The Commitment Criterion

Fact Production as Relational-Readiness Selection in a Finite-Capacity Commitment Contest

Keith Taylor — VERSF Theoretical Physics Programme

General Reader Summary

The idea this paper explores

Earlier work argued that reality is built from commitments *— events in which one possibility becomes a fact. But it left one question open: what decides which possibility becomes the fact, and with what odds? This paper proposes an answer with a familiar shape — a* race *among the candidate facts — but is careful about one thing the racing image hides. A race usually means "who is fastest," and "fast" needs a clock. Yet this programme says commitment is what* creates *time. So the deepest version of the picture cannot be about speed at all. It must be about which candidate has the most* relational support *— and time is what emerges once those contests start resolving, one after another.*

The companion paper *Measurement as Commitment* identified measurement with commitment and then admitted, plainly, that it could not yet say what physical condition selects one possibility as the fact, or why the outcomes come out with the frequencies the Born rule assigns. It named those two gaps and located them in one object: the commitment dynamics. This paper is an attempt at that object.

The proposal is that commitment is a contest among candidate facts — the admissible ways a single unresolved structure could be completed into one record. Each candidate carries a *relational readiness*: a measure of how much pairwise relational support, in the sense the probability programme makes precise, reinforces it. One candidate is selected to complete — *not* always the one with the most readiness, but each with a chance equal to its share of the total readiness, so a strongly supported candidate is likelier, not certain, to be the one that completes. Completing consumes the limited capacity available, and the others can no longer commit. Thresholds established by earlier work — the conditions a record must meet to become permanent — are the *finish line* of the contest, not the commitment itself; commitment is the event of being the completion that the readiness contest selects.

One subtlety is the heart of the paper. It is tempting to describe this contest as a *rate* race — each candidate completing at some rate per unit time, the fastest winning. But a rate presupposes a clock, and the programme insists time is downstream of commitment, not upstream. So the substrate-level primitive cannot be a rate. It is the readiness *ordering* itself — an atemporal

comparison of relational support among candidates — and the familiar rate-and-first-passage description is its *continuum shadow*, valid only once enough commitments have accumulated to constitute a local sense of time. We keep both levels visible: the readiness ordering as the primitive, the first-passage race as its emergent-time image.

Two things make this more than a metaphor. First, a clean fact about contests of this kind: when the readiness measure carries no memory of a candidate's failed past attempts, the probability that a given candidate wins is exactly its *share* of the total readiness. If readiness is proportional to the quantum weight, the winning frequencies come out as the Born rule, the overall scale cancelling. Second, that no-memory condition is not arbitrary: it is the dynamical face of a principle the companion paper had already isolated separately — that commitment uses no information beyond the present configuration. The two papers were circling one assumption from two directions.

The honest limits are stated throughout. That readiness is proportional to the quantum weight, and that it carries no accumulation, are the inputs the result rests on; neither is proved here, and both are marked. And the proportionality, made precise, *is* the old question of why quantum probability is quadratic — the paper does not hide that this is where the real difficulty now lives, but shows that everything else reduces to it.

Abstract

Prior VERSF work established that physical facts arise through irreversible commitment events: the Fold Interface programme derived a commitment boundary separating reversible possibility dynamics from committed reality, and the Fact Production programme identified three necessary conditions for commitment — distinguishability, environmental amplification, and available commitment capacity **(inherited)**. None of these, individually or together, specifies the event that selects one alternative as the fact. This is the *commitment criterion* the companion paper *Measurement as Commitment* marked **(open)** as the principal dynamical input.

We propose a dynamical completion. Commitment is a **contest among candidate facts** — the admissible completions of a single unresolved structure — decided by a **relational readiness** R_i, a cardinal measure of the pairwise relational support reinforcing each candidate. A candidate is selected to complete with probability equal to its readiness share, $R_i / \Sigma_j R_j$; finite capacity makes the outcome winner-take-all; established thresholds (S_commit, $\chi(L)$, amplification criteria) are reinterpreted as the *finish line* rather than the commitment event itself. The competing objects are not separate local alternatives but candidate resolutions of a single global constraint structure (§2.1), so the contest is constraint-resolution, not signalling between branches.

A foundational point governs the formalism. A *rate* $\lambda_i = dN_i/dt$ presupposes a time parameter, but this programme holds that commitment *constitutes* temporal order; a rate-driven race would therefore assume the very thing it exists to explain. We avoid the circularity by taking the substrate primitive to be the atemporal readiness *ordering* among candidates, and treating the

first-passage rate description — exponential competition, the winning law — as its **continuum shadow**, valid only once accumulated commitments constitute a local temporal order (§6, §11). The correspondence between the substrate readiness ordering and the emergent-time rate description is itself marked **(open)**: it is what licenses reading the proven first-passage Lemma as the continuum image of substrate readiness-selection.

The technical core is an elementary selection result together with its honest scope. For candidates whose readiness carries no accumulation from unresolved attempts, the probability that candidate i is the completion is its readiness share, $P(i) = R_i / \Sigma_j R_j$ — a *posited* linear-share selection rule at the substrate level, proven only in the continuum shadow (below), not a substrate theorem. Granting readiness proportional to the quantum weight, $R_i \propto |c_i|^2$ **(conditional)**, this yields $P(i) = |c_i|^2$ with the overall scale cancelling — Born reproduction. The result holds *only* under the no-accumulation condition; were readiness to build from a candidate's past near-misses, share-winning would fail and the frequencies need not be Born. This no-accumulation condition is the readiness-level form of the **Commitment Minimality Principle** of the companion paper — that commitment access no information beyond the present admissibility structure — since accumulated-attempt history is precisely such forbidden information. The two papers' open inputs are thereby identified as one. In the emergent-time continuum, the readiness-share law is the shadow of the first-passage Lemma $P(i) = \lambda_i / \Sigma_j \lambda_j$ for memoryless competition, which the paper proves and reads as the continuum image of substrate readiness-selection — the proven content being the shadow-lemma, the substrate share rule a posit it carries back via an open correspondence.

A second cancellation addresses capacity. The effective finish condition is modulated by available distinguishability capacity (bit density, substrate saturation), which could couple to candidates unequally and distort the weighting. We isolate the **Capacity Neutrality Principle** — every primitive commitment consumes the same branch-common unit of capacity, with outcome-specific elaboration written downstream of the commitment event — grounded in the inherited account of the minimal commitment as a single structural record (a one-fold, one-record, one-fact identity — a count of records, not the outcome's information content; §5). Under Neutrality the capacity factor cancels from selection: capacity governs *how readily* and *whether* the contest resolves, readiness-share governs *which* candidate wins. We are careful about standing: since selection sees only the *total* readiness R_i (incorporating any common capacity factor), Born-consistency imposes one constraint, $R_i \propto |c_i|^2$, of which Capacity Neutrality and the readiness law are two coupled faces — given Born, fixing either fixes the other — with Born-consistency deriving neither. They are coupled but not equally supported: Capacity Neutrality carries a Born-free operational argument (the one-fold structural event) that the readiness law lacks, and the readiness law carries the exponent. But the exponent is not a debt this paper owes: **the quadratic measure $|c_i|^2$ is inherited from the probability-reconstruction programme as derived under stated admissibility conditions, and inherited with its conditional status — this paper does not re-derive the exponent and neither strengthens nor weakens those conditions (inherited; conditional at programme level)**. What is open *inside this paper* is the narrower identification — that a candidate's readiness is set by its own relational-support measure (§4). We are nonetheless explicit that the inherited target is the deep one: **the proportionality $R_i \propto |c_i|^2$ is, at its core, the quantum-probability problem itself** — the question of why the admissible weight is quadratic, addressed by Double Square and Physical Necessity under assumptions argued physically necessary rather than mathematically inevitable. The contribution is not to close that question but to show that the entire measurement problem reduces to it, and that the reconstruction programme has already narrowed the escape routes around it.

The framework supplies outcome uniqueness from finite-capacity winner-take-all rather than from a collapse postulate, reconciles with the companion paper's "forced, not stochastic" criterion by separating *that* commitment occurs (forced by saturation) from *which* candidate wins (the readiness contest), and links record formation, persistence, and the advance of temporal order to one substrate process — with the readiness-ordering primitive making temporal order genuinely downstream of commitment rather than presupposed. It is deliberately agnostic about how committed history is represented after the contest resolves — the phase-as-memory question of companion work — answering only what causes a fact, not how it is remembered. It modifies no evolution law and is observationally equivalent to standard quantum mechanics wherever the single Born-consistency constraint on readiness and the no-accumulation condition hold.

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

Table of Contents

---

1. The Remaining Question

The commitment programme has reduced the measurement problem in stages. *Measurement as Commitment* identified measurement with commitment, removing the need for a separate collapse postulate; the admissibility programme fixed the probability *weighting* to the unique quadratic measure; and §5.5 of that paper reduced Born *reproduction* — whether commitment selects with those weights — to a single named principle. But the reduction left a hole at its centre, and the companion paper was explicit about it: it could establish the *form* the commitment criterion must take, but not the criterion itself.

The form is constrained, and the constraint is sharp. Because physical time is constituted by committed fact, the criterion cannot read "commitment occurs at time t" — there is no pre-existing clock against which to time it, since the clock is built from the commitments being timed **(inherited** from the temporal-emergence results**)**. This cuts deeper than ruling out clock-indexed collapse rates: it warns against the word *rate* altogether. A rate $\lambda = dN/dt$ is a count per unit time, and "per unit time" is exactly the structure the programme says commitment must *produce*. A criterion built on rates would therefore assume the temporal order it is meant to explain. The criterion must instead be an intrinsic, atemporal condition on the substrate configuration, one whose holding *is* an increment of temporal order rather than an event within an already-running time. That requirement rules out the rate-indexed stochastic-hitting criteria of standard objective-collapse models — and, as §4.3 develops, it pushes the positive criterion toward a *relational* rather than a *temporal* primitive.

The remaining question is therefore twofold, and the companion paper noted that its two halves are facets of one object:

What physical condition selects one admissible possibility as the fact, and does that selection occur with the Born frequencies?

A threshold condition alone is insufficient for either half. A threshold may state *what makes a record permanent enough to count*; it does not state *which* admissible possibility achieves such a record, nor *with what frequency*. The gap between "a threshold exists" and "this possibility, this often" is exactly the gap this paper addresses — and it addresses it, crucially, without reaching for a clock.

---

2. The Race Interpretation

Consider an admissible set of alternatives

$\mathcal{C} = \{C_1, C_2, \ldots, C_n\}$, each a way the current record could be extended by one committed fact while preserving admissibility. After decoherence the alternatives are distinguishable: interference between them is suppressed, and they evolve as independent candidate extensions **(inherited**, decoherence/distinguishability**)**. At this stage there is no fact — only distinguishable possibilities competing for a status none yet holds.

We propose **(conjectural)** that commitment is a contest among these candidates, decided by *relational readiness*. Each candidate carries a readiness R_i — a measure of how much admissible relational support reinforces it (made precise in §4); a candidate is selected to complete with probability equal to its readiness share (§6), so the more strongly supported a candidate, the likelier — but not the certain — completion. The proposal is not that a threshold is crossed — thresholds are retained, with their role re-specified in §3 — but that commitment is the *completion the readiness contest selects*. The distinction is the whole content of the picture: a threshold defines what completion requires, and commitment is the candidate the readiness contest picks out to complete it.

A word on language, developed fully in §4.3. It is natural to describe this as a *race* — candidates "completing first," a "first-passage" contest — and we keep that language because it is vivid and because, once temporal order exists, it is accurate. But the racing image carries a clock the substrate does not have, so the race is the *continuum shadow* of the primitive, not the primitive itself. The primitive is the atemporal readiness ordering; "first" is what that ordering looks like after temporal order has emerged. We use both registers deliberately and mark which is which.

This is a candidate answer to the §1 question, and like the companion paper's offerings it is a candidate, not a theorem. What recommends it is that, given the inputs made explicit below, it delivers the second half of the question — the Born frequencies — from elementary structure, and delivers outcome uniqueness without further assumption. What it costs is those inputs, each marked where it enters.

2.1 Commitment Races as Constraint Resolution

The language of competing branches may suggest that physically separate alternatives — distinct local objects, each already in some sense *there* — are racing against one another. That is not the intended picture, and the misreading matters, because it would import a localist ontology the rest of the programme rejects and would invite a spurious worry about signals propagating between separated branches.

Companion work on entanglement and pre-spatial structure argues that an unresolved quantum state is a *global constraint structure*, not a collection of independently existing local facts **(inherited)**. Before commitment there is no set of completed records sitting in different places; there is a single admissible constraint structure that admits multiple resolutions. The race is therefore not among separate classical alternatives but among *candidate completions of one shared constraint*:

global constraint structure → admissible resolutions → commitment race → committed record.

What competes is not a collection of already-existing facts but a collection of *candidate* facts — admissible ways the single constraint could be discharged into one record. The winner becomes the record; the others were never facts and remain unrealised possibilities, in the strict sense of §8: a losing branch is not an outcome that happened elsewhere but an attempt that did not complete.

This reading carries three advantages and costs nothing mathematically. It removes any suggestion that commitment requires propagation between spatially separated branches — there are no separated branches to signal between, only one constraint resolving. It stays compatible with the pre-spatial account of entanglement, in which correlations are properties of a single global mode rather than interactions among local systems, so the Bell-relevant globality of §11 is the *same* globality, here named at its source. And it aligns with the ontology of fact production: the contest is among candidate record-completions of a single unresolved constraint, won by exactly one. None of the selection mathematics changes — the readiness-share law of §6, its continuum first-passage shadow, the cancellations all stand. What changes is only the answer to *what the competing objects physically are*: not branches in space, but resolutions of a constraint.

---

3. Thresholds as Finish Lines

Earlier work introduced several threshold constructs — the commitment entropy scale S_{commit}, the capacity-response function $\chi(L)$, the amplification and barrier criteria — and the present interpretation retains all of them unchanged **(inherited)**. What it adds is a dynamical reading of what they *are*.

In the standard telling a threshold is the commitment: reach it and the fact exists. In the contest telling a threshold is the *finish line*: it defines the condition a candidate must complete, but which candidate completes it is settled by relational readiness, drawn according to readiness share (§6). The analogy is exact rather than decorative. A finish line does not run the contest or pick the winner; it defines what completion means. Likewise the thresholds define what commitment requires; the readiness contest among candidates is what selects which one achieves it.

This reassignment is conservative: it changes no threshold value and adds no new threshold. It changes only the claim about what a threshold *does* — from "constitutes commitment" to "defines completion" — and that change is what makes room for a selection mechanism, since a finish line, unlike a commitment, is something several candidates can be contending toward at once.

A simplifying assumption is made explicit here. We take the finish line to be *common* across candidates — every candidate contends toward the same irreversible-record condition — so that any candidate-dependent barrier is absorbed into the candidate's *effective* readiness rather than appearing as a different threshold per candidate **(conditional)**. Genuinely candidate-dependent thresholds would complicate the winning law of §6 and are flagged **(open)** here, then resolved in

§5: the natural physical source of candidate dependence is capacity cost, and §5 argues that cost is branch-common, so the clean result holds. The clean result assumes the common finish line.

---

4. Relational Readiness, Not Rate

The contest needs a quantity that decides which candidate completes. The Tick-Bit framework supplies the raw material: each admissible candidate generates microscopic commitment attempts — *ticks* — that may or may not complete the finish-line condition **(inherited)**. The naive move is to assign each candidate a *rate* λ_i at which it attempts completion and let the fastest win. We resist that move at the substrate level, for the reason §1 gave and §4.3 develops: a rate is a count per unit time, and time is what the programme says commitment produces. So we take the deciding quantity to be not a rate but a **relational readiness**:

R_i = the cardinal measure of admissible relational support reinforcing candidate C_i.

Readiness is a *magnitude*, not a mere ranking — it must be cardinal, because the winning law selects *shares* of total readiness (§6), and shares require quantities, not just an order. But it is defined relationally and atemporally: R_i measures how much pairwise relational support (§4.2) stands behind candidate i, with no reference to seconds or to a rate of anything. Ticks remain the substrate events; readiness is what their relational structure amounts to. A candidate of greater readiness is not one that "runs faster" but one that is *more strongly supported* among the admissible relations.

The selection law the construction needs is then

$R_i \propto |c_i|^2$ **(identification open here; exponent inherited-conditional from the reconstruction programme)**.

Its status must be stated exactly, and more carefully than "readiness tracks the weight" would suggest. The naive reading takes $|c_i|^2$ to be an external dial — a static weight that readiness is calibrated against. That misdescribes what the quantum weight *is*. On the Double Square ontology (§4.2), $|c_i|^2$ is not a speed, an energy, or a clock-reading; it is itself a *measure of relational support* — how strongly the pairwise distinguishability network reinforces candidate i. So the readiness law is not "readiness is proportional to an external weight" but a *relational identity*: the relational support a candidate carries, which is what its readiness *is*, is proportional to the quantum weight *because the weight is already the admissible measure of that same relational support*. Readiness and the weight are the same kind of quantity — relational support — not two things one of which tracks the other.

This relocation strengthens the *motivation* for the law without discharging its open component, and the distinction must be kept exact. What is inherited is the quadratic relational-support measure: the probability-reconstruction programme derives $|c_i|^2$ as the admissible measure under stated conditions — positivity, normalisation, relabelling invariance, geometric dependence, compositional factorisation (Double Square), argued physically necessary by the finite-distinguishability, irreversible-commitment, temporal-consistency, and compositional-stability considerations of Physical Necessity. That derivation is a *physical-necessity* result under those admissibility assumptions, not a mathematical inevitability from nothing, and **this paper inherits both the quadratic measure and its conditional status — it does not re-derive the exponent, and it neither strengthens nor weakens the conditions under which the exponent holds (inherited; conditional at programme level on the admissibility assumptions)**. What remains genuinely open *inside this paper* is therefore not the exponent but the *identification* — that a candidate's readiness is set by its own relational-support measure and no other admissibility datum (the single-candidate and functional-form layers below). As §5 makes precise, that identification is one face of a single Born-consistency constraint on total readiness; the substrate-level account that would close it sits on this face, while the exponent it would deploy is inherited from Double Square, conditional as Double Square leaves it.

A robustness point worth recording in advance: only the *proportionality* matters, not the scale. The winning law of §6 depends on ratios $R_i / \Sigma_j R_j$, in which any common factor cancels. The mechanism therefore need not fix the absolute scale of readiness to reproduce the Born frequencies. §5 shows that a second, larger common factor cancels for the same reason — the capacity cost of commitment — provided that cost is branch-common.

The debt is fenced by the paper's other conjecture, but it is important not to overstate by how much. The no-accumulation condition (§7) constrains the *form* of readiness — it carries nothing from a candidate's history of failed attempts. Minimality (§7) constrains what readiness may *depend on* — admissibility data alone, which after decoherence excludes phase and any dynamical quantity such as energy or coupling. Together they fix the *form-class*: readiness is a history-free function of admissibility data only. They do *not* collapse the residue to a single number, and three distinct things remain open inside that form-class.

First, *single-candidate reduction*. The present admissibility structure is the whole configuration $\{c_j\}$, not c_i in isolation, and Minimality — "depend only on the present admissibility structure" — does not forbid R_i from reading the other candidates' amplitudes: $R_i \propto |c_i|^2 \cdot f(\{|c_j|^2\})$ is Minimality-compliant. That readiness reads only its own candidate's relational support is a further assumption, doing real work. Second, *functional form*: even granting single-candidate dependence, that readiness is set by that support *linearly* — $R_i \propto$ (support of i), rather than by some monotone re-scaling of it — is an identification, not a theorem. These two are the genuinely paper-local open pieces: they concern whether *readiness* is the relational-support measure, candidate by candidate. Third, *the exponent*: why the relational-support measure goes as $|c_i|^2$ rather than $|c_i|^4$. This is the deepest layer and it is *the quantum-probability problem itself* — but it is not a debt this paper owes, because it is not re-derived here. It is inherited from Double Square as derived-under-admissibility-conditions, and inherited with its conditional status intact: closed modulo the admissibility assumptions, which would reopen the exponent if they failed (§4.1). So the honest residue sorts into two registers. *Inside this paper*: the identification layers — single-candidate reduction and linear functional form — by which readiness is the relational-support measure. *Inherited, conditional at programme level*: the quadratic exponent, the hardest layer, addressed by Double Square under its physically-necessary-but-not-inevitable assumptions, neither re-derived nor re-opened here. We mark the readiness law **(open** for the identification; **inherited-conditional** for the exponent**)**, at that true size — not collapsing the identification to "one number," and not claiming the exponent as a debt this paper either owes or discharges.

4.1 The Inherited Weighting

It is worth saying plainly what this paper does and does not take on regarding the quadratic weight, because the distinction is what keeps the Commitment Criterion a *measurement-mechanism* paper rather than a rival account of the Born rule.

The *weighting* — that admissible alternatives carry the quadratic measure $w_i \propto |c_i|^2$ — is not derived here. It is inherited from the probability-reconstruction programme **(inherited)**. The Double Square Rule programme establishes that irreversible selection acts on pairwise correlation structures rather than on individual reversible paths, and that under positivity, normalisation, relabelling invariance, geometric dependence, and compositional factorisation the resulting probability structure is uniquely quadratic. The Physical Necessity analysis independently shows that alternative probability structures fail admissibility tests grounded in finite distinguishability, irreversible commitment, temporal consistency, and compositional stability. The quadratic weight thus enters the present paper as a fixed external result; the contest does *not* determine it, and the paper makes no claim to.

The division of labour is therefore clean. The probability programme asks: *why is the admissible weighting quadratic?* The commitment programme asks: *how does one outcome become the fact while preserving that weighting?* Complementary, logically distinct.

One seam must be kept sharp, or this subsection would appear to discharge §4's open item by fiat. What is inherited is the *static weighting* $w_i \propto |c_i|^2$ on admissible alternatives. What §4 leaves open is the *readiness law* $R_i \propto |c_i|^2$ — the claim that the dynamical readiness of a candidate tracks that inherited static weight. These are not the same statement: the first is a fact about which probability measure is admissible, settled elsewhere; the second is a claim about what sets the *relational support* a candidate carries, and that this support should follow the inherited measure is precisely the identification §4 marks open and §5 places as one face of the Born-consistency constraint. So inheriting the weighting does not close the readiness law. It means the contest is not searching for a probability law — the law is fixed — but must still explain why its *dynamics* reproduce the fixed law during selection. The remaining task is not to derive a weighting but to show the mechanism *preserves* the inherited one. One consequence, drawn out in §4.2: because the readiness law is *not* the same statement as the weighting, the quadratic *exponent* is inherited only *contingently* — it transfers from the weighting to readiness exactly insofar as readiness is identified with the relational-support magnitude Double Square measured. The exponent's inheritance therefore rides on that identification rather than standing free of it.

4.2 Pairwise Selection and Relational Readiness

The interpretation of the quantum weight that makes readiness the natural primitive deserves to be stated once, sharply, as a named principle the rest of the paper leans on.

Relational Support Principle (inherited, interpretive)

The quantity $|c_i|^2$ is *not* interpreted as a speed, an energy, or a temporal rate. It is interpreted as the admissible measure of the *pairwise relational support* associated with candidate completion i — how strongly i is reinforced by the network of pairwise distinguishability relations. The quantum weight is, in this sense, already a relational quantity; it is inherited as such from the Double Square / Physical Necessity reconstruction (§4.1).

This single reading is what licenses a *relational* readiness as the substrate primitive rather than an arbitrary dynamical scalar, and the rest of §4.2 is its unpacking.

The inherited weighting comes with an *ontology*, and readiness is where the contest inherits it most directly. The Double Square programme does not merely assert that the measure is quadratic; it locates *why*: irreversible selection acts on pairwise correlation structures — relational "handshakes" — rather than on individual reversible paths, and the weighting is quadratic because the objects being selected are correlations, whose natural measure is bilinear. This is exactly the material relational readiness is built from. R_i is not "how fast candidate i runs" but "how much pairwise relational support reinforces candidate i" — and pairwise relational support is the Double Square primitive. The readiness ontology and the quadratic weighting therefore have *one* source: relational support among admissible completions.

This makes the alignment more than compatibility — it is identity of primitive. The chain is:

pairwise relational support → relational readiness R_i → [readiness-share selection] → outcome frequency.

The first arrow is the Double Square ontology **(inherited)**: readiness *is* accumulated pairwise support. The step $R_i \propto |c_i|^2$ of §4 is then the claim that this relational support is quadratic in amplitude — and here a contingency must be stated exactly, because §4.1 was careful to say the weighting and the readiness law are "not the same statement," and this subsection's "one question viewed twice" must be reconciled with that, not left to seem to contradict it. The reconciliation is this: the two are the same question *only on the condition that the selection-driving magnitude (readiness) is the very quantity Double Square measured*. Double Square is, in the first instance, a theorem about which *probability measure* on admissible alternatives is admissible. If readiness is identified with that measure's underlying relational-support magnitude, then "why is the weighting quadratic" and "why is readiness $\propto |c_i|^2$" are indeed one question, and its answer — the quadratic exponent — is inherited. But that identification is exactly the open piece §4.1 names: it is *not* separately given that the dynamical magnitude driving selection is the same quantity the measure-theorem quantified. So the honest statement is layered: **the exponent is inherited (conditional on Double Square's assumptions) *given* the identification; the identification — that readiness is Double Square's relational-support magnitude — is open here.** §4.1's "not the same statement" and this subsection's "one question viewed twice" are therefore both correct and consistent: they are the same question *after* the identification is granted, and two different statements *before* it. The exponent's inheritance rides on the identification; it is not free-standing. With that contingency marked, the final arrow — that readiness-share selection reproduces the weighting as *frequencies* — is the no-accumulation-conditional, linear-share-posited result of §6, not delivered by the ontology alone. So the contest inherits the quadratic weighting and the relational ontology that grounds it *conditionally on the identification*, and adds the selection dynamics by which one supported completion becomes the fact while the inherited weighting is preserved.

The ontological point underneath this — now carried by the Relational Support Principle above — is what makes a *relational* readiness the right primitive rather than an arbitrary choice dressed in relational language. Because $|c_i|^2$ is already a measure of relational support and not an absolute per-candidate scalar, a mechanism that selects on a *relational readiness* is ontologically matched to the weight it must reproduce, whereas one selecting on a rate, an energy, or any absolute scalar would be reproducing a relational measure with a non-relational dynamics — a coincidence to be explained rather than a structure to be inherited. This is the force of "handshakes, not absolutes": the weight counts handshakes, so the mechanism selects on handshakes. Readiness $\propto |c_i|^2$ is then not a calibration against an external dial but the statement that the dynamics selects on the same relational support the weight already measures — with the *exponent* of that support inherited, conditionally, from the programme that derived it.

4.3 Why Not a Rate: Readiness as the Atemporal Primitive

This subsection states explicitly the foundational reason readiness, not rate, is the primitive — because getting this wrong would have the paper assume the very thing it sets out to explain.

A rate is $\lambda = dN/dt$: a count of attempts per unit time. It presupposes a time parameter t against which to count. But the programme's deepest claim is that physical time is *constituted* by the accumulation of committed fact (§1, §14): there is no t prior to commitment. A commitment criterion phrased as a rate would therefore define commitment using a quantity that itself presupposes the temporal order commitment is supposed to produce. That is circular, and the circularity is not cosmetic — it is the difference between explaining time and quietly assuming it.

Relational readiness escapes the circle because it is *atemporal*. R_i measures how strongly the pairwise relational network supports candidate i — a magnitude attaching to the constraint structure, comparable across candidates, that requires no clock to be true. The primitive content of the contest is these cardinal support magnitudes $\{R_j\}$ and the shares they define, *not* a set of rates and *not* a bare ordering (shares need magnitudes, §4, §6). Nothing about "support" counts anything per second.

Where, then, does the rate language of a "race" come from? It is the *continuum shadow* of relational readiness, in the same sense the programme treats other continuum objects as shadows of substrate primitives. Once enough commitments have resolved to constitute a local temporal order, the readiness magnitudes can be *re-described* as if candidates were completing at rates $\lambda_i \propto R_i$ per unit of that emergent time, and in that description the first-passage formalism of §6 applies and the rate-share law holds. The rate is real at the emergent level and fictitious at the substrate level. We therefore read the proven first-passage Lemma (§6) as the *continuum image* of substrate readiness-share selection — and we mark the correspondence between the atemporal readiness magnitudes and the emergent-time rate description as **(open)**: that the continuum shadow faithfully represents the substrate quantity is an assumption the paper makes and a fuller theory would have to establish.

The payoff is that the mechanism is now consistent with the programme's deepest claim rather than in tension with it. Commitment does not occur *at a rate*; it occurs as the resolution of a readiness contest, and rates — like time itself — are what emerge once those resolutions accumulate. The order is: relational support → readiness ordering → commitment → temporal order → (only now) rates. The paper that began by promising not to smuggle in a clock keeps that promise here.

---

5. Capacity Neutrality: the Branch-Common Commitment Unit

The §3 finish line was taken to be common across candidates, with candidate-dependent barriers absorbed into effective readiness and the genuinely candidate-dependent case flagged **(open)**. That open case has a concrete physical origin worth confronting directly, because it is where Born reproduction is most easily lost: the effective finish condition a candidate must meet is modulated by available distinguishability capacity — bit density and substrate saturation — and capacity could in principle couple to candidates unequally. This section isolates the assumption under which it does not, states it as a principle, and shows the principle is what the inherited Born weighting requires rather than an independent convenience.

The two cases. Suppose candidate i, on completing, consumes capacity ΔN_i. Write the total readiness as a bare readiness times a capacity factor,

$R_i = R°_i \cdot F(\Delta N_i, N/N\_BCB)$,

where $R°_i$ is the bare relational readiness of §4 and N/N_BCB measures substrate saturation. Two cases arise, and they decide everything.

Case A — branch-common cost. If every primitive commitment consumes the same capacity, $\Delta N_i = \Delta N_0$ independent of i, then F is one common factor $F(N/N\_BCB)$ and the winning law (§6) gives

$P(i) = R_i / \Sigma_j R_j = R°_i F / \Sigma_j R°_j F = R°_i / \Sigma_j R°_j = |c_i|^2$ (with the readiness law of §4).

The capacity factor cancels exactly as the overall scale did. Capacity then sets *how readily* fact production proceeds — abundant when capacity is plentiful, slowing as the substrate saturates, unable to resolve once capacity is exhausted — but it does not touch *which* candidate wins. **(proven**, given branch-common cost and the readiness law**)**.

Case B — outcome-dependent cost. If instead ΔN_i depends on the outcome — say $\Delta N_1 = 1$ but $\Delta N_2 = 5$, so outcome 2 demands five times the commitment resource and faces a correspondingly higher effective barrier — then F does not factor out:

$P(i) \propto |c_i|^2 \cdot G(\Delta N_i)$, G branch-dependent,

and the extra factor distorts the weighting: $P(i) \neq |c_i|^2$. The readiness-share selection still picks a unique winner and still produces facts, but the frequencies tilt toward outcomes with favourable capacity cost, and Born fails.

The principle that selects Case A. The discriminating question — does committing outcome i consume an i-dependent amount of capacity, or a common unit? — is answered twice over, by two arguments that run in opposite directions and converge. The first is *operational*, from the minimal structure of a commitment event: it says Neutrality is *true*. The second is by *consistency*, from the independently established Born weighting: it says Neutrality *must hold*. The two are worth keeping distinct, because together they make Neutrality not a convenient stipulation but an overdetermined one.

Capacity Neutrality Principle (conjectural)

Every primitive commitment event consumes the same minimal commitment capacity — the structural cost of writing **one fold, one record** — regardless of which candidate completes. This is a count of records, not a measure of the information content of the outcome: the surprisal $-\log|c_i|^2$ differs across candidates, but the structural cost of committing one record does not. Outcome-specific elaboration — memory content, geometric structure, particle content, macroscopic consequence — is written *downstream* of commitment and does not form part of the commitment event itself.

The operational ground (Neutrality is true). The support is the inherited account of what a commitment minimally *is* — but it requires one sharp distinction, because the word that names the minimal event has two readings that point in opposite directions, and only one grounds Neutrality. The Fold Interface programme derives the minimal commitment event as a single fold — one distinction written, one record formed — and the Fact Production programme treats commitment as the formation of one stable record **(inherited)**. The neutral quantity is this *structural* one: the **one-fold / one-record count**, which is 1 per commitment for every candidate regardless of amplitude. On that reading every completing candidate writes the *same* primitive structure — one fold, one record — and so consumes the same primitive unit of capacity; candidates differ in relational readiness, not in what it structurally costs to write one record.

This must be stated against the reading that would *destroy* Neutrality. If "the cost of committing outcome i" were read as its **information content** — the surprisal $-\log|c_i|^2$, the number of bits needed to *specify* outcome i among the alternatives — then the cost would be branch-*dependent*, larger for low-amplitude outcomes, and this is precisely Case B: a candidate-dependent ΔN_i that tilts the frequencies away from Born. The information reading does not merely fail to support Neutrality; it *is* the violation Neutrality must exclude. So the capacity that must be branch-common is the structural cost of writing one record, *not* the information content of the outcome written. We disavow the information-theoretic reading explicitly: Neutrality is the claim that committing is structurally one fold per fact, and it carries no commitment to the surprisal being equal across outcomes — indeed the surprisal is *not* equal, and Neutrality's content is exactly that the surprisal is not what the capacity cost tracks.

Two clarifications guard the slogan. First, the One Fold identity is best stated as *one fold = one record = one fact* (a structural count), not *one fold = one bit = one fact*, since "bit" invites the information reading just disavowed; where "one bit" appears it means one structural distinction, the binary special case of one record. Second, for a contest among $n > 2$ candidates the structural cost is still one fold / one record per commitment — branch-common for any n — whereas "one bit" would be imprecise, since *specifying* one of n outcomes is $\log_2 n$ bits; this is a further reason the neutral quantity is the one-record count, not a bit-count. With "one fold = one record" understood structurally, every completing candidate pays the same primitive cost, and the operational ground holds for any number of candidates.

The consistency ground (Neutrality must hold if Born does). The second argument does not appeal to the structure of commitment at all, only to a result fixed elsewhere. The Born weighting $P(i) = |c_i|^2$ is established independently of this paper, by the reconstruction programme **(inherited)**. Run the chain:

Born weighting fixed independently ⟹ any branch-dependent capacity cost would multiply it by a branch-dependent factor $G(\Delta N_i)$ ⟹ the selection frequencies would be $P(i) \propto |c_i|^2 \cdot G(\Delta N_i) \neq |c_i|^2$ ⟹ contradiction with the fixed weighting.

The only way to discharge the contradiction is G branch-independent — capacity cost branch-neutral. The direction of inference matters and must not be reversed: Born is the premise and Neutrality the consequence, not the other way round; the argument is "Born is fixed elsewhere, and a non-neutral capacity dynamics would contradict it, so capacity must be neutral," not "Neutrality holds because it yields the Born answer we wanted" (which would assume its conclusion).

But the exact reach of this argument must be stated carefully, because an earlier version of this section overstated it. Born-consistency is a *constraint*, not a *derivation*. It tells us what must hold *if* Born holds and the mechanism is right; it does not explain why the substrate produces a branch-neutral cost without already importing Born. And — this is the point that disciplines the whole section — the constraint is indifferent to how we split readiness. The winning law (§6) sees only the *total* readiness R_i. Born requires the single condition $R_i \propto |c_i|^2$. Our decomposition of R_i into a "bare readiness" R°_i and a "capacity factor" F is narrative bookkeeping; the selection structure does not recognise it. So the very same Born-consistency move that constrains F to be branch-neutral *equally* constrains R°_i to carry $|c_i|^2$ and not $|c_i|^4$ — if consistency forces Neutrality, it forces the exponent too, by identical logic. We cannot consistently treat Born-consistency as *forcing* the capacity piece while merely *motivating* the readiness piece. It does the same thing to both: it constrains them, given Born; it derives neither, absent a substrate account that does not presuppose Born.

So we decline the asymmetry an earlier draft drew — the one that had Born-consistency *forcing* the capacity piece while merely *motivating* the readiness piece. That was wrong: Born-consistency does the same thing to both, constraining each and deriving neither. But we must not over-correct into the opposite error of calling the two faces *equally supported*, because they are not, and the asymmetry runs in the paper's favour. Two things must be kept separate.

First, the faces *are* coupled, by the product constraint. There is a single **Born-consistency constraint** on total readiness, $R_i \propto |c_i|^2$, and since $R_i = R^\circ_i \cdot F$, fixing either face hands you the other: given Born, F branch-neutral $\Longrightarrow R^\circ_i \propto |c_i|^2$, and $R^\circ_i \propto |c_i|^2 \Longrightarrow$ F branch-neutral. In that logical sense they go together, and Born-consistency derives neither from scratch.

Second, and despite that coupling, the faces are *not* equally well supported, nor equally hard. The F-face has a **Born-free argument** — the one-fold/one-record operational ground below, which concludes branch-neutral cost from the structure of the commitment event without invoking Born at all (conjectural through its elaboration premise, but Born-free in form). The R°-face has **no analogue**: the "only admissibility-covariant scalar" gesture of §4 is not one, since $|c_i|^4$ is built from the same scalar, so naturalness fixes neither the functional form nor the exponent. And the hard residue — the exponent, which §4 rightly calls the quantum-probability problem — lives *entirely* on the R°-face; there is no exponent-like freedom on the F-face, which is branch-neutral or it is not. So the genuine open difficulty is concentrated on the readiness law, not split evenly, and Capacity Neutrality is the better-supported of the two. The honest statement is therefore: the faces are coupled by the product constraint and Born-consistency derives neither, but the F-face carries a Born-free operational argument the R°-face lacks, and the R°-face carries the exponent. "Two faces of one constraint" is right about the coupling; it would be wrong to read it as equal standing.

Can the operational ground discharge Neutrality outright? There is a route by which Capacity Neutrality could cease to be conjectural and become inherited. The One Fold programme advances the identity *one fold = one record = one fact*: that the primitive commitment event is a single fold, writing a single record, constituting a single fact. If that identity is *derived* in the One Fold work — established as a substrate theorem rather than posited as a definition — then branch-dependent primitive cost is impossible by construction: every commitment is one fold, one fold is one record, one record is one unit of structural capacity, so $\Delta N_i = \Delta N_0$ necessarily, and Capacity Neutrality is **(inherited)** rather than conjectural. We flag this **(inherited if One Fold derives the one-fold–one-record identity; conjectural if it posits it)**, declining to assert the stronger reading without confirmation that the identity is derived, in the same dual-marked spirit the companion paper used for the closure-admissible basis. One caution even under the strong reading: the identity secures that the *event* is one record, branch-common; it does not by itself secure the separate downstream-elaboration premise below — that no candidate must write *more than* the primitive record to win. So One Fold, even fully derived, discharges the branch-commonality of the event, and the elaboration premise remains to be carried.

With the consistency ground put in its place, the operational ground bears the real weight, and it is worth seeing that it is the non-circular one: it argues Neutrality from the structure of the commitment event — one fold, one record, common to all candidates — not from Born, so it offers a reason the substrate cost *is* branch-neutral rather than merely that it *must be* if Born holds. That argument has one load-bearing premise, the place Case B would re-enter: record *elaboration* is downstream of and separate from the commitment event. ΔN_i in the contest is the capacity consumed *to write the primitive record*, not the eventual size of the structure that record seeds. An outcome that ultimately involves a larger or richer record does not pay for that richness *at the finish line*; it pays the common primitive cost to become a fact, and the elaboration is written afterward, in territory this paper does not enter (§14). We mark the premise **(conjectural)** unless discharged by a derived One Fold identity: the operational one-fold/one-record argument supports it, the consistency constraint shows that *if Born holds* nothing else is admissible — but absent a derived One Fold result, neither proves from the substrate up that no candidate must write more than the primitive record to win. A candidate that did would falsify Neutrality and, by the consistency chain, tilt Born. So Capacity Neutrality stands as *strongly motivated* — operationally grounded in the one-fold/one-record event, required by Born-consistency, and inherited outright if One Fold derives its identity. It is in this respect in *better* shape than the §4 readiness law, not equal standing with it: the readiness law has no comparable Born-free argument and carries the exponent that is the hard residue, whereas Capacity Neutrality has the operational ground and no exponent-like freedom of its own. The two faces are coupled, but the difficulty is not where the symmetric phrasing would put it.

Proposition 5 (Capacity modulates readiness uniformly, not the winner) (conditional on Capacity Neutrality and the readiness law**)**

Under Capacity Neutrality the capacity factor is branch-common, $F(N/N\_BCB)$, so by the Case A computation the relative outcome probabilities are $P(i) = R°_i / \Sigma_j R°_j = |c_i|^2$. Capacity therefore modifies *how readily* fact production proceeds — whether and how abundantly the contest resolves — but leaves the winning frequencies at their Born values.

The separation this buys is clean and worth stating as the section's output:

capacity → *whether* and *how readily* commitment resolves; relational readiness → *which* candidate wins.

These are now decoupled. Saturation, bit density, and the capacity-modulated finish condition all bear on *whether and how readily* fact production proceeds; none bears on *outcome selection*, because the resource each fact consumes is branch-neutral and cancels.

Independence and the shared factor. One technical point reconciles this section with the continuum Lemma of §6, which in its rate image assumes the candidate first-passage times are independent. The capacity factor $F(N/N\_BCB)$ depends on global saturation — a *shared* resource — which might seem to correlate the candidates. It does not, for two reasons that should be surfaced as model assumptions. First, attempts are *free*: only the completing candidate consumes capacity (the winner-take-all of §8), so candidates do not deplete one another before the contest resolves, and their readiness contributions remain independent up to resolution. Second, F enters as a *common* modulation of every candidate's readiness, not as independent per-candidate noise — it rescales all readinesses by the same factor, and a common rescaling leaves the readiness-share law invariant (the ratios $R°_iF / \Sigma_j R°_jF$ are unchanged). So the shared capacity factor is consistent with the independence the continuum Lemma requires: it modulates the contest uniformly without coupling the candidates.

With Capacity Neutrality understood as the F-face of the single Born-consistency constraint, Proposition 5 closes the §3 candidate-dependent-threshold **(open)** item in the Born-safe direction — contingent on Neutrality, which the operational one-fold/one-record argument grounds, the consistency constraint requires given Born, a derived One Fold identity would inherit outright, and a single over-costly candidate would refute.

5.1 One Fold, One Finish Line

Capacity Neutrality has a compact form worth stating, because it answers in one line the question a reader poses immediately — *why does one candidate not need more commitment than another?*

Every primitive commitment event is the writing of one fold — one record, one distinction made permanent. So every candidate contends toward the same primitive completion condition: the finish line is **one fold, one record, one fact**, identical for all candidates. The quantity that is common is *structural* — a count of one record per commitment — and not the information content of the outcome: a low-amplitude outcome carries more surprisal ($-\log|c_i|^2$ is larger), but it does not cost more *to write one record*, and it is the cost of writing the record, not the surprisal, that Capacity Neutrality holds branch-common. A candidate may go on to seed a richer record, greater geometric elaboration, or larger downstream consequence, but none of that is part of the primitive commitment event; it is written *after* the line is crossed (§2 of *Measurement as Commitment*; §14 here). The contest therefore terminates at the same one-fold condition for every competitor — for any number of competitors, since one record per commitment is branch-common whether the contest is binary or n-way — and that is all Capacity Neutrality says: the structural finish line is candidate-independent. A candidate wins by being the completion selected — drawn on its readiness share — of one record, not by completing a *cheaper* or *larger* one; there is only one size of structural finish, and it is one fold. If the One Fold programme derives *one fold = one record = one fact*, this is not an assumption but a structural fact: there is no other size a primitive commitment could have. (The slogan is best stated with "record" rather than "bit": "bit" reads correctly only for a binary contest and invites the information-content reading the principle must exclude.)

One scope question is worth acknowledging even though its resolution lies outside this paper. If every commitment is one fold, one record (here), yet §6 writes $P(i) = |c_i|^2$ as though one contest decided one macroscopic measurement outcome, then a macroscopic outcome — itself the accumulation of very many primitive commitments — must arise from many micro-contests whose shares aggregate to the observed $|c_i|^2$. That aggregation (how micro-contest readiness-shares compose into the macroscopic Born weight, and why the composition preserves the exponent) is a genuine reconciliation the programme owes; it is adjacent to the decoherence-and-amplification story the contest inherits rather than something the present paper resolves. We flag it **(open)** and treat §6's single-contest writing as the coarse-grained description whose micro-foundation is left to that account.

6. The Winning Law and Its Continuum Shadow

The technical core has two levels, and keeping them straight is what makes the mechanism non-circular: a substrate-level *readiness-share* law, and its continuum-level *first-passage* image. We state the substrate law first, because it is the primitive; the first-passage Lemma is its shadow.

The substrate law (a posited selection rule). Among the candidates, the one selected to complete is drawn according to readiness share: the probability that candidate i is the completion is

$P(i) = R_i / \Sigma_j R_j$.

It is essential to be honest about the standing of this equation, because the reframing has changed it. In the rate picture, the analogous rate-share law was a *theorem* — exponential racing *forces* P = rate-share, via the integral below. At the substrate level there is no such forcing: $P \propto R_i$ is a *posited selection rule*, the statement that selection is drawn *linearly* in readiness share. This is one choice among many — selection could in principle go as R_i^2, as a softmax in R_i, or by any other share function — and nothing at the substrate level yet singles out the linear one. So the linear-share rule is a third marked input of the mechanism, alongside the §4 identification and §7's no-accumulation, and we list it as such. What recovers it as *more* than an arbitrary stipulation is not a substrate proof but two things: the continuum shadow below, in which exponential first-passage *does* force linearity (so the linear rule is what the emergent-time description delivers, the open §4.3 correspondence then carrying it back); and the surrounding ontology of §§8–14, which fixes that selection is among candidate facts, winner-take-all, forced-and-stochastic, one temporal increment per resolution. The linear-share rule is the minimal rule consistent with that ontology, not a theorem of it.

A candid statement of the residual worry belongs here, since the reviewer's question is the right one. With $R_i \propto |c_i|^2$ (the §4 identification) and $P \propto R_i$ (the linear-share rule), "selection ∝ relational support" amounts to "outcomes occur with probability $|c_i|^2$" — and one may fairly ask how far that is from simply *assuming* the Born frequencies it sets out to explain. The honest answer is that the selection rule *alone* is close to assuming what it explains; what does the explanatory work is not the rule but everything around it — that the objects selected among are candidate completions of one constraint (§2.1), that uniqueness comes from finite capacity rather than a collapse postulate (§8), that *that* a commitment occurs is forced while *which* is stochastic (§10), that each resolution is one temporal increment (§14), and that no-accumulation is the companion paper's Minimality met from the other side (§7). The contribution is an *ontology within which Born selection is the minimally-assumed rule*, not a derivation of Born from nothing. We claim the former and disclaim the latter.

This is a statement about *shares of relational support*, not about times: the candidate selected is drawn with probability $R_i / \Sigma_j R_j$, so over an ensemble the selection frequency of i equals i's readiness share. Selection is stochastic, not a maximisation: a lower-readiness candidate can be the one selected; it is simply selected less often, in proportion to its share. No clock appears. It is here that *cardinality* of readiness is essential — shares require magnitudes, which is why §4 insisted R_i be a quantity, not a bare ordering. Composing with the readiness law $R_i \propto |c_i|^2$,

$P(i) = |c_i|^2 / \Sigma_j |c_j|^2 = |c_i|^2$ (normalised),

which is **Born reproduction** — not the weighting (inherited) but the *frequencies*: the selection frequency of outcome i converges to $|c_i|^2$. This is the quantity *Measurement as Commitment* §5.3 marked **(open)** as the reproduction condition, here delivered **(conditional**, on the readiness law of §4 and the no-accumulation condition of §7**)**.

The continuum shadow. Once accumulated commitments constitute a local temporal order, the readiness-share law can be re-described as a race in emergent time, and in that description it coincides with a standard first-passage result — which we state and prove, then read as the continuum image of the substrate law above.

Lemma (winning law for memoryless first-passage) (proven). *In the emergent-time description, model each candidate's completion as an independent first-passage time T_i with constant hazard $\lambda_i \propto R_i$ (constant in accumulated progress, so that a candidate gains no advantage from past failed attempts — the emergent-time image of no-accumulation). Then the probability that candidate i completes first is its rate share, $P(i) = \lambda_i / \Sigma_j \lambda_j$, and the time to commitment is exponential, $T \sim Exp(\Sigma_j \lambda_j)$.*

The derivation is immediate. With $T_i \sim Exp(\lambda_i)$ independent,

$$P(T_i < T_j \text{ for all } j \neq i) = \int_0^\infty \lambda_i e^{-\lambda_i t} \prod_{j\neq i} e^{-\lambda_j t}\, dt = \int_0^\infty \lambda_i e^{-(\Sigma_j \lambda_j) t}\, dt = \lambda_i / \Sigma_j \lambda_j,$$

and the integrand summed over i gives the $Exp(\Sigma\lambda)$ density for T. Since $\lambda_i \propto R_i$, the rate share $\lambda_i / \Sigma_j \lambda_j$ equals the readiness share $R_i / \Sigma_j R_j$ — so the first-passage Lemma reproduces the substrate law exactly, as its continuum image. The Lemma is genuinely *proven*, but what it is proven *about* is the emergent-time description; that this description faithfully represents the atemporal substrate ordering is the correspondence §4.3 marks **(open)**. We rely on the Lemma for its content — the share law — while locating its standing precisely: a theorem about the continuum shadow, valid for the substrate primitive insofar as the shadow is faithful.

The conditionality is not a formality, and the honest weight of the section rests on naming it. The share law is exact *only* under no-accumulation. If readiness builds from a candidate's past near-misses — the substrate analogue of non-exponential, accumulating first-passage — then a candidate that has accumulated support is *ahead*, the selection depends on history, and P(i) is in general **not** the readiness share. In that regime the contest still selects a unique winner and still produces facts, but the frequencies need not be Born. So the mechanism does not reproduce quantum statistics automatically; it reproduces them precisely when readiness carries no accumulation. That condition is one of the mechanism's marked inputs — the others being the Born-consistency constraint on readiness (§§4–5) and the linear-share selection rule posited above — and §7 argues it is not arbitrary.

One directional point, since it is easy to invert: what is established is *no-accumulation ⟹ share-winning ⟹ Born* — no-accumulation is *sufficient*. The converse is not claimed; Born statistics do not by themselves entail no-accumulation, since special symmetries can give share-winning from accumulating processes too. The paper takes the sufficient direction only, and the case for no-accumulation is made in §7 on independent (Minimality) grounds, not retrofitted from the Born result it delivers.

---

7. No Accumulation as Dynamical Minimality

The two inputs the §6 result rests on are the readiness law (§4) and the no-accumulation condition (§6). The second is not, on inspection, a new assumption: it is the dynamical form of a principle the companion paper had already isolated on independent grounds. (We use "no-accumulation" for the substrate condition — readiness carries nothing from failed attempts — and "memorylessness" for its emergent-time image, the exponential first-passage of the §6 Lemma; they are the same condition at the two levels.)

Measurement as Commitment §5.5 introduced the **Commitment Minimality Principle (conjectural)**: a commitment event may depend only on information contained within the admissibility structure of the present pre-commitment state — no hidden selector, no accumulated external variable, participates in selection. There it was an abstract constraint on *what information selection may use*, argued by analogy to the single-source theorem and left as the load-bearing conjecture of the reproduction reduction.

In the contest picture that abstract constraint acquires a concrete reading. No-accumulation is exactly the condition that a candidate's readiness depends on the present configuration and *not* on its history of prior failed attempts. Accumulated-attempt history is information beyond the present admissibility structure. So:

Minimality forbids commitment from using information beyond the present configuration ⟹ commitment may not use accumulated-attempt history ⟹ candidate readiness is history-independent ⟹ readiness carries no accumulation (its emergent-time image is memoryless, exponential first-passage) ⟹ (with the readiness law) the winning law is the readiness share, and Born follows.

The second arrow carries a premise that must be stated explicitly, because it is load-bearing and not automatic. Minimality forbids the selection from using information *beyond the present configuration*; it forbids using accumulated-attempt history *only if that history is not itself part of the present configuration*. If failed attempts left a residue in the substrate state, a selection reading that residue would be fully Minimality-compliant — it reads only the present state — yet accumulation-dependent, and the chain would break at the second arrow. The implication therefore depends on a substantive claim about what the substrate records: **uncommitted attempts leave no configurational trace** — a failed tick is not written into the state, so a candidate that has "tried and failed" repeatedly is, given the same present configuration, in exactly the position of one just entering the contest. This is the dynamical analogue of the committed/uncommitted asymmetry the programme insists on elsewhere, and it is argued, not assumed, in §7.1, where it is also what protects the whole picture from the "no memory" misreading. We flag it here as the premise on which Minimality ⟹ no-accumulation actually turns, and carry the argument for it in §7.1.

We mark the identification of no-accumulation with dynamical Minimality **(conjectural)** — it is an identification, not a derivation, and the same care the companion paper applied to Minimality applies to its dynamical face. But the identification has real consequences in both directions. It tells the abstract principle of §5.5 what it *means* dynamically: Minimality is not merely "no hidden variables" in the abstract, but specifically "no accumulation of failed-attempt history" in the contest. And it tells the present mechanism that its no-accumulation assumption is not free-floating: it is the same conjecture the reproduction reduction already rested on, now wearing dynamical clothing. The two papers, pursuing the criterion (§5.2) and the reproduction (§5.5) as separate open inputs, were circling one object — and this is where that object shows its single face. What was, across the two papers, "an open commitment criterion" and "an open reproduction principle" is, under this identification, **one** conjecture: that unresolved commitment attempts carry no accumulation. Establish it, and both halves of §1's question close together; refute it, and the §6 result localises exactly where the failure enters — an accumulation term in the selection, a candidate carrying history.

This does not prove Minimality, and the paper does not claim to. It claims something narrower and, within the programme, more useful: that the criterion gap and the reproduction gap are not two debts but one, and that the one has a sharp, falsifiable dynamical statement — *unresolved commitment attempts carry no memory*.

7.1 What No-Accumulation Does Not Mean

A clarification is required, because this condition — *no-accumulation* at the substrate, *memorylessness* in its emergent-time image (§7) — carries the single most misreadable name in this paper, and the natural misreading collides head-on with the rest of the programme. The slogan "unresolved commitment attempts carry no memory" must not be heard as "reality carries no memory." The programme is built on the opposite: persistent records, committed history, commitment density, fact-momentum, the κ-field, the arrow of temporal order, and the conjectured phase-as-memory structure all depend on commitment leaving a residue that endures. The no-accumulation condition of §7 says nothing against any of that. It is far narrower, and the narrowness is the whole point.

The condition applies only to the *selection governing an unresolved contest* — to candidates that have not yet become facts. It does not apply to the *committed record* a contest produces — to a fact that already exists. These are different objects on opposite sides of the commitment boundary, and the condition speaks only to the first:

No-accumulation: P(candidate i commits in the next increment | unresolved so far) depends only on the present configuration, not on i's history of prior *failed, uncommitted* attempts within the current contest.

One distinction does the protective work, and it is the Markov distinction stated for this setting. "Depends only on the present configuration" does *not* mean the contest is blind to prior commitments. Prior committed facts are precisely what set the current admissible candidates and their amplitudes; the present configuration *encodes the entire committed past*, and the contest reads that past in full, because the past is in the state. What no-accumulation forbids is dependence on a *different* kind of history — the candidate's accumulated record of *unsuccessful* attempts within the still-unresolved contest. Past *commitments* are in the present state and are used; past *non-commitments* are not recorded and confer no advantage. A candidate that has "tried and failed" many times in the current contest is in exactly the same position as one just entering it, given the same present configuration. That is all no-accumulation asserts.

This is more than a guard against misreading; it is the premise on which the §7 derivation actually turns, and it should be read as load-bearing rather than merely clarifying. The chain Minimality ⟹ no-accumulation breaks at its second step unless uncommitted attempts leave *no configurational trace* — for if a failed tick were written into the substrate state, a selection reading that trace would respect Minimality (it reads only the present state) while still carrying memory. So the claim "past non-commitments are not recorded" is exactly what makes Minimality deliver no-accumulation, and it is a substantive claim about what the substrate stores, of the kind the programme argues rather than assumes.

The argument has two steps, and honesty requires separating them, because the first is suggestive and only the second closes it. *First*, the committed/uncommitted asymmetry: a commitment is, by construction, the irreversible writing of a record, and an *un*committed attempt is one that has *not* written anything. But this much shows only that one cannot record a failed attempt *without making some commitment* — for the record "candidate i tried and failed at step k" would itself be a fact, a *different* commitment from "candidate i won." The bare asymmetry therefore establishes that recording an attempt costs a commitment; it does not yet establish that the substrate declines to pay. *Second*, the closing step, which supplies what the first lacks: there is no commitment available to spend on facts-about-attempts during an unresolved contest. Commitment capacity is finite **(inherited)**, and a record-about-an-attempt would consume the same one-record unit any commitment costs (§5); but while the contest is unresolved that capacity has not yet been spent on *anything* — winner-take-all reserves it for the resolution — so there is no spare capacity to write side-records of the attempts. And once the contest resolves, there are no pending attempts left to record: the losers simply did not complete. Facts-about-attempts are thus committable neither during the contest (no spare capacity) nor after it (nothing pending). It is this — finite capacity with no allocation for attempt-records — not the bare asymmetry alone, that grounds the no-trace premise and hence the §7 chain. The winner-take-all this step invokes is the *structural* one — capacity is finite and is consumed by whichever resolution occurs, which follows from finite capacity and irreversibility alone — not the *frequency-correct* selection of §6, which depends on no-accumulation downstream; so §7.1 leans only on the structural fact and not on the result it is helping to ground, and the apparent circle does not close. We mark the combined argument **(conjectural)**, as the section's status requires: it draws only on the inherited finite-capacity premise rather than on the §8 uniqueness result it resembles, so it is not circular, but it does rest on the claim that the substrate has no attempt-record allocation, which we state as a premise rather than derive from below.

So the residue of a *successful* commitment is untouched by the condition. When a contest resolves, the selected candidate writes a persistent record, and that record contributes to committed history, temporal ordering, commitment density, fact-momentum, and — if the companion programme is right — phase-memory. The selection carries no accumulation; the fact it produces is not erased. The architecture is clean precisely because the two live on opposite sides of the boundary:

before commitment → no-accumulation selection among unresolved candidates; the commitment event → one candidate is selected, one record forms; after commitment → persistent memory, accumulated and transported.

The phase-as-memory conjecture is therefore not in tension with no-accumulation but complementary to it, as §14 already noted from the other direction: the contest determines *which* commitment occurs (pre-commitment, no-accumulation selection), and the phase-memory structure, if it holds, records *that* it occurred and carries it forward (post-commitment, persistent). One concerns fact *selection*, the other fact *persistence*. They are not the same mechanism and must not be conflated — and once they are kept apart, "unresolved attempts carry no memory" and "committed facts are the memory of the universe" are not merely compatible but two halves of one picture.

---

8. Outcome Uniqueness from Finite Capacity

A standing puzzle of measurement is the appearance of a single outcome where the formalism offered many. The contest framework answers it from a resource already in the inherited structure rather than from a new postulate.

Commitment capacity is finite **(inherited**, Commitment Capacity programme**)**. There is enough capacity to complete one irreversible record, not several at once. The candidate selected by readiness consumes that capacity in completing one record; the others, not selected, find the capacity spent and can no longer commit. Uniqueness is therefore **winner-take-all** and follows from readiness-selection plus finite capacity:

many candidates admissible → one is the readiness-selected completion → it consumes the capacity → one fact, others excluded.

Two features are worth marking. First, it requires no additional collapse rule: uniqueness is a consequence of the contest and the capacity bound, both already present, not an imposed selection **(conditional**, on the contest interpretation and the capacity bound**)**. Second, it explains not only *that* one outcome appears but *why the others leave no committed trace* — they did not complete a record, and a non-completed candidate is not a fact. The unselected candidates are not "elsewhere" in the manner of unrealised worlds; they are candidate completions that were not selected, and the framework's ontology has no category of committed-but-losing fact for them to occupy.

---

9. The Master Reduction of the Commitment Problem

The results of §§4–8 can be collected into a single statement, because their value is cumulative: each is a step in reducing the measurement problem to one place. We state the reduction as a theorem with its hypotheses marked, so that what is assumed, what is proven, and what remains open are visible at once.

Theorem (Master Reduction) (conditional, on the hypotheses as marked**)**

Assume:

1. **(no-accumulation, conjectural)** Relational readiness carries no contribution from a candidate's history of unresolved attempts — equivalently (§7), commitment uses no information beyond the present admissibility structure (the Commitment Minimality Principle).
2. **(Capacity Neutrality, conjectural; inherited if One Fold derives one-fold = one-record = one-fact)** Every primitive commitment consumes the same branch-common unit of structural capacity — one record written, not the outcome's information content — outcome-specific elaboration being downstream of the commitment event.
3. **(readiness law — identification open here; exponent inherited-conditional, contingent on the identification)** Relational readiness is set by a candidate's own relational-support measure, which is the inherited quadratic weight: $R_i \propto |c_i|^2$. The *identification* (that readiness is that measure) is the paper-local open piece; the *quadratic exponent* is inherited from Double Square / Physical Necessity as derived-under-admissibility-conditions — but only *given* the identification, since the exponent transfers to readiness exactly insofar as readiness is the magnitude Double Square measured (§4.2), not re-derived here.
4. **(linear-share selection rule, posited)** Selection is drawn *linearly* in readiness share, $P(i) = R_i / \Sigma_j R_j$, rather than as R_i^2 or any other share function. A theorem in the continuum shadow (exponential first-passage forces it), a posit at the substrate level (§6), carried back by the open §4.3 correspondence.

Then:

1. **Outcome uniqueness** follows from finite-capacity winner-take-all (§8) — one candidate is selected, the rest excluded, with no collapse postulate.
2. **Outcome frequencies reproduce the Born weighting**, $P(i) = |c_i|^2$, by readiness-share selection (§6).
3. **Commitment advances temporal order**, one resolved contest at a time (§14), with no clock presupposed (§4.3).

4. **Persistent records arise** from completing commitments (§14), the substrate of memory and history.

Crucially, hypotheses 2 and 3 are *not independent*: §5 established that Capacity Neutrality and the readiness law are the two faces of a *single* Born-consistency constraint on total readiness, $R_i \propto |c_i|^2$. So the four hypotheses collapse toward three genuinely distinct objects — no-accumulation, the single readiness constraint, and the linear-share selection rule — with hypothesis 1 reducing further (§7) to the companion paper's Minimality, met from the other side. The honest count of *paper-local* open inputs is therefore: one conjecture (no-accumulation = Minimality), one identification (that readiness is the relational-support measure, relieved on its capacity face by a derived One Fold identity), and one posited selection rule (linear share, a theorem in the shadow but posited at the substrate). The quadratic exponent the constraint deploys is not a paper-local input at all — it is inherited from the reconstruction programme, conditional as that programme leaves it (§4).

The consequence for the measurement problem. The theorem reduces measurement to its hypotheses. With no-accumulation supplied by Minimality and Capacity Neutrality potentially inherited from One Fold, the readiness law,

$R_i \propto |c_i|^2$,

is the *single remaining target* — but it must be read at its true status, which §§4–4.2 fixed. It has two components, and only one is open here. The *identification* — that a candidate's readiness is set by its own relational-support measure — is the paper-local open piece. The *exponent* — that the relational-support measure is quadratic — is inherited from the probability-reconstruction programme as derived-under-admissibility-conditions, conditional at programme level on those assumptions, and *not re-derived in this paper*. So the measurement problem, on this account, is no longer "why does one outcome occur, and with what odds?" but the sharper pair: *is readiness the relational-support measure* (open here), and *why is that measure quadratic* (inherited, the quantum-probability problem, conditional as Double Square leaves it). The reduction is real and the inherited target is the deep one — but the paper claims a reduction, not a closure, and does not represent the exponent as a debt it owes or discharges.

Naming the hard part. We do not present the exponent in $R_i \propto |c_i|^2$ as a small thing. It is not. *Why the weight is quadratic — why $|c_i|^2$ and not $|c_i|$^α for some other α — is the quantum-probability problem*, the question on which every reconstruction in this area stands or falls. The contribution of this paper is not to answer it but to show that everything else in measurement reduces to it: given the quadratic readiness law, uniqueness, Born frequencies, temporal advance, and record formation all follow. And the probability programme (Double Square, Physical Necessity) has independently narrowed the escape routes around that exponent — showing pairwise selection forces bilinearity and that non-quadratic rules fail admissibility. The measurement problem is thereby reduced to a target the surrounding programme was already besieging from the other side. That reduction, not a closure, is the claim.

---

10. Forced and Stochastic: the Two Halves of the Criterion

The companion paper characterised the commitment criterion as a **closure-saturation** condition and described commitment as *forced, not stochastic* — what closure does when admissibility can no longer be maintained. The present paper makes commitment a *race*, which is manifestly stochastic in outcome. The tension is apparent and must be resolved directly, because two companion papers may not describe the same event as both forced and stochastic without saying how.

The resolution is that "forced" and "stochastic" attach to different questions, exactly the two §1 distinguished. That *a* commitment occurs is forced: closure saturation and finite capacity together guarantee that the open admissibility cannot be sustained indefinitely, so some candidate *will* be selected and consume the capacity — there is no admissible history in which the contest simply never resolves. *Which* candidate wins is stochastic: it is the readiness-selection outcome of §6, governed by the readiness law and the share law. The companion paper's "forced, not stochastic" was a claim about *whether and when* commitment happens — and on that the present paper agrees, supplying saturation-plus-capacity as the forcing. The stochasticity the present paper adds is located entirely in the *selection of the winner*, which the companion paper had marked separately as the reproduction question.

So the two papers partition the commitment dynamics cleanly: closure saturation supplies the *that/when* (forced), and the readiness contest supplies the *which/how-often* (stochastic). This is precisely the partition *Measurement as Commitment* anticipated when it wrote that its §5.2 and §5.3 were "facets of one underlying object — the commitment dynamics — examined once for *when* it fires and once for *with what frequencies*." The present paper is that object, viewed from both facets at once: forced to fire, stochastic in what fires.

---

11. Relativistic Status: "First" Without a Clock

A contest has a winner, and the language of this paper leans on ordering throughout — "completes first," "the ordered sequence of resolutions." A reader trained on relativity will hear a *preferred temporal ordering*, which Lorentz invariance forbids at the fundamental level. The objection must be met, because left unmet it reads as a contradiction the paper failed to notice rather than an open problem it has bounded. It is the second — and the readiness reframing answers it more cleanly than a rate-race ever could.

Why "first" does not posit a preferred frame — and why readiness makes this sharp. A referee's reflex is: *first according to which clock?* On a rate-race that question bites, because rates presuppose a clock. But the substrate primitive of this paper is *not* a rate and *not* a "first in time"; it is relational readiness (§4.3) — the cardinal magnitudes of pairwise relational support standing behind the candidates, *how strongly each is supported*. These are facts about the constraint structure that require no clock to be true, and selection is drawn from their *shares* (§6), not from a temporal order of completion. The primitive content is not "candidate i completes before candidate j" but the support magnitudes $\{R_j\}$ themselves, from which one candidate is selected with probability $R_i / \Sigma_j R_j$ — an atemporal, relational, and stochastic statement, with no "before" in it. Temporal order is recovered only *after* readiness contests resolve and accumulate (§14); "first" is what the emergent-time re-description of that selection looks like once a clock exists. So the answer to "first according to which clock?" is: *no clock* — the primitive is relational support, and the clock is downstream of it. The framework had already rejected background time (§1); the readiness reframing is what lets the contest honour that rejection at its own foundation rather than only in its conclusions.

What genuinely remains open, stated sharply. Dissolving the naive objection is not the whole job, and the residual problem is more pointed than "reformulate covariantly." It arises from a tension between two things the paper already needs:

- *Bell-compatibility wants globality.* The standard reproduction of Bell correlations in this picture assigns readiness to *global* candidates in configuration space, not to local factors — readiness $R_i \propto |c_i|^2$ attaches to the global amplitude (§2.1), so correlations come out right without a local hidden ordering. This is the same configuration-space globality every amplitude-respecting account uses, and is not in itself a relativity problem; quantum correlations are nonlocal in exactly this sense.
- *Relativity wants no preferred global slicing.* But a single readiness resolution over globally-assigned candidates is a *global* commitment event, and globality of the event sits uneasily with the absence of a preferred slicing unless the emergent-order story does real work to make that globality frame-independent.

So the open problem is not the naive one (which the readiness-ordering point dissolves) and not the vague one ("RQFT dislikes orderings"). It is the specific reconstruction task: **show that the globally-assigned readiness contest — whose Bell-compatibility requires configuration-space globality — nonetheless reconstructs local Lorentz causal structure in the continuum limit, with the emergent commitment-order aggregating into a Lorentz-invariant causal order rather than a preferred foliation.** That is a reconstruction, not a repair: there is no contradiction to remove, only a continuum-limit result to establish — that many discrete, globally-resolved commitments coarse-grain to the causal structure relativity observes. We mark it **(open)**, as a programme-level reconstruction task rather than a defect internal to the mechanism. The non-relativistic content — Born reproduction, uniqueness, the forced/stochastic partition — stands independently of how this reconstruction resolves, exactly as the base measurement proposal stands independently of the Gate-3 residue.

---

12. What Would Falsify the Commitment Contest

A foundational mechanism earns its standing partly by saying what would defeat it. The commitment contest is conjectural and conditional, but it is not unfalsifiable — and the falsifiers are sharp. The picture fails if any of the following is established.

1. **Pre-commitment residue.** If unresolved commitment attempts are shown to leave a measurable trace in the substrate state — so that a candidate's failed attempts confer advantage — then readiness is *not* history-free, the no-accumulation condition fails, and the readiness-share law (and with it Born reproduction) is lost (§6, §7.1).
2. **Branch-dependent capacity cost.** If a primitive commitment is shown to consume an outcome-dependent amount of capacity — a candidate that must write more than the primitive distinction to win — then Capacity Neutrality fails, the capacity factor no longer cancels, and the frequencies tilt away from Born (§5).
3. **A non-quadratic readiness law.** If a substrate-level derivation forces relational readiness $R_i \propto |c_i|^\alpha$ with $\alpha \neq 2$, the mechanism contradicts the inherited Born weighting and fails. (Equivalently: if the probability-reconstruction programme's quadratic result were overturned, the target the contest preserves would move.) This is the deepest and most consequential falsifier, because $\alpha = 2$ is the quantum-probability question itself (§9).
4. **Failure of covariant reconstruction.** If the globally-resolved readiness order provably *cannot* reconstruct Lorentz causal structure in the continuum limit — if it necessarily induces a preferred foliation — then the covariant *extension* fails. This falsifies the programme-level reconstruction (§11), not the non-relativistic mechanism, which stands regardless.

Two points about the character of this list. First, these are *internal* falsifiers — they kill the mechanism by violating its own stated conditions, not by clashing with a current experiment. Because the contest imports Born and modifies no evolution law, no measurement statistic on the books can refute it; its only experimental contact is the conditional Gate-3 residue it shares with the companion programme. The mechanism is observationally safe and internally exposed — vulnerabilities located in its own structure, which is the right shape for a foundational proposal. Second, the falsifiers are *ordered* by depth: (1) and (2) are conditions the substrate might violate; (3) is the one that, if it broke, would take the whole probability programme with it; (4) bounds only the relativistic extension. Naming them at their true sizes is part of the honesty the paper aims at.

---

13. Quantum Computation as Race Engineering

The contest interpretation gives a clean reading of what a quantum algorithm does, offered here as an *illustration* of the picture's reach rather than as a load-bearing result **(conjectural, illustrative)**.

Quantum algorithms operate entirely before commitment, on the reversible possibility dynamics. In the contest reading their function is to reshape the relational-support landscape that sets the candidates' readiness: constructive interference raises the readiness R_i of target candidates, destructive interference suppresses competitors. The algorithm does not select the winner — selection is the readiness contest, resolved at measurement — but it engineers the contest conditions so that the desired candidate enters with an overwhelming readiness share, making it the near-certain completion.

Read this way, the familiar slogan that a quantum computer "explores all paths" is sharpened: it does not commit all paths, which finite capacity forbids; it *re-weights the relational support among them* so that committing the wanted path becomes likely. Measurement then resolves the engineered contest. Quantum computation is, in this picture, commitment-readiness engineering — a reshaping of relational support upstream of a selection it does not itself perform. The reading is consistent with the rest of the framework and pleasingly concrete, but nothing later depends on it.

---

14. Persistence and the Advance of Temporal Order

A winning candidate is selected only by completing a *stable* record; a completion that does not persist has not produced a fact. Commitment and persistence are therefore inseparable: the winning condition *is* successful persistent record formation, and "fact created" and "record persists" are two descriptions of one event rather than a cause and its effect (**conditional**, on the contest interpretation**)**.

This connects to the programme's account of temporal order, and it is where the readiness primitive pays off most directly. Prior work identified the advance of physical time with the accumulation of committed fact **(inherited)**; the contest supplies the increment. Each resolved readiness contest contributes exactly one new irreversible record, and the ordered succession of resolutions *is* the succession of temporal increments — not events occurring within time, but the constitution of temporal order itself, since each resolution is what a temporal increment consists in. This is exactly why the primitive had to be a readiness ordering and not a rate (§4.3): a rate would have presupposed the temporal order this section derives. The directedness of temporal order then has a mechanical source: a resolution consumes capacity and writes an irreversible record, neither undoable, so the sequence admits no reversal. The arrow of temporal order is the arrow of accumulated commitment resolutions (**conjectural**, as an identification; **inherited** for the time–commitment link it rests on**)**.

We resist overstating this. The contest supplies a *mechanism* for the increment whose existence was inherited; it does not independently derive that temporal order must accumulate, which remains the prior result. What is new is the granularity and the non-circularity: temporal order advances one resolved readiness contest at a time, and because the contest is decided by atemporal relational support, the account derives temporal order without presupposing it.

This paper is deliberately agnostic about the *representation* of committed history after a contest resolves — about how a fact, once made, is carried forward. Companion work explores the possibility that persistent phase structure is the transport representation of accumulated commitment history, so that phase is the continuum shadow of the surviving closure history rather than the substrate object itself. If that direction holds, the division of labour is clean: the contest mechanism of this paper determines *which* commitment occurs, while the phase-memory structure records *that it persists*. The two questions — what causes a fact, and how a fact is remembered — are complementary but logically distinct, and the present paper answers only the first. The second is left to that companion line, and nothing here depends on its outcome.

The separation deserves emphasis, because the two are easily heard as rival explanations of measurement when they are in fact successive stages of one process. The probability papers repeatedly tie phase to holonomy and to persistent distinguishability structure — that is, to what *survives* and is *transported* once facts exist. The contest, by contrast, operates entirely *before* a fact exists, selecting which admissible record becomes factual. So if phase is ultimately identified with the transport-memory of committed history, phase does *not* participate in selecting the winner; it records the surviving commitment history *after* the contest has resolved. The order is strict:

contest → commitment → persistence → phase memory.

Selection and memory are therefore not two competing accounts of what measurement is, but two consecutive stages of one process — the first choosing the fact, the second carrying it forward. A reader who took phase-as-memory to be a rival to the contest, or the contest to be a rival to phase, would be conflating stages that the architecture keeps in sequence: the contest ends exactly where memory begins, at the commitment boundary.

---

15. Relation to Existing VERSF Results

The paper replaces none of the existing commitment theory; it supplies the dynamical interpretation that connects several previously separate results into one process, and its relation to each should be stated at the right size.

The **Fold Interface Law** defines the commitment boundary separating reversible possibility dynamics from committed reality; the contest takes place at and across this boundary, and its derivation of the minimal commitment event as a single fold — one record — is what grounds the Capacity Neutrality Principle of §5 — and, if it derives one-fold = one-record = one-fact, inherits Neutrality outright **(inherited)**. **Fact Production** supplies the necessary conditions — distinguishability, amplification, capacity — which the contest reading recasts as, respectively, what makes candidates distinguishable competitors, what a completing commitment initiates, and what makes the contest winner-take-all **(inherited)**. **Commitment Capacity** supplies the finite resource that §8 turns into uniqueness and §5 shows to be branch-neutral in its bearing on outcome selection **(inherited)**. **Tick-Bit** supplies the substrate attempts — ticks as attempts, Bits as completed records — whose relational structure §4 reads as readiness and §6 builds the share law upon **(inherited)**. **Double Square** and **Physical Necessity** fix the quadratic weighting and the pairwise-relational ontology that readiness is built from (§§4.1–4.2) **(inherited)**. **Measurement as Commitment** identifies measurement with commitment and, in §§5.2 and 5.5, posed the criterion and reproduction as its two open inputs; the present paper is a candidate for both, unified in §7 **(this paper)**. The temporal-emergence results identify accumulated commitment with the advance of temporal order, to which §14 supplies the increment — non-circularly, since readiness is atemporal **(inherited)**.

The contribution is the unification, and it is a conditional one: distinguishability, capacity, thresholds, pairwise support, and the attempt mechanism are recast as the ingredients of a single readiness contest whose resolution is fact production, and the contest's reproduction of Born statistics is shown to rest on the same conjecture (no-accumulation = Minimality) the companion paper's reproduction reduction already required. The paper does not add a primitive; it proposes that the existing primitives describe a readiness contest, and follows the consequences — taking care that the primitive is a relational, atemporal readiness rather than a rate, so the account does not presuppose the temporal order it derives.

---

16. What This Paper Does and Does Not Claim

The paper does **not** prove that commitment occurs, nor derive the necessary conditions, nor the quadratic weighting; these are inherited **(inherited)**. It does **not** derive the readiness law $R_i \propto |c_i|^2$; §4 sizes that residue honestly as three-layered — single-candidate reduction, functional form, and the exponent — of which only the first two are paper-local (the identification that readiness is the relational-support measure), the exponent being inherited-conditional from the reconstruction programme and contingent on that identification, not a single number and not a paper-local debt (below). It does **not** prove no-accumulation; that readiness carries nothing from failed attempts is marked **(conjectural)** — shown in §7 to be the dynamical form of the companion paper's Commitment Minimality Principle, hence not an independent debt but the same one, and resting on the no-configurational-trace premise argued in §7.1. It does **not** prove Capacity Neutrality unconditionally; §5 shows it is not a separate debt but the capacity-face of the same single Born-consistency constraint on readiness of which the readiness law is the other face — and that it becomes **(inherited)** if the One Fold programme derives the one-fold = one-record = one-fact identity, **(conjectural)** if that identity is posited rather than derived. It does **not** supply a covariant formulation; the contest generates rather than presupposes temporal order

(§§4.3, 11), so it posits no clock, but the reconstruction of local Lorentz causal structure from the globally-resolved readiness order is marked **(open)** as a programme-level task (§§11–12). It modifies no evolution law and claims no empirical signature beyond those the companion programme already locates in the conditional residue.

The paper **does** propose, **(conjectural)**, a definite dynamical reading of the commitment criterion: commitment is the readiness-share resolution of a contest among candidate completions of one constraint structure, with thresholds as finish lines and finite capacity as the winner-take-all resource — and with relational readiness, not a rate, as the atemporal primitive (§4.3). It **does** prove, **(proven)**, the first-passage winning law $P(i) = \lambda_i / \Sigma_j \lambda_j$ in the continuum shadow, and — *granting the posited linear-share selection rule* at the substrate level (§6), not as a substrate theorem — show, **(conditional)** on the named inputs, that readiness-share selection reproduces the Born frequencies, the overall scale cancelling. The proven content is the shadow-lemma; the substrate share law it images is a posit, carried back by the open §4.3 correspondence, and the paper banks this rather than letting the proven Lemma stand in for a substrate proof. It **does** show, **(conditional** on Capacity Neutrality**)**, that the capacity cost of commitment cancels from selection, so capacity governs *whether and how readily* fact production proceeds while readiness governs *which* candidate wins — clarifying that Capacity Neutrality and the §4 readiness law are two coupled faces of a single Born-consistency constraint $R_i \propto |c_i|^2$ — coupled by the product $R_i = R^\circ_i \cdot F$, though not equally supported, since Capacity Neutrality has a Born-free operational argument the readiness law lacks. It **does** supply outcome uniqueness from finite capacity rather than a collapse postulate **(conditional)**. It **does** reconcile with the companion paper's "forced, not stochastic" criterion by partitioning the dynamics into a forced *that/when* (saturation) and a stochastic *which* (the contest). It **does**, in §7, identify the criterion gap and the reproduction gap as one conjecture — *unresolved attempts carry no memory*. It **does**, in §9, collect these into a Master Reduction showing the measurement problem reduces to the single target $R_i \propto |c_i|^2$. And it **does**, in §12, name explicitly what would falsify the proposal.

The honest summary is that the paper converts the open commitment criterion into a concrete readiness-contest mechanism whose Born reproduction is exact under named inputs, shows that the no-accumulation input is the dynamical face of a principle the programme had already isolated, and recasts the whole as a contest decided by an atemporal relational readiness rather than a rate — so that temporal order is genuinely derived, not presupposed. Its residue is, honestly, one paper-local open item, one inherited-conditional target, and two programme-level matters. The paper-local open item is the **identification** (§4): that a candidate's readiness is set by its own relational-support measure — the single-candidate and functional-form layers — together with **no-accumulation (conjectural)**, the companion paper's Minimality met from the other side, turning on the no-configurational-trace premise of §7.1. The inherited-conditional target is the **quadratic exponent** of $R_i \propto |c_i|^2$: the quantum-probability problem itself, but *not re-derived here* — inherited from Double Square and Physical Necessity as derived-under-admissibility-conditions, and inherited with that conditional status, so it would reopen only if those assumptions failed. The two programme-level matters are a hoped-for *substrate* re-derivation of the exponent that would not import Born (an aspiration of the wider programme, not a debt this paper undertakes) and the **covariant-reconstruction task (open)** of §§11–12. There is also a third paper-local input not to be elided: the **linear-share selection rule (posited)** of §6 — that selection is drawn linearly in readiness share rather than as some other share function — a theorem in the continuum shadow but a posit at the substrate, and the place where, taken alone, the rule comes closest to assuming the frequencies it yields; what carries it beyond stipulation is the surrounding ontology (§§8–14), not the rule itself. On the single Born-consistency constraint $R_i \propto |c_i|^2$, the readiness-face and capacity-face are coupled by the product $R_i = R°_i \cdot F$ — given Born, fixing either fixes the other — but not equally supported: the capacity-face has a Born-free operational ground (relieved entirely if One Fold derives its identity), while the readiness-face carries the identification and inherits the exponent. The paper neither undersells the identification as "one number" nor oversells the capacity-face as "forced," nor flattens the two faces into equal standing, nor claims the exponent as a debt it either owes or discharges: it inherits the exponent, conditional as the reconstruction programme leaves it, and names it for what it is — the question of why quantum probability is quadratic.

The *character* of these open items is what §12 makes precise: none is a prediction that could clash with experiment, because the paper imports rather than modifies the Born weighting and adds no term to the evolution law. They are debts the mechanism owes, internal to its own structure — refuted by a found accumulation term, a found branch-specific cost, a forced non-quadratic exponent, or a failure of the continuum reconstruction, not by any measurement statistic on the books. The proposal is, in that precise sense, observationally safe and internally exposed: the right shape for an interpretive mechanism.

---

17. Conclusion

The commitment programme had reduced measurement to commitment and fixed the probability weighting, but it could not say what physical condition produces a fact or why the outcomes carry the Born frequencies. It named those two gaps and located them in one object — the commitment dynamics — without supplying the object. This paper proposes the object: commitment is a contest among candidate completions of one constraint structure, decided by *relational readiness* — a cardinal measure of pairwise relational support — in which established thresholds serve as the finish line, finite capacity makes the contest winner-take-all, and the committed fact is the candidate the readiness ordering selects.

One conceptual decision shapes the whole. A rate presupposes a clock, and the programme says commitment makes the clock; so the substrate primitive cannot be a rate. It is the atemporal relational-support magnitudes, selection drawn from their shares, and the familiar first-passage race is the continuum shadow of that selection, valid once temporal order has emerged. This is what lets the paper derive temporal order rather than assume it, and it is what answers the relativist's "first according to which clock?" with: no clock — relational support, with the clock downstream.

The picture earns its keep at one technical point and one structural one. Technically, the readiness-share law $P(i) = R_i / \Sigma_j R_j$ — posited at the substrate as the linear-share selection rule, proven in the continuum shadow — composes with the readiness law $R_i \propto |c_i|^2$ to give the Born frequencies exactly, the overall scale cancelling. Structurally, the no-accumulation this requires is the dynamical form of the companion paper's Commitment Minimality Principle, so the criterion gap and the reproduction gap collapse to one conjecture — *unresolved commitment attempts carry no memory* — whose failure the construction localises to a single place. Capacity does not threaten either: the winner depends only on readiness ratios, so a branch-common capacity cost cancels from selection while still governing whether and how readily the contest resolves. That capacity-face and the readiness law are coupled — two faces of one Born-consistency constraint $R_i \propto |c_i|^2$ — but not equally exposed: the capacity-face has a Born-free operational ground (and is relieved entirely if One Fold derives one-fold = one-record = one-fact), while the readiness-face carries the identification of readiness with relational support and inherits the quadratic exponent from Double Square — derived there under admissibility assumptions argued physically necessary, inherited here with that conditional status rather than re-derived. A substrate re-derivation that would re-ground the exponent without importing Born is a horizon for the wider programme, not a debt this paper owes.

What stands is conditional and says so. Outcome uniqueness follows from finite capacity without a collapse postulate; the "forced, not stochastic" criterion of the companion paper is recovered as the forced *that*-it-occurs against which the stochastic *which*-candidate-wins plays out; record formation, persistence, and the advance of temporal order become facets of one contest-resolution process. The §9 Master Reduction makes the shape explicit, and §16 sizes the residue: given no-accumulation and Capacity Neutrality, the entire measurement problem reduces to the single target $R_i \propto |c_i|^2$ — and that target is not a technical residue but the quantum-probability problem itself, the question of why the weight is quadratic. The paper does not answer it; it shows everything else reduces to it, and that the probability-reconstruction programme has already narrowed the escape routes around it. No evolution law is modified, and the proposal is observationally equivalent to standard quantum mechanics wherever that single constraint and no-accumulation hold.

One guard against the most natural misreading, stated once more because the term invites it. That unresolved attempts carry no memory refers *only* to candidate selection before a fact exists — to readiness that gains nothing from a candidate's history of failed attempts (§7.1). It does *not* imply that successful commitments leave no residue. The opposite is the programme's foundation: committed facts persist, accumulate, and are carried forward as records, temporal order, commitment density, fact-momentum, and conjecturally phase-memory. The contest is memoryless; the universe of committed facts is precisely the memory. The two live on opposite sides of the commitment boundary, and the architecture is clean only when they are kept there.

A threshold does not make a fact.

A fact is the candidate the readiness ordering selects — and the unresolved attempts it passes over leave no trace, while the one it selects is written into reality forever.