having an idea,” but a three-way resonance:

- right person: deep constraint-internalization, technical skill, creative temperament, and disciplined skepticism
- right tools: data, compute, pipeline maturity, simulation capacity, cross-check infrastructure, collaboration access
- right moment: a live tension, a fresh dataset release, a methodological impasse, an anomaly that has ripened, or a conceptual bottleneck that many have failed to pass

In ASCII math, you can think of breakthrough emergence as a conjunction event:

breakthrough_event ~= person_ready AND tools_ready AND moment_ready

Each factor is rare. Their conjunction is rarer. That conjunction view explains why even extremely talented people can go years without a field-shaping spike: one of the conjuncts is missing. Maybe the data needed does not exist yet. Maybe the pipeline is not mature enough to test the hypothesis. Maybe the anomaly has not reached "forced choice" status. Maybe the person is buried in administration or teaching. Maybe the collaboration gate is closed. Breakthroughs are not merely intellectual; they are phase transitions of a coupled human-infrastructure system.

The rare-alignment view also explains why breakthroughs can appear "sudden" to outsiders. The public sees the final moment—the paper, the announcement, the new parameter constraint—and imagines a sudden flash. But what often happened is:

- long invisible preparation (constraint accumulation)
- long invisible toolchain maturation (pipelines, calibration, simulation)
- long social incubation (topic becomes fundable and legible)
- long private verification (error bars hammered into place)
- then a sudden crossing of a threshold (a decisive plot, a decisive consistency check, a decisive residual)

When the threshold is crossed, the output can feel instantaneous. But the spike is the last step of a long, coupled process.

Finally, the rare-alignment view ties directly back to the paper's headline statistic. If breakthroughs require rare conjunctions, then breakthrough-mode cognition cannot be evenly distributed. It must be sparse and bursty. Under a humanity-wide denominator, sparseness and burstiness naturally push the measured share toward the ppb regime or below, without implying anything mystical about the field—only that cosmology is a frontier practice embedded in constraints, tools, institutions, and disciplined error culture that collectively make genuine field-shaping cognition rare.

8. Implications

8.1 Public understanding: why "frontier science" feels omnipresent but isn't

The public experiences cosmology as omnipresent because cosmology is a high-symbol domain. It answers the biggest questions in the shortest phrases: "origin," "beginning," "expansion," "dark matter," "black holes," "multiverse." These phrases compress enormous technical work into emotionally resonant tokens. Media systems love tokens that carry awe and existential charge, because they travel well. A single image (a deep field, a CMB map, a black hole shadow) can be replayed for years as if it were continuous novelty.

This creates an illusion of continuous frontier motion. In reality, frontier science is not omnipresent in human cognition, and even within science it is not continuous. The field moves in bursts, while media reproduces the burst outputs as constant presence. The same few frontier stories appear everywhere, so the mind infers a large generating engine behind them. But reproduction is not generation. The paper's ppb framing is a way to name the difference: the generating stream is tiny, and the reproduction stream is enormous.

The implication is not cynicism about popular science. It is clarity about how collective attention works. When the public says "science says," they often mean "the media loop says," which is an amplifier, not the lab. Understanding this helps reduce mistrust: it is not that "science keeps changing its mind" because scientists are fickle; it is that the public is receiving amplified snapshots of a process whose true rhythm is slow, adversarial, and discontinuous.

8.2 Funding narratives: why rarity is not an argument for defunding

A tempting misuse of the ppb claim is: "If breakthrough cosmology occupies so little cognition, why fund it?" That inference is false because it mistakes *frequency* for *value*.

Many valuable processes are rare by design. Emergency medicine is not "most of life," yet we fund it because when it is needed it matters. Bridge engineering failures are rare; we still fund bridge inspection because rare failures can be catastrophic. Likewise, the rarity of field-shaping cognition does not imply low value. It often implies high difficulty, high overhead, and high verification cost.

Cosmology in particular is infrastructure-heavy: instruments, surveys, calibration regimes, data archives, and compute. These systems are not funded to maximize daily "breakthrough thoughts." They are funded to make claims about the universe *responsibly possible*. The overhead is not waste; it is the price of truth.

Moreover, the value of cosmology is not only "new parameters." It includes:

- training in precision reasoning under uncertainty
- toolchain innovations that transfer elsewhere (methods, detectors, inference frameworks)
- the cultural role of disciplined awe: expanding imagination without abandoning rigor
- the long-term civilizational benefit of maintaining institutions that can pursue non-immediate truth

So the right funding narrative is not "fund what is frequent," but "fund what is hard, honest, and has high option value." Rarity is often the signature of a frontier that cannot be mass-produced.

8.3 Epistemic humility: resisting oracle politics in cosmology claims

A tiny generating stream invites a dangerous social pattern: oracle politics. If only a few people can produce frontier knowledge, it becomes tempting for others to outsource judgment entirely: "The experts have spoken; therefore you must obey." Cosmology is not usually a policy lever in the way epidemiology or climate science can be, but the oracle pattern still appears as a cultural temptation: turning cosmological claims into metaphysical authority.

The ppb framing is useful precisely because it can cut two ways:

- It can be used to mystify cosmology: "So rare, so special, therefore unquestionable."
- Or it can be used to de-mystify it: "So rare, therefore we must demand transparency, methods, and error bars."

The paper endorses the second. Cosmology's authority is methodological, not priestly. Its credibility rests on reproducibility, adversarial checking, and explicit uncertainty. When cosmological claims are imported into cultural battles ("science disproves meaning," "the universe proves X"), the right response is epistemic humility: distinguish what the data actually constrain from what people wish to claim.

Resisting oracle politics has practical consequences:

- never treat cosmology as a substitute for ethics, meaning, or theology
- never treat a frontier result as a final metaphysical verdict
- demand that uncertainty be carried forward honestly rather than rhetorically erased
- keep domain boundaries clear: cosmology informs some questions; it does not own all questions

The ppb claim can become a discipline: it reminds us that the generating stream is tiny, and therefore vulnerable to over-interpretation, over-amplification, and authority transfer. If we care about truth, we should resist that transfer.

8.4 Education: training for deep comprehension versus training for spectacle

If frontier cognition is rare, education faces a choice. Do we train the public to consume spectacular cosmology narratives, or do we train them to understand the structure of cosmological inference?

Spectacle education produces:

- awe without method
- slogans without error bars
- authority transfer ("science says") without contestability
- a brittle relationship to uncertainty (either worship or rejection)

Deep comprehension education produces:

- comfort with uncertainty and probabilistic claims
- awareness of systematics, selection effects, and model dependence
- understanding that science is an adversarial process, not a revelation event
- respect for the difference between data and narrative

The ppb framing suggests an educational ethic: if most people will never participate directly in frontier production, they can still be trained to be competent witnesses of it. That means learning how to read claims, not just headlines: what was measured, what was assumed, what would falsify it, what are the alternative explanations, what are the error bars, what is the scope of the conclusion?

This also reframes what "science literacy" should mean. It should not mean memorizing cosmology facts. It should mean being able to reason about inference under uncertainty, because that is the transferable core. A society trained for deep comprehension is less susceptible to both hype and nihilism: it can appreciate wonder while staying disciplined.

8.5 AI and centaur cognition: does AI increase g, C, or both?

AI is the most interesting perturbation to the paper's headline claim because it can change the funnel at multiple stages.

There are two primary levers:

- Increase C: expand the number of active contributors by lowering barriers to entry.
- Increase g: increase the fraction of time that contributors can spend in true breakthrough-mode by collapsing overhead.

AI could increase C by:

- tutoring more people into technical competence faster
- translating across subfields (physics <-> statistics <-> computation)
- enabling smaller teams to do work previously requiring large groups
- lowering the friction of literature navigation and method acquisition

AI could increase g by:

- automating routine pipeline work (data cleaning, regression tests, monitoring)
- accelerating error checking and systematic exploration
- generating candidate models and diagnostics quickly, allowing more time for conceptual integration
- helping maintain reproducible analysis environments and audit trails
- compressing “verification time” so that human effort shifts toward higher-level judgment

But there is a catch: AI can also *decrease* g_true even while increasing g_attempt. It can make it easier to generate plausible-sounding novelty, which increases attempted-breakthrough activity. If that novelty is not constrained by strong verification, the survival fraction s can drop. In the decomposition:

g_true = g_attempt * s

AI may increase g_attempt while decreasing s unless the field builds stronger gates.

So the net effect depends on governance and practice:

- If AI primarily amplifies ideation without strengthening verification, the field may experience more churn, more false leads, and more narrative volatility. That increases apparent “frontier motion” but not true breakthroughs.
- If AI primarily collapses overhead and strengthens audits, then g_true can rise and the breakthrough process can become less spiky, or at least more frequent.

A second-order possibility is that AI changes what counts as a contributor. If AI enables robust, distributed participation—through shared pipelines, open data, reproducible notebooks, standardized tests—then the breakthrough process could become more networked. C could rise substantially while B (breakthrough-capable) also rises because access to tools and validation frameworks becomes more universal.

The “centaur” framing captures the most realistic near-term outcome: humans remain responsible for judgment, framing, and epistemic hygiene, while AI accelerates the toolchain. Under centaur conditions, the likely outcome is:

- C increases modestly
- g_attempt increases substantially
- g_true increases only if verification gates are improved
- the field becomes more bursty in output unless it deliberately counteracts churn

This has a direct consequence for the paper’s headline. The sub-ppb claim is not an eternal constant. It is a statement about a regime: strict definitions, current institutional structures, and current overhead burdens. AI is the most plausible force that could push the regime upward. Whether it does so in a truth-preserving way is the central epistemic challenge of the next era: increasing the numerator without turning the field into a novelty factory.

In that sense, the paper’s ppb framing becomes not only descriptive but also prescriptive: it invites cosmology to adopt practices that make any AI-driven increase in C or g translate into genuine, surviving breakthroughs rather than amplified illusion.

9. Objections and Replies

9.1 "This is elitist." Response: denominators are not value judgments

Objection: The paper sounds like it is saying only a tiny elite "matters," because it highlights how little of humanity’s cognition is devoted to breakthrough cosmology. The phrase "less than one thought in a billion" can feel like a rhetorical pedestal.

Reply: This is a category error. The paper’s central statistic is a *scaling statement*, not a value statement. It is about the size of a numerator under strict definitions divided by a species-wide denominator. Denominators are not moral rankings. Saying "sub-ppb of cognition" is not saying "sub-ppb of people are important."

If anything, the denominator choice expresses the opposite of elitism: it refuses to shrink humanity to "educated adults" or "internet discourse" or "scientists." It counts the whole species. The statistic becomes small because human life is vast and filled with real cognition—work, love, survival, care, play—not because the paper is assigning worth.

A good check is to swap in other rare activities. "Less than one part per billion of human awake time is spent doing neurosurgery" would also be true without implying that neurosurgeons are more worthy than parents, nurses, teachers, or plumbers. Rarity is not rank. The paper's point is that "breakthrough cosmology" is a rare cognitive mode, not that the humans outside it are lesser.

9.2 "Breakthroughs happen outside academia." Response: include them; still rare
Objection: Academia is not the only source of breakthroughs. Independent thinkers, cross-discipline physicists, and outsiders sometimes produce transformative ideas. The funnel model may be biased toward institutional cosmology.

Reply: Agreed—breakthroughs can occur outside academia, and the paper's framework can incorporate them cleanly. The definition of the numerator is not "credentialed cosmologists." It is "new-to-world, field-shaping cosmology that survives constraint checking." If an outsider produces such a contribution, the model counts it.

But including outsiders does not necessarily inflate the numerator enough to escape the sub-ppb regime, for two structural reasons.

First, the *translation burden* remains. Cosmology breakthroughs must hook into the field's constraints: data, systematics, inference methods, reproducibility. Outsiders can propose ideas, but ideas become field-shaping only if they are tested, integrated, and accepted (or at least must-be-addressed) by the field. That integration is expensive and rare.

Second, the *rate* of outsider breakthroughs that survive is limited by the same bottleneck that limits insider breakthroughs: the number of ideas that can survive adversarial checks is small. Outsider contribution changes who might be in B (the breakthrough-capable subset), but it does not erase the fact that g_true is small and N is enormous.

So the best reply is: widen the model to include outsiders explicitly, but do not assume that doing so changes the order of magnitude. It may nudge parameters upward, but the structural rarity remains.

9.3 "Cognition about cosmology includes contemplation, not just research."
Objection: You are missing vast amounts of cosmology cognition. People contemplate the cosmos, pray about creation, wonder about infinity, read popular books, and have genuine metaphysical engagement. That is "about cosmology" in a meaningful sense, even if it is not research.

Reply: This objection is correct if "cosmology" is defined broadly as "thinking about the universe." Under that definition, the numerator grows dramatically and the headline becomes false.

But that is not a refutation; it is a definitional fork. The paper is explicitly about *breakthrough-level cosmology in the research sense*: contributions that change the field’s constraint-structured knowledge. Contemplation may be profound and spiritually important, but it is not what the paper is measuring.

There is a useful reconciliation: the word "cosmology" names at least two distinct domains:

- cosmology_science: constraint-heavy inference about the universe’s structure, history, and parameters
- cosmology_contemplation: existential, theological, philosophical, and aesthetic engagement with the cosmos

The paper’s claim is about cosmology_science and, within that, about the breakthrough tail. It is compatible with the idea that cosmology_contemplation is widespread. In fact, one could argue that cosmology_contemplation is precisely why a tiny amount of cosmology_science can have disproportionate cultural impact: the public is already primed with wonder; the field supplies disciplined constraints that shape the wonder’s narrative.

So the correct reply is: yes, contemplation is real and widespread, but it is a different numerator. If the goal is to measure contemplation, we need a different model and probably different proxies (media consumption, self-report surveys, religious reflection measures). The present paper’s statistic should not be imported into the contemplative domain.

9.4 "Time-share is not thought-share." Response: propose measurable alternatives
Objection: A minute of cognition is not a minute of cognition. Some minutes contain dozens of discrete thoughts; others contain few. Breakthroughs can occur in seconds. Counting time undervalues high-density ideation and overvalues slow, routine work. Therefore the proxy fraction_of_thoughts ~= fraction_of_cognitive_time is invalid.

Reply: The objection is philosophically fair, but operationally incomplete. If time-share is rejected, we need an alternative that is both more faithful and more measurable. Otherwise we have replaced a tractable proxy with an ungrounded intuition.

Here are several alternative metrics, each with tradeoffs:

(1) **Event-count proxy**: Count "breakthrough events" per year and assign each event a standardized cognitive weight.

Problem: defining an event and its weight reintroduces arbitrariness. Also, events have long preparation tails.

(2) **Information-weighted proxy**: Weight cognition by model update magnitude ("bits of posterior shift").
Problem: in practice, quantifying posterior shift across models is complex, and it risks overweighting wrong-but-seductive novelty unless validation is included.

(3) **Output-weighted proxy**: Use published outputs as tokens (papers, code releases, catalogs) and estimate the share that qualifies as field-shaping.
Problem: publication is not cognition; it is a filtered artifact. But it is measurable and can bound the process.

(4) **Attention-weighted proxy**: Weight by how many other researchers' hours are redirected by a contribution (e.g., follow-on projects).
Problem: this becomes impact-share, not cognition-share. It answers a different question.

The paper's stance is pragmatic: time-share is imperfect, but it is the least arbitrary first-order proxy, and it can be cross-checked by other proxies. The goal is not to claim "time equals thought." The goal is to claim "time gives us a defensible sampling measure" and then show that multiple independent proxies still indicate extreme rarity under strict definitions.

If a critic prefers an alternative metric, the paper welcomes it—but insists it must be defined clearly enough to generate bounds, not just rhetoric.

9.5 "The claim can't be tested." Response: partial tests and bounding arguments
Objection: You cannot measure "fraction of human thoughts that are breakthrough cosmology." Therefore the claim is unfalsifiable and closer to poetry than science.

Reply: The claim is not directly measurable to high precision, but it is not untestable. It is a *bounded* claim. We can partially test it by triangulation and by establishing upper bounds and plausibility ranges.

Several partial tests are available:

(1) **Bounding C and B by observable headcounts.**
We can estimate the size of the active cosmology community using author lists, arXiv categories, major collaborations, and departmental rosters. This provides plausible ranges for C and B. Even crude bounds constrain the numerator strongly because N is enormous.

(2) **Bounding g by time-use studies and workflow audits.**
Within cosmology groups, one can estimate the fraction of time spent on deep ideation versus pipeline work, meetings, writing, and verification. Even if "breakthrough-mode" is hard to label,

one can bound it above: it cannot exceed total deep work time, and deep work time itself is often a minority fraction.

(3) **Bounding breakthrough frequency via bibliometrics and field-turn indicators.**
Count how many contributions per decade plausibly qualify as field-shaping by agreed proxies (top-tail uptake, canonical review inclusion, method adoption). This bounds how much true breakthrough cognition can exist, because breakthroughs are not continuous.

(4) **Upper-bound argument from "breakthrough-hours required."**
If we can estimate the approximate human-hours embodied in a breakthrough (including ideation + verification + integration), then the product of (breakthroughs per year) * (hours per breakthrough) yields an estimate of total breakthrough-mode hours per year. Divide by global awake-hours. This can generate a quantitative bound even if the hours estimate is crude.

These are not precise measurements, but they are not arbitrary either. They can falsify extremes. For example, if someone claimed "breakthrough cosmology is 1 ppm of human cognition" (1e-6), that would imply an enormous number of breakthrough-hours worldwide—so large that it would predict either a vastly larger cosmology workforce or vastly more frequent field turns than we observe. Conversely, if someone claimed "1e-15," they would be implying essentially no effective breakthrough cognition, which conflicts with the observable existence of field-shaping progress over time.

So the right framing is: the claim is a regime claim supported by bounded reasoning and cross-check proxies. It is not a point estimate with three decimals. It can be contested and refined, but it is not immune to evidence. The paper invites exactly that: sharpen the definitions, choose proxies, set bounds, and see where the weight of constrained inference places the fraction.

10. A Small Empirical Program to Anchor the Numbers

10.1 Estimate C: triangulate from memberships, author lists, and departments

The goal of this program is not to pretend we can measure “thought fractions” directly. The goal is to anchor the key parameters (C, B, g, and breakthrough-event frequency) with observable traces. We start with C: the number of active cosmology contributors under the strict, research-facing definition.

A practical triangulation strategy uses three partially overlapping lenses:

(1) **Membership and collaboration rosters (coarse headcount bounds).**
Professional societies and major collaboration rosters provide an outer envelope. This is not a clean “cosmology list,” but it bounds the scale of the relevant workforce. The value is that rosters are curated and persistent, so they help avoid double-counting and transient noise.

(2) **Author-list extraction from a defined cosmology corpus (behavioral inclusion).**
Define a cosmology corpus with clear criteria (e.g., arXiv categories + selected journals + selected keywords with manual filtering). Then extract all authors in a rolling time window (e.g., last 3 years) and count those who meet an activity threshold (e.g., at least k contributions in the window, or at least one significant role paper plus ongoing involvement). This can estimate active contributors rather than one-off participation.

(3) **Departmental and research-group mapping (institutional cross-check).**
Sample a representative set of institutions across geographies and tiers. Estimate the number of cosmology-facing faculty, postdocs, graduate students, research staff, and pipeline engineers. Extrapolate with appropriate uncertainty. This provides a sanity check against author-based counts, which can be distorted by mega-collaboration authorship practices.

The key is not to force agreement. It is to generate a credible range:

- C_low: a strict interpretation (frontier analysis + core theory + cosmology-facing pipeline leads)
- C_high: a generous interpretation (includes broader survey teams and cosmology-facing infrastructure roles)

The output of 10.1 should be a range, not a point: C in [C_low, C_high], with explicit notes about boundary choices.

A refinement: introduce B explicitly. Many contributors will not be positioned for field-shaping breakthroughs. Use proxies for "breakthrough-capable positioning" such as sustained centrality in the author graph, leadership roles in major analyses, and repeated presence in high-impact method or result papers. This yields an estimated fraction f = B/C and therefore a range for B as well.

10.2 Estimate g: time-use surveys for researchers (deep work vs overhead)
Estimating g is harder because it is about cognitive mode, not mere activity. The empirical program therefore proceeds in layers, each providing a bound.

Layer 1: **Work fragmentation and overhead accounting**
Survey or diary studies can estimate time spent in:

- meetings, email, administration, teaching, proposal writing
- code maintenance, debugging, documentation
- running standard analyses and pipelines
- reading and literature review

- writing papers and responding to reviews
- deep work blocks (uninterrupted analysis/derivation/design)

Even if we never label “breakthrough-mode,” this already bounds it:

g_true <= fraction_of_time_in_deep_work

Because true breakthrough-mode cannot exceed time spent in deep work.

Layer 2: **Attempted breakthrough-mode tagging (self-report, with skepticism)**
Within deep work blocks, researchers can self-report what fraction felt like:

- incremental execution (“finish pipeline,” “run standard fit,” “write section”)
- exploratory ideation (“new diagnostic,” “new model comparison,” “new systematic hypothesis,” “new observable idea”)

This yields g_attempt (attempted breakthrough-mode share) as a fraction of deep work time. It is imperfect—people overestimate the novelty of their own ideas—but it provides a structured estimate.

Layer 3: **Survival calibration (external validation)**
To estimate g_true from g_attempt, introduce a survival fraction s: the fraction of attempted breakthrough-mode time that results in contributions later judged to be field-shaping under agreed criteria. This can be calibrated retrospectively by taking known field-shaping events and reconstructing the approximate time investment and the number of failed attempts that preceded them.

This yields:

g_true ~= g_attempt * s

The program should not pretend s is stable across people or eras. It should treat s as heavy-tailed and uncertain, providing ranges rather than fixed values.

Crucially, the program must separate:

- *personal novelty* (new-to-me)
- *publishable novelty* (new-to-literature but incremental)
- *field-shaping novelty* (breakthrough-level)

g_true targets only the last category.

10.3 Estimate breakthrough events: field-turn metrics (method and caveats)
The “ppb cognition share” claim depends on how often the field produces events that truly reshape it. We therefore need a way to estimate breakthrough events per unit time.

Define a “field turn” metric set. No single metric is sufficient, so use a portfolio:

(1) **Method adoption metric**
Count methods/pipelines that become standard practice (e.g., widely used likelihood frameworks, standard map-making procedures, canonical systematics corrections). Identify adoption by widespread reuse across groups and by inclusion in review/tutorial literature.

(2) **Parameter/regime shift metric**
Identify events where the accepted parameter regime or baseline model narrative shifts (or where a persistent tension becomes organizing). Use review articles, major survey synthesis papers, and consensus statements as signals.

(3) **Canonical reference metric (top-tail bibliometrics)**
Use a field-normalized citation tail as a proxy, but only as one signal, because citations can reflect sociology.

(4) **Instrumentation-driven turn metric**
Count turns triggered by new observational regimes (new CMB experiments, new large surveys) where the inference landscape changes materially.

Each candidate “turn” should be graded with a rubric aligned to the strict criteria in Section 2.5:

- new_to_world
- constraint_surviving
- field_shaping
- irreducible_impact
- counterfactual dependence

The output is not “the exact number of breakthroughs.” It is a bounded range:

breakthrough_events_per_year in [E_low, E_high]

Caveats must be explicit:

- turns can be distributed, not attributable to one paper
- recognition lags and can be biased
- some breakthroughs are “quiet infrastructure” and under-recognized

- some highly visible events are narrative churn, not genuine turns

But even with caveats, a plausible range for E is constraining because it limits the total amount of true breakthrough cognition the field can contain.

10.4 Combine into confidence intervals: ranges, not point estimates

Once we have ranges for C (and possibly B), g_true, and breakthrough-event rates, we can combine them into an interval for the population-wide cognition share.

The core expression is:

share ~= (B / N) * g_true

But because B and g_true are uncertain and heavy-tailed, the right output is an uncertainty band, not a point:

share in [share_low, share_high]

A practical way to do this without pretending to precision is:

- choose conservative low/high values for B and g_true consistent with the empirical anchors
- compute the implied share range
- report the range on a log scale (orders of magnitude), not as a single decimal number

For example, an output might look like:

- share ~ 1e-11 to 1e-9 under strict definitions (illustrative)

The program can also output conditional ranges:

- under strict "field-shaping" criteria
- under looser "publishable novelty" criteria
- under different denominator choices (all awake time vs working adult time)

This is important because it makes clear that the ppb claim is conditional on definitions rather than a metaphysical absolute. The payoff is transparency: readers can see exactly which assumptions move the number.

A further improvement is to treat the model explicitly as heavy-tailed:

- do not assume uniform g across B
- model g as a distribution (many near-zero, few high)

- compute median and tail bounds rather than only mean

Even if the paper does not implement full statistical modeling, it should conceptually acknowledge that “confidence intervals” here are best understood as robust bounds rather than classical frequentist intervals.

10.5 Falsification conditions: what would force revision upward

A claim that cannot be forced upward by evidence is rhetoric, not analysis. The paper therefore specifies conditions under which the sub-ppb conclusion would need to be revised upward (toward or above 1e-9).

Because:

share ~= (B / N) * g_true

and N is large and relatively fixed, upward revision requires credible evidence that B is much larger, or g_true is much larger, or both.

Here are concrete falsification / revision triggers:

(1) **A much larger B is demonstrated**

If empirical triangulation shows that the number of breakthrough-capable contributors is not “a few thousand” but “tens of thousands” globally—meaning many more people are repeatedly positioned to produce field-shaping work (data access, pipeline authority, sustained frontier activity)—then the numerator increases substantially. Under that condition, even modest g_true could push the share above 1 ppb.

Operational evidence: sustained core-author graph size at the frontier is far larger than assumed, and instrument/survey access is widely distributed rather than concentrated.

(2) **g_true is empirically shown to be at least an order of magnitude larger**

If time-use studies plus retrospective calibration show that true breakthrough-mode cognition occupies, say, ~0.5% to 2% of breakthrough-capable researchers’ cognitive work (rather than ~0.01% to 0.1%), then share rises by 1–2 orders of magnitude.

Operational evidence: deep work time is high, exploratory ideation is a large fraction of deep work, and the survival fraction s is not tiny (i.e., many exploratory hours routinely produce field-shaping outcomes).

(3) **Field-turn frequency is far higher than the model allows**

If an agreed “field-turn” rubric indicates that cosmology experiences genuine field-shaping turns frequently—say, dozens per year—then either B or g_true (or both) must be larger than assumed because the system is producing many more surviving breakthroughs than the model predicts.

Operational evidence: multiple independent proxies (method adoption, regime shifts, instrumentation turns, canonical uptake) all indicate high turn-rate, and those turns are not merely narrative churn.

(4) **AI-driven overhead collapse is observed with preserved epistemic hygiene**
If AI systems demonstrably reduce overhead while maintaining or improving adversarial verification—so that researchers spend substantially more time in true breakthrough-mode and the survival fraction does not drop—then g_true could rise materially. This is a real pathway to upward revision.

Operational evidence: measurable decline in time spent on maintenance/verification per published claim, without increased retraction/correction rates or rising instability; plus faster convergence on robust results.

(5) **Denominator redefinition is justified and adopted**
If the community agrees that the relevant denominator is not all human awake cognition but a narrower, defensible subset (e.g., "global scientific cognition time"), then the share rises by definition. This is not falsification of the original claim; it is a different claim. But it would force the paper to restate its headline with the new denominator.

Finally, note a crucial asymmetry: many plausible criticisms weaken the precision of the estimate, but fewer plausibly force a large upward revision under strict definitions. That is the robustness thesis of this paper. Still, the falsification conditions above are real. If future empirical anchors land there, the paper's headline must move. That is not a failure; it is the point of offering a small empirical program: to keep the claim contestable rather than rhetorical.

11. Conclusion

11.1 Thesis restated: sub-ppb is the natural regime under strict definitions

This paper has argued for a claim that initially sounds paradoxical: if we define our target narrowly as breakthrough-level / new-to-world / field-shaping cosmology, and if we define our denominator broadly as the cognitive life of the entire human population, then the share of human cognition devoted to that target naturally falls below one part per billion. The conclusion is not a fragile point estimate. It is a regime statement produced by multiplicative filters.

Under strict definitions, the numerator collapses in stages:

- only a small fraction of humans are cosmology-capable in the research sense
- only a smaller fraction are active contributors at the frontier
- only a smaller subset are positioned to generate field-shaping breakthroughs

- only a small fraction of even that subset's cognitive work is true breakthrough-mode rather than overhead, verification, and incremental progress

When these filters are made explicit, sub-ppb does not require cynical assumptions. It follows from the scale mismatch between a species-wide denominator and a narrow, constraint-heavy frontier practice. Even generous parameter choices often remain in the sub-ppb region unless one broadens "cosmology," broadens "breakthrough," shrinks the denominator, or assumes a future regime where overhead collapses without loss of rigor.

11.2 What is gained: clarity about rarity without mystique

The value of the ppb framing is not the number itself. It is the kind of clarity it forces.

First, it separates **generation** from **reproduction**. Cosmology's cultural presence can be vast while the generating cognitive stream remains tiny. This helps explain why frontier science can feel omnipresent and yet be produced by a relatively small community operating under heavy constraints.

Second, it separates **rarity** from **rank**. A sub-ppb share is not a moral hierarchy and not a declaration of who matters. It is a scaling description. The cognitive life of humanity is full; cosmology is rare within it because most human cognition is rightly spent on living.

Third, it separates **awe** from **method**. Cosmology is naturally tempting as a cultural oracle because it speaks to ultimate questions. The ppb framing can de-mystify that temptation by reminding us that cosmology's authority is not priestly but methodological. When a generating stream is tiny, we should demand more transparency, not more deference.

Finally, it provides a disciplined way to talk about what "breakthrough" really means. Many ideas feel new; fewer survive. The paper's strictness is a form of respect: it honors the field's adversarial commitment to error bars, systematics, and reproducibility by not confusing attempted novelty with surviving change.

11.3 Closing charge: honor the work, police the hype, measure what we can

If breakthrough cosmology is rare in the strict sense, the right response is not mystique and not cynicism. It is stewardship.

Honor the work by recognizing what it costs: years of constraint accumulation, expensive instruments, brittle pipelines, careful inference, and the ethical discipline of not claiming more than the data allow. Honor also the invisible labor that makes breakthroughs possible: calibration, verification, documentation, and the slow grind of reproducibility. Those are not distractions from truth; they are how truth is paid for.

Police the hype because rarity invites amplification. Media, institutions, and even scientists can be tempted to convert provisional frontier claims into existential headlines. The public deserves better. The integrity of cosmology depends on carrying uncertainty forward honestly, resisting oracle politics, and refusing to let “big questions” become shortcuts around rigor.

Measure what we can because contestability is a virtue. The ppb claim should not be treated as a slogan. It should be treated as a bounded estimate under explicit definitions. The empirical program sketched here—triangulating active contributor counts, estimating overhead versus deep work, and tracking genuine field turns—will never yield perfect precision, but it can keep the conversation anchored. We do not need exactness to avoid nonsense; we need transparent assumptions and honest bounds.

The closing ethic is simple: treat cosmology as a disciplined witness, not as an oracle; treat rarity as a scaling fact, not as a hierarchy; and treat measurement as an act of humility. If we do that, we can hold two truths at once without contradiction: breakthrough cosmology may occupy less than one thought in a billion, and yet it can still reshape how humanity understands its place in the universe.

References

American Astronomical Society (AAS). (2025). *About the AAS*. (Accessed Dec 28, 2025.) aas.org

American Astronomical Society (AAS). (2025). *AAS Names 24 New Fellows for 2025*. Press release (Jan 13, 2025). aas.org

Barabasi, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512. doi:10.1126/science.286.5439.509

Bloom, N., Jones, C. I., Van Reenen, J., & Webb, M. (2020). Are ideas getting harder to find? *American Economic Review*, 110(4), 1104–1144.

Clauset, A., Shalizi, C. R., & Newman, M. E. J. (2009). Power-law distributions in empirical data. *SIAM Review*, 51(4), 661–703. (arXiv preprint: arXiv:0706.1062.) SNAP

Decker, W. H., et al. (2007). *Faculty Workload Survey: Survey Report with Appendices* (Federal Demonstration Partnership). The Federal Demonstration Partnership

Fortunato, S., Bergstrom, C. T., Borner, K., Evans, J. A., Helbing, D., Milojevic, S., et al. (2018). Science of science. *Science*, 359(6379), eaao0185. doi:10.1126/science.aao0185 research.ceu.edu

International Astronomical Union (IAU). (2025). *Individual and Junior Members* (Directory summary counts). (Accessed Dec 28, 2025.) iau.org

Jones, B. F. (2009). The burden of knowledge and the "death of the renaissance man": Is innovation getting harder? *Review of Economic Studies*, 76(1), 283–317.

Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.

Lakatos, I. (1978). *The Methodology of Scientific Research Programmes*. Cambridge University Press.

Lotka, A. J. (1926). The frequency distribution of scientific productivity. *Journal of the Washington Academy of Sciences*, 16(12), 317–323.

Mark, G., Gudith, D., & Klocke, U. (2008). The cost of interrupted work: More speed and stress. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. doi:10.1145/1357054.1357072 ACM Digital Library+1

Merton, R. K. (1968). The Matthew effect in science. *Science*, 159(3810), 56–63. doi:10.1126/science.159.3810.56

OECD. (2015). *Frascati Manual 2015: Guidelines for Collecting and Reporting Data on Research and Experimental Development*. OECD Publishing. OECD+1

Planck Collaboration. (2020). Planck 2018 results. VI. Cosmological parameters. *Astronomy & Astrophysics*, 641, A6.

Popper, K. (1959). *The Logic of Scientific Discovery*. Hutchinson.

Price, D. J. de Solla. (1963). *Little Science, Big Science*. Columbia University Press.

Redner, S. (1998). How popular is your paper? An empirical study of the citation distribution. *European Physical Journal B*, 4, 131–134. (arXiv:cond-mat/9804163.) arXiv

Rockwell, S. (2009). The FDP faculty burden survey. *Research Management Review*, 16(2). The Federal Demonstration Partnership+1

Simms, C. (2025). Pressure to publish is rising as research time shrinks, finds survey of scientists. *Nature*, 647(8090), 565–566. doi:10.1038/d41586-025-03623-2 Nature+1

United Nations, Department of Economic and Social Affairs, Population Division. (2024). *World Population Prospects 2024: Summary of Results*. World Population Prospects

United Nations, Department of Economic and Social Affairs, Population Division. (2024). *World Population Prospects 2024 (online portal)*. World Population Prospects

United Nations Statistics Division. (2024). *SDG Indicator Metadata: 9.5.2 Researchers (in full-time equivalent) per million inhabitants*. UNSD

U.S. Census Bureau. (2025). *Population Clock: World*. Census.gov

World Bank. (2025). *Population, total (SP.POP.TOTL)*. World Bank Open Data+1

World Bank. (2025). *Researchers in R&D (per million people) (SP.POP.SCIE.RD.P6): indicator definition and metadata*. DataBank+1

Wu, L., Wang, D., & Evans, J. A. (2019). Large teams develop and small teams disrupt science and technology. *Nature*, 566, 378–382. doi:10.1038/s41586-019-0941-9 Nature+1

York, D. G., et al. (2000). The Sloan Digital Sky Survey: Technical summary. *The Astronomical Journal*, 120(3), 1579–1587.

Uzzi, B., Mukherjee, S., Stringer, M., & Jones, B. (2013). Atypical combinations and scientific impact. *Science*, 342(6157), 468–472. doi:10.1126/science.1240474 researchgate.net