Finite Commitment Mathematics (FCM): Value Demonstration Package

---

Table of Contents

General Reader Summary

Imagine you can only keep track of a limited number of distinctions at once. You might be able to remember that Alice is different from Bob, and Bob is different from Carol—but if you try to track too many such distinctions simultaneously, something has to give. In everyday life, we handle this by lumping things together: "those people from the conference" rather than remembering each individual.

Finite Commitment Mathematics (FCM) asks: what if this isn't just a practical limitation, but a fundamental feature of how identity works? More precisely: **FCM explores what mathematics looks like if identity is treated as a finite resource.** What if there's a maximum number of distinctions any system can sustain, and when that limit is exceeded, things that were previously distinct *become the same thing*—not approximately, but actually?

This paper develops the mathematics of such a world. We define precisely what it means for a system to exceed its "distinction budget," and we show that there's a principled way to restore balance: a collapse operation that merges identities until the system is back within its limits. This collapse is irreversible—once two things become one, they can't be separated again.

The surprising consequences include:

- **Processes that oscillate forever can still have well-defined outcomes** (because after collapse, the oscillation becomes invisible)
- **You can't organise these structures into standard mathematical categories** (because the very objects you're working with can change identity mid-operation)
- **Different priorities lead to different collapses** (if you care more about preserving relationships than preserving identities, you'll merge differently)

None of this requires infinity, contradiction, or exotic physics. It's a self-contained mathematical framework that may illuminate situations—from quantum mechanics to cognition—where identity is not as stable as we usually assume.

Technical Abstract

Finite Commitment Mathematics (FCM) is an axiomatic framework in which identity is neither primitive nor stable, but emerges from—and is regulated by—finite distinguishability constraints. When commitment structures exceed their capacity bounds, identity itself changes through a canonical collapse projection. This document presents FCM's axiomatic foundations, defines the collapse projection operator, proves its core properties, and establishes that FCM's structures and dynamics are not representable within category theory, domain theory, quotient constructions, or rewriting systems without violating their axioms. This non-embeddability—not claimed importance or applicability—is the precise sense in which FCM constitutes new mathematics. We demonstrate FCM's mathematical usefulness: it renders overloaded constraint

systems well-posed, provides semantics for update sequences despite categorical failure, supplies limits for identity-changing dynamics, and defines collapse-invariants for classification. We show that identity-fluid systems are mathematically inevitable under finite resources and already implicit in coarse graining, quantum indistinguishability, entity resolution, and cognitive categorisation. Finally, we demonstrate problem-relative usefulness through two applications: (1) entity resolution under persistent constraints, where FCM provides canonical solutions and a stability-radius theorem guaranteeing when identity assignments are robust under streaming updates; and (2) tensor-network truncation, where FCM provides foundation-level semantics and proves path-independent error certificates for iterative coarse-graining.

---

1. Primitive Objects and Axioms

1.1 Basic Definitions

Definition 1.1 (Carrier Set). Let E be a finite set of primitive entities.

Definition 1.2 (Commitment Structure). A commitment structure over E is a pair K = (P, R), where:

- P is a partition of E into identity blocks
- $R \subseteq P \times P$ is an irreflexive, symmetric set of commitments between distinct blocks

Remark. Commitments are taken to be undirected (symmetric) here for minimality. Directed variants are possible without altering the collapse logic; the essential features—capacity bounds, admissibility, and identity-changing projection—transfer directly.

Definition 1.3 (Coarsening). Let P and P′ be partitions of E. We write $P' \preceq P$ ("P′ is coarser than P") if for every block $B \in P$ there exists a block $B' \in P'$ such that $B \subseteq B'$. Equivalently, every block of P′ is a union of blocks of P.

Definition 1.4 (Block Map). Given $P' \preceq P$, define the block map $\pi_{P \to P'} : P \to P'$ by $\pi_{P \to P'}(B)$ = the unique block $B^* \in P'$ such that $B \subseteq B^*$.

Definition 1.5 (Degree). The degree of a block $B \in P$ with respect to commitments R is:

$$\deg_R(B) := |\{ B' \in P : (B, B') \in R \text{ or } (B', B) \in R \}|$$

This counts the number of distinct external identity blocks distinguished from B.

Definition 1.6 (Capacity Functional). A capacity functional is a map $\mathcal{C} : 2^E \to \mathbb{N}$ assigning to each possible identity block a maximum sustainable degree.

Remark. No monotonicity of $\mathcal{C}$ under set inclusion is assumed. Capacity may increase, decrease, or remain constant as blocks merge. This generality allows FCM to model systems where larger aggregates have different (not necessarily greater) distinguishability resources.

Definition 1.7 (Admissibility). A commitment structure K = (P, R) is admissible if and only if:

$\deg\_R(B) \leq \mathcal{C}(B)$ for all $B \in P$

1.2 Axioms

FCM is governed by three axioms concerning admissibility rather than classical consistency.

Axiom A1 (Finite Distinguishability). Every identity block B has finite capacity $\mathcal{C}(B) < \infty$. No block may participate in more external commitments than its capacity permits.

Axiom A2 (Persistence of Commitment; Collapse Internalisation). Once established, commitments persist as constraints: they are not deleted by admissible updates. Under collapse, commitments between blocks that merge become internal and are therefore definitionally satisfied within the coarsened identity block. Thus commitments may cease to appear as external edges only by becoming internal under identity collapse.

Remark. Axiom A2 is the separatrix: weakening it permits provenance retention and collapses FCM back into quotient-style formalisms. The framework's distinctiveness depends on this persistence-plus-internalisation semantics.

Axiom A3 (Existence of Collapse Projection). For every commitment structure K, there exists at least one admissible structure K′ reachable from K by identity collapse (coarsening of the partition P).

Remark. Axiom A3 is an existence axiom, not an operational rule. The collapse projection Π_L defined below provides a canonical construction satisfying A3.

1.3 General Reader Explanation

Think of the carrier set E as a collection of things you're trying to keep track of—say, five people at a party. Initially, you can distinguish all five: that's your partition P, where each person is in their own "identity block."

Now suppose you make mental notes about who's different from whom: "Alice is definitely not Bob" (that's a commitment). The capacity $\mathcal{C}$ represents how many such distinctions any one person can anchor. If Alice can only be "the one who's not X" for two other people, her capacity is 2.

The problem arises when you try to distinguish Alice from *everyone*. If her capacity is 2 but you've committed to her being distinct from four others, you've exceeded the budget. Something has to give.

Axiom A1 says: budgets are finite and real. Axiom A2 says: you can't just forget a distinction you've made. Axiom A3 says: there's always a way out—by *merging identities* until the budget balances.

The mathematics that follows makes "merging identities" precise.

---

2. Interpretation and Role of the Capacity Functional

A natural question arises: what determines $\mathcal{C}$? This section clarifies the status of the capacity functional within FCM.

2.1 Status of $\mathcal{C}$

The capacity functional $\mathcal{C}$ is an input constraint, not a derived quantity. FCM does not assume a universal value of $\mathcal{C}$; instead, it studies the structural consequences of finite distinguishability under arbitrary but fixed capacity bounds. This is analogous to:

- Degree bounds in graph theory (where one studies k-regular graphs without deriving k)
- Arity limits in logic (where predicate arity is specified, not proven)
- Capacity constraints in network flow (where edge capacities are given data)

2.2 Possible Origins

Without committing to any particular interpretation, capacity bounds may arise from:

- **Cognitive resolution limits**: an observer can track only finitely many simultaneous distinctions
- **Physical interaction constraints**: a system can sustain only finitely many independent couplings
- **Computational bandwidth**: finite memory or channel capacity
- **Thermodynamic cost**: maintaining distinctions requires free energy; capacity reflects available resources

FCM remains agnostic on which (if any) of these applies in a given context.

2.3 Sensitivity and Robustness

Different choices of $\mathcal{C}$ induce different collapse dynamics, but do not alter the qualitative necessity of collapse, identity instability, or admissibility enforcement. The framework's core results—idempotence of Π_L, failure of categorical structure, existence of collapse-convergent limits—hold for any valid capacity functional.

2.4 A Toy Derivation: Capacity from a Distinction Budget

While FCM treats $\mathcal{C}$ as an input, we can show how capacity might emerge from more primitive constraints.

Toy Model (Budgeted Distinguishability). Assume each identity block $B \subseteq E$ is associated with a nonnegative resource $w(B) \in \mathbb{R}_{\geq 0}$ representing the "distinction budget" available to that block. Assume that sustaining one external commitment incident to B consumes at least $\alpha > 0$ units of budget. Then admissibility requires:

$$\deg\_R(B) \cdot \alpha \leq w(B)$$

Define the induced capacity functional:

$$\mathcal{C}(B) := \lfloor w(B) / \alpha \rfloor$$

This yields a derived capacity rule rather than an oracle.

Aggregation rule (one natural choice). If blocks merge, budgets add:

$$w(B \cup B') = w(B) + w(B')$$

Then $\mathcal{C}$ is (weakly) superadditive up to flooring effects, but need not be monotone under inclusion if other aggregation rules are chosen (e.g., diminishing returns or interference costs).

Example. Setting $w(\{a\}) = 2\alpha$ for a singleton block yields $\mathcal{C}(\{a\}) = 2$, recovering the uniform capacity used in Appendix A.

Edge case. If $w(B) = 0$ for some block, then $\mathcal{C}(B) = 0$, meaning that block can sustain zero external commitments. This is permitted by the framework (Lemma 3.1 handles $\mathcal{C}(E) = 0$). A zero-capacity block forces immediate collapse if any commitment involving it is attempted.

Remark (Predictive content). Under this toy model, $\mathcal{C}$ is no longer arbitrary; it is determined by α and the block budgets $w(\cdot)$. Different physical/cognitive/computational interpretations correspond to different choices of w and aggregation rules, while the FCM collapse machinery remains unchanged. FCM's predictive content is conditional: given a resource model that induces $\mathcal{C}$ and a choice of L, the collapse outcome and collapse-convergent limits are determined up to the induced equivalence class.

2.5 General Reader Explanation

You might wonder: where does the capacity limit come from? The honest answer is that FCM doesn't say. It's like asking where the speed limit on a road comes from—FCM studies what happens *given* a speed limit, not why the limit is 65 rather than 70.

This is actually a strength. The same mathematics applies whether the limit comes from:

- How many things a brain can track simultaneously

- How many particles can interact at once
- How much memory a computer has
- How much energy is available to maintain distinctions

By not committing to a source, FCM can model all of these. The results we prove—that collapse happens, that it's irreversible, that it breaks normal mathematical structure—follow from the *existence* of any finite limit, not from its specific value.

---

3. The Collapse Projection Π_L

3.1 Loss Functionals

To select among possible coarsenings, we introduce a loss functional.

Definition 3.1 (Loss Functional). A loss functional L is a map $L : (P, P') \to \mathbb{R}_{\geq 0}$ defined on pairs of partitions where P′ is coarser than P, satisfying:

1. **Domain**: L(P, P′) is defined only when $P' \preceq P$ (P′ coarser than P)
2. **Identity**: L(P, P) = 0
3. **Positivity**: L(P, P′) > 0 for P′ strictly coarser than P
4. **Monotonicity**: If $P' \preceq P'' \preceq P$, then $L(P, P') \geq L(P, P'')$
5. **Finiteness**: For finite E, the minimum of L over admissible coarsenings exists

3.2 Definition of Π_L

Definition 3.2 (Pushforward of Commitments). Let $P' \preceq P$ and let $\pi = \pi_{\{P \to P'\}}$ be the block map. The pushforward of $R \subseteq P \times P$ along π is:

$R' := \{ (\pi(B_1), \pi(B_2)) : (B_1, B_2) \in R, \pi(B_1) \neq \pi(B_2) \}$

That is, R′ consists of those commitments from R that remain external under P′, with duplicates removed.

Definition 3.3 (Collapse Projection). Given a commitment structure K = (P, R) and loss functional L, define:

$\Pi\_L(K) := \{ K' = (P', R') : P' \preceq P$, R′ is the pushforward of R, K′ is admissible, and L(P, P′) is minimal $\}$

Remark. When multiple admissible partitions achieve the same minimal loss, Π_L(K) is the equivalence class of all such minimisers. This set-valued definition is mathematically honest and aligns with the universality class interpretation developed in Appendix B.

3.3 Core Properties

Lemma 3.1 (Existence of Minimal Collapse). For finite carrier E and loss functional L satisfying Definition 3.1, the set of admissible coarsenings of any K contains at least one L-minimiser.

Proof. The set of partitions of E coarser than P is finite. Among these, the subset of admissible partitions is non-empty: the single-block partition $P' = \{E\}$ is always admissible, since all commitments become internal under P′, giving $R' = \emptyset$ and $\deg_{R'}(E) = 0 \leq \mathcal{C}(E)$ (including the edge case $\mathcal{C}(E) = 0$). The minimum of L over a non-empty finite set exists.

Lemma 3.2 (Idempotence). If K is admissible, then $\Pi_L(K) = \{K\}$. If K is inadmissible and $K' \in \Pi_L(K)$, then $\Pi_L(K') = \{K'\}$.

Proof. If K is admissible, then K itself achieves $L(P, P) = 0$, which is minimal by positivity. Hence $\Pi_L(K) = \{K\}$. If $K' \in \Pi_L(K)$, then K′ is admissible by definition, so $\Pi_L(K') = \{K'\}$ by the first case.

Lemma 3.3 (Non-Invertibility). Collapse projection is not injective: distinct inadmissible structures may collapse to the same admissible structure.

Proof. Let $K = (P, R)$ and $K' = (P, R \cup S)$ where every commitment in S becomes internal under any minimising coarsening P^*. Then the admissible coarsenings of K and K′ coincide, so $\Pi_L(K) = \Pi_L(K')$ despite $K \neq K'$.

Lemma 3.4 (Representative Independence). The structural properties of Π_L—idempotence, non-invertibility, and identity change—hold for all representatives in the equivalence class $\Pi_L(K)$.

Proof. These properties depend only on admissibility and the coarsening relation, which are shared by all minimisers.

3.5 Computational Note

Remark (Computability and complexity of Π_L). For finite E, $\Pi_L(K)$ is computable for any admissible loss functional L: one may enumerate all coarsenings of P, compute the induced R′, test admissibility, and select the L-minimisers. This brute-force procedure is finite but may scale with Bell numbers $B(|E|)$. Whether Π_L is *efficiently* computable depends on the representation of R, the form of $\mathcal{C}$, and the choice of L. Many natural cases admit greedy or graph-clustering approaches; identifying tractable subclasses and complexity bounds is left open.

3.6 Not a Quotient Construction

Remark (Collapse vs. Quotient). A quotient construction begins with a fixed underlying set E and an equivalence relation ∼, producing a quotient set $E/\sim$ together with a canonical surjection q : E

$\rightarrow E/\sim$. The original elements remain available as a domain, and q provides a well-defined map from pre-quotient to post-quotient structure.

In FCM, collapse does not produce a quotient of a fixed structure under a fixed relation. Instead, collapse produces a new commitment structure whose identity blocks are the ontological units; there is, in general, no internal map from the post-collapse structure back to a pre-collapse domain because the framework forbids retaining lineage/provenance (Axiom A2). The domain has genuinely changed—collapse is not "working modulo an equivalence" but a rule that changes what counts as an object.

3.7 General Reader Explanation

The collapse projection Π_L is FCM's central tool. Here's what it does in plain terms:

Given a system that's "over budget" on distinctions, Π_L finds the cheapest way to merge identities until the budget balances. "Cheapest" is determined by the loss functional L—a measure of how much you dislike merging.

The key properties:

Idempotence (Lemma 3.2): If a system is already within budget, collapse does nothing. If you collapse an over-budget system, the result is within budget and won't collapse further. Collapse stabilises.

Non-invertibility (Lemma 3.3): Different starting points can collapse to the same endpoint. Once merged, you can't tell what was separate before. Information is genuinely lost.

Set-valued output: Sometimes there are multiple equally-cheap ways to collapse. Rather than arbitrarily picking one, Π_L returns all of them as equivalent options. This honesty about non-uniqueness becomes important later (see Appendix B).

The loss functional L captures *what you care about preserving*. If L penalises losing identity blocks, you'll merge as few things as possible. If L penalises losing relationships, you might merge more aggressively to preserve the connections that remain.

---

4. Collapse-Convergence

Classical convergence requires a fixed underlying space and metric. When identity or resolution changes dynamically, classical convergence often fails. FCM introduces an alternative notion.

Definition 4.1 (Collapse-Convergence). Let $\{K_t\}$ be a sequence of commitment structures. The sequence is collapse-convergent if there exists $N \in \mathbb{N}$ such that:

$$\Pi_L(K_t) = \Pi_L(K_{t+1}) \text{ for all } t \geq N$$

where equality is of sets of admissible structures.

Remark. No topology or metric is assumed. Convergence is defined purely as eventual stability of the projected equivalence class under admissibility. By using set equality of $\Pi_L(K_t)$ rather than representative comparison, we avoid any dependence on choice functions. This notion applies precisely where classical convergence fails due to excessive distinguishability demands.

4.2 General Reader Explanation

In ordinary mathematics, a sequence "converges" if its terms get closer and closer to some limit—think of 1, 1/2, 1/4, 1/8, ... approaching zero. This requires a notion of "distance" and a fixed space in which distances are measured.

But what if the space itself is changing? What if the very things you're measuring can merge into each other? Classical convergence breaks down.

Collapse-convergence offers an alternative. Instead of asking "do the raw values settle down?" (they might not), we ask "do the *collapsed* values settle down?"

Imagine a process that forever oscillates between trying to distinguish A from B, then A from C, then A from B again. Classically, this never converges. But if the system's capacity forces A and B to merge early on, then all subsequent attempts to distinguish them become meaningless—they're the same thing now. The collapsed sequence *does* settle down, even though the raw sequence doesn't.

This is powerful: it lets us assign meaningful limits to processes that look chaotic from the outside.

---

5. Why Collapse-Projected Updates Do Not Form a Category

One might attempt to organise admissible FCM structures into a category by taking morphisms to be collapse-projected updates:

- Objects: admissible commitment structures K
- Morphisms: update maps followed by collapse projection

Here we analyse the naïve categorical attempt in which morphisms are specified as block-addressed updates on current identity blocks (not entity-lifted updates as in Appendix C), so that collapse may change the domains of subsequent updates.

This attempt fails for fundamental reasons.

Theorem 5.1 (Failure of Categorical Structure). The collection of FCM commitment structures under collapse-projected updates does not form a category.

Proof. We exhibit failure of composition. Let E = {a, b, c} with $\mathcal{C}(B) = 1$ for all blocks.

Define $K_1 = (\{\{a\}, \{b\}, \{c\}\}, \emptyset)$.

Let f be the update adding commitment (a, b). After collapse:

$K_2 = \Pi\_L(f(K_1)) = (\{\{a, b\}, \{c\}\}, \emptyset)$

Let g be the update adding commitment (a, c). The update g is not well-defined on K_2: its domain presupposes the existence of a block containing a as a singleton, but collapse has replaced that block with the strictly coarser block {a, b}.

Therefore g ∘ f does not exist, and composition is not closed.

Remark (Provenance and Information Loss). One might attempt to preserve compositionality by attaching provenance labels to elements through collapse—tracking which original entities merged into which blocks. FCM explicitly disallows such enrichment: collapse is defined as a loss of identity resolution, not as a quotient with traceable lineage. Introducing provenance would reintroduce the very distinctions collapse eliminates, violating admissibility. This is why Axiom A2 (Persistence of Commitment) is fundamental: if provenance were retained, collapse could in principle be reversed by unpacking the lineage, contradicting the framework's core premise.

---

6. Distinction from Existing Frameworks

FCM is not a notational variant of existing finitary mathematics. The following table summarises the key distinctions:

Framework	Fixed Identity?	Internal Collapse?	Admissibility-Based Failure?
Ultrafinitism	Yes	No	No
Domain theory	Yes	No	Partial
Rough sets	Yes	No	No
Finite model theory	Yes	No	No
FCM	**No**	**Yes**	**Yes**

6.1 Ultrafinitism

Ultrafinitism rejects completed infinities and restricts attention to "feasibly constructible" numbers, but operates on fixed finite objects with stable identity. There is no mechanism by which an ultrafinitist object changes identity in response to constraint violations.

6.2 Domain Theory

Domain theory uses approximation and partial information to handle computation over infinite structures. Objects retain identity under refinement; one approaches a limit from below. There is no collapse operator—approximations grow more refined, never coarser.

6.3 Rough Sets

Rough set theory allows indiscernibility classes, but these are externally imposed equivalences based on available attributes. The equivalence relation is given, not emergent from internal constraint violations. Identity within the underlying universe remains fixed.

6.4 Finite Model Theory

Finite model theory studies logical properties of finite structures under fixed signatures. Identity is determined by the model specification and does not change during analysis. There is no capacity bound regulating how many distinctions a structure can sustain.

6.5 The FCM Distinction

FCM alone treats identity as dynamically regulated by finite capacity. Collapse is not an approximation, not an external quotient, and not a modelling choice—it is an internal structural response to admissibility violation. This makes FCM's failure modes (non-categorical structure, non-invertible projection) intrinsic rather than artefactual.

Remark (Not graph theory with extra steps). While (P, R) is graph-like—a graph on identity blocks—and admissibility resembles degree bounds, the novelty is that the vertex set itself is endogenous: collapse changes the object set by merging vertices. The underlying graph is not fixed. Standard graph-theoretic constructions assume stable vertices; FCM does not.

6.6 Relation to Rewriting Systems and Dynamic Algebras

Remark. Abstract rewriting systems and dynamic/evolving algebra frameworks (e.g., state-based transition formalisms) allow structures to change over time. However, they typically preserve a fixed underlying carrier (or preserve element identity via persistent names/addresses), with evolution occurring in relations, functions, or predicates defined on that carrier.

FCM differs in that the carrier of "objects" is the partition itself: collapse changes what counts as an object by merging identity blocks, and the framework explicitly forbids retaining lineage/provenance (Axiom A2). Thus FCM is not merely a dynamic algebra on a fixed domain; it is a *dynamics of domains*.

FCM can be viewed as a partial dynamics on partitions with a projection operator, rather than a rewriting system with confluent normal forms.

6.7 General Reader Explanation

Why develop a new framework when mathematicians have studied finite structures for decades?

The key difference is where identity lives. In all the frameworks listed above, identity is *given from outside*: you start with a set of distinct things, and they stay distinct throughout your analysis. You might group them, approximate them, or study their logical properties—but the underlying things remain the same things.

FCM makes identity *internal and dynamic*. Things start distinct, but if the system's capacity is exceeded, they *become the same thing*—not as a modelling convenience, but as a fundamental operation within the mathematics.

This is like the difference between:

- A filing system where you can put documents in folders (rough sets, finite model theory)
- A filing system where documents can *merge into each other* and become a single document

The second kind of system behaves very differently. It doesn't form a category. Classical convergence doesn't apply. Standard proof techniques break. But it may be exactly what's needed to model situations—in physics, cognition, or computation—where identity is genuinely fluid.

---

7. Non-Representability: Why FCM Is New Mathematics

The claim that FCM constitutes "new mathematics" requires precise justification. We do not claim novelty in the sense of importance or applicability, but in the sense of *non-embeddability*: FCM's central structures and operators cannot be faithfully represented within existing mathematical frameworks without violating their axioms.

Definition 7.1 (External Rule). An external rule is any meta-level stipulation that changes the object set, morphism domains, or composition behavior mid-derivation, rather than being derivable from the framework's axioms. Examples include: stipulating that certain objects "disappear" after an operation, adding provenance tracking not present in the original signature, or declaring certain compositions undefined by fiat.

Theorem 7.2 (Non-Representability). There exists no faithful representation of Finite Commitment Mathematics within any framework satisfying all of the following:

1. Fixed underlying object identity
2. Total composition of updates

3. Preservation of provenance under equivalence or quotienting

In particular, FCM collapse dynamics cannot be represented as:

- a category (Theorem 5.1; the categorical obstruction is made explicit in Proposition D.1),
- a preorder or lattice,
- a quotient construction,
- a confluent rewriting system, or
- a domain-theoretic approximation process.

Proof.

(i) **Categories require fixed objects and closed composition.** Theorem 5.1 shows that collapse-projected updates fail composition closure because the domain of subsequent morphisms depends on prior collapse outcomes. No category with fixed objects can model this.

(ii) **Quotient constructions preserve surjections.** A quotient $E \to E/\sim$ retains a canonical map from the original domain. Axiom A2 forbids provenance retention; there is no internal map from post-collapse structure to pre-collapse domain. Hence collapse is not a quotient.

(iii) **Rewriting systems preserve term identity.** Standard rewriting tracks positions through reduction; the "same subterm" is well-defined before and after rewriting. FCM collapse destroys positional identity—merged entities have no individual positions. Hence FCM is not a rewriting system.

(iv) **Domain theory requires refinement monotonicity.** In domain-theoretic semantics, approximations grow more refined (more informative) over time. FCM collapse is *anti-refinement*: it reduces resolution, coarsening the partition. The direction is reversed.

(v) **Preorders and lattices require an internal, fixed-domain ordering compatible with dynamics.**

A natural attempt is to define an order relation from the collapse dynamics itself. Let $K = (P, R)$ and $K' = (P', R')$ be commitment structures over the same carrier E. Write $K \to_{\Pi_L} K'$ to mean $K' \in \Pi_L(K)$.

We consider two canonical candidates for an "FCM order" and show both fail to capture the dynamics without external rules.

Definition 7.3 (Collapse-step relation). Define a binary relation $\trianglelefteq$ on structures by:

$$K \trianglelefteq K' \Leftrightarrow K' \in \Pi_L(K)$$

This is the most direct attempt to turn collapse into an order.

Lemma 7.4 (Non-reflexivity of the collapse-step relation). If K is inadmissible, then $K \ntrianglelefteq K$. Hence $\trianglelefteq$ is not reflexive on the full class of commitment structures.

Proof. If K is inadmissible, then every element of $\Pi\_L(K)$ is admissible by definition, so $K \notin \Pi\_L(K)$. Therefore $K \ntrianglelefteq K$.

A preorder must be reflexive; thus the collapse-step relation cannot be a preorder on all structures. One may restrict objects to the admissible subcollection, but then $\trianglelefteq$ becomes trivial: by Lemma 3.2, if K is admissible then $\Pi\_L(K) = \{K\}$, so $K \trianglelefteq K'$ implies $K' = K$.

Definition 7.5 (Coarsening-only relation). A second natural attempt is to ignore commitments and order only by partition coarsening:

$$(P, R) \leq (P', R') \Leftrightarrow P' \preceq P$$

(This uses only identity resolution.)

Lemma 7.6 (Coarsening-only order collapses commitment structure). The relation $\leq$ fails to distinguish commitment states: if $R_1 \neq R_2$ on the same partition P, then $(P, R_1) \leq (P, R_2)$ and $(P, R_2) \leq (P, R_1)$, so antisymmetry fails and the induced equivalence relation identifies structurally distinct states.

Proof. If both structures have the same partition P, then $P \preceq P$, hence both inequalities hold. But $R_1 \neq R_2$, so the structures are distinct.

Thus $\leq$ can define a preorder only by quotienting out commitment information—contrary to the intent of FCM, where R carries persistent commitments (Axiom A2).

Proposition 7.7 (No preorder/lattice without external closure or loss of structure). There is no "natural" preorder on the class of FCM commitment structures that simultaneously:

1. is defined internally from (P, R) and the collapse operator Π_L (without adding external closure rules),
2. is reflexive and transitive on the intended class of objects (admissible and inadmissible), and
3. preserves commitment structure up to isomorphism (i.e., does not identify distinct R-states on the same P).

Proof. Any preorder induced directly by collapse steps fails reflexivity on inadmissible structures (Lemma 7.4), and restricting to admissible structures trivialises the relation (Lemma 3.2). Any preorder induced only by coarsening ignores commitments and collapses distinct states (Lemma 7.6). Forcing reflexivity by reflexive closure or recovering commitments by quotienting introduces precisely the kind of meta-level modification disallowed as an external rule (Definition 7.1).

Consequence. Since lattices and most order-theoretic constructions presuppose a stable underlying set equipped with a compatible preorder, FCM collapse dynamics is not representable as a preorder- or lattice-based semantics without either (a) adding external closure rules or (b) discarding commitment information.

Therefore, FCM defines mathematical structures and operators not expressible within these frameworks without adding external machinery. By standard usage in foundations, this constitutes new mathematics.

Remark. The novelty of FCM does not rest on applications but on non-reducibility: its central operator and dynamics provably cannot be internalised within existing mathematical frameworks without violating their axioms.

---

8. Mathematical Usefulness of Finite Commitment Mathematics

The usefulness of a mathematical framework does not depend on whether it ultimately describes physical, cognitive, or computational systems. In pure mathematics, a framework is useful if it renders previously ill-posed problems well-posed, supplies canonical outcomes or invariants where none existed, or provides a semantics for processes that otherwise lack one.

In this section, we show that FCM is useful in precisely this sense. Its utility follows directly from its axioms and does not rely on external interpretation.

8.1 Well-Posedness of Identity-Overloaded Constraint Systems

Classical constraint systems assume a fixed domain of objects. When constraints exceed what the domain can satisfy, the system becomes inconsistent or undefined.

FCM replaces inconsistency with collapse.

Theorem 8.1 (Canonical Resolution of Overloaded Constraints). Let $K = (P, R)$ be any finite commitment structure with capacity functional $\mathcal{C}$ and loss functional L. Then:

1. $\Pi\_L(K)$ is non-empty.
2. Every $K' \in \Pi\_L(K)$ is admissible.
3. If K is admissible, then $\Pi\_L(K) = \{K\}$.

Proof. Existence and admissibility follow from Lemma 3.1. If K is admissible, then $L(P, P) = 0$ is minimal by Definition 3.1, hence $\Pi\_L(K) = \{K\}$.

Interpretation. FCM replaces "constraint failure" with a canonical resolution operator acting on identity itself. Any finite commitment structure admits a well-defined admissible outcome with minimal loss. This is a strict strengthening over classical constraint satisfaction, which offers no outcome once constraints become incompatible.

8.2 Resolution Semantics for Update Sequences

Although FCM structures do not form a category (Section 5), finite sequences of updates still admit a well-defined semantics once collapse is taken into account.

Definition 8.2 (Update Semantics). Let $U = (u_1, \ldots, u_n)$ be a finite sequence of update rules, each interpreted element-wise and lifted to blocks as described in Appendix C. Define the collapse-resolved semantics of U by:

$\mathrm{Sem\_L}(U)(K) := \Pi\_L(u_n(\cdots u_1(K)\cdots))$

Theorem 8.2 (Well-Defined Semantics). For any finite update sequence U and any commitment structure K, the set Sem_L(U)(K) is non-empty and consists of admissible structures.

Proof. The composite update yields a finite commitment structure. Applying Π_L yields a non-empty set of admissible structures by Lemma 3.1.

Interpretation. Even though individual updates may become undefined or vacuous due to identity collapse, every finite update sequence has a well-defined resolved outcome. FCM thus provides a semantics for identity-changing processes where categorical composition fails.

8.3 Existence of Limits for Identity-Changing Dynamics

Many dynamical systems fail to converge because their state space changes over time. FCM supplies a notion of limit that survives such instability.

Theorem 8.3 (Existence of Collapse-Convergent Limits). Fix finite E, capacity functional $\mathcal{C}$, and loss functional L. Let (K_t)*{t≥0} be any sequence of commitment structures generated by iterated updates. Then the sequence (Π_L(K_t))*{t≥0} is eventually constant or eventually periodic. In particular, if Π_L is applied after each update, the projected sequence stabilises by idempotence (Lemma 3.2).

Proof. The number of admissible commitment structures over finite E is finite. The projected sequence Π_L(K_t) therefore takes values in a finite set. By finiteness, the sequence is eventually periodic. If idempotence applies (Lemma 3.2), periodicity collapses to constancy.

Interpretation. FCM assigns meaningful limits to processes that are classically divergent or oscillatory, without invoking topology, metrics, or continuity.

8.4 Collapse-Invariants and Classification

Collapse produces canonical structural summaries that persist under further evolution.

Definition 8.3 (Collapse-Invariant Signature). Let K* ∈ Π_L(K). Any function I of K* invariant under further collapse (e.g., number of blocks, block-size multiset, degree distribution of (P*, R*)) is called a collapse-invariant signature of K.

Proposition 8.4 (Stability of Collapse-Invariants). If I is a collapse-invariant signature and K′ ∈ Π_L(K), then I(K′) = I(K″) for all K″ ∈ Π_L(K′).

Proof. By idempotence, Π_L(K′) = {K′}.

Interpretation. FCM provides classification tools for unstable systems: distinct raw structures can be compared via their collapsed invariants even when their uncollapsed dynamics are ill-behaved.

8.5 Summary of Mathematical Usefulness

FCM is mathematically useful because it:

1. **Renders overloaded constraint systems well-posed** via canonical collapse
2. **Provides semantics for update sequences** despite failure of categorical composition
3. **Supplies a notion of limit** for identity-changing dynamics where classical convergence fails
4. **Defines collapse-invariants** enabling classification and comparison

These properties follow directly from the axioms of FCM and do not depend on external interpretation. Whether identity-fluid systems arise in physics, cognition, or computation is an empirical question. The mathematical utility of FCM—understood as the provision of canonical outcomes, limits, and invariants where standard frameworks fail—is independent of that question.

---

9. Why Identity-Fluid Systems Exist and Merit Mathematics

A natural question arises: do identity-fluid systems actually exist, or is FCM mathematics in search of a subject? This section argues that identity-fluidity is mathematically inevitable under finite resources and already implicit in core scientific disciplines.

9.1 Two Notions of Identity

Two claims are often conflated:

- **Metaphysical identity**: the world has perfectly individuated objects "in themselves"
- **Operational identity**: the identity a theory can support given finite observation/representation resources

FCM concerns operational identity. And operational identity is provably resource-dependent.

This isn't philosophy—it's built into standard practice in statistics, physics, and computation: indistinguishability is defined relative to tests, resolution, or representational capacity.

Theorem 9.1 (Finite Resolution Forces Identity Collapse). Let E be a finite carrier of "microstates." Let M be a finite family of observables/tests available to an agent or theory. Define an observational signature:

$$\sigma_M(x) := (m(x))_{m \in M}$$

Define operational indistinguishability:

$$x \sim_M y \Leftrightarrow \sigma_M(x) = \sigma_M(y)$$

Then $\sim_M$ is an equivalence relation whose classes are the identities the theory can actually sustain.

Claim. If the available tests/resources shrink ($M' \subset M$), then operational identity coarsens:

$$x \sim_M y \Rightarrow x \sim_{M'} y$$

and in general there exist x, y such that $x \nsim_M y$ *but* $x \sim_{M'} y$. Distinct identities collapse when resolution decreases.

Proof. If $\sigma_M(x) = \sigma_M(y)$, then restricting coordinates to M' preserves equality, so $x \sim_{M'} y$. For strict coarsening, pick x, y that differ only on some test in $M \setminus M'$.

Interpretation. This shows identity-fluidity is mathematically inevitable whenever identity is operational and resources are finite. It is exactly the logic behind coarse-grained ensembles in statistical mechanics and information-loss under coarse graining and renormalisation. FCM generalises this from "microstates collapse under coarse graining" to "identity blocks collapse under capacity overload."

9.2 Identity-Fluidity Already Exists in Core Disciplines

Even setting aside metaphysics, identity-fluidity is already standard in at least three mature areas:

(A) Physics: Indistinguishability is not optional. Quantum theory treats many particles as fundamentally indistinguishable; labels are not physical observables. This is identity-fluidity at the level of "which particle is which"—and it drives non-classical state spaces (Fock space, symmetrisation). FCM doesn't replace quantum mechanics; it supplies a general mathematical language for identity regulated by finite distinguishability.

(B) Computation/Data: Entity resolution literally merges identities. In data systems, entity resolution and deduplication is the explicit act of merging records into a single identity when the evidence cannot sustain distinctness. This is a practical, widely studied, mathematically nontrivial problem. What makes it hard is precisely what FCM formalises: finite evidence and constraints force identity merges.

(C) Cognition: Categories exist because memory is limited. Human working memory is limited and categorisation depends on capacity constraints; what counts as "the same kind of thing" is resolution-dependent. Identity is not an externally fixed primitive in practice—it is a resource-managed construct.

9.3 Inevitability Under Finite Resources

Theorem 9.2 (Finite Distinguishability Forces Identity-Fluidity). Assume a domain in which (i) representational capacity for distinctions is finite, formalised by a capacity functional $\mathcal{C}$, and (ii) distinctions persist as commitments (Axiom A2). Assume further that there exists an update process capable of generating, for some initial structure $K_0 = (P_0, R_0)$, arbitrarily many distinct external commitments incident to a single identity block prior to collapse. Then there exist attainable states in which maintaining all distinctions is impossible without coarsening identity.

Proof. Let $B \in P_0$ be a block for which the assumed update process can generate n distinct external neighbours in R (i.e., $\deg\_R(B) \geq n$) for arbitrarily large n, before applying collapse. By Axiom A1, $\mathcal{C}(B) < \infty$. Choose $n > \mathcal{C}(B)$. After the updates, the resulting structure $K = (P, R)$ satisfies $\deg\_R(B) > \mathcal{C}(B)$ and is therefore inadmissible.

By Axiom A2, the excess commitments cannot be removed by revocation. Hence admissibility cannot be restored while keeping the partition P fixed and the commitments external. By Axiom A3, there exists an admissible structure reachable by identity collapse, i.e., by coarsening P to some $P' \preceq P$ and pushing commitments forward. Such a coarsening necessarily merges blocks, thereby changing identity.

Therefore, there exist attainable states (those produced by the update process with $n > \mathcal{C}(B)$) for which restoring admissibility requires identity coarsening.

Remark. This result is conditional only on (a) finite distinguishability capacity and (b) the existence of processes that can demand more distinctions than a finite capacity allows—both of which are generic in resource-limited settings. In this precise sense, identity-fluidity is not an optional philosophical stance but a structural consequence of finite resources plus persistent commitments.

9.4 Why Explicit Mathematics Is Needed

Even if one maintains that "identity is stable in reality," the moment one does any of the following, one is already working in an identity-fluid regime:

- Coarse-grain states (physics, statistical mechanics)
- Treat indistinguishable particles (quantum theory)
- Deduplicate or entity-resolve (data systems)

- Use categories because memory is limited (cognition)

What's missing in the literature is a unified, axiomatic treatment where identity change is internal and capacity-regulated, and where failure modes (non-categoricity, new convergence notions) are derived rather than patched. That is the mathematical niche FCM fills.

9.5 Summary

Identity-fluid systems exist at least in the operational sense: whenever distinguishability is finite and resource-limited, the equivalence classes induced by available tests or representations necessarily coarsen as resources tighten, forcing identity collapse (Theorem 9.1). This phenomenon is not exotic; it is implicit in coarse graining and renormalisation, in indistinguishability in quantum theory, and in entity resolution and categorisation under finite memory. These domains therefore motivate, and arguably require, explicit mathematics of identity change. FCM provides such a framework by internalising capacity-regulated collapse as a primitive operator and deriving its structural consequences.

---

10. Application: Entity Resolution Under Finite Distinguishability

To demonstrate that FCM addresses a problem people already work on—and handles it more cleanly than existing approaches—we consider entity resolution with persistent distinction constraints.

10.1 Problem Statement

Given a growing set of entities E (database records, knowledge graph nodes, sensor readings) and persistent pairwise distinction constraints R, determine a stable identity assignment under finite resolution constraints.

This is a real problem in:

- **Databases**: Record deduplication with audit trails
- **Knowledge graphs**: Entity linking with provenance requirements
- **Record linkage**: Matching under legal constraints that forbid revoking distinctions
- **Data integration**: Merging sources where some distinctions are contractually mandated

The core tension: evidence accumulates over time and cannot be revoked (logs, audits, regulatory constraints), but systems have finite resolution capacity. At some point, identities must merge.

Commitment semantics for entity resolution. In this application, a commitment $(X, Y) \in R$ does not mean "X and Y can never be merged under any circumstances." It means: the system has recorded a persistent distinction constraint—a fact, rule, or audit requirement that must not be revoked or deleted.

Under identity collapse, such commitments are not removed; rather, if two committed blocks become part of the same identity block, the commitment becomes internal and is thereby satisfied as an audit constraint: the record that a distinction once existed is retained, but no longer appears as an external separability constraint because the identity units have changed.

Remark. If an application requires "hard cannot-link" constraints that must remain external forever, then that requirement is incompatible with finite distinguishability under overload; the system becomes inconsistent rather than collapsible. FCM formalises the alternative: accept finite identity resources and allow principled collapse.

10.2 How Existing Tools Handle It

Current approaches fall into four categories, each importing external rules:

(A) Procedural heuristics. Thresholds, clustering, greedy merges. These produce outcomes, but the outcome depends on threshold choices, merge order, and stopping criteria. No canonical result; different runs may yield different identity assignments.

(B) Probabilistic models. Posterior distributions over identity assignments. Elegant, but identity never actually *changes*—the model maintains uncertainty rather than committing to a resolution. When forced to output a single assignment, external decision rules (MAP, sampling) must be imported.

(C) Graph clustering. Cut the commitment graph by hand-chosen criteria (modularity, min-cut, spectral methods). Again, the outcome depends on the clustering objective—an external optimisation rule not intrinsic to the data.

(D) Quotient constructions. Declare equivalence relations based on evidence. This retains provenance: the original entities remain available, and the quotient map is invertible in principle. Collapse is reversible, which may violate audit or regulatory constraints requiring genuine identity consolidation.

Common deficiency: All four approaches require external decision rules to resolve what happens when constraints exceed capacity. The mathematics does not *force* a particular outcome; humans must choose.

Explicit failure mode. Given persistent distinctions and a hard capacity bound, the classical CSP "satisfy all cannot-link edges" becomes infeasible; clustering methods must either violate edges, soften them into penalties, or change the constraint set. FCM formalises that inevitability as identity collapse rather than treating it as an error condition.

10.3 FCM Formulation

Model the system as a commitment structure:

- **Carrier E**: the set of records/entities
- **Partition P**: current identity assignment (which records are "the same entity")
- **Commitments R**: persistent distinction constraints (see 10.1 for semantics)
- **Capacity $\mathcal{C}$**: finite resolution capacity (how many external commitments each identity block can sustain)
- **Loss functional L**: encodes priority (prefer preserving identity resolution vs. minimising commitment internalisation)

When commitments exceed capacity, collapse is *forced*, not chosen:

$\Pi\_L(K)$ yields the canonical admissible identity structure with minimal loss.

Remark. FCM does not eliminate modelling choices; it relocates them from the moment of overload to the prior specification of admissibility and loss. This mirrors how variational principles relocate choice from dynamics to action functionals. Once ($\mathcal{C}$, L) are fixed, the collapse outcome is determined up to the minimiser equivalence class.

10.4 What FCM Provides That Existing Tools Do Not

Theorem 10.1 (Canonical Resolution). For any finite entity set E, persistent commitment set R, capacity functional $\mathcal{C}$, and loss functional L, the collapse projection $\Pi\_L(K)$ yields:

1. A non-empty set of admissible identity structures
2. Minimal loss among all admissible resolutions
3. A well-defined outcome without external decision rules

Proof. Direct application of Theorem 8.1.

Theorem 10.2 (Collapse-Convergence for Streaming Updates). Let (u_t) be a stream of updates adding new entities or commitments. The projected identity sequence ($\Pi\_L(K\_t)$) is collapse-convergent: it stabilises or cycles within a finite equivalence class.

Proof. Direct application of Theorem 8.3.

Key differences from existing approaches:

Property	Heuristics	Probabilistic	Clustering	Quotients	FCM
Canonical outcome	No	No	No	Yes*	Yes
Identity actually changes	Yes	No	Yes	No	Yes
No external resolution rules	No	No	No	No	Yes
Handles streaming updates	Varies	Yes	No	Yes	Yes

Property	Heuristics	Probabilistic	Clustering	Quotients	FCM
Provenance discarded	Varies	No	Varies	No	Yes

*Quotients are canonical given a fixed equivalence relation, but choosing which relation to impose is itself an external decision.

FCM's distinction is not that other methods cannot output a result, but that they require extra decision structure (thresholds, tie-breakers, selection rules, optimisation objectives, or posterior decision criteria) that is not forced by the persistent-constraint-plus-capacity setting itself. FCM makes the resolution step internal: given ($\mathcal{C}$, L), collapse is a canonical admissibility projection rather than an external adjudication.

10.5 Concrete Example

Let E = {a, b, c, d, e} be database records. Suppose auditing requirements establish persistent distinction constraints:

R = {(a,b), (a,c), (a,d), (a,e), (b,c)}

Interpretation: records a and b, a and c, a and d, a and e, and b and c have been flagged as requiring distinction (e.g., different customer accounts, separate audit trails). These constraints persist but may become internal under collapse (see 10.1).

Suppose $\mathcal{C}$(B) = 2 for all blocks (each identity block can sustain at most 2 external commitments).

Initial state: P_0 = {{a}, {b}, {c}, {d}, {e}} (all records distinct)

Check admissibility: deg_R({a}) = 4 > 2 = $\mathcal{C}$({a}). Inadmissible.

Apply Π_L: Must coarsen P until admissible. Consider candidate merges:

Why not merge a with another block? Any merge involving a (e.g., {a,b}) would internalise one commitment but a would still be incident to three external commitments from the remaining pairs. This does not efficiently reduce deg_R(a).

Preferred first merge: Merge d and e (not committed to be distinct from each other). This reduces |P| while internalising zero commitments:

P_1 = {{a}, {b}, {c}, {d,e}} R_1 = {({a}, {b}), ({a}, {c}), ({a}, {d,e}), ({b}, {c})}

Now deg_R({a}) = 3 > 2. Still inadmissible. Continue:

Two candidate paths from P_1:

Path A: Merge b and c, internalising (b,c): $P_2^A = \{\{a\}, \{b,c\}, \{d,e\}\}$ $R_2^A = \{(\{a\}, \{b,c\}), (\{a\}, \{d,e\})\}$ deg_R({a}) = 2. Admissible. Cost: 1 commitment internalised.

Path B: Merge c, d, e (keeping b separate): $P_2^B = \{\{a\}, \{b\}, \{c,d,e\}\}$ $R_2^B = \{(\{a\}, \{b\}), (\{a\}, \{c,d,e\}), (\{b\}, \{c,d,e\})\}$ deg_R({a}) = 2. Admissible. Cost: 0 commitments internalised (but more identity merging).

Under L_1 (minimal identity merging), Path A is preferred. Under L_2 (minimal commitment internalisation), Path B is preferred. Both are admissible; the choice of L determines which is canonical.

Result: The system canonically determines an identity assignment based on ($\mathcal{C}$, L). No threshold was chosen; no clustering objective was imported; no probabilistic cutoff was applied. The collapse was forced by capacity constraints and resolved by the loss functional.

10.6 New Result: Stability Radius for Identity Assignments

The FCM formulation yields a theorem that goes beyond semantic reinterpretation.

Theorem 10.3 (Stability Radius for Identity Assignments). Let K = (P, R) be admissible. For each block B ∈ P, define the margin:

$m(B) := \mathcal{C}(B) - deg\_R(B),\ m\_min := min\_\{B \in P\}\ m(B)$

Let R_new be any set of new external commitments added to R such that for every B ∈ P:

$deg\_\{R \cup R\_new\}(B) - deg\_R(B) \leq m\_min$

Then K_updated = (P, R ∪ R_new) remains admissible, hence Π_L(K_updated) = {K_updated} and the partition P is unchanged.

Proof. An admissible structure K = (P, R) satisfies deg_R(B) ≤ $\mathcal{C}$(B) for all B ∈ P. By assumption, for every block B:

$deg\_\{R \cup R\_new\}(B) \leq deg\_R(B) + m\_min \leq deg\_R(B) + (\mathcal{C}(B) - deg\_R(B)) = \mathcal{C}(B)$

Hence K_updated remains admissible. By Lemma 3.2, Π_L(K_updated) = {K_updated}, so no collapse occurs and the partition P is preserved.

Remark on novelty. The stability criterion itself is elementary; what is new is that FCM makes it a certified consequence of the admissibility axioms and ties it to computable margins rather than algorithm-specific thresholds.

Corollary (Streaming Stability Certificate). Given an admissible identity assignment K, compute m_min. Any stream of updates that maintains the invariant "no block's degree increases by more than m_min" preserves the identity assignment without requiring re-resolution.

Significance. This is a hard, checkable robustness theorem. It tells practitioners exactly when identity assignments are stable under streaming updates, without rerunning the resolution algorithm. The margin m_min is computable from the current state and provides a concrete stability certificate.

Example. In the worked example (10.5), suppose we reach the admissible state:

P = {{a}, {b}, {c,d,e}}, with deg_R({a}) = 2, deg_R({b}) = 2, deg_R({c,d,e}) = 2

If $\mathcal{C}$(B) = 2 for all blocks, then m(B) = 0 for all blocks, so m_min = 0. Any new commitment incident to any block will trigger collapse. But if instead $\mathcal{C}$(B) = 3, then m_min = 1, and the system can absorb one new commitment per block before identity must change.

10.7 Significance

FCM resolves a known problem—entity resolution under persistent constraints—without introducing external decision rules. The identity outcome is determined by the axioms, not by human choices about thresholds, clustering objectives, or probabilistic cutoffs.

This demonstrates *problem-relative usefulness*: FCM is not merely foundationally interesting but provides cleaner solutions to problems that existing mathematics handles awkwardly.

---

11. Application: Tensor-Network Truncation as Finite Commitment Collapse

To demonstrate that FCM's scope extends beyond data systems into computational physics, we show that tensor-network truncation—a standard technique in quantum many-body simulation—is naturally expressed as finite commitment collapse.

11.1 Problem Statement: Truncation Under Finite Bond Dimension

Tensor-network methods (e.g., MPS, PEPS, MERA) represent quantum many-body states using structured factorisations. Exact representations typically require bond dimensions that grow exponentially with system size. In practice, one imposes a finite bond-dimension cutoff χ, discarding degrees of freedom to remain computationally tractable.

This gives rise to a well-known problem:

How should degrees of freedom be merged or discarded when representational capacity is exceeded, and in what sense is the resulting approximation canonical?

Existing tensor-network algorithms answer this procedurally (e.g., truncate small singular values), but the truncation step is external to the mathematical structure of the state representation.

11.2 How Existing Tools Handle Truncation

Current tensor-network methods resolve capacity overload by externally specified projection rules:

- **SVD truncation**: retain the largest χ singular values
- **Energy-based truncation**: preserve low-energy subspaces
- **Entropy-based criteria**: retain modes contributing most to entanglement
- **Algorithm-specific prescriptions**: TRG vs TNR vs MERA variants

These approaches work well empirically, but they share three features:

1. Resolution loss follows an algorithmically specified criterion chosen externally to the state representation
2. Multiple inequivalent truncations exist for the same bond-dimension cutoff
3. No canonical equivalence class of approximations is defined beyond algorithm-specific choices

In other words, the mathematics does not force a unique outcome when capacity is exceeded.

11.3 FCM Reformulation

Tensor-network truncation can be expressed as an instance of finite commitment collapse.

Entities and Commitments:

- **Entities E**: distinguishable internal modes or Schmidt components across a bipartition
- **Partition P**: grouping of modes into effective degrees of freedom retained by the network
- **Commitments R**: operational distinction demands induced by the target fidelity criterion—i.e., the requirement to keep certain mode contributions distinguishable to preserve a chosen set of quantities (e.g., state norm, reduced density matrices up to a given rank, or expectation values of a fixed operator set). These demands persist as constraints, but may become internalised under collapse when representational capacity is exceeded.
- **Capacity $\mathcal{C}$**: maximum number of distinct modes a bond can support, identified with bond dimension χ
- **Loss functional L**: specifies what is prioritised during truncation (preserve norm, preserve entanglement entropy, preserve local observables)

Collapse Interpretation:

When the number of distinct modes incident to a bond exceeds χ, the structure becomes inadmissible:

$$\deg_R(B) > \mathcal{C}(B) = \chi$$

By Axioms A1–A3, admissibility must be restored by identity collapse: modes are merged into coarser identity blocks until capacity constraints are satisfied.

The truncation step is therefore not an arbitrary deletion of modes, but a canonical projection:

$$\Pi_L(K) = \text{minimal admissible coarsening under loss } L$$

11.4 What FCM Adds Beyond Existing Tensor-Network Formalism

(A) Canonical Semantics for Truncation.

Given a fixed bond dimension χ and loss functional L, FCM assigns a canonical equivalence class of admissible truncations.

Different tensor-network algorithms correspond to different choices of L, but within each choice the collapse outcome is canonical up to the minimiser class.

This replaces "We truncated using algorithm X" with "We projected to the admissible identity structure under capacity χ with loss L."

Remark. The choice of L is not hidden arbitrariness; it is the same structural relocation of choice that variational principles perform when replacing dynamical rules with action functionals.

(B) Structured Non-Uniqueness.

FCM clarifies why multiple truncation schemes exist: non-uniqueness is not algorithmic ambiguity but reflects distinct universality classes indexed by L.

This mirrors Appendix B's universality-class analysis: SVD truncation and entropy-preserving truncation are not competing heuristics but different collapse priorities.

(C) Collapse-Convergence Under Iterative Coarse-Graining.

Renormalisation procedures repeatedly truncate after coarse-graining. By Theorem 8.3, for fixed ($\mathcal{C}$, L), the sequence of projected identity structures is collapse-convergent:

$$\Pi_L(K_t) = \Pi_L(K_{t+1}) \text{ for } t \geq N$$

This provides a clean explanation for why iterative tensor-network methods reach stable fixed points despite repeated information loss.

11.5 Toy Example: Schmidt Mode Truncation

Consider a bipartition with Schmidt decomposition:

$|\psi\rangle = \Sigma_{i=1}^{n} \lambda_i |i_l\rangle|i_r\rangle$

with $n > \chi$.

- **Entities**: Schmidt pairs i
- **Commitments**: operational demands to preserve contributions from each mode
- **Capacity**: $\mathcal{C} = \chi$

The standard truncation keeps the largest χ coefficients.

A canonical loss functional for norm preservation. Let $w_i := \lambda_i^2$ be the squared Schmidt weights. Define L_norm(P, P′) to be the total weight internalised into the remainder block (i.e., the sum of w_i over modes that are no longer externally distinguished after coarsening). Minimising L_norm under capacity χ retains the χ largest weights, matching standard truncated SVD.

In FCM terms:

- Loss functional L_norm penalises loss of squared norm as defined above
- Collapse merges the $n - \chi$ smallest modes into a single internal block
- The distinction that "these modes were once separate" persists internally but is no longer externally represented

Formally, "discarding" corresponds here to collapsing all sub-threshold modes into an undifferentiated remainder block that is not externally addressable by the truncated network. This matches the operational effect of truncation: the theory no longer distinguishes those modes individually, and no provenance of their separate contributions is retained at the representational level.

Thus standard SVD truncation is precisely $\Pi_{\{L\_norm\}}(K)$.

Other truncation criteria correspond to different L.

11.6 New Results: Truncation-Error Certificates

The FCM formulation yields theorems that go beyond reinterpretation of existing practice.

Theorem 11.1 (Path-Independent Truncation-Error Certificate). Consider an iterative pipeline producing normalised states $|\psi_0\rangle, |\psi_1\rangle, ..., |\psi\_T\rangle$, where each step applies a truncation enforcing capacity χ. Let Δ_t denote the discarded squared Schmidt weight at step t, i.e., the mass internalised into the remainder block under L_norm (as defined in 11.5). Assume each truncation realises Π_{L_norm} at that step. Then:

$$\| |\psi_0\rangle - |\psi\_T\rangle \| \leq \Sigma_{t=1}^{T} \sqrt{(2\Delta_t)}$$

Moreover, among all admissible truncations satisfying the same per-step capacity constraints, the FCM choice minimises Δ_t at each step, hence yields the smallest bound of this form.

Proof sketch. For a single truncation that discards squared weight Δ, the truncated and renormalised state satisfies $\| |\psi\rangle - |\tilde{\psi}\rangle \| \leq \sqrt{(2\Delta)}$ (standard). Apply triangle inequality over steps:

$$\| |\psi_0\rangle - |\psi\_T\rangle \| \leq \Sigma_t \| |\psi_{t-1}\rangle - |\psi_t\rangle \| \leq \Sigma_t \sqrt{(2\Delta_t)}$$

Pointwise minimality of Δ_t follows from the definition of Π_{L_norm}.

Remark on novelty. The single-step inequality $\|\psi - \tilde{\psi}\| \leq \sqrt{(2\Delta)}$ is standard truncation geometry. What is new here is that FCM makes it a certified consequence of an internal projection operator, ties the bound to collapse losses rather than algorithm-specific truncation steps, and guarantees minimality through the collapse axioms.

Significance. FCM turns iterative truncation into a certified process whose error bound is computed directly from collapse losses. The bound is independent of algorithmic details beyond "this step realises Π_{L_norm}."

Theorem 11.2 (Multi-Bond Optimality in the FCM Truncation Model). Consider a multi-bond truncation model in which each bond e has an associated set of distinguishable mode-entities E_e, capacity χ_e, and discarded-weight loss L_norm^(e). Let the global carrier be the disjoint union $E := \sqcup_e E_e$, and let the global loss functional be:

L(P, P′) := $\Sigma_e\, \alpha_e$ L_norm^(e)(P|*{E_e}, P′*|{E_e})

with weights $\alpha_e > 0$.

Then Π_L(K) yields an admissible coarsening that minimises the weighted total discarded weight among all admissible coarsenings in this model. In particular, no other admissible coarsening can reduce the discarded weight on any bond e without increasing it on some bond e′ (Pareto optimality with respect to (α_e)).

Proof. By definition of Π_L as the set of L-minimising admissible coarsenings. Pareto optimality follows: if some admissible K′ improved on bond e without worsening others, it would have strictly lower L(K′) < L(Π_L(K)), contradicting minimality.

Remark on novelty. The single-cut "keep the top χ" result is classical. What is new is that FCM provides a canonical notion of multi-bond global optimality as a projection property, giving a well-posed target even when the optimisation is computationally hard.

Significance. The multi-bond globally constrained statement provides a canonical notion of optimality that tensor-network practice does not currently have in unified form. FCM supplies the axiomatic foundation.

11.7 Significance

FCM does not replace tensor-network algorithms. It provides a complementary perspective:

1. An explicit foundation-level semantics for truncation
2. A principled way to compare truncation schemes
3. A language for discussing approximation without invoking external choice at the moment of overload

In particular, FCM makes precise the sense in which "truncation is forced by finite distinguishability, not chosen by algorithmic taste."

Tensor-network truncation is already an instance of identity-fluid mathematics in practice. FCM does not introduce collapse into quantum computing; it makes explicit and axiomatic the collapse that is already occurring.

This demonstrates problem-relative usefulness: FCM provides an explicit axiomatic account of identity collapse that is implicit in existing truncation practice, offering canonical projection semantics for capacity-limited representations—without altering the underlying physics or algorithms.

Tensor-network truncation shows that identity-fluid mathematics is not speculative: it is already embedded in successful quantum-computing techniques. FCM provides an axiomatic identity-collapse semantics for truncation, complementary to variational and algorithmic accounts, clarifying what is preserved, what is lost, and why the result is stable.

---

12. Conclusion

FCM introduces a new class of finitary mathematical structures—capacity-regulated commitment structures equipped with an identity-changing collapse projection—and proves that the resulting dynamics cannot be faithfully internalised within standard representational frameworks (categories with total composition, quotient constructions preserving provenance, confluent rewriting systems, or refinement-monotone domain semantics) without adding meta-level rules that are not derivable from the axioms. In this precise, foundation-theoretic sense of non-embeddability, FCM constitutes new mathematics.

Its core contributions are:

1. **Identity as regulated, not primitive**: Distinguishability is finite and capacity-bounded

2. **Collapse as canonical projection**: The operator Π_L enforces admissibility with minimal loss
3. **Collapse-convergence**: A notion of limit applicable where classical convergence fails
4. **Structural non-categoricity**: Fundamental failure of composition arising from identity instability

These features are not expressible within classical frameworks without external rules. FCM internalises them as primitive operations.

FCM is motivated by any setting where distinguishability is resource-limited—cognitive, computational, or physical—so that maintaining too many simultaneous distinctions forces resolution loss.

Remark. FCM is developed here as a standalone mathematical framework. In related work on physical and informational foundations (VERSF, BCB, TPB), similar admissibility and collapse mechanisms appear in different contexts. Those connections are deliberately omitted here to preserve mathematical independence.

The Bottom Line (General Reader)

If you've made it this far, here's what to take away:

What FCM proposes: Identity isn't fundamental. It's a resource that can be exhausted. When you try to maintain too many distinctions, reality pushes back by *merging things together*.

What FCM delivers mathematically:

- A precise definition of "too many distinctions" (capacity bounds)
- A principled way to restore balance (collapse projection)
- A notion of convergence that works even when identity is unstable
- A proof that standard mathematical tools (categories) don't apply—and why

Why it might matter: Many hard problems in physics (quantum measurement), philosophy (personal identity), and AI (concept formation) involve situations where identity seems to change or merge. FCM provides rigorous tools for thinking about such situations, without requiring infinity, contradiction, or hand-waving.

Whether these tools illuminate real phenomena is an empirical question. That they constitute a coherent mathematical framework is what this paper establishes.

Appendix A: Canonical Worked Example of Collapse Projection

For general readers: This appendix walks through a complete example, step by step. No prior knowledge beyond the main text is assumed. If you want to see exactly how collapse works with actual numbers, this is the place.

This appendix provides a fully explicit, minimal worked example of the collapse projection Π_L. Every object is finite, every operation is defined, and identity change is visible.

A.1 Carrier and Initial Structure

Let the carrier set be:

$E = \{a, b, c, d, e\}$

Initial identity partition:

$P_0 = \{\{a\}, \{b\}, \{c\}, \{d\}, \{e\}\}$

All entities are initially distinguishable.

Initial commitments:

$R_0 = \{(a, b), (a, c), (a, d), (a, e)\}$

These commitments express distinctions between block {a} and all others.

A.2 Capacity Functional

Define a uniform capacity functional:

$\mathcal{C}(B) = 2$ for all blocks B

No identity block may participate in more than two external commitments.

A.3 Inadmissibility

In $K_0 = (P_0, R_0)$, the block {a} participates in four commitments:

$\deg_R(\{a\}) = 4 > \mathcal{C}(\{a\}) = 2$

Therefore K_0 is inadmissible.

A.4 Loss Functional

Define the loss functional L as the number of merged singleton blocks:

$L(P, P') = |P| - |P'|$

This measures resolution loss as the reduction in identity classes.

A.5 First Collapse Step

To restore admissibility, block {a} must reduce its external degree. The minimal-loss solution is to merge {a} with exactly one other block.

All single-merge choices achieve equal loss $L(P_0, P_1) = 1$, but further collapse is required for admissibility. We select one representative for concreteness. The resulting non-uniqueness (multiple minimisers under the same L) is treated systematically in Appendix B. Choose:

$P_1 = \{\{a, b\}, \{c\}, \{d\}, \{e\}\}$

Commitment (a, b) becomes internal and is removed from R:

$R_1 = \{(\{a,b\}, \{c\}), (\{a,b\}, \{d\}), (\{a,b\}, \{e\})\}$

Now:

$deg\_R(\{a, b\}) = 3 > \mathcal{C}(\{a, b\}) = 2$

Still inadmissible.

A.6 Second Collapse Step

A second merge is required. Minimal additional loss yields:

$P_2 = \{\{a, b, c\}, \{d\}, \{e\}\}$

Remaining commitments:

$R_2 = \{(\{a,b,c\}, \{d\}), (\{a,b,c\}, \{e\})\}$

Now:

$deg\_R(\{a, b, c\}) = 2 = \mathcal{C}(\{a, b, c\})$

Structure $K_2 = (P_2, R_2)$ is admissible.

A.7 Result and Properties

Thus:

$\Pi\_L(K_0) = \{K_2\} = \{(\{\{a,b,c\}, \{d\}, \{e\}\}, \{(\{a,b,c\}, \{d\}), (\{a,b,c\}, \{e\})\})\}$

Properties illustrated:

- Identity has changed irreversibly
- Collapse is idempotent: $\Pi\_L(K_2) = \{K_2\}$
- Loss is minimal under L
- No contradictions or infinities appear

A.8 Significance

This example cannot be expressed internally in classical mathematics without external rules governing when and how to coarsen. In FCM, identity change is the result of a canonical operator acting within the system.

Why This Matters (General Reader)

In ordinary mathematics, if you start with five distinct objects, you end with five distinct objects. You can *group* them, *label* them, or *ignore* some—but they remain five separate things underneath.

This example shows something different. We started with five entities (a, b, c, d, e) and ended with three identity blocks ({a,b,c}, {d}, {e}). The entities a, b, and c didn't just get grouped together—they *became the same thing*. There's no hidden variable tracking their original distinctness. The information is gone.

This is what makes FCM different from a filing system or a database with grouped records. Collapse isn't organisation—it's ontological change.

Appendix B: Non-Uniqueness and Universality Classes

For general readers: You might worry that collapse is arbitrary—that different choices lead to different answers with no way to decide between them. This appendix shows that the non-uniqueness is structured and meaningful, not chaotic. Different priorities lead to different collapses, but within each priority system, everything is well-defined.

This appendix demonstrates that while Π_L is canonical for a chosen loss functional L, different choices of L lead to systematically different collapse outcomes defining universality classes.

B.1 Purpose

A common objection to identity-changing projections is underdetermination. FCM addresses this by showing that non-uniqueness is structured: different loss functionals correspond to different but well-defined collapse behaviours, analogous to:

- Choosing different norms in optimisation
- Different renormalisation schemes in physics
- Different metrics inducing different topologies

B.2 Setup

Reusing the carrier from Appendix A:

$E = \{a, b, c, d, e\}$ $P_0 = \{\{a\}, \{b\}, \{c\}, \{d\}, \{e\}\}$ $R_0 = \{(a, b), (a, c), (a, d), (a, e)\}$ $\mathcal{C}(B) = 2$ for all B

B.3 Loss Functional L_1: Minimal Identity Merging

$L_1(P, P') = |P| - |P'|$

This penalises identity loss directly. As shown in Appendix A:

$\Pi\_\{L_1\}(K_0)$ yields $P_2 = \{\{a, b, c\}, \{d\}, \{e\}\}$

Three identity blocks preserved.

B.4 Loss Functional L_2: Minimal Commitment Loss

$L_2(P, P')$ = number of commitments internalised

This prioritises relational structure over identity resolution.

Under L_2, the preferred collapse is:

$P' = \{\{a, b, c, d\}, \{e\}\}$

Merging four elements internalises three commitments simultaneously, leaving one external commitment ({a,b,c,d}, e) satisfying capacity.

B.5 Comparison

Loss Functional	Collapse Result	Blocks Remaining	Commitments Remaining
L_1 (identity)	{{a,b,c}, {d}, {e}}	3	2
L_2 (commitment)	{{a,b,c,d}, {e}}	2	1

Both are admissible. Both satisfy irreversibility. Both are canonical relative to their loss functionals.

B.6 Universality Classes

FCM does not define a single collapse outcome, but a family of universality classes indexed by loss functional choice. Within each class:

- Π_L is idempotent
- Collapse-convergence is well-defined
- Qualitative behaviour is stable under perturbations

Across classes, qualitative behaviour differs in controlled, explainable ways.

B.7 Conclusion

Structured non-uniqueness is a strength. It allows FCM to model different resolution priorities while maintaining mathematical rigour. Non-uniqueness is internal, explicit, and governed—not arbitrary.

Appendix C: Collapse-Convergence in an Iterative Process

For general readers: This appendix shows something surprising—a process that bounces back and forth forever can still have a meaningful "answer" once we apply collapse. It's like a heated argument that never ends, but where the participants eventually become the same person and the argument becomes moot.

This appendix demonstrates that an iterative process divergent in classical mathematics can possess a well-defined collapse-convergent limit.

C.1 Motivation

Classical convergence requires fixed identity and metric structure. FCM's collapse-convergence applies where these fail.

C.2 Update Semantics

Remark (Update semantics after collapse). Update rules are specified on underlying entities (e.g., "add a commitment between a and b"). When applied to a structure K = (P, R), such a rule is interpreted as adding a commitment between the current blocks containing those entities: add (B_a, B_b) where B_a, B_b ∈ P are the blocks with a ∈ B_a, b ∈ B_b. If B_a = B_b, the attempted commitment is vacuous (it would be internal) and has no effect. If the commitment (B_a, B_b) already exists in R, the update has no effect (R is a set).

This semantics ensures that updates remain well-defined after collapse, even when the named entities no longer form singleton blocks.

C.3 Setup

$E = \{a, b, c, d\}$ $P_0 = \{\{a\}, \{b\}, \{c\}, \{d\}\}$ $\mathcal{C}(B) = 1$ for all B

Each block may participate in at most one external commitment.

C.4 Oscillatory Update Rule

Define an update process alternating commitment attempts:

- Step 1: add (a, b)
- Step 2: add (a, c)
- Step 3: add (a, b)
- Step 4: add (a, c)
- ...

Classically, this has no convergent limit—the structure continually attempts incompatible distinctions.

C.5 Collapse Projection at Each Step

After Step 1: $R_1 = \{(a, b)\}$ — admissible ($\deg(\{a\}) = 1 = \mathcal{C}(\{a\})$)

After Step 2: $R_2 = \{(a, b), (a, c)\}$ — inadmissible ($\deg(\{a\}) = 2 > 1$)

Collapse merges {a} with one neighbour. Under minimal identity loss:

$P_2^* = \{\{a, b\}, \{c\}, \{d\}\}$ $R_2^* = \{(\{a,b\}, \{c\})\}$

Now $\deg(\{a,b\}) = 1$, admissible.

C.6 Subsequent Iterations

Further steps attempt to reintroduce distinctions involving 'a'. However, 'a' is now internal to {a, b}. Attempts to distinguish {a} from {c} are either redundant or trigger no further collapse.

For all $t \geq 2$:

$\Pi\_L(K\_t) = (\{\{a, b\}, \{c\}, \{d\}\}, \{(\{a,b\}, \{c\})\})$

C.7 Collapse-Convergent Limit

The unprojected sequence oscillates indefinitely. The projected sequence stabilises at $t = 2$.

The process is collapse-convergent with limit:

$K^* = (\{\{a, b\}, \{c\}, \{d\}\}, \{(\{a,b\}, \{c\})\})$

C.8 Significance

FCM assigns well-defined limits to processes that are classically divergent. Collapse-convergence depends on admissibility rather than metrics, making it applicable to identity-changing systems.

Appendix D: Explicit Failure of Categorical Structure

For general readers: Category theory is the standard toolkit mathematicians use to organise structures. This appendix shows, with a concrete example, why that toolkit breaks when applied to FCM. It's not that we haven't tried hard enough—it's that the very notion of "chaining operations" doesn't survive identity collapse.

This appendix provides the detailed counterexample for Theorem 5.1.

D.1 Attempted Category

Objects: admissible commitment structures K Morphisms: $f : K \to K'$ where $K' \in \Pi\_L(f(K))$ for some update f

D.2 Failure of Composition

Let $E = \{a, b, c\}$ with $\mathcal{C}(B) = 1$ for all blocks.

Initial structure: $K_1 = (\{\{a\}, \{b\}, \{c\}\}, \emptyset)$

Morphism f: add commitment (a, b)

After f and collapse: $K_2 = (\{\{a, b\}, \{c\}\}, \emptyset)$

Morphism g: add commitment (a, c)

The update g is not well-defined on K_2: its domain presupposes the existence of a block containing a as a singleton, but collapse has replaced that block with the strictly coarser block {a, b}.

Therefore $g \circ f$ is undefined. Composition is not closed.

D.3 Failure of Associativity

Even where partial composition exists, let h add commitment (b, c).

The order of collapse affects which blocks exist:

$$\Pi\_L(h(\Pi\_L(g(f(K_1))))) \neq \Pi\_L((h \circ g)(\Pi\_L(f(K_1))))$$

because the domain of h depends on prior collapse outcomes.

D.4 Information-Destroying Collapse

One might propose tracking provenance through collapse. FCM explicitly disallows this: collapse is information-destroying, not merely relabelling.

If provenance were retained:

- Collapsed blocks would carry internal structure
- This structure would reintroduce the distinctions collapse eliminated
- Admissibility would be violated or collapse would become reversible

Axiom A2 (Persistence of Commitment) enforces this: collapse cannot be undone by "remembering" what merged.

Proposition D.1 (No Faithful Embedding into Fixed-Object Categories). There is no faithful functor from the partial-composition structure induced by collapse-projected updates (where some compositions are undefined) into any category in which composition is total.

Proof. In any category, if $f : A \to B$ and $g : B \to C$ exist, then $g \circ f$ exists. In FCM, there exist updates f and g such that f and g are individually well-defined but $g \circ f$ is undefined (Theorem 5.1). A functor must preserve sources/targets and preserve composition wherever defined. But mapping both f and g into a category forces the existence of a composite image $F(g) \circ F(f)$, contradicting preservation of undefined composition. Hence no faithful functor exists. ■

Remark (Relation to Rewriting Systems). FCM collapse superficially resembles term rewriting, but differs in two essential ways. First, rewriting systems are typically confluent (different reduction paths reach the same normal form); FCM collapse depends on loss functional choice and is deliberately non-confluent across universality classes. Second, rewriting preserves the identity of term positions; FCM collapse destroys identity, making "the same position after rewriting" undefined. FCM is thus not a rewriting system in the standard sense.

General Reader Explanation

Category theory is the mathematician's universal language for structure. Almost everything in modern mathematics—sets, groups, topological spaces, programming languages—can be organised into categories where you have objects, arrows between them (morphisms), and the ability to chain arrows together (composition).

FCM doesn't fit. Here's why, in intuitive terms:

Suppose you have three people: Alice, Bob, and Carol. You perform an operation that merges Alice and Bob into a single entity "AliceBob." Now you want to perform another operation that was defined in terms of Alice alone—say, "distinguish Alice from Carol."

But Alice doesn't exist anymore. She's part of AliceBob now. The second operation isn't just difficult to perform—it's *undefined*. Its target has ceased to exist as an independent thing.

In a category, you must always be able to chain operations: if f goes from A to B, and g goes from B to C, then g∘f goes from A to C. But in FCM, performing f can *change what B is*, making g undefined. The chain breaks.

This isn't a bug in FCM—it's the whole point. We're modelling systems where identity genuinely changes, and that means the usual algebraic tools don't apply. We need different tools, which is what collapse-convergence and universality classes provide.

D.5 Conclusion

FCM commitment structures do not form a category (Theorem 5.1, this appendix), preorder, or lattice (Proposition 7.7). This is not a deficiency but a reflection of the framework's core premise: identity instability is fundamental.

Notation Summary

Symbol	Meaning
E	Finite carrier set of primitive entities
$K = (P, R)$	Commitment structure: partition P and commitments R
$B \in P$	Identity block (element of partition)
$P' \preceq P$	P' is coarser than P (every block of P' is a union of blocks of P)
$\pi_{P \to P'}$	Block map under coarsening: $\pi(B)$ = unique $B^* \in P'$ with $B \subseteq B^*$
R'	Pushforward of R under coarsening (external commitments that survive)
$\mathcal{C}(B)$	Capacity of block B
$\deg_R(B)$	Degree of B (number of external commitments)
$L(P, P')$	Loss functional measuring coarsening cost
$\Pi_L(K)$	Collapse projection (set of minimal-loss admissible coarsenings)
$K \trianglelefteq K'$	Collapse-step relation: $K' \in \Pi_L(K)$
$\sigma_M(x)$	Observational signature under test family M
$x \sim_M y$	Operational indistinguishability under M
$m(B)$	Admissibility margin: $\mathcal{C}(B) - \deg_R(B)$
Δ_t	Per-step discarded weight in iterative truncation