MINIMAX REGRET AGAINST AN ADVERSARIAL EXPERIMENT: CERTIFIED VALUES FOR THREE TO EIGHT BINARY SIGNALS AND THE TWO-POINT SUPPORT OF NATURE'S OPTIMAL MIXTURE

KOREA SUPERINTELLIGENCE LABS

ABSTRACT. A decision-maker who will observe n i.i.d. binary signals of unknown precision $\pi \in [\frac{1}{2}, 1]$ commits in advance to a belief a_k for every count k of high signals; an adversarial Nature mixes over π; the payoff is the mean-squared-error regret against an oracle who knows π, under a uniform prior on the binary state. Che, Li and Luo (arXiv:2602.15246) solve this game for $n \le 3$ — Nature mixes the two endpoint experiments $\{\frac{1}{2}, 1\}$ for $n \le 2$ and $\{\frac{1}{2}, \pi^*\}$ with an interior π^* at $n = 3$ — state that a closed-form characterisation for $n > 3$ is intractable by their approach, conjecture the two-point structure $\{\frac{1}{2}, \pi_n^*\}$ for every n, and plot the values that the conjecture predicts for $3 \le n \le 18$. We prove, without assuming anything about the structure of Nature's strategy, rational brackets $L_n \le V_n \le U_n$ on the minimax regret for $n = 3, \ldots, 8$, of width at most $2.1 \cdot 10^{-11}$: L_n is what an explicit rational two-point mixture guarantees against every belief rule, and U_n is what an explicit rational belief rule guarantees against every precision, proved by a Bernstein subdivision certificate. The numbers agree with the source's conjecture-based numerics; what is new is that they are bounds. A second certificate, with a margin vanishing exactly at the two conjectured atoms, shows that every finitely supported mixture guaranteeing L_n — in particular every finitely supported optimal one — satisfies $\sum_i w_i\,(2p_i - 1)^2(2p_i - 1 - \hat{d}_n)^2 \le 8.4 \cdot 10^{-11}$, which localises Nature's optimal support at the two conjectured atoms with an explicit error bar; at $n = 4$ at most $5.25 \cdot 10^{-4}$ of the mass lies further than $1/100$ from both. Alongside, the game in which Nature is confined to $\{\frac{1}{2}, 1\}$ has the closed-form value $1/(4(1 + 2^{(n-1)/2})^2)$ for every n, which reproduces the source's $n \le 2$ strategies and is strictly worse for Nature from $n = 3$ on, by a factor that grows from below 2 to above 21; and $V_3 > V_4 > \cdots > V_8$. Every statement is verified in Lean 4 against *Mathlib*.

1. INTRODUCTION

The objects. A binary state $\theta \in \{0, 1\}$ carries the uniform prior. A decision-maker (DM) will observe n conditionally i.i.d. binary signals from a symmetric experiment of *precision* $\pi = \Pr(s = h \mid \theta = 1) = \Pr(s = l \mid \theta = 0)$, and knows only that $\pi \in \Pi = [\frac{1}{2}, 1]$. The count K of high signals is sufficient, so a belief-formation rule is a vector $a = (a_0, \ldots, a_n)$, a_k being the belief the DM commits to, before the data arrive, for the event $K = k$. With quadratic loss the ex-ante regret of a against the oracle who knows π is

$$(1)\qquad R_n(a, \pi) = \sum_{k=0}^{n} P_k(\pi)\,\big(q_k(\pi) - a_k\big)^2, \qquad P_k(\pi) = \tfrac{1}{2}\tbinom{n}{k}\big[\pi^k(1-\pi)^{n-k} + (1-\pi)^k\pi^{n-k}\big],$$

where $q_k(\pi) = \Pr(\theta = 1 \mid \pi, K = k)$ is the oracle's posterior. Nature chooses π, and may randomise; the DM's minimax regret is

$$(2)\qquad V_n = \inf_a \sup_{\sigma \in \Delta(\Pi)} \mathbb{E}_\sigma R_n(a, \pi) = \inf_a \sup_{\pi \in [1/2, 1]} R_n(a, \pi),$$

the second equality because a supremum of a linear functional over $\Delta(\Pi)$ is a supremum over point masses (an elementary reduction, and the point at which the formal definition starts). This is the "special case: MSE with binary signals" of Che, Li and Luo [1], Section 3, whose general framework (an arbitrary strictly proper scoring rule, an arbitrary signal space) reduces to it under quadratic loss. The criterion is Savage's minimax regret [2], not Wald's minimax risk [3], and the difference matters here in a concrete way: the regret (1) is the DM's mean squared error in estimating θ *minus* the oracle's own Bayes risk $\mathbb{E}\operatorname{Var}(\theta \mid K; \pi)$ (Remark 2.2), so an uninformative experiment, which maximises the

DM's error, gives Nature nothing — under minimax *risk* Nature would simply play $\pi = \frac{1}{2}$ and the DM would ignore the data. It is the subtraction of the oracle's risk that makes Nature's optimal strategy a nontrivial mixture of experiments and makes the value depend on n.

The source, and where it stops. The source solves the game for $n \leq 3$ and stops. Its Proposition 1 (p. 9): for $n = 1$ Nature randomises between $\frac{1}{2}$ and 1 with equal probability and the DM plays $(a_0, a_1) = (\frac{1}{4}, \frac{3}{4})$. Its Theorem 1 (p. 10): for $n = 2$ Nature randomises between $\frac{1}{2}$ and 1 with probabilities $2 - \sqrt{2}$ and $\sqrt{2} - 1$; for $n = 3$ Nature randomises between $\frac{1}{2}$ and some $\pi^* \in (\frac{1}{2}, 1)$. In both statements only the *strategies* are given; no value is printed for any n, and no value or bound is given for π^*. Then, verbatim (p. 11):

> "A closed-form characterization for $n > 3$ is computationally intractable using our approach. Instead, we provide numerical evidence for the conjectured equilibrium by imposing the necessary equilibrium conditions implied by its structure. [...] Specifically, we conjecture that Nature randomizes between the uninformative experiment $\pi = \frac{1}{2}$ and an informative experiment with precision π_n^*, assigning probability $w \in [0, 1]$ to the latter."

Under this conjecture the equilibrium is pinned down by three conditions the source prints in full — the DM's first-order condition, Nature's indifference between $\frac{1}{2}$ and π_n^*, and local optimality of π_n^* — and the source solves them numerically. Its Figure 2 (p. 12) plots the regret curve $\pi \mapsto R_n(a_n^*, \pi)$ for $3 \leq n \leq 5$, and its Figure 3 (p. 13) plots π_n^*, the equilibrium regret $R_n(a_n^*, \pi_n^*)$, the weight w and the rule a_n^* for $3 \leq n \leq 18$; the text (p. 12) reads the trends off the plots: π_n^* decreases with n, and "the equilibrium regret $R(a_n^*, \pi_n^*)$ also decreases with n". The source then moves to the large-n limit (its Theorem 2, p. 15, and Theorem 3, p. 17), where the two-point structure is proved in a Gaussian limit game. Nothing in the source bounds V_n for any $n \geq 3$ in either direction, and for $n \geq 4$ the plotted numbers are conditional on the conjectured structure. (Statement, page and figure numbers throughout are those of the e-print compiled from the arXiv source, version 1, the only version as of 2026-09-07.)

What is proved here. Everything below holds for the game exactly as the source defines it, with the DM restricted to pure rules (the source's own reduction, Section 2: the regret is strictly convex in a).

(i) *Certified brackets* (Theorem 4.1). For $n = 3, \ldots, 8$,

$$L_n \leq V_n \leq U_n$$

with explicit rationals L_n, U_n of width $8 \cdot 10^{-12}$, $2.1 \cdot 10^{-11}$, $7 \cdot 10^{-12}$, $5 \cdot 10^{-12}$, $2 \cdot 10^{-12}$, $8 \cdot 10^{-13}$ (Table 2). The lower bound is what one explicit rational two-point mixture guarantees against every rule; the upper bound is what one explicit rational rule guarantees against every precision. Neither half assumes anything about the structure of Nature's optimal strategy, so the bracket is an unconditional enclosure of the true minimax regret.

(ii) *Localisation of Nature's optimal support* (Theorem 5.2). For $n = 3, \ldots, 8$, every finitely supported mixture $\sum_i w_i \delta_{p_i}$ that guarantees Nature at least L_n — every optimal mixture does — satisfies

$$\sum_i w_i \,(2p_i - 1)^2\,(2p_i - 1 - \hat{d}_n)^2 \;\leq\; \varepsilon_n, \qquad \varepsilon_n \leq 8.4 \cdot 10^{-11},$$

where $\hat{d}_n = 2\hat{\pi}_n - 1$ and $\hat{\pi}_n$ is the rational interior atom of the mixture in (i). The weight function vanishes exactly at the two conjectured atoms $\frac{1}{2}$ and $\hat{\pi}_n$, so this is the source's conjectured two-point structure *derived* rather than assumed, with an explicit error bar: at $n = 4$ at most $5.25 \cdot 10^{-4}$ of Nature's mass can lie further than $1/100$ from both atoms.

(iii) *The restricted game in closed form* (Proposition 3.1). If Nature is confined to mixtures of the two endpoint experiments $\{\frac{1}{2}, 1\}$ — the support the source proves optimal for $n \leq 2$ — the value is exactly $V_n^{\mathrm{r}} = 1/\big(4(1 + 2^{(n-1)/2})^2\big)$ for every $n \geq 1$, with an explicit saddle point. At $n = 1, 2$ this reproduces the source's Proposition 1 and Theorem 1(i). This is a routine equaliser computation and is not claimed as new; its role is to give (iv) an exact benchmark.

(iv) *The endpoint support is strictly suboptimal from $n = 3$ on* (Corollary 4.6): $V_n^{\mathrm{r}} < V_n$ for $n = 3, \dots, 8$, with $V_3 < 2V_3^{\mathrm{r}}$ and $21V_8^{\mathrm{r}} < V_8$; and against the certified rule at $n = 4$ the fully revealing experiment is not even a best reply.

(v) *Monotonicity* (Corollary 4.5): $V_3 > V_4 > \cdots > V_8$, because the six brackets are pairwise disjoint and ordered.

What is, and is not, new. The *numbers* in (i) are the source's own. Its Figure 3 plots the equilibrium regret predicted by the two-point conjecture for $3 \le n \le 18$, and our brackets agree with those plotted values to the precision of a plot — indeed to eleven digits, as the equilibrium computed from the source's three conditions lies inside each bracket (Remark 4.3). We do not claim the values. What is new is that they are now *proved* to be the value, unconditionally, in the range where the source states it cannot characterise the equilibrium. Statement (ii) is new in kind: the source proves the two-point structure only at $n = 3$ (Theorem 1(ii)) and nowhere quantifies it; here the structure is a consequence of a certificate, for $n = 3, \dots, 8$, and comes with a bound on how much mass can sit away from the two atoms. Statement (v) is the source's own observation (its Figure 3, p. 13), proved here for $3 \le n \le 8$ without assuming the structure. Statements (iii) and (iv) are routine once (i) is available. We do *not* obtain a closed form for π_n^* or V_n, do not prove that the optimal support is *exactly* two points, and prove nothing for $n \ge 9$ or for general n; Section 6 says what stands in the way.

Method. Once the DM's rule a is fixed, $R_n(a, \cdot)$ is a rational function of π whose denominators are positive on $[\frac{1}{2}, 1]$ (Lemma 2.3); in the coordinate $d = 2\pi - 1 \in [0, 1]$ and after clearing denominators, "$R_n(a, \pi) \le U$ for every π" is the nonnegativity of an integer-coefficient polynomial of degree at most 28 on $[0, 1]$, which a Bernstein subdivision certificate proves (Section 2.3). The rules $a^{(n)}$ are the 10-digit roundings of the equilibrium rules obtained from the source's three conditions, and U_n is a decimal threshold just above the candidate value at which the subdivision search succeeded; the controls of Proposition 4.4 show that U_n is within 10^{-11} of the least upper bound available for that rule. Lower bounds need no certificate: against a two-point mixture the DM's best reply decouples over k, and each coordinate is a harmonic-mean inequality (Lemma 2.4), a rational number. The localisation (ii) is a second certificate on the same polynomial with a quartic margin subtracted, plus three lines of averaging. The shape of (iii) coincides with the classical binomial minimax *risk* of Hodges and Lehmann [4], $1/(4(1 + \sqrt{n})^2)$, under $\sqrt{n} \mapsto 2^{(n-1)/2}$; Remark 3.3 explains why this is a shared one-dimensional equaliser identity and not a correspondence of problems.

Verification. Every theorem, proposition and corollary in Sections 2–5 is a declaration formally verified in Lean 4 against *Mathlib*, with the statement reproduced verbatim in Appendix A; Section 7 describes the development, says exactly where compiled evaluation (`native_decide`) enters — only in the twelve Boolean acceptance checks of the Bernstein trees — and reports the axiom audit. Every number in the paper was recomputed independently in exact rational arithmetic for this paper.

2. Preliminaries

2.1. **The game.** Fix $n \ge 1$. For $\pi \in [0, 1]$ and $0 \le k \le n$ put $A_k(\pi) = \pi^k (1 - \pi)^{n-k}$ and $B_k(\pi) = (1 - \pi)^k \pi^{n-k}$, so that

$$P_k(\pi) = \tfrac{1}{2}\tbinom{n}{k}\big(A_k + B_k\big), \qquad q_k(\pi) = \frac{A_k}{A_k + B_k},$$

with the convention $q_k = 0$ when $A_k + B_k = 0$ (this happens only at $\pi \in \{0, 1\}$ and $0 < k < n$, where $P_k = 0$ and the summand of (1) vanishes whatever the convention). A rule is any $a \in \mathbb{R}^{n+1}$; we write $R_n(a, \pi)$ for (1), and

$$\Gamma_n(a) = \sup_{\pi \in [1/2, 1]} R_n(a, \pi), \qquad V_n = \inf_{a \in \mathbb{R}^{n+1}} \Gamma_n(a).$$

Since $0 \le q_k \le 1$ and $0 \le P_k \le \binom{n}{k}$ on $[0, 1]$, $R_n(a, \cdot)$ is bounded on $[\frac{1}{2}, 1]$ and Γ_n is finite; $\Gamma_n \ge 0$, so $V_n \ge 0$. Two elementary facts are used throughout: any rule gives an upper bound, $V_n \le \Gamma_n(a)$; and any L with $L \le \Gamma_n(a)$ for every a is a lower bound, $L \le V_n$.

The source lets the DM choose $a \in [0, 1]^{n+1}$; allowing all of $\mathbb{R}^{n+1}$ changes nothing:

Proposition 2.1 (the action space)**.** *For every $a \in \mathbb{R}^{n+1}$ there is $b \in [0,1]^{n+1}$ with $R_n(b,\pi) \le R_n(a,\pi)$ for every $\pi \in [0,1]$, and hence $\Gamma_n(b) \le \Gamma_n(a)$. Consequently V_n is unchanged when the DM is confined to the source's action set,*

$$V_n = \inf_{a\in[0,1]^{n+1}} \sup_{\sigma\in\Delta(\Pi)} \mathbb{E}_\sigma R_n(a,\pi)$$

(attainment is not part of the statement). Moreover the regret is not identically zero: $R_1\big((\frac{1}{2},\frac{1}{2}),1\big) > 0$.

Proof. Take $b_k = \max(0, \min(1, a_k))$, the metric projection of a_k onto $[0,1]$. Every oracle posterior $q_k(\pi)$ lies in $[0,1]$, so $(q_k - b_k)^2 \le (q_k - a_k)^2$ termwise, and $P_k \ge 0$. The last claim is $R_1((\frac{1}{2},\frac{1}{2}),1) = \frac{1}{8} + \frac{1}{8}$. □

Remark 2.2 (regret versus risk). Under the uniform prior, $q_K(\pi) = \mathbb{E}[\theta \mid K]$ when the precision is π, so by the orthogonality of conditional expectation

$$\mathbb{E}\big[(a_K-\theta)^2\big] = \mathbb{E}\big[(a_K-q_K)^2\big] + \mathbb{E}\big[(q_K-\theta)^2\big] = R_n(a,\pi) + \mathbb{E}\operatorname{Var}(\theta \mid K;\pi).$$

Thus $R_n(a,\pi)$ is the DM's mean squared error in estimating θ minus the oracle's Bayes risk, the price of not knowing π. Under minimax *risk* (no subtraction) the game is trivial: at $\pi = \frac{1}{2}$, K is independent of θ and $\mathbb{E}[(a_K-\theta)^2] = \sum_k \binom{n}{k} 2^{-n}\big[\frac{1}{2}a_k^2 + \frac{1}{2}(1-a_k)^2\big] \ge \frac{1}{4}$ with equality only for $a \equiv \frac{1}{2}$, which has risk exactly $\frac{1}{4}$ at every π; so the minimax-risk value is $\frac{1}{4}$, $\pi = \frac{1}{2}$ is least favourable and the DM ignores the data. This is the "triviality" the source avoids by choosing regret — its footnote 1 (p. 2) puts it as a *maxmin utility* criterion, under which "an adversarial Nature would select a completely uninformative DGP, leading the DM to simply ignore the data", which for this zero-sum problem is the minimax-risk computation above — and the classical distinction between Wald [3] and Savage [2]. (This remark is elementary and outside the formal development.)

2.2. **Two reductions.** The relabelling $\theta \mapsto 1-\theta$, $k \mapsto n-k$ fixes $R_n(\cdot,\pi)$ for every π, so the unique minimiser of the strictly convex Γ_n (uniqueness is the source's observation) satisfies $a_{n-k} = 1 - a_k$; every rule certified below has this symmetry, and for it the regret collapses to the counts $2k < n$.

Lemma 2.3 (symmetric reduction)**.** *Let $a \in \mathbb{R}^{n+1}$ satisfy $a_{n-k} = 1 - a_k$ for all k (so $a_{n/2} = \frac{1}{2}$ when n is even), and let $\pi \in (0,1)$. With $S_m(\pi) = \pi^m + (1-\pi)^m$,*

$$R_n(a,\pi) = \sum_{2k<n} \binom{n}{k} \big(\pi(1-\pi)\big)^k \frac{\big((1-\pi)^{n-2k} - a_k\, S_{n-2k}(\pi)\big)^2}{S_{n-2k}(\pi)}.$$

Moreover $R_n(a,1) = \frac{1}{2}a_0^2 + \frac{1}{2}(1-a_n)^2$ for every a, and $R_n(a,\frac{1}{2}) = \sum_k \binom{n}{k} 2^{-n}(\frac{1}{2} - a_k)^2$.

Proof. For $2k < n$, $A_k + B_k = (\pi(1-\pi))^k S_{n-2k}$ and $A_k - a_k(A_k+B_k) = (\pi(1-\pi))^k\big((1-\pi)^{n-2k} - a_k S_{n-2k}\big)$, so the k-th summand of (1) is $\frac{1}{2}\binom{n}{k}(\pi(1-\pi))^k((1-\pi)^{n-2k} - a_k S_{n-2k})^2/S_{n-2k}$. The $(n-k)$-th summand has $A_{n-k} = B_k$, $B_{n-k} = A_k$ and $A_{n-k} - a_{n-k}(A_{n-k}+B_{n-k}) = -(A_k - a_k(A_k+B_k))$, so it equals the k-th; the pair contributes twice, cancelling the $\frac{1}{2}$. For even n and $k = n/2$, $q_k = \frac{1}{2} = a_k$ and the summand vanishes. At $\pi = 1$ only $k \in \{0,n\}$ have $P_k \ne 0$, with $P_k = \frac{1}{2}$ and $q_0 = 0$, $q_n = 1$; at $\pi = \frac{1}{2}$ every $q_k = \frac{1}{2}$ and $P_k = \binom{n}{k}2^{-n}$. □

In the development the identity of Lemma 2.3 is established separately for each of the six certified rules (by `ring` after clearing denominators), and the two endpoint formulas for general n and a.

Lemma 2.4 (harmonic mean)**.** *For $c_1, c_2 > 0$ and all real x, y, $c_1x^2 + c_2y^2 \;\ge\; \dfrac{c_1c_2}{c_1+c_2}(x-y)^2$.*

Proof. $(c_1+c_2)(c_1x^2+c_2y^2) - c_1c_2(x-y)^2 = (c_1x+c_2y)^2 \ge 0$. □

Lemma 2.4 is the whole content of every lower bound in this paper: against a mixture $(1-w)\delta_{1/2} + w\delta_p$ the expected regret of a rule a is $\sum_k \big[wP_k(p)(q_k(p)-a_k)^2 + (1-w)P_k(\frac{1}{2})(\frac{1}{2}-a_k)^2\big]$, each summand is bounded below by the lemma with $x = q_k(p) - a_k$, $y = \frac{1}{2} - a_k$, and the bound $\sum_k \frac{c_{1,k}c_{2,k}}{c_{1,k}+c_{2,k}}(q_k(p) - \frac{1}{2})^2$ no longer depends on a. When p and w are rational it is a rational number.

2.3. **Bernstein certificates.** A polynomial F of degree at most m has a unique expansion $F(t) = \sum_{i=0}^{m} b_i \binom{m}{i} t^i (1-t)^{m-i}$ on $[0,1]$; if all Bernstein coefficients b_i are nonnegative then $F \geq 0$ on $[0,1]$. If $F > 0$ on $[0,1]$, then after finitely many midpoint subdivisions every piece, re-expanded on its own interval, has nonnegative Bernstein coefficients (this is the effective form of Bernstein's theorem on polynomials positive on an interval; see Powers and Reznick [7] and, for certificates in this basis, Boudaoud, Caruso and Roy [8]). A *certificate* for $F \geq 0$ on $[0,1]$ is therefore a finite binary tree of dyadic subintervals at each of whose leaves the Bernstein coefficients are nonnegative; checking a certificate is exact rational arithmetic, and its soundness is a theorem. The certificates here were found by a search that bisects at the midpoint whenever the root coefficients are not all nonnegative, and were then checked in the formal development by the shared Bernstein layer of the repository, whose soundness theorem is proved in Lean and whose arithmetic is run by compiled evaluation (Section 7).

3. The restricted game: Nature confined to the endpoint experiments

For $w \in \mathbb{R}$ write

$$R_n^{\mathrm{r}}(a, w) = w\, R_n(a, 1) + (1-w)\, R_n(a, \tfrac{1}{2}),$$

Nature's payoff when she puts weight w on the fully revealing experiment and $1-w$ on the uninformative one; for $w \in [0,1]$ this is $\mathbb{E}_\sigma R_n(a, \cdot)$ with $\sigma = w\,\delta_1 + (1-w)\,\delta_{1/2}$. Put

$$c_n = 2^{(n-1)/2}, \qquad w_n = \frac{1}{1+c_n}, \qquad V_n^{\mathrm{r}} = \frac{1}{4(1+c_n)^2},$$

and let $a^\circ \in \mathbb{R}^{n+1}$ be the rule $a_0^\circ = \dfrac{1}{2(1+c_n)}$, $a_n^\circ = 1 - a_0^\circ$, $a_k^\circ = \frac{1}{2}$ for $0 < k < n$.

Proposition 3.1 (the restricted game, solved for every n)**.** *Let $n \geq 1$.*

(a) *For every rule $a \in \mathbb{R}^{n+1}$, $R_n^{\mathrm{r}}(a, w_n) \geq V_n^{\mathrm{r}}$.*

(b) *For every $w \in \mathbb{R}$, $R_n^{\mathrm{r}}(a^\circ, w) = V_n^{\mathrm{r}}$; that is, a° is an equaliser, $R_n(a^\circ, 1) = R_n(a^\circ, \frac{1}{2}) = V_n^{\mathrm{r}}$.*

Consequently the game in which Nature may only mix over $\{\frac{1}{2}, 1\}$ has value $V_n^{\mathrm{r}} = 1/\big(4(1+2^{(n-1)/2})^2\big)$, with saddle point (a°, w_n).

Proof. By Lemma 2.3, $R_n(a,1) = \frac{1}{2}a_0^2 + \frac{1}{2}(1-a_n)^2$ involves only the two extreme coordinates, and $R_n(a, \frac{1}{2}) \geq 2^{-n}(\frac{1}{2} - a_0)^2 + 2^{-n}(\frac{1}{2} - a_n)^2$ by dropping the interior terms. With $c_1 = w/2$ and $c_2 = (1-w)/2^n$, Lemma 2.4 applied to $(x,y) = (a_0, a_0 - \frac{1}{2})$ and to $(1-a_n, \frac{1}{2} - a_n)$ gives $R_n^{\mathrm{r}}(a,w) \geq 2 \cdot \frac{c_1 c_2}{c_1 + c_2} \cdot \frac{1}{4}$. At $w = w_n$ one computes $c_1 + c_2 = c_n/2^n$ and $c_1 c_2 = c_n/(2^{n+1}(1+c_n)^2)$ using $c_n^2 = 2^{n-1}$, so the harmonic mean is $1/(2(1+c_n)^2)$ and the bound is V_n^{r}: this is (a). For (b), both $R_n(a^\circ, 1) = a_0^{\circ 2}$ and $R_n(a^\circ, \frac{1}{2}) = 2 \cdot 2^{-n}(\frac{1}{2} - a_0^\circ)^2 = 2^{1-n} c_n^2/(4(1+c_n)^2)$ equal V_n^{r}, and a convex combination of two equal numbers is that number. The two halves give $\max_w \min_a \geq V_n^{\mathrm{r}} \geq \min_a \max_w$, and $\max_w \min_a \leq \min_a \max_w$ always. $\square$

Corollary 3.2 (the source's two solved cases)**.** *At $n = 1$: $V_1^{\mathrm{r}} = \frac{1}{16}$, $w_1 = \frac{1}{2}$ and $a^\circ = (\frac{1}{4}, \frac{3}{4})$. At $n = 2$: $V_2^{\mathrm{r}} = (3 - 2\sqrt{2})/4$ and $w_2 = \sqrt{2} - 1$.*

These are exactly the strategies of the source's Proposition 1 and Theorem 1(i) — Nature's weight w_n on $\pi = 1$ and the DM's rule — so for $n = 1, 2$ the restriction to $\{\frac{1}{2}, 1\}$ does not bind and $V_n^{\mathrm{r}} = V_n$. The source prints neither value; $\frac{1}{16}$ and $(3 - 2\sqrt{2})/4$ follow from its strategies in one line. We record the corollary as a control on the model rather than as a result.

Remark 3.3 (the Hodges–Lehmann shape). The classical minimax estimator of a binomial proportion, $X \sim \mathrm{Bin}(n,p)$ with squared error, is $(X + \sqrt{n}/2)/(n + \sqrt{n})$, an equaliser with constant risk $1/(4(1 + \sqrt{n})^2)$, Bayes against the $\mathrm{Beta}(\sqrt{n}/2, \sqrt{n}/2)$ prior (Hodges and Lehmann [4], Section 5, Problem 1, who credit the result to H. Rubin). Proposition 3.1 has the same shape under $\sqrt{n} \mapsto 2^{(n-1)/2}$, and so does the extreme-cell action ($1/(2(1+\sqrt{n}))$ there, $1/(2(1+c_n))$ here). We checked this rather than waving at it: both are the one-dimensional identity $\min_\lambda [A/(1-\lambda) + B/\lambda] = (\sqrt{A} + \sqrt{B})^2$ applied to a ratio, and the ratios differ. Classically the ratio is n, a Fisher-information scale; here it is $2^{n-1} = P_0(1)/P_0(\frac{1}{2})$, the likelihood ratio at the extreme count between the two experiments Nature is allowed to mix — an informativeness gap in Blackwell's sense [5], exponential where the classical one is linear. This is exactly why $V_n^{\mathrm{r}} \to 0$ geometrically while V_n stays bounded away from zero (the source's Theorem 2 and Corollary 1): the exponential ratio is an artefact of forbidding Nature her interior atom. The problems also differ in what Nature indexes (the estimand and the law, versus the experiment only), in the loss (risk versus regret against a π-oracle) and in what the DM estimates (the parameter versus the functional $q_K(\pi)$ of a nuisance parameter); in their problem the oracle who knows p makes no error, so risk and regret coincide, whereas here the oracle's own Bayes risk is positive and its subtraction is what makes the game nondegenerate (Remark 2.2). The two-point least-favourable device on a constrained binomial parameter has its own literature, e.g. Marchand and MacGibbon [6]; none of it is this problem, and none of it writes down the formula of Proposition 3.1. It is, however, the reason the proposition is presented as routine.

4. Certified brackets on the minimax regret, $n = 3, \ldots, 8$

4.1. **The data.** For each $n = 3, \ldots, 8$ the certified objects are: a rational rule $a^{(n)} \in \mathbb{Q}^{n+1}$ with $a_k^{(n)} = A_{n,k}/10^{10}$ for $2k < n$, $a_{n/2}^{(n)} = \frac{1}{2}$ for even n, and $a_{n-k}^{(n)} = 1 - a_k^{(n)}$; a rational two-point mixture $\hat\sigma_n = (1-\hat w_n)\delta_{1/2} + \hat w_n \delta_{\hat\pi_n}$; and the rationals $L_n < U_n$. Table 1 lists them. The integers $A_{n,k}$ are the 10-digit roundings, and $(\hat\pi_n, \hat w_n)$ the 6-digit roundings, of the equilibrium candidate computed from the source's three conditions (Remark 4.3); L_n is the guarantee of $\hat\sigma_n$ rounded down to twelve decimals. U_n was chosen by scanning decimal resolutions $10^{-14}, 10^{-13}, \ldots$ and, at each resolution, the ten grid points just above the candidate value, and taking the first threshold at which the subdivision search (depth at most 40, at most $6 \cdot 10^5$ nodes explored) found a tree; Proposition 4.4(a) shows that the result is within 10^{-11} of the least upper bound available for the rule $a^{(n)}$.

n	$A_{n,k}$ for $2k<n$ (rule $a_k^{(n)} = A_{n,k}/10^{10}$)	$\hat\pi_n$	$\hat w_n$
3	1411751726, 4036110563	$899641/10^6$	$468117/10^6$
4	1054984660, 3149265253	17151/20000	$465223/10^6$
5	756398073, 2358019394, 4269897244	$830217/10^6$	235857/500000
6	581133721, 1762845674, 3552583731	$803367/10^6$	94011/200000
7	433555020, 1296091334, 2858415834, 4384816653	$785269/10^6$	$472581/10^6$
8	337583999, 970263753, 2244296003, 3774777348	383947/500000	94419/200000

Table 1. The certified rules and Nature's certified mixtures.

n	L_n	U_n	$U_n - L_n$	deg	tree (nodes/depth)	margin tree
3	0.039156935533	0.039156935541	$8 \cdot 10^{-12}$	7	31/15	31/15
4	0.036580028069	0.036580028090	$2.1 \cdot 10^{-11}$	10	31/15	31/15
5	0.036399353288	0.036399353295	$7 \cdot 10^{-12}$	14	31/15	31/15
6	0.035570811186	0.035570811191	$5 \cdot 10^{-12}$	18	31/15	31/15
7	0.035382000376	0.035382000378	$2 \cdot 10^{-12}$	23	35/17	33/16
8	0.035026812023	0.0350268120238	$8 \cdot 10^{-13}$	28	37/18	39/19

Table 2. The brackets $L_n \le V_n \le U_n$ (exact decimals), the degree in $d = 2\pi - 1$ in which the certificate polynomial is expressed, and the sizes of the two subdivision trees walked by the formal checker (Sections 4 and 5).

4.2. **The theorem.**

Theorem 4.1 (certified brackets). *For $n = 3,4,5,6,7,8$ the minimax regret satisfies $L_n \le V_n \le U_n$ with the exact rationals of Table 2. More precisely:*

(a) *(*lower*) for every rule $a \in \mathbb{R}^{n+1}$, $\hat w_n\, R_n(a, \hat\pi_n) + (1-\hat w_n)\, R_n(a, \tfrac12) \;\ge\; L_n$;*
(b) *(*upper*) for every $\pi \in [\tfrac12, 1]$, $R_n(a^{(n)}, \pi) \le U_n$.*

The formally verified statements are the six brackets; (a) and (b) are the lemmas from which each bracket is assembled in the development (`lower`n, `upper`n), and the Boolean acceptance of each subdivision tree is recorded as its own declaration (`bOk`n`_true`).

Proof. (a) Apply Lemma 2.4 termwise as explained after it, with $c_{1,k} = \hat{w}_n P_k(\hat{\pi}_n)$ and $c_{2,k} = (1-\hat{w}_n)\binom{n}{k}2^{-n}$; all quantities are rational, the sum is evaluated exactly, and it exceeds L_n (in the development each of the $n+1$ termwise bounds and the final comparison are discharged by `norm_num` on $\mathbb{Q}$). Then $R_n(a, \hat{\pi}_n) \le \Gamma_n(a)$ and $R_n(a, \frac{1}{2}) \le \Gamma_n(a)$ give $L_n \le \Gamma_n(a)$ for every a, hence $L_n \le V_n$.

(b) At $\pi = 1$, Lemma 2.3 gives $R_n(a^{(n)}, 1) = (a_0^{(n)})^2 \le U_n$ directly. For $\pi \in [\frac{1}{2}, 1)$ put $d = 2\pi - 1 \in [0, 1)$ and $T_m(d) = (1+d)^m + (1-d)^m = 2^m S_m(\pi)$. Multiplying the identity of Lemma 2.3 by $Q^2\, 2^n \prod_{2k<n} T_{n-2k}(d)$, $Q = 10^{10}$, turns it into an integer-coefficient polynomial identity

$$\begin{aligned} &R_n(a^{(n)}, \pi) \cdot Q^2\, 2^n \prod_{2k<n} T_{n-2k}(d) \\ &\qquad = \sum_{2k<n} \binom{n}{k} (1-d^2)^k \big(Q(1-d)^{n-2k} - A_{n,k} T_{n-2k}(d)\big)^2 \prod_{\substack{2j<n \\ j \ne k}} T_{n-2j}(d) =: N_n(d), \end{aligned}$$

of degree at most $n + \sum_{2k<n}(n-2k)$, the degree in which the certificate is expressed (column "deg" of Table 2); for odd n the odd factors T_m have degree $m-1$ and the true degree drops to 6, 12, 20 at $n = 3, 5, 7$. Since every $T_m > 0$ on $[0, 1]$, the claim is equivalent to $C_n(d) = U_n^{\mathrm{num}}\, 2^n Q^2 \prod_{2k<n} T_{n-2k}(d) - U_n^{\mathrm{den}} N_n(d) \ge 0$ on $[0, 1]$, where $U_n = U_n^{\mathrm{num}}/U_n^{\mathrm{den}}$. This is what the Bernstein subdivision certificate of Table 2 proves: a binary tree of 31 to 37 nodes — 16 to 19 dyadic subintervals — of depth at most 18, at each leaf of which all Bernstein coefficients of C_n are nonnegative. The formal development evaluates the certificate by compiled exact arithmetic (this is the only use of `native_decide`), passes from the accepted Boolean to $C_n \ge 0$ on $[0, 1]$ by the kernel-checked soundness theorem of the shared Bernstein layer, and from there to (b) by the identity above (`field_simp`, `ring`) and one `nlinarith` step. Then $V_n \le \Gamma_n(a^{(n)}) \le U_n$. □

Remark 4.2 (both orderings of the quantifiers). Part (a) says more than $L_n \le V_n$: it says that Nature's max-min value $\sup_\sigma \inf_a \mathbb{E}_\sigma R_n(a, \cdot)$ is at least L_n, since $\hat{\sigma}_n$ alone guarantees it. With $\sup_\sigma \inf_a \le \inf_a \sup_\sigma$ always, the brackets enclose both $\sup_\sigma \inf_a$ and $\inf_a \sup_\sigma$ without appeal to a minimax theorem; the source invokes one (its Section 2) to speak of an equilibrium, and Theorem 5.2 is phrased so as not to need it.

Remark 4.3 (the equilibrium candidate; outside the formal development). Solving the source's three conditions — a_k the Bayes posterior against $(1-w)\delta_{1/2} + w\delta_\pi$ (its (FOC)), indifference $R_n(a, \pi) = R_n(a, \frac{1}{2})$ (its (Indif)) and $\partial_\pi R_n(a, \pi) = 0$ (its (Local Opt)) — by Newton's method in (π, w) at 60 digits gives, for $n = 3, \ldots, 8$,

n	π_n^*	w_n	$R_n(a_n^*, \pi_n^*)$
3	0.899641450569	0.468117397717	0.039156935533286
4	0.857549873223	0.465222632432	0.036580028069353
5	0.830216735043	0.471713851439	0.036399353289027
6	0.803366728676	0.470054978985	0.035570811186940
7	0.785269174944	0.472581348862	0.035382000376493
8	0.767894457742	0.472095147168	0.035026812023353

and a 4001-point grid finds no π at which $R_n(a_n^*, \pi)$ exceeds the value at π_n^*. These are the numbers the source plots in its Figure 3 (it prints none of them; the agreement with the plot is to the precision of a plot). Each value lies inside the bracket of Theorem 4.1, and the 10-digit roundings of a_n^* are exactly the rules $a^{(n)}$ of Table 1. Nothing in this remark is used in any proof: the candidate only tells the certificate where to look.

4.3. **Controls.** Three kinds of control were built in, as the development's discipline requires: a too-small claim is refuted, a too-large claim is refuted, and no hypothesis is vacuous.

Proposition 4.4 (controls)**.** (a) *(*sharpness of each U_n*) For $n = 3, \dots, 8$ there is a precision $\pi_n^\bullet \in [\frac{1}{2}, 1]$ at which $R_n(a^{(n)}, \pi_n^\bullet) > T_n$, where $T_n < U_n$ and $U_n - T_n \le 10^{-11}$; explicitly*

n	$\pi_n^\bullet$	T_n	$U_n - T_n$
3	$\frac{1}{2}$	0.039156935540	10^{-12}
4	$\hat{\pi}_4$	0.03658002808	10^{-11}
5	$\frac{1}{2}$	0.036399353294	10^{-12}
6	$\frac{1}{2}$	0.03557081119	10^{-12}
7	$\hat{\pi}_7$	0.035382000377	10^{-12}
8	$\hat{\pi}_8$	0.035026812023	$8 \cdot 10^{-13}$

so no bound below T_n holds for the rule $a^{(n)}$ on all of $[\frac{1}{2}, 1]$.

(b) *(*the value is pinned on both sides*) $V_4 > 0.036580028068$ and $V_4 < 0.036580028091$.*

(c) *(*non-vacuity*) $\frac{1}{2} < \hat{\pi}_4 < 1$, $0 < \hat{w}_4 < 1$, and $R_4(a^{(4)}, \frac{1}{2}) > 0$.*

Proof. (a) is a rational evaluation of Lemma 2.3 at the named point; (b) is Theorem 4.1 at $n = 4$ and arithmetic; (c) is arithmetic. Part (a) is the reason the brackets cannot be narrowed by more than 10^{-11} without changing the rules: the certified rules are 10-digit roundings, and Nature's best replies to them at $\frac{1}{2}$ and near $\hat{\pi}_n$ are already within 10^{-11} of U_n. $\square$

4.4. **Two consequences.**

Corollary 4.5 (the minimax regret falls with the sample size)**.** *$V_3 > V_4 > V_5 > V_6 > V_7 > V_8$.*

Proof. $U_{n+1} < L_n$ for each $n = 3, \dots, 7$ (Table 2), so the six brackets are pairwise disjoint and ordered. $\square$

The source observes this in its Figure 3 (p. 13); its p. 12 reads off the upper-right panel that "the equilibrium regret $R(a_n^*, \pi_n^*)$ also decreases with n", over the wider range $3 \le n \le 18$ and from numerics that impose the two-point structure. The observation is the source's; the corollary proves it for $3 \le n \le 8$ without that assumption. The differences $L_n - U_{n+1}$ are $2.6 \cdot 10^{-3}$, $1.8 \cdot 10^{-4}$, $8.3 \cdot 10^{-4}$, $1.9 \cdot 10^{-4}$ and $3.6 \cdot 10^{-4}$, so the ordering is not a near thing at any step.

Corollary 4.6 (the endpoint support is strictly suboptimal from $n = 3$ on)**.** *For $n = 3, \dots, 8$, $V_n^{\mathrm{r}} < V_n$. Moreover $V_3 < 2V_3^{\mathrm{r}}$ and $21V_8^{\mathrm{r}} < V_8$; and against the certified rule at $n = 4$ the fully revealing experiment is not a best reply: $R_4(a^{(4)}, 1) < R_4(a^{(4)}, \frac{1}{2})$.*

Proof. $V_3^{\mathrm{r}} = \frac{1}{36} < L_3$, $V_5^{\mathrm{r}} = \frac{1}{100} < L_5$, $V_7^{\mathrm{r}} = \frac{1}{324} < L_7$ directly; for even n, $c_n = 2^{(n-2)/2}\sqrt{2}$ and $\sqrt{2} > 1.414$ give $4(1 + c_n)^2 \ge 58, 177, 606$ for $n = 4, 6, 8$, hence $V_n^{\mathrm{r}} \le \frac{1}{58}, \frac{1}{177}, \frac{1}{606}$, each below L_n. Next $U_3 < \frac{2}{36}$ and $21 \cdot \frac{1}{606} < L_8$. The last claim is a rational evaluation: $R_4(a^{(4)}, 1) = (a_0^{(4)})^2 \approx 0.01113$ and $R_4(a^{(4)}, \frac{1}{2}) \approx 0.03658$. $\square$

At $n = 3$ confining Nature to $\{\frac{1}{2}, 1\}$ costs her a factor $V_3/V_3^{\mathrm{r}} \in [1.409, 1.410]$; at $n = 8$ a factor in $[21.24, 21.25]$. The mechanism is visible in the last claim: against the certified rule the revealing experiment is worth less to Nature than the uninformative one, so the $\pi = 1$ atom of the source's $n \le 2$ equilibrium has left the support entirely, and the interior atom of Theorem 5.2 has replaced it. The source proves the existence of an interior atom only at $n = 3$ (its Theorem 1(ii)) and conjectures it beyond; the corollary is unconditional for $n = 3, \dots, 8$ and quantified, and is routine given Proposition 3.1 and Theorem 4.1(a).

5. Where an optimal Nature must put her mass

Let $\hat{d}_n = 2\hat{\pi}_n - 1$ and consider the quartic weight

$$\omega_n(d) = d^2\,(d - \hat{d}_n)^2, \qquad d = 2\pi - 1,$$

which vanishes exactly at the two conjectured atoms $\pi = \frac{1}{2}$ and $\pi = \hat{\pi}_n$ and is positive elsewhere. The second certificate strengthens Theorem 4.1(b) by a multiple of ω_n:

Lemma 5.1 (certificate with a margin)**.** *For $n = 3, \ldots, 8$ and every $\pi \in [\frac{1}{2}, 1]$,*

$$R_n(a^{(n)}, \pi) \;\le\; U_n - g_n\,\omega_n(2\pi - 1), \qquad (g_3, \ldots, g_8) = \big(\tfrac{2}{5}, \tfrac{1}{4}, \tfrac{1}{5}, \tfrac{1}{6}, \tfrac{1}{7}, \tfrac{1}{8}\big).$$

Proof. As in Theorem 4.1(b): at $\pi = 1$ the claim is $(a_0^{(n)})^2 \le U_n - g_n(1 - \hat{d}_n)^2$, a rational inequality; on $[\frac{1}{2}, 1)$ the claim is, after multiplying by the positive denominator, the nonnegativity on $[0, 1]$ of an integer-coefficient polynomial expressed in the same degree as C_n (one higher at $n = 3$), and a second Bernstein subdivision certificate proves it (column "margin tree" of Table 2). The constants g_n were found by trial with the same search; no optimality is claimed for them. □

Theorem 5.2 (support localisation)**.** *Let $n \in \{3, \ldots, 8\}$ and put $\varepsilon_n = (U_n - L_n)/g_n$, i.e.*

$$(\varepsilon_3, \ldots, \varepsilon_8) = \big(2 \cdot 10^{-11},\; 8.4 \cdot 10^{-11},\; 3.5 \cdot 10^{-11},\; 3 \cdot 10^{-11},\; 1.4 \cdot 10^{-11},\; 6.4 \cdot 10^{-12}\big).$$

Let $m \ge 0$, let $w_0, \ldots, w_{m-1} \ge 0$ with $\sum_{i<m} w_i = 1$, and let $p_0, \ldots, p_{m-1} \in [\frac{1}{2}, 1]$, so that $\sigma = \sum_{i<m} w_i \delta_{p_i}$ is a finitely supported mixture of experiments. If σ guarantees Nature at least L_n, that is

$$\sum_{i<m} w_i\, R_n(a, p_i) \;\ge\; L_n \qquad \textit{for every rule } a \in \mathbb{R}^{n+1},$$

then

$$\sum_{i<m} w_i\,(2p_i - 1)^2\,(2p_i - 1 - \hat{d}_n)^2 \;\le\; \varepsilon_n.$$

In particular, at $n = 4$: if $S \subseteq \{0, \ldots, m-1\}$ indexes the atoms at distance at least $1/100$ from both $\frac{1}{2}$ and $\hat{\pi}_4$, i.e. $|2p_i - 1| \ge \frac{1}{50}$ and $|2p_i - 1 - \hat{d}_4| \ge \frac{1}{50}$ for $i \in S$, then $\sum_{i \in S} w_i \le 525 \cdot 10^{-6}$.

Proof. Apply the hypothesis to $a = a^{(n)}$ and Lemma 5.1 to each atom:

$$L_n \;\le\; \sum_i w_i\, R_n(a^{(n)}, p_i) \;\le\; \sum_i w_i\big(U_n - g_n\,\omega_n(2p_i - 1)\big) = U_n - g_n \sum_i w_i\,\omega_n(2p_i - 1),$$

using $\sum_i w_i = 1$; rearrange. For the last claim, $\omega_4(2p_i - 1) \ge (\frac{1}{50})^4$ on S and the summands off S are nonnegative, so $\sum_{i \in S} w_i \le \varepsilon_4 \cdot 50^4 = \frac{21}{25 \cdot 10^{10}} \cdot 6.25 \cdot 10^6 = 5.25 \cdot 10^{-4}$ (Markov's inequality). □

The hypothesis is satisfied by every optimal mixture, in the max-min sense: if σ^* maximises Nature's guarantee $\inf_a \mathbb{E}_\sigma R_n(a, \cdot)$ among finitely supported mixtures, then its guarantee is at least that of $\hat{\sigma}_n$, which is at least L_n by Theorem 4.1(a). It is likewise satisfied by Nature's component of any saddle point, whose guarantee is the value. No minimax theorem is needed, and no assumption on σ^* beyond the guarantee. Restricting to finitely supported mixtures keeps the statement measure-free; the identical three lines work for a general $\sigma \in \Delta([\frac{1}{2}, 1])$ with the sum replaced by an integral (not formalised).

Remark 5.3 (reading the bound). Theorem 5.2 is the source's conjectured two-point structure derived rather than assumed, with an error bar, for the sample sizes at which the source states that a closed-form characterisation is intractable and imposes the structure as a necessary condition instead. It does not say the support is exactly $\{\frac{1}{2}, \hat{\pi}_n\}$ — it cannot, since $\hat{\pi}_n$ is a 6-digit rounding of the true atom — but it says every guaranteeing mixture has ω_n-weighted mass at most $\varepsilon_n \approx 10^{-11}$, and Markov's inequality converts this into a mass bound at any scale: at π-distance δ from both atoms the mass is at most $\varepsilon_n/(2\delta)^4$, which for $\delta = 1/100$ is $1.25 \cdot 10^{-4}$, $5.25 \cdot 10^{-4}$, $2.19 \cdot 10^{-4}$, $1.88 \cdot 10^{-4}$, $8.75 \cdot 10^{-5}$, $4.00 \cdot 10^{-5}$ for $n = 3, \ldots, 8$ (only the $n = 4$ instance is formalised; the others are the same one line). At $\delta = 10^{-3}$ the bound is vacuous, so the localisation is sharp to about three decimals in π, not to the eleven decimals of the bracket: the error bar on the value is quartic in the error bar on the atoms.

Remark 5.4 (the interior atom at $n = 3$; outside the formal development). At $n = 3$ the source *proves* that Nature's equilibrium strategy is $\sigma^* = (1 - w^*)\delta_{1/2} + w^*\delta_{\pi^*}$ for some $\pi^* \in (\frac{1}{2}, 1)$ (its Theorem 1(ii)), but gives no value for π^*. Granting that theorem, Theorem 5.2 locates π^*. Nature's component of a saddle point guarantees the value, so σ^* guarantees at least L_3; against the uninformed rule $a \equiv \frac{1}{2}$, whose regret is $\sum_k P_k(q_k - \frac{1}{2})^2 \le \frac{1}{4}$ at every π, this gives $w^* R_3(\frac{1}{2}, \pi^*) \ge L_3$, hence $w^* \ge 4L_3 > 0.1566$ and $R_3(\frac{1}{2}, \pi^*) \ge L_3$. By Lemma 2.3, $R_3(\frac{1}{2}, \pi) = \frac{d^2}{16}\big[\frac{(3+d^2)^2}{1+3d^2} + 3(1-d^2)\big] \le \frac{3}{4}d^2$ with $d = 2\pi - 1$, so $2\pi^* - 1 \ge \sqrt{4L_3/3} > 0.2284$. Theorem 5.2 then gives $w^*(2\pi^*-1)^2(2\pi^* - 1 - \hat{d}_3)^2 \le 2 \cdot 10^{-11}$, whence

$$|2\pi^* - 1 - \hat{d}_3| \le \sqrt{\frac{2 \cdot 10^{-11}}{0.1566 \cdot 0.2284^2}} < 5 \cdot 10^{-5}, \qquad \text{i.e.} \qquad |\pi^* - 0.899641| < 2.5 \cdot 10^{-5}.$$

Under the source's conjecture the same computation locates π_n^* for $n = 4, \dots, 8$; without it, Theorem 5.2 is what remains, and it is unconditional.

6. What is not claimed

The interior atom itself. A closed form, or even a certified bracket, for π_n^* is not proved. A bracket would need to show that Nature's best reply to $a^{(n)}$ is unique near $\hat{\pi}_n$ — a statement about the regret curve near its second peak, not about its maximum — and Theorem 5.2 deliberately avoids it: the localisation is of the mass, not of the atom. The minimal polynomial of π_n^*, obtainable in principle by eliminating (a, w) from the source's three conditions, was not computed.

Exactness of the two-point support. Theorem 5.2 bounds the mass away from the two atoms by ε_n; it does not show that it is zero. Showing that would require a certificate at the exact equilibrium, whose data are algebraic rather than rational.

General n. The brackets, the localisation and the strict suboptimality of the endpoint support are proved one certificate per n, for $n \le 8$. The pipeline is generic in n — the certificate degree is at most $n + \sum_{2k<n}(n-2k) \approx n^2/4 + 3n/2$ and the trees stayed near 31–39 nodes — so extending to the range $n \le 18$ the source plots is a matter of cost, not of method; a uniform statement in n is a different matter, and would need a lower bound on V_n that does not vanish as $n \to \infty$, which is the content of the source's limit game (its Theorem 2 and Corollary 1) rather than of a certificate.

Other losses. Under logarithmic loss (the source's Section 4) the regret is no longer rational in π, and the Bernstein route needs an enclosure of the logarithm; nothing here applies.

Prior work. We looked for prior bounds on this value and found none. A three-pass full-text scan of an arXiv corpus of 890,739 papers (including all of econ.TH and all of math.ST, through early September 2026), anchored on "minimax regret", "price of robustness", "Hodges–Lehmann", "least favourable prior" and "minimax risk" and rescored on the model's vocabulary, returns the source as the top paper by every rule; the nearest neighbours — Yata [15] on minimax-regret rules under partial identification, Lim and Park [16] on dynamically consistent statistical decisions, Halpern and Leung [17] on minimax weighted expected regret, Takeuchi and Barron [18] on minimax regret under logarithmic loss — have no Nature choosing the experiment's precision. A fixed-string pass on the decimal expansions of every constant in this paper, and on $4(1+\sqrt{n})^2$ and $2^{(n-1)/2}$, returns nothing. A further pass run for this paper — every line of the corpus containing "minimax regret" that also mentions binary signals, a two-point support or the precision of the signal or experiment — returns only the source. The source has no recorded citations (OpenAlex and Semantic Scholar, 2026-09-07). The minimax-regret treatment-choice literature the source positions itself against — Manski [9], Stoye [10], Hirano and Porter [11], Tetenov [12] — has fixed precision and linear regret, and the ambiguity-learning papers the source compares itself with (Frick, Iijima and Ishii [13]; Reshidi, Thereze and Zhang [14]) use a maxmin criterion. For the pre-arXiv literature, which the corpus cannot see, three further checks were run for this paper. A sweep of every work indexed in OpenAlex carrying "minimax regret" in its title or abstract before 2003 finds that strand divided between pre-test estimation, acceptance sampling, facility location and universal prediction under logarithmic loss, with no quadratic-loss value over a family of experiments. Chernoff and Moses [19], the pre-arXiv textbook that does treat minimax regret with binomial data, defines its squared-error regret function — its (10.2), $r(\theta, T) = c(\theta)(T - \theta)^2$ — against the true parameter, so the oracle's baseline is zero and regret collapses to risk; that is the structural reason the classical literature has nothing of this shape. The closest analogue we found is outside statistics: the competitive-minimax estimation of Eldar and Merhav [20], whose criterion is likewise a worst-case difference between the mean squared error of an estimator ignorant of a nuisance parameter and the optimal error of one that knows it, but in a linear-Gaussian model and with no small-sample values. MathSciNet and the zbMATH search interface were not available to us, so the pre-arXiv negative rests on OpenAlex, OEIS and the sources named and remains weaker than the arXiv one; if a prior computation of a value of this game exists, that is where it would be.

7. Verification

Propositions 2.1, 3.1 and 4.4, Corollaries 3.2, 4.5 and 4.6, and Theorems 4.1 and 5.2 rest on declarations formally verified in Lean 4 against *Mathlib* [21, 22] (Lean 4.32.0, Mathlib at its `v4.32.0` tag), whose statements are reproduced verbatim in Appendix A. The development is five modules, 2,946 lines in all: 50 theorems, 78 lemmas and 83 definitions. The first (465 lines) defines P_k, q_k, R_n, Γ_n and V_n exactly as in Section 2 (`pr`, `post`, `regret`, `guarantee`, `value`), proves the elementary bounds, Proposition 2.1 (`exists_unit_dominating`, `regret_pos_example`), the endpoint formulas of Lemma 2.3 (`regret_at_one`, `regret_at_half`), Lemma 2.4 (`quad_lb`), Proposition 3.1 for every $n \geq 1$ (`restrValue_le`, `restr_aRestr`, `restrValue_eq`) and Corollary 3.2 (`n_one`, `n_two`). The second (204 lines) holds the twelve certificate polynomials and subdivision trees as data, importing nothing but the shared Bernstein layer. The third (1,505 lines) proves, for each n, the identity of Lemma 2.3 for $a^{(n)}$ (`regret`n`_closed`), the acceptance of the tree (`bOk`n`_true`), Theorem 4.1(a) and (b) (`lower`n, `upper`n), the bracket (`value`n`_bracket`) and $V_n^{\mathrm{r}} < V_n$ (`restr_lt_value`n). The fourth (586 lines) proves Lemma 5.1 (`bOkM`n`_true`, `upper`n`_margin`) and Theorem 5.2 (`localise`n, `localise4_mass`). The fifth (186 lines) proves Proposition 4.4 (`upper`n`_sharp`, `value4_not_below`, `value4_not_above`, `hypotheses_nonvacuous`), the last claim of Corollary 4.6 (`one_not_best_reply4`, `restriction_price`) and Corollary 4.5 (`value_strict_anti`).

Compiled evaluation (`native_decide`) is used in exactly twelve declarations, the Boolean acceptance checks `bOk3_true`, …, `bOk8_true`, `bOkM3_true`, …, `bOkM8_true`, each of the form `checkHP tree (cert ()) = true` on the shared layer's arithmetic; the implication from that Boolean to the real inequality on the box is the kernel-checked theorem `Cert.checkHP_sound` of that layer, which represents a polynomial through its homogenisation so that the Bernstein coefficients are the output of the arithmetic rather than an input to be trusted. The axiom print (`#print axioms`) of all 50 theorems was re-run read-only for this paper. Fifteen declarations — all of Proposition 2.1, Proposition 3.1, Corollary 3.2, the six sharpness controls, `hypotheses_nonvacuous` and `one_not_best_reply4` — depend on `propext`, `Classical.choice`, `Quot.sound` only. The twelve Boolean checks depend on `propext` and one named axiom each, `bOk`n`_true._native.native_decide.ax_1_1` (respectively `bOkM`n`_true...`), and the remaining 23 declarations carry the standard three axioms plus exactly the native axioms of the certificates they use: one for each bracket, each $V_n^{\mathrm{r}} < V_n$, each of the seven localisation statements and the two two-sided controls at $n = 4$; two (`bOk3`, `bOk8`) for `restriction_price`; six for `value_strict_anti`. There is no `sorry` and no axiom declared by hand. As measured when the development was written, the five modules check cold in 28, 3, 96, 32 and 12 seconds of wall-clock time, 171 seconds in all and 336 seconds of CPU, the third module — whose six `ring` identities dominate — accounting for more than half; the axiom audit itself takes 15 seconds.

Outside the formal development, and labelled as such where they occur, are Remarks 2.2, 4.3 and 5.4, the five non-formalised instances of the Markov consequence, and the literature search of Section 6. Every number in the paper was recomputed for this paper in exact rational arithmetic (Python `fractions`) independently of the development's own scripts: the identity of Lemma 2.3 against the raw definition at six rational precisions for each n, the closed form of Proposition 3.1 against a direct maximisation of the definition for $n \leq 10$, the twelve certificate polynomials by an independent de Casteljau subdivision (which needed 16 to 20 leaves), the six lower bounds, the constants ε_n, the sharpness evaluations, the orderings, and the equilibrium candidate of Remark 4.3 at 60 digits. The certificate search that produced the data of Table 1 took under a minute; nothing in the development cost more than a few processor minutes. The development is currently in a private repository: repository reference to be added at submission.

Appendix A. The formal statements

The Lean statements corresponding to the results above, verbatim (statements only; proofs, docstrings and attributes omitted). The namespace `LeanProblemSpec.NatureMinimaxN4` is omitted. `pr n k` `π`, `post n k π`, `regret n a π`, `guarantee n a` and `value n` are $P_k(\pi)$, $q_k(\pi)$, $R_n(a,\pi)$, $\Gamma_n(a)$ and V_n; a rule is `a : ℕ → ℝ`, of which only `a 0`, …, `a n` are read; `restr n a w` is $R^{\mathrm{r}}_n(a,w)$, `cN n` is $c_n = \sqrt{2}^{\,n-1}$, `restrValue n` is V^{r}_n and `aRestr n` is a°; `a3`, …, `a8` are the rules $a^{(n)}$ of Table 1; `bOk`n and `bOkM`n are the Boolean Bernstein checks of the bracket and margin certificates; `∑ i ∈ range m, ...` is the finite sum over $i < m$.

The game and the restricted game (Propositions 2.1 and 3.1, Corollary 3.2).

```
theorem exists_unit_dominating {n : ℕ} (a : ℕ → ℝ) :
    ∃ b : ℕ → ℝ, (∀ k, 0 ≤ b k ∧ b k ≤ 1)
      ∧ (∀ π, 0 ≤ π → π ≤ 1 → regret n b π ≤ regret n a π)
      ∧ guarantee n b ≤ guarantee n a

theorem restrValue_le {n : ℕ} (hn : 1 ≤ n) (a : ℕ → ℝ) :
    restrValue n ≤ restr n a (1 / (1 + cN n))

theorem restr_aRestr {n : ℕ} (hn : 1 ≤ n) (w : ℝ) :
    restr n (aRestr n) w = restrValue n

theorem restrValue_eq {n : ℕ} (hn : 1 ≤ n) :
    (∀ a : ℕ → ℝ, restrValue n ≤ restr n a (1 / (1 + cN n)))
      ∧ (∀ w : ℝ, restr n (aRestr n) w = restrValue n)

theorem n_one : restrValue 1 = 1/16 ∧ (1 : ℝ) / (1 + cN 1) = 1/2
    ∧ aRestr 1 0 = 1/4 ∧ aRestr 1 1 = 3/4

theorem n_two : restrValue 2 = (3 - 2 * Real.sqrt 2) / 4
    ∧ (1 : ℝ) / (1 + cN 2) = Real.sqrt 2 - 1

theorem regret_pos_example : 0 < regret 1 (fun _ => 1/2) 1
```

The brackets (Theorem 4.1) and $V^{\mathrm{r}}_n < V_n$ (Corollary 4.6).

```
theorem bOk3_true : bOk3 () = true

theorem value3_bracket :
    (39156935533/1000000000000 : ℝ) ≤ value 3
      ∧ value 3 ≤ 39156935541/1000000000000

theorem restr_lt_value3 : restrValue 3 < value 3

theorem bOk4_true : bOk4 () = true

theorem value4_bracket :
    (36580028069/1000000000000 : ℝ) ≤ value 4
      ∧ value 4 ≤ 3658002809/100000000000

theorem restr_lt_value4 : restrValue 4 < value 4

theorem bOk5_true : bOk5 () = true

theorem value5_bracket :
    (4549919161/125000000000 : ℝ) ≤ value 5
      ∧ value 5 ≤ 7279870659/200000000000

theorem restr_lt_value5 : restrValue 5 < value 5

theorem bOk6_true : bOk6 () = true

theorem value6_bracket :
    (17785405593/500000000000 : ℝ) ≤ value 6
      ∧ value 6 ≤ 35570811191/1000000000000

theorem restr_lt_value6 : restrValue 6 < value 6
```

```
theorem bOk7_true : bOk7 () = true

theorem value7_bracket :
    (4422750047/125000000000 : ℝ) ≤ value 7
      ∧ value 7 ≤ 17691000189/500000000000

theorem restr_lt_value7 : restrValue 7 < value 7

theorem bOk8_true : bOk8 () = true

theorem value8_bracket :
    (35026812023/1000000000000 : ℝ) ≤ value 8
      ∧ value 8 ≤ 175134060119/5000000000000

theorem restr_lt_value8 : restrValue 8 < value 8
```

Support localisation (Theorem 5.2).

```
theorem bOkM3_true : bOkM3 () = true

theorem localise3 {m : ℕ} (w p : ℕ → ℝ) (hw : ∀ i, 0 ≤ w i)
    (hsum : ∑ i ∈ range m, w i = 1)
    (hp : ∀ i ∈ range m, 1/2 ≤ p i ∧ p i ≤ 1)
    (hgood : ∀ a : ℕ → ℝ,
      (39156935533/1000000000000 : ℝ) ≤ ∑ i ∈ range m, w i * regret 3 a (p i)) :
    ∑ i ∈ range m, w i * ((2 * p i - 1)^2 * (2 * p i - 1 - 399641/500000)^2)
      ≤ 1/50000000000

theorem bOkM4_true : bOkM4 () = true

theorem localise4 {m : ℕ} (w p : ℕ → ℝ) (hw : ∀ i, 0 ≤ w i)
    (hsum : ∑ i ∈ range m, w i = 1)
    (hp : ∀ i ∈ range m, 1/2 ≤ p i ∧ p i ≤ 1)
    (hgood : ∀ a : ℕ → ℝ,
      (36580028069/1000000000000 : ℝ) ≤ ∑ i ∈ range m, w i * regret 4 a (p i)) :
    ∑ i ∈ range m, w i * ((2 * p i - 1)^2 * (2 * p i - 1 - 7151/10000)^2)
      ≤ 21/250000000000

theorem bOkM5_true : bOkM5 () = true

theorem localise5 {m : ℕ} (w p : ℕ → ℝ) (hw : ∀ i, 0 ≤ w i)
    (hsum : ∑ i ∈ range m, w i = 1)
    (hp : ∀ i ∈ range m, 1/2 ≤ p i ∧ p i ≤ 1)
    (hgood : ∀ a : ℕ → ℝ,
      (4549919161/125000000000 : ℝ) ≤ ∑ i ∈ range m, w i * regret 5 a (p i)) :
    ∑ i ∈ range m, w i * ((2 * p i - 1)^2 * (2 * p i - 1 - 330217/500000)^2)
      ≤ 7/200000000000

theorem bOkM6_true : bOkM6 () = true

theorem localise6 {m : ℕ} (w p : ℕ → ℝ) (hw : ∀ i, 0 ≤ w i)
    (hsum : ∑ i ∈ range m, w i = 1)
    (hp : ∀ i ∈ range m, 1/2 ≤ p i ∧ p i ≤ 1)
    (hgood : ∀ a : ℕ → ℝ,
      (17785405593/500000000000 : ℝ) ≤ ∑ i ∈ range m, w i * regret 6 a (p i)) :
    ∑ i ∈ range m, w i * ((2 * p i - 1)^2 * (2 * p i - 1 - 303367/500000)^2)
      ≤ 3/100000000000

theorem bOkM7_true : bOkM7 () = true

theorem localise7 {m : ℕ} (w p : ℕ → ℝ) (hw : ∀ i, 0 ≤ w i)
    (hsum : ∑ i ∈ range m, w i = 1)
    (hp : ∀ i ∈ range m, 1/2 ≤ p i ∧ p i ≤ 1)
```

```
    (hgood : ∀ a : ℕ → ℝ,
      (4422750047/125000000000 : ℝ) ≤ ∑ i ∈ range m, w i * regret 7 a (p i)) :
    ∑ i ∈ range m, w i * ((2 * p i - 1)^2 * (2 * p i - 1 - 285269/500000)^2)
      ≤ 7/500000000000

theorem bOkM8_true : bOkM8 () = true

theorem localise8 {m : ℕ} (w p : ℕ → ℝ) (hw : ∀ i, 0 ≤ w i)
    (hsum : ∑ i ∈ range m, w i = 1)
    (hp : ∀ i ∈ range m, 1/2 ≤ p i ∧ p i ≤ 1)
    (hgood : ∀ a : ℕ → ℝ,
      (35026812023/1000000000000 : ℝ) ≤ ∑ i ∈ range m, w i * regret 8 a (p i)) :
    ∑ i ∈ range m, w i * ((2 * p i - 1)^2 * (2 * p i - 1 - 133947/250000)^2)
      ≤ 1/156250000000

theorem localise4_mass {m : ℕ} (w p : ℕ → ℝ) (S : Finset ℕ) (hS : S ⊆ range m)
    (hw : ∀ i, 0 ≤ w i) (hsum : ∑ i ∈ range m, w i = 1)
    (hp : ∀ i ∈ range m, 1/2 ≤ p i ∧ p i ≤ 1)
    (hgood : ∀ a : ℕ → ℝ,
      (36580028069/1000000000000 : ℝ) ≤ ∑ i ∈ range m, w i * regret 4 a (p i))
    (hfar : ∀ i ∈ S, (1/50 : ℝ) ≤ |2 * p i - 1| ∧ (1/50 : ℝ) ≤ |2 * p i - 1 - 7151/10000|) :
    ∑ i ∈ S, w i ≤ 525/1000000
```

Controls (Proposition 4.4), the price of the restriction (Corollary 4.6) and monotonicity (Corollary 4.5).

```
theorem upper3_sharp :
    ¬ (∀ π : ℝ, 1/2 ≤ π → π ≤ 1 → regret 3 a3 π ≤ 39156935540/1000000000000)

theorem upper4_sharp :
    ¬ (∀ π : ℝ, 1/2 ≤ π → π ≤ 1 → regret 4 a4 π ≤ 3658002808/100000000000)

theorem upper5_sharp :
    ¬ (∀ π : ℝ, 1/2 ≤ π → π ≤ 1 → regret 5 a5 π ≤ 36399353294/1000000000000)

theorem upper6_sharp :
    ¬ (∀ π : ℝ, 1/2 ≤ π → π ≤ 1 → regret 6 a6 π ≤ 3557081119/100000000000)

theorem upper7_sharp :
    ¬ (∀ π : ℝ, 1/2 ≤ π → π ≤ 1 → regret 7 a7 π ≤ 35382000377/1000000000000)

theorem upper8_sharp :
    ¬ (∀ π : ℝ, 1/2 ≤ π → π ≤ 1 → regret 8 a8 π ≤ 35026812023/1000000000000)

theorem value4_not_below : ¬ (value 4 ≤ 36580028068/1000000000000)

theorem value4_not_above : ¬ ((36580028091/1000000000000 : ℝ) ≤ value 4)

theorem hypotheses_nonvacuous :
    (1/2 : ℝ) < 17151/20000 ∧ (17151/20000 : ℝ) < 1
      ∧ (0:ℝ) < 465223/1000000 ∧ (465223/1000000 : ℝ) < 1
      ∧ 0 < regret 4 a4 (1/2)

theorem one_not_best_reply4 : regret 4 a4 1 < regret 4 a4 (1/2)

theorem restriction_price :
    value 3 < 2 * restrValue 3 ∧ 21 * restrValue 8 < value 8

theorem value_strict_anti :
    value 4 < value 3 ∧ value 5 < value 4 ∧ value 6 < value 5
      ∧ value 7 < value 6 ∧ value 8 < value 7
```

References

[1] Y.-K. Che, L. Li and T. Luo, *Learning Against Nature: Minimax Regret and the Price of Robustness*, arXiv:2602.15246v1 [econ.TH], 16 February 2026. Read in full. Section, statement, page and figure numbers are those of the e-print compiled from the arXiv source of version 1 (the only version as of 2026-09-07; the live record was checked on that date): Section 3 "Special Case: MSE with Binary Signals" (p. 7), Proposition 1 (p. 9), Theorem 1 (p. 10), the conditions (FOC), (Indif), (Local Opt) and the quoted passage (p. 11), Figure 2 (p. 12), Figure 3 (p. 13), Theorem 2 (p. 15), Theorem 3 (p. 17), Corollary 1 (p. 18) and footnote 1 (p. 2); the quotations in Section 1 are verbatim from that compilation, which runs to 72 pages.

[2] L. J. Savage, *The theory of statistical decision*, Journal of the American Statistical Association **46** (1951), no. 253, 55–67. Bibliographic data checked against Crossref (doi:10.1080/01621459.1951.10500768). The origin of the minimax-regret criterion; not read for this work.

[3] A. Wald, *Statistical Decision Functions*, John Wiley & Sons, New York, 1950. Bibliographic data checked against zbMATH (Zbl 0040.36402, which records "New York: John Wiley & Sons, Inc. 179 p. (1950)"). The minimax-risk criterion; not read for this work.

[4] J. L. Hodges, Jr. and E. L. Lehmann, *Some problems in minimax point estimation*, Annals of Mathematical Statistics **21** (1950), no. 2, 182–197. Bibliographic data checked against Crossref (doi:10.1214/aoms/1177729838); the article is open access on Project Euclid and was read there. The binomial estimator quoted in Remark 3.3 is its Section 5, Problem 1 (p. 190), printed as $X/(\sqrt{n}(\sqrt{n}+1)) + 1/(2(\sqrt{n}+1))$, which is the same estimator; the constant risk $1/(4(1+\sqrt{n})^2)$ and the prior, written there as the density $C(pq)^{\sqrt{n}/2-1}$, are displayed on the same page. The authors add that the result "had been obtained earlier by H. Rubin, to whom, therefore, the priority belongs".

[5] D. Blackwell, *Equivalent comparisons of experiments*, Annals of Mathematical Statistics **24** (1953), no. 2, 265–272. Bibliographic data checked against Crossref (doi:10.1214/aoms/1177729032). Cited for the notion of informativeness only; not read for this work.

[6] É. Marchand and B. MacGibbon, *Minimax estimation of a constrained binomial proportion*, Statistics & Decisions **18** (2000), no. 2, 129–167. Bibliographic data checked against zbMATH (Zbl 0953.62010, "Stat. Decis. 18, No. 2, 129-167 (2000)") and against Crossref (doi:10.1524/strm.2000.18.2.129), which returns no page range and indexes the journal under its later name *Statistics & Risk Modeling*. Cited as representative of the two-point least-favourable-prior literature on constrained binomial problems; not read for this work.

[7] V. Powers and B. Reznick, *Polynomials that are positive on an interval*, Transactions of the American Mathematical Society **352** (2000), no. 10, 4677–4692. Bibliographic data checked against Crossref (doi:10.1090/S0002-9947-00-02595-2).

[8] F. Boudaoud, F. Caruso and M.-F. Roy, *Certificates of positivity in the Bernstein basis*, Discrete & Computational Geometry **39** (2008), no. 4, 639–655. Bibliographic data checked against Crossref (doi:10.1007/s00454-007-9042-x).

[9] C. F. Manski, *Statistical treatment rules for heterogeneous populations*, Econometrica **72** (2004), no. 4, 1221–1246. Bibliographic data checked against Crossref (doi:10.1111/j.1468-0262.2004.00530.x). Cited from the source's bibliography; not read for this work.

[10] J. Stoye, *Minimax regret treatment choice with finite samples*, Journal of Econometrics **151** (2009), no. 1, 70–81. Bibliographic data checked against Crossref (doi:10.1016/j.jeconom.2009.02.013). Cited from the source's bibliography; not read for this work.

[11] K. Hirano and J. R. Porter, *Asymptotics for statistical treatment rules*, Econometrica **77** (2009), no. 5, 1683–1701. Bibliographic data checked against Crossref (doi:10.3982/ECTA6630). Cited from the source's bibliography; not read for this work.

[12] A. Tetenov, *Statistical treatment choice based on asymmetric minimax regret criteria*, Journal of Econometrics **166** (2012), no. 1, 157–165. Bibliographic data checked against Crossref (doi:10.1016/j.jeconom.2011.06.013). Cited from the source's bibliography; not read for this work.

[13] M. Frick, R. Iijima and Y. Ishii, *Efficient learning from ambiguous information*, working paper, 2023. Title, authors and year are those of the source's own bibliography entry, which records no journal; the authors' current title for the paper is *Efficient learning from ambiguously generated data* (February 2023), and we found it unpublished, with no DOI, arXiv identifier or discussion-paper number. It is *not* the same authors' *Learning efficiency of multiagent information structures* (J. Polit. Econ. **131** (2023), 3377–3414). Cited from the source's bibliography; not read for this work.

[14] P. Reshidi, J. Thereze and M. Zhang, *Asymptotic learning with ambiguous information*, American Economic Journal: Microeconomics **17** (2025), no. 3, 244–288. Bibliographic data checked against Crossref (doi:10.1257/mic.20230142). Cited from the source's bibliography; not read for this work.

[15] K. Yata, *Optimal decision rules under partial identification*, arXiv:2111.04926 [econ.EM], 2021 (v4). Read for the literature search only.

[16] C. Lim and Y. Park, *Dynamically consistent statistical decisions*, arXiv:2607.10519 [econ.EM], 2026. Read for the literature search only.

[17] J. Y. Halpern and S. Leung, *Weighted sets of probabilities and minimax weighted expected regret: new approaches for representing uncertainty and making decisions*, arXiv:1302.5681 [cs.GT], 2013. Read for the literature search only.

[18] J. Takeuchi and A. R. Barron, *Asymptotically minimax regret by Bayes mixtures*, arXiv:2406.17929 [cs.IT], 2024 (v2). Read for the literature search only.

[19] H. Chernoff and L. E. Moses, *Elementary Decision Theory*, John Wiley & Sons, New York, 1959. Bibliographic data checked against zbMATH (Zbl 0089.14406). Its Chapter 10 was read for the squared-error regret function (10.2) quoted in Section 6.

[20] Y. C. Eldar and N. Merhav, *A competitive minimax approach to robust estimation of random parameters*, IEEE Transactions on Signal Processing **52** (2004), no. 7, 1931–1946. Bibliographic data checked against Crossref (doi:10.1109/TSP.2004.828931). Cited for the shape of its regret criterion only; not read in full for this work.

[21] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction – CADE 28, Lecture Notes in Computer Science **12699**, Springer, 2021, pp. 625–635.

[22] 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.