Uniform discrepancy bound for equal-law alternating shoe pairs

Myeongjin Kang

Abstract

For every discrete equal shoe-color law p, under literal independent strict left/right alternating arrivals, the total variation discrepancy between the independent conditional matching-pair color and actual first chronological opposite-side match color is at most255/256. Consequently the supremum over all countable laws is strictly below1; no finite-support restriction, positive-support-only premise, simultaneous-pair/random-side replacement or unproved law identity. Actual product/conditioning/stopping/fallback/TV/countable and finite-extension bridges proved.

AI-assisted research preprint. The complete accepted technical proof follows. Independent internal AI reviews are not external journal peer review. Exact mathematical hypotheses, formal coverage, source-version boundaries, trust limits, and bounded prior-art qualifications remain controlling. No worldwide-priority certification is asserted.

mathdb:2611.00115v1 17 Nov 2026

A UNIFORM DISCREPANCY BOUND FOR ALTERNATING SHOE PAIRS MathDB 324571 - implementation report, 8 November 2026 UTC

1. Statement and exact source experiment For every discrete probability law p, the total variation distance between the conditional matching-pair color X and the first chronological left-right matching color Y satisfies D(p,p) <= 255/256. Consequently sup_p D(p,p) <= 255/256 < 1. This addresses the strict bound, not the optimal-constant conjecture.

The copied full MathDB literal asks sup_p D(p,p)<1 without a support-size restriction. Arratia and DeSalvo, On the Random Sampling of Pairs, with Pedestrian Examples, arXiv:1211.6486v2 (2013), Sections 2, 4, 9.1-9.2, Conjecture 4 equation (30), specifies the model. The final version, Advances in Applied Probability 47 (2015), 292-305, DOI 10.1239/aap/1427814592, Sections 9.1-9.2, Conjecture 4 equation (19), retains it. Their Theorem 3 gives a limiting-family formula, not a universal bound. Original and final copied bodies supply the actual definitions.

Colors are numbered by Nat in the source, with finite alphabets padded by zero. Independent colors arrive as L1,R1,L2,R2,... . A match requires one shoe of each side. If the first match occurs on a new arrival, its color is the new shoe's color; multiple matching partners cannot create a color tie. Simultaneous-pair and independently randomized-side rules are different experiments and are not used.

Put s=sum_i p_i^2. Then 0<s<=1, p_i<=sqrt(s), and x_i=P(X=i)=p_i^2/s. With iid infinite coordinates, no match through 2N arrivals implies all N aligned left-right pairs are unequal. Thus P(T>2N)<=(1-s)^N. This tends to zero, proving almost-sure termination. Arbitrary fallback colors on the null nontermination event do not change the law.

The formal model uses PMF, Measure.infinitePi, a prefix HasMatch predicate, Nat.find for the least matching prefix, and measurable winner pushforwards. ConditionalColorMeasure is the normalized restriction of the genuine iid space to equality of its first two coordinates, then projection to the first color. No desired law relation or bound is an assumption.

mathdb:2611.00115v1 17 Nov 2026

2. Ordinary proof of the uniform constant Write y_i=P(Y=i). We prove y_i>=x_i/256 pointwise. If s>=1/64, L1=R1=i forces Y=i, so y_i>=p_iˆ2=s*x_i>=x_i/256. This includes point masses and the threshold equality.

If 0<s<1/64, set K=floor(1/(4sqrt(s))). Since the quantity floored exceeds 2, K>=1/(8sqrt(s)), K<=1/(4sqrt(s)), Kˆ2>=1/(64s), and Kˆ2*s<=1/16. A zero coordinate is immediate. For a positive coordinate, p_i<=sqrt(s)<1/8, so 1-p_i>7/8 and K*p_i<=1/4.

For a,b among the first K positions on each side, let E_ab specify that exactly L_a and R_b have color i within this block. Its probability is p_iˆ2*(1-p_i)ˆ(2K-2). These Kˆ2 events are disjoint. Let B_ab be the event of another cross-side match, avoiding the two designated positions. Conditional on E_ab, the remaining coordinates are independent with masses p_j/(1-p_i) for j != i. A competing pair has equality probability at most s/(1-p_i)ˆ2. There are (K-1)ˆ2 such pairs. The union bound therefore gives P(B_ab|E_ab)<=4/49<1/2.

The implementation bounds E_ab intersect each competing collision by the unrestricted four-coordinate event, of probability at most p_iˆ2*s after summation over colors. Thus P(F_ab)>=p_iˆ2/2-Kˆ2*p_iˆ2*s>=7*p_iˆ2/16>=p_iˆ2/4. This is the actual ShoeCollision and ShoeMinorization route; the preceding conditional argument is a separate valid ordinary proof. No independence of competing-pair events is assumed.

On F_ab=E_ab minus B_ab, i is the only common color on the two sides. The later designated arrival is the chronological first match, in either arrival order. Thus F_ab implies Y=i. Bernoulli's inequality gives (1-p_i)ˆ(2K-2)>=1-2K*p_i>=1/2. Summing disjoint F_ab gives y_i>=(Kˆ2/4)*p_iˆ2>=p_iˆ2/(256s)=x_i/256.

The countable series sum_i |x_i-y_i| converges, being bounded termwise by x_i+y_i. Hence dTV(X,Y)=1-sum_i min(x_i,y_i)<=255/256. The subset-supremum convention agrees with half the l1 sum: the positive-difference subset attains it, and every subset is bounded by that value. The same constant for every p bounds the supremum strictly below 1.

mathdb:2611.00115v1 17 Nov 2026

3. Arbitrary countable colors and finite zero extension Let A be any countable discrete measurable alphabet and p a PMF on A. Countability supplies an injective code e:A->Nat; no ordering of colors or positive-support assumption is imposed. Push p forward as p.map e. Coordinatewise relabeling pushes the actual infinite iid product measure to the Nat product measure. This follows from the product-measure mapping theorem and the PMF-to-measure mapping identity.

Injectivity preserves HasMatch at each prefix, hence FirstMatch, termination, and the least prefix length. On terminating paths the winner is mapped by e. With the fallback mapped as well, the explicitly defined colorWinner commutes pathwise with coding even on nonterminating paths. Its measurability is proved by countable winner-fiber decomposition. The resulting winner pushforward measure is exactly the Nat sourceWinner measure after coding.

For the conditional experiment, equality of the first two coordinates is preserved by injectivity. Its probability becomes the Nat collisionMass. Mapping the normalized restriction and first-coordinate projection gives exactly conditionalColorMeasure of p.map e. This proves transport of the actual conditional law rather than postulating a relation between arbitrary PMFs. Separate normalization certificates show the conditional and winner measures are probability measures and that their toPMF pushforwards are exactly sourceX and sourceY.

For any measures mu,nu on the discrete countable alphabet, the range of absolute subset-probability differences agrees with that of their injective pushforwards. A source subset is represented by its image; a Nat subset pulls back. This proves equality of the TV suprema. The source-defined colorDiscrepancy therefore equals sourceDiscrepancy(p.map e,e(c)). The existing Nat bound now proves countable_color_discrepancy_bound and countable_color_strict_universal_supremum; no duplicate elementary inequality proof is introduced.

Finite alphabets use Fin.val. The iid law, chronological winner, conditional law and discrepancy are preserved. Every Nat index outside the finite range has PMF mass exactly zero. Zero coordinates, finite support, singleton laws and countably infinite support are permitted. The empty alphabet admits no PMF or fallback; no probability law is excluded by requiring a fallback element. Generic almost-sure termination follows by pulling back the Nat null event, and fallback independence follows on this full-measure event.

Formal map: ShoeRelabel handles chronology; ShoeColorMeasure handles iid coding; ShoeColorWinner handles measurable winner transport; ShoeColorConditional handles normalized conditioning; ShoeColorEndpoint handles TV and the universal supremum; ShoeColorNormalization and ShoeColorPMF provide probability and PMF certificates. The original Nat core is unchanged.

mathdb:2611.00115v1 17 Nov 2026

4. Kernel checks, body audit and extracted replay Pinned Lean is 4.34.1; mathlib revision is d13f23b723b8a846827a245b89c10fc7d3f11612. The actual Lean version, binary hashes and mathlib HEAD metadata were checked through the lane runner. Strict builds use -j1 -M3500 -DwarningAsError=true. All 20 baseline sources were cold-compiled from an empty build before the first new proof. All 28 final proof/control/named-audit modules passed strict compilation and leanchecker.

Classification is explicit in results/PRODUCTION-PLAN.json: 23 core modules, three controls, two named-axiom audit modules. ShoeReachableAudit is a tool, never imported into a mathematical endpoint. These are modules supporting one candidate, not separate credited results.

Three inherited deep-audit attempts failed with exit -6 while serializing the whole imported environment. Their logs and receipts remain immutable inputs. The repaired tool first enumerates actual production-module origins, including private/generated declarations, then traverses the shared reachable type/body graph with a visited set. Each node records origin, dependencies, body presence, axiom status and unsafe flag once. Unrelated imported declarations are not dumped. Documentary Python only merges already exported metadata and checks consistency and edge completeness; it performs no mathematics.

The expanded audit completed with real exit 0 under unchanged caps. It enumerates 259 production-origin declarations and 29,367 reachable nodes. The Nat supremum closure has 29,263 nodes; the new countable supremum closure has 29,308. Exact closures for the minorization, TV identity, sourceY identity, conditional law, normalization and transport endpoints are saved. The entire graph and every listed endpoint closure have only propext, Classical.choice and Quot.sound as axioms, with no unsafe flags. Named print-axioms output alone is not this audit.

The fresh extracted replay began with an empty replay/build. All 28 sources were extracted as text; no oleans were copied. Its Lean path contains replay/build and pinned upstream libraries only. Each module passed strict compilation followed by leanchecker sequentially. Source and extracted-source hashes were compared after checker execution, and fresh olean hashes recorded in results/REPLAY.json.

Positive controls include singleton instantiation, unused finite colors of mass zero, finite discrepancy invariance and distinct-coordinate chronology. Intentional semantic-mutation failures are retained separately and never imported by proofs. Logs preserve actual exits. Further isolated-control checks and the final artifact inventory are recorded alongside this report.

mathdb:2611.00115v1 17 Nov 2026

5. Scientific limits, process history and disposition This report proves a uniform bound, not an optimal constant or worldwide priority. The independent final review passed literal/final-source fidelity, all 23 core bodies, actual countable transport, all 28 strict/checker builds, separate source-only replay and the complete reachable dependency audit. Root acceptance and campaign credit are recorded separately in shoe-discrepancy-acceptance.json. The implementation lane assigned itself no credit.

The independent primary comparison distinguishes the source's limiting-family theorem and Bhattacharya's random-side, diffuse-array collision-time results from deterministic alternating marked-color laws. Galbraith-Holmes 2012 returned 404 and remains unread. No specific applicable marked-color theorem was identified. This is a bounded access limit, not worldwide novelty clearance. Same-day registry-head and indexed-search nonmatches do not establish complete proof-body coverage. All access denials remain binding; no alternative retrieval was attempted.

Historical uncapped/runpy displays, a brief replay/control overlap, truncations, checker chronology, omitted session display, and PDF/render failures are preserved in author/root records. Earlier no-endpoint/no-paper statements describe earlier segments. Later replay and frozen passes do not erase those failures. Author r4 used an own-cwd runner and a prospective self-capped log-display exception. Bulk receipt metadata were machine-read, not every entry personally displayed. Final-source Section4 TV was personally displayed late after compilation; this does not claim personal prebuild reading. Independent r6 read the final TV and chronology before compilation.

The incomplete r5 review's two guessed own-log paths failed before content display. The completed r6 review instead used receipt-derived log access. Its failed audit strategies and corrected syntax are retained; successful metadata partition and body traversal used unchanged caps. Root accepts certificate content on the independent source/body/kernel evidence with these procedural qualifications, not a blanket historical compliance claim. No production proof was edited for this reporting revision.

Caps are per runner group: two CPUs, CPU120, AS12GiB, sampled RSS6GiB, development guard130 and wall140. Sampling is not a cgroup guarantee; native image and compiler-internal peaks are unavailable. Actual exits, postchecker hashes, complete dependency graph/closures, isolated controls, frozen checks and process notes accompany the independent review. Only propext, Classical.choice and Quot.sound are transitive axioms. No added axiom, sorry, unsafe or native_decide discharges a proof.

Retained evidence: shoe-pair-discrepancy-formal-independent-review-sol-r6/FINAL-REVIEW.md and review.json; results/strict-replay-evidence.json; audit/own-nodes.json, own-closures.json and origin-validation.json. The root archive includes source texts, reproducibility instructions, full evidence references and the reporting corrections. There is one mathematical result; finite/countable transport and the constant earn no extra credit.

mathdb:2611.00115v1 17 Nov 2026