GREGO'S LIMBO MATCH™

The Maths of the Feast

The feasible tally, the worth of a category, and the merit that ranks a champion

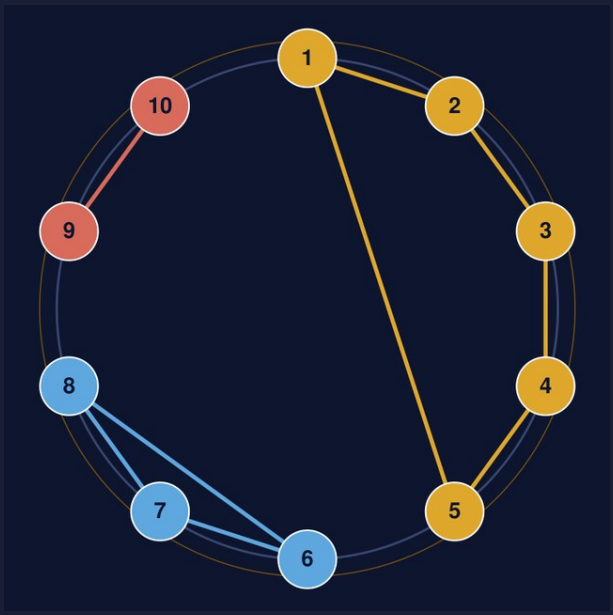

one shuffle of P = 10 · loops {5, 3, 2} · a partition of 10

grego.hu/limbo · OEIS A389262 (CST)

The Feast · 220,407 categories

The object, and what a category is

Grego's Limbo Match is a memory–strategy–lottery game. P blue chips numbered 1…P sit face-down in a middle stack; their P matching red chips lie face-down around a marked dial. The player flips the top blue to read its value, then searches the reds for its pair within a flip limit, or gambles the clover for a blind guess and banks points for unspent flips. A missed search sends the blue to the bottom of the stack to resurface last with a better chance: the Limbo. Jackpot Jack, the score ceiling of a category, is P×(F+1).

A category fixes the setup with four dials

P — pairs (2 to 21). F — the flip limit per search (1 to P). L — the limbo come-back chances (0 to P). LL — the loop-length window [MIN, MAX] of the shuffle. A shuffle lays the reds in a secret order; it falls into loops, and the set of loop lengths is the whole story that the pages below count and weigh.

The F = P edge

When F = P the flip limit already lets every red be revealed, so a miss is possible only by re-flipping a red, which is non-virgin and denies the limbo. Limbo fires only when F < P. At F = P the L dial is inert, so an F = P category exists only at L = 0.

Dial	Meaning	Range	Note
P	pairs — one blue + one red	2 … 21	board size
F	flip limit per search	1 … P	F = P makes every red reachable
L	limbo come-back chances	0 … P	fixed at 0 when F = P
LL	loop-length window [MIN,MAX]	1 … P	selects the shuffle pool

The dial face also carries the clover (a blind-guess lottery for the full flip bonus), the pin (rotate all reds until a chosen red sits on its own dial number) and the eye (peek every red, at the cost of the record). These are strategy, not category axes.

The canon — the World Record line

The World Record is one fixed category per board. On every WR rung the limbo count is L = 1 and the loop range is MIXED (any length); the flip limit is the CST, A389262(P) — the Critical Survival Threshold, the least flip-cap at which more than half of all P-shuffles are winnable under the 100-prisoners loop strategy (Golomb & Gaal, 1998). Each WR rung sits on the fair, better-than-even line.

The WR line runs P = 3 to 21 — nineteen rungs. At P = 2 the CST equals P, so P = 2 carries no WR rung; every P = 2 category is a Custom Record (F = 1 with any L, or F = 2 with L = 0). The world shuffle may still deal a P = 2 board.

P	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
CST	2	3	3	4	5	5	6	6	7	8	8	9	9	10	11	11	12	12	13

WR flip limit per rung — the CST sequence A389262, by Grego (Gergely Földvári).

The field — Custom Records

Custom Records open the whole field around that calibrated centre. A player cooks a category by choosing F (1 to P), L (0 to P) and a loop window LL = [MIN, MAX]. The nineteen WR rungs are the fair midline; CR is everything else reachable from it.

The maths — shuffles, loops, partitions

A shuffle of P reds is a permutation, and its loops are its cycles. A loop pattern is a partition of P — a way to write P as a sum of parts (for P = 10, one pattern is 5 + 3 + 2). The number of shuffles sharing one loop pattern is the cycle-type count; summed over every partition of P these counts equal P!.

A category's shuffle pool

The window [MIN, MAX] gathers every partition whose parts all fall inside it; the pool size is the sum of those partitions' shuffle counts. At P = 10 the window [1, 4] gathers 23 partitions (632,736 shuffles) and [2, 5] gathers 7 (345,681 shuffles). The window [1, k] — all loops at most k — is where the CST lives: the smallest k whose pool passes P!/2.

These pools are large: a full P = 13 board has 13! = 6,227,020,800 shuffles, and every category's pool is an exact sub-sum of that total, so a score is placed against a field whose size is known to the last permutation.

Distinct pools

Redundant bounds collapse. At P = 10 with MIN = 3, raising MAX from 7 to 8 or 9 selects the same partitions (7 is the longest loop that fits), so 3–7, 3–8 and 3–9 are one category; 3–10 adds the single 10-loop and stays distinct. A window that no shuffle can satisfy is empty and is not a category. The distinct, non-empty pools per P are the LL categories column below.

The count — 220,407 categories

For each P the categories multiply out as Feast(P) = (distinct LL categories) × (valid F, L pairs). F runs 1 to P and L runs 0 to P, and F = P forces L = 0, which leaves P^2 valid (F, L) pairs per board:

P	partitions	LL cats	F,L pairs (P^2)	Feast(P)
2	2	3	4	12
3	3	4	9	36
4	5	7	16	112
5	7	8	25	200
6	11	13	36	468
7	15	14	49	686
8	22	19	64	1,216
9	30	22	81	1,782
10	42	28	100	2,800
11	56	30	121	3,630
12	77	40	144	5,760
13	101	41	169	6,929
14	135	49	196	9,604
15	176	55	225	12,375
16	231	63	256	16,128
17	297	67	289	19,363
18	385	80	324	25,920
19	490	83	361	29,963
20	627	95	400	38,000
21	792	103	441	45,423
2–21				**220,407**

LL categories are distinct shuffle pools after collapsing redundant bounds; every column of partition counts sums to P!.

Summed over P = 2 to 21 the feasible tally is 220,407 categories: 19 World Records (the MIXED slice with F = CST and L = 1) and 220,388 Custom Records. Weighted by the shuffle counts above, the Feast is the ground the ranking stands on — every score placed against a category whose size and difficulty are known exactly.

Worth — difficulty in the game's own currency

A category's worth is the most merit any deal inside it can ever hold, on the same 1…441 scale as the merit itself. It answers one question: how much does knowing the shuffle buy here? Worth is window-aware. A 13-pair Unit

board and a 13-pair Longest board share the same P, F and L, yet their worth is 1 and 103.

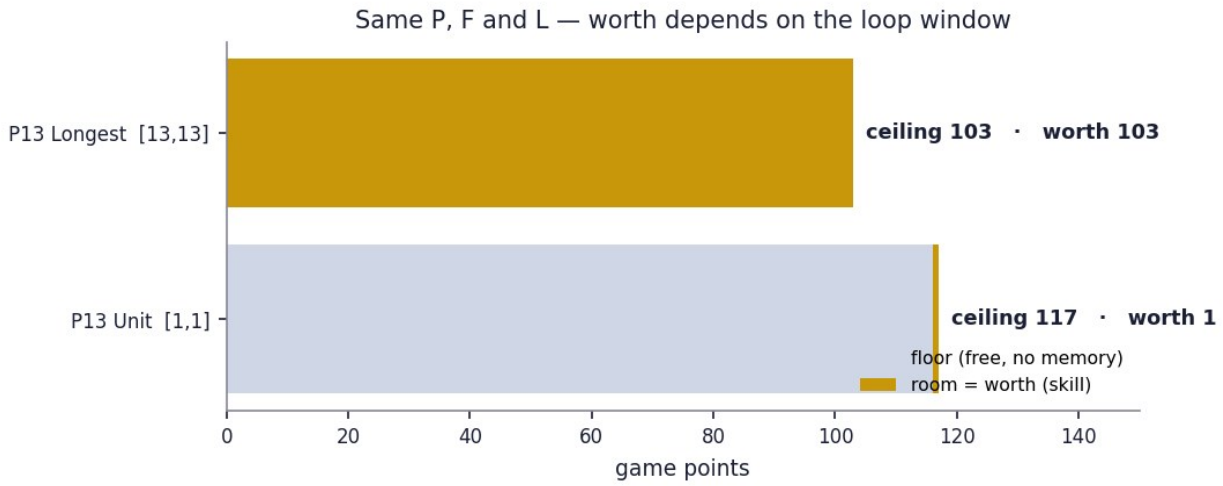

Worth is the room a shuffle leaves for memory, maximised over the shapes the window allows. A Unit board hands everything over free (room 1); a Longest board hides a full trail (room 103).

How worth is computed

For every loop shape the window allows, worth takes room = max(1, ceiling – floor) and keeps the largest. The ceiling depends only on the shape, never on the blue order, so worth is exactly computable by sweeping the shapes — at most 792 at P = 21. Worth orders the Champions' Table house, the empties sort, GLSX prestige and GLTX difficulty.

Floor and ceiling — the two honest bounds

Merit is measured against the actual shuffle the nick played, between two bounds that both err on the safe side.

The floor — what the shuffle hands over free

The floor is the no-memory trail-follower. To find red v it walks the trail from dial position v, which takes exactly the length of v's loop; a match on flip k scores 1 + (F – k), and a trail longer than F is never found. The floor depends only on the shuffle and F. It is the luck the board gives everyone.

The ceiling — the most a perfect memory can take, without gambling

The clover replaces flipping the blue, so a full 1 + F needs knowing both blue and red; a known red otherwise costs one tap for F. Limboed blues return in order, yellow-ringed and fanned, so they are all knowable; trail-following unseen reds is normal play, and the last unseen red is free by elimination. The deliberate-limbo harvest lets a spare slot force a known blue down to be clovered on return for 1 + F instead of F, at the cost of carrying its value and position across the wait. The ceiling takes the best of these perfect-memory lines. The virgin rule bans only re-flips; burning matched or known reds is legal.

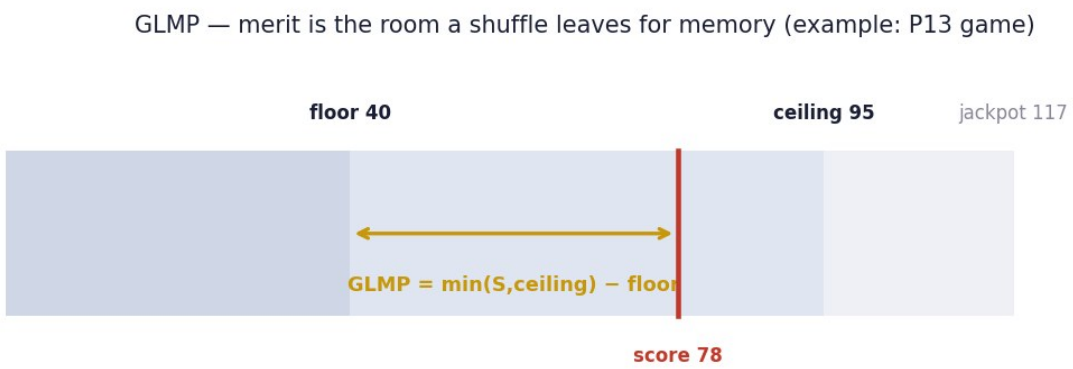

Between the free floor and the risk-free ceiling lies the room a memory can convert into points; capped at the score, that room is the merit.

The Feast, by worth

Weighted by their shuffle counts, the categories spread across the worth scale as below. The mass sits in the Serious band (100–199) at 36% of the Feast; the Summit (400–441) is razor-thin at 506 categories, reachable only where a long-loop window and a high L let a perfect memory harvest deliberate limbos. Worth and board size are independent axes — most low-worth categories still carry ten or more pairs.

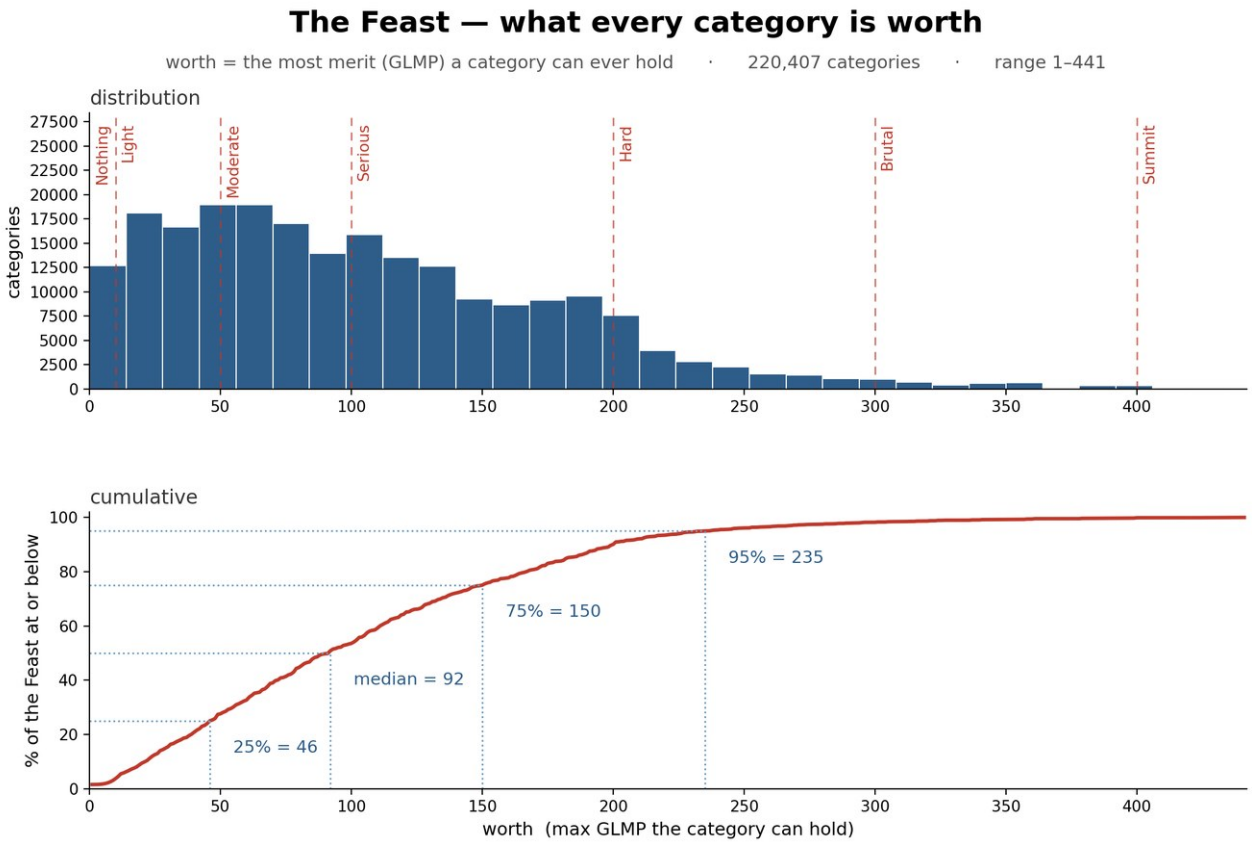

Distribution and cumulative worth across all 220,407 categories.

statistic	value
mean	104.5
25th percentile	46
median	92
75th percentile	150
90th percentile	200
95th percentile	235
99th percentile	337

worth	name	categories	share
1	Nothing to prove	3,315	1.5%
2–9	Trivial	3,366	1.5%
10–49	Light	53,587	24.3%
50–99	Moderate	57,214	26.0%
100–199	Serious	80,165	36.4%
200–299	Hard	18,818	8.5%
300–399	Brutal	3,436	1.6%
400–441	Summit	506	0.2%

Left: worth percentiles across the Feast. Right: the eight worth bands.

GLMP, GLMR and GLM100 — merit and its ranking

GLMP — the merit of a game

GLMP = max(1, min(S, ceiling) − floor). It reads the raw game points S against the shuffle's own floor and ceiling — nothing else. The floor of 1 is the point for understanding and playing. Across the Feast GLMP runs 1 … 441, the same scale as worth, because worth is the largest GLMP a category can hold.

On a P = 13 Longest board every trail is length 13 — longer than the flip cap — so a no-memory player finds nothing and the floor is 0, while a perfect memory reaches a ceiling of 103; a nick who scores 90 there earns GLMP 90. On a P = 13 Unit board (floor 116, ceiling 117) the same nick might score 116 and earn GLMP 1, because the board asks nothing.

GLMR — where a nick stands

A nick is ranked by its best single GLMP. GLMR = 1 + (the number of nicks with a strictly higher GLMP): a peer-versus-peer ladder with honest joint ranks and no tiebreak, so the best players sit at the top rather than against the 441 ceiling.

GLM100 — the certified hundred

The GLM100 is the head of the ladder: the nicks at GLMR 1 to 100. Matched reds stay face-down, visually identical to unmatched ones, so remembering which red is already spent is provable skill that even a two-pair game asks — the record score itself encodes merit. Within a single category the record is the best score, faster time breaking a tie.

Where worth and merit are used

The Champions' Table

The house is ordered by worth — lighter categories first — and the empty categories are sorted the same way. Each category holds one record, its best score with faster time breaking a tie; across the house the champions are ranked by GLMP then GLMR, with the GLM100 as the certified head.

GLTX — tournaments buy difficulty by percentile

An organiser sets difficulty as a percentile band, and the slider reads the LUT below, mapping 0–100 to worth so “harder than 75% of the game” is plain. The presets are Gentle (worth 1–49), Club (50–95), Competitive (95–198) and Geeks (198–441).

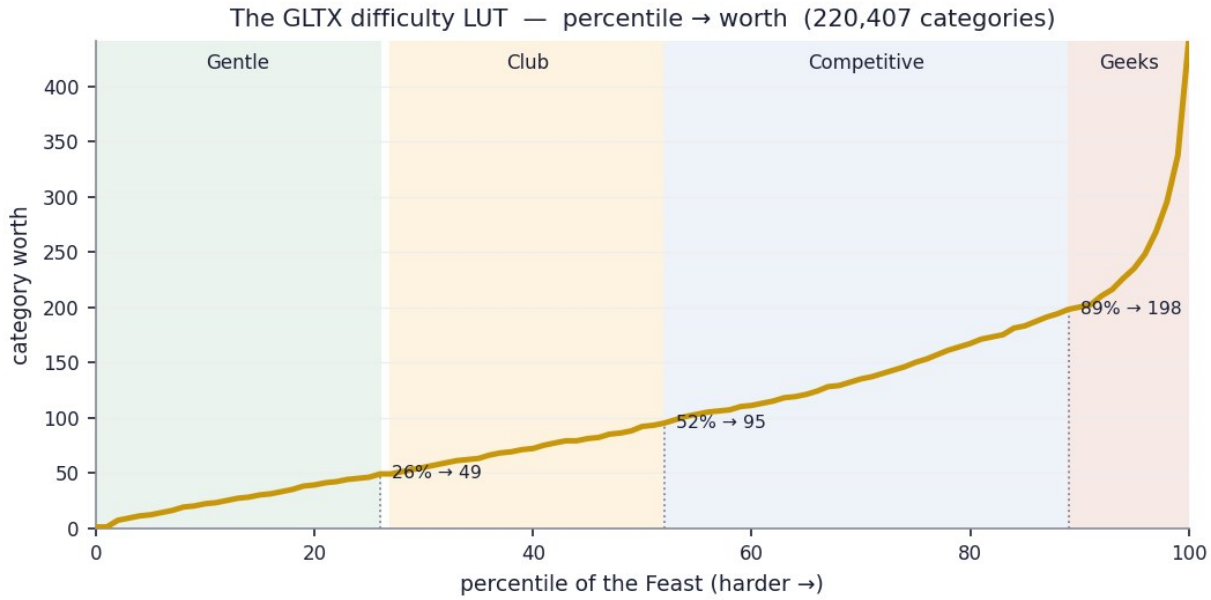

The GLTX difficulty LUT — percentile of the Feast → category worth.

GLSX — sponsorship price rests on worth

A category's sponsorship price combines its prestige — its worth, 1…441 — with its share, how busy it is against an average category, floored at a small per-category-day minimum.

The summa — the numbers this rests on

summa	value
Feast — total categories	220,407
World Record rungs	19

summa	value
Custom Record categories	220,388
Worth and GLMP range	1 … 441
Median worth	92
Summit categories (400–441)	506
Jackpot ceiling	J = P × (F + 1)
Shuffles at P = 13	13! = 6,227,020,800

The nineteen World Record rungs

Each rung fixes L = 1 and the MIXED window; F is the CST. Jackpot Jack is J = P × (F + 1), and the worth is the most merit the rung can hold.

P	F	J	Worth
3	2	9	8
4	3	16	14
5	3	20	16
6	4	30	25
7	5	42	36
8	5	48	41
9	6	63	55
10	6	70	60
11	7	88	77
12	8	108	96

P	F	J	Worth
13	8	117	103
14	9	140	125
15	9	150	133
16	10	176	158
17	11	204	185
18	11	216	195
19	12	247	225
20	12	260	236
21	13	294	269

P = 2 carries no rung: its CST equals P.

The tools on the dial

Clover — a blind-guess lottery: tap any red for 1 point plus the full flip bonus when it matches.

Pin — rotate all reds, keeping the shuffle, until a chosen red sits on its own dial number.

Eye — reveal every red face-up; while used, the game does not qualify for a record.

Limbo and the virgin rule — a missed search sends the blue to the bottom to return last; the limbo is granted only when no red is re-flipped during the search.

The Feast is the ground the ranking stands on

Every score is placed against a category whose size and difficulty are known exactly: its shuffle pool counted to the last permutation, its worth swept over every loop shape, its merit measured between an honest floor and an honest ceiling. The Feast holds 220,407 categories.

Worth alone is the lever. It tells a trivial board from a brutal one wherever the two share a size, and it lets a small board played perfectly outrank a large one shown up but not solved. A prodigy needs no age, and memory checks none.

References

A389262 (CST) — Grego (Gergely Földvári), 2025. Golomb & Gaal, *On the Number of Permutations on n Objects with Greatest Cycle Length k*, Adv. Appl. Math. 20 (1998). Wikipedia, *100 prisoners problem*.