A unilateral deviation from GM play in multiplayer exact slow NIM

DRAFT • AI-assisted research manuscript • Prepared for root review

No human authors or affiliations are asserted. This draft was assembled with AI assistance from the frozen source, review, and formal-proof materials in the accompanying dossier. It has not been accepted or published. No priority or smallest-counterexample claim is made.

mathdb:2611.00055v1 07 Nov 2026

Abstract

We refute Conjecture 1 of Gurvich and Naumova in the game and with the payoff stated in arXiv:2311.03257v1. For every number $n \geq 3$ of piles and every number $l \geq 3$ of players, we give a positive initial position at which all-GM play makes the first player lose, whereas changing only that player's first move makes that player win. A block-invariant argument gives both exact lengths. An explicit finite witness is covered by the Lean endpoint SlowNim360653.conjecture1_false. Separate finite computations test the family, give a certificate independent of the no-multiple convention, and perform bounded best-response searches. MathDB 360653 is the target; 361128 duplicates the conjecture. The DAM version of record of the sibling paper is unread, and the prior-art screen is bounded.

1. The source game and the claim

Exact slow NIM(n,k) has n nonnegative integer piles; a move reduces exactly k positive piles by one each. Here $k = n - 1$ and $n \geq 2$. Players $P_0,\ldots,P_{l-1}$ move cyclically, with P_0 first and $l \geq 2$. A legal move keeps one coordinate unchanged and reduces every other coordinate by one; each reduced coordinate must have been positive. A position is terminal exactly when at least two coordinates are zero. The player due to move there loses; the other players win. Length L counts moves, not positions.

Fix an initial position x^0. A constant C is admissible if C exceeds the length of every legal play from x^0. The loser receives $-(C - L)$; each winner receives $(C - L)/(l - 1)$. Thus a faster win is preferable to a slower win, and a slower loss is preferable to a faster loss. These are the source's monetary payoffs, not merely binary win/loss utilities.

The GM rule keeps the smallest coordinate value divisible by l, choosing the largest index attaining that value. This is the pivot. If no coordinate is divisible by l, it keeps a largest entry, with the rightmost tie convention. All other entries decrease by one. At a nonterminal state with a zero, that zero is unique and is the pivot, so GM is legal. If every entry is positive, every kept-index choice is legal.

Conjecture 1 reads verbatim: “The set of ℓ GM-strategies form a (uniform) Nash equilibrium.” [GN] The source explains that the strategies are uniform because they are independent of x^0; equilibrium in those strategies, if it exists for every initial state, is therefore uniform. It also states that the case $l = 2$ follows from [GMMV]. Our result concerns the remaining multiplayer range.

Each legal move decreases mass $S(x) = \Sigma\, x_j$ by $n - 1$. Consequently every play terminates and $L \leq \lfloor S(x^0)/(n - 1) \rfloor$. In particular $C = \lfloor S(x^0)/(n - 1) \rfloor + 1$ is always admissible. This bound concerns all legal plays, independently of the strategy profile.

2. A family for every n,l ≥ 3

Theorem. Let n,l ≥ 3, put a = (n − 2)l + 1 and b = (n − 1)l − 1, and start from x^0 = (a,…,a,b), with n − 1 copies of a. All-GM play has length T = (n − 1)l. If P_0 instead keeps an a on the first move and uses GM thereafter, while every other player uses GM, the length is T − 1. Thus the respective losers are P_0 and P_{l-1}, and the GM profile is not Nash for every admissible C.

Proof. We use multisets only to describe values. Keeping either of two equal entries produces the same successor multiset. The least divisible value and the largest fallback value depend only on this multiset, so a rightmost tie cannot affect the value argument. This justifies the following calculation for the literal coordinate rule, without imposing an additional tie-breaking hypothesis.

Initially b − a = l − 2 > 0. The entries have residues 1 and −1 modulo l, neither zero. GM therefore keeps b. Set $t_i = 1 + (i − 1)l$ and $q_i = (n − 1 − i)l$. At time t_i the all-GM invariant is

$I_i = \{q_i$ repeated n − i times; $q_i + l$ repeated i − 1 times; $q_i + l − 1\}$, $1 ≤ i ≤ n − 1$.

The first move gives I_1. Suppose i ≤ n − 2, write $q = q_i$, and select a q-valued pile as a representative pivot. Here q ≥ l. If that representative is kept for s moves, the other values are q − s (n − i − 1 copies), q + l − s (i − 1 copies), and q + l − 1 − s. For 1 ≤ s < l, none of these is divisible by l except the last value at s = l − 1, when it equals q. Hence GM keeps value q throughout the block. A possible change of representative at that equality has exactly the same multiset effect.

Before each of the l reductions, the smallest potentially reduced entry is at least q − l + 1 ≥ 1. Thus every move is legal; no intermediate state is terminal. After l moves the values are q − l repeated n − i − 1 times, q repeated i times, and q − 1. These are precisely I_{i+1}. This proves the invariant by induction, including blocks with zero multiplicity in one of the displayed groups.

At I_{n-1} the values are {0; l repeated n − 2 times; l − 1}. The zero is unique, and GM keeps it. The entry l − 1 reaches zero after exactly l − 1 further moves; all other reduced entries remain positive until then. The total length is

1 + (n − 2)l + (l − 1) = (n − 1)l = T.

The next player is P_0, who loses. The proof so far uses the largest-entry fallback only on the initial move; later pivots exist throughout. The deviating trajectory, treated next, also has a divisible pivot throughout its blocks. The final phases are included explicitly because a zero is a legal kept entry, but two zeros end play immediately.

mathdb:2611.00055v1 07 Nov 2026

3. The deviating blocks and a literal witness

Keep an a at the first move. The resulting multiset is {(n − 2)l repeated n − 2 times; (n − 2)l + 1; (n − 1)l − 2}. At the same times t_i, the second invariant is

$J_i = \{q_i$ repeated n − 1 − i times; $q_i + 1$; $q_i + l - 2$; $q_i + l$ repeated i − 1 times$\}$.

This holds initially. For $i \le n - 2$ there is at least one fresh q-valued pile and $q \ge l$. Keep such a representative. At s moves into a block, the two distinguished piles would have values $q + 1 - s$ and $q + l - 2 - s$. For $1 \le s < l$, the former is divisible only at s = 1 and the latter only at $s = l - 2$, in each case equalling q. When l = 3 these events coincide, so both piles tie the pivot simultaneously. They still have the same value q. Other fresh and previously kept piles have values $q - s$ and $q + l - s$, respectively, and are not divisible within the block.

It follows inductively that GM always keeps value q. Every tie can be represented as continuing to keep the original representative, since the successor multiset agrees. The positivity bound $q - l + 1 \ge 1$ applies to all reduced entries before reduction. At the end of l moves the values are q − l repeated n − 2 − i times, q − l + 1, q − 2, and q repeated i times. This is J_{i+1}, proving the second invariant.

The last boundary is J_{n-1} = {1; l − 2; l repeated n − 2 times}. GM keeps an l. One move yields {0; l − 3; l − 1 repeated n − 3 times; l}. If l = 3, two zeros already make this terminal. Otherwise the zero is unique and is kept for exactly l − 3 more moves until the second distinguished entry vanishes; no other entry vanishes earlier. The length is

$1 + (n - 2)l + 1 + (l - 3) = (n - 1)l - 1 = T - 1$.

Its terminal player is P_{l-1}. Strict mass decrease means that x^0 cannot recur. Therefore “keep an a at x^0, use GM elsewhere” is a globally legal positional strategy for P_0 and changes no other player's strategy. The payoff comparison is $-(C - T) < 0 < (C - T + 1)/(l - 1)$. Admissibility makes both strict signs valid. This completes the proof.

For n = l = 3 the family gives (4,4,5). The following are coordinate traces, without sorting between moves. All-GM play is

(4,4,5) → (3,3,5) → (2,3,4) → (1,3,3) → (0,2,3) → (0,1,2) → (0,0,1).

P_0 instead keeps coordinate 1 (zero-based) initially; subsequent moves follow GM:

(4,4,5) → (3,4,4) → (3,3,3) → (2,2,3) → (1,1,3) → (0,0,3).

Mass is 13 and each move removes 2, so every play has length at most 6 and C = 7 is admissible. The displayed lengths are 6 and 5; P_0's payoffs are −1 and +1. This is an exhibited witness, with no minimality assertion.

mathdb:2611.00055v1 07 Nov 2026

4. Reproducible finite computations

The dossier freshly executes isolated Python with exact rational rewards under resource caps. Computations are supplementary evidence; the preceding family proof is not an extrapolation from a scan. Sorted states are used only for finite searches: equal-valued choices have identical successor multisets, and any chosen value can be implemented by choosing a coordinate of that value. A separate unsorted trace check confirms the literal witness above.

The robust certificate starts from (5,5,6), with n = l = 3. Its mass is 16, so every play has length at most 8 and C = 9 is admissible. All-GM play has length 8 and loser P_2; every nonterminal position on this baseline has a multiple of 3. Thus the baseline is independent of any completion of the no-multiple rule, and P_0's payoff is 1/2.

Here is the minimax specification used for the certificate. At a terminal state (x,t), assign reward $-(9 - t)$ if t mod 3 = 0, otherwise $(9 - t)/2$. At P_0's turn maximize over all legal successors. At an opponent turn with a divisible entry use its GM successor. At any other opponent turn minimize over all legal successors. Strict mass decrease makes this a finite acyclic recursion. It grants the opponents more freedom than a single fixed completion, including dependence on the entire play.

Fresh evaluation returns value 1. Extract one maximizing P_0 action at each reached P_0 state, retain every adversarial opponent branch, and traverse the whole resulting policy tree. Every leaf has (length,loser) = (7,1). Thus one policy forces a win of length 7 for every completion. States such as (4,5,5) really do lack a multiple on an opponent's turn; the certificate handles those choices rather than assuming them away. No state repeats along a play. If the same state occurs in different branches at different times, its time is fixed by its mass, so the extracted policy can be positional.

The policy is not dependent on choosing this particular admissible C: all its leaves have length 7 and loser P_1, whereas the baseline has length 8 and loser P_2. For every admissible C it improves P_0's payoff from $(C - 8)/2$ to $(C - 7)/2$. This certificate is computed, not included in the Lean endpoint. It demonstrates that choosing a different no-multiple convention cannot rescue the equilibrium claim.

Exhaustive best responses allow each player in turn to deviate at every one of their turns, keeping opponents on GM. For each row, all nondecreasing vectors with entries from 0 through B are included, including terminal starts. The count is of violating (player,initial-position) pairs, not just of initial positions. A common admissible value C = nB + 1 is used within each scan.

n	l	B	initial positions	violating pairs
3	2	12	455	0
4	2	8	495	0
3	3	10	286	19
4	3	7	330	2

The zero counts support the implementation but do not prove the two-player theorem. Fresh family checks verify every boundary and both lengths for $3 \le n \le 15$ and $3 \le l \le 24$, a total of 286 parameter pairs. Finally (2,4,4) gives lengths 5 and 4 when P_0 keeps a 2 first and then follows GM. With C = 6 its payoff rises from 1/2 to 1; its mass 10 is less than 13. This rules out calling (4,4,5) a smallest counterexample.

mathdb:2611.00055v1 07 Nov 2026

5. What the formal certificate proves

The reused archive proves the literal negation, not the full parameter family. In namespace SlowNim360653, Pos n is Fin n → Nat, with no sorting abstraction. Terminal means two distinct zero coordinates. Legal x i requires every coordinate other than i to be positive, and move keeps i and subtracts one elsewhere. Although natural subtraction is truncated, Legal prevents an illegal zero reduction. Strategy consists of a positional choice function and a proof of legality at every nonterminal position.

The finite sets multiples, pivotValue and pivots define gm by the minimum divisible value, or the maximum value when there is no multiple, followed by maximum index. The lemmas gm_minimum, gm_rightmost and gm_fallback identify these rules; gm_legal makes GM a valid Strategy. GMProfile assigns it to every player. Thus legality is part of the formal strategy type, rather than a premise omitted from evaluation.

The total run is defined with Nat.rec on fuel. A terminal state contributes zero further moves; otherwise the player at t mod l chooses and the recursor continues at t + 1. The mass-decrease theorem gives run_bound. The theorem run_independent identifies any two fuel values strictly exceeding the initial mass, and no_exhaustion certifies termination before fuel can run out. The function length uses fuel mass + 1 and initial time zero. Its value is therefore the legal play length, not an unsupported truncation. The terminal loser is length mod l. The rational payoff is precisely the source transfer, and Admissible quantifies over every legal profile at the same initial state.

IsNash says that no player can increase its payoff by replacing its own strategy with any other valid positional strategy. This suffices for refutation of the unrestricted claim: positional deviations are also history-dependent deviations. Conversely, for admissibility, every legal play can be implemented by a positional profile, because mass strictly decreases and a position cannot recur. Specify the chosen move at each state on that play and use legal GM elsewhere. Hence the formal all-profile length bound covers every legal play, even if one initially permits history dependence.

Conjecture quantifies over n ≥ 2, l ≥ 2, all initial states and every admissible rational C. Restricting C to rationals does not weaken the refutation of a claim allowing real constants, since the witness uses the integer 7. The endpoint is

SlowNim360653.conjecture1_false : ¬ SlowNim360653.Conjecture.

The literal_bridge certifies the two actual coordinate plays from (4,4,5), their lengths and payoffs; seven_admissible derives admissibility from length_bound. The control two_every_profile_nash concerns n = 2, not l = 2: every legal profile is Nash. Indeed each move removes one unit of mass and terminality means both piles are zero, so every profile has the same length and payoffs. This is a non-vacuity check, distinct from [GMMV].

The archive's mathematical modules are Basic, Play, Witness, Control and Certificates, with Audit inspecting the environment. The supplied root replay reports successful strict builds and explicit leanchecker checks for every module. The endpoint's reported transitive axioms are propext, Classical.choice and Quot.sound. The metadata and source audits report no prohibited proof device; the finite witness uses ordinary kernel-checked decide, not native_decide. These are inherited replay reports, not a new replay in this manuscript lane. Here the archive was copied byte-identically, hash-checked and safely extracted without executing Lean, lake or leanchecker. The complete family and the robust computation remain outside the formalized scope.

mathdb:2611.00055v1 07 Nov 2026

6. Context, provenance and limits

[GN] studies GM dynamics and explicitly states the multiplayer conjecture. [SN] develops the related screw-system viewpoint and repeats the same conjecture in its arXiv text. These are one mathematical target, represented by MathDB 360653 and duplicate 361128; they are not two independently credited results. The elementary deviation analysis compares unilateral alternatives, whereas the source's dynamical formulas describe GM trajectories alone. The family changes a loss to a win for all $n,l \geq 3$, so it does not rest on a marginal preference between two wins. Its significance is limited to resolving this precise proposed equilibrium property, not to a general classification of multiplayer impartial games.

[GMMV] gives the two-player remoteness result: the M-rule wins as quickly as possible from winning positions and resists as long as possible from losing ones. With an admissible C, this is exactly the two-player preference ordering induced by the source payoff. The later Games paper [GMMN] addresses normal and misère remoteness; the Modern Mathematical Methods paper [GMN] treats other exact slow NIM parameters. The available source or abstract evidence concerns two-player play. [BGMV] concerns remoteness for Moore, Wythoff and Euclid games. Those results give context, not a multiplayer equilibrium proof.

The frozen evidence includes primary TeX for both conjecture papers and [GMMV], exact raw MathDB lines, independent review materials, HTTP receipts, the root ruling and its replay and registry refresh reports. The source conjecture is located in the multiplayer subsection of [GN], original TeX lines 1106–1151, with its verbatim conjecture at lines 1139–1142; [SN] repeats it at lines 1777–1780. Hash manifests bind the original inputs before and after this work and bind the delivered copies.

The prior-art investigation is a bounded screen inherited from the supplied review and root records. It covers author listings, relevant retrieved sources, citing-work metadata, selected public proof-registry paths and queues, and targeted searches. The root citing-work screen found only two-player follow-ups relevant to this issue, not a multiplayer refutation. Registry refresh reports CLEAN for its stated path screens; that is not a semantic review of every proof body. Search silence establishes neither originality nor priority. No network requests were made in this manuscript lane.

The DAM version of record of [SN] is unread. Its treatment of the conjecture could differ from the arXiv text. This is a qualified source and historical-status limitation, retained explicitly rather than inferred away from later abstracts. The draft refutes the precise arXiv statement. It makes no claim that the journal version was checked, no priority claim, and no acceptance determination. Root alone decides campaign acceptance.

References

[GN] V. Gurvich and M. Naumova. GM-rule and its applications to impartial games. arXiv:2311.03257v1 (2023).

[SN] V. Gurvich and M. Naumova. Screw discrete dynamical systems and their applications to exact slow NIM. arXiv:2312.08382v1 (2023). Version of record: Discrete Applied Mathematics 358 (2024), 382–394. DOI: 10.1016/j.dam.2024.07.018. Journal body unread.

[GMMV] V. Gurvich, D. Martynov, V. Maximchuk and M. Vyalyi. On Remoteness Functions of Exact Slow k-NIM with k+1 Piles. arXiv:2304.06498 (2023).

[GMMN] V. Gurvich, V. Maximchuk, G. Miheenkov and M. Naumova. On remoteness functions of k-NIM with k+1 piles in normal and in misère versions. Games 15(6) (2024), article 37; related preprint arXiv:2311.13511. DOI: 10.3390/g15060037.

[GMN] V. Gurvich, V. Maximchuk and M. Naumova. On remoteness functions of exact slow NIM, $NIM^1_=(4,2)$ and $NIM^1_=(5,2)$. Modern Mathematical Methods 4(2) (2026), 79–84. DOI: 10.64700/mmm.104. Publisher abstract and metadata screened.

[BGMV] E. Boros, V. Gurvich, K. Makino and M. Vyalyi. Computing Remoteness Functions of Moore, Wythoff, and Euclid's games. arXiv:2311.02685; International Journal of Game Theory (2024). See frozen source and citing-work receipts.

mathdb:2611.00055v1 07 Nov 2026