Coalescence-Class Cardinalities in Finite Grand Couplings Need Not Be Deterministic

An explicit family and a computer-assisted state-minimality theorem

Luke Kabbash | GPT-5.6 Sol

20 July 2026

Abstract.

Let $F_1, F_2, \ldots$ be independent random self-maps of a finite set, and apply each map simultaneously to trajectories started from every state. Grimmett and Holmes asked whether, for a fixed consistent grand coupling, the cardinalities of the eventual coalescence classes must be deterministic. We answer in the negative. For every $m \geq 1$ and $0 < p < 1$, we construct a two-map coupling on $4m+2$ states whose terminal profile is $(2m, 2m+2)$ with probability $1/(2-p)$ and $(2m+1, 2m+1)$ with probability $(1-p)/(2-p)$. The induced Markov chain is irreducible and aperiodic, and exactly two classes remain almost surely. The proof is symbolic. A separate exhaustive kernel-graph census establishes, with computer assistance, that no example exists on fewer than six states; thus the six-state instance is state-minimal without a restriction on support size. Reproducible Python and C++ implementations accompany the manuscript.

Keywords. grand coupling; random mapping representation; coalescence; transformation monoid; kernel graph; synchronizing automaton; computer-assisted proof

Mathematics Subject Classification (2020). Primary 60J10; secondary 20M20, 05C15, 68Q45.

Theorem (six-state counterexample).

The answer to Open Problem 3.11 of Grimmett and Holmes is negative: on six states, let $A=(0,1,0,0,1,0)$ and $C=(3,4,5,1,2,1)$, and choose each successive map independently to be A or C with probability 1/2 each; the induced Markov chain is irreducible and aperiodic, and the grand coupling has exactly two terminal coalescence classes almost surely, with size profile (2,4) with probability 2/3 and (3,3) with probability 1/3.

---

1. Introduction

A random mapping representation writes a finite Markov chain as repeated application of random self-maps. When one map is applied to every initial state at each time, the resulting process is a grand coupling: trajectories that meet remain together. This representation is classical in the study of iterated random functions [2], synchronization, and exact sampling [9].

Grimmett and Holmes [1, Open Problem 3.11] ask whether the multiset of cardinalities of the forward coalescence classes is deterministic. Their number is almost surely constant, but the classes themselves may be random. Their four-state example randomizes the paired labels while leaving the size profile equal to (2,2) in every outcome. The question is whether the sizes can also vary under one fixed transition matrix and one fixed map law.

They can. The counterexample below uses two maps on six states and extends to a family on every state count $4m+2$. Its mechanism is elementary: one map collapses two odd paths according to vertex parity, while a second map changes the relative parity of the paths. The first occurrence of the collapsing map fixes a rank-two kernel forever. The parity of the preceding run chooses between an unequal and an equal pair of block sizes.

The construction lies at the intersection of finite random dynamical systems and transformation semigroups. Minimum-rank maps form the minimal ideal of a finite transformation monoid [3]. Kernel graphs connect non-synchronizing monoids with graph colorings and endomorphisms [4-6], while the integer partition of kernel-block sizes is already known as kernel type [7]. The contribution claimed here is narrower: an explicit negative answer to the probability question in [1], its exact infinite family and terminal law, and a computer-assisted state-minimality result.

A targeted search of exact phrases, the source title and DOI, arXiv, publisher records, and author publication pages through 20 July 2026 found no public resolution of Open Problem 3.11. A negative search is not proof of priority. The complete reproduction pipeline passes, but the manuscript has not received independent peer review.

2. Setting and notation

Let S be finite, let $Map(S)=S^S$ be the set of self-maps of S, and let μ be a probability measure on Map(S). Let $F_1, F_2, \ldots$ be i.i.d. with law μ. Products are forward products

$$Q_t = F_t F_{t-1} \cdots F_1, \quad \text{where} \quad (fg)(x)=f(g(x)).$$

The law μ is consistent with a transition matrix P when

$$P(x,y) = \mu\{f \in Map(S) : f(x)=y\}.$$

Two states x and y coalesce if $Q_t(x)=Q_t(y)$ for some t. The common maps then keep them together. For a map f, write ker(f) for its partition into fibers and rank(f)=|f(S)|. The *profile* of f is the sorted vector of its fiber cardinalities. The terminal profile is the profile of the eventual coalescence partition.

Lemma 1 (kernel monotonicity).

For self-maps f and q, ker(q) refines ker(fq). Consequently, a forward product can merge kernel blocks but cannot split them. If rank(q) is already the minimum rank of the support-generated monoid, then every later product has the same kernel as q.

Proof. If q(x)=q(y), then f(q(x))=f(q(y)). Thus every block of ker(q) lies in a block of ker(fq). At minimum rank, a strict merger would decrease the number of blocks and contradict minimality.

3. The counterexample family

3.1 Construction

Fix m ≥ 1 and set d=2m+1. Take two disjoint paths

$$L_0\text{-}L_1\text{-}\cdots\text{-}L_{2m} \quad \text{and} \quad R_0\text{-}R_1\text{-}\cdots\text{-}R_{2m}.$$

Their 4m+2 vertices form S_m. Define A and C by

$$A(L_i)=A(R_i)=L_{i \bmod 2},$$

$$C(L_i)=R_i, \quad C(R_i)=L_{i+1} \text{ for } i<2m, \quad \text{and} \quad C(R_{2m})=L_{2m-1}.$$

Let $\mu=p\delta_A+(1-p)\delta_C$, where 0 < p < 1, and let P be the consistent transition matrix induced by μ.

Theorem 1 (counterexample family).

For every integer $m \geq 1$ and $p \in (0,1)$, the pair (P,μ) above is a consistent grand coupling of an irreducible, aperiodic Markov chain on $4m+2$ states. Its forward coalescence partition has two blocks almost surely. Its sorted block-size vector is

$(2m, 2m+2)$ with probability $1/(2-p)$, and $(2m+1, 2m+1)$ with probability $(1-p)/(2-p)$.

Both outcomes therefore have positive probability under the same P and μ.

3.2 Proof of Theorem 1

Proof. Consistency is immediate from $P(x,y)=p\,1\{A(x)=y\}+(1-p)\,1\{C(x)=y\}$. The C-orbit of L0 visits every vertex before entering the final four-cycle. From any vertex, C reaches an even-indexed vertex in that cycle, and A sends it to L0. Thus the positive-transition digraph is strongly connected.

Since A(L0)=L0 with positive probability, the irreducible chain has a self-loop and is aperiodic.

Both A and C are endomorphisms of the loopless graph formed by the two paths. For A this is the parity coloring of a bipartite graph. For C, an edge L_iL_{i+1} maps to R_iR_{i+1}; an edge R_iR_{i+1} maps to $L_{i+1}L_{i+2}$, except at the endpoint, where it maps to $L_{2m}L_{2m-1}$. Every supported word is therefore a graph endomorphism. A constant map would send an edge to a loop, so every supported word has rank at least two.

The first A occurs almost surely. If T copies of C precede it, the prefix product is $A\,C^T$, whose rank is at most two and hence exactly two. All later prefixes are still graph endomorphisms, so their rank cannot fall below two. Lemma 1 then makes the rank-two kernel permanent.

It remains to count its fibers. The map C^2 stays on the same path and reverses index parity at every vertex, including the reflected endpoint. Hence an even power of C changes parity in the same way on both paths. Applying A then gives the original parity counts: each path contains m+1 even and m odd indices, so the sorted profile is (2m,2m+2). An odd power of C puts the two paths in opposite parity phases before A. Each fiber then receives m vertices from one path and m+1 from the other, giving (2m+1,2m+1).

Finally, $\Pr(T=t)=p(1-p)^t$. Summing the even and odd terms gives

$$\Pr(T\text{ even}) = p\,\Sigma_{k\geq 0}(1-p)^{2k} = 1/(2-p),$$
$$\Pr(T\text{ odd}) = p\,\Sigma_{k\geq 0}(1-p)^{2k+1} = (1-p)/(2-p).$$

This proves the theorem.

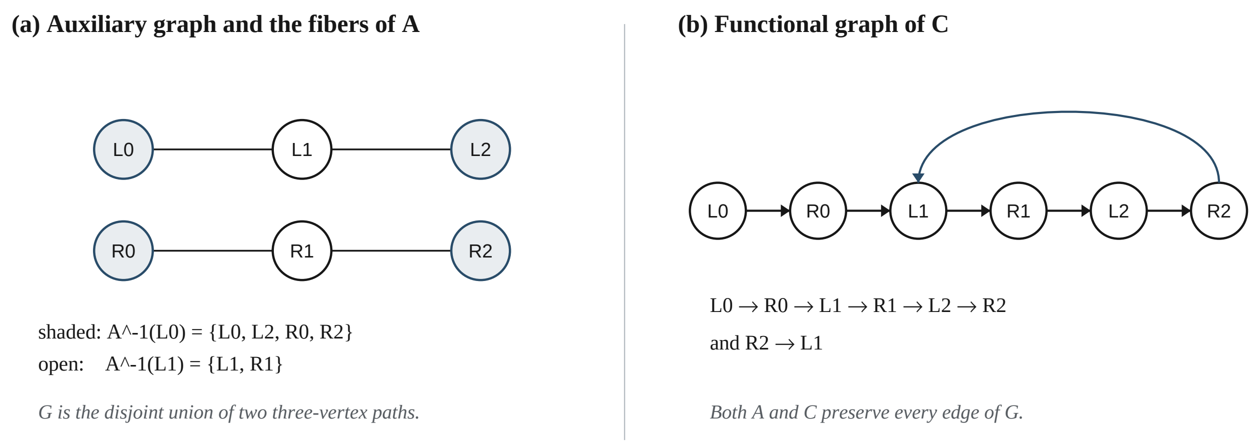

Figure 1. The six-state instance m=1. Panel (a) shows the auxiliary graph and the two fibers of A. Panel (b) shows the directed functional graph of C. The undirected path edges in (a) are proof devices, not Markov transitions. Both maps preserve those edges; because the graph is loopless, no supported word can be constant.

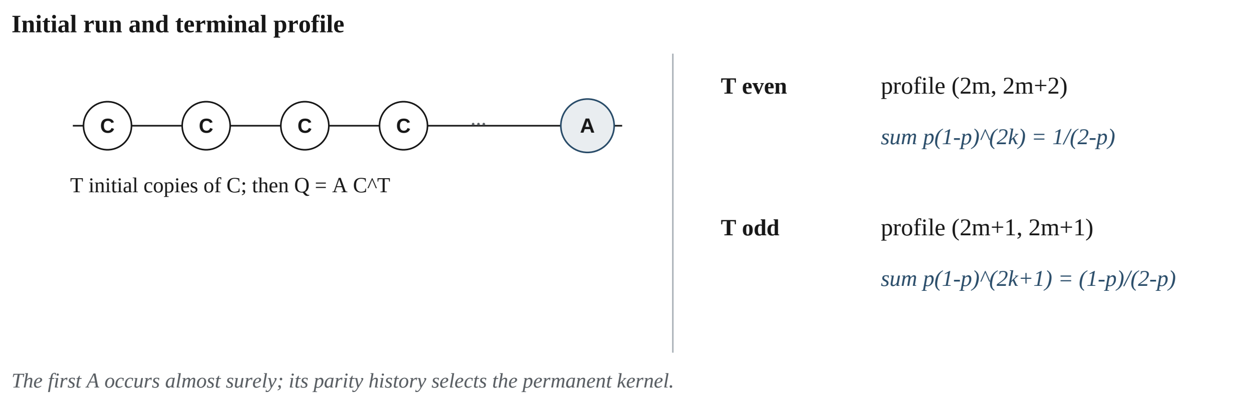

Figure 2. The exact terminal law. Composition is read chronologically from left to right in the displayed draws and algebraically as A C^T. Simulation is unnecessary: the two probabilities are geometric sums.

4. The six-state instance

For m=1, identify (L0,L1,L2,R0,R1,R2) with (0,1,2,3,4,5). The maps are

$$A=(0,1,0,0,1,0), \quad C=(3,4,5,1,2,1).$$

state	L0	L1	L2	R0	R1	R2
A	L0	L1	L0	L0	L1	L0
C	R0	R1	R2	L1	L2	L1

If the first draw is A, the terminal blocks are {L0,L2,R0,R2} and {L1,R1}, with profile (2,4). If the first two draws are C,A, the terminal blocks are {L0,L2,R1} and {L1,R0,R2}, with profile (3,3). These cylinder events have probabilities p and (1-p)p. For p=1/2, the complete terminal law is (2,4) with probability 2/3 and (3,3) with probability 1/3.

Exact closure supplies a compact cross-check. The support-generated semigroup has 21 nonidentity elements; the prefix-product chain has 22 reachable states when the identity is included; its minimum rank is two; and every rank-two kernel is permanent. These computations corroborate the proof but are not used by it.

5. State minimality

Theorem 2 (computer-assisted state minimality).

No irreducible grand coupling on at most five states has a nonconstant terminal block-size profile. Consequently, the six-state construction is state-minimal, with no restriction on the number of maps in the support.

5.1 Reduction to a finite graph census

Let M be the transformation monoid generated by the support of μ, and let r be its minimum rank. Since every support word has positive probability, a fixed minimum-rank word appears almost surely in disjoint i.i.d. blocks. At that occurrence the forward prefix reaches rank r, and Lemma 1 fixes its kernel thereafter. Irreducibility of P is equivalent to transitivity of the action of M on S.

Define the kernel graph Gr(M) on S by joining distinct x and y exactly when no element of M collapses them. Every element of M is an endomorphism of Gr(M). If f has rank r, then f(S) is an r-clique: otherwise a later map could collapse two image points and produce rank below r. The fibers of f form a proper r-coloring. It follows that

$$\omega(\mathrm{Gr}(M)) = \chi(\mathrm{Gr}(M)) = r.$$

If terminal profiles vary, this graph has at least two optimal coloring-size profiles. It also has no isolated vertex. Indeed, $r \geq 2$ gives an edge; if v were isolated, transitivity would map one endpoint of an edge to v, while endomorphism preservation would require an edge incident with v. Finally, End(Gr(M)) is transitive because it contains M. Thus any smaller counterexample would yield a simple graph with all four properties: no isolates, equal clique and chromatic numbers, multiple optimal color profiles, and a transitive endomorphism monoid.

5.2 Exhaustive result

The supplied census enumerates every labelled simple graph on $n \leq 5$ vertices, all set partitions needed for exact chromatic profiles, all vertex subsets needed for the clique number, and all n^n endofunctions needed for the endomorphism monoid. A separately structured implementation recomputes colorings by surjective color assignments, canonicalizes graphs under every relabelling, and independently tests endomorphism transitivity.

order	graphs after coloring filter	isomorphism types	transitive End(G)
1-4	0	0	0
5	150 labelled	3	0
6, rank 2	90 labelled	1	1

At order five, all three surviving isomorphism types have optimal profiles (1,1,3) and (1,2,2), with 48, 42, and 50 endomorphisms respectively; none acts transitively. At order six and rank two, the unique surviving isomorphism type is the disjoint union of two three-vertex paths. It has 144 endomorphisms, a transitive endomorphism monoid, and profiles (2,4) and (3,3). Only the $n \leq 5$ census is needed for the lower bound; the $n=6$ line is a consistency check against the construction.

The theorem is computer-assisted because its last step depends on exhaustive finite enumeration. The mathematical reduction is independent of the program. A redundant C++ search over all unordered two-map supports on four and five states also finds no example, but that narrower search is not the basis for the all-support statement.

6. Computational verification and reproducibility

The release separates theorem-bearing checks from exploration. All mathematical verifiers use only the Python standard library or C++20. The document builders are separate and require ReportLab and pypdf. From the repository root, the full verification is

```
./reproduce.sh --trials 100000
```

The script first checks every file against MANIFEST.sha256. It then runs independent exact Python and C++ checks of the six-state example, exact family closure, two independently organized graph censuses, the exhaustive two-map searches for n=4 and n=5, sanitizer builds, deterministic property sweeps, exact rational stationarity checks, and seeded Monte Carlo smoke tests. A temporary build directory is removed on exit.

claim	decisive basis	status
nonconstant terminal profile	explicit construction and symbolic proof	proved
exact two-atom law	geometric series	proved
six states are minimal	graph reduction and exhaustive census	computer-assisted
implementation agrees	independent Python/C++, sanitizers, properties	verified
publication priority	targeted public search only	not established

The exact six-state programs report 22 reachable prefix states including the identity, 21 nonidentity semigroup elements, minimum rank two, and precisely the profiles (2,4) and (3,3). The family is closed exactly in Python for m=1,...,8 and independently in C++ for m=1,...,24, reaching 98 states. The observed monoid size 20m+2 is asserted over that tested range but is not claimed here as a theorem for all m.

The exhaustive two-map program freezes its complete counts: 32,640 pairs, 10,476 strongly connected pairs, and 1,152 nonsynchronizing pairs at n=4; 4,881,250, 1,277,160, and 60,000 respectively at n=5. A changed count makes the program fail rather than merely print a different result.

Monte Carlo is deliberately subordinate. Its fixed-seed trials compare empirical frequencies with the exact law and fail only as a statistical smoke test. No simulation, floating-point calculation, or large search is needed to establish Theorem 1.

7. Consequences and research directions

7.1 What rank omits

Synchronization records whether rank one is attainable; minimum-rank theory records how many images survive. This example shows that rank does not determine how the original states are distributed among the surviving fibers. The natural additional statistic is the hitting distribution of kernel types when the random prefix product first enters the minimum-rank ideal. In the present family this terminal kernel-type law has two atoms and is explicit.

The distinction concerns cardinality, not stationary block mass. In the transitive setting, Steinberg [3, Theorem 10.4] gives stationary mass 1/r to every block of a minimum-rank kernel. Thus both blocks in the six-state example have stationary mass 1/2 under either profile, even though a uniformly chosen initial label sees different block sizes.

7.2 A concrete computational problem

For a finite support, one can construct the reachable prefix-product chain, identify its minimum-rank states, label those states by kernel type, and solve an absorbing Markov-chain system for the exact hitting probabilities. The result is a coupling profiler that reports more than coalescence number: terminal profile law, time to minimum rank, largest-block distribution, and collision or imbalance costs. This is immediately implementable for small transformation supports; scalable quotienting and complexity bounds remain research problems.

7.3 Coupling design

A transition matrix generally admits many random mapping representations. Their weights satisfy linear marginal constraints, but the induced semigroup and its terminal kernels depend combinatorially on which maps receive positive mass. This suggests an optimization problem: among couplings consistent with a fixed P, minimize time to a target rank, constrain the terminal profile law, or optimize a balance objective. Mixed-integer programming, SAT or SMT search, and column generation over maps are plausible proposal engines; exact closure and rational arithmetic can then certify a candidate. No complexity or performance advantage is established here.

7.4 Limits of immediate application

The construction is a sharp test case for systems driven by common random instructions, including finite automata and replicated-state simulations. It shows that identical marginal dynamics and a fixed number of surviving groups do not force fixed group cardinalities. It does not, by itself, improve coupling from the past, distributed load balancing, or any production protocol. Forward and backward products must be distinguished carefully [1,9].

7.5 Open problems

- Characterize the sets of kernel types that can occur at minimum rank in a transitive transformation monoid.
- For fixed n and minimum rank r, maximize the number or entropy of attainable terminal profiles.
- Characterize transition matrices P for which every consistent grand coupling has a deterministic terminal profile.
- Optimize the terminal kernel-type law over random mapping representations of a fixed P.
- Find a structural, non-computational proof of the obstruction on at most five states.
- Prove or refute the observed monoid-size formula $|M_m^1|=20m+2$ for the family.
- Determine which forward profile statements have valid backward or coupling-from-the-past analogues.

Appendix A. Census method

For each $n \le 5$, enumerate the $2^{n(n-1)/2}$ labelled simple graphs. Discard graphs with isolated vertices. Enumerate all set partitions of the vertex set, retain the proper colorings with the fewest blocks, and record their sorted block-size profiles. Enumerate all vertex subsets to compute the clique number. Retain only graphs with $\omega=\chi$ and at least two optimal profiles. Finally, enumerate all n^n self-maps, retain exactly the graph endomorphisms, and test reachability of every target vertex from every source under their action.

The independent classifier replaces set partitions by surjective color assignments and canonicalizes each graph under all n! vertex permutations. It freezes the labelled multiplicities, canonical codes, endomorphism counts, minimum ranks, minimum profiles, and transitivity results. Agreement is therefore structural, not merely agreement on a final Boolean answer.

Appendix B. Independent checks

The primary six-state Python verifier closes the generated semigroup in both multiplication directions and checks every minimum-rank continuation. The C++ verifier instead performs breadth-first search on the actual prefix-product chain from the identity. The family implementations use separately written constructors and closure logic. Assertions are required: Python refuses optimized execution under -O, and C++ refuses compilation with NDEBUG.

AddressSanitizer and UndefinedBehaviorSanitizer builds are part of the canonical reproduction. Users whose compiler lacks those facilities may pass --no-sanitizers, but the script labels that run as reduced verification. The manifest detects accidental drift; it is not a signature and does not establish authorship or publication priority. A signed Git tag or archival DOI should be added at public release.

Author's note on provenance

This paper is the output of a mostly autonomous experiment in AI-assisted mathematical research. OpenAI's GPT-5.6 Sol searched for a tractable problem, developed the claimed result, wrote the supporting software and verification design, and drafted the manuscript with minimal human intervention. The human project lead initiated the experiment through an OpenAI subscription, prompted occasionally when the process reached stopping points, set disclosure and release standards, and chose to make the result available for review. The shared byline is a credit line naming the human project lead and the model; it does not treat the model as a legal or accountable human author. The work is not presented as primarily human mathematical research or as evidence of conventional mathematical training. Readers should treat the result as a reproducible research claim awaiting specialist review. Deterministic code and frozen artifacts are supplied so the claims can be assessed on their evidence. Specialist criticism is invited. Demonstrated errors will be documented and corrected, or the affected claims withdrawn.

Repository and review materials. github.com/lukekabbash/finite-grand-couplings

References

[1] G. R. Grimmett and M. Holmes, “Coalescence in Markov Chains,” *Journal of Theoretical Probability* 39 (2026), Article 49. doi:10.1007/s10959-026-01500-w.

[2] P. Diaconis and D. Freedman, “Iterated Random Functions,” *SIAM Review* 41(1) (1999), 45-76. doi:10.1137/S0036144598338446.

[3] B. Steinberg, “A Theory of Transformation Monoids: Combinatorics and Representation Theory,” *Electronic Journal of Combinatorics* 17(1) (2010), R164, 56 pp. doi:10.37236/436.

[4] P. J. Cameron and P. A. Kazanidis, “Cores of Symmetric Graphs,” *Journal of the Australian Mathematical Society* 85(2) (2008), 145-154. doi:10.1017/S1446788708000815.

[5] P. J. Cameron, “Dixon's Theorem and Random Synchronization,” *Discrete Mathematics* 313(11) (2013), 1233-1236. doi:10.1016/j.disc.2012.06.002.

[6] A. Schaefer, “Generating Sets of the Kernel Graph and the Inverse Problem in Synchronization Theory,” arXiv:1601.04295v2 (2016). doi:10.48550/arXiv.1601.04295.

[7] J. Araújo and P. J. Cameron, “Primitive Groups Synchronize Non-Uniform Maps of Extreme Ranks,” *Journal of Combinatorial Theory, Series B* 106 (2014), 98-114. doi:10.1016/j.jctb.2014.01.006.

[8] G. Budzban and P. Feinsilver, “A Hierarchical Structure of Transformation Semigroups with Applications to Probability Limit Measures,” *Journal of Difference Equations and Applications* 18(8) (2012), 1405-1434. doi:10.1080/10236198.2011.639365.

[9] J. G. Propp and D. B. Wilson, “Exact Sampling with Coupled Markov Chains and Applications to Statistical Mechanics,” *Random Structures & Algorithms* 9(1-2) (1996), 223-252. doi:10.1002/(SICI)1098-2418(199608/09)9:1/2<223::AID-RSA14>3.0.CO;2-O.