THE EXACT FINITE-HORIZON MINIMAX REGRET OF PREDICTION WITH THREE AND FOUR EXPERTS, AND THE EXACT DEFICIT OF THE COMB ADVERSARY

KOREA SUPERINTELLIGENCE LABS

Abstract. In the classical game of prediction with expert advice — k experts, T rounds, binary gains, full information, an adaptive adversary — the minimax expected regret $V_T(k)$ is a rational number, but to our knowledge only the case $k = 2$ (Cover, 1965) is in print exactly: for $k \geq 3$ the literature gives asymptotics, the geometric-horizon variant, or a continuum limit. We determine $V_T(3)$ for every $1 \leq T \leq 12$ and $V_T(4)$ for every $1 \leq T \leq 10$. For three experts the values obey the closed form

$$V_T(3) \;=\; \frac{\big(4\lceil T/2\rceil + [T \text{ even}]\big)\binom{T}{\lfloor T/2\rfloor}}{3\cdot 2^T} \;=\; \frac{2\,\mathbb{E}|S_T| + \mathbb{P}(S_T = 0)}{3},$$

S a simple random walk — the three-expert companion of Cover's $V_T(2) = \mathbb{E}|S_T|/2$ — and the approximating expression published by Abbasi-Yadkori, Bartlett and Gabillon for even T exceeds the exact value by exactly $\mathbb{P}(S_T = 0)/3$. We also compute exactly what the *comb adversary* of Gravin, Peres and Sivan — conjectured by them to be asymptotically optimal for every k — inflicts at finite horizon: for $k = 3$ it is strictly suboptimal at every $1 \leq T \leq 12$, with deficit $1/6, 1/12, 1/8, 1/16, \ldots, 191/4096$; for $k = 4$ it is strictly suboptimal at $T = 1$ and at every $3 \leq T \leq 10$, and exactly optimal at $T = 2$. The bound holds for the whole class of "two teams by rank, fair coin" adversaries, of which the comb adversary is the best member. Every value is certified on both sides by an explicit per-state certificate (a player mixed strategy and a balanced adversary distribution), glued by weak duality, and machine-checked in Lean 4; Section 8 says exactly what is and what is not formalised.

1. Introduction

The game. Fix $k \geq 2$ experts and a horizon $T \geq 1$. On each of T rounds an adversary chooses a gain vector $g_t \in \{0,1\}^k$, as a function of everything that happened before; at the same time the player, who has seen the gains of all previous rounds but not of the current one, chooses an expert $j(t)$, possibly at random. The player receives $g_{j(t),t}$, and then all k gains are revealed. Writing $G_{i,T} = \sum_{t\leq T} g_{i,t}$ for the cumulative gain of expert i, the player's regret is

$$R_T \;=\; \max_{1\leq i\leq k} G_{i,T} \;-\; \sum_{t=1}^{T} g_{j(t),t},$$

and the object of this paper is the minimax expected regret

$$V_T(k) \;=\; \min_{\text{player}} \; \max_{\text{adversary}} \; \mathbb{E}[R_T].$$

This is the game of Gravin, Peres and Sivan [9] (in the fixed-horizon form of their Section 3; their main results are for a geometric horizon, see below), of Abbasi-Yadkori, Bartlett and Gabillon [1], and of Chase [5]: binary gains, full information, an adaptive adversary, a fixed and known horizon, expected regret; see [6] for the general theory of the game. The restriction of the adversary to $\{0,1\}$ rather than $[0,1]$ is without loss of generality; [9], Section 1: "We show that in the worst-case instance there is no benefit in using gains other than 0 and 1, so we restrict to $g_{it} \in \{0,1\}$." The game is a finite tree of matrix games with rational payoffs, so $V_T(k)$ is a rational number for every k and T; the question is which rational number.

What is in print. For two experts the answer is Cover's [7]: $V_T(2) = \lceil T/2\rceil \binom{T}{\lfloor T/2\rfloor}/2^T = \frac{1}{2}\mathbb{E}|S_T|$ for a simple symmetric random walk S, restated in [9], Section 3, as "the optimal regret in the finite horizon model with T steps is exactly half the expected distance travelled by a simple random walk in T steps", with the optimal adversary "at every step, choose an expert uniformly at random (with probability 1/2) and set him to 1, and the other expert to 0."

For three experts the exact solution of Gravin, Peres and Sivan is in a different model. Their abstract: "We study the classical problem of prediction with expert advice in the adversarial setting with a geometric stopping time. In 1965, Cover gave the optimal algorithm for the case of 2 experts. In this paper, we design the optimal algorithm, adversary and regret for the case of 3 experts." Their Theorem 4.2 gives, for the geometric horizon with stopping parameter δ, the optimal regret $\frac{2}{3}(1-\delta)/\sqrt{1-(1-\delta)^2} \to \frac{2}{3}\cdot\frac{1}{\sqrt{2\delta}}$ and an optimal adversary as a function of the gaps d_{ij} between the ranked cumulative gains; at the all-tied state $0 = d_{12} = d_{13}$, which is the initial state, the adversary listed is $\{1\}\{2\}\{3\}$ (advance one expert chosen uniformly at random). For the finite horizon their Section 3 treats $k = 2$ and, for general k, proves structure: a minimax-optimal adversary can be taken *balanced* (every expert receives the same expected gain at every step), the balanced distributions form a convex polytope (their Claim 2), and a fixed finite set of them suffices (their Claim 3), which "enables us to write a mundane dynamic program of size $O(T^k)$." No finite-horizon value for $k \geq 3$ appears in the paper. In their Appendix B.1 they record that the finite-horizon optimal adversary for three experts is time-dependent: with one step left it plays $\{1\}\{2\}\{3\}$ if all three experts are tied, $\{1\}\{23\}$ or $\{13\}\{2\}$ if exactly the two leaders are tied, and "any balanced strategy" otherwise.

Two sentences of [9], Section 1, frame everything that follows. On four or more experts: "Through this dynamic program, we found that the optimal adversary for $k \geq 4$ does not always have a simple description like for $k = 2, 3$. Further, we observe that unlike $k = 2, 3$, for 4 experts, the optimal adversary is already δ-dependent, and its actions at a given configuration of cumulative gains depend on the exact values of cumulative gains and not just their order." And the conjecture: "Nevertheless, inspired by results from this dynamic program, we conjecture that there is a simple adversary ('comb adversary') which is asymptotically (in δ or T) optimal: split experts into two teams $\{1, 3, 5, \dots\}$ and $\{2, 4, 6, \dots\}$ and increment the gains of all experts in exactly one of these teams chosen uniformly at random." Here the experts are numbered by current rank. Their Theorem 5.1 shows that the three-expert comb adversary ("advances experts 1 and 3 together with probability 1/2, and, expert 2 with probability 1/2") inflicts $\frac{2}{3}\cdot\frac{1}{\sqrt{2\delta}}$ as $\delta \to 0$, the same constant as the optimum of Theorem 4.2, and their Theorem 5.2 shows that the four-expert comb adversary inflicts $\frac{\pi}{4}\cdot\frac{1}{\sqrt{2\delta}}$.

Abbasi-Yadkori, Bartlett and Gabillon [1] return to the finite horizon for three experts and prove that "up to a small additive term, COMB is minimax optimal in the finite-time three expert problem." Their Theorem 1 bounds the comb adversary's regret from below by $V_T - 12\log^2(T+1)$, and their Theorem 2 reads: "Let $K = 3$, for even T, we have that

$$\left| V_T - \binom{T+2}{T/2+1}\frac{T/2+1}{3\cdot 2^T} \right| \;\leq\; 12\log^2(T+1), \qquad \text{with} \qquad \binom{T+2}{T/2+1}\frac{T/2+1}{3\cdot 2^T} \;\sim\; \sqrt{\frac{8T}{9\pi}}\text{''}.$$

They also write: "As noted in [9], and noticeable from solving exactly small instances of the game with a computer, in the finite-time case, the exact optimal adversary seems to depend in a complex manner on time and state" — small instances were solved, and no value is printed.

The later literature is asymptotic. Bayraktar, Ekren and Zhang prove the comb conjecture for four experts in the geometric-horizon model [2] and solve the continuum (PDE) limit of the four-expert finite-horizon problem [3]; Drenska and Kohn [8] and Calder, Drenska and Mosaphir [4] study the same PDE limit, the latter recording, from the explicit solution, that for $n = 3$ the comb strategy is globally optimal in that limit (as is one other strategy, which is not the comb). In a short experimental note, Chase [5] compares, for fixed rank-subsets $A \subseteq [k]$, the adversary "giving a gain of 1 to the experts ranked in A and 0 to the rest with probability 1/2, and doing the opposite with probability 1/2", and writes: "It is straightforward to compute that for $T = 5$, stategy $[1, 3]$ is strictly better than strategy

[1, 3, 5]. This is in contrast to the cases of $k = 2, 3, 4$ in which doing comb appears by code to be always exactly optimal (for any T)." (The spelling is the original's.)

What this paper settles. Everything below is at the finite horizon, in the game defined above.

(i) For every $1 \le T \le 12$, $V_T(3) = (4\lceil T/2\rceil + [T \text{ even}])\binom{T}{\lfloor T/2\rfloor}/(3\cdot 2^T) = (2\,\mathbb{E}|S_T| + \mathbb{P}(S_T = 0))/3$ (Theorem 4.1). The scaled integers $3 \cdot 2^T V_T(3) = 4, 10, 24, 54, 120, \ldots$ are the OEIS sequence A152548 [10], whose entry has no connection to learning theory.

(ii) For every even $2 \le T \le 12$ the expression of [1], Theorem 2, exceeds $V_T(3)$ by exactly $\binom{T}{T/2}/(3\cdot 2^T) = \mathbb{P}(S_T = 0)/3$ (Theorem 4.3). This is a refinement of their theorem, not a contradiction: their error bar is $12\log^2(T+1)$.

(iii) For three experts and every $1 \le T \le 12$, the comb adversary — and every adversary that splits the experts into two teams by rank and advances one team chosen by a fair coin, under every tie-breaking rule — inflicts at most $c_T = 1/2, 3/4, 7/8, 17/16, \ldots, 7509/4096$, the comb adversary inflicts exactly c_T against every player, and $c_T < V_T(3)$ with the exact deficit $1/6, 1/12, 1/8, 1/16, 3/32, \ldots, 191/4096$ (Theorem 4.5).

(iv) For four experts and $T = 1, \ldots, 10$ (Theorem 5.1),

$$V_T(4) \;=\; \tfrac{3}{4},\; 1,\; \tfrac{23}{20},\; \tfrac{4}{3},\; \tfrac{29}{20},\; \tfrac{8}{5},\; \tfrac{1283}{750},\; \tfrac{329}{180},\; \tfrac{15433}{8000},\; \tfrac{21923}{10800}.$$

(v) For four experts the two-team class is worth exactly $c_T = \frac{1}{2}, 1, 1, \frac{5}{4}, \frac{11}{8}, \frac{3}{2}, \frac{53}{32}, \frac{111}{64}, \frac{241}{128}, \frac{499}{256}$ for $T = 1, \ldots, 10$; this is strictly below $V_T(4)$ at $T = 1$ and at every $3 \le T \le 10$, and equal to $V_2(4) = 1$ at $T = 2$ (Theorem 5.3).

Items (iii) and (v) separate two readings of Chase's sentence. Read as minimax optimality it fails for $k = 3$ at every horizon considered and for $k = 4$ at every horizon but $T = 2$; read within the class his note compares — adversaries of the "two teams, fair coin" form — it holds, and both readings are certified. Nothing here contradicts the asymptotic statements in print: the deficit at $T = 12$ is $191/4096 \approx 0.047$ against $V_{12}(3) \approx 1.880$, and Section 7 reports it positive but decaying through $T = 60$.

Beyond the certified range, Section 7 reports computations that are exact but not machine-checked: the closed form (i) holds for every $1 \le T \le 60$; the five-expert values $V_T(5)$ for $T \le 10$; and a reproduction of the numbers printed in [5].

Method. The value function is a backward induction over the vector of gaps $d_i = \max_j G_j - G_i$, with a $k \times 2^k$ matrix game at every state. We solved these games exactly over $\mathbb{Q}$ and kept, for every state, a *certificate*: the value, an optimal mixed strategy of the player, and an optimal balanced distribution of the adversary. A player strategy bounds the value from above by a one-step inequality; a balanced adversary bounds it from below by another; and the two bounds meet. The whole argument therefore rests on weak duality for the finite game, which we prove directly by induction with no minimax theorem, plus a finite check of the one-step inequalities at every reachable state. That finite check — all states, all gain vectors, all experts, at every level — is what the machine verification performs (Section 8); it is the only place where exhaustive computation enters a proof. Cover's two-expert values come out of the same machinery, which we use as a control (Section 6).

2. Preliminaries

Notation. $[k] = \{1, \ldots, k\}$; Δ_k is the set of probability vectors on $[k]$; for $g \in \{0,1\}^k$, $\bar{g} = \mathbf{1} - g$ is its complement. A *state* is a pair (n, G) of the number n of rounds left and the vector $G \in \mathbb{Z}^k$ of cumulative gains; the initial state is $(T, 0)$. For a state we write $W_n(G)$ for the minimax expected value of $\max_i G_{i,\text{final}} -$ (the player's gain over the remaining rounds), so that $V_T(k) = W_T(0)$. Player and adversary strategies are behavioural: at each state the player chooses $p \in \Delta_k$ and the adversary a distribution q on $\{0,1\}^k$, both as functions of the state. Since the adversary sees the history and the player's past choices affect nothing but the player's own gain, it is enough (and it is what the certificates exhibit) to let the adversary react to the player's mixed strategy and the player react to the adversary's distribution; the two one-step bounds below are then the two halves of the matrix game at a state.

Gap coordinates. For $G \in \mathbb{Z}^k$ put $d_i = \max_j G_j - G_i \in \mathbb{N}$; then $\min_i d_i = 0$, and W_n depends on G only through d (shifting every G_i by the same c shifts W_n by c). For a gap vector d and a gain vector g, let $\delta(d,g) = 1$ if some leader ($d_i = 0$) receives gain 1 and $\delta(d,g) = 0$ otherwise — the increase of $\max_i G_i$ — and let $d' = d'(d,g)$, $d'_i = d_i + \delta(d,g) - g_i$, be the new gap vector. Writing $W_n(d)$ for the value at gap vector d,

$$
\begin{aligned}
W_0(d) &= 0,\\
W_{n+1}(d) &= \min_{p\in\Delta_k} \max_{g\in\{0,1\}^k} \Big[\delta(d,g) + W_n(d'(d,g)) - \langle p, g\rangle\Big] \\
&= \max_{q} \min_{i\in[k]} \sum_g q_g \Big[\delta(d,g) + W_n(d'(d,g)) - g_i\Big],
\end{aligned}
\tag{1}
$$

the last expression by the minimax theorem for the $k \times 2^k$ matrix game at the state. We never use the minimax theorem: the certificates exhibit both sides. Only the multiset of gaps matters, and the gaps reachable with n rounds left from the initial state of a T-round game are bounded by $T - n$; the number of pairs (rounds left, *sorted* gap vector) reachable in a T-round game with k experts, the terminal level included, is $\binom{T+k}{k}$ — 91, 455 and 1001 for $(k,T) = (2,12), (3,12), (4,10)$.

Balanced adversaries. A distribution q on $\{0,1\}^k$ is *balanced* if $\sum_g q_g g_i$ is the same for every i. Against a balanced adversary every player receives the same expected gain, so the expected regret is player-independent and equals $\mathbb{E}[\max_i G_{i,T}] - \mathbb{E}[\mathrm{mean}_i G_{i,T}]$; [9] prove (their Claim 5) that a minimax-optimal adversary can be taken balanced. The balanced distributions form a polytope Bal_k whose vertices have support at most k; it has 3, 7 and 43 vertices for $k = 2, 3, 4$. The seven vertices for $k = 3$ are the seven "atomic distributions" of [1], Section 5 and Figure 1. Our lower-bound certificate at every state is one of these vertices. (Remark 2 in Appendix B of [9] lists the vertices for $k = 3$ and $k = 4$; its $k = 4$ list has 29 entries besides the trivial $\{\}$ and $\{1234\}$, and omits the twelve vertices of the form $\{1\}\{134\}\{23\}\{24\}$ — one expert alone, the same expert together with two others, and the fourth expert paired with each of those two, each with probability 1/4 — which are balanced, have support 4, and are extreme. Nothing here depends on that list; the vertex sets used are computed and stored as data.)

Two-team adversaries and the comb adversary. An adversary is a *two-team adversary* if at every state it chooses a set $A \subseteq [k]$ (as a function of the state) and plays the gain vector $\mathbf{1}_A$ or its complement $\mathbf{1}_{[k]\setminus A}$ with probability 1/2 each. Every two-team adversary is balanced (each expert receives expected gain 1/2 per round). The class contains every adversary of the form compared by [5] — a fixed rank-subset A applied at every state — under every rule for breaking ties in the ranking, and, because the player's expected regret is linear in the adversary's one-round distribution, its value is unchanged by mixing such adversaries within a round. Write

$$c_T(k) \;=\; \min_{\text{player}} \; \max_{\text{two-team adversaries}} \; \mathbb{E}[R_T]$$

for the value of the class; since each member is balanced, this is also $\max_{\text{two-team}} \mathbb{E}[\max_i G_{i,T}] - T/2$, and the player is irrelevant. The *comb adversary* of [9] is the two-team adversary with A the set of experts of odd rank ($1, 3, 5, \ldots$, the leader having rank 1); in the formal statements, ties in the ranking are broken by the expert's index. For $k = 3$ this is the adversary of [9], Theorem 5.1 — $\{13\}\{2\}$ in their bracket notation — and for $k = 4$ it is the adversary of their Theorem 5.2, $\{13\}\{24\}$. For $k = 2$ the comb adversary is Cover's adversary.

Random-walk quantities. For a simple symmetric random walk S on $\mathbb{Z}$ started at 0,

$$\mathbb{E}|S_T| = \frac{2\lceil T/2\rceil\binom{T}{\lfloor T/2\rfloor}}{2^T}, \qquad \mathbb{P}(S_T = 0) = [T \text{ even}]\,\frac{\binom{T}{T/2}}{2^T};$$

these classical closed forms are how the random-walk quantities enter the formal statements. Cover's value is $V_T(2) = \mathbb{E}|S_T|/2$.

3. The certificates

The two halves of (1) become two predicates on states, and the comb comparison adds two more. For $n \in \mathbb{N}$, $G \in \mathbb{Z}^k$, $v \in \mathbb{Q}$, define by induction on n:

- $\mathrm{Guar}(n, G, v)$: $n = 0$ and $\max_i G_i \le v$; or $n = m + 1$ and there is $p \in \Delta_k$ with $\mathrm{Guar}(m, G + g, v + \langle p, g\rangle)$ for every $g \in \{0, 1\}^k$. (The player can guarantee at most v against every adversary.)
- $\mathrm{Force}(n, G, v)$: $n = 0$ and $v \le \max_i G_i$; or $n = m+1$ and there are a distribution q on $\{0, 1\}^k$ and values w_g with $\mathrm{Force}(m, G + g, w_g)$ for every g and $v \le \sum_g q_g(w_g - g_i)$ for every i. (The adversary can force at least v against every player.)
- $\mathrm{AntiHold}(n, G, v)$: as Guar, but the one-step condition is $\frac{1}{2}(w_g - \langle p, g\rangle) + \frac{1}{2}(w_{\bar g} - \langle p, \bar g\rangle) \le v$ for every g, with $\mathrm{AntiHold}(m, G + g, w_g)$. (The player holds every two-team adversary to at most v.)
- $\mathrm{Play}_\pi(n, G, v)$, for a fixed adversary strategy π assigning to each state a distribution on $\{0, 1\}^k$: $n = 0$ and $v = \max_i G_i$; or $n = m + 1$ and there are w_g with $\mathrm{Play}_\pi(m, G + g, w_g)$ and $v = \sum_g \pi(n, G)_g(w_g - g_i)$ for every i. (Against π the expected regret is exactly v, whatever the player does.)

$\mathrm{IsValue}(n, G, v)$ means $\mathrm{Guar}(n, G, v) \wedge \mathrm{Force}(n, G, v)$.

Proposition 3.1 (Weak duality and its consequences)**.** *For all n, G, v, v':*

(a) $\mathrm{Guar}(n, G, v)$ *and* $\mathrm{Force}(n, G, v')$ *imply* $v' \le v$*; hence* $\mathrm{IsValue}(n, G, v)$ *and* $\mathrm{IsValue}(n, G, v')$ *imply* $v = v'$.

(b) *If π is a probability distribution at every state,* $\mathrm{Play}_\pi(n, G, v)$ *implies* $\mathrm{Force}(n, G, v)$*; and* $\mathrm{Play}_\pi(n, G, v)$, $\mathrm{Play}_\pi(n, G, v')$ *imply* $v = v'$.

(c) *If π is a two-team adversary,* $\mathrm{AntiHold}(n, G, v)$ *and* $\mathrm{Play}_\pi(n, G, c)$ *imply* $c \le v$.

(d) *Each of the four predicates is invariant under shifting G by a constant c and v by c.*

Proof sketch. (a) is the standard exchange of the two sums: with p from Guar and (q, w) from Force, $v' = \sum_i p_i v' \le \sum_i p_i \sum_g q_g(w_g - g_i) = \sum_g q_g(w_g - \langle p, g\rangle) \le \sum_g q_g v = v$, the last step by induction. (b) and (d) are immediate inductions. (c) averages the exact one-step identity of Play_π over the player's mixed strategy from AntiHold and uses that $\pi(n, G)$ is supported on a complementary pair. No minimax theorem, no linear-programming duality and no game-value API is used — none is available in the library we build on, and none is needed. □

Proposition 3.2 (From a finite table to the four predicates)**.** *Let $T \in \mathbb{N}$ and let $W, C : \mathbb{N} \times \mathbb{N}^k \to \mathbb{Q}$, $P : \mathbb{N} \times \mathbb{N}^k \to \mathbb{Q}^k$, $Q : \mathbb{N} \times \mathbb{N}^k \to \mathbb{Q}^{\{0,1\}^k}$ be tables indexed by (rounds left, gap vector). Suppose that for every $n < T$ and every gap vector d with $\max_i d_i \le T - n - 1$ and $\min_i d_i = 0$:*

(G) *$P(n+1, d) \in \Delta_k$ and, for every g, $\delta(d, g) + W(n, d'(d, g)) \le W(n + 1, d) + \langle P(n + 1, d), g\rangle$;*

(F) *$Q(n+1, d)$ is a probability distribution and, for every i, $W(n+1, d) \le \sum_g Q(n+1, d)_g\big(\delta(d, g)+ W(n, d'(d, g)) - g_i\big)$;*

(A) *for every g, $\frac{1}{2}\big(\delta(d, g) + C(n, d'(d, g))\big) + \frac{1}{2}\big(\delta(d, \bar g) + C(n, d'(d, \bar g))\big) - \frac{1}{2} \le C(n + 1, d)$;*

(P) *for a strategy π that is a probability distribution at every state, and every i, $C(n + 1, d) = \sum_g \pi(n + 1, d)_g\big(\delta(d, g) + C(n, d'(d, g)) - g_i\big)$;*

and that $W(0, d) = C(0, d) = 0$ for every d with $\max_i d_i \le T$, $\min_i d_i = 0$. Then for every $n \le T$ and every such d with $\max_i d_i \le T - n$: (G) gives $\mathrm{Guar}(n, -d, W(n, d))$*, (F) gives* $\mathrm{Force}(n, -d, W(n, d))$*, (A) gives* $\mathrm{AntiHold}(n, -d, C(n, d))$*, and (P) gives* $\mathrm{Play}_\pi(n, -d, C(n, d))$.

Proof sketch. Induction on n, using Proposition 3.1(d) to pass from the absolute state $-d+g$ to the normalised gap vector $d'(d,g)$ shifted by $\delta(d,g)$. In (A) the player's mixed strategy can be taken uniform: against a complementary pair the player's average gain is $1/2$ whatever p is, which is the balanced-adversary property in this special case. $\square$

Propositions 3.1 and 3.2 are routine; they are stated because every value in this paper rests on them and because the hypotheses of Proposition 3.2 are exactly what was computed. The tables are stored on *sorted* gap vectors only (455 rows for $k=3$, 1001 for $k=4$), and a lookup sorts its argument, so the finite check ranges over every gap vector in $\{0,\dots,T\}^k$ with a zero entry at every level ($12\cdot 13^3$ candidate pairs (n,d) for $k=3$, $10\cdot 11^4$ for $k=4$, each against all 2^k gain vectors and all k experts) while the data stays the size of the sorted state space. The adversary certificate at a state is a single index into the list of vertices of Bal_k.

4. Three experts

Theorem 4.1 (The exact value, $k=3$)**.** *For every* $1\le T\le 12$*,*

$$V_T(3) \;=\; \frac{\bigl(4\lceil T/2\rceil + [T\ \mathit{even}]\bigr)\binom{T}{\lfloor T/2\rfloor}}{3\cdot 2^T} \;=\; \frac{2\,\mathbb{E}|S_T| + \mathbb{P}(S_T=0)}{3}.$$

In particular $V_T(3) = 2/3, 5/6, 1, 9/8, 5/4, 65/48, 35/24, 595/384, 105/64, 441/256, 231/128, 1925/1024$ *for* $T=1,\dots,12$.

Proof sketch. Let W,P,Q be the computed tables for $k=3$, $T=12$. The hypotheses (G) and (F) of Proposition 3.2, together with the base case, were checked exhaustively at every level $n<12$, every gap vector $d\in\{0,\dots,12\}^3$ with a zero entry and $\max d\le 11-n$, every $g\in\{0,1\}^3$ and every $i\in[3]$; so $\mathrm{IsValue}(T,0,W(T,0))$ for every $T\le 12$. A second finite check identifies $W(T,0)$ with the displayed closed form for $1\le T\le 12$, and a third identifies the closed form with $(2\,\mathbb{E}|S_T|+\mathbb{P}(S_T=0))/3$ through the random-walk closed forms. By Proposition 3.1(a) the value is unique. $\square$

Remark 4.2. The closed form holds at the *initial* state only. Values at other states of the same certified table have denominators with high powers of 3: $W_6(0,2,3) = 1147/3888 = 1147/(2^4\cdot 3^5)$, $W_6(0,4,4) = 235/3888$, $W_5(0,3,3) = 77/648 = 77/(2^3\cdot 3^4)$, $W_6(0,2,2) = 173/432$, against $3\cdot 2^T$ at the initial state. So the usual route to a proof for all T — guess a closed form for the whole value function and verify (1) symbolically — is closed: there is nothing to guess. A proof for all T would have to be an argument about the initial state alone; see Section 9.

Theorem 4.3 (The published approximation against the exact value)**.** *Let* $A(T) = \binom{T+2}{T/2+1}(T/2+1)/(3\cdot 2^T)$ *be the expression of* [1]*, Theorem 2. For every even* T *with* $2\le T\le 12$*,*

$$A(T) - V_T(3) \;=\; \frac{\binom{T}{T/2}}{3\cdot 2^T} \;=\; \frac{\mathbb{P}(S_T=0)}{3}.$$

Proof sketch. A finite check of the identity $A(T) - \mathrm{cf}(T) = \binom{T}{T/2}/(3\cdot 2^T)$ for even $T\le 12$, where cf is the closed form of Theorem 4.1, which is $V_T(3)$ by that theorem. $\square$

Remark 4.4. For $T=2m$ one has $\binom{2m+2}{m+1}(m+1) = 2(2m+1)\binom{2m}{m}$, so $A(T) = (2\,\mathbb{E}|S_T| + 2\,\mathbb{P}(S_T=0))/3$ for every even T, while Theorem 4.1 gives $(2\,\mathbb{E}|S_T|+\mathbb{P}(S_T=0))/3$: the published expression carries the term $\mathbb{P}(S_T=0)/3$ twice. (This elementary identity is not part of the formal development; the theorem is.) At $T=4$, $A(4)=5/4$ against $V_4(3)=9/8$; at $T=12$, $A(12)=1001/512$ against $1925/1024$. Both discrepancies are far inside the bound $12\log^2(T+1)$ of [1], so Theorem 4.3 sharpens their Theorem 2 and contradicts nothing; the leading constant $\sqrt{8T/(9\pi)}$ is common to both expressions. For odd T no approximating expression is in print.

Theorem 4.5 (The comb adversary is strictly suboptimal, $k=3$)**.** *For every* $1\le T\le 12$ *let* c_T *be the* T*-th entry of*

$$\tfrac{1}{2},\ \tfrac{3}{4},\ \tfrac{7}{8},\ \tfrac{17}{16},\ \tfrac{37}{32},\ \tfrac{83}{64},\ \tfrac{177}{128},\ \tfrac{383}{256},\ \tfrac{807}{512},\ \tfrac{1713}{1024},\ \tfrac{3579}{2048},\ \tfrac{7509}{4096}.$$

Then, in the three-expert T-round game:

(a) *the player has a strategy whose expected regret against every two-team adversary is at most c_T;*
(b) *against the comb adversary, every player strategy has expected regret exactly c_T;*
(c) *c_T is the least number for which (a) holds; hence $c_T = c_T(3)$ is the exact value of the two-team class, and the comb adversary is an optimal member of it;*
(d) *$c_T < V_T(3)$, with $V_T(3) - c_T = \frac{1}{6}, \frac{1}{12}, \frac{1}{8}, \frac{1}{16}, \frac{3}{32}, \frac{11}{192}, \frac{29}{384}, \frac{41}{768}, \frac{33}{512}, \frac{51}{1024}, \frac{117}{2048}, \frac{191}{4096}$ for $T = 1, \ldots, 12$.*

In particular no two-team adversary — the comb adversary under any tie-breaking rule included — is minimax optimal for three experts at any horizon $1 \le T \le 12$.

Proof sketch. Let C be the computed table of two-team values for $k = 3$, $T = 12$, and π the comb strategy. Hypotheses (A) and (P) of Proposition 3.2 were checked exhaustively over the same range as in Theorem 4.1; this gives $\mathrm{AntiHold}(T, 0, C(T, 0))$ and $\mathrm{Play}_\pi(T, 0, C(T, 0))$, which are (a) and (b) with $c_T = C(T, 0)$ (a finite check identifies $C(T, 0)$ with the list). (c) is Proposition 3.1(c). (d) is a finite comparison of c_T with the closed form of Theorem 4.1, and the differences are a finite check. The last sentence follows from (a) and (d): any two-team adversary is held to $c_T < V_T(3)$. □

T	$V_T(3)$	$3 \cdot 2^T V_T(3)$	$c_T(3)$	$V_T(3) - c_T(3)$	$A(T)$
1	2/3	4	1/2	1/6	—
2	5/6	10	3/4	1/12	1
3	1	24	7/8	1/8	—
4	9/8	54	17/16	1/16	5/4
5	5/4	120	37/32	3/32	—
6	65/48	260	83/64	11/192	35/24
7	35/24	560	177/128	29/384	—
8	595/384	1190	383/256	41/768	105/64
9	105/64	2520	807/512	33/512	—
10	441/256	5292	1713/1024	51/1024	231/128
11	231/128	11088	3579/2048	117/2048	—
12	1925/1024	23100	7509/4096	191/4096	1001/512

TABLE 1. Three experts. $V_T(3)$ and $c_T(3)$ are certified (Theorems 4.1 and 4.5); the third column is OEIS A152548; $A(T)$ is the expression of [1], Theorem 2, defined for even T.

Remark 4.6. The exact deficit is not monotone in T ($1/16 < 3/32$, $41/768 < 33/512$), and its numerators $1, 1, 1, 1, 3, 11, 29, 41, 33, 51, 117, 191$ do not appear in the OEIS. Relative to the value, the deficit is 25% at $T = 1$ and about 2.5% at $T = 12$.

5. Four experts

Theorem 5.1 (The exact value, $k = 4$)**.** *For $T = 1, \ldots, 10$,*

$$V_T(4) \;=\; \tfrac{3}{4},\; 1,\; \tfrac{23}{20},\; \tfrac{4}{3},\; \tfrac{29}{20},\; \tfrac{8}{5},\; \tfrac{1283}{750},\; \tfrac{329}{180},\; \tfrac{15433}{8000},\; \tfrac{21923}{10800}.$$

Proof sketch. As for Theorem 4.1, with the computed tables for $k = 4$, $T = 10$: hypotheses (G) and (F) of Proposition 3.2 were checked exhaustively at every level $n < 10$, every gap vector $d \in \{0, \ldots, 10\}^4$ with a zero entry and $\max d \le 9 - n$, every $g \in \{0, 1\}^4$ and every $i \in [4]$, and $W(T, 0)$ was identified with the list by a finite check. The adversary certificate at each state is one of the 43 vertices of Bal_4. □

Remark 5.2. The denominators carry 5^3 and 3^3, so no closed form of the shape of Theorem 4.1 is to be expected; neither the numerators nor the denominators appear in the OEIS. This is the finite-horizon counterpart of the sentence of [9] quoted in the introduction: for four experts the optimal adversary's "actions at a given configuration of cumulative gains depend on the exact values of cumulative gains and not just their order."

Theorem 5.3 (The comb adversary, $k = 4$)**.** *For $T = 1, \dots, 10$ let c_T be the T-th entry of*

$$\tfrac{1}{2},\ 1,\ 1,\ \tfrac{5}{4},\ \tfrac{11}{8},\ \tfrac{3}{2},\ \tfrac{53}{32},\ \tfrac{111}{64},\ \tfrac{241}{128},\ \tfrac{499}{256}.$$

Then, in the four-expert T-round game, for $T = 1$ and for every $3 \le T \le 10$:

(a) *the player has a strategy whose expected regret against every two-team adversary is at most c_T;*
(b) *against the comb adversary every player strategy has expected regret exactly c_T;*
(c) *c_T is the least number for which (a) holds, so $c_T = c_T(4)$ and the comb adversary is an optimal member of the two-team class;*
(d) *$c_T < V_T(4)$.*

The exact deficits $V_T(4) - c_T$ for $T = 1, \dots, 10$ are

$$\tfrac{1}{4},\ 0,\ \tfrac{3}{20},\ \tfrac{1}{12},\ \tfrac{3}{40},\ \tfrac{1}{10},\ \tfrac{653}{12000},\ \tfrac{269}{2880},\ \tfrac{741}{16000},\ \tfrac{13943}{172800}.$$

At $T = 2$ the deficit is zero: $c_2 = V_2(4) = 1$, so the comb adversary is exactly minimax optimal for four experts at horizon two.

Proof sketch. As for Theorem 4.5, with the tables for $k = 4$, $T = 10$. The strict inequality (d) is a finite check for $T \neq 2$; the equality at $T = 2$ is a separate finite check of the two list entries, and statements (a) and (b) at $T = 2$ come from the same table lemmas as for the other horizons. The exclusion of $T = 2$ from (d) is a fact about the game, not an artefact of the certificate: the equality is certified, not merely unproved. $\square$

T	$V_T(4)$	$c_T(4)$	$V_T(4) - c_T(4)$	$V_T(4)$ (decimal)
1	3/4	1/2	1/4	0.75
2	1	1	0	1
3	23/20	1	3/20	1.15
4	4/3	5/4	1/12	1.333...
5	29/20	11/8	3/40	1.45
6	8/5	3/2	1/10	1.6
7	1283/750	53/32	653/12000	1.7106...
8	329/180	111/64	269/2880	1.8277...
9	15433/8000	241/128	741/16000	1.929125
10	21923/10800	499/256	13943/172800	2.0299...

TABLE 2. Four experts. All entries are certified (Theorems 5.1 and 5.3).

Remark 5.4 (Chase's sentence). Theorems 4.5(c) and 5.3(c) say that, within the class of two-team adversaries, the comb adversary is exactly optimal for $k = 3$ ($T \le 12$) and $k = 4$ ($T \le 10$); this is the reading under which the sentence of [5] quoted in the introduction — "doing comb appears by code to be always exactly optimal" — is true, and it is the class his note compares within. Theorems 4.5(d) and 5.3(d) say that as a minimax adversary the comb adversary is strictly suboptimal at every horizon considered except $(k, T) = (4, 2)$. The two statements are about different comparison classes and both are certified; no erratum is being claimed. (For $k = 2$ the two classes give the same value at every $T \le 12$: Proposition 6.1.)

6. Controls

The machinery of Section 3 is run at $k = 2$ against a published theorem before it is trusted at $k = 3, 4$.

Proposition 6.1 (Cover's theorem as a control)**.** *For every $1 \le T \le 12$, $V_T(2) = \lceil T/2 \rceil \binom{T}{\lfloor T/2 \rfloor}/2^T$ (Cover [7]), certified by the same tables and the same Proposition 3.2 as Theorems 4.1 and 5.1. Moreover, at $k = 2$ the comb adversary is exactly minimax optimal at every such T: against it every player strategy has expected regret exactly $V_T(2)$, and the player holds every two-team adversary to $V_T(2)$.*

The second half is the control for Theorems 4.5 and 5.3: the same formalisation of the two-team class that separates at $k = 3, 4$ is tight at $k = 2$, so the separations are not an artefact of how the class is defined. The value $V_T(2)$ is Cover's; only its certification here is ours.

Proposition 6.2 (Negative controls)**.** *For $k = 3$ and $1 \le T \le 12$, and for $k = 4$ and $1 \le T \le 10$, with v_T the value of Theorem 4.1 or 5.1:*

(a) *no player strategy guarantees expected regret $< v_T$ against every adversary, and no adversary forces expected regret $> v_T$ against every player;*
(b) *for $k = 3$, no player strategy holds every two-team adversary below $c_T(3)$;*
(c) *against the comb adversary the expected regret is not v_T (for $k = 4$, $T \neq 2$);*
(d) *the normalisation hypothesis $\min_i d_i = 0$ carried by every table lemma is satisfied by the initial state, and $v_T > 0$, $c_T(3) > 0$ throughout.*

These are corollaries of Proposition 3.1 and the two-sided certificates; they are recorded so that "the value is v_T" cannot be read as "v_T is one bound among many", and so that no statement above is about an empty set of states or the trivial value.

7. Computations outside the formal development

Everything in this section was computed in exact rational arithmetic and is *not* machine-checked; it is reported as computation.

How the tables were produced, and checked. Two independent implementations produced the values. The first is the backward induction (1) itself: at every state the $k \times 2^k$ matrix game is solved by an exact simplex over $\mathbb{Q}$ (Bland's rule), and every solve is self-certified before its value is used — the primal and dual solutions are checked to satisfy $\sum_i p_i M_{ig} \le v$ for all g and $\sum_g q_g M_{ig} \ge v$ for all i — so its output depends on the correctness of the certificate check, not of the simplex. The second uses the structural reduction of [9]: since a minimax-optimal adversary is balanced and the regret against a balanced adversary is the player-free quantity $\mathbb{E}[\max_i G_{i,T} - \mathrm{mean}_i G_{i,T}]$, it maximises that expectation over the vertices of Bal_k by a different recursion with no matrix game at all. The two agree at $k = 2$ ($T \le 10$), $k = 3$ ($T \le 16$), $k = 4$ ($T \le 13$) and $k = 5$ ($T \le 8$). A third check solved 120 randomly chosen states of the $k = 3$, $T = 9$ run by support enumeration with exact Gaussian elimination, with no mismatch. The first implementation's per-state certificates are the tables of Section 3. Producing and re-verifying them took 2 seconds for $k = 3$, $T = 12$, and 22 seconds for $k = 4$, $T = 10$.

Calibration against printed numbers. Besides Cover's values (Proposition 6.1), the code reproduces: the seven atomic distributions of [1], Section 5 and Figure 1, as the vertex set of Bal_3; their Lemma 5 (the comb and twin-comb adversaries $\{13\}\{2\}$ and $\{1\}\{23\}$ have the same value); and the two computations printed in [5]. For $k = 5$, $T = 5$, the rank-subset $[1, 3]$ inflicts $25/16$ against $49/32$ for $[1, 3, 5]$, strictly more as he says, and the two agree for $T \le 4$. For $k = 6$, $T = 13$, he prints 9.14453125 for the adversary allowed to choose between $[1, 3, 6]$ and $[1, 4, 6]$ at each state and 9.143310546875 for the best single subset; we obtain $677/256 = 2.64453125$ and $10827/4096 = 2.643310546875$, each his number minus $T/2 = 6.5$: his "expected regret" is $\mathbb{E}[\max_i G_{i,T}]$, which against a balanced adversary exceeds the regret by exactly $T/2$, and the constant cancels in every comparison he makes. Every digit matches, and $[1, 3, 6]$ is the best single subset as he reports. Both tie-breaking conventions for the ranking (adversarially best and uniformly random) were run and give the same comb values as the index tie-break of the formal statements.

The closed form beyond $T = 12$. The closed form of Theorem 4.1 holds for every $1 \le T \le 60$: the backward induction was carried to $T = 60$ in 57 CPU-seconds, over the $\binom{63}{3} = 39\,711$ pairs (rounds left, sorted gap vector) of all levels, each of the $37\,820$ with a round left carrying a verified primal/dual pair as above, so the values are certificate-checked even though the check ran outside the formal development. Over the same range the deficit $V_T(3) - c_T(3)$ of Theorem 4.5 stays strictly positive. Equivalently, $3 \cdot 2^T V_T(3) = a(T)$ for $1 \le T \le 60$, where a is OEIS A152548 [10] ("Sum of squared terms in rows of triangle A152547", with offset 0 and $a(0) = 1$, so the identity fails at $T = 0$ where $V_0 = 0$); the entry gives the generating function $\sqrt{(1 + 2x)/(1 - 2x)^3}$, the asymptotic $a(n) \sim 2^{n+3/2}\sqrt{n/\pi}$, which is $V_T(3) \sim \sqrt{8T/(9\pi)}$, and the recurrence $(n + 1)a(n + 1) = 4a(n) + 4n\, a(n - 1)$. Whether the identity holds for all T is open (Question 9.1).

Five experts. The same computation gives, for $T = 1, \ldots, 10$,

$$V_T(5) \;=\; \tfrac{4}{5},\ \tfrac{11}{10},\ \tfrac{5}{4},\ \tfrac{103}{70},\ \tfrac{193}{120},\ \tfrac{53}{30},\ \tfrac{2117}{1120},\ \tfrac{13351}{6615},\ \tfrac{150359}{70560},\ \tfrac{667189}{297675},$$

computed by the self-certified simplex route to $T = 10$ (21 CPU-seconds) and by the balanced-vertex route to $T = 8$, the two agreeing where both ran. These values are not certified in the formal development; the cost of doing so is estimated in Section 9. Chase's evidence against the comb conjecture is at $k = 5$, which makes this table the natural next datum. For four experts the same computation reaches $T = 20$ in 30 seconds, with $V_{20}(4) = 70531379980777/24883200000000$.

8. Formal verification

Every theorem and proposition of Sections 3–6 is machine-checked in Lean 4 [11] (v4.32.0) with Mathlib [12]; the repository reference is to be added at submission. The game is formalised as the four inductive predicates of Section 3 on (rounds left, cumulative gains, a rational); Proposition 3.1 and Proposition 3.2 are proved from first principles, and the fifteen declarations behind them depend on the axioms `propext`, `Classical.choice` and `Quot.sound` only. The tables of rationals are Lean data, and the hypotheses of Proposition 3.2 — the base case, that P and Q are probability distributions, and the one-step conditions (G), (F), (A), (P) — are each discharged by a single `native_decide` ranging over every level, every gap vector in $\{0, \ldots, T\}^k$ with a zero entry, every gain vector and every expert; the identifications of the table entries with the closed forms, lists and differences stated in the theorems are further `native_decide` checks. Accordingly, `#print axioms` reports for each value theorem the three standard axioms plus the native-decision axioms of its own finite checks, named `<check>._native.native_decide.ax_1_1`: for Theorem 4.1 those of `base3`, `fstep3`, `gstep3`, `k3_closed_form`, `p3_ok`, `q3_ok` (and, for its random-walk form, `k3_random_walk`); for Theorem 4.3, `k3_abg_excess`; for Theorem 4.5, `astep3`, `base3cv`, `k3_comb_strict`, `k3_comb_values`, `pstep3` (and `k3_comb_gap`); for Theorem 5.1, `base4`, `fstep4`, `gstep4`, `k4_values`, `p4_ok`, `q4_ok`; for Theorem 5.3, `astep4`, `base4cv`, `k4_comb_strict`, `k4_comb_values`, `pstep4` (and `k4_comb_tight_at_two`, `k4_comb_gap`); for Proposition 6.1, `base2`, `fstep2`, `gstep2`, `k2_cover_values`, `p2_ok`, `q2_ok`, `astep2`, `base2cv`, `k2_comb_tight`, `pstep2`; the controls of Proposition 6.2 inherit the axioms of the theorems they are corollaries of, and `hyp_not_vacuous` carries three of its own. No `sorry` is used anywhere. The $k = 4$ checks take about five minutes; the $k = 3$ checks about twenty seconds. Appendix A lists the formal statements verbatim. Nothing in Section 7 is formalised.

9. Open questions

Question 9.1. Does $V_T(3) = (4\lceil T/2\rceil + [T \text{ even}])\binom{T}{\lfloor T/2\rfloor}/(3 \cdot 2^T)$ hold for every $T \ge 1$? It holds for $T \le 12$ by Theorem 4.1 and for $T \le 60$ by computation. As the remark after Theorem 4.1 explains, the value function away from the initial state has no such shape, so a proof cannot verify (1) by a symbolic ansatz; the recurrence $(n + 1)a(n + 1) = 4a(n) + 4n\, a(n - 1)$ of A152548 is the shape an argument would have to produce. Extending the certified range is cheap by comparison — the finite check scales as $T(T+1)^3 2^3$, about 3 times the $T = 12$ cost at $T = 16$ and 7 times at $T = 20$ — but no finite range settles the question.

Question 9.2. Four experts beyond $T = 10$, and five experts. The exact values are cheap to compute ($T \le 20$ for $k = 4$ in 30 seconds); the certification cost scales as $T(T+1)^k 2^k$, which against the measured $k = 4$, $T = 10$ check gives about 2.3 times at $T = 12$ and 3.4 times at $T = 13$; for $k = 5$ about 1.4 times at $T = 6$, 6.5 times at $T = 8$ and 22 times at $T = 10$. The five-expert values of Section 7 are therefore parked for time, not for difficulty.

Question 9.3. How does the deficit $V_T(k) - c_T(k)$ decay? For $k = 3$ it is 1/6, 1/12, 1/8, 1/16, 3/32, 11/192, …, not monotone, positive through $T = 60$ by computation, and about 2.5% of the value at $T = 12$. Since the comb adversary is asymptotically optimal for $k = 3$ ([9], Theorems 4.2 and 5.1) and optimal up to $12\log^2(T+1)$ at finite horizon ([1], Theorem 1), the deficit is $o(\sqrt{T})$; is it $O(1)$, or does it grow?

Question 9.4. What is the finite-horizon optimal adversary for four experts? [9] report, for the geometric horizon, that it depends on the exact gaps and not only on their order. The certificate tables behind Theorem 5.1 record an optimal balanced vertex at every state of every game with $T \le 10$; reading a pattern off them is a question in its own right.

Appendix A. The formal statements

The following are the Lean statements, verbatim, of the declarations that carry the results above. Throughout, ι is a finite nonempty type of experts (`variable {ι : Type*} [Fintype ι] [DecidableEq ι] [Nonempty ι]`); `gq g i` is the gain $g_i \in \{0, 1\}$ as a rational; `cneg g` is the complement of `g`; `addG G g` is $G + g$; `shiftG G c` is $G + c\mathbf{1}$; `top G` is $\max_i G_i$; `gapsOf G` is the gap vector; `antiDist a` plays `a` or `cneg a` with probability 1/2 each. `T3 = 12`, `T4 = 10`; `W3`, `CV3`, `P3`, `Q3` (and the `4` versions) are the tables W, C, P, Q of Section 3 read from the stored sorted rows; `pi3`, `pi4`, `pi2` are the comb strategy `fun _ d => antiDist (combMark d)`, where `combMark d` marks the experts of even 0-based rank with ties broken by index.

The predicates.

```
def Guar : ℕ → (ι → ℤ) → ℚ → Prop
  | 0, G, v => (top G : ℚ) ≤ v
  | n + 1, G, v => ∃ p : ι → ℚ, (∀ i, 0 ≤ p i) ∧ (∑ i, p i = 1) ∧
      ∀ g : ι → Bool, Guar n (addG G g) (v + ∑ i, p i * gq g i)

def Force : ℕ → (ι → ℤ) → ℚ → Prop
  | 0, G, v => v ≤ (top G : ℚ)
  | n + 1, G, v => ∃ q : (ι → Bool) → ℚ, ∃ w : (ι → Bool) → ℚ,
      (∀ g, 0 ≤ q g) ∧ (∑ g, q g = 1) ∧ (∀ g, Force n (addG G g) (w g)) ∧
      (∀ i, v ≤ ∑ g, q g * (w g - gq g i))

def IsValue (n : ℕ) (G : ι → ℤ) (v : ℚ) : Prop := Guar n G v ∧ Force n G v

def AntiHold : ℕ → (ι → ℤ) → ℚ → Prop
  | 0, G, v => (top G : ℚ) ≤ v
  | n + 1, G, v => ∃ p : ι → ℚ, (∀ i, 0 ≤ p i) ∧ (∑ i, p i = 1) ∧
      ∃ w : (ι → Bool) → ℚ, (∀ g, AntiHold n (addG G g) (w g)) ∧
      ∀ g : ι → Bool,
        (w g - ∑ i, p i * gq g i) / 2 + (w (cneg g) - ∑ i, p i * gq (cneg g) i) / 2 ≤ v

def Play (π : ℕ → (ι → ℕ) → ((ι → Bool) → ℚ)) : ℕ → (ι → ℤ) → ℚ → Prop
  | 0, G, v => v = (top G : ℚ)
  | n + 1, G, v => ∃ w : (ι → Bool) → ℚ, (∀ g, Play π n (addG G g) (w g)) ∧
      ∀ i, v = ∑ g, π (n + 1) (gapsOf G) g * (w g - gq g i)
```

Proposition 3.1.

```
theorem force_le_guar :
    ∀ (n : ℕ) (G : ι → ℤ) (v v' : ℚ), Guar n G v → Force n G v' → v' ≤ v

theorem IsValue.unique {n : ℕ} {G : ι → ℤ} {v v' : ℚ}
    (h : IsValue n G v) (h' : IsValue n G v') : v = v'
```

```
theorem Play.force (π : ℕ → (ι → ℕ) → ((ι → Bool) → ℚ))
    (h0 : ∀ n d g, 0 ≤ π n d g) (h1 : ∀ n d, ∑ g, π n d g = 1) :
    ∀ (n : ℕ) (G : ι → ℤ) (v : ℚ), Play π n G v → Force n G v

theorem Play.unique (π : ℕ → (ι → ℕ) → ((ι → Bool) → ℚ)) :
    ∀ (n : ℕ) (G : ι → ℤ) (v v' : ℚ), Play π n G v → Play π n G v' → v = v'

theorem play_le_antihold (π : ℕ → (ι → ℕ) → ((ι → Bool) → ℚ))
    (hanti : ∀ n d, ∃ a : ι → Bool, π n d = antiDist a) :
    ∀ (n : ℕ) (G : ι → ℤ) (v c : ℚ), AntiHold n G v → Play π n G c → c ≤ v

theorem guar_shift : ∀ (n : ℕ) (G : ι → ℤ) (c : ℤ) (v : ℚ),
    Guar n (shiftG G c) v ↔ Guar n G (v - c)

theorem force_shift : ∀ (n : ℕ) (G : ι → ℤ) (c : ℤ) (v : ℚ),
    Force n (shiftG G c) v ↔ Force n G (v - c)

theorem antihold_shift : ∀ (n : ℕ) (G : ι → ℤ) (c : ℤ) (v : ℚ),
    AntiHold n (shiftG G c) v ↔ AntiHold n G (v - c)

theorem play_shift (π : ℕ → (ι → ℕ) → ((ι → Bool) → ℚ)) (h1 : ∀ n d, ∑ g, π n d g = 1) :
    ∀ (n : ℕ) (G : ι → ℤ) (c : ℤ) (v : ℚ), Play π n (shiftG G c) v ↔ Play π n G (v - c)

lemma sum_antiDist (a : ι → Bool) : ∑ g, antiDist a g = 1

lemma sum_antiDist_mul (a : ι → Bool) (h : (ι → Bool) → ℚ) :
    ∑ g, antiDist a g * h g = h a / 2 + h (cneg a) / 2
```

Proposition 3.2.

```
theorem guar_of_table (T : ℕ) (W : ℕ → (ι → ℕ) → ℚ) (P : ℕ → (ι → ℕ) → ι → ℚ)
    (hbase : ∀ d : ι → ℕ, (∀ i, d i ≤ T) → (∃ i, d i = 0) → 0 ≤ W 0 d)
    (hp0 : ∀ n, n < T → ∀ d : ι → ℕ, (∀ i, d i ≤ T - (n + 1)) → (∃ i, d i = 0) →
        ∀ i, 0 ≤ P (n + 1) d i)
    (hp1 : ∀ n, n < T → ∀ d : ι → ℕ, (∀ i, d i ≤ T - (n + 1)) → (∃ i, d i = 0) →
        ∑ i, P (n + 1) d i = 1)
    (hstep : ∀ n, n < T → ∀ d : ι → ℕ, (∀ i, d i ≤ T - (n + 1)) → (∃ i, d i = 0) →
        ∀ g : ι → Bool,
          W n (nxt d g) + (dlt d g : ℚ) ≤ W (n + 1) d + ∑ i, P (n + 1) d i * gq g i) :
    ∀ n, n ≤ T → ∀ d : ι → ℕ, (∀ i, d i ≤ T - n) → (∃ i, d i = 0) →
      Guar n (stG d) (W n d)

theorem force_of_table (T : ℕ) (W : ℕ → (ι → ℕ) → ℚ) (Q : ℕ → (ι → ℕ) → (ι → Bool) → ℚ)
    (hbase : ∀ d : ι → ℕ, (∀ i, d i ≤ T) → (∃ i, d i = 0) → W 0 d ≤ 0)
    (hq0 : ∀ n, n < T → ∀ d : ι → ℕ, (∀ i, d i ≤ T - (n + 1)) → (∃ i, d i = 0) →
        ∀ g, 0 ≤ Q (n + 1) d g)
    (hq1 : ∀ n, n < T → ∀ d : ι → ℕ, (∀ i, d i ≤ T - (n + 1)) → (∃ i, d i = 0) →
        ∑ g, Q (n + 1) d g = 1)
    (hstep : ∀ n, n < T → ∀ d : ι → ℕ, (∀ i, d i ≤ T - (n + 1)) → (∃ i, d i = 0) →
        ∀ i : ι, W (n + 1) d
          ≤ ∑ g, Q (n + 1) d g * ((W n (nxt d g) + (dlt d g : ℚ)) - gq g i)) :
    ∀ n, n ≤ T → ∀ d : ι → ℕ, (∀ i, d i ≤ T - n) → (∃ i, d i = 0) →
      Force n (stG d) (W n d)

theorem antihold_of_table (T : ℕ) (V : ℕ → (ι → ℕ) → ℚ)
    (hbase : ∀ d : ι → ℕ, (∀ i, d i ≤ T) → (∃ i, d i = 0) → 0 ≤ V 0 d)
    (hstep : ∀ n, n < T → ∀ d : ι → ℕ, (∀ i, d i ≤ T - (n + 1)) → (∃ i, d i = 0) →
        ∀ g : ι → Bool,
          (V n (nxt d g) + (dlt d g : ℚ)) / 2
            + (V n (nxt d (cneg g)) + (dlt d (cneg g) : ℚ)) / 2 - 1 / 2 ≤ V (n + 1) d) :
    ∀ n, n ≤ T → ∀ d : ι → ℕ, (∀ i, d i ≤ T - n) → (∃ i, d i = 0) →
      AntiHold n (stG d) (V n d)

theorem play_of_table (T : ℕ) (V : ℕ → (ι → ℕ) → ℚ) (π : ℕ → (ι → ℕ) → ((ι → Bool) → ℚ))
    (h1 : ∀ n (d : ι → ℕ), ∑ g, π n d g = 1)
    (hbase : ∀ d : ι → ℕ, (∀ i, d i ≤ T) → (∃ i, d i = 0) → V 0 d = 0)
    (hstep : ∀ n, n < T → ∀ d : ι → ℕ, (∀ i, d i ≤ T - (n + 1)) → (∃ i, d i = 0) →
        ∀ i : ι, V (n + 1) d
          = ∑ g, π (n + 1) d g * ((V n (nxt d g) + (dlt d g : ℚ)) - gq g i)) :
    ∀ n, n ≤ T → ∀ d : ι → ℕ, (∀ i, d i ≤ T - n) → (∃ i, d i = 0) →
      Play π n (stG d) (V n d)
```

Here `dlt d g` is $\delta(d, g)$, `nxt d g` is $d'(d, g)$, and `stG d` is $-d$.

Theorems 4.1 and 4.3.

```
def cf3 (T : ℕ) : ℚ :=
  (((4 * ((T + 1) / 2) + (if T % 2 = 0 then 1 else 0)) * Nat.choose T (T / 2) : ℕ) : ℚ)
    / ((3 * 2 ^ T : ℕ) : ℚ)
```

```
def ew3 (T : ℕ) : ℚ := ((2 * ((T + 1) / 2) * Nat.choose T (T / 2) : ℕ) : ℚ) / ((2 ^ T : ℕ) : ℚ)
def p03 (T : ℕ) : ℚ :=
  if T % 2 = 0 then ((Nat.choose T (T / 2) : ℕ) : ℚ) / ((2 ^ T : ℕ) : ℚ) else 0
def abg3 (T : ℕ) : ℚ :=
  ((Nat.choose (T + 2) (T / 2 + 1) * (T / 2 + 1) : ℕ) : ℚ) / ((3 * 2 ^ T : ℕ) : ℚ)

theorem k3_exact (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 12) :
    IsValue T (fun _ : Fin 3 => (0 : ℤ)) (cf3 T)

theorem k3_random_walk : ∀ T, 1 ≤ T → T ≤ T3 → cf3 T = (2 * ew3 T + p03 T) / 3

theorem k3_abg_excess : ∀ T, 1 ≤ T → T ≤ T3 → T % 2 = 0 →
    abg3 T - cf3 T = ((Nat.choose T (T / 2) : ℕ) : ℚ) / ((3 * 2 ^ T : ℕ) : ℚ)
```

Theorem 4.5.

```
def c3list : List ℚ :=
  [1/2, 3/4, 7/8, 17/16, 37/32, 83/64, 177/128, 383/256, 807/512, 1713/1024, 3579/2048,
   7509/4096]
def gap3list : List ℚ :=
  [1/6, 1/12, 1/8, 1/16, 3/32, 11/192, 29/384, 41/768, 33/512, 51/1024, 117/2048, 191/4096]

theorem k3_comb_strict (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 12) :
    AntiHold T (fun _ : Fin 3 => (0 : ℤ)) (c3list.getD (T - 1) 0)
      ∧ Play pi3 T (fun _ : Fin 3 => (0 : ℤ)) (c3list.getD (T - 1) 0)
      ∧ c3list.getD (T - 1) 0 < cf3 T

theorem k3_comb_exact (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 12) :
    AntiHold T (fun _ : Fin 3 => (0 : ℤ)) (c3list.getD (T - 1) 0)
      ∧ Play pi3 T (fun _ : Fin 3 => (0 : ℤ)) (c3list.getD (T - 1) 0)
      ∧ (∀ v : ℚ, AntiHold T (fun _ : Fin 3 => (0 : ℤ)) v → c3list.getD (T - 1) 0 ≤ v)
      ∧ c3list.getD (T - 1) 0 < cf3 T

theorem k3_comb_gap : ∀ T, 1 ≤ T → T ≤ T3 →
    cf3 T - c3list.getD (T - 1) 0 = gap3list.getD (T - 1) 0
```

Theorems 5.1 and 5.3.

```
def v4list : List ℚ :=
  [3/4, 1, 23/20, 4/3, 29/20, 8/5, 1283/750, 329/180, 15433/8000, 21923/10800]
def c4list : List ℚ := [1/2, 1, 1, 5/4, 11/8, 3/2, 53/32, 111/64, 241/128, 499/256]
def gap4list : List ℚ :=
  [1/4, 0, 3/20, 1/12, 3/40, 1/10, 653/12000, 269/2880, 741/16000, 13943/172800]

theorem k4_exact (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 10) :
    IsValue T (fun _ : Fin 4 => (0 : ℤ)) (v4list.getD (T - 1) 0)

theorem k4_comb_strict (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 10) (h3 : T ≠ 2) :
    AntiHold T (fun _ : Fin 4 => (0 : ℤ)) (c4list.getD (T - 1) 0)
      ∧ Play pi4 T (fun _ : Fin 4 => (0 : ℤ)) (c4list.getD (T - 1) 0)
      ∧ c4list.getD (T - 1) 0 < v4list.getD (T - 1) 0

theorem k4_comb_exact (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 10) (h3 : T ≠ 2) :
    AntiHold T (fun _ : Fin 4 => (0 : ℤ)) (c4list.getD (T - 1) 0)
      ∧ Play pi4 T (fun _ : Fin 4 => (0 : ℤ)) (c4list.getD (T - 1) 0)
      ∧ (∀ v : ℚ, AntiHold T (fun _ : Fin 4 => (0 : ℤ)) v → c4list.getD (T - 1) 0 ≤ v)
      ∧ c4list.getD (T - 1) 0 < v4list.getD (T - 1) 0

theorem k4_comb_tight_at_two :
    c4list.getD 1 0 = v4list.getD 1 0 ∧ v4list.getD 1 0 = (1 : ℚ)

theorem k4_comb_gap : ∀ T, 1 ≤ T → T ≤ T4 →
    v4list.getD (T - 1) 0 - c4list.getD (T - 1) 0 = gap4list.getD (T - 1) 0
```

Proposition 6.1.

```
def cover2 (T : ℕ) : ℚ :=
  (((T + 1) / 2 * Nat.choose T (T / 2) : ℕ) : ℚ) / ((2 ^ T : ℕ) : ℚ)

theorem k2_cover (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 12) :
    IsValue T (fun _ : Fin 2 => (0 : ℤ)) (cover2 T)

theorem k2_comb_tight (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 12) :
    Play pi2 T (fun _ : Fin 2 => (0 : ℤ)) (cover2 T)
      ∧ AntiHold T (fun _ : Fin 2 => (0 : ℤ)) (cover2 T)
```

Proposition 6.2.

```
theorem k3_no_guar_below (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 12) (v : ℚ) (hv : v < cf3 T) :
    ¬ Guar T (fun _ : Fin 3 => (0 : ℤ)) v

theorem k3_no_force_above (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 12) (v : ℚ) (hv : cf3 T < v) :
    ¬ Force T (fun _ : Fin 3 => (0 : ℤ)) v

theorem k3_no_antihold_below (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 12) (v : ℚ)
    (hv : v < c3list.getD (T - 1) 0) : ¬ AntiHold T (fun _ : Fin 3 => (0 : ℤ)) v

theorem k3_comb_not_minimax (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 12) :
    ¬ Play pi3 T (fun _ : Fin 3 => (0 : ℤ)) (cf3 T)

theorem k4_no_guar_below (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 10) (v : ℚ)
    (hv : v < v4list.getD (T - 1) 0) : ¬ Guar T (fun _ : Fin 4 => (0 : ℤ)) v

theorem k4_no_force_above (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 10) (v : ℚ)
    (hv : v4list.getD (T - 1) 0 < v) : ¬ Force T (fun _ : Fin 4 => (0 : ℤ)) v

theorem k4_comb_not_minimax (T : ℕ) (h1 : 1 ≤ T) (h2 : T ≤ 10) (h3 : T ≠ 2) :
    ¬ Play pi4 T (fun _ : Fin 4 => (0 : ℤ)) (v4list.getD (T - 1) 0)

theorem hyp_not_vacuous :
    (∃ i : Fin 3, (fun _ : Fin 3 => (0 : ℕ)) i = 0)
      ∧ (∃ i : Fin 4, (fun _ : Fin 4 => (0 : ℕ)) i = 0)
      ∧ (∀ T, 1 ≤ T → T ≤ 12 → 0 < cf3 T)
      ∧ (∀ T, 1 ≤ T → T ≤ 10 → 0 < v4list.getD (T - 1) 0)
      ∧ (∀ T, 1 ≤ T → T ≤ 12 → 0 < c3list.getD (T - 1) 0)
```

References

[1] Y. Abbasi-Yadkori, P. L. Bartlett and V. Gabillon, *Near minimax optimal players for the finite-time 3-expert prediction problem*, in: Advances in Neural Information Processing Systems 30 (NIPS 2017), pp. 3033–3042. Not on arXiv; read from the proceedings PDF at `papers.nips.cc`, whose theorem numbering is used here.

[2] E. Bayraktar, I. Ekren and Y. Zhang, *On the asymptotic optimality of the comb strategy for prediction with expert advice*, Ann. Appl. Probab. 30 (2020), no. 6, 2517–2546; DOI `10.1214/20-AAP1565`; arXiv:1902.02368.

[3] E. Bayraktar, I. Ekren and X. Zhang, *Finite-time 4-expert prediction problem*, Comm. Partial Differential Equations 45 (2020), no. 7, 714–757; DOI `10.1080/03605302.2020.1712418`; arXiv:1911.10936.

[4] J. Calder, N. Drenska and D. Mosaphir, *Numerical solution of a PDE arising from prediction with expert advice*, European J. Appl. Math. 37 (2026), no. 1, 96–122 (published online 2025); DOI `10.1017/S0956792525000075`; arXiv:2406.05754.

[5] Z. Chase, *Experimental evidence for asymptotic non-optimality of comb adversary strategy*, arXiv:1912.01548.

[6] N. Cesa-Bianchi and G. Lugosi, *Prediction, Learning, and Games*, Cambridge University Press, 2006.

[7] T. M. Cover, *Behavior of sequential predictors of binary sequences*, in: Transactions of the Fourth Prague Conference on Information Theory, Statistical Decision Functions, Random Processes (Prague, 31 August – 11 September 1965), Publishing House of the Czechoslovak Academy of Sciences, Prague, 1967, pp. 263–272.

[8] N. Drenska and R. V. Kohn, *Prediction with expert advice: a PDE perspective*, J. Nonlinear Sci. 30 (2020), no. 1, 137–173; DOI `10.1007/s00332-019-09570-3`; arXiv:1904.11401.

[9] N. Gravin, Y. Peres and B. Sivan, *Towards optimal algorithms for prediction with expert advice*, in: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016), SIAM, 2016, pp. 528–547; DOI `10.1137/1.9781611974331.ch39`; arXiv:1409.3040v5. Theorem numbers and quotations are those of the arXiv version v5 (11 July 2016, "accepted to SODA 2016"), read from its PDF; the SODA version was not consulted and its numbering may differ.

[10] The OEIS Foundation, *The On-Line Encyclopedia of Integer Sequences*, sequence A152548, `https://oeis.org/A152548`.

[11] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction — CADE 28, Lecture Notes in Computer Science, Springer, 2021, 625–635; DOI `10.1007/978-3-030-79876-5_37`.

[12] The mathlib Community, *The Lean mathematical library*, in: Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020), ACM, 2020, pp. 367–381; DOI `10.1145/3372885.3373824`.