Measurement as Commitment

A Conditional Resolution of the Quantum Measurement Problem from Closure Dynamics

Keith Taylor — VERSF Theoretical Physics Programme

General Reader Summary

The idea this paper explores

The famous "measurement problem" of quantum theory may not be a problem about measurement at all. It may be the same question the rest of this programme already answers under a different name: what physical event turns a possibility into a fact? If measuring just is *that event, there is no separate mystery of collapse to solve — only the one process the framework was built on all along.*

Quantum theory describes two motions that do not sit easily together. Between observations, a system spreads into a haze of coexisting possibilities, evolving smoothly and predictably. At an observation, exactly one possibility is found, and the haze is gone. Nothing in the standard formalism says what physical event flips the first picture into the second. That gap is the measurement problem, and a century of interpretations have answered it in incompatible ways — some by declaring measurement a brute primitive, some by denying that the haze ever collapses at all, some by adding new physical laws to force the collapse.

This paper makes a different move. In the VERSF programme, reality is built from *commitments* — events in which one possibility becomes a fact that can never afterward be undone. Commitment is not introduced here for measurement; it was already needed to explain how physical time advances, how records form, and how different observers can agree about a shared past. The proposal is that measurement is simply an instance of that same commitment. There is no extra collapse law, because the thing collapse was invented to describe — a possibility becoming a fact — is the process the framework already runs on.

That reframing is the easy part, and the paper does not pretend it is the whole job. The real measurement problem has three teeth, and a serious proposal has to bite on all three. *Which* possibility becomes the fact — and in which set of alternatives, since a haze can be sliced many ways? *When* does commitment happen, if no observer triggers it? And does the commitment process actually reproduce the precise odds quantum theory predicts, or merely leave room for them? The paper takes each in turn, says exactly what the framework can offer toward it, and marks plainly the part that remains unsolved. On the first, it offers a real candidate the standard theory lacks: the set of alternatives is fixed by what the closure structure makes admissible, not chosen by hand. On the second, it notes that since time itself is built from commitment, "when" cannot mean a reading on a pre-existing clock — which sharpens the question rather than dissolving it. On the third, it is careful to separate the odds (imported, unchanged) from the selection event (the new content), and to flag that whether the one governs the other is an open consistency step, not a settled result.

The honest bottom line is modest and exact. If the framework's prior results hold — commitment, the probability rule, the closure structure — then measurement needs no separate postulate; it is commitment seen in the laboratory. What survives unchanged is every experimental prediction quantum theory already makes. What is newly offered is a physical reading of where the definite outcome comes from. And what is newly admitted, rather than hidden, is that three specific hinges remain open, each named and located, with the deepest of them — what fixes the set of possible outcomes — given a candidate answer the framework alone can supply. On the question of the odds, the paper narrows things twice over. First, several convergent lines of the programme agree that the precise weighting can only be the one quantum theory already uses — so the question is no longer *which* odds. Second, the question of whether the commitment process actually *delivers* those odds is reduced, in turn, to a single structural assumption: that commitment draws on nothing beyond what is already physically present in the situation, with no hidden extra ingredient steering the outcome. If that one assumption holds, the odds follow; if it fails, there is a hidden ingredient to be found, and the paper says exactly where it would have to enter. So what began as an open question of unbounded shape is now the truth of one isolated, stated assumption about the nature of commitment — still unproven, but no longer open-ended.

Abstract

The quantum measurement problem arises from the coexistence of two facts: quantum states evolve continuously under deterministic laws, and measurements yield definite outcomes. The standard formalism predicts outcome probabilities accurately but specifies no physical mechanism by which possibilities become facts.

This paper proposes a conditional resolution within the VERSF framework. The central claim is that measurement is not a primitive process and requires no independent collapse postulate; it is identified with **irreversible commitment** — the physical transition by which one admissible possibility becomes a committed fact. The measurement problem is thereby transformed from *"why does the wavefunction collapse?"* into *"what physical process creates a fact?"* — a question the framework must already answer to construct physical time, record formation, and observer comparability. Measurement is then not an additional mystery but an instance of the same commitment process.

The proposal does not stop at that reframing, because the measurement problem's genuine difficulty lies in three places the reframing alone does not reach. We confront each explicitly: the **preferred-basis problem** (which decomposition of the state supplies the admissible

possibilities), the **commitment criterion** (what physical condition constitutes a commitment event, given that time is emergent and cannot supply a trigger), and **Born reproduction** (whether the commitment dynamics select outcomes with the frequencies the probability rule assigns, rather than merely carrying those probabilities as labels). For the first we offer a candidate the standard theory lacks — the admissible set is fixed by closure admissibility, not chosen — while marking its identification with the empirically observed pointer basis as open. The commitment criterion is isolated as the principal open dynamical input. On Born reproduction the paper marks a constraint and then a reduction: several convergent strands of the programme, sharing its operational geometry, identify the quadratic measure $P(i) = |c_i|^2$ as the unique admissible *weighting* (§5.4), and a minimality argument then reduces reproduction itself — whether commitment selects with that measure's frequencies — to a single conjectural principle, that commitment uses no information beyond the admissibility structure, set against two standard structural assumptions and an inherited uniqueness theorem (§5.5). The open problem narrows from *which* probability rule commitment should follow, through reproduction *of a fixed measure*, to the truth of one isolated, structurally falsifiable principle about commitment — conditional, throughout, on the surrounding reconstruction results.

We then position the proposal among existing interpretations: it is *objective* in ontology (facts are observer-independent and definite) yet, unlike dynamical-collapse models, introduces no modification of unitary evolution — commitment is the substrate event of which unitary evolution is the pre-commitment description, not a stochastic perturbation of it. Consequently the proposal is observationally equivalent to standard quantum mechanics except where commitment leaves a structural trace, and the natural location of any such trace is the conditional Gate-3 residue of the companion programme.

The resulting picture leaves quantum probabilities unchanged and supplies a candidate physical interpretation for outcome selection and record formation. It is conditional on the previously developed closure and probability structures and should be read as an interpretive consequence of those results, not a standalone derivation.

Epistemic markers: ***(inherited)*** *for results imported from prior VERSF papers;* ***(conditional)*** *for results holding under stated inputs;* ***(conjectural)*** *for interpretive identifications;* ***(open)*** *for what remains undecided.*

Table of Contents

---

1. Introduction

Few questions in modern physics have generated more debate than the quantum measurement problem, and the reason is that quantum mechanics appears to describe two incompatible kinds of evolution while predicting experiment with extraordinary accuracy. Between measurements, a system evolves continuously and deterministically through a superposition of possible outcomes,

$$|\psi\rangle = \Sigma_i \, c_i \, |i\rangle,$$

each branch retained, none privileged. At a measurement, a single outcome $|i\rangle$ is found, and the others are gone. The standard formalism does not identify the physical event that separates these two regimes; it supplies a rule for the *statistics* of the outcome, the Born rule $P(i) = |c_i|^2$, but no account of the transition that produces an outcome at all.

The major interpretations answer the gap differently, and it is worth stating their answers plainly, because the proposal of this paper is best understood by where it agrees and disagrees with each. Copenhagen treats measurement as a primitive, a step outside the dynamics invoked when an apparatus is involved. Many-Worlds denies that any collapse occurs, retaining every branch and relocating definiteness into the observer's perspective. Objective-collapse theories of the Ghirardi–Rimini–Weber type add a stochastic term to the evolution law so that superpositions of large systems decay dynamically into definite states. Relational and information-theoretic approaches reinterpret the quantum state as a catalogue of relations or of an agent's information, so that "outcome" is always outcome-relative-to-something.

Despite their differences these share a single underlying question:

What physically distinguishes a possibility from a fact?

The present paper argues that the VERSF framework already contains the materials to answer this question, because the possibility/fact distinction is not introduced for measurement — it is the distinction on which the rest of the framework is built. The argument proceeds in three parts. First, the reframing: measurement is the creation of a fact, and fact-creation is commitment (§§2–4). Second, and this is where the paper does its real work, the confrontation: the measurement problem's genuine difficulty survives the reframing in three specific forms, each of which the framework must address rather than dissolve (§5). Third, the positioning and honest accounting: where the proposal sits among existing interpretations, what it leaves unchanged, and what it leaves open (§§6–13).

The objective is not to prove that commitment occurs — that is inherited from prior work — but to determine whether the measurement problem is, under the framework's existing commitments, an additional mystery or an instance of a process already required. We argue it is the latter, conditionally, and we mark exactly which conditions the conclusion rests on.

---

2. The Central Reframing: Possibility and Fact

Within VERSF the fundamental ontological distinction is not between the wavefunction and the measurement that acts on it, but between **admissible possibilities** and **committed facts**. A possibility is a state compatible with the current closure structure; a fact is a record that has become irreversibly committed. The wavefunction, in this reading, is a description of the admissible possibilities — what remains open — and measurement is not an operation performed on that description but the physical event in which the description's openness is resolved into a fact.

This distinction is not imported for the occasion. It already does load-bearing work elsewhere in the framework: physical time is constituted through committed facts, records are committed facts, and the comparability of observations across observers depends on a shared body of committed fact. The framework therefore already contains a notion of "becoming definite," and it is the same notion at issue in measurement. The measurement problem, on this view, is a special case of a question the framework must answer regardless:

How does commitment occur?

Stating the reframing precisely: where standard quantum theory poses a transition between two *descriptions* (unitary evolution and collapse) and asks what triggers the switch, VERSF poses a transition between two *ontological statuses* (admissible and committed) and identifies measurement with the second. The shift is from a problem about the formalism's two modes to a claim about reality's two statuses. Whether that shift is progress or relabelling depends entirely on whether the framework can say something the standard treatment cannot about the three hard components of the transition — which is the burden of §5, and the reason this paper does not end at the reframing.

---

3. Why Facts Are Required, Not Optional

A world containing only possibilities has no distinguished history. Multiple alternatives coexist, no particular outcome is privileged, and no record exists. Such a world cannot support the structures the framework treats as physically basic: history is undefined, because there is no fact of the matter about what occurred; record formation is impossible, because a record is a committed fact and there are none; and physical time cannot advance, because in this framework the advance of time *is* the accumulation of committed fact.

The consequence is that the existence of facts is not an optional addition to a world of possibilities — it is a precondition for there being a physical world of the kind we inhabit at all. This matters for the measurement problem because it removes a tempting evasion. One might hope to take the possibility-only description as complete and treat definite outcomes as an illusion or a perspectival artefact, as Many-Worlds does. The framework forecloses that route not by argument against Many-Worlds but by its own architecture: time, records, and observer comparability are built from committed fact, so a framework that dispensed with commitment would have to rebuild all three on another foundation. Within VERSF, commitment is already spent; the only question is how it occurs, not whether.

What this section establishes is therefore a necessity, not a mechanism. Facts must arise **(inherited** from the framework's construction of time and records**)**. The mechanism by which a particular fact arises from particular possibilities is the content of the sections that follow, and is where the genuine difficulty lives.

---

4. Measurement as Commitment

The proposal of this paper is stated in one line:

Measurement = Commitment.

A measurement occurs when an admissible possibility becomes irreversibly committed. Before commitment, multiple outcomes remain admissible, probability applies to them, and alternative futures remain open; after commitment, one outcome is selected, a record exists, the alternatives are excluded, and physical time has advanced. Schematically the event is the transition

$$|\psi\rangle = \Sigma_i c_i |i\rangle \longmapsto |i\star\rangle \text{ (one committed fact, } i\star\text{)},$$

but the arrow here is not a new dynamical law added to the Schrödinger equation. It is the substrate event of which the left-hand superposition is the *pre-commitment description* — the catalogue of what was admissible — and the right-hand fact is what commitment produces. The measurement event is not an additional rule imposed upon quantum mechanics; it is the physical creation of a fact, and quantum mechanics' unitary evolution is the framework's description of the system while no commitment has yet occurred.

This is the whole of the positive proposal, and by itself it is almost trivial — which is the point. The reframing converts "why does collapse happen?" into "how does commitment happen?" at no cost, precisely because it adds nothing: it identifies an existing mystery with an existing primitive. The substance of the paper is not this identification but the test of whether the identification *helps* — whether, having relocated the measurement problem into the theory of commitment, the framework can say more about the relocated problem than standard quantum theory could say about the original. If it cannot, the identification is a relabelling. If it can, the relabelling was the right move. The three hard problems of §5 are exactly the places this is decided.

---

5. The Three Hard Problems the Proposal Must Face

The reframing of §4 leaves the measurement problem's genuine content untouched in three places. Each is a question that "measurement = commitment" does not by itself answer, and a serious proposal must say, for each, what the framework offers and what remains open. We take them in turn — the preferred basis, where the framework offers a structural candidate; the commitment criterion, the least supplied and most open of the three; and Born reproduction, where the weighting is now constrained even though reproduction itself is not — and the amount the framework can currently supply differs sharply between them, which is itself part of the honest picture.

5.1 The preferred-basis problem

A superposition has no intrinsic decomposition. The same state

$$|\psi\rangle = \Sigma_i c_i |i\rangle$$

can be rewritten in any basis, and "the admissible possibilities" are well-defined only once a basis is chosen. Saying that commitment selects one admissible possibility therefore presupposes a privileged set of possibilities — a preferred basis. This is the deepest tooth of the measurement problem: standard collapse postulates put the basis in by hand (the eigenbasis of the measured observable), and any account claiming to *explain* outcome production rather than stipulate it owes an account of why commitment occurs in one basis and not another.

The framework has a genuine candidate here, and it is the principal thing it can offer that the standard treatment cannot. The admissible possibilities are not an arbitrary basis choice; they are fixed by the **closure structure** — by which decompositions are admissible under the framework's closure dynamics. Commitment occurs in the closure-admissible decomposition, and that decomposition is determined by physics (the structure of closure) rather than selected by the experimenter's choice of observable. In slogan form: the preferred basis is the closure-admissible basis, and it is preferred because closure prefers it, not because measurement stipulates it. This is a candidate dissolution of the preferred-basis problem rather than a restatement of it, because it names a physical structure that does the selecting.

The honesty required here is exact, and it is the same shape as the identification gaps marked throughout the companion papers. That the closure-admissible decomposition *exists and is unique* is a structural claim about closure — **(inherited** if the closure programme establishes it, **conjectural** if not yet; the present paper takes it as a claim about closure to be discharged there, not proved here, and the candidate dissolution rests on it**)**. That this decomposition *coincides with the empirically observed pointer basis* — the position-like, decoherence-selected basis in which laboratory outcomes are actually definite — is a further and separate identification, and it is **(open)**. The two must not be run together: the first is whether closure singles out a basis at all, the second is whether the basis it singles out is the observed one. The framework offers a principled origin for a preferred basis; whether the basis it singles out is the one experiment singles out is not established here and must be shown. We therefore claim **(conjectural)** that closure admissibility supplies the preferred basis, and mark **(open)** the identification of that basis with the observed pointer basis. The candidate is real — it is more than standard quantum theory offers, which is no origin at all — but it is a candidate, and the identification step is precisely where future work on the closure structure would have to deliver.

5.2 The commitment criterion: when, without a clock

If commitment is not triggered by an observer — and §8 will argue the observer is secondary — then some physical condition must determine when an admissible possibility becomes a committed fact. Standard collapse models answer with a rate (GRW's stochastic hitting process); Copenhagen answers with the presence of an apparatus; both supply a *when*. The present proposal owes one too, and here the framework's own architecture makes the question sharper rather than easier.

The sharpening is this. In VERSF, physical time is constituted by committed fact (§3); the advance of time *is* the accumulation of commitment. So the commitment criterion cannot be a condition of the form "commitment occurs at time t," because there is no pre-existing time t against which to time it — time is downstream of commitment, not upstream. The question "when does commitment happen?" cannot mean "at what clock reading," because the clock is built from the answer. What the criterion must instead specify is an intrinsic, atemporal condition on the closure structure: a condition such that, when it holds of an admissible configuration, that configuration constitutes a commitment event — and the holding of the condition is what the advance of time consists in, not something that happens within an already-running time.

The natural form for such a criterion, given the rest of the framework, is a **closure-saturation condition**: commitment occurs when the closure structure reaches a configuration in which an admissible possibility can no longer be sustained as merely admissible — when continuing to hold the alternatives open would violate closure. On this reading commitment is forced, not stochastic: it is what closure does when admissibility can no longer be maintained, the analogue of a constraint that must be discharged. This is the framework's candidate answer, and it is **(open)**: the paper does not supply the saturation condition explicitly, and identifying it — stating precisely which closure configurations force commitment — is one of the two load-bearing open inputs on the commitment dynamics, the other being Born reproduction (§5.3). What the paper does establish is the *form* the answer must take (an intrinsic closure condition, not a clock-indexed rate) and *why* it must take that form (emergent time forbids the clock-indexed alternative). That is a constraint on the solution, not the solution; but it is a non-trivial constraint, and it rules out the entire family of rate-based criteria that the standard objective-collapse models use.

5.3 Born reproduction: selection versus weighting

The proposal imports the Born rule and does not modify it: $P(i) = |c_i|^2$ **(inherited** from the ODG/OIP programme**)**. But importing the rule is not the same as showing the commitment dynamics obey it, and the distinction is the same one drawn sharply in the companion phase paper between an inherited structure and the demonstration that the new mechanism reproduces it.

Two things must be kept apart. **Weighting** is the assignment of probabilities to admissible possibilities — what the Born rule does, and what is imported unchanged. **Selection** is the commitment event itself — which admissible possibility actually becomes the fact $i^\star$. The proposal's positive content is that commitment *is* the selection event. The open question is whether selection, as the framework realises it, occurs with the frequencies weighting assigns: whether, over an ensemble of commitment events on identically prepared systems, the relative frequency of outcome i converges to $|c_i|^2$,

$\mathrm{freq}(i^\star = i) \longrightarrow |c_i|^2$ (the reproduction condition),

or whether the commitment dynamics select with some other statistics, leaving the Born weights as labels the selection does not actually follow. If the former, commitment reproduces Born statistics and the account is consistent; if the latter, the proposal contradicts experiment and fails.

This is the analogue, for measurement, of the reproduction gap the phase paper isolates for dynamical phase: it is not enough that the right probability structure is available; the new mechanism must be shown to *generate* it, not merely to coexist with it. We mark the reproduction condition **(open)** and state plainly what would close it: a demonstration that the closure-saturation selection of §5.2, applied across an ensemble, yields frequencies governed by the imported Born weights. The proposal does not establish this. It is consistent with it — nothing in the commitment picture obviously contradicts Born statistics — but consistency is not reproduction, and the honest status is that Born reproduction is a condition the proposal requires and does not yet prove. Notably, §5.2 and §5.3 are linked: the reproduction condition is a statement about how the §5.2 commitment criterion behaves statistically, so the two open inputs are facets of one underlying object — the commitment dynamics — examined once for *when* it fires and once for *with what frequencies*. The section that follows constrains the first half of §5.3 — the weighting — while leaving the reproduction half, the contested one, where this section leaves it.

5.4 The weighting is now constrained: convergence on the Born measure

Four convergent strands of the VERSF programme identify the same probability weighting, and the result bears directly on §5.3 — though on its first half only, and the distinction must be held throughout. A word first on how these relate to the inherited Born rule. §5.3 imports $P(i) = |c_i|^2$ from the ODG/OIP programme as an established input; the four strands below are not that single import but the wider body of reconstruction results — admissibility, pairwise selection, physical-necessity elimination, and record consistency — that by different routes arrive at the same measure and so establish its *uniqueness* rather than merely its availability. ODG/OIP is the inheritance §5.3 leaned on; the four routes are what show that inheritance was forced, not chosen. Recall that §5.3 separated *weighting* (which probabilities attach to the admissible possibilities) from *selection* (which possibility the commitment event produces) and isolated *reproduction* (whether selection frequencies match the weighting) as the open condition. The results here constrain the weighting; they do not, by themselves, touch reproduction, and the section is written to keep the two apart.

The convergent results are four, and they share a common dependence that must be stated first, because four routes are only as independent as their shared premises are few. Each takes the operational Hilbert geometry of the framework as given; what differs is the additional structure each adds to it. The convergence is therefore evidence about what that shared geometry forces, not four wholly independent discoveries — a qualification that matters when weighing the cumulative force below. **(conditional** on the shared operational geometry, and on each route's own further inputs as noted**)**.

The first and strongest is the **admissibility reconstruction programme**. The Admissible Measure and Admissibility Closure papers show that once operational Hilbert geometry, reversible transport structure, finite distinguishability, and compositional consistency are accepted, the probability layer is no longer free: the only admissible measure compatible with the inherited geometry is the quadratic one, $P(i) = |c_i|^2$, with every alternative violating admissible transport invariance, compositional consistency, or cross-channel universality. The significance for the present paper is that commitment is not selecting among an unrestricted family of candidate measures — the admissibility structure already isolates a unique weighting. **(conditional** on the four stated inputs**)**.

The second route, from the **Double Square programme**, is the one that may bear on more than weighting, and it is flagged here precisely because its scope decides how much this section can claim. There probability is derived from pairwise correlation structure rather than individual-path selection: because distinguishability is relational, irreversible selection acts on path-correlation structures rather than on isolated paths, and the resulting bilinear structure uniquely yields $P(A) = |\psi_A|^2$. The scope question is decisive. *If* this programme derives the measure from the structure of the selection act itself — from what fact-producing selection does to relational distinguishability — then it bears on the *reproduction* half of §5.3, not merely the weighting half, because it is then an argument about selection dynamics rather than about which measure is abstractly admissible. *If* instead it derives the measure abstractly and only afterward observes that selection respects it, it constrains the weighting alone, like the other three. The paper marks this distinction **(open)**: the double-square route is the natural candidate path toward reproduction, but whether it reaches reproduction or only weighting turns on which of these two it does — and that should be settled before the section leans on it for anything beyond convergence. As presently stated it is counted toward the weighting convergence only; its possible reach into reproduction is noted, not used.

The third route, the **Physical Necessity analysis**, is an elimination argument: it systematically removes the alternatives — finite holonomy structures, individual-path selection rules, higher-order selection kernels, non-factorising probability assignments, non-quadratic exponents — and finds the standard quantum probability architecture as the unique survivor. Its force is methodological: the quadratic measure is not one successful candidate among many but the sole survivor of the admissibility filter. **(conditional** on the filter being exhaustive — that the enumerated alternatives are in fact all of them**)**.

The fourth route, the **consistency programme**, interprets probability as a measure of compatibility between candidate extensions of the committed record and the existing record structure, and shows that under non-negativity, non-contextuality, unitary invariance, and compositional multiplicativity the unique consistency measure is again $P(i) = |c_i|^2$. This route is especially apt here because commitment is itself a record-producing event, so the consistency interpretation links the weighting directly to the geometry of fact creation. **(conditional** on the four stated axioms**)**.

What the convergence establishes, and what it does not, must be stated at its exact size — and stated in §5.3's own vocabulary, or the cumulative weight of four routes will read as having closed more than it has. It establishes that the *weighting* is not arbitrary: operational geometry, admissibility, pairwise selection, record consistency, and physical-admissibility elimination all identify the same measure $P(i) = |c_i|^2$, so commitment, whatever its dynamics, is not choosing among candidate probability laws — the law is forced to be quadratic. That is a genuine narrowing of §5.3's first half. It does *not* establish *reproduction*: that the commitment dynamics of §5.2 *select* with quadratic frequencies, rather than merely that the quadratic measure is the one they would have to inherit. Four derivations of the same weighting are evidence about the weighting's uniqueness, not about the selection's frequencies; the convergence narrows the *target* commitment must hit without showing that it hits it. This is the same caution the criterion paper applies to its own convergence evidence — co-occurrence of several routes to one structure constrains that structure but does not, by itself, identify the dynamics with it — and it applies here with equal force.

The remaining problem is therefore narrower than §5.3 first posed it, and that narrowing is the contribution of this section. The open question is no longer *"which probability rule should commitment follow?"* — the admissibility programme has constrained that to the unique quadratic measure. It is the strictly smaller *"do the commitment dynamics inherit that already-identified measure?"* — a consistency problem linking the §5.2 commitment criterion to a now-fixed target, rather than a search over possible laws. Reproduction remains open; what has changed is that it is reproduction *of a known measure*, not discovery of an unknown one. (Should the double-square scope question above resolve in favour of selection, even this residual gap would begin to close from that direction — but that is contingent on the flagged identification, and is not claimed here.)

The compression can be stated as a proposition, since it is the section's load-bearing output.

Proposition 5.4 (Reduction of the Born-reproduction problem) (conditional on the reconstruction inputs of §5.4**)**

§5.3 already distinguished two questions:

- the **weighting problem** — which probability measure is compatible with the framework's admissibility structure;
- the **reproduction problem** — whether the commitment dynamics select outcomes with that measure's frequencies.

The results surveyed in §5.4 — admissibility reconstruction, record consistency, pairwise selection, and physical-necessity elimination — converge on the same answer to the first, $P(i) = |c_i|^2$, *conditional on their shared operational geometry and each route's stated inputs*. Granting those inputs, the weighting problem is no longer open as a search among candidate laws: the admissible measure is uniquely fixed. The reproduction problem is untouched by this and remains open. Consequently the two problems, distinct since §5.3, now stand at different statuses — weighting settled-given-inputs, reproduction open — and the residual measurement burden is the strictly smaller question:

Does the commitment mechanism inherit the unique admissible weighting already identified by the surrounding quantum structure?

Scope, stated exactly. The proposition claims a *conditional* settling of the weighting, not an unconditional one: every route of §5.4 carries its own antecedents, and the convergence is evidence about what the shared geometry forces, not four independent proofs (§5.4, opening). It claims nothing about reproduction beyond what §5.3 and the paragraph above state. And it counts the double-square route toward the weighting convergence only — its possible reach into reproduction is the flagged **(open)** identification of §5.4, not an input here. With those bounds, the proposition's content is exact: the measurement programme is reduced from a search over candidate probability laws to a consistency problem linking commitment dynamics to an already-constrained measure — a substantial compression of the original problem, and a conditional one.

5.5 Born reproduction from commitment consistency

Section 5.3 isolated the reproduction question — *does the commitment process select with the Born frequencies, or merely coexist with them?* — and §5.4 narrowed it by fixing the *weighting* uniquely, conditional on the reconstruction inputs. This section attacks the reproduction half directly and reduces it, in turn, to a single structural principle about commitment. The reduction is built on the **consistency** route specifically; the pairwise-selection (double-square) route's possible reach into reproduction remains the flagged open question of §5.4 and is not recruited here, so that the firewall §5.4 maintained between the two routes is preserved.

Commitment as selection on admissible extensions. Let $\mathcal{C} = \{C_1, C_2, \ldots, C_n\}$ be the admissible commitment extensions available to a system in a given pre-commitment configuration — the ways the current record can be extended by one committed fact while preserving admissibility. A commitment event selects exactly one member of $\mathcal{C}$. The consistency programme assigns each extension C_i a weighting W_i measuring the degree to which C_i preserves the admissible structure of the existing record, and its consistency-weighting theorem establishes $W_i \propto |c_i|^2$ **(inherited)**. What §5.5 must address is not *that* W_i is the admissible weighting — §5.4 settled that — but *why commitment should select according to it*, which is the reproduction question in its sharpest form.

The principle the reduction turns on.

Commitment Minimality Principle (conjectural)

A commitment event may depend only on information contained within the admissibility structure of the pre-commitment state. No hidden selector, external weighting, or independent fact-producing variable participates in outcome selection.

This is the measurement-theoretic analogue of the single-source theorem: if all physically meaningful structure is already carried by the admissibility geometry, commitment cannot legitimately access information not present within that geometry. One scope point should be stated at the outset, to forestall over-reading: the principle constrains *which information* commitment may use, not the contextual independence of the weights it assigns — it does not by itself deliver non-contextuality, which the proof below imports as a separate premise for exactly this reason. We mark it **(conjectural)** — it is introduced here and argued by analogy to single-source, not proved — because it is the load-bearing new assumption, and the honesty of the section depends on its status being visible rather than buried in a proof.

Motivating Minimality: the Completeness Principle

The Minimality Principle can be motivated by — and is arguably the operational consequence of — a stronger structural claim, though one whose argument must be read carefully, since it is *transcendental* rather than a derivation from independent premises, and presenting it as the latter would smuggle in a circularity.

Commitment Completeness Principle (conjectural — resting on the completeness of admissibility as the framework's inventory of physically relevant distinguishability**)**

Every variable capable of influencing fact production is already represented within the admissibility structure of the pre-commitment state. Equivalently: there exists no admissibility-external fact selector.

The motivation runs as follows. Commitment is not a process occurring within an already-completed reality; it is the mechanism by which new facts are produced. Any variable capable of influencing which fact becomes real therefore participates directly in fact production. Suppose an admissibility-external selector X existed, so that outcome selection took the form

Fact = F(Admissibility, X).

If changing X changes which fact becomes real, then X carries physically relevant distinguishability — it contributes to the structure of reality. And here is the step that must be labelled for what it is: in this framework, *to carry physically relevant distinguishability is what it means to belong to the admissibility structure*, because admissibility is defined as the framework's complete inventory of the physically relevant alternatives available to the system. So X cannot consistently remain external to admissibility — the moment it influences fact production it is physically significant, and the admissibility structure is precisely the encoding of physical significance. X was assumed external, yet its efficacy forces it inside: a contradiction.

The honesty this requires is explicit, because the contradiction is *not* derived from premises independent of what "admissibility" means. It is transcendental: it turns on admissibility being, by the framework's construction, the complete register of physically relevant distinguishability, so that "efficacious but external" is close to a contradiction in terms. Read that way the argument is sound, and it is the same *form* as the single-source theorem, which excludes a second independent source not by experiment but by what the framework takes a source to be. Read instead as a derivation from an *independently* characterised admissibility structure, the middle step — efficacy $\Longrightarrow$ membership in *that specific* structure — would itself require proof and is not supplied here. We intend the first reading and mark the principle accordingly: it rests on the completeness of admissibility as defined, and its content is the claim that that definition is the right one — that reality contains no second, hidden register of fact-relevant structure.

Corollary 5.5a (Minimality from Completeness) (conditional on the Completeness Principle**)**

If the Commitment Completeness Principle holds, the Commitment Minimality Principle follows. Since all fact-producing structure is already contained within admissibility, commitment has no admissibility-external information available to access; hence it may depend only on admissibility structure.

Proof. Immediate. Completeness asserts that no fact-relevant variable lies outside admissibility; Minimality asserts that commitment uses no information outside admissibility. If there is no such information to be had (Completeness), commitment trivially uses none (Minimality). ■

Two qualifications keep this from over-reaching. First, the reduction *relocates* the conjectural burden rather than removing it: Minimality is not made cheaper by being proved, but shown to be the same conjecture as Completeness seen at the operational level — Completeness the claim that admissibility is the complete register of fact-relevant structure, Minimality its consequence for what commitment can use. What one buys, granting either, is a single structural commitment — that there is no second source of physically effective structure outside admissibility, exactly as single-source asserts there is none outside the committed record. Second, the proof of Proposition 5.5 uses Minimality *directly* as premise 4; Completeness is offered as its possible grounding, not substituted for it in the proof, so the premise structure of §5.5 is unchanged. The status is therefore unchanged in kind — Completeness is **(conjectural)**, as Minimality was — but the *content* is clarified: the open question is not a free-floating restriction on commitment but a definite claim about the completeness of admissibility, of a piece with the programme's existing single-source commitment.

Proposition 5.5 (Commitment inherits the consistency weighting) (conditional on the Commitment Minimality Principle, the composition and non-contextuality premises, and the consistency programme's uniqueness theorem**)**

Assume:

1. **(selection)** commitment selects exactly one admissible extension;
2. **(admissibility preservation)** the selected extension preserves admissibility, so that R_i is a weighting on the admissibility structure at all — the premise that makes "admissibility data" the domain of R_i, rather than supplying one of the four axioms;
3. **(relabelling invariance) (structural assumption)** selection is invariant under relabelling of admissible alternatives — admissibility-equivalent extensions receive equal weight;
4. **(Minimality)** commitment obeys the Commitment Minimality Principle;
5. **(composition)** commitment weightings on independent subsystems multiply;
6. **(non-contextuality) (structural assumption)** an extension's weight does not depend on which other admissible alternatives it is committed alongside;

together with the inherited

7. **(consistency uniqueness, inherited)** the consistency weighting W_i is the *unique* weighting on admissibility data satisfying non-negativity, non-contextuality, unitary invariance, and multiplicativity (consistency programme).

Then the commitment-selection weighting R_i satisfies $R_i \propto W_i$, hence — after normalisation — $R_i = |c_i|^2$.

Proof. The premises divide into those that confine R_i to a domain and those that supply the axioms the uniqueness theorem requires, and the proof keeps the two roles distinct. Premise 2 (admissibility preservation) is what makes R_i a weighting *on the admissibility structure* in the first place — it fixes the domain on which R_i is defined, and supplies none of the four axioms directly; without it "admissibility data" would not be the thing R_i ranges over. Premise 3 (relabelling invariance) then forces R_i to treat admissibility-equivalent extensions identically: R_i depends on the admissibility *status* of C_i and is constant across relabellings preserving that status. Premise 4 (Minimality) forces R_i to depend on *nothing beyond* that status: no variable outside the admissibility structure enters. Together, premises 2–4 confine R_i to the **admissibility-covariant weightings** — functions of the admissibility data alone, equivariant under relabelling.

This is as far as premises 2–4 reach, and it is *not yet* W_i. The admissibility-covariant weightings are a family, not a single function: Minimality and relabelling-invariance exclude any weighting using information beyond admissibility, but they do not, by themselves, single out the consistency functional from other admissibility-covariant assignments (its square, say, or any other functional of the same data). The identification with W_i specifically requires the consistency programme's uniqueness result, premise 7 — and to invoke it, R_i must satisfy the four consistency axioms. Their sourcing must be stated exactly, because not all four fall out of the commitment premises, and pretending they do would repeat the error this proof is built to avoid:

- **non-negativity** — genuinely derived: R_i is a selection frequency and is ≥ 0;
- **unitary invariance** — genuinely derived: admissibility is unitarily defined, so an admissibility-covariant weighting (premises 2–4) inherits it;
- **multiplicativity** — *assumed*, as premise 5; it is the composition property of weightings on independent subsystems, not a consequence of Minimality, and is flagged as a premise rather than claimed for free;
- **non-contextuality** — *assumed*, as premise 6, and this is the axiom whose sourcing must be most careful. It does **not** follow from Minimality together with relabelling-invariance, and earlier framings that suggested it did were mistaken. Non-contextuality is the claim that an extension's weight is independent of the co-admissible set it is committed alongside. But the co-admissible context is plausibly *internal* to the admissibility structure — it is itself admissibility data — so Minimality, which excludes only information *beyond* admissibility, is simply silent on contextual dependence and cannot exclude it; and relabelling-invariance is a permutation symmetry, a different proposition entirely. Non-contextuality is therefore imported as its own structural premise (premise 6), matching the consistency programme's own use of it as an axiom rather than being manufactured from the commitment side.

With non-negativity and unitary invariance derived, and multiplicativity and non-contextuality assumed as premises 5 and 6, R_i is a weighting on admissibility data satisfying all four axioms — and by premise 7 there is exactly one such weighting, W_i. Hence $R_i \propto W_i$. Since $W_i \propto |c_i|^2$ (inherited), $R_i \propto |c_i|^2$, and after normalisation $R_i = |c_i|^2$. ■

Remark (what kind of result this is). Proposition 5.5 is **not** an independent derivation of the Born rule, and reading it as one would make it circular against the consistency programme it draws on. The Born weighting is inherited, fixed elsewhere; the proposition is a **transfer result**. It establishes that *if* commitment is restricted to admissibility-covariant information (Minimality, premise 4) and the commitment weighting satisfies the stated structural assumptions (composition and non-contextuality, premises 5–6), *then* commitment selection inherits the unique admissible weighting already fixed by the consistency programme (premise 7) — carrying that weighting from the abstract probability layer onto the dynamics of fact-production. This is the same kind of move as the criterion paper's transport-natural transfer theorem: not a fresh derivation of the transferred quantity, but a demonstration that a structure established in one place is inherited by another under stated conditions. The contribution is the transfer, not the measure; the measure is upstream.

What the proof does and does not establish. It does **not** derive Born reproduction from the Minimality Principle alone — that was the temptation, and premises 2–4 reach only "admissibility-covariant," not "consistency-weighting." The step to the consistency functional specifically is carried by the inherited uniqueness theorem (premise 7), and of the four axioms that theorem requires, only two (non-negativity, unitary invariance) are derived from the commitment premises; the other two (multiplicativity, non-contextuality) are themselves assumed, as premises 5 and 6. So the result is genuinely conditional, on a set of inputs of distinct kinds: the Commitment Minimality Principle (premise 4, **conjectural**, the load-bearing novelty); two structural assumptions on the commitment weighting (composition, premise 5, standard for independent subsystems; and non-contextuality, premise 6, which — as the proof stresses — does *not* fall out of Minimality and must be assumed); and the consistency programme's uniqueness theorem (premise 7, **inherited**). The proof derives reproduction not from nothing, and not from Minimality alone, but from Minimality plus two structural assumptions plus the inherited consistency structure — which is the reduction the section claims, and no more. The honest headline is therefore not "one principle" simpliciter but "one conjectural principle, two standard structural assumptions, and one inherited theorem"; where the summaries compress this to the Minimality Principle, that compression is exact only about the *novel* and *conjectural* part, with the rest understood as the assumed and inherited scaffolding the proof names explicitly.

Consequence

Under the Commitment Minimality Principle, the Born-reproduction problem reduces. The commitment mechanism requires no independent probability law and no separate stochastic postulate: the admissibility structure determines a unique weighting (conditional on the reconstruction inputs of §5.4 and the consistency uniqueness theorem), and commitment inherits it because — by Minimality — admissibility is the only structure available to it. The resulting picture is

admissibility → consistency weighting → $|c_i|^2$ → commitment selection.

The remaining question is therefore no longer *"why does commitment obey the Born rule?"* but the single structural one:

Is the Commitment Minimality Principle true?

If it holds, Born reproduction follows from the existing admissibility structure (given the inherited uniqueness result); if it fails, the additional selector — the information commitment uses beyond admissibility — must be identified explicitly, and the proof above localises exactly where it would enter (it would be a violation of premise 4, an admissibility-external variable in R_i). This is the sense in which the principle is *falsifiable*, and it should be read precisely: it is refutable in principle by exhibiting such a hidden selector — a structural disproof internal to the framework — not by an experiment, since §12 establishes that the base proposal has no measurement-distinguishing test. "Falsifiable" here means the principle makes a definite structural claim whose negation has definite content (some admissibility-external variable participates in selection), not that a laboratory could decide it. Reproduction is thereby reduced from a search over probability laws to a single structural question about the nature of commitment, in the same way §5.4 reduced weighting from a search to an inheritance. The two reductions compose: §5.4 fixes the measure, §5.5 reduces commitment's adherence to that measure to one principle — and what was, in §5.3, an open reproduction problem of unbounded shape is now the conjunction of one inherited uniqueness theorem, two structural assumptions, and one named, isolated, structurally-refutable principle about commitment.

One inheritance asymmetry must be flagged, because it bears on how secure §5.5 is. §5.4's narrowing of the *weighting* drew strength from four-way convergence — four routes agreeing on the same measure. §5.5's reduction of *reproduction* does not inherit that convergence: it rests on the consistency route *singly*, because premise 7 is the consistency programme's uniqueness theorem specifically, and the other three routes of §5.4 play no part in it. A reader carrying the "four convergent routes" impression from §5.4 into §5.5 should therefore not over-rate §5.5's security: its conditional status is the consistency programme's alone, not the convergence's. Should the consistency uniqueness result fail, the other three weighting-routes would not rescue the reproduction reduction, which is built on consistency and would have to be rebuilt on whichever route, if any, also supplies a uniqueness theorem of the required form.

---

6. The Role of Probability

With the three hard problems located, the division of labour between probability and commitment can be stated cleanly, and it is a genuine division rather than a redundancy. Probability governs *which* admissible outcome is weighted how; commitment governs *that* an outcome becomes a fact, and (via §5.2) under what closure condition. Probability selects among possibilities in the sense of assigning them weights; commitment creates the fact. Neither does the other's work: a probability distribution over possibilities never, by itself, produces a fact — it remains a distribution — and a commitment event without the Born weighting would produce facts with the wrong statistics.

The relationship to the Born-rule programme is therefore complementary in the strict sense developed for the phase and Born structures elsewhere in the programme. The ODG/OIP results explain why probabilities attach to amplitudes as they do; the present proposal explains what physical event the probabilities are probabilities *of* — namely, commitment. Standard quantum theory has the weighting and not the event; Many-Worlds has neither event nor genuine weighting (only branch-relative frequencies it must work to recover); the present proposal has the event and imports the weighting, with the open reproduction condition (§5.3) being exactly the requirement that the imported weighting governs the event. The probability structure is unchanged; what is added is the physical referent for the probability — the commitment whose outcome the probability describes.

---

7. Position Among Interpretations

It is worth saying precisely where this proposal sits, because it is easily misfiled, and its empirical commitments depend on the filing.

On **ontology** it is *objective*: facts are observer-independent and definite. A committed fact is committed whether or not anyone observes it; the alternatives are really excluded, not merely excluded-relative-to-a-branch or relative-to-an-agent. This separates it sharply from Many-Worlds (which retains all branches) and from agent-relative information-theoretic readings (in which outcomes are relative to an information-holder). On this axis the proposal belongs to the objective-collapse *family* — it holds that superpositions really resolve into definite facts.

On **dynamics**, however, it is *not* a collapse model in the GRW sense, and this is the crucial distinction. Dynamical-collapse theories modify the Schrödinger equation, adding a stochastic nonlinear term so that superpositions of macroscopic systems decay. The present proposal adds *no term to the evolution law*. Unitary evolution is retained, unmodified, as the framework's description of the system while no commitment has occurred; commitment is not a perturbation of that evolution but the substrate event of which the unitary description is the pre-commitment account. The arrow $|\psi\rangle \longmapsto |i^\star\rangle$ of §4 is not a dynamical process competing with Schrödinger evolution at some rate; it is the change of ontological status that the framework's notion of fact-creation already supplies.

The consequence is significant and is the subject of §12: because the proposal does not modify unitary evolution, it does *not* inherit the empirical signatures of GRW-type models (the predicted deviations from the Schrödinger equation, the spontaneous energy gain, the relativistic tensions). It is, in its dynamics, as conservative as Copenhagen — it changes no equation — while being, in its ontology, as definite as objective collapse. Its nearest relations are therefore: Copenhagen on dynamical conservatism (no modified evolution) but against Copenhagen on the observer (the observer is not what triggers definiteness, §8); objective collapse on objective definite outcomes but against it on dynamics (no modified evolution); relational accounts on the secondary role of the observer but against them on the absoluteness of facts (commitment is not relative). It is its own position: *objective and definite in ontology, unmodified-unitary in dynamics, with the definiteness supplied by commitment rather than by a collapse term or an observer.*

---

8. Records and Observers

A successful measurement leaves a record, and in this framework a record is not merely stored information — it is a committed fact. The chain is therefore

measurement → commitment → record,

with each arrow an identity of the same underlying event rather than a causal sequence of distinct ones: the measurement *is* the commitment, and the record *is* the committed fact, viewed as something that can be subsequently read. The appearance of stable classical outcomes follows from the irreversibility of commitment — a committed fact cannot be uncommitted, so the record, once formed, persists and is available for later comparison.

The status of the observer follows, and it is deflationary. Observers do not create facts; they encounter facts that are already committed. The role of the observer in this account is secondary rather than fundamental — an observer is a system that reads committed facts and forms further committed facts (its own records of them), not a privileged agent whose attention precipitates definiteness. This is what allows the proposal to discharge the Copenhagen dependence on a measuring agent: the agent is replaced by the closure condition of §5.2, which is observer-independent. Whatever triggers commitment, it is not being-observed; it is the closure-saturation condition holding of the configuration, and that condition holds or fails regardless of whether the configuration is part of an apparatus or a laboratory or an unwatched corner of the world. Measurement is not special among commitments; an apparatus is simply a system arranged so that closure saturation occurs reliably and leaves a legible record.

---

9. Schrödinger's Cat Revisited

The cat thought experiment is, in this framework, a confusion of two categories rather than a paradox. Before commitment, both *alive* and *dead* are admissible possibilities — the closure structure sustains both — and the joint system is in the corresponding superposition. The puzzle in the standard telling is that the cat seems to be both at once until observed, which offends intuition. The framework's reply is that "both at once" describes the *admissible possibilities*, not a committed fact: there is no fact of the matter about the cat's state precisely because no commitment has yet occurred, and admissibility of both is not the simultaneous truth of both.

Commitment resolves this without an observer. When the closure-saturation condition is met — which, for a system as large and richly coupled as a cat, occurs overwhelmingly before any human opens the box — one outcome becomes a committed fact and the other is excluded. The observer who opens the box discovers the fact; the observer does not create it, and the fact was committed at saturation, not at the door. The paradox arises only when possibility and fact are treated as the same category, so that "both admissible" is misread as "both factual." In the commitment framework they are distinct categories, and the cat is in a definite committed state well before observation — exactly the intuition the standard telling violates, here recovered. (This is, transparently, conditional on §5.2: it is the closure-saturation criterion that does the work of committing the cat before the door opens, and that criterion is the open input. The cat's resolution is therefore as secure as the commitment criterion is, and no more — which is the honest place to leave it.)

---

10. The Apparent Circularity of Time and Commitment

An objection must be faced directly, because it threatens the whole account. The framework defines physical time through committed fact (§3); it identifies measurement with commitment (§4); and §5.2 argues that the commitment criterion cannot be clock-indexed because time is downstream of commitment. Does this circle? If time is defined by commitment and commitment is what measurement is, has anything been explained, or has a single primitive been described in three ways that lean on each other?

The reply is that the account is a *non-reductive identification*, not a definition chain, and the distinction is what saves it from circularity. The proposal does not define commitment in terms of time, nor time in terms of measurement, in a sequence that closes on itself. It identifies three things — the advance of time, the formation of a record, and the occurrence of a measurement — as three aspects of one primitive event, commitment, which is itself *not* defined in terms of any of them. Commitment is taken as primitive **(inherited)**; time, records, and measurement are then understood as facets of it. There is no circle because nothing in the chain is being reduced to anything else in the chain; rather, several phenomena previously treated as independent are being unified under one primitive that stands outside the list.

This has a cost the paper accepts openly: commitment itself is not explained. The account does not derive commitment from something more basic; it takes commitment as the substrate primitive and shows that measurement, time, and records are not separate mysteries but the same one. That is a genuine explanatory gain — it reduces the number of independent primitives — but it is not a derivation of commitment from nothing, and the paper does not pretend otherwise. The measurement problem is dissolved into the commitment primitive; the commitment primitive remains primitive. Whether that is satisfying depends on whether one regards a unifying primitive as an explanation or as a relocation of the mystery — but it is, at minimum, *one* mystery where there were several, and the unification is the claim.

---

11. Relation to the Conditional Gate-3 Closure

The Gate-3 transport branch suggests that commitment may leave a persistent topological residue — that the creation of a fact marks the structure of reality in a way no later motion can erase.

The relation between that branch and the present proposal must be stated carefully, because the dependence runs in only one direction.

Measurement-as-commitment does **not** require the residue. The identification of measurement with commitment, the resolution of the preferred basis into closure admissibility (conditionally), the deflation of the observer, and the cat's resolution all stand whether or not a surviving transport residue exists. The residue is downstream: it concerns what commitment *leaves behind* in the transport structure, not what commitment *is*. So the present paper rests on commitment, not on the residue, and the conditional Gate-3 closure can fail entirely without touching the measurement proposal.

If the residue survives, however — if the Gate-3 branch closes positively, which is itself conditional on the assumption set of the companion closure paper — then the interpretation gains a layer it does not otherwise have. Commitment would then do more than create facts; it would leave a persistent memory of fact-creation itself, and measurement would be not only the event that produces an outcome but the event that writes that outcome into the permanent topological record the phase programme identifies with quantum phase. On that stronger picture, the same commitment event that resolves a measurement is the one that contributes to the memory residue, and measurement and phase-as-memory become two consequences of one process. This is an *additional layer, conditional on the Gate-3 closure*, and the paper claims it only in that conditional form. The base proposal — measurement is commitment — is independent of it.

---

12. Empirical Status

The honest baseline, following directly from §7, is that the proposal is interpretive and observationally equivalent to standard quantum mechanics in every regime where commitment reproduces Born statistics (§5.3) and modifies no evolution law (§7). Because no term is added to the Schrödinger equation, there is no predicted deviation from unitary evolution of the GRW kind to look for; because the Born rule is imported unchanged, the outcome statistics are exactly those quantum theory already predicts. On this baseline the proposal claims an *explanatory* gain — measurement acquires a physical referent, the preferred basis acquires a candidate origin, the observer is deflated — and no *empirical* gain, which is a coherent and common status for a foundational reinterpretation.

There is, however, one place a genuine signature could appear, and it is the same place the companion programme locates empirical contact, which is consistent rather than coincidental. If commitment leaves the Gate-3 residue (§11), and if that residue has the discrete ($\mathbb{Z}_7$-graded) structure the closure programme assigns it — the $\mathbb{Z}_7$ here being the transport-residue grading, a distinct object from the $\mathbb{Z}_2$ fold register settled elsewhere in the Gate-3 arc, not to be conflated with it — then near the substrate scale, where the continuum description of phase breaks down, measurement outcomes could carry a trace of the discrete commitment structure that a primitive-collapse or primitive-phase account has no reason to produce. This routes the empirical question for measurement-as-commitment through the same conditional residue that the phase programme's signatures depend on: there is no separate measurement-specific prediction, but there is a shared one, and it is **(open, conjectural)** exactly as it is there. The proposal is interpretive *so far*, with its potential empirical content concentrated in the conditional residue rather than in any modification of measurement dynamics.

One sharper point distinguishes this proposal from Copenhagen despite their shared dynamical conservatism. Copenhagen is interpretive with *no* route to empirical content — the measurement primitive is, by construction, untestable. The present proposal is interpretive *with a conditional route*: if the Gate-3 residue survives and has observable structure, measurement-as-commitment acquires a signature Copenhagen cannot have, because the signature lives in the commitment residue that Copenhagen's primitive measurement does not posit. So the two are equally conservative dynamically but not equally sterile empirically — one has a conditional handle, the other none.

---

13. What This Paper Does and Does Not Claim

The paper does **not** prove that commitment occurs; commitment is inherited from prior VERSF work **(inherited)**. It does **not** derive the Born rule; that result is imported from the ODG/OIP programme **(inherited)**. It leaves open one *identification* — the preferred basis, where a candidate exists (§5.1) but its coincidence with the observed pointer basis is unproven — and two *commitment-dynamics inputs*: the commitment criterion (§5.2), fully open, and Born reproduction (§5.3), open but narrowed. On Born reproduction it claims a narrowing, then a reduction, not a closure: §5.4 constrains the weighting to the unique quadratic measure by four convergent routes sharing the operational geometry, and §5.5 reduces reproduction itself — whether commitment *selects* with that measure's frequencies — to the Commitment Minimality Principle, conditional on which (together with two standard structural assumptions on the commitment weighting, composition and non-contextuality, and the inherited consistency uniqueness theorem) reproduction follows. The novel and conjectural part of that conjunction is the single Minimality Principle — which §5.5 in turn argues may be grounded in a deeper structural claim, the Commitment Completeness Principle, by a transcendental argument that leaves the conjectural status unchanged in kind while clarifying its content. The remainder of the conjunction is assumed or inherited scaffolding the §5.5 proof names explicitly. What remains open is therefore not which measure commitment should follow, nor an unbounded reproduction problem, but the truth of one isolated, structurally falsifiable principle (or equivalently the completeness claim grounding it), against a stated backing of standard and inherited inputs. It does **not** require the Gate-3 residue; that relation is a conditional additional layer **(§11)**.

The paper **does** argue, conditionally on the inherited commitment and probability structures, that measurement is naturally identified with commitment, so that the measurement problem is transformed from *"why does the wavefunction collapse?"* into *"how are facts created?"* — a question the framework must answer regardless of measurement. It **does** offer a candidate origin for the preferred basis (closure admissibility) that standard quantum theory lacks, while marking the identification with the pointer basis open. It **does** establish the *form* the commitment criterion must take (an intrinsic closure condition, not a clock-indexed rate) and *why* (emergent time), ruling out the rate-based criteria of standard collapse models. It **does** separate selection from weighting cleanly enough to state the Born-reproduction condition precisely. And it **does** position the proposal exactly: objective and definite in ontology, unmodified-unitary in dynamics, with definiteness supplied by commitment rather than by an observer or a collapse term.

The honest summary is that the paper converts the measurement problem into a problem about the commitment dynamics, shows that this conversion is more than relabelling at one point (the preferred basis, where the framework offers a real candidate), and leaves two named gaps: the commitment criterion, fully open, and Born reproduction, now reduced — via §5.4's fixing of the measure and §5.5's reduction of adherence — to whether the Commitment Minimality Principle holds, against a backing of two standard structural assumptions and one inherited uniqueness theorem. Those are the concrete open inputs a fuller theory of commitment would have to supply — the first untouched, the second compressed from an unbounded reproduction problem into one conjectural principle plus named, assumed-or-inherited scaffolding.

---

14. Conclusion

The quantum measurement problem has traditionally been framed as a conflict between continuous quantum evolution and definite observational outcomes, with no physical event identified as the bridge. The VERSF framework suggests a different framing: measurements are not primitive — facts are — and a measurement is the physical transition by which one admissible possibility becomes an irreversibly committed fact. The appearance of definite outcomes, records, observers, and physical time then emerge from the single underlying process of commitment, and measurement is not an exception to the rules governing reality but one of the mechanisms by which reality is built.

The reframing is cheap, and the paper does not rest on it. Its substance is the confrontation with the three components of the measurement problem the reframing does not dissolve: the preferred basis, for which the framework offers closure admissibility as a candidate origin (conjectural, with the pointer-basis identification open); the commitment criterion, which emergent time forces into an intrinsic closure condition rather than a clock-indexed rate, and which remains the principal open dynamical question; and Born reproduction, which separates the imported weighting from the commitment selection and asks whether the second obeys the first.

On the last, the probability problem changed status in two steps, set out in §§5.4–5.5 and §13 and not re-derived here. First, §5.4 fixes the *weighting* to the unique quadratic measure $P(i) = |c_i|^2$, by convergent reconstruction results outside this paper, so commitment no longer chooses among candidate probability laws. Second, §5.5 reduces commitment's *adherence* to that measure — reproduction proper — to the Commitment Minimality Principle, against two standard structural assumptions and an inherited uniqueness theorem. The novelty in that conjunction is the single conjectural principle; the rest is scaffolding the §5.5 proof names rather than hides, and the reduction rests — unlike the weighting — on the consistency route singly, not on §5.4's convergence. So reproduction is neither determining the measure nor an open search, but the truth of one isolated, structurally refutable claim about the nature of commitment. These are the honest residue of the proposal, and naming them precisely — at their true and now-unequal sizes — is most of what the paper does.

What stands, conditionally, is this. If commitment, the Born rule, and the closure structure are granted — each inherited, none proved here — then measurement requires no separate collapse postulate, because the event collapse was invented to name is the commitment the framework already runs on. Quantum probabilities are unchanged; unitary evolution is unmodified; the observer is deflated from trigger to witness; and the definiteness of outcomes is supplied by commitment, observer-independently and without a new dynamical law. The proposal is interpretive, observationally equivalent to standard quantum mechanics except where the conditional Gate-3 residue might leave a trace, and explicit about the two commitment-dynamics inputs it leaves open.

The measurement problem is thereby not solved but *relocated* — moved out of the formalism's two modes and into the framework's theory of commitment, where it becomes the question of what physical condition forces a possibility to become a fact, and whether that forcing obeys the Born weights. That is a sharper and more tractable question than the one it replaces, and locating it is the contribution.

The wavefunction does not collapse.

A fact is committed — and the possibilities that were not committed were never anything more.