Can the Theory Carry Gravity? — An Honest Plain-Language Account — Fable_Version rendered package. Rendered from `QG_HUMAN_CLOSURE_COMPENDIUM.md` ; frozen technical content unchanged by rendering.

Can the Theory Carry Gravity? — An Honest Plain-Language Account

Article 2: A Quantum Theory of Gravity, in the spirit of *In Search of Nature's Elegance*

> **What this is.** A plain-language companion to the technical "Quantum Theory of Gravity" closure compendium. It carries the *same* honest conclusions as the technical version, in words a careful non-physicist can follow start to finish. It proves nothing new, computes no new number, and upgrades no claim. It reports a *status* for each of twelve questions — and it reports each one exactly as cautiously as the technical document does. **The frozen object it discusses** (the "13-dimensional geometry," labeled `dcc66f1b2685`) is **read-only**. Nothing here touches it. **Promotions: 0.** Translation: nothing in this document gets moved from "we think" to "we proved," or from "we declined the question" to "we answered it."

The one move we refuse to make

There is a tempting move that a less careful version of this work would make, and it sits at the very heart of the gravity story. So we name it up front, in plain words, and promise not to make it:

> *We will never let "we stopped the question from being asked" sound like "we answered the question."*

This document leans hard on a single idea — call it the **graininess rule** — that makes a whole class of famously hard problems *go away*. But "go away" is doing very specific, very limited work here, and the entire honesty of the document is in keeping that work limited. When a problem goes away under the graininess rule, it does **not** mean we solved it. It means we declined a particular impossible version of it — and a smaller, finite version of the same problem is usually still sitting there, unsolved, in plain sight. We will say "dissolved" for the first thing and "still open" for the second, and we will never let the two blur together.

If you remember only one sentence from this entire document, make it this one:

> **Dissolved is not solved. Declining the impossible half of a question is honest progress; it is not an answer to the half that remains.**

This document is written to survive a *hostile* reading — the kind of reviewer who trusts nobody and hunts for the exact place where "we dodged it" got dressed up as "we cracked it." So we will flag, for every single one of the twelve questions, which of these it is:

- **closed by an axiom** ("this part just *is*, we posited it"),
- **derived, but only once we accept which particles exist** (it follows from the geometry — given the menu of real particles as input),
- **dissolved** (we decline the impossible idealization — *not* the same as solving it),
- **partly there** (one half done, the other half named and open), or
- **still open** (genuinely unfinished — and here's exactly what's missing).

And at the end, four questions remain **flatly, honestly open** even after we grant ourselves every tool we have. We list them by name. They are the ones we are most tempted to fudge, which is exactly why we don't.

---

1. The graininess rule, and what it can and cannot do

To follow the rest of this document you need one idea, and only one. Here it is in plain words.

The rule

The theory is built on a small handful of starting assumptions ("posits" — things we accept without proof because the whole structure rests on them). One of them is the important one for gravity:

> **Graininess.** The world is not infinitely divisible down to a perfect mathematical point. Physics is always evaluated on a smallest *operational cell* — a finite grain. There is no free pass to "let the grain shrink to nothing."

Physicists have a name for the move this rule forbids. In the usual way of doing things, you check that a theory makes sense by imagining the grain shrinking all the way to zero — a perfect, seamless continuum — and demanding the theory stay well-behaved the whole way down. That "shrink to zero" demand (`a → 0`, in the trade) is the source of most of the deepest trouble in quantum gravity. The graininess rule simply declines to make that demand. With a finite grain, there is no "all the way to zero" to worry about.

What the rule *can* do — and what it absolutely cannot

This is the crux of the whole document, so let's be painfully precise.

What graininess CAN do. It can make the *infinite* version of a hard question stop being a live demand. "Does the theory still make sense as the grain shrinks to nothing?" is a question about an infinite limit. If you've decided the grain never shrinks to nothing, that question is no longer something you're on the hook to answer. It is **dissolved** — removed from the table.

What graininess CANNOT do. It cannot answer the *finite* version of the question, the version that's still about real, grainy physics. Removing "what happens in the perfect continuum?" does nothing about "does the leftover finite quantity actually come out right?" That finite leftover is a definite question with a definite answer that someone still has to work out — and the graininess rule does not work it out. It is **still open**.

So almost every hard gravity question in this document splits cleanly into two halves:

- the **continuum half** ("does it survive shrinking the grain to nothing?") — **dissolved** by graininess; and
- the **finite half** ("is the actual finite quantity right?") — **still open**, untouched by the rule.

A hostile reviewer's first instinct will be: *"You've just defined your hardest problems out of existence."* The honest answer is: no — we've defined *one half* of them out of existence (the impossible-idealization half), and we keep the other half (the finite, real half) on the table, openly unsolved. Watch us do exactly that, twelve times.

One more guard, because it matters: there are a few questions where graininess is simply **the wrong tool**. The graininess rule is about *cost and finiteness*. Some questions aren't about cost or finiteness at all — they're about probabilities, or about a specific cancellation, or about a structural bookkeeping check. For those, the rule has nothing to say, and we will not pretend it does. It neither dissolves them nor solves them. They stay exactly as open as they were.

The twelve questions

A real quantum theory of gravity has to clear twelve distinct hurdles. We take them one at a time. For each, we give: what's being asked, where it stands *on its own*, where it stands *once we grant the graininess rule and the rest of the theory's assumptions*, and the exact catch.

Question 1 — Is there a graviton? (the carrier of gravity)

What's asked. Gravity, if it's a quantum force, needs a carrier particle — the "graviton" — just as light has the photon. It has to be massless, travel at the right speed, have exactly the two allowed spin-states, and reproduce ordinary Newtonian gravity when things are weak and slow.

On its own. Partly there — a conditional certificate. For *weak* gravity (the gentle, "linearized" regime), this is genuinely settled and matches the textbooks: the construction produces exactly two physical spin-states, massless, with the right speed and the correct Newtonian pull. This is a real success — it's the place where the theory shakes hands with known physics. The catch is in the word *conditional*: the weak-gravity certificate leans on Question 2 (what happens at the highest energies). If Question 2 fails, this certificate loses its standing.

Once we grant the theory. Same — still a conditional certificate, no upgrade. The two-spin-state count and the Newtonian limit hold in the weak regime, full stop. The *strong*, interacting graviton — gravitons slamming into each other hard — routes straight into Question 2. The graininess rule dissolves the impossible-continuum half of Question 2, but it leaves a specific finite number uncomputed (we'll meet it as "the a_6 number"), and the strong graviton needs that number. So nothing here gets upgraded: it's a weak-gravity certificate resting on a finite leftover the theory has not supplied.

The exact catch. Weak gravity only. The interacting, high-energy graviton is **not** certified. No claim is made about gravitons scattering above the theory's cutoff energy.

Question 2 — Does the theory make sense at the highest energies? (UV completion)

This is the headline question of quantum gravity, and the cleanest illustration of the whole dissolved-vs-solved idea. Read this one slowly.

What's asked. Show the theory stays sensible at *arbitrarily* high energy — that it doesn't blow up or become gibberish no matter how hard you push it. The standard way to demand this is to require a special well-behaved "fixed point" that survives all the way down to a perfect continuum.

On its own. Still open. Flatly. Whether that special high-energy fixed point even exists is a famous *unsolved* problem in all of physics, not just here. The paper does not claim to have it.

Once we grant the theory. Dissolved (the impossible half) — and still open (the finite half). Here is exactly what that means, with nothing hidden:

- The *demand* was: "prove the theory is well-defined at arbitrarily high energy, all the way down to a perfect zero-grain continuum." The graininess rule **declines that demand outright** — if the grain never shrinks to zero, there is no "arbitrarily high energy / perfect continuum" regime to police. So the question "is there a continuum fixed point?" stops being a live requirement. It is **dissolved.** *Note carefully what this is not:* we have **not** exhibited the fixed point, and we have **not** proven the continuum theory is fine. We have **removed the question**, which is a different and weaker thing.
- But a **finite leftover survives**, and it is real. Even a grainy theory must pass a specific finite consistency check — a single, definite number (physicists call it the **a_6 coefficient**) that the theory's own machinery demands. That number is **not computed.** It can in principle be computed — it's a definite, if heavy, calculation — but no one has done it, and the theory does not hand it to you for free. So the finite half is **still open**, with a name and a known method but no result.

The exact catch. Dissolved is not solved: no high-energy fixed point is exhibited; we only removed the demand for a perfect continuum. The finite a_6 number is a definite, heavy calculation that has **not** been done — and we never invent it. In fact a falsifier rides on it: if someone computes a_6 and it comes out the wrong sign, a load-bearing claim of the theory is refuted on the spot. Also, a strict guard: nothing in this high-energy story may be read as solving the separate "cosmological constant" puzzle — that one stays open by its own admission, and two attempts to close it have actually *failed*. We keep that failure visible.

Question 3 — Do all physical states have sensible (non-negative) probabilities? (positivity)

What's asked. In a sane quantum theory, every genuine physical state must have a non-negative "size" — informally, you can't have a state with negative probability. The machinery has to guarantee this, and weed out the bookkeeping ghosts (helper states that aren't physical) so only real states remain.

On its own. Partly there. For the free and gently-interacting sectors, this is genuinely demonstrated — the ghost states are properly removed and what's left has sensible sizes. The hard, *strongly-interacting* sector (the same beast as the confined-quark sector of nuclear physics) is not settled on its own; it leans on Question 11.

Once we grant the theory. Conditional certificate for the gentle part; the strong part stays open even here. The free and perturbative positivity is a real certificate (conditional on Questions 4 and 9). But — and this is important — the graininess rule does **not** rescue the strongly-interacting half. Proving positivity for a full nonperturbative theory is a genuinely different kind of object, exactly as hard as the corresponding open problem in ordinary 4-dimensional nuclear physics. A finite grain does not, by itself, prove the grainy theory has sensible probabilities in the strong regime. So that half is **still open at both levels**.

The exact catch. Gentle-regime grade only. If a negative-probability state ever survived the cleanup, this question and several others lose their certificates. The strongly-interacting positivity certificate is **not** supplied by the theory.

Question 4 — The a_6 number, on its own (the finite consistency coefficient)

What's asked. Compute the one specific finite number that Question 2 leaves behind — the "a_6 coefficient" — for this exact 13-dimensional geometry, and check that it has the right (positive) sign for the theory's probabilities to stay sensible.

On its own. Still open. The corpus states plainly: it is *not computed.*

Once we grant the theory. Still open — at both levels — for the actual number. The continuum-flavored question ("must infinitely many higher coefficients all also behave?") is dissolved by graininess, same as Question 2. But the a_6 number itself is one specific, finite quantity, and **graininess does not compute it.** No axiom hands you a number. So the headline stays open.

The exact catch — and why we never fake it. Here is the honest texture. The general *recipe* for this number exists and is standard: it's a known combination of a fixed list of building blocks with known fractions in front of them. What's missing is the *specialization* — actually grinding that recipe through this particular 13-dimensional geometry, which is a long, heavy symbolic computation no one has run. The corpus is explicit that fabricating an answer here "would be the worst outcome," so the number is **never invented**. And a separate honesty note: even if someone computes it and it's positive, that is *one* consistency check passing — whether one passing check *amounts to* full high-energy sanity is a further judgment the corpus has deliberately **not** made. The a_6 number is the missing piece *for* the high-energy question, not the high-energy answer itself.

Question 5 — Black hole entropy and the Page curve

What's asked. Two things. (a) Count the hidden internal states of a black hole and show the count matches its famous "area / 4" entropy — *and for the right reason.* (b) Show that when a black hole evaporates, information isn't destroyed — the radiation's entropy follows the expected rise-then-fall "Page curve."

On its own. Partly there — but the honest part is small. There is a banked *consistency check*: a particular sector of the theory reproduces the "area / 4" entropy value to within 0.0028%, and this was checked under freeze-before-compare discipline (you commit to the prediction before looking at the answer). That's a real and pretty number-match. But it is explicitly graded as only a *consistency datum* — it says nothing about *why* there are that many states (no actual count) and nothing about evaporation (no Page curve). Those two are flatly not claimed.

Once we grant the theory. The high-energy leg is dissolved; the count and the evaporation mechanism stay open. The part of this question that flows in from the high-energy problem (Question 2) has its continuum half dissolved by graininess. But the actual *microstate count* (a genuine theorem about *why* that many states) and the *Page-curve mechanism* (how information escapes) are entirely different objects that graininess does not supply. They stay **still open**.

The exact catch. The 0.0028% match is a consistency over-determination, **not** a derivation of black hole entropy, and the disclaimers forbid reading it as a state count. No Page-curve mechanism is claimed. No microstate number is invented. (This question doesn't even have a falsifier of its own yet — precisely because the disclaimers stop it short of a claim that could be falsified.)

Question 6 — Do the curled-up extra dimensions stay stable? (compactification)

This one is the *strongest* result in the gravity story, so we have to be especially careful not to oversell it.

What's asked. The theory hides several extra dimensions, curled up small. Quantum effects could in principle make them unravel, collapse, or sprout unwanted near-massless "wobble" fields that would violate gravity experiments. Show they sit in a stable configuration — a genuine valley they don't roll out of.

On its own. Partly there. The paper's own row defers to Question 2: without the high-energy certificate, it says, the higher-order stability calculation isn't fully controllable.

Once we grant the theory. Derived — the *existence* of a stable valley — though only at "decision grade," and the full certificate still depends on the open a_6 number. This needs unpacking, because "derived" is a strong word and we mean something specific by it.

A *separate* line of the corpus — not the route the gravity paper itself uses — builds the stabilizing valley directly and finds that it exists *unconditionally*: the configuration clears the relevant bound comfortably (by a wide margin in the most demanding variant), the stabilizing well is genuinely there, and an internal sign-check that could have killed it instead confirmed it. By *that independent route*, the question "is there a stable valley?" is answered at "decision grade" — meaning: good enough to act on as established, independent of the unsolved high-energy fixed point.

Now the discipline, stated flatly so no one misreads the win:

- This is the **existence** of stability only — *is there a stable valley?* — answered yes at decision grade.

- It is **not** the paper's own full certificate-grade verification, which is a different, more demanding object that still inherits the open a_6 number from Question 2.
- It is **not a promotion.** The paper *uses* the stable background as a trusted input but does **not** convert it into its own top-grade result. The paper's banked status stays exactly where it was (audit grade). "Decision grade" is a real and useful standing, but it is **not** full-proof-grade or certificate-grade, and we don't let it pretend to be.

So: stronger than the paper's pessimistic "blocked by Question 2" framing — the stable valley genuinely exists by an independent route — but still short of a full certificate, and carrying zero promotions.

The exact catch. Decision-grade *existence* only, via the independent route — **not** the paper's own top-grade trajectory calculation, which remains the paper's open obligation and still carries the unsolved a_6 number. The paper's tier is unchanged.

Question 7 — No faster-than-light, no broken cause-and-effect (unitarity, causality, locality)

What's asked. Probabilities must always add to one (nothing leaks); causes must precede effects; distant things shouldn't instantaneously affect each other. Check these across each sector, including graviton collisions.

On its own. Partly there. Where a sector is gentle, these checks are routine and pass. The graviton sector *above* the cutoff energy leans on Question 2.

Once we grant the theory. Conditional certificate below the cutoff; the high-energy half stays open. Below the cutoff, the checks are a real conditional certificate. Above it, the graviton sector inherits Question 2's situation: the impossible-continuum demand is dissolved by graininess, but a *finite* proof that very-high-energy graviton collisions keep probabilities adding to one is a different object the rule does not deliver — the same finite-leftover class as the a_6 number. So that half is **still open**.

The exact catch. Below-cutoff grade only. High-energy graviton-sector probability conservation is **open**. (And the same guard as before: no high-energy reading is allowed to sneak in a solution to the cosmological-constant puzzle.)

Question 8 — Does the bookkeeping close consistently? (BRST / anomaly cancellation)

What's asked. Two structural checks. First, an internal consistency operation has to "square to zero" (a technical but checkable bookkeeping identity). Second, certain quantum inconsistencies called "anomalies" must all cancel — including after the theory is projected down from 13 dimensions to our 4, through its boundary and its twisted internal structure.

On its own. Partly there. The parent theory's anomaly ledger already cancels; whether that cancellation *survives* the projection down to 4 dimensions and the quantization here requires explicit checking that has been deferred.

Once we grant the theory. Partly there — graininess does not touch this one. This is *not* a continuum-idealization question; it's a finite, structural bookkeeping check. So the graininess rule offers no dissolution at all. The bulk part (the squaring-to-zero and the inherited cancellation) is the closed half; the **boundary** and **twisted-internal** parts of the descent remain **still open**, with a named path, at both levels.

The exact catch. Bulk bookkeeping plus inherited cancellation only; the boundary and twist parts are **not** verified. If an anomaly failed to cancel, several other questions drop a grade and the whole paper is capped at its current grade until it's resolved.

Question 9 — One generation of matter, with the right handedness (fermion chirality)

What's asked. Our universe's matter particles are "left-handed" in a specific lopsided way, and there are exactly three families. Reproduce one full family with the correct handedness, with **no** unwanted mirror-image partners, and with the bookkeeping still closing.

On its own. Partly there. The no-mirror inventory is inherited from the classical level; showing the quantum projection down to 4 dimensions doesn't secretly spawn a light mirror partner needs explicit checking.

Once we grant the theory. The classical count is derived (given the particle menu); the quantum check stays open. The classical inventory is genuinely *derived from the frozen geometry* — read straight off a geometric index that comes out to exactly three families with no mirrors, and cross-checked two independent ways. We flag this as **derived, but only once we accept which particles exist**: it follows from the geometry *given* the input data, and it's twice-confirmed. So the *count* is not the open item. The open item is the *residual*: that the full quantum projection introduces no light mirror partner beyond the classical count. That's a finite, structural check — graininess doesn't dissolve it — so it stays **still open**, with a named path.

The exact catch. Only the *classical* inventory is closed (and only given the observed particle content as input). The *quantum* no-mirror check is open. The status is unchanged — narrowed to the quantum check, **not** promoted. A surviving mirror would downgrade the whole unification claim to "non-chiral candidate."

Question 10 — Where do the Born-rule probabilities come from?

This is one of the four that stays flatly open. We are especially careful here.

What's asked. Quantum mechanics assigns probabilities to outcomes by a specific rule (the "Born rule" — probability goes as the square of a certain quantity). Derive *that rule* from the framework, rather than just assuming it.

On its own. Still open. The corpus honestly halts here — it stops and declares "not implemented" rather than faking a derivation, and notes the route would itself depend on the open a_6 number.

Once we grant the theory. Still open — and here graininess is simply the wrong tool. This is the cleanest example of "the rule has nothing to say." The graininess rule is about *cost and finiteness*. The Born rule is about *probability weights* — a different mathematical animal entirely. You cannot manufacture a probability measure out of a finiteness-of-cost assumption; it would take a *separate* posit we have not made. So graininess neither dissolves nor solves this. It is **still open at both levels**, with a named (decoherence-based) path and no fabricated answer.

The exact catch. The graininess rule cannot be stretched to produce the probability rule — wrong shape. The proposed route is an honest halt that itself depends on the open a_6 number. No probability rule is invented.

Question 11 — The Yang-Mills mass gap (why the strong force has a floor)

This is a long-standing open problem of mathematical physics — one of the hardest recognized open problems in mathematics. Tread carefully.

What's asked. Show, rigorously, that the strong nuclear force confines quarks and has a "mass gap" — a strictly positive minimum energy — and crucially that this gap *survives* the idealized continuum limit. That last clause is the whole difficulty and the whole prize.

On its own. Still open. The corpus explicitly does **not** claim to have solved the recognized open continuum problem. It banks only a weaker thing: the geometry is *structurally compatible* with a strong force that has such a gap — which is a candidate, **not** a proof.

Once we grant the theory. The continuum half is dissolved; a finite leftover stays open. The hard core of the continuum mass-gap existence problem is a conjunction: build the theory in a perfect continuum *and* show the gap survives shrinking the grain to zero and the volume to infinity. The graininess rule **declines the shrink-to-zero half** — with a finite grain there's no continuum limit to control, so "does the gap survive shrinking the grain to nothing?" is **dissolved.** (Again: dissolved, not solved — no proof is produced; the demand is removed.) What graininess does **not** give you is the **finite leftover**: showing the gap is *uniform* — bounded below by one and the same positive floor across the whole finite family of grainy cells. That finite uniform-floor statement is exactly the open object. The geometry contributes only the *input specification* — which symmetry group, which particle content, no structural obstruction — and it explicitly does **not** supply the actual strong-force dynamics.

The exact catch. Dissolved is not solved: no continuum construction, no uniform floor, no positivity is proven; graininess only removes the shrink-to-zero demand. The finite uniform-floor leftover is **still open** and is never fabricated — no measure, bound, or estimate is invented. The owner of any real closure is external (constructive quantum field theory). And the theory's own internal version of this problem is, if anything, *harder* than the standard one, since it inherits all the ordinary difficulty plus a boundary/parity sector on top.

Question 12 — Why is the strong force so well-behaved under time-reversal? (strong CP)

The last of the four flatly-open questions, and the one the corpus calls the most tempting to fudge.

What's asked. There's a quantity (call it $\bar{\theta}$) that *could* make the strong force violate time-reversal symmetry. Experiment says it's astonishingly tiny — essentially zero. Either explain *why* it's so small, or eliminate it, with a named mechanism and a named way to be proven wrong.

On its own. Still open. The smallness of $\bar{\theta}$ is honestly **underived**. The paper offers "a candidate and a falsifier, not a solution." And the natural geometric route to explaining it has been **proven blocked** by an internal theorem.

Once we grant the theory. Still open — graininess is, again, the wrong tool. This is a finite, structural symmetry question with a *banked no-go result*, not a continuum idealization. So graininess neither dissolves nor solves it. It is **still open at both levels**, never fabricated.

The honest detail — a confirmed dead end, not progress. The geometric route was *refuted* by a representation-theory theorem (owner-ratified). In plain terms: the mechanism that would have explained $\bar{\theta}$ requires two groups of particles to sit on geometric structures with *opposite* "sign-character"; but the relevant building blocks all carry the *same* sign-character — so the explanation and the *three-families* fact turn out to be **mutually exclusive**. A committed computer scan looking for an escape found exactly zero. Any way out would require structure *outside* the allowed setting, which the discipline declines (fail-closed). The graininess rule has no bearing on this obstruction whatsoever.

The exact catch. A candidate plus a falsifier, **not** a solution; $\bar{\theta}$-smallness is honestly underived. The geometric route is **proven blocked** — this is a banked *negative result*, not progress toward closure. The neutron's electric dipole moment is the experiment that would adjudicate it. No $\bar{\theta}$ value or mechanism is invented.

The honest ledger, all twelve at a glance

The question	On its own	Once we grant the theory's assumptions
1. Is there a graviton?	Partly there (weak-gravity certificate, conditional)	Same — no upgrade; strong graviton waits on Q2's open number
2. Sense at highest energies?	Still open	Continuum half **dissolved**; finite a_6 number **still open**
3. Sensible probabilities?	Partly there	Gentle half conditional-certificate; strong half **still open**
4. The a_6 number itself	Still open	**Still open at both levels** — never invented
5. Black hole entropy / Page	Partly there (a consistency match only)	High-energy leg dissolved; count + evaporation **still open**
6. Extra dimensions stable?	Partly there	*Existence* of stable valley **derived** (decision grade); full certificate still waits on Q2; **no promotion**
7. No faster-than-light?	Partly there	Below-cutoff conditional certificate; high-energy half **still open**
8. Bookkeeping closes?	Partly there	Partly there — graininess doesn't touch it; boundary + twist **still open**
9. Right-handed matter?	Partly there	Classical count **derived (given the particles)**; quantum check **still open**
10. Where do probabilities come from?	Still open	**Still open** — graininess is the wrong tool
11. Yang-Mills mass gap	Still open	Continuum half **dissolved**; finite uniform-floor **still open**
12. Strong-force time symmetry	Still open	**Still open** — geometric route proven blocked

The four that stay flatly open — named out loud

When we grant ourselves *every* assumption the theory has, four things remain genuinely, irreducibly open. The graininess rule neither dissolves them nor manufactures them. We name them so no reader has to dig for them:

1. **The a_6 number** (Question 4) — a single finite quantity, definite and computable in principle, that simply has not been computed. Never invented.
2. **The finite mass-gap leftover** (Question 11) — what's left of the continuum mass-gap existence problem after the impossible shrink-to-zero half is declined: showing the strong force's energy floor is uniform across the finite grainy cells. Never fabricated.
3. **The Born rule** (Question 10) — where quantum probabilities come from. The graininess rule is the wrong shape to produce a probability; it would take a separate posit we haven't made.
4. **Strong-force time symmetry / $\bar{\theta}$** (Question 12) — why the strong force respects time-reversal so precisely. The natural geometric explanation has been *proven blocked*, which is a confirmed dead end, not a step forward.

And two more honest halves, for completeness: the **strongly-interacting** versions of Question 3 (sensible probabilities) and Question 7 (no faster-than-light) stay open even with every assumption granted — only their *gentle* halves earn conditional certificates. The boundary-and-twist part of Question 8 and the quantum part of Question 9 are likewise finite structural checks the graininess rule does not touch.

The bottom line

Across twelve hurdles, the honest tally is this. Two hurdles (the graviton, and probabilities-and-causality below the cutoff) earn real but *conditional* certificates in their gentle regime. One (right-handed matter) is *derived* at the classical level — given the observed particles as input — with a finite quantum check still open. One (the stability of the extra dimensions) reaches *decision-grade existence* by an independent route, with the full certificate still owed and **zero promotion** claimed. Several hurdles have their *impossible-continuum half dissolved* by the graininess rule while their *finite half stays open*. And four — the a_6 number, the finite mass-gap leftover, the Born rule, and strong-force time symmetry — stay **flatly open**, by the theory's own honest account.

The single discipline running through all of it: **dissolved is not solved.** Declining the demand for a perfect infinite continuum is honest, legitimate, and useful — but it is a *narrowing* of the problem, never an answer to what remains. Every finite leftover, every uncomputed number, every blocked route is kept in plain sight, with its name attached, never papered over and never invented.

> **Nothing was promoted. Nothing was fabricated. The overall grade stays exactly where it was. The honesty is the point.**