A Latin strategy that is not optimal for the finite field / uniform-target objective

Draft mathematical manuscript and Lean dossier for MathDB 347712. Prepared 1 November 2026. Author attribution and publication decisions are reserved to the project owner.

Abstract

For the exact finite field objective of Definition 3 in Czumaj, Kontogeorgiou and Paterson, Haystack Hunting Hints and Locker Room Communication, we give a Latin maximum-agreement strategy on five symbols that is not optimal. With the first maximizing row as the tie rule, the cyclic Latin square has field 290, whereas an explicit second Latin square has field 302. Both strategies partition all 120 permutations into five nonempty fibres. The corresponding optimized success probabilities for an independent uniform target and uniform permutation are 29/60 and 151/300, with gap 1/50. This refutes the Latin-implies-optimal direction of the published Conjecture 3 under its Definition 3 interpretation. Complete finite count certificates and unchanged Lean proofs accompany the argument.

Scope qualification. The raw MathDB statement omits the target distribution. The source also introduces a worst-target objective in §2.1; that is a different objective and is not settled here. Indeed, the competing strategy has a smaller worst-target success probability. We do not determine the exact optimum, prove minimality of order five, settle the necessity direction, exclude the existence of an optimal Latin strategy, or change the asymptotic conclusions of the source.

Review and acceptance status

The independent source/significance review gives a qualified PASS for the exact finite field refutation and bounded novelty. The independent fidelity/formal review gives a qualified PASS for Definition 3, while retaining HOLD for an unqualified or worst-target reading of the MathDB wording. Their full reports and machine-readable verdicts are included in the dossier. These qualifications are part of the claim. This manuscript and replay do not constitute a new independent acceptance: root acceptance and root final replay remain pending, with accepted credit zero and campaign context unchanged at 2/20 beyond 39.

mathdb:2611.00003v1 03 Nov 2026

1. Protocol, field, and the precise conjecture

Fix $n \geq 1$ and write $[n] = \{0,...,n-1\}$. A permutation σ maps a position i to the symbol $\sigma(i)$. Alice sees σ and sends a message $h = s(\sigma)$ in $[n]$. Bob knows the target symbol k and the message, and chooses a position $b(h,k)$. Success means $\sigma(b(s(\sigma),k)) = k$. Bob need not use an injective response map; each message-target pair may be optimized independently. A deterministic Alice strategy is a function $s : S_n \to [n]$. Its fibres $C_h = \{\sigma : s(\sigma)=h\}$ form a labelled partition. We will exhibit nonempty fibres, so the convention on empty parts is immaterial to this counterexample.

Following Definition 2 of [1], define magneticity, intensity, and the field of a fixed strategy by

```
mag(s,h,i,k) = #{σ  in  S_n : s(σ)=h and σ(i)=k},
I(s,h,k) = max_{i  in  [n]} mag(s,h,i,k),
Φ(s) = Σ_{h  in  [n]} Σ_{k  in  [n]} I(s,h,k).
```

The source writes $F(n)$ for the maximum of this sum over strategies. Thus Definition 3 calls s optimal exactly when $\Phi(t) \leq \Phi(s)$ for every admissible t. Our notation Φ distinguishes a fixed partition's field from the global optimum $F(n)$. All counts use the full symmetric group, rather than a sample of permutations.

If σ and k are independent and uniform, then for a fixed decoder b the success probability is $(1/(n\cdot n!))$ times the sum of $mag(s,h,b(h,k),k)$ over h,k. This follows by counting the $n\cdot n!$ equiprobable pairs (σ,k) and grouping by message and target. Each summand is bounded by $I(s,h,k)$, and equality is attained by choosing a maximizing position separately for each (h,k). Consequently optimized uniform-target success is exactly $\Phi(s)/(n\cdot n!)$. This normalization is an ordinary counting proof; the supplied Lean files formalize the finite fields and counts, not a probability-space theorem.

An n by n Latin square L has each symbol exactly once in every row and column. Put $a(L,\sigma,h) = \#\{i : \sigma(i)=L(h,i)\}$. A Latin strategy chooses a row of maximum agreement for every input: $a(L,\sigma,j) \leq a(L,\sigma,s(\sigma))$ for all σ and j. The square is arbitrary, as are choices between maximizing rows. Section 7.2 of [1] explicitly recognizes legitimate tie choices for shifts, and imposes no extra tie optimization in the Latin construction.

The literal published Conjecture 3 is: "Optimal strategies are exactly the latin strategies." Under Definition 3 its universal form is: for every positive n and every strategy s, s is optimal if and only if s is a Latin strategy. We refute the implication from Latin to optimal. A restriction requiring all fibres nonempty is refuted as well; no general equivalence between the two partition conventions is asserted.

mathdb:2611.00003v1 03 Nov 2026

2. Two explicit partitions of S_5

Use row and column indices 0 through 4. Let A(h,i)=h+i modulo 5, and let B be the following second square. The complete arrays are displayed for direct checking.

```
A = 0 1 2 3 4       B = 4 0 3 2 1
    1 2 3 4 0           3 2 4 1 0
    2 3 4 0 1           1 4 2 0 3
    3 4 0 1 2           2 1 0 3 4
    4 0 1 2 3           0 3 1 4 2
```

Each row and each column in each array is a permutation of [5]. Define s_L(σ) to be the least row index attaining max_h a(L,σ,h). This first-maximum rule is used for both arrays. It defines a total strategy on all 5! = 120 permutations. In particular both s_A and s_B are Latin strategies under the source construction.

Strategy	Fibre sizes in row order	Sum
s_A	46, 26, 21, 16, 11	120
s_B	33, 27, 25, 20, 15	120

Every fibre is nonempty. The first-maximum rule is substantive: it fixes an actual partition and avoids leaving a family of unspecified tie choices. No claim is made that the fields of all tie choices are equal. Nor does the conjecture require optimizing ties before a strategy counts as Latin.

Theorem 1 (finite field / uniform-target refutation)

There exist two nonempty-fibre Latin maximum-agreement strategies s_A,s_B on S_5 such that Φ(s_A)=290 < 302=Φ(s_B). Thus s_A is not field-optimal, even when optimality is compared only with nonempty-fibre strategies. The universal Latin characterization in the exact finite Definition 3 sense is false. Optimized uniform-target success differs by 1/50.

Proof. The Latin property and maximum-selection rule were checked above. The exact intensities in Section 3 yield the two field equalities by summation. As s_B is an admissible nonempty-fibre competitor, the inequality contradicts optimality of s_A. Substituting n=5 and s=s_A into the Latin-implies-optimal implication therefore refutes the universal biconditional. Finally divide by 5·120=600: 290/600=29/60, 302/600=151/300, and 12/600=1/50. The full finite computation establishing the intensities is specified below and is independently kernel checked in the accompanying Lean files. □

mathdb:2611.00003v1 03 Nov 2026

3. Exact count certificate and its verification

The entries in each row below are the five intensities I(s,h,k), in target order k=0,...,4. The last column sums that row. These are maxima of magneticities for the actual fibres, not merely agreement scores.

h	s_A: intensities	Sum	s_B: intensities	Sum
0	18,18,18,18,18	90	16,16,16,14,14	76
1	13,13,13,13,13	65	14,14,14,14,11	67
2	11,11,11,11,11	55	13,13,12,13,12	63
3	9,9,9,9,9	45	10,11,11,11,10	53
4	7,7,7,7,7	35	9,8,9,9,8	43
Total		290		302

For a referee wishing to verify these tables without trusting Lean, enumerate every tuple $(p_0,...,p_4)$ of distinct elements of [5] in lexicographic order. There are $5\cdot4\cdot3\cdot2\cdot1=120$ tuples, each representing exactly one permutation and none omitted. For each tuple, compute the five agreements with L and take the first largest. Increment the count of that fibre, and increment $M[h,i,p_i]$ for each i. After all tuples are processed, take $\max_i M[h,i,k]$. This is precisely magneticity followed by intensity. The algorithm has no probabilistic approximation or search cutoff.

```
M[h,i,k] := 0 for all h,i,k
for p in all permutations of (0,1,2,3,4):
    score[h] := number of i with p[i] = L[h,i]
    h := first index with score[h] = max(score)
    for i in (0,1,2,3,4): M[h,i,p[i]] += 1
I[h,k] := max(M[h,i,k] for i in (0,1,2,3,4))
```

The dossier provides enumerate.py, partition.csv with all 120 assignments to both strategies and their agreement vectors, and certificate.json with both arrays of all 125 magneticities, all intensities, decoders, fibre sizes, and agreement histograms. Running python3 enumerate.py reconstructs the certificate from the two squares alone. A hand or independent implementation can check each CSV row, reconstruct the two 5 by 5 by 5 count arrays, and recover every displayed maximum. This is a complete finite count certificate, not an assertion that a large unreported experiment found a candidate.

For these particular partitions the inverse-row decoder d_L(h,k), the unique i with L(h,i)=k, attains every magneticity maximum. The certificate checks all 25 equalities per square. Thus these inverse decoders realize the actual fields. The Lean field computation does not assume this inverse-decoder fact: it maximizes over all positions directly. Equality between field and inverse-row agreement total here is verified, not used as an unproved surrogate.

mathdb:2611.00003v1 03 Nov 2026

4. What the comparison does and does not explain

For any Latin square with the selected row as decoder, the number of successful targets for a permutation equals its agreement with that row. With maximum selection, summing over permutations gives the sum of maximum row agreements. In these examples inverse-row decoding is optimal within each message-target pair, so that sum is also the field. The agreement histogram provides a second check on the totals:

Maximum agreement	1	2	3	4	5	Weighted total
A: number of inputs	15	50	50	0	5	290
B: number of inputs	3	62	50	0	5	302

Each row separately has the same fixed-point agreement law under a uniform permutation; relabelling its symbols identifies agreement with fixed points. Distinct rows of a Latin square disagree in every coordinate. These facts do not determine the joint distribution of the five agreements. The two histograms differ in twelve inputs at levels one and two, in aggregate, and give a total difference of twelve. Latinness alone therefore does not fix the field or ensure exact field optimality. This is an explanation of the displayed comparison, not a theorem that every square has an optimal inverse-row decoder or that ties never affect field.

Explicit worst-target limitation

For a fixed target k and a uniform permutation, optimized success is $q_s(k)=(\Sigma_h I(s,h,k))/120$. The same independently chosen maxima optimize each target, so the optimized worst-target value of this fixed partition is $\min_k q_s(k)$. The exact target totals are:

Strategy	Target totals k=0,...,4	Minimum /120
s_A	58, 58, 58, 58, 58	58/120
s_B	62, 62, 62, 61, 55	55/120

The field is the sum of these target totals, while the worst-target objective uses their minimum. Hence the competitor improves the uniform average while worsening the minimum. Source §2.1 states a worst-target objective, whereas §2.2 and Definition 3 use the uniform-target field. We rely on the explicit Definition 3 meaning of “optimal” in the classification and do not identify the objectives. An interpretation of the MathDB sentence as a worst-target classification remains unresolved by this comparison.

The theorem does not compute F(5), assert optimality of B, find the smallest counterexample order, refute optimal-implies-Latin, rule out all optimal Latin strategies, or establish a new asymptotic rate. Each would require a separate argument. Neither the robustness supplement nor the numerical identities are separate discoveries for credit.

mathdb:2611.00003v1 03 Nov 2026

5. Formal correspondence and reproducible replay

The source-only packet contains byte-identical Latin.lean from the author lane and Supplement.lean from the fidelity review. Latin.lean defines P n as Equiv.Perm (Fin n) and Strategy n as P n → Fin n. agreement and magneticity are filtered full-universe cardinalities; intensity is Finset.univ.sup over positions; field is the double sum. At positive n the finite supremum is the maximum. Latin requires row and column bijections. IsLatinStrategy existentially supplies a Latin square and verifies the maximum-agreement inequality for every permutation and every row. IsOptimal quantifies over every strategy.

```
MathDB347712.SourceClaim :=
  forall n, 0 < n → forall s : Strategy n,
    IsOptimal s ↔ IsLatinStrategy s
MathDB347712.source_refutation : ¬SourceClaim
```

The cyclic_latin and cyclic_selects_maximum theorems establish Latin admissibility. cyclic_field and better_field compute 290 and 302 by ordinary decide, producing kernel-checkable proof terms. cyclic_not_optimal instantiates universal optimality with choose better and derives the numerical contradiction. source_refutation applies the conjectured implication at n=5. The contradiction is the negation of the full universal statement, rather than an assertion proved only for a bounded sample of n.

The general function type permits empty fibres. Supplement.lean explicitly introduces NonemptyFibres, RestrictedOptimal comparing only with nonempty-fibre strategies, and RestrictedClaim quantifying only such strategies. Its restricted_refutation and ordinary_partition_counterexample establish the nonempty robustness claimed in Theorem 1. It also proves that B is Latin, that its selection is maximal, both fibre-size vectors, card(P 5)=120, and the two actual target-count vectors. The ordinary probability normalization and inverse-decoder explanation are not claimed as additional formal probability theorems.

Reproduce with python3 replay.py after inspecting the script. It validates proof-source.zip, rejects traversal, absolute paths, symlinks, unexpected members and prebuilt local objects, and extracts into a newly created directory in this lane. The process environment is sanitized. LEAN_PATH contains only that fresh local directory and the specified package caches, omitting the upstream project’s own build directory. It runs Lean 4.34.1 with -j1 -DwarningAsError=true, then audits every declaration by module origin, including generated helpers, with full types and transitive axioms. Expected counts are 22 for Latin and 17 for Supplement. Only propext, Classical.choice and Quot.sound are permitted; unsafe declarations and new axioms are rejected.

Both modules are independently checked with leanchecker. Each Lean/checker process is restricted to at most two CPUs, 8 GiB address space, 100 CPU seconds and 110 wall seconds. Invalid-False and corrupt-module negative controls must fail. The pinned mathlib revision is d13f23b723b8a846827a245b89c10fc7d3f11612. All stdout, stderr, commands, timings, binary hashes, dependency manifest and fresh audit sources are retained. Imported prebuilt package oleans are trusted: this is not a complete upstream rebuild or source-to-olean attestation. A further root replay remains a separate gate.

mathdb:2611.00003v1 03 Nov 2026

6. Source versions, related work, and bounded significance

The exact Latin biconditional appears in arXiv:2008.11448v2 (3 June 2021), §7.2 Conjecture 3, in the ICALP 2021 version as Conjecture 25, and in the publisher-typeset article with DOI 10.1002/rsa.21114 as Conjecture 3. The arXiv v1 (2020) instead states an optimal-implies-shift conjecture. It is a different claim and direction. The raw MathDB record's 2020 field therefore must not be taken to date this exact Latin biconditional. We give no refutation of that earlier necessity assertion.

Both independent reviewers inspected the lawful Warwick institutional copy of the publisher-typeset body, including Definition 3 and §7.2. Its SHA256 is recorded in the source index and README. That frozen body supports the exact formulation used here. It does not authenticate today's Wiley-hosted bytes or establish a complete erratum search. The previously denied Wiley endpoint was not retried. Frozen source originals are referenced by absolute path and hash, rather than duplicated into the dossier.

The source's asymptotic bounds and shift construction concern the order of the optimum and do not assert the exact finite equality needed to make every Latin strategy optimal. Its Alice-In-Chains result imposes an extra exclusion condition and gives the naive value 2/n. At n=5 this is 2/5, below both displayed uniform-target values. The independent source/significance review read the actual relevant proof bodies and found no complete deduction from them to this finite separation. The present claim does not strengthen the asymptotic theorem or remove the AIC restriction.

Hendrey and Wanless [2] study covering radius in Hamming permutation space. The distance identity $d(\sigma,\text{row})=n-\text{agreement}(\sigma,\text{row})$ connects the vocabulary, but covering radius determines the minimum over σ of the maximum agreement, whereas this comparison sums that maximum over all σ. Both displayed squares have minimum maximum agreement one, hence covering radius four, while their sums differ. The inspected covering-radius results alone therefore do not provide the full distribution or the needed field comparison. This distinction does not certify that no other literature contains such a result.

The source/significance review conducted a bounded primary-source and public-registry screen on 1 November 2026 and located no matching public refutation. Its qualified PASS recommends at most one conditional result. It records incomplete specialist/citing-literature and registry-body coverage, inaccessible live MathDB/Palomar and direct Wiley body, and no worldwide priority certification. The fidelity review independently checked semantics and Lean but did not repeat the originality search. This manuscript reuses those frozen qualified findings and makes no new worldwide novelty claim. The contribution is the explicit exact field counterexample to the written finite classification; its elementary size neither establishes nor negates originality.

References and evidence

[1] Artur Czumaj, George Kontogeorgiou and Mike Paterson. Haystack Hunting Hints and Locker Room Communication. arXiv:2008.11448v2; ICALP 2021, LIPIcs 198, article 58; published DOI: 10.1002/rsa.21114. Definitions 1–3 and §7.2 are the primary scope anchors.

[2] Kevin Hendrey and Ian Wanless. Covering radius in the Hamming permutation space. arXiv:1811.09040v3. Related-work comparison is bounded to the primary material inspected in the independent review.

[3] Frozen local reviews: latin-source-significance-review-r1/REVIEW.md and review.json; latin-fidelity-review-r1/REVIEW.md and review.json. Full reports are included under evidence/. source-reference-index.json and frozen-inputs.json bind original sources and receipts. README.md describes reproduction; manifest.json binds the final dossier components and archives.

mathdb:2611.00003v1 03 Nov 2026