Opponent Modeling in Incomplete-Information Continuous Colonel Blotto

Yuanyuan Zhang[1] **Gang Xiao**[1] **Feng Ye**[1] **Lingtao Xue**[1,2] **Zhipeng Du**[1]
[1]National Key Laboratory for Complex Systems Simulation, Beijing 100101, China
[2]School of Computer Science and Technology, Xidian University, Xi'an 710071, China
yuan18334701755@outlook.com, searchware@qq.com, 13701384998@139.com,
lingtaoxue@stu.xidian.edu.cn, duzhipeng1107@163.com

Abstract

We study when a compact summary of the opponent's posterior suffices for the interim decision problem of a realized type in incomplete-information continuous Colonel Blotto. For summaries that retain the distributions of finitely many continuous statistics of the opponent's allocation, such as per-battlefield marginals, we prove an exact ambiguity identity: the worst-case payoff gap between posterior laws that the summary cannot distinguish equals twice the sup-norm distance of a continuous payoff slice from the observable-additive span, the closed class of sums of functions of the retained statistics. Thus, for continuous payoffs, a summary is payoff-exact, preserving the value of every allocation, if and only if every payoff slice lies in that span. In exact-budget Blotto, first-order marginals are exact for continuous slices precisely when the slices lie in the closed coordinate-additive span, and battlefield-wise winner-take-all payoffs are exact by direct integration. The two-battlefield case is degenerate, whereas from three battlefields onward marginals can discard payoff-relevant dependence; more generally, exact-budget allocation games exhibit a strict r-order marginal hierarchy. The same criterion separates representation from estimation: the structured exact states studied here admit vanishing estimation error, whereas a non-exact state leaves an ambiguity that no sample size removes for rules based only on that state. Finally, approximate Bayesian Nash equilibrium remains PPAD-hard in a regularized diagonal exact-budget allocation subclass, even though the first-order states in our reduction are exact and their payoff-relevant means are estimable from polynomially many type-labeled samples.

1 Introduction

A player facing an unknown opponent rarely models the opponent's full strategy. In Colonel Blotto [Borel, 1921, Friedman, 1958], where two players split fixed budgets across K battlefields, a compact model may retain only the distribution of the opponent's allocation on each battlefield. These marginals describe one battlefield at a time but discard dependence between allocations across battlefields. When does the compact state still give the right value to every allocation the player might choose? For continuous payoffs, we show that this holds exactly when every opponent-side payoff slice can be uniformly approximated by sums of functions of the retained statistics. Without continuity, the state still preserves the value of such sums when their terms are integrable. When a continuous slice falls outside that class, the worst-case payoff gap between two posterior laws sharing the state equals exactly twice its uniform distance from the class, and any rule restricted to that state faces a worst-case criterion error that no sample size removes. We call a state with zero gap *payoff-exact*, the additive class the *observable-additive span*, and the equality between gap and distance the *ambiguity identity*.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

Example 1 (Three battlefields)**.** Let μ^- be uniform on the three cyclic permutations of $(0,1,2)$, namely $(0,1,2)$, $(1,2,0)$, $(2,0,1)$, and let μ^+ be uniform on the other three, $(0,2,1)$, $(2,1,0)$, $(1,0,2)$. Both live on the opponent simplex with budget 3, and under both every single battlefield is uniform on $\{0,1,2\}$: no per-battlefield model can tell the two opponents apart. Let the player's budget be $9/2$. For the additive winner-take-all payoff $u(\mathbf{x},\mathbf{y}) = \sum_a v_a \mathbf{1}\{x_a > y_a\}$, each fixed $\mathbf{x}$ gives a sum of functions of single coordinates, so the two laws give the same value at every feasible $\mathbf{x}$. For the majoritarian variant $u(\mathbf{x},\mathbf{y}) = \mathbf{1}\{|\{a : x_a > y_a\}| \geq 2\}$, the allocation $(1/2, 3/2, 5/2)$ wins with probability $2/3$ against μ^- and 1 against μ^+, whereas $(1/2, 5/2, 3/2)$, the same allocation with two battlefields swapped, wins with probability 1 and $2/3$. The strict ranking of the two allocations depends on information the marginals cannot carry, and the two laws have different best-response correspondences. Half-integer allocations keep every comparison strict, so tie-breaking never enters.

The non-additivity of the majoritarian payoff comes only from the final threshold on the number of battlefields won; the win indicators underneath are the standard battlefield-wise ones. The example does not force regret, however: $(3/2, 3/2, 3/2)$ wins two battlefields against every atom of both laws and is optimal under both.

Contributions. Table 1 records the structure each result uses. **Representation** (Sections 3–4): a criterion for when a reduced opponent state is lossless and the exact loss otherwise (Theorem 1); in exact-budget Blotto, first-order exactness of continuous slices is equivalent to membership in the closed coordinate-additive span, with a sharp $K = 2/K \geq 3$ boundary and a strict interaction-order hierarchy of block states (Theorems 2–4). **Learning** (Section 5): the structured exact states of Table 2 have vanishing statistical error at explicit rates, whereas non-exact states leave a criterion floor for rules restricted to that state (Theorem 5, Table 2). **Equilibrium** (Section 6): losslessness and learnability do not by themselves make Bayesian equilibrium tractable; approximate Bayesian Nash equilibrium (BNE) is PPAD-hard in a regularized diagonal allocation subclass whose first-order states are exact and learnable (Theorem 6).

Why Bayesian Blotto. We do not claim that Bayesian Blotto is the correct model for incomplete-information applications; our claim is narrower. Private types represent budgets, valuations, or signals. With correlated types, a realized type faces a posterior mixture over opponent types and type-contingent allocation laws, and per-battlefield marginals are the natural compact state to test. The exact-budget simplex makes the compression question sharp: at $K = 2$ one marginal determines the whole allocation because $y_2 = B - y_1$, whereas from $K = 3$ onward distinct feasible joint laws can share all one-dimensional marginals. A product action space would not by itself remove the difficulty, since marginals need not determine a correlated joint law there either; independence would be an added modeling restriction.

Table 1: Main results, the structure each uses, and the role of types. "None" means any compact metric opponent action space with finitely many continuous observables.

Result	Statement	Structure used	Types
Thm. 1, Cors. 2–3	Span criterion and ambiguity identity	None	One suffices; typed case type by type
Thm. 2, Cor. 4	Continuous slices: exactness iff membership in $\mathcal{A}_s^{(1)}$	Coordinate observables only	One suffices
Thms. 3–4	$K = 2/K \geq 3$ boundary; strict r-order hierarchy	Exact-budget simplex (permutation orbits, paired bits)	One suffices
Thm. 5	Criterion and regret floors for state-based rules	None	One suffices
Table 2	Rates for exact states	Bounded supports and local payoffs; the sharp rate and minimax bounds use the exact-budget simplex	One suffices
Thm. 6	PPAD-hard approximate BNE	Allocation simplex with a regularized diagonal bilinear payoff, not a contest success function	Growing type set

Related work. The continuous exact-budget Blotto game was introduced by Borel [1921] and solved in classical complete-information settings by Gross and Wagner [1950], Roberson [2006]; it models resource competition in political campaigns, military strategy, cybersecurity, and auctions. Asymmetric, heterogeneous-value, and general Lotto variants are studied in Kovenock and Roberson [2021], payoffs with externalities between battlefields in Golman and Page [2009], and multiplayer extensions in Boix-Adserà et al. [2021]. Incomplete-information Blotto, where private types encode budgets, valuations, or signals, has been studied for equilibrium existence and structural characterization in Adamo and Matros [2009], Kovenock and Roberson [2011], Ewerhart and Kovenock [2021], Kim and Kim [2024]. These papers characterize equilibria of specified Bayesian games; we study which posterior summaries preserve interim payoff functions.

On algorithms, Behnezhad et al. [2023] give fast methods for complete-information Blotto, Afiouni et al. [2026] analyze equilibrium existence and computation when each battlefield allocation is followed by a parametrized subgame, and Kontogiannis et al. [2025] develop kernelized no-regret algorithms for learning coarse correlated equilibria (CCE) in polyhedral games including Blotto and note that computing Nash equilibria in their general Blotto setting is PPAD-hard; none of the three involves exogenous private types. Opponent-modeling algorithms [Jing et al., 2024, 2025] and information compression in dynamic games with asymmetric information [Tang et al., 2026] concern sequential settings; we study the decision-sufficiency of a reduced state in a one-shot interim problem and derive the exact ambiguity when the state is not sufficient.

Our notion is related to decision-theoretic sufficiency and comparison of experiments [Blackwell, 1951, 1953], which are payoff-family-independent. We fix the payoff family, so our notion can be strictly coarser: distinct posterior laws may be payoff-equivalent while remaining statistically distinguishable. Nor is it identifiability: the state must determine the criterion, not the posterior. Le Cam's deficiency and approximate sufficiency [Le Cam, 1964] quantify information loss in general; for a fixed payoff slice, we identify the loss with a uniform approximation distance rather than bound it. Theorem 1 applies classical annihilator duality for uniform approximation by sums of univariate functions and pullback algebras [Ismailov, 2009, Asgarova and Ismailov, 2017]. It expresses this duality through opponent laws with identical retained observable laws; normalization of the Jordan parts yields the factor two.

PPAD-completeness of Nash equilibrium is established by Daskalakis et al. [2009]; tight inapproximability results for graphical games and PPAD appear in Deligkas et al. [2023, 2024]. Peng and Rubinstein [2024] study learning approximate correlated equilibria (CE) and CCE in incomplete-information games. Our hard instances additionally enforce exact and learnable first-order representations.

2 The opponent-modeling problem

2.1 Bayesian continuous Blotto

There are two players and K battlefields. Player i has type $t_i \in T_i$, and the joint prior over types is $F \in \Delta(T_1 \times T_2)$. Types may be correlated; for each realized t_i, let $F_{-i}(\cdot \mid t_i)$ denote the posterior over the opponent's type.

Main typed results use finite T_i. "Continuous Blotto" refers to the action space. The exact-state characterizations apply type by type on arbitrary compact opponent action spaces.

Player i's budget at type t_i is $b_i := \mathsf{B}_i(t_i) > 0$. The feasible allocations are the exact-budget Blotto simplex

$$\mathcal{X}_i(b_i) := \left\{ \mathbf{x}_i \in \mathbb{R}_+^K : \sum_{k=1}^{K} x_{ik} = b_i \right\}.$$

Write $\mathcal{Y}_K(B) := \{\mathbf{y} \in \mathbb{R}_+^K : \sum_{k=1}^K y_k = B\}$ for exact-budget opponent simplices (used in witness constructions). A behavioral strategy is a measurable kernel $\pi_i(\cdot \mid t_i) \in \Delta(\mathcal{X}_i(b_i))$. The primitive payoff $u_i(\mathbf{x}_i, \mathbf{y}, t_i, s)$ depends on the player's allocation $\mathbf{x}_i$, the opponent allocation $\mathbf{y}$, and the type realization (t_i, s). The interim payoff of realized type t_i against opponent strategy π_{-i} is

$$U_i(\mathbf{x}_i, t_i; \pi_{-i}) = \sum_{s \in T_{-i}} F_{-i}(s \mid t_i) \int_{\mathcal{X}_{-i}(b_{-i}(s))} u_i(\mathbf{x}_i, \mathbf{y}, t_i, s)\, \pi_{-i}(d\mathbf{y} \mid s).$$

2.2 Payoff-exact opponent models

Fix a realized type and, in this subsection, a single opponent type, and suppress both. For an opponent law $\mu \in \Delta(\mathcal{Y})$ on a compact opponent action space $\mathcal{Y}$, write $h_{\mathbf{x}_i}(\mathbf{y}) := u_i(\mathbf{x}_i, \mathbf{y})$ for the *payoff slice* at own action $\mathbf{x}_i$ and $F_\mu(\mathbf{x}_i) := \int h_{\mathbf{x}_i}\, d\mu$ for the criterion. A *reduction* R maps opponent laws to a lower-dimensional state.

Definition 1 (Three levels of exactness)**.** *A reduction R is* payoff-exact *if $R(\mu) = R(\nu)$ implies $F_\mu \equiv F_\nu$ on $\mathcal{X}_i(b_i)$; difference-exact if it implies that $F_\mu - F_\nu$ is constant on $\mathcal{X}_i(b_i)$; and* argmax-exact *if it implies* $\arg\max_{\mathcal{X}_i(b_i)} F_\mu = \arg\max_{\mathcal{X}_i(b_i)} F_\nu$.

Each level implies the next. Payoff-exactness preserves payoff values, comparisons, and best responses, and it is the notion used throughout, because the ambiguity identity, the criterion floor of Theorem 5(a), and the uniform estimation bounds behind Table 2 concern criterion values, not only the identity of an optimizer. Decision-level conclusions (best responses, ε-best-response sets, regret) need only difference-exactness. Difference-exactness, which ignores action-independent shifts, is characterized by the same theorem (Corollary 1); argmax-exactness is nonlinear and is not characterized here. We write $\mu \sim_{b_i} \nu$ when $F_\mu \equiv F_\nu$. Typed versions replace μ by a typed posterior opponent law (Section 3.2).

3 Exact opponent-state reductions

3.1 Observable-state exactness

Let $\mathcal{Y}$ be a compact metric opponent action space. Let $\rho_\ell : \mathcal{Y} \to \mathcal{Z}_\ell$, $\ell = 1, \dots, L$, be continuous maps into compact metric spaces. For an opponent law $\mu \in \Delta(\mathcal{Y})$, the *observable opponent state* is

$$R_\rho(\mu) := \big((\rho_1)_\#\mu, \dots, (\rho_L)_\#\mu\big).$$

In Blotto, first-order battlefield marginals correspond to the coordinate observables $\rho_k(\mathbf{y}) = y_k$.

The *closed observable-additive span* is

$$\mathcal{A}_\rho := \overline{\left\{h \in C(\mathcal{Y}) : h(\mathbf{y}) = \textstyle\sum_{\ell=1}^{L} \phi_\ell(\rho_\ell(\mathbf{y})),\ \phi_\ell \in C(\mathcal{Z}_\ell)\right\}}^{\|\cdot\|_\infty}.$$

The *observable ambiguity* is

$$\mathrm{Amb}_\rho(h) := \sup_{\substack{\mu,\nu \in \Delta(\mathcal{Y}) \\ R_\rho(\mu) = R_\rho(\nu)}} \left| \int h\, d\mu - \int h\, d\nu \right|.$$

Theorem 1 (Quantitative exact opponent-state theorem)**.** *Let $\mathcal{Y}$ be compact metric and $\rho_\ell : \mathcal{Y} \to \mathcal{Z}_\ell$ continuous. For every $h \in C(\mathcal{Y})$,*

$$\mathrm{Amb}_\rho(h) = 2\ \mathrm{dist}_\infty(h, \mathcal{A}_\rho).$$

Consequently, R_ρ is exact for h if and only if $h \in \mathcal{A}_\rho$.

Proof sketch. Let $d = \mathrm{dist}_\infty(h, \mathcal{A}_\rho)$. *Upper bound*: if $a \in \mathcal{A}_\rho$ and $R_\rho(\mu) = R_\rho(\nu)$, then $\int a\, d\mu = \int a\, d\nu$, so $|\int h\, d\mu - \int h\, d\nu| \le 2\|h - a\|_\infty$; infimize over a. *Lower bound* (for $d > 0$): Hahn–Banach gives a norm-one signed measure γ annihilating $\mathcal{A}_\rho$ with $\int h\, d\gamma = d$. Since constants lie in $\mathcal{A}_\rho$, $\gamma(\mathcal{Y}) = 0$; normalizing the Jordan parts gives probability measures μ, ν in the same observable fiber with payoff gap $2d$. Full details are in Appendix B. □

The upper bound uses only boundedness and measurability: for bounded measurable h, the same argument gives $\mathrm{Amb}_\rho(h) \le 2\|h - a\|_\infty$ for every sum $a = \sum_\ell \phi_\ell \circ \rho_\ell$ of bounded measurable ϕ_ℓ. Continuity enters through the equality and the necessity direction, where Hahn–Banach and Riesz act on $C(\mathcal{Y})$.

Remark 1 (Reading the span criterion)**.** Among continuous payoffs, R_ρ determines the expectations of exactly those that sums of functions of the retained observables can uniformly approximate. Whether a compact opponent model is lossless thus reduces to whether the payoff is an additive model in the retained statistics: $\mathcal{A}_\rho$ is the main-effects-only (generalized additive) model, and $\mathrm{dist}_\infty(h, \mathcal{A}_\rho)$ is the interaction it cannot capture. The distance measures criterion ambiguity, which need not cause decision regret. On a finite opponent action space, this distance is the value of the linear program $\min t$ subject to $|h(\mathbf{y}) - \sum_\ell \phi_\ell(\rho_\ell(\mathbf{y}))| \le t$ for all $\mathbf{y}$, the residual of an ℓ_∞ additive-model fit; on a grid in a continuous space the same program returns a grid-restricted certificate, and the full distance needs discretization control.

Corollary 1 (Difference-exactness)**.** *Under the hypotheses of Theorem 1, let the slices $h_{\mathbf{x}_i}$ be continuous, and let $\Delta_{\mathrm{diff}} := \sup_{\mathbf{x}_i, \mathbf{x}_i'} \sup_{R_\rho(\mu)=R_\rho(\nu)} |(F_\mu - F_\nu)(\mathbf{x}_i) - (F_\mu - F_\nu)(\mathbf{x}_i')|$ be the worst same-fiber change in a payoff difference. Then*

$$\Delta_{\mathrm{diff}} = 2 \sup_{\mathbf{x}_i, \mathbf{x}_i' \in \mathcal{X}_i(b_i)} \mathrm{dist}_\infty(h_{\mathbf{x}_i} - h_{\mathbf{x}_i'}, \mathcal{A}_\rho).$$

Hence R_ρ is difference-exact iff $h_{\mathbf{x}_i} - h_{\mathbf{x}_i^0} \in \mathcal{A}_\rho$ for every feasible $\mathbf{x}_i$ and one fixed reference $\mathbf{x}_i^0$. Difference-exactness is sufficient, but not necessary, for argmax-exactness.

Proof. Apply Theorem 1 to $h_{\mathbf{x}_i} - h_{\mathbf{x}_i'}$; one reference suffices because $\mathcal{A}_\rho$ is a linear subspace. Appendix B.2 gives the bounds for a single reference and the non-necessity example. □

The gap between the two upper levels is an action-independent background term and can be nonzero. For $u(\mathbf{x}, \mathbf{y}) = \sum_k x_k y_k + \kappa y_1^2 y_2$ all differences are additive, yet on the pair of Example 1 every action's value shifts by exactly $2\kappa/3$: best responses, ε-best-response sets, and regret coincide, while the payoff ambiguity is at least $2|\kappa|/3$.

3.2 Typed exact states

With types, write $\mathcal{Y}_s := \mathcal{X}_{-i}(b_{-i}(s))$. The opponent type posterior and strategy induce the *typed posterior opponent law* $\eta_i^\pi(ds, d\mathbf{y} \mid t_i) := F_{-i}(ds \mid t_i)\, \pi_{-i}(d\mathbf{y} \mid s)$, with criterion $F_\eta(\mathbf{x}_i) := \sum_s \int u_i(\mathbf{x}_i, \mathbf{y}, t_i, s)\, \eta(s, d\mathbf{y})$. For $s \in T_{-i}^+ := \{s : F_{-i}(s \mid t_i) > 0\}$, write η_s for the conditional law of $\mathbf{y}$ given s and $\eta_{s,k}$ for its kth coordinate marginal. Two laws are *payoff-equivalent*, $\eta \sim_{t_i} \tilde{\eta}$, if $F_\eta \equiv F_{\tilde{\eta}}$ on $\mathcal{X}_i(b_i)$, and Definition 1 applies with η in place of μ. The main low-order summary is the typed first-order opponent state.

Definition 2 (Typed first-order opponent state)**.**

$$\mathsf{M}_i^{(1)}(t_i; \pi_{-i}) := \big(F_{-i}(s \mid t_i),\, \pi_{-i,k}(\cdot \mid s) : s \in T_{-i}^+,\, k \in [K]\big),$$

where $\pi_{-i,k}(\cdot \mid s)$ is the kth coordinate marginal of $\pi_{-i}(\cdot \mid s)$.

Fix t_i; let $\mathcal{A}_{\rho,s}$ be the closed observable-additive span of $\rho_{\ell,s} : \mathcal{Y}_s \to \mathcal{Z}_{\ell,s}$, and let the *typed observable opponent state* be $\mathsf{R}_\rho(t_i; \pi_{-i}) := \big(F_{-i}(s \mid t_i),\, (\rho_{\ell,s})_\# \pi_{-i}(\cdot \mid s) : s \in T_{-i}^+,\, \ell \in [L]\big)$.

Corollary 2 (Typed exact-state characterization)**.** *Fix a realized type t_i with T_{-i} finite and unrestricted behavioral strategies, and assume every opponent slice $\mathbf{y} \mapsto u_i(\mathbf{x}_i, \mathbf{y}, t_i, s)$ is continuous on $\mathcal{Y}_s$ for feasible $\mathbf{x}_i$ and s with $F_{-i}(s \mid t_i) > 0$. The typed observable state $\mathsf{R}_\rho(t_i; \pi_{-i})$ is payoff-exact for t_i if and only if each such slice belongs to $\mathcal{A}_{\rho,s}$.*

Proof. Apply Theorem 1 type by type; failure at any s^* gives two equal-state laws with different criteria. □

Corollary 3 (r-order opponent states)**.** *With r-block observables $\rho_A(\mathbf{y}) = \mathbf{y}_A$, under the hypotheses of Corollary 2 the typed r-order state (Definition 3) is exact iff each positive-posterior slice lies in the closed r-block span. Independently, every r-local payoff with integrable block kernels is r-order exact by direct integration, without a continuity assumption.*

Remark 2 (Where multiple types matter)**.** Theorem 1 has no types, and Corollary 2 says that typed exactness is exactness type by type, so multiple types do not alter the membership criterion. All obstruction witnesses below are single-type objects, and Theorems 2–5 and Table 2 hold with one type. Types matter in two ways. The typed ambiguity aggregates per-type distances against

the posterior weights (Section 5.1), and a state that keeps only the marginals of the type mixture, $\sum_s F_{-i}(s \mid t_i)\pi_{-i,k}(\cdot \mid s)$, discarding the type label, can fail to be exact even for additive payoffs once the payoff depends on the opponent's type (Appendix B.3). Separately, Theorem 6 needs a growing type set (Remark 4).

3.3 Typed first-order Blotto states

The key observable is the coordinate projection $\rho_k(\mathbf{y}) = y_k$; write

$$\mathcal{A}_s^{(1)} := \overline{\left\{\mathbf{y} \mapsto \sum_{k=1}^{K} \phi_k(y_k) : \phi_k \in C([0, b_{-i}(s)])\right\}}^{\|\cdot\|_\infty}$$

for the *closed coordinate-additive span* of the opponent allocation. The typed first-order state $\mathsf{M}_i^{(1)}$ is Definition 2.

Theorem 2 (Typed first-order Blotto exactness)**.** *Fix a realized type t_i with T_{-i} finite and unrestricted behavioral strategies, and assume every positive-posterior opponent slice is continuous on its $\mathcal{Y}_s$. The typed first-order opponent state $\mathsf{M}_i^{(1)}(t_i; \pi_{-i})$ is payoff-exact for t_i if and only if, for every s with $F_{-i}(s \mid t_i) > 0$ and every feasible $\mathbf{x}_i$,*

$$\mathbf{y} \mapsto u_i(\mathbf{x}_i, \mathbf{y}, t_i, s) \in \mathcal{A}_s^{(1)}.$$

Proof. Corollary 2 with $\rho_{k,s}(\mathbf{y}) = y_k$: exactness equals slice membership in $\mathcal{A}_s^{(1)}$. □

Corollary 4 (Typed marginal-state invariance under componentwise locality)**.** *Assume componentwise additive-locality: $u_i(\mathbf{x}_i, \mathbf{y}, t_i, s) = \sum_{k=1}^{K} g_{k,t_i,s}(x_{ik}, y_k)$, with each kernel measurable and integrable under every relevant type-conditional battlefield marginal. Fix a realized type t_i. Let η and $\tilde{\eta}$ be two typed posterior opponent laws. If they have the same posterior type weights and the same type-conditional battlefield marginals,*

$$F_{-i}(s \mid t_i) = \tilde{F}_{-i}(s \mid t_i), \qquad \eta_{s,k} = \tilde{\eta}_{s,k} \quad (F_{-i}(s \mid t_i) > 0,\ k \in [K]),$$

then $\eta \sim_{t_i} \tilde{\eta}$.

Proof. Immediate by integrating each component against the corresponding type-conditional battlefield marginal. □

This includes measurable winner-take-all and threshold kernels defined battlefield by battlefield.

4 Battlefield marginals

4.1 The $K = 2$ boundary

At $K = 2$ the opponent simplex is one-dimensional: $y_2 = B - y_1$. Hence every continuous opponent slice is a function of one coordinate and lies in the first-order coordinate-additive span. First-order exactness is therefore automatic, while additive decompositions remain nonunique: $q(y_1) - q(B - y_2) = 0$ on the simplex for every measurable q. For $K \geq 3$ they are unique up to an affine gauge (Proposition 3). The nondegenerate marginal theory begins at $K = 3$, where Theorem 3 shows first-order exactness can fail.

4.2 The $K \geq 3$ regime

For a single opponent type and componentwise additive-local payoffs, under an identification condition on the kernels, the marginal tuple is the coarsest payoff-exact state (Appendix C). For interaction payoffs, first-order marginals can be too coarse.

Theorem 3 (First-order marginal obstruction for every $K \geq 3$)**.** *For every $K \geq 3$, there exist two probability measures μ^+, μ^- supported on a single exact-budget opponent simplex $\mathcal{Y}_K(B)$ with identical first-order battlefield marginals and a two-local payoff $u(\mathbf{x}, \mathbf{y})$ such that $F_{\mu^+} \not\equiv F_{\mu^-}$ on the player's exact-budget simplex. The obstruction is already realized at the minimal battlefield count $K = 3$.*

Proof. For the laws $\mu^{\pm}$ of Example 1, $\mathbb{E}_{\mu^+}[Y_1^2 Y_2] = 4/3$ and $\mathbb{E}_{\mu^-}[Y_1^2 Y_2] = 2/3$ although their first-order marginals agree. Thus $u(\mathbf{x}, \mathbf{y}) = x_1 y_1^2 y_2$ gives criteria $\frac{4}{3}x_1$ and $\frac{2}{3}x_1$. For $K > 3$, pad by $\delta_0^{\otimes(K-3)}$. □

The obstruction does not need an algebraic cross-term: the majoritarian payoff of Example 1 separates the same pair (Proposition 6), while the additive winner-take-all payoff lies on the exact side (Corollary 4).

4.3 The r-order marginal hierarchy

First-order states are the first level of a hierarchy indexed by interaction order.

Definition 3 (Typed r-order opponent state)**.**

$$\mathsf{M}_i^{(r)}(t_i; \pi_{-i}) := \big(F_{-i}(s \mid t_i),\ \pi_{-i,A}(\cdot \mid s) : s \in T_{-i}^+,\ 1 \le |A| \le r\big),$$

where $\pi_{-i,A}(\cdot \mid s)$ is the A-coordinate marginal of type s's law.

Definition 4 (r-local payoff)**.** *u_i is* r-local *if $u_i = \sum_{A \subseteq [K],\, 1 \le |A| \le r} g_{A,t_i,s}(\mathbf{x}_{i,A}, \mathbf{y}_A)$ for measurable kernels $g_{A,t_i,s}$.*

The 1-local case is componentwise additive-locality.

Theorem 4 (Exact-budget marginal hierarchy)**.** *Let $\mathsf{M}_i^{(r)}$ be the typed r-order opponent state. Under the hypotheses of Theorem 2, it is exact iff every positive-posterior slice lies in the closed additive span of coordinate-block observables of size at most r. Independently, without assuming continuity, every r-local payoff whose displayed kernels are integrable under every positive-posterior type-conditional opponent law is $\mathsf{M}_i^{(r)}$-exact by direct integration. Conversely, for every $r \ge 1$ and $K \ge 2(r+1)$, no r-order state is payoff-exact for all $(r+1)$-local payoffs.*

Proof sketch. The continuous characterization follows from Corollary 3. Direct sufficiency holds because each block-local term integrates against the retained block marginal, so all payoff comparisons are preserved. Strictness is witnessed by the paired-bit construction of Proposition 7 (Appendix D): any r observed coordinates reveal at most r of the $r+1$ latent bits, so the even- and odd-parity laws share all r-order marginals; yet an $(r+1)$-local product payoff separates them. □

5 Learning exact opponent states

Opponent modeling can fail in two ways: statistical (the chosen model is correct but must be estimated from data) and structural (the chosen model is not payoff-exact, so even perfect knowledge of that model leaves payoff-relevant ambiguity).

The span test checks a state for exactness before any data are collected: coordinate-additive slices call for first-order states and r-local slices for block states. If the test fails, better estimation of the same state cannot recover the omitted criterion. In Example 1, each displayed allocation's majoritarian value differs by $1/3$ between two laws with the same marginals, so any rule that reports a value from the first-order state alone is wrong by at least $1/6$ on one of them at every sample size (Proposition 6). This is a criterion floor, not a regret lower bound (Example 1).

5.1 Representation error versus estimation error

Fix t_i, $\mathcal{X}_i(b_i)$, and $F_{-i}(\cdot \mid t_i)$, and let $\mathcal{E}_{t_i}$ be the admissible typed posterior laws. Let R_ρ have spans $\mathcal{A}_{\rho,s}$ and continuous slices $h_{\mathbf{x}_i,s}(\mathbf{y}) := u_i(\mathbf{x}_i, \mathbf{y}, t_i, s)$; for $\eta \in \mathcal{E}_{t_i}$ write $F_\eta(\mathbf{x}_i) := \sum_s \int h_{\mathbf{x}_i,s}(\mathbf{y})\, \eta(s, d\mathbf{y})$. The *typed ambiguity* is

$$\Delta_\rho(t_i) := \sup_{\mathbf{x}_i \in \mathcal{X}_i(b_i)} \sup_{\substack{\eta, \tilde{\eta} \in \mathcal{E}_{t_i}: \\ \mathsf{R}_\rho(\eta) = \mathsf{R}_\rho(\tilde{\eta})}} \left| F_\eta(\mathbf{x}_i) - F_{\tilde{\eta}}(\mathbf{x}_i) \right|.$$

When $\mathcal{E}_{t_i}$ contains every typed law with type posterior $F_{-i}(\cdot \mid t_i)$, Theorem 1 applied type by type gives

$$\Delta_\rho(t_i) = 2 \sup_{\mathbf{x}_i} \sum_{s \in T_{-i}^+} F_{-i}(s \mid t_i)\, \mathrm{dist}_\infty(h_{\mathbf{x}_i,s}, \mathcal{A}_{\rho,s}),$$

Table 2: Learning rates after exact opponent-state reduction (T samples), for one opponent type or type-independent payoff kernels. The last row holds with probability at least $1 - \delta$.

Payoff class	Opponent model to estimate	Rate
Lipschitz componentwise (L_g-Lip.)	First-order marginals $(\mu_1, \dots, \mu_K)$	$O(KL_gM/\sqrt{T})$
Linearly modulated ($g_k = t\phi_k(y)$)	First-order marginals (mean modulators)	$O(b_i R_\phi \sqrt{\log K/T})$
r-local block-separable with bounded factors	Payoff-relevant block means ($O(K^r)$)	$O(K^r/\sqrt{T})$
Battlefield-wise WTA ($g_k = v_k \mathbf{1}\{t > y\}$)	First-order CDFs	$2\sum_k \lvert v_k\rvert \sqrt{\log(2K/\delta)/(2T)}$

where the separating signed gaps are oriented type by type. Under the same condition, the state is exact iff $h_{\mathbf{x}_i,s} \in \mathcal{A}_{\rho,s}$ for all feasible $\mathbf{x}_i$ and positive-posterior s.

Write $\mathcal{S}_{\rho,t_i} := \{\mathsf{R}_\rho(\eta) : \eta \in \mathcal{E}_{t_i}\}$. An R_ρ*-based reconstruction* is a map $\Gamma : \mathcal{S}_{\rho,t_i} \to \mathbb{R}^{\mathcal{X}_i(b_i)}$, and an R_ρ*-based decision rule* is a map $\Lambda : \mathcal{S}_{\rho,t_i} \to \mathcal{X}_i(b_i)$; both receive only the state.

Theorem 5 (Lower bounds for state-based rules)**.** (a) *If every relevant $h_{\mathbf{x}_i,s} \in \mathcal{A}_{\rho,s}$, the criterion is reconstructible from R_ρ; the converse holds when $\mathcal{E}_{t_i}$ contains every typed law with type posterior $F_{-i}(\cdot \mid t_i)$. For any $\mathcal{E}_{t_i}$, if $\Delta_\rho(t_i) > 0$, every R_ρ-based reconstruction satisfies*

$$\sup_{\eta \in \mathcal{E}_{t_i}} \sup_{\mathbf{x}_i \in \mathcal{X}_i(b_i)} |\Gamma(\mathsf{R}_\rho(\eta))(\mathbf{x}_i) - F_\eta(\mathbf{x}_i)| \;\geq\; \tfrac{1}{2}\Delta_\rho(t_i).$$

(b) *Let $\eta, \tilde{\eta} \in \mathcal{E}_{t_i}$ lie in one R_ρ-fiber and have α-strongly concave criteria with maximizers $\mathbf{x}^*_\eta$ and $\mathbf{x}^*_{\tilde{\eta}}$. Every R_ρ-based decision rule incurs regret at least $\frac{\alpha}{8}\|\mathbf{x}^*_\eta - \mathbf{x}^*_{\tilde{\eta}}\|_2^2$ on at least one of the two laws.*

Proof. (a) Sufficiency, and necessity for the full posterior class, follow from Corollary 2. For two laws in one fiber with gap approaching $\Delta_\rho(t_i)$, the reconstructions coincide; the triangle inequality gives half the gap at one law. (b) The rule chooses the same allocation at both laws; by the triangle inequality it lies at distance at least $\|\mathbf{x}^*_\eta - \mathbf{x}^*_{\tilde{\eta}}\|_2/2$ from one maximizer, and strong concavity gives the stated regret at that law. □

The floor in part (a) uses only the fiber. The converse in part (a) can fail on smaller classes: on a class of Dirac laws, for instance, first-order marginals identify the law. Part (b) additionally needs strong concavity and distinct maximizers, so a non-exact state need not force positive regret (Example 1). Both parts concern rules whose input is the retained state, not learners that estimate the full joint law.

5.2 Rates for exact states

Under the sampling model in Appendix E, an exact state leaves only statistical error. For a sample-dependent allocation $\widehat{\mathbf{x}}_T$ built from T i.i.d. samples, let $\mathrm{Regret}_T := \mathbb{E}[\max_{\mathbf{x}_i \in \mathcal{X}_i(b_i)} F_\mu(\mathbf{x}_i) - F_\mu(\widehat{\mathbf{x}}_T)]$. Table 2 collects the rates of plug-in rules that maximize the empirical criterion of the retained state; formal statements, constants, high-probability versions, and proofs are in Appendix E. The componentwise row assumes L_g-Lipschitz kernels on $[0, M]$ (Theorem 9). The linearly modulated row, with R_ϕ the largest modulator range width, uses the player's own exact-budget simplex through $\sup_{\mathcal{X}_i(b_i)} |\sum_k \alpha_k x_{ik}| = b_i \max_k |\alpha_k|$, the source of the factor b_i, and is matched by a minimax lower bound on a paired-bit exact-budget subclass (Theorems 11 and 13). The r-order row assumes factorized block kernels $g_A = f_A\phi_A$ with bounded factors (Theorem 15). Winner-take-all kernels are not Lipschitz, so their row replaces the Wasserstein bound by the Dvoretzky–Kiefer–Wolfowitz inequality on battlefield CDFs (Proposition 8). Bandit-feedback variants are in Appendix E.

6 Hardness of equilibrium opponent modeling

A BNE requires each type's opponent model to be generated by the opponent's equilibrium strategy, which is itself a best response to its own model of the other player. We show this self-consistency condition remains PPAD-hard even when every individual model is first-order exact and efficiently learnable from type-labeled samples.

6.1 Regularized diagonal exact-budget allocation games

Write $\Delta_K := \{\mathbf{x} \in \mathbb{R}_+^K : \sum_{a=0}^{K-1} x_a = 1\}$ for the unit simplex over battlefield indices $\{0, \dots, K-1\}$.

Definition 5 (Regularized diagonal bilinear payoff)**.** *On the battlefield index set $\{0, \dots, K-1\}$ with $\mathbf{x}_i, \mathbf{x}_{-i} \in \Delta_K$, $\lambda > 0$, and player-specific rational coefficients $\theta^i_{t_i,t_{-i},a}, d^i_{t_i,a} \in \mathbb{Q}$,*

$$u_i(\mathbf{x}_i, \mathbf{x}_{-i}, t_i, t_{-i}) = \sum_{a=0}^{K-1} \Big(\theta^i_{t_i,t_{-i},a}\, x_{i,a}\, x_{-i,a} + d^i_{t_i,a}\, x_{i,a} \Big) - \tfrac{\lambda}{2} \|\mathbf{x}_i\|_2^2.$$

Proposition 1 (Typed first-order sufficiency in the regularized class)**.** *In a finite-type game with payoffs in Definition 5, the interim payoff of realized type t_i depends on the opponent strategy only through the typed first-order opponent state $\mathsf{M}_i^{(1)}(t_i; \pi_{-i})$.*

Proof. For each opponent type s, the payoff slice is the linear function $\mathbf{y} \mapsto \sum_a \theta^i_{t_i,s,a} x_{i,a} y_a$ plus terms independent of $\mathbf{y}$, hence lies in $\mathcal{A}_s^{(1)}$. Theorem 2 reduces the criterion to posterior weights and battlefield means. □

6.2 The BNE computation problem

Definition 6 (Regularized diagonal BNE computation problem)**.** *An instance consists of finite type sets T_1, T_2 with $|T_i| = n$; rational prior weights giving every type positive marginal probability; player-specific rational coefficients $\theta^i_{t_i,t_{-i},a}, d^i_{t_i,a}$ and regularization $\lambda \in \mathbb{Q}_{>0}$ as in Definition 5; battlefield count $K \geq 2$; and rational tolerance $\gamma = 1/\mathrm{poly}(n)$. The output is $\{p_{i,t_i,a}\}$, rationals of polynomial bit length, where $p_{i,t_i,a} := \int x_a \, \pi_i(d\mathbf{x} \mid t_i)$ is type t_i's mean allocation on battlefield a under some mixed γ-BNE π; such outputs exist by Proposition 10.*

For a tolerance $\beta \geq 0$, a mixed profile π is a *β-BNE* if, for every player i, type t_i, and alternative mixed action $\sigma_i \in \Delta(\Delta_K)$,

$$\int U_i(\mathbf{x}_i, t_i; \pi_{-i})\, \pi_i(d\mathbf{x}_i \mid t_i) \geq \int U_i(\mathbf{x}_i, t_i; \pi_{-i})\, \sigma_i(d\mathbf{x}_i) - \beta.$$

In this class, opponent strategies enter each interim payoff only through $p_{-i,s,a} := \int y_a \, \pi_{-i}(d\mathbf{y} \mid s)$. For a mixed β-best response with mean $\bar{\mathbf{x}}$ and exact best response $\mathbf{x}^*$, Jensen's inequality and λ-strong concavity give $\frac{\lambda}{2}\|\bar{\mathbf{x}} - \mathbf{x}^*\|_2^2 \leq \beta$, hence $\|\bar{\mathbf{x}} - \mathbf{x}^*\|_2 \leq e := \sqrt{2\beta/\lambda}$.

Lemma 1 (Dominated dummy battlefields and two-action response)**.** *Fix $K \geq 3$, $\lambda > 0$, $\tau > 3\lambda/2$, and consider $\max_{\mathbf{x} \in \Delta_K} \{S_1 x_1 - \tau \sum_{a \geq 2} x_a - \frac{\lambda}{2}\|\mathbf{x}\|_2^2\}$. Its unique maximizer has $x_a^* = 0$ for all $a \geq 2$, $x_1^* = \mathrm{clip}(\frac{1}{2} + \frac{S_1}{2\lambda}, 0, 1)$, and if a mixed response has expected payoff within β of optimum with mean $\bar{\mathbf{x}}$, then $|\bar{x}_1 - x_1^*| \leq \sqrt{2\beta/\lambda}$.*

6.3 PPAD-hardness in the regularized diagonal subclass

The hardness parameter is $n = |V|$ with K fixed: each node of a threshold-equilibrium instance $G = (V, E)$ becomes a type, player 2 runs a copy gadget, player 1 runs a threshold gadget, and battlefields $a \geq 2$ are dominated dummies.

Remark 3 (Scope and distance to contest-success payoffs)**.** Definition 5 uses regularized diagonal bilinear payoffs; the reduction establishes no hardness result for canonical contest-success payoffs. At fixed $\mathbf{x}$, both $x_a/(x_a + y_a + \varepsilon)$ and $\mathbf{1}\{x_a > y_a\}$ depend on y_a alone and are first-order exact by Corollary 4, but their expectations generally require the full battlefield marginal. Our copy and threshold gadgets instead require linear dependence on battlefield means, supplied by $\theta_a x_a y_a$; no choice of θ, d, λ turns this payoff into either canonical form above. Exactness therefore extends beyond this subclass, while the hardness theorem is specific to it.

Theorem 6 (PPAD-hardness in the regularized diagonal subclass)**.** *Fix an integer $K \geq 3$, a constant $c > 0$, and a rational constant $\varepsilon_0 \in (0, 1/6)$. For every polynomial-bit rational $\lambda \in [n^{-c}, 1]$, the regularized diagonal BNE computation problem (Definition 6) with target tolerance γ is PPAD-hard for every rational inverse-polynomial $0 < \gamma \leq \lambda\varepsilon_0^2/128$. The hardness holds even though, for every realized type and every own action, the opponent slice is coordinate-additive in the opponent allocation, so the typed first-order opponent state is exact throughout the reduction.*

Proof sketch. We analyze the baseline tolerance $\beta := \lambda\varepsilon_0^2/128$; every γ-BNE with $0 < \gamma \le \beta$ is also a β-BNE. We reduce from ε_0-threshold equilibrium, PPAD-hard for every $\varepsilon_0 < 1/6$ even with in- and out-degrees at most two [Deligkas et al., 2024, Theorem 4.23]: $z \in [0,1]^V$ satisfies $\sum_{u\in N_v} z_u > \frac{1}{2} + \varepsilon_0 \Rightarrow z_v \le \varepsilon_0$ (and the lower-threshold analogue).

Set $T_1 = T_2 = V$, decode $z_v := p_{1,v,1}$, abbreviate $p_{i,v} := p_{i,v,1}$; battlefields $0/1$ active, $a \ge 2$ dominated ($\tau = 4\lambda$); $D := 2$ (the source in-degree bound), $A := 4\lambda/\varepsilon_0$, and $e := \sqrt{2\beta/\lambda} = \varepsilon_0/(4D)$.

P2 copy gadget: $d^2_{v,1} := -\lambda$, $\theta^2_{v,u,1} := 2\lambda n\, \mathbf{1}_{u=v}$. Score advantage $2\lambda p_{1,v} - \lambda$; Lemma 1 forces $|p_{2,v} - p_{1,v}| \le e$.

P1 threshold gadget: $d^1_{v,1} := A/2$, $\theta^1_{v,u,1} := -An\, \mathbf{1}_{u\in N_v}$. Copy error $|\sum_{u\in N_v}(p_{2,u} - z_u)| \le De \le \varepsilon_0/4$, so both threshold-equilibrium conditions follow. Thus decoding a target solution solves the source problem, proving PPAD-hardness. Full reduction details and bit-complexity checks are in Appendix F. □

Remark 4 (Source of the hardness)**.** Affine gradients and λ-strong concavity yield a rational mixed complementarity system with K pairs per player-type, plus simplex equalities and free multipliers. Enumerating its active-set branches and solving a rational linear feasibility problem per branch returns an exact BNE in time $2^{K(|T_1|+|T_2|)}\mathrm{poly}(|\mathcal{I}|)$ for an instance $\mathcal{I}$ of bit-size $|\mathcal{I}|$ (Proposition 10). This is polynomial when $K(|T_1| + |T_2|) = O(\log|\mathcal{I}|)$, so PPAD-hardness in that regime would make every PPAD problem solvable in polynomial time. Our reduction keeps K fixed and grows the type set, outside that regime; the exact-budget simplex also couples the coordinates.

6.4 Learnable representation, hard equilibrium

Corollary 5 (Learnability–hardness separation)**.** *The class in Definition 5 simultaneously has first-order exact states (Theorem 2), $O(R_{t_i}\sqrt{\log K/T})$ regret from type-labeled samples, where R_{t_i} bounds the range width of the type-weighted score samples (Proposition 9), and PPAD-hard approximate BNE for every fixed $K \ge 3$ (Theorem 6). First-order exactness and learnability are not sufficient for equilibrium tractability.*

Conclusion. For continuous opponent-side slices, an observable opponent state is payoff-exact if and only if every slice lies in its closed observable-additive span; the worst-case same-state payoff gap is twice the uniform distance to that span. In exact-budget Blotto, first-order states are exact for additive payoffs, including winner-take-all, but can fail for interaction payoffs when $K \ge 3$; block states form a strict hierarchy. The structured exact states studied here admit vanishing estimation error, whereas a non-exact state leaves a criterion floor for rules restricted to that state. Exactness and learnability nevertheless do not make Bayesian equilibrium tractable: approximate BNE stays PPAD-hard in a regularized diagonal allocation subclass.

Limitations. The ambiguity identity and its necessity direction assume continuous payoff slices on compact action spaces; discontinuous payoffs are covered through the upper bound and direct integration only. Typed results assume finitely many types. When payoffs depend on the opponent's type, the learning guarantees assume samples labeled with that type, since allocation samples alone need not identify the criterion. The main theorem characterizes payoff- and difference-exactness; argmax-only preservation is nonlinear and is not characterized. The setting is a one-shot interim decision problem, and extending the criterion to nonstationary or sequential opponents, as in multi-agent reinforcement learning, requires a different object. Our reduction covers only the regularized diagonal subclass and does not settle the complexity of approximate BNE for two-player continuous Bayesian Blotto with canonical contest-success payoffs.

Acknowledgments and Disclosure of Funding

We thank the anonymous reviewers, the area chair, and the NeurIPS 2026 program committee for their constructive feedback. This work was supported by the National Key Laboratory for Complex Systems Simulation. The authors declare no competing interests.

References

Tim Adamo and Alexander Matros. A Blotto game with incomplete information. *Economics Letters*, 105(1):100–102, 2009. doi: 10.1016/j.econlet.2009.06.025.

Salam Afiouni, Jakub Černý, Chun Kai Ling, and Christian Kroer. Colonel Blotto with battlefield games. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(20):16612–16620, 2026. doi: 10.1609/aaai.v40i20.38702.

Aida Kh. Asgarova and Vugar E. Ismailov. On the representation by sums of algebras of continuous functions. *Comptes Rendus. Mathématique*, 355(9):949–955, 2017. doi: 10.1016/j.crma.2017.09.015.

Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002. doi: 10.1137/S0097539701398375.

Soheil Behnezhad, Sina Dehghani, Mahsa Derakhshan, Mohammedtaghi Hajiaghayi, and Saeed Seddighin. Fast and simple solutions of Blotto games. *Operations Research*, 71(2):506–516, 2023. doi: 10.1287/opre.2022.2261.

David Blackwell. Comparison of experiments. In Jerzy Neyman, editor, *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102, Berkeley and Los Angeles, 1951. University of California Press. URL `https://projecteuclid.org/euclid.bsmsp/1200500222`.

David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24 (2):265–272, 1953. doi: 10.1214/aoms/1177729032.

Enric Boix-Adserà, Benjamin L. Edelman, and Siddhartha Jayanti. The multiplayer Colonel Blotto game. *Games and Economic Behavior*, 129:15–31, 2021. doi: 10.1016/j.geb.2021.05.002.

Émile Borel. La théorie du jeu et les équations intégrales à noyau symétrique. *Comptes rendus hebdomadaires des séances de l'Académie des sciences*, 173:1304–1308, 1921. URL `https://www.academie-sciences.fr/pdf/dossiers/Borel/Borel_pubd1.htm`.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009. doi: 10.1137/070699652.

Argyrios Deligkas, John Fearnley, Alexandros Hollender, and Themistoklis Melissourgos. Tight inapproximability for graphical games. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(5):5600–5607, 2023. doi: 10.1609/aaai.v37i5.25695.

Argyrios Deligkas, John Fearnley, Alexandros Hollender, and Themistoklis Melissourgos. Pure-Circuit: Tight inapproximability for PPAD. *Journal of the ACM*, 71(5):31:1–31:48, 2024. doi: 10.1145/3678166.

Aryeh Dvoretzky, Jack Kiefer, and Jacob Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics*, 27(3):642–669, 1956. doi: 10.1214/aoms/1177728174.

Christian Ewerhart and Dan Kovenock. A class of N-player Colonel Blotto games with multidimensional private information. *Operations Research Letters*, 49(3):418–425, 2021. doi: 10.1016/j.orl.2021.03.010.

Lawrence Friedman. Game-theory models in the allocation of advertising expenditures. *Operations Research*, 6(5):699–709, 1958. doi: 10.1287/opre.6.5.699.

Russell Golman and Scott E. Page. General Blotto: Games of allocative strategic mismatch. *Public Choice*, 138(3–4):279–299, 2009. doi: 10.1007/s11127-008-9359-x.

Oliver Alfred Gross and R. A. Wagner. A continuous Colonel Blotto game. Research Memorandum RM-408, RAND Corporation, Santa Monica, CA, 1950. URL `https://www.rand.org/pubs/research_memoranda/RM408.html`.

Vugar E. Ismailov. On the theorem of M. Golomb. *Proceedings – Mathematical Sciences*, 119(1): 45–52, 2009. doi: 10.1007/s12044-009-0005-4.

Yuheng Jing, Bingyun Liu, Kai Li, Yifan Zang, Haobo Fu, Qiang Fu, Junliang Xing, and Jian Cheng. Opponent modeling with in-context search. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 61549–61591. Curran Associates, Inc., 2024. doi: 10.52202/079017-1967.

Yuheng Jing, Kai Li, Bingyun Liu, Ziwen Zhang, Haobo Fu, Qiang Fu, Junliang Xing, and Jian Cheng. Offline opponent modeling with truncated Q-driven instant policy refinement. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu, editors, *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 28256–28287. PMLR, 2025. URL `https://proceedings.mlr.press/v267/jing25a.html`.

Geofferey Jiyun Kim and Jerim Kim. A sequentially variant Blotto game with one-sided and incomplete information. *Decision Analytics Journal*, 13:100524, 2024. doi: 10.1016/j.dajour.2024.100524.

Andreas Kontogiannis, Vasilis Pollatos, Gabriele Farina, Panayotis Mertikopoulos, and Ioannis Panageas. Efficient kernelized learning in polyhedral games beyond full information: From Colonel Blotto to congestion games. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen, editors, *Advances in Neural Information Processing Systems*, volume 38, pages 145808–145858. Curran Associates, Inc., 2025. doi: 10.52202/085713-4879.

Dan Kovenock and Brian Roberson. A Blotto game with multi-dimensional incomplete information. *Economics Letters*, 113(3):273–275, 2011. doi: 10.1016/j.econlet.2011.08.009.

Dan Kovenock and Brian Roberson. Generalizations of the General Lotto and Colonel Blotto games. *Economic Theory*, 71(3):997–1032, 2021. doi: 10.1007/s00199-020-01272-2.

Marek Kuczma. *An Introduction to the Theory of Functional Equations and Inequalities: Cauchy's Equation and Jensen's Inequality*. Birkhäuser, Basel, second edition, 2009. doi: 10.1007/978-3-7643-8749-5. Edited by Attila Gilányi.

Lucien Le Cam. Sufficiency and approximate sufficiency. *The Annals of Mathematical Statistics*, 35 (4):1419–1455, 1964. doi: 10.1214/aoms/1177700372.

Pascal Massart. The tight constant in the Dvoretzky–Kiefer–Wolfowitz inequality. *The Annals of Probability*, 18(3):1269–1283, 1990. doi: 10.1214/aop/1176990746.

T. S. Motzkin and E. G. Straus. Maxima for graphs and a new proof of a theorem of Turán. *Canadian Journal of Mathematics*, 17:533–540, 1965. doi: 10.4153/CJM-1965-053-6.

Binghui Peng and Aviad Rubinstein. The complexity of approximate (coarse) correlated equilibrium for incomplete information games. In Shipra Agrawal and Aaron Roth, editors, *Proceedings of Thirty Seventh Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, pages 4158–4184. PMLR, 2024. URL `https://proceedings.mlr.press/v247/peng24a.html`.

Brian Roberson. The Colonel Blotto game. *Economic Theory*, 29(1):1–24, 2006. doi: 10.1007/s00199-005-0071-5.

Dengwang Tang, Vijay Subramanian, and Demosthenis Teneketzis. Information compression in dynamic games. *Dynamic Games and Applications*, 16(4):1279–1336, 2026. doi: 10.1007/s13235-025-00663-1.

Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer, New York, 2009. doi: 10.1007/b13794.

A Notation and standing assumptions

The appendix follows the order of Table 1. Appendix B proves the exact opponent-state reduction results. Appendix B.4 treats componentwise Blotto programs. Appendices C and D prove the Blotto marginal geometry and hierarchy results. Appendix E proves the learning rates, the marginal-state regret floor, and the minimax lower bounds. Appendix F proves the PPAD lower bound.

Symbol	Meaning
K	number of battlefields
$T_i,\ t_i$	player-i type space and realized type
n	number of types $\lvert T_i\rvert$ (PPAD section)
$F,\ F_{-i}(\cdot \mid t_i)$	joint prior and posterior over opponent type
T_{-i}^{+}	positive-posterior opponent types $\{s : F_{-i}(s \mid t_i) > 0\}$
$b_i := \mathsf{B}_i(t_i)$	realized type's budget
$\mathcal{X}_i(b_i)$	realized type's exact-budget allocation simplex
$\mathcal{Y}_s,\ \mathcal{Y}_K(B)$	opponent type-s simplex; exact-budget simplex with budget B
Δ_K	unit simplex over battlefield indices $\{0, \dots, K-1\}$
$\pi_i(\cdot \mid t_i)$	behavioral strategy kernel
$\pi_{-i,k}(\cdot \mid s)$	kth coordinate marginal of opponent type s's law
$\pi_{-i,A}(\cdot \mid s)$	A-coordinate block marginal of opponent type s's law
$u_i,\ U_i$	primitive and interim payoffs
$F_\mu(\mathbf{x}_i)$	payoff functional integrated against opponent law μ
$\mathbf{x}_\mu^*$	Bayes-optimal allocation under posterior μ
η_i^π	typed posterior opponent law
$\eta_s,\ \eta_{s,k}$	type-s conditional opponent law and its kth coordinate marginal
$L,\ \ell$	number of observables; observable index $\ell \in [L]$
$R_\rho,\ \rho_\ell$	observable-state reduction and observables
$\mathcal{Z}_\ell$	compact metric codomain of observable ρ_ℓ
$\mathcal{A}_\rho,\ \mathcal{A}_s^{(1)}$	closed observable-additive span; closed coordinate-additive span
$\mathcal{A}_{\rho,s}$	closed observable-additive span for opponent type s
$\mathrm{Amb}_\rho(h)$	worst-case payoff ambiguity of slice h
$h_{\mathbf{x}_i,s}$	opponent payoff slice at own action $\mathbf{x}_i$, type s
R_ρ	typed observable opponent state
$\mathsf{M}_i^{(1)},\ \mathsf{M}_i^{(r)}$	typed first- and r-order opponent states
$\mu \sim_{b_i} \nu,\ \eta \sim_{t_i} \tilde{\eta}$	payoff-equivalence (single type; typed)
Δ_{diff}	worst same-fiber change in a payoff difference (Corollary 1)
$\Delta_\rho(t_i)$	typed ambiguity for state R_ρ
$g_k,\ g_A$	componentwise and block-local payoff kernels
$\mathcal{G}_k(b_i)$	battlefield kernel family $\{y \mapsto g_k(t,y) : t \in [0, b_i]\}$
$\mathfrak{N}_k$	marginal class for battlefield k
$L_g,\ M$	Lipschitz constant, support radius
$R_\phi = R,\ \bar{R}_\phi$	largest individual modulator width $\max_k R_k$; common width of all modulator values (bandit)
α	strong concavity constant
W_1	Wasserstein-1 distance
$\mathrm{reg}_T,\ \mathrm{Regret}_T$	one-shot regret and its expectation
$\mathfrak{R}_T$	cumulative regret against the best fixed allocation over T rounds
$\hat{\mu}_k^{(T)},\ \hat{\phi}_k$	empirical marginal and empirical modulator mean
$\mathfrak{M}_i(b_i)$	admissible posterior-law class
$\mathcal{M}$	codomain (state space) of a marginal-based opponent model R
v_{k,t_i}^π	componentwise expected battlefield-k payoff function
$\mathfrak{M}_i^{\mathrm{pair}}(b_i)$	paired-bit exact-budget minimax subclass
$\theta^i,\ d^i,\ \lambda$	regularized diagonal payoff coefficients
$p_{i,t_i,a},\ p_{-i,s,a}$	mean allocation: player i type t_i / opponent type s on battlefield a
$q_{t_i,a}^i,\ \widehat{q}_{t_i,a}^i$	interim score and its empirical estimator (learning proposition)
R_{t_i}	range-width bound for score samples Z_a^m
$A,\ D,\ e$	PPAD reduction parameters: score scale, in-degree bound, approximation error
$G = (V,E),\ N_v,\ z_v$	threshold-equilibrium source graph and node variables
$\varepsilon_0,\ \beta,\ \gamma$	source tolerance, baseline BNE tolerance, and target BNE tolerance

Symbols used only within a proof are defined locally.

Whenever an appendix proof fixes a realized type or type pair, we suppress only the fixed indices. For typed criteria, retain the opponent-type-dependent kernel inside each conditional integral and sum against $F_{-i}(\cdot \mid t_i)$. This algebraic reduction does not remove the observation requirements of the learning guarantees in Appendix E.

B Supplementary details for exact opponent-state reductions

This appendix gives complete proofs of the exact-state results stated in Section 3; the main text provides proof sketches only.

Full proof of Theorem 1. Let $d := \operatorname{dist}_\infty(h, \mathcal{A}_\rho)$.

Upper bound. For any $a \in \mathcal{A}_\rho$ and any μ, ν with $R_\rho(\mu) = R_\rho(\nu)$,

$$\left| \int h \, d\mu - \int h \, d\nu \right| \le \left| \int (h - a) \, d\mu \right| + \left| \int (h - a) \, d\nu \right| \le 2\|h - a\|_\infty.$$

Infimizing over $a \in \mathcal{A}_\rho$ gives $\operatorname{Amb}_\rho(h) \le 2d$.

Lower bound. If $d = 0$, the upper bound already gives $\operatorname{Amb}_\rho(h) = 0 = 2d$, so assume $d > 0$. We construct $\mu, \nu \in \Delta(\mathcal{Y})$ with $R_\rho(\mu) = R_\rho(\nu)$ and $\int h \, d\mu - \int h \, d\nu = 2d$.

Step 1 (Annihilator functional). Consider the coset $[h] = h + \mathcal{A}_\rho$ in $C(\mathcal{Y})/\mathcal{A}_\rho$. Since $\mathcal{A}_\rho$ is closed, the quotient norm equals $\|[h]\| = d$. By the Hahn–Banach theorem, there exists a norm-one linear functional $\Psi : C(\mathcal{Y}) \to \mathbb{R}$ with $\Psi\big|_{\mathcal{A}_\rho} = 0$ and $\Psi(h) = d$.

Step 2 (Riesz representation). By the Riesz–Markov theorem, $\Psi(f) = \int f \, d\gamma$ for a regular Borel signed measure γ on $\mathcal{Y}$ with $\|\gamma\|_{\mathrm{TV}} = \|\Psi\| = 1$.

Step 3 (Zero mass). Since constant functions lie in $\mathcal{A}_\rho$ (take $\phi_\ell \equiv c_\ell$ and sum), $\Psi(\mathbf{1}) = 0$, so $\gamma(\mathcal{Y}) = 0$.

Step 4 (Jordan decomposition). Write $\gamma = \gamma^+ - \gamma^-$ (Jordan parts). From $\gamma(\mathcal{Y}) = 0$ and $\|\gamma\|_{\mathrm{TV}} = \gamma^+(\mathcal{Y}) + \gamma^-(\mathcal{Y}) = 1$, we get $\gamma^+(\mathcal{Y}) = \gamma^-(\mathcal{Y}) = \frac{1}{2}$.

Step 5 (Probability measures). Set $\mu := 2\gamma^+$ and $\nu := 2\gamma^-$, both probability measures on $\mathcal{Y}$.

Step 6 (Same observable fiber). For each ℓ and every $\phi_\ell \in C(\mathcal{Z}_\ell)$, the function $\phi_\ell \circ \rho_\ell \in \mathcal{A}_\rho$, so $\Psi(\phi_\ell \circ \rho_\ell) = 0$, i.e., $\int \phi_\ell \circ \rho_\ell \, d\gamma = 0$. Hence $\int \phi_\ell \, d(\rho_\ell)_\# \mu = 2 \int \phi_\ell \circ \rho_\ell \, d\gamma^+ = 2 \int \phi_\ell \circ \rho_\ell \, d\gamma^- = \int \phi_\ell \, d(\rho_\ell)_\# \nu$. Since $C(\mathcal{Z}_\ell)$ separates measures on compact metric $\mathcal{Z}_\ell$, $(\rho_\ell)_\# \mu = (\rho_\ell)_\# \nu$ for every ℓ, i.e., $R_\rho(\mu) = R_\rho(\nu)$.

Step 7 (Payoff gap). $\int h \, d\mu - \int h \, d\nu = 2 \int h \, d\gamma = 2\Psi(h) = 2d$.

Step 8 (Conclusion). $\operatorname{Amb}_\rho(h) \ge |\int h \, d\mu - \int h \, d\nu| = 2d$, so together $\operatorname{Amb}_\rho(h) = 2d = 2 \operatorname{dist}_\infty(h, \mathcal{A}_\rho)$. The exactness characterization $\operatorname{Amb}_\rho(h) = 0 \Leftrightarrow h \in \mathcal{A}_\rho$ is immediate. □

The typed, r-order, and first-order specializations (Corollaries 2, 3 and Theorem 2) follow by applying the theorem type by type with, respectively, general observables $\rho_{\ell,s}$, block projections $\mathbf{y} \mapsto \mathbf{y}_A$ with $1 \le |A| \le r$, and coordinate projections $\mathbf{y} \mapsto y_k$; exactness then reduces to slice membership in $\mathcal{A}_{\rho,s}$, in the closed r-block span, and in $\mathcal{A}_s^{(1)}$, respectively.

The componentwise additive-locality proposition (Proposition 2) is in Section B.4 below.

B.1 Bounded measurable payoffs

For bounded measurable h and any sum $a = \sum_\ell \phi_\ell \circ \rho_\ell$ with bounded measurable ϕ_ℓ, the integral $\int a \, d\mu = \sum_\ell \int \phi_\ell \, d(\rho_\ell)_\# \mu$ is a function of $R_\rho(\mu)$. The upper-bound argument above therefore gives $|\int h \, d\mu - \int h \, d\nu| \le 2\|h - a\|_\infty$ whenever $R_\rho(\mu) = R_\rho(\nu)$. Only the matching lower bound uses $h \in C(\mathcal{Y})$, through Steps 1 and 2.

B.2 Levels of exactness

Proof of Corollary 1. For fixed $\mathbf{x}_i, \mathbf{x}_i'$, the inner supremum in Δ_{diff} is $\mathrm{Amb}_\rho(h_{\mathbf{x}_i} - h_{\mathbf{x}_i'})$ with $h_{\mathbf{x}_i} - h_{\mathbf{x}_i'} \in C(\mathcal{Y})$, so Theorem 1 gives the identity. If $h_{\mathbf{x}_i} - h_{\mathbf{x}_i^0} \in \mathcal{A}_\rho$ for all $\mathbf{x}_i$, then $h_{\mathbf{x}_i} - h_{\mathbf{x}_i'} = (h_{\mathbf{x}_i} - h_{\mathbf{x}_i^0}) - (h_{\mathbf{x}_i'} - h_{\mathbf{x}_i^0}) \in \mathcal{A}_\rho$ and $\Delta_{\text{diff}} = 0$; the converse is the case $\mathbf{x}_i' = \mathbf{x}_i^0$. With $d_0 := \sup_{\mathbf{x}_i} \mathrm{dist}_\infty(h_{\mathbf{x}_i} - h_{\mathbf{x}_i^0}, \mathcal{A}_\rho)$, the same decomposition and the triangle inequality for dist_∞ give

$$2d_0 \le \Delta_{\text{diff}} \le 4d_0,$$

so a single reference determines Δ_{diff} only up to a factor of two, and the identity needs the pairwise supremum.

Difference-exactness gives $F_\nu = F_\mu + c$ on each fiber, which preserves the argmax. It is not necessary. Take $K = 3$, opponent budget 3, and $u(\mathbf{x}, \mathbf{y}) = x_1(1 + y_1^2 y_2)$. Every criterion is $F_\mu(\mathbf{x}) = x_1\, \mathbb{E}_\mu[1 + Y_1^2 Y_2]$ with a positive coefficient, so the argmax $\{\mathbf{x} \in \mathcal{X}_i(b_i) : x_1 = b_i\}$ is the same for every law and the first-order state is argmax-exact. On the pair of Example 1 the coefficients are $7/3$ and $5/3$, so payoff differences change within the fiber. □

Background-term example. For $u(\mathbf{x}, \mathbf{y}) = \sum_k x_k y_k + \kappa y_1^2 y_2$, the difference $h_{\mathbf{x}} - h_{\mathbf{x}'} = \sum_k (x_k - x_k') y_k$ lies in $\mathcal{A}^{(1)}$, so the first-order state is difference-exact. Under the laws of Example 1, the coordinate means agree while $\mathbb{E}_{\mu^+}[Y_1^2 Y_2] = 4/3$ and $\mathbb{E}_{\mu^-}[Y_1^2 Y_2] = 2/3$. Hence $F_{\mu^+} = F_{\mu^-} + 2\kappa/3$ identically, and the payoff ambiguity of every slice is at least $2|\kappa|/3$.

B.3 States that discard the type label

Let $T_{-i} = \{s_1, s_2\}$ with $F_{-i}(s_j \mid t_i) = \frac{1}{2}$, a common opponent budget B, $K \ge 2$, and $u_i(\mathbf{x}_i, \mathbf{y}, t_i, s_1) = x_{i1} y_1 = -u_i(\mathbf{x}_i, \mathbf{y}, t_i, s_2)$, which is coordinate-additive in $\mathbf{y}$ for each type. Under η, type s_1 plays $B\mathbf{e}_1$ and type s_2 plays $B\mathbf{e}_2$; under $\tilde{\eta}$ the two types swap. The battlefield marginals of the type mixture, $\sum_s F_{-i}(s \mid t_i)\pi_{-i,k}(\cdot \mid s)$, coincide (each is $\frac{1}{2}\delta_0 + \frac{1}{2}\delta_B$ on battlefields $1, 2$ and δ_0 elsewhere), yet $F_\eta(\mathbf{x}_i) = \frac{1}{2} B x_{i1}$ and $F_{\tilde{\eta}}(\mathbf{x}_i) = -\frac{1}{2} B x_{i1}$. The typed first-order state, which keeps the type label, separates the two laws.

B.4 Componentwise Blotto programs

Assumption 1 (Componentwise additive-local payoff)**.** *$u_i(\mathbf{x}_i, \mathbf{y}, t_i, s) = \sum_{k=1}^K g_{k,t_i,s}(x_{ik}, y_k)$ for jointly measurable kernels $g_{k,t_i,s}$. (Type indices suppressed when irrelevant: g_k.)*

Assumption 2 (Integrability of battlefield payoffs)**.** *$y \mapsto g_{k,t_i,s}(t, y)$ is integrable under every relevant type-conditional battlefield marginal, for every k, t, t_i, and s.*

Componentwise locality implies first-order exactness and separates own-side optimization:

$$U_i(\mathbf{x}_i, t_i; \pi_{-i}) = \sum_{k=1}^K v^\pi_{k,t_i}(x_{ik}) \quad \text{where} \quad v^\pi_{k,t_i}(t) := \sum_s F_{-i}(s \mid t_i) \int g_{k,t_i,s}(t, y)\, \pi_{-i,k}(dy \mid s),$$

so best response reduces to $\max_{\mathbf{x}_i \in \mathcal{X}_i(b_i)} \sum_k v^\pi_{k,t_i}(x_{ik})$.

Proposition 2 (Componentwise marginal programs and additive-locality)**.** *Fix $K \ge 2$, an opponent budget $B > 0$, and an exact-budget posterior class $\mathfrak{M}_i(b_i)$ containing $\delta_{\mathbf{y}}$ for every $\mathbf{y} \in \mathcal{Y}_K(B)$. For measurable kernels g_k satisfying Assumption 2, the componentwise exact opponent model $F_\mu(\mathbf{x}_i) = \sum_{k=1}^K \int g_k(x_{ik}, y)\, d\mu_k(y)$ holds for all $\mu \in \mathfrak{M}_i(b_i)$ and $\mathbf{x}_i \in \mathcal{X}_i(b_i)$* if and only if *$u(\mathbf{x}_i, \mathbf{y}) = \sum_{k=1}^K g_k(x_{ik}, y_k)$ on $\mathcal{Y}_K(B)$.*

Proof. Sufficiency follows by integrating each term against its coordinate marginal. For necessity, evaluating at $\mu = \delta_{\mathbf{y}}$ gives $u = \sum_k g_k$. □

Remark 5 (Exactness is not separability)**.** Theorem 2 concerns the opponent state only. It requires that, for fixed $(\mathbf{x}_i, t_i, s)$, the slice $\mathbf{y} \mapsto u_i(\mathbf{x}_i, \mathbf{y}, t_i, s)$ lie in the closed coordinate-additive span $\mathcal{A}_s^{(1)}$. It does not require the map $\mathbf{x}_i \mapsto u_i(\mathbf{x}_i, \mathbf{y}, t_i, s)$ to be separable. Componentwise additive-locality is therefore strictly stronger: it gives first-order exactness *and* turns the reduced best response into a

separable program. For example, for $K \geq 3$, $u(\mathbf{x}, \mathbf{y}) = x_1^2 x_2 y_1$ is first-order exact in the opponent state, but when $\mathbb{E}_\mu[Y_1] > 0$ its reduced best-response objective $\mathbb{E}_\mu[Y_1] x_1^2 x_2$ is not componentwise separable on the own simplex.

C Battlefield marginal geometry

Proposition 3 (Simplex affine gauge, $K \geq 3$)**.** *Assume $K \geq 3$ and $b_i > 0$. Let $\phi_k, \psi_k : [0, b_i] \to \mathbb{R}$ be measurable with $\sum_k \phi_k(x_{ik}) = \sum_k \psi_k(x_{ik})$ for all $\mathbf{x}_i \in \mathcal{X}_i(b_i)$. Then there exist a_k, c with $\sum_k a_k = -cb_i$ such that $\psi_k(t) = \phi_k(t) + a_k + ct$ for all t and all k.*

The identification and minimality results below concern a single opponent type, or the pooled allocation law when the payoff kernel is opponent-type independent. They do not assert minimality of the typed conditional-marginal tuple.

Definition 7 (Affine-determining battlefield family)**.** *The family $\mathcal{G}_k(b_i) := \{y \mapsto g_k(t, y) : t \in [0, b_i]\}$ is* affine-determining *on $\mathfrak{N}_k$ if, whenever $\int g_k(t, y)\, d\nu(y) = \int g_k(t, y)\, d\tilde{\nu}(y) + a + ct$ for all t, it follows that $\nu = \tilde{\nu}$ and $a = c = 0$.*

Theorem 7 (Marginal identification, $K \geq 3$)**.** *Under Assumptions 1, 2, $K \geq 3$, and kernel families $\mathcal{G}_k(b_i)$ that are affine-determining on classes $\mathfrak{N}_k$, two posterior measures $\mu, \tilde{\mu}$ with $\mu_k, \tilde{\mu}_k \in \mathfrak{N}_k$ for every k are payoff-equivalent on $\mathcal{X}_i(b_i)$ iff $\mu_k = \tilde{\mu}_k$ for all k.*

Theorem 8 (Marginal minimality, $K \geq 3$)**.** *Under the hypotheses of Theorem 7, the battlefield-marginal tuple is the coarsest payoff-exact statistic: any T with $T(\mu) = T(\tilde{\mu}) \Rightarrow \mu \sim_{b_i} \tilde{\mu}$ satisfies $T(\mu) = T(\tilde{\mu}) \Rightarrow (\mu_k) = (\tilde{\mu}_k)$.*

Corollary 6 (Primitive affine-determining classes)**.** *Assume $K \geq 3$, $b_i > 0$, and Assumptions 1 and 2. Suppose, for each battlefield k, that one of the sufficient conditions from Proposition 4 holds on the relevant class of marginal laws; in particular, this includes moment-determining polynomial families and the regularized Tullock kernel on compact support. Then for posterior-weighted opponent measures $\mu, \tilde{\mu}$ whose battlefield marginals lie in those classes,*

$$\mu \text{ and } \tilde{\mu} \text{ are payoff-equivalent on } \mathcal{X}_i(b_i) \iff \mu_k = \tilde{\mu}_k \text{ for every } k = 1, \ldots, K,$$

and the battlefield marginal tuple is the coarsest exact statistic preserving payoff-equivalence.

Proof. Proposition 4 provides primitive sufficient conditions under which each battlefield family is affine-determining on the relevant class of marginal laws. The conclusion then follows immediately from Theorems 7 and 8. □

Proof of Proposition 3. Set $\delta_k := \psi_k - \phi_k$. Then

$$\sum_{k=1}^{K} \delta_k(x_{ik}) = 0 \qquad \text{for all } \mathbf{x}_i \in \mathcal{X}_i(b_i).$$

We freeze all but three coordinates and read the residual identity on the resulting two-dimensional chart of the simplex.

Fix coordinates $4, \ldots, K$ at zero and set

$$C := \sum_{m=4}^{K} \delta_m(0),$$

with the convention $C = 0$ when $K = 3$. For every $t, s \geq 0$ with $t + s \leq b_i$,

$$\delta_1(t) + \delta_2(s) + \delta_3(b_i - t - s) + C = 0.$$

Subtracting the identity at $(t, s) = (0, 0)$ gives

$$f(t) + g(s) + q(t + s) = 0,$$

where

$$f(t) := \delta_1(t) - \delta_1(0), \quad g(s) := \delta_2(s) - \delta_2(0), \quad q(r) := \delta_3(b_i - r) - \delta_3(b_i).$$

Setting $s = 0$ yields $q(t) = -f(t)$, and setting $t = 0$ yields $g(s) = f(s)$. Hence

$$f(t+s) = f(t) + f(s) \qquad \text{for all } t, s \geq 0 \text{ with } t + s \leq b_i.$$

This restricted Cauchy equation extends to $\mathbb{R}$. Repeated use of it gives $f(x) = nf(x/n)$ for $0 \leq x \leq b_i$ and every integer $n \geq 1$. For $x \geq 0$ set $F(x) := nf(x/n)$ with any positive integer $n \geq x/b_i$, and $F(-x) := -F(x)$. Then F is well defined and $F = f$ on $[0, b_i]$. For $x, y \geq 0$ and n large, $F(x+y) = nf((x+y)/n) = nf(x/n) + nf(y/n) = F(x) + F(y)$, and oddness extends additivity to all of $\mathbb{R}$. On each $[0, nb_i]$, $F(x) = nf(x/n)$ is measurable, so F is a measurable additive function on $\mathbb{R}$ and hence linear by Fréchet's theorem [see, e.g., Kuczma, 2009]. Only measurability, inherited from the representing functions, is used. Hence $f(t) = ct$ on $[0, b_i]$ for some constant $c \in \mathbb{R}$. Therefore

$$\delta_1(t) = \delta_1(0) + ct, \qquad \delta_2(t) = \delta_2(0) + ct, \qquad \delta_3(t) = \big(\delta_3(b_i) - cb_i\big) + ct.$$

Repeating the same argument with the triple $(1, 2, k)$ in place of $(1, 2, 3)$ for each $k \in \{4, \ldots, K\}$, using allocations of the form $(t, s, 0, \ldots, 0, b_i - t - s, 0, \ldots, 0)$ with $b_i - t - s$ in coordinate k, shows that every δ_k has the form $\delta_k(t) = a_k + ct$, with the same constant c. Finally, evaluating the identity $\sum_{k=1}^{K} \delta_k(x_{ik}) = 0$ on any $\mathbf{x}_i \in \mathcal{X}_i(b_i)$ gives

$$\sum_{k=1}^{K} a_k + c \sum_{k=1}^{K} x_{ik} = \sum_{k=1}^{K} a_k + c\, b_i = 0,$$

which proves the stated constraint on the constants. □

Proof of Theorem 7. Equal marginals imply payoff-equivalence by Corollary 4. For the converse, assume that μ and $\tilde{\mu}$ are payoff-equivalent on $\mathcal{X}_i(b_i)$ and define

$$v_k^{\mu}(t) := \int_{\mathbb{R}_+} g_k(t, y)\, d\mu_k(y), \qquad v_k^{\tilde{\mu}}(t) := \int_{\mathbb{R}_+} g_k(t, y)\, d\tilde{\mu}_k(y).$$

Under Assumption 1, $F_{\mu}(\mathbf{x}_i) = \sum_k v_k^{\mu}(x_{ik})$ and $F_{\tilde{\mu}}(\mathbf{x}_i) = \sum_k v_k^{\tilde{\mu}}(x_{ik})$, so payoff-equivalence says that the two families represent the same criterion on $\mathcal{X}_i(b_i)$; each v_k is measurable in t because the kernels are jointly measurable and integrable. Proposition 3 therefore yields constants $a_1, \ldots, a_K, c \in \mathbb{R}$ such that

$$v_k^{\mu}(t) = v_k^{\tilde{\mu}}(t) + a_k + ct \qquad \text{for all } t \in [0, b_i],\ k = 1, \ldots, K.$$

The slope c is common across all battlefields; this is the simplex-level affine gauge of Proposition 3. Now fix a battlefield k. Because $\mu_k, \tilde{\mu}_k \in \mathfrak{N}_k$ and $\mathcal{G}_k(b_i)$ is affine-determining on $\mathfrak{N}_k$, the displayed identity forces

$$\mu_k = \tilde{\mu}_k \qquad \text{and} \qquad a_k = c = 0.$$

Since this argument applies for every battlefield, we obtain $\mu_k = \tilde{\mu}_k$ for all $k = 1, \ldots, K$. This proves the converse. □

Proof of Theorem 8. If $T(\mu) = T(\tilde{\mu})$, then the defining property of T gives payoff-equivalence of μ and $\tilde{\mu}$. By Theorem 7, payoff-equivalence is equivalent to equality of battlefield marginals. Thus every fiber of T lies inside a marginal-equality class. Conversely, direct integration under componentwise locality shows that the marginal tuple itself is an exact statistic. Hence the battlefield marginals are the coarsest exact statistic. □

Proposition 4 (Sufficient conditions for affine-determining families)**.** *The following battlefield kernel families are affine-determining on the indicated classes of marginal measures:* moment-determining polynomial families*, where* $g_k(t, y) = \sum_{j=0}^{J} a_j(t)\, y^j$ *with* $a_j \in C([0, b_i])$*, the quotient classes* $[a_1], \ldots, [a_J]$ *are linearly independent in* $C([0, b_i])/\operatorname{span}\{1, t\}$*, and equality of the first* J *moments determines measures within the class under consideration; and the* regularized Tullock contest success function*,* $g_k(t, y) = t/(t + y + \varepsilon)$ *for a fixed* $\varepsilon > 0$*, which on any compact support interval* $[0, M]$ *is affine-determining on the class of Borel probability measures supported on* $[0, M]$.

Proof of Proposition 4. For the polynomial family, let $\eta := \nu - \tilde{\nu}$ and $m_j := \int y^j \, d\eta(y)$. Since ν and $\tilde{\nu}$ are probability measures, $m_0 = 0$. If $\int g_k(t, y) \, d\eta(y) = a + ct$ for every t, then

$$\sum_{j=1}^{J} m_j a_j(t) \in \operatorname{span}\{1, t\}.$$

Linear independence of $[a_1], \ldots, [a_J]$ in the quotient space gives $m_1 = \cdots = m_J = 0$. The moment-determining assumption on the indicated measure class then gives $\nu = \tilde{\nu}$, after which the displayed affine identity forces $a = c = 0$. The bilinear kernel $g_k(t, y) = ty$ is excluded: it identifies only battlefield means and cannot remove the affine gauge.

Moreover, fix $\varepsilon > 0$ and $0 < M < \infty$, and consider the regularized Tullock contest success kernel

$$g_k(t, y) = \frac{t}{t + y + \varepsilon} \qquad \text{on } [0, b_i] \times [0, M].$$

Then g_k is affine-determining on the class of Borel probability measures supported on $[0, M]$.

For this kernel, let $\eta := \nu - \tilde{\nu}$ and suppose

$$\int_{[0,M]} \frac{t}{t + y + \varepsilon} \, d\nu(y) = \int_{[0,M]} \frac{t}{t + y + \varepsilon} \, d\tilde{\nu}(y) + a + ct \qquad \text{for all } t \in [0, b_i].$$

Evaluating at $t = 0$ gives $a = 0$. For $t > 0$, divide by t to obtain

$$\int_{[0,M]} \frac{1}{t + y + \varepsilon} \, d\eta(y) = c \qquad \text{for all } t \in (0, b_i].$$

The left-hand side is the Stieltjes transform of the finite signed measure η, hence a real-analytic function of t on $(0, \infty)$. Since it is constant on the interval $(0, b_i]$, it is constant on all of $(0, \infty)$. Letting $t \to \infty$ yields $c = 0$, because

$$\left| \int_{[0,M]} \frac{1}{t + y + \varepsilon} \, d\eta(y) \right| \le \frac{\|\eta\|_{\mathrm{TV}}}{t + \varepsilon} \to 0.$$

Thus the Stieltjes transform of η vanishes on $(0, \infty)$, and by uniqueness of the Stieltjes transform for finite Borel measures on a compact interval, $\eta = 0$. Hence $\nu = \tilde{\nu}$ and the affine correction is zero, proving the affine-determining property. □

Proposition 5 (Degeneracy of the two-battlefield case)**.** *Assume $K = 2$ and fix a realized budget b_i. Then every measurable interim criterion $F : \mathcal{X}_i(b_i) \to \mathbb{R}$ admits a measurable one-dimensional payoff-preserving representation. In particular, every deterministic slice $G_{\mathbf{y}}$ on $\mathcal{X}_i(b_i)$ is additively separable, but its additive decomposition need not be unique even up to an affine gauge.*

Proof. Every feasible allocation has the form $(t, b_i - t)$ with $t \in [0, b_i]$. Define

$$h(t) := F(t, b_i - t).$$

Then for every $(x_{i1}, x_{i2}) \in \mathcal{X}_i(b_i)$,

$$F(x_{i1}, x_{i2}) = h(x_{i1}) = h(x_{i1}) + 0(x_{i2}),$$

which is a measurable one-dimensional representation on the simplex. Applying the same argument to any deterministic slice $G_{\mathbf{y}}(\mathbf{x}_i) = u_i(\mathbf{x}_i, \mathbf{y})$ proves separability. Since $q(x_{i1}) - q(b_i - x_{i2}) = 0$ for every measurable q, the components admit non-affine gauge changes. □

Remark 6 (The threshold $K = 3$ is sharp for both directions)**.** At $K = 2$, first-order marginals already identify the opponent allocation law, while additive decompositions admit non-affine gauges. The $K = 3$ boundary concerns affine-gauge uniqueness and the first-order obstruction. At $K = 3$, the obstruction of Theorem 3 is already available. Degree-two mixed monomials are too low-order to witness failure on the three-battlefield simplex:

$$x_{i1} x_{i2} = -\frac{1}{2} x_{i1}^2 - \frac{1}{2} x_{i2}^2 + \frac{1}{2} (b_i - x_{i3})^2,$$

so they are additively separable on $\mathcal{X}_i(b_i)$. Cubic terms such as $x_{i1}^2 x_{i2}$ are not: on the chart $(z_1, z_2) \mapsto (z_1, z_2, b_i - z_1 - z_2)$, their mixed partial $\partial^2_{z_1 z_2}(z_1^2 z_2) = 2z_1$ fails to depend only on the residual budget. Hence $K = 3$ is the smallest number of battlefields at which the positive results are nondegenerate and first-order exactness can already fail.

D Exact-budget marginal hierarchy witnesses

Example 1 supplies the $K = 3$ example used in Theorem 3.

Proposition 6 (Majoritarian separation at $K = 3$)**.** *Let $\mu^{\pm}$ be as in Example 1, let the player's budget be $9/2$, and write $\mathbf{x} = (\frac{1}{2}, \frac{3}{2}, \frac{5}{2})$, $\mathbf{x}' = (\frac{1}{2}, \frac{5}{2}, \frac{3}{2})$, and $\bar{\mathbf{x}} = (\frac{3}{2}, \frac{3}{2}, \frac{3}{2})$.* (i) *For $u = \sum_a v_a \mathbf{1}\{x_a > y_a\}$, $F_{\mu^+} \equiv F_{\mu^-}$.* (ii) *For $u_{\mathrm{maj}} = \mathbf{1}\{|\{a : x_a > y_a\}| \geq 2\}$, $F_{\mu^-}(\mathbf{x}) = F_{\mu^+}(\mathbf{x}') = \frac{2}{3}$ and $F_{\mu^+}(\mathbf{x}) = F_{\mu^-}(\mathbf{x}') = 1$. Thus $\mathbf{x}$ is optimal under μ^+ but not under μ^-, $\mathbf{x}'$ the reverse, and $\bar{\mathbf{x}}$ is optimal under both.* (iii) *Any value at $\mathbf{x}$ reported from the first-order state errs by at least $\frac{1}{6}$ under one of the two laws.*

Proof. (i) $F_\mu(\mathbf{x}) = \sum_a v_a \mu(Y_a < x_a)$ depends on μ only through its battlefield marginals, which coincide. (ii) Against the atoms $(0,1,2), (1,2,0), (2,0,1)$ of μ^-, the allocation $\mathbf{x}$ wins $3, 1, 2$ battlefields, and against the atoms $(0,2,1), (2,1,0), (1,0,2)$ of μ^+ it wins $2, 2, 2$; for $\mathbf{x}'$ the counts are $2, 2, 2$ and $3, 2, 1$. Since $u_{\mathrm{maj}} \leq 1$, value 1 is optimal. Against any permutation of $(0,1,2)$, $\bar{\mathbf{x}}$ wins exactly the two battlefields with $y_a \leq 1$. All coordinates of $\mathbf{x}, \mathbf{x}', \bar{\mathbf{x}}$ are half-integers and all atoms are integers, so no comparison is tied. (iii) A value computed from the common first-order state is one number c under both laws, and $\max\{|c - \frac{2}{3}|, |c - 1|\} \geq \frac{1}{6}$. □

A binary $K = 4$ witness follows as the $r = 1$ case of the paired-bit construction below.

Proposition 7 (Paired-bit exact-budget hierarchy witnesses)**.** *For every $r \geq 1$ and $K \geq 2(r+1)$, there exist two probability measures on a single exact-budget opponent simplex $\mathcal{Y}_K(r+1)$ with identical r-order marginals but different induced criteria under an $(r+1)$-local payoff.*

Proof. Set $K = 2(r+1)$ battlefields and $B_{-i} = r+1$. Embed $r+1$ latent bits $Z_1, \ldots, Z_{r+1} \in \{0,1\}$ into K coordinates via

$$Y_{2j-1} := Z_j, \quad Y_{2j} := 1 - Z_j, \qquad j = 1, \ldots, r+1.$$

Every realization satisfies $\sum_{k=1}^K Y_k = (r+1) = B_{-i}$ (exact budget). Define

$$\mu^+ := \text{uniform over } \{\mathbf{y} : Z_1 \oplus \cdots \oplus Z_{r+1} = 0\},$$
$$\mu^- := \text{uniform over } \{\mathbf{y} : Z_1 \oplus \cdots \oplus Z_{r+1} = 1\},$$

each supported on $\mathcal{Y}_K(r+1)$ with 2^r atoms. *Marginal identity.* Let $A \subseteq [K]$ with $|A| \leq r$. The coordinates in A reveal values of at most r latent bits among $Z_1, \ldots, Z_{r+1}$, because each coordinate belongs to one paired encoding of one bit. Fix any assignment of the revealed bits. Under the even-parity law and under the odd-parity law, the unrevealed bits include at least one free bit. Among the completions of the unrevealed bits, exactly half have even parity and half have odd parity. Therefore the probability of the fixed revealed assignment is the same under μ^+ and μ^-. Since this holds for every assignment of coordinates in A, the A-marginals coincide. *Separation.* The $(r+1)$-local payoff

$$u(\mathbf{x}, \mathbf{y}) = x_1 \prod_{j=1}^{r+1} Z_j = x_1 \prod_{j=1}^{r+1} y_{2j-1}$$

separates the two laws. The product $\prod_{j=1}^{r+1} Z_j$ equals 1 only at the all-ones vector and equals 0 elsewhere. The all-ones vector has parity $(r+1) \bmod 2$. Therefore exactly one of μ^+ and μ^- assigns probability 2^{-r} to this event, and the other assigns probability 0. Hence

$$\left| \mathbb{E}_{\mu^+} \prod_{j=1}^{r+1} Z_j - \mathbb{E}_{\mu^-} \prod_{j=1}^{r+1} Z_j \right| = 2^{-r}.$$

Consequently,

$$|F_{\mu^+}(\mathbf{x}) - F_{\mu^-}(\mathbf{x})| = 2^{-r} x_1,$$

which is positive whenever $x_1 > 0$. This confirms the strict hierarchy for $K = 2(r+1)$.

For $K > 2(r+1)$, append $K - 2(r+1)$ additional coordinates fixed at zero: $\bar{\mu}^{\pm} := \mu^{\pm} \otimes \delta_0^{\otimes(K-2(r+1))}$. The added coordinates have identical degenerate marginals under both laws. Any marginal of order at most r involving added coordinates is therefore identical under $\bar{\mu}^+$ and $\bar{\mu}^-$. The separating payoff ignores the added coordinates: $u(\mathbf{x}, \mathbf{y}) = x_1 \prod_{j=1}^{r+1} y_{2j-1}$. Thus the same separation holds for every $K \geq 2(r+1)$. □

E Learning exact opponent states

The structural lower bounds use the typed ambiguity $\Delta_\rho(t_i)$ (Section 5.1).

The representation lower bounds below apply to reconstruction or decision rules whose input is the retained population state. They do not apply to a learner allowed to estimate the full joint posterior law from vector-valued samples.

E.1 Empirical marginal plug-in rule

Maximizers below are assumed to exist, as for continuous criteria on the compact feasible simplex or bounded componentwise kernels that are upper semicontinuous in the player's action. For the plug-in upper bounds, if a finite supremum is not attained, use that supremum as the optimal value and an ε'-maximizer of the empirical criterion; the upper bounds increase by ε'.

The allocation-sample rates below assume a single opponent type or payoff kernels independent of the opponent type. In the latter case, μ is the posterior mixture of allocation laws. For general type-dependent kernels, allocation samples alone need not identify the criterion (Appendix B.3); Proposition 9 instead uses type-labeled samples $(S^m, \mathbf{Y}^m)$. Fix a realized budget $b_i \geq 0$ and assume the opponent's relevant battlefield allocations lie in a compact interval $[0, M]$. A learner observes T i.i.d. samples $\mathbf{Y}^{(1)}, \ldots, \mathbf{Y}^{(T)}$ from the posterior-weighted opponent law $\mu \in \mathfrak{M}_i(b_i)$, then plays an allocation $\widehat{\mathbf{x}}_T \in \mathcal{X}_i(b_i)$. For a realized sample, define the sample-dependent one-shot regret

$$\mathrm{reg}_T(\mu; \widehat{\mathbf{x}}_T) := F_\mu(\mathbf{x}^*_\mu) - F_\mu(\widehat{\mathbf{x}}_T), \qquad \mathbf{x}^*_\mu \in \arg\max_{\mathcal{X}_i(b_i)} F_\mu.$$

Its expected regret is

$$\mathrm{Regret}_T(\mu; \widehat{\mathbf{x}}_T) := \mathbb{E}_{\mathbf{Y}^{(1:T)} \sim \mu^{\otimes T}}\big[\mathrm{reg}_T(\mu; \widehat{\mathbf{x}}_T)\big].$$

Definition 8 (Empirical marginal plug-in rule)**.** *The empirical marginal plug-in rule $\widehat{\mathbf{x}}^{\mathrm{marg}}_T$ is defined by estimating each battlefield marginal by its univariate empirical distribution $\hat{\mu}^{(T)}_k := T^{-1}\sum_{j=1}^T \delta_{Y^{(j)}_k}$, forming the empirical battlefield value functions $\widehat{v}^{(T)}_k(t) := \int g_k(t,y)\, d\hat{\mu}^{(T)}_k(y)$, and playing $\widehat{\mathbf{x}}^{\mathrm{marg}}_T \in \arg\max_{\mathbf{x}_i \in \mathcal{X}_i(b_i)} \sum_{k=1}^K \widehat{v}^{(T)}_k(x_{ik})$.*

E.2 Upper bound under additive-local payoffs

Theorem 9 (Regret bound for the empirical marginal plug-in rule)**.** *Suppose Assumption 1 holds and each kernel $g_k : [0, b_i] \times [0, M] \to \mathbb{R}$ is L_g-Lipschitz in its second argument, uniformly in the first. Then for every posterior law μ supported on $[0, M]^K$,*

$$\mathrm{Regret}_T\big(\mu; \widehat{\mathbf{x}}^{\mathrm{marg}}_T\big) \;\leq\; 2KL_gM/\sqrt{T}.$$

Proof. Under Assumption 1, the realized-type criterion is $F_\mu(\mathbf{x}_i) = \sum_k v_k(x_{ik})$ where $v_k(t) := \int g_k(t,y)\, d\mu_k(y)$. The Kantorovich–Rubinstein duality and the L_g-Lipschitz hypothesis give, for every $t \in [0, b_i]$,

$$\big|\widehat{v}^{(T)}_k(t) - v_k(t)\big| = \left|\int g_k(t,y)\, d(\hat{\mu}^{(T)}_k - \mu_k)(y)\right| \leq L_g\, W_1(\hat{\mu}^{(T)}_k, \mu_k),$$

so $\sup_t |\widehat{v}^{(T)}_k(t) - v_k(t)| \leq L_g\, W_1(\hat{\mu}^{(T)}_k, \mu_k)$. For a univariate Borel measure on $[0, M]$, let $H_k, \widehat{H}_k$ be its population and empirical CDFs. The identity $W_1(\hat{\mu}^{(T)}_k, \mu_k) = \int_0^M |\widehat{H}_k(y) - H_k(y)|\, dy$ and the bound $\mathbb{E}|\widehat{H}_k(y) - H_k(y)| \leq \sqrt{H_k(y)(1 - H_k(y))/T} \leq 1/(2\sqrt{T})$ give $\mathbb{E}[W_1(\hat{\mu}^{(T)}_k, \mu_k)] \leq M/\sqrt{T}$ and hence $\mathbb{E}[\sup_t |\widehat{v}^{(T)}_k(t) - v_k(t)|] \leq L_gM/\sqrt{T}$. The standard empirical-maximization decomposition gives

$$\mathrm{reg}_T \leq 2\sum_{k=1}^K \sup_t |\widehat{v}^{(T)}_k(t) - v_k(t)|.$$

Taking expectations and summing the K univariate bounds yields the claimed $2KL_gM/\sqrt{T}$. □

The K factor comes from the generic Wasserstein–Lipschitz bound. On the bilinear additive-local subclass it improves to $\sqrt{\log K}$ by maximal concentration.

Theorem 10 (High-probability version of Theorem 9)**.** *Under the hypotheses of Theorem 9, for every $\delta \in (0,1)$, with probability at least $1-\delta$ over the i.i.d. sample $\mathbf{Y}^{(1:T)} \sim \mu^{\otimes T}$,*

$$\mathrm{reg}_T(\mu; \widehat{\mathbf{x}}_T^{\mathrm{marg}}) \;\le\; 2KL_gM\Big(\frac{1}{\sqrt{T}} + \sqrt{\frac{\log(2K/\delta)}{2T}}\Big) \;=\; O\Big(KL_gM\sqrt{\log(K/\delta)/T}\Big).$$

Proof. Changing one sample changes $W_1(\hat{\mu}_k^{(T)}, \mu_k)$ by at most M/T, so McDiarmid's inequality plus $\mathbb{E}[W_1(\hat{\mu}_k^{(T)}, \mu_k)] \le M/\sqrt{T}$ gives, for each battlefield,

$$\mathbb{P}\Big[\, W_1(\hat{\mu}_k^{(T)}, \mu_k) \ge \frac{M}{\sqrt{T}} + M\sqrt{\frac{\log(2K/\delta)}{2T}}\,\Big] \;\le\; \frac{\delta}{K}.$$

A union bound controls the maximum over k. The Kantorovich–Rubinstein step from Theorem 9 gives $\sup_t |\widehat{v}_k^{(T)} - v_k| \le L_g W_1(\hat{\mu}_k^{(T)}, \mu_k)$; linearity of the criterion gives $\mathrm{reg}_T \le 2K \max_k \sup_t |\widehat{v}_k^{(T)} - v_k|$. □

E.3 Sharper bound on linearly modulated additive-local payoffs

The bilinear additive-local family is a special case of a broader class for which the K factor in Theorem 9 can be replaced by $\sqrt{\log K}$.

Definition 9 (Linearly modulated additive-local payoff)**.** *The primitive payoff u_i is* linearly modulated additive-local *if there exist bounded Borel measurable maps $\phi_k : [0,M] \to \mathbb{R}$, $k = 1, \dots, K$, such that $g_k(t,y) = t\,\phi_k(y)$ for all $t \in [0, b_i]$ and $y \in [0,M]$. Write $R_k := \sup_{[0,M]} \phi_k - \inf_{[0,M]} \phi_k$ for the range width and $R := \max_k R_k$.*

Theorem 11 (Near-tight bound on linearly modulated additive-local payoffs)**.** *Under Definition 9, the empirical marginal plug-in rule $\widehat{\mathbf{x}}_T^{\mathrm{marg}}$ satisfies*

$$\mathrm{Regret}_T(\mu; \widehat{\mathbf{x}}_T^{\mathrm{marg}}) \;\le\; b_i R\sqrt{\frac{2\log(2K)}{T}},$$

and for every $\delta \in (0,1)$, with probability at least $1-\delta$,

$$\mathrm{reg}_T(\mu; \widehat{\mathbf{x}}_T^{\mathrm{marg}}) \;\le\; b_i R\sqrt{\frac{2\log(2K/\delta)}{T}}.$$

Proof. Under Definition 9, the realized-type criterion reduces to a linear functional of the modulator means $\bar{\phi}_k(\mu) := \int \phi_k(y)\, d\mu_k(y)$: $F_\mu(\mathbf{x}_i) = \sum_k \bar{\phi}_k(\mu)\, x_{ik}$. The empirical marginal plug-in rule replaces $\bar{\phi}_k(\mu)$ by the univariate plug-in $\hat{\phi}_k := T^{-1}\sum_{j=1}^T \phi_k(Y_k^{(j)})$ and plays $\widehat{\mathbf{x}}_T^{\mathrm{marg}} \in \arg\max_{\mathcal{X}_i(b_i)} \widehat{F}_T$, where $\widehat{F}_T(\mathbf{x}_i) := \sum_k \hat{\phi}_k\, x_{ik}$.

For any residual $\boldsymbol{\alpha} \in \mathbb{R}^K$, the elementary identity $\sup_{\mathcal{X}_i(b_i)} |\sum_k \alpha_k x_{ik}| = b_i \max_k |\alpha_k|$ applied to $\alpha_k := \hat{\phi}_k - \bar{\phi}_k(\mu)$ gives

$$\sup_{\mathbf{x}_i \in \mathcal{X}_i(b_i)} \big|\widehat{F}_T(\mathbf{x}_i) - F_\mu(\mathbf{x}_i)\big| \;=\; b_i \max_k \big|\hat{\phi}_k - \bar{\phi}_k(\mu)\big|.$$

By the standard empirical-maximization decomposition,

$$\mathrm{reg}_T \le 2\sup_{\mathbf{x}_i} |\widehat{F}_T - F_\mu| \le 2b_i \max_k |\hat{\phi}_k - \bar{\phi}_k(\mu)|.$$

If $R = 0$, every modulator is constant and the empirical criterion is exact, so regret is zero. Assume $R > 0$ below. Each $\phi_k(Y_k^{(j)})$ takes values in an interval of width $\le R_k \le R$, so by Hoeffding's lemma $\hat{\phi}_k - \bar{\phi}_k(\mu)$ is sub-Gaussian with parameter $R^2/(4T)$. The maximal sub-Gaussian inequality gives $\mathbb{E}\max_k |\hat{\phi}_k - \bar{\phi}_k(\mu)| \le R\sqrt{\log(2K)/(2T)}$, which yields the expectation bound. For the high-probability statement, Hoeffding's inequality combined with a union bound over k gives $\mathbb{P}[\max_k |\hat{\phi}_k - \bar{\phi}_k(\mu)| \ge t] \le 2K\exp(-2Tt^2/R^2)$, and setting the right-hand side equal to δ produces $t = R\sqrt{\log(2K/\delta)/(2T)}$. □

Example 2 (Linearly modulated shapes)**.** Theorem 11 covers bilinear rewards $\phi_k(y) = c_k y$ ($R \le CM$ when $|c_k| \le C$), polynomial feedback $\phi_k(y) = c_k y^p$ ($R \le CM^p$ for integers $p \ge 1$ and $|c_k| \le C$), threshold and smoothed-threshold scores of width at most C, and truncated linear scores $\phi_k(y) = \max\{-\zeta, \min\{\zeta, c_k y\}\}$ for $\zeta > 0$ ($R \le 2\zeta$). The regularized Tullock kernel $g_k(t, y) = t/(t + y + \varepsilon)$, $\varepsilon > 0$, is nonlinear in t and is covered instead by the generic Wasserstein bounds.

Corollary 7 (Bilinear specialization)**.** *If $g_k(t, y) = c_k t y$ for $|c_k| \le C$ and $y \in [0, M]$ (Definition 9 with $\phi_k(y) = c_k y$, so $R \le CM$), then*

$$\begin{aligned}\mathrm{Regret}_T(\mu; \widehat{\mathbf{x}}_T^{\mathrm{marg}}) &\le C b_i M \sqrt{2\log(2K)/T},\\ \mathrm{reg}_T(\mu; \widehat{\mathbf{x}}_T^{\mathrm{marg}}) &\le C b_i M \sqrt{2\log(2K/\delta)/T} \quad \textit{with probability at least } 1-\delta.\end{aligned}$$

On the paired-bit exact-budget bilinear subclass $\mathfrak{M}_i^{\mathrm{pair}}(b_i)$, Theorem 11 matches the Fano lower bound of Theorem 13 up to universal constants.

E.4 Winner-take-all payoffs

Proposition 8 (DKW rate for battlefield-wise winner-take-all)**.** *Let $u_i(\mathbf{x}_i, \mathbf{y}) = \sum_{k=1}^K v_k \mathbf{1}\{x_{ik} > y_k\}$ and let $\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(T)}$ be i.i.d. from μ. Write $F_k^-(t) := \mu(Y_k < t)$ and $\widehat{F}_k^-(t) := T^{-1}\sum_j \mathbf{1}\{Y_k^{(j)} < t\}$, so that $U(\mathbf{x}_i) = \sum_k v_k F_k^-(x_{ik})$, and let $\widehat{U}(\mathbf{x}_i) := \sum_k v_k \widehat{F}_k^-(x_{ik})$. For every $\delta \in (0,1)$, with probability at least $1-\delta$,*

$$\sup_{\mathbf{x}_i \in \mathcal{X}_i(b_i)} |\widehat{U}(\mathbf{x}_i) - U(\mathbf{x}_i)| \le \Big(\sum_k |v_k|\Big)\sqrt{\frac{\log(2K/\delta)}{2T}},$$

and every maximizer of $\widehat{U}$ over $\mathcal{X}_i(b_i)$ has regret at most twice this bound. No continuity or Lipschitz condition on the kernel is needed.

Proof. Strict winning gives $\mathbb{E}\,\mathbf{1}\{x > Y_k\} = \mu(Y_k < x) = F_k^-(x)$, the left limit of the CDF F_k at x, and likewise for the empirical law. Left limits are limits of values at points $s \uparrow t$, so $\sup_t |\widehat{F}_k^-(t) - F_k^-(t)| \le \sup_t |\widehat{F}_k(t) - F_k(t)|$. The Dvoretzky–Kiefer–Wolfowitz inequality with Massart's constant [Dvoretzky et al., 1956, Massart, 1990] gives $\mathbb{P}[\sup_t |\widehat{F}_k(t) - F_k(t)| > \varepsilon] \le 2e^{-2T\varepsilon^2}$. A union bound over k with $\varepsilon = \sqrt{\log(2K/\delta)/(2T)}$ and $|\widehat{U} - U| \le \sum_k |v_k|\,|\widehat{F}_k^- - F_k^-|$ give the first claim. The regret claim is the empirical-maximization decomposition $U(\mathbf{x}^*) - U(\widehat{\mathbf{x}}) \le 2\sup|\widehat{U} - U|$; if $\widehat{U}$ does not attain its supremum, an ε'-maximizer satisfies the same bound plus ε'. □

E.5 Lower bound for non-exact reduced states

Theorem 12 (Constant regret floor for marginal-state decision rules)**.** *Let $\mathfrak{M}_i(b_i)$ be a posterior class on which every F_μ is α-strongly concave with a uniform constant $\alpha > 0$. Let $\mu^A, \mu^B \in \mathfrak{M}_i(b_i)$ have identical first-order marginals. For every model $R : \mathfrak{M}_i(b_i) \to \mathcal{M}$ satisfying $R(\mu^A) = R(\mu^B)$ and every allocation rule $\Lambda : \mathcal{M} \to \mathcal{X}_i(b_i)$,*

$$\sup_{\mu \in \mathfrak{M}_i(b_i)} \Big(F_\mu(\mathbf{x}_\mu^*) - F_\mu(\Lambda(R(\mu)))\Big) \;\ge\; \frac{\alpha}{8}\,\big\|\mathbf{x}_{\mu^A}^* - \mathbf{x}_{\mu^B}^*\big\|_2^2.$$

Hence any rule based only on the first-order marginal state obeys the bound for every such pair.

Proof. Since $R(\mu^A) = R(\mu^B)$, the rule's allocation $\Lambda(R(\cdot))$ is identical at μ^A and μ^B. By the triangle inequality, this common allocation is at distance at least $\frac{1}{2}\|\mathbf{x}_{\mu^A}^* - \mathbf{x}_{\mu^B}^*\|_2$ from at least one of the two maximizers. Strong concavity gives a regret of at least $\frac{\alpha}{8}\|\mathbf{x}_{\mu^A}^* - \mathbf{x}_{\mu^B}^*\|_2^2$ at that law. □

E.6 Information-theoretic minimax lower bound

The structural lower bound of Theorem 12 vanishes inside the exact-state regime, where $\Delta_\rho(t_i) = 0$. A Fano multiple-hypothesis argument shows that the $\sqrt{\log K/T}$ rate of Theorem 11 is then minimax-optimal up to constants on the bilinear subclass, for every learner, marginal-based or not.

Theorem 13 (Sharp minimax lower bound: bilinear payoffs on paired-bit exact-budget subclass)**.** *Fix $K = 2m \geq 4$, $b_i > 0$, $M > 0$, and the bilinear additive-local payoff $g_k(t, y) = ty$. Let $\mathfrak{M}_i^{\mathrm{pair}}(b_i)$ denote the class of laws supported on the paired-bit exact-budget set*

$$\mathcal{Y}_{\mathrm{pair}} := \big\{(MZ_1, M(1-Z_1), \ldots, MZ_m, M(1-Z_m)) : Z_j \in \{0,1\}\big\} \subset \mathcal{Y}_K(mM),$$

which is contained in the exact-budget simplex with budget mM. There are universal constants $c, C > 0$ such that for every $T \geq C \log m$ and every estimator $\widehat{\mathbf{x}}_T$,

$$\sup_{\mu \in \mathfrak{M}_i^{\mathrm{pair}}(b_i)} \mathbb{E}\big[F_\mu(\mathbf{x}_\mu^*) - F_\mu(\widehat{\mathbf{x}}_T)\big] \geq c\, b_i M \sqrt{\frac{\log K}{T}}.$$

Proof. Write $b = b_i$. For each $j \in [m]$, define $\mu^{(j)}$ on $\mathcal{Y}_{\mathrm{pair}}$ with independent $Z_1, \ldots, Z_m$, where $Z_j \sim \mathrm{Ber}(\frac{1}{2} + \gamma)$, $Z_\ell \sim \mathrm{Ber}(\frac{1}{2})$ for $\ell \neq j$, $Y_{2\ell-1} := MZ_\ell$, $Y_{2\ell} := M(1-Z_\ell)$. Every sample satisfies $\sum_k Y_k = mM$ (exact budget). Battlefield $2j-1$ is uniquely optimal. For $j \neq \ell$, the one-sample KL is at most $C_0\gamma^2$; tensorization gives $\mathrm{KL} \leq C_0 T\gamma^2$. Fano's inequality [Tsybakov, 2009] gives $\inf_{\widehat{j}} \frac{1}{m}\sum_j \mathbb{P}_j(\widehat{j} \neq j) \geq 1 - (C_0 T\gamma^2 + \log 2)/(\log m)$. Choose $\gamma = c_0\sqrt{\log m / T}$; for $T \geq C \log m$ and $m \geq m_0$, with m_0 an absolute constant (e.g., $m_0 = 16$ once $C_0 c_0^2 \leq \frac{1}{4}$), the bound gives universal constant error. For the finitely many $m < m_0$, Le Cam's two-point method on $\mu^{(1)}, \mu^{(2)}$ with $\gamma = c_0/\sqrt{T}$ gives the same conclusion, because $\log K$ is then bounded by an absolute constant. Decode $\widehat{j} \in \arg\max_{\ell \in [m]} \widehat{x}_{T,2\ell-1}$. If $\widehat{j} \neq j$, then $\widehat{x}_{T,2j-1} \leq b/2$, so regret is at least $\gamma M b/2$. Averaging over j yields $\sup_\mu \mathbb{E}[F_\mu(\mathbf{x}^*) - F_\mu(\widehat{\mathbf{x}}_T)] \geq c\, b_i M \sqrt{\log K / T}$. □

Theorem 14 (Minimax lower bound in terms of the range width)**.** *Fix $K = 2m \geq 4$, $b_i > 0$, $M > 0$, and $R > 0$. Consider the following linearly modulated additive-local payoff and opponent laws. For each pair $j \in [m]$, let*

$$Y_{2j-1} = MZ_j, \qquad Y_{2j} = M(1 - Z_j), \qquad Z_j \in \{0, 1\}.$$

Thus every realization lies on the exact-budget simplex

$$\sum_{k=1}^{2m} Y_k = mM.$$

Let the odd battlefields have modulators

$$\phi_{2j-1}(y) = \frac{R}{M} y,$$

and let the even battlefields have zero modulators,

$$\phi_{2j} \equiv 0.$$

Then there exist universal constants $c, C > 0$ such that, for every $T \geq C \log m$ and every estimator $\widehat{\mathbf{x}}_T$,

$$\sup_\mu \mathbb{E}\big[F_\mu(\mathbf{x}_\mu^*) - F_\mu(\widehat{\mathbf{x}}_T)\big] \geq c\, b_i R \sqrt{\frac{\log K}{T}},$$

where the supremum is over laws supported on the paired-bit set above.

Proof. Use the independent paired-bit laws $\mu^{(j)}$, $j \in [m]$, from the proof of Theorem 13. The observation laws and pairwise KL bounds are unchanged. Under $\mu^{(j)}$, the mean coefficient is $R(1/2 + \gamma)$ on battlefield $2j-1$, $R/2$ on every other odd battlefield, and zero on even battlefields. Thus $b_i \mathbf{e}_{2j-1}$ is uniquely optimal.

Decode $\widehat{j} \in \arg\max_{\ell \in [m]} \widehat{x}_{T,2\ell-1}$. If $\widehat{j} \neq j$, then $\widehat{x}_{T,2j-1} \leq b_i/2$, and

$$F_{\mu^{(j)}}(b_i \mathbf{e}_{2j-1}) - F_{\mu^{(j)}}(\widehat{\mathbf{x}}_T) \geq R\gamma(b_i - \widehat{x}_{T,2j-1}) \geq b_i R\gamma/2.$$

The same Fano argument for large m, and the same two-point argument for the remaining m, give a constant lower bound on the decoding error. Using the choices of γ from that proof yields $c\, b_i R\sqrt{\log K / T}$. □

Proof of the typed ambiguity formula in Section 5.1. Fix a realized type t_i and an action $\mathbf{x}_i$. For each positive-posterior opponent type s, define $d_s(\mathbf{x}_i) := \mathrm{dist}_\infty(h_{\mathbf{x}_i,s}, \mathcal{A}_{\rho,s})$. By Theorem 1, the maximum payoff gap between two type-s conditional laws with the same observable state is $2d_s(\mathbf{x}_i)$.

Upper bound. Take any two typed posterior laws $\eta, \tilde{\eta}$ with the same typed observable state. Type by type, $\left|\int h_{\mathbf{x}_i,s}\, d\eta_s - \int h_{\mathbf{x}_i,s}\, d\tilde{\eta}_s\right| \le 2d_s(\mathbf{x}_i)$. Multiplying by posterior weights and summing gives

$$\left|\sum_{s\in T^+_{-i}} F_{-i}(s \mid t_i)\left(\int h_{\mathbf{x}_i,s}\, d\eta_s - \int h_{\mathbf{x}_i,s}\, d\tilde{\eta}_s\right)\right| \le 2\sum_{s\in T^+_{-i}} F_{-i}(s \mid t_i) d_s(\mathbf{x}_i).$$

Taking the supremum over $\eta, \tilde{\eta}$ and $\mathbf{x}_i$ gives the upper bound.

Lower bound. Fix $\mathbf{x}_i$. For each positive-posterior s, Theorem 1 gives two type-s laws μ_s, ν_s with the same observable state and payoff gap at least $2d_s(\mathbf{x}_i) - \varepsilon$. Swap if necessary so that the difference is nonnegative at $\mathbf{x}_i$. Construct two typed posterior laws using μ_s and ν_s at each type s with the same posterior type weights. Their typed observable states are identical. Their criterion gap is at least $\sum_{s\in T^+_{-i}} F_{-i}(s \mid t_i)(2d_s(\mathbf{x}_i) - \varepsilon)$. Letting $\varepsilon \downarrow 0$ and taking the supremum over $\mathbf{x}_i$ completes the lower bound. □

E.7 Rates for r-local payoffs

Theorem 15 (Plug-in rate for r-local block-separable payoffs)**.** *Let u_i be r-local with $g_A = f_A\phi_A$, $\sup_{\mathbf{x}_i} |f_A(\mathbf{x}_{i,A})| \le C_{|A|}$, and ϕ_A taking values in an interval of width at most $R_{|A|}$. From T i.i.d. samples, the plug-in rule $\widehat{\mathbf{x}}_T^{(r)}$ that maximizes the empirical criterion built from the $N_{K,r} := \sum_{j=1}^{r}\binom{K}{j}$ scalar block means satisfies*

$$\mathrm{Regret}_T \le \sqrt{\frac{2\pi}{T}}\sum_{j=1}^{r}\binom{K}{j} C_j R_j = O(K^r/\sqrt{T})$$

for fixed r and bounded C_jR_j. For every $\delta \in (0,1)$, with probability at least $1-\delta$,

$$\mathrm{reg}_T \le \sqrt{\frac{2\log(2N_{K,r}/\delta)}{T}}\sum_{j=1}^{r}\binom{K}{j} C_j R_j.$$

Proof of Theorem 15. Let $\hat{\phi}_A^{(T)} := T^{-1}\sum_{m=1}^{T}\phi_A(\mathbf{Y}_A^{(m)})$ denote the scalar empirical estimate of $\bar{\phi}_A(\mu_A) := \mathbb{E}_\mu[\phi_A(\mathbf{Y}_A)]$ for each $A \subseteq [K]$ with $1 \le |A| \le r$. Define the population and empirical criteria

$$F_\mu(\mathbf{x}_i) := \sum_{A:\, 1\le|A|\le r} f_A(\mathbf{x}_{i,A})\,\bar{\phi}_A(\mu_A), \qquad \hat{F}_T(\mathbf{x}_i) := \sum_{A:\, 1\le|A|\le r} f_A(\mathbf{x}_{i,A})\,\hat{\phi}_A^{(T)}.$$

Since $\widehat{\mathbf{x}}_T^{(r)} \in \arg\max_{\mathbf{x}_i}\hat{F}_T(\mathbf{x}_i)$, the standard empirical-maximization decomposition gives

$$\begin{aligned} F_\mu(\mathbf{x}^*) - F_\mu(\widehat{\mathbf{x}}_T^{(r)}) &= \big[F_\mu(\mathbf{x}^*) - \hat{F}_T(\mathbf{x}^*)\big] + \big[\hat{F}_T(\mathbf{x}^*) - \hat{F}_T(\widehat{\mathbf{x}}_T^{(r)})\big] + \big[\hat{F}_T(\widehat{\mathbf{x}}_T^{(r)}) - F_\mu(\widehat{\mathbf{x}}_T^{(r)})\big] \\ &\le |F_\mu(\mathbf{x}^*) - \hat{F}_T(\mathbf{x}^*)| + 0 + |\hat{F}_T(\widehat{\mathbf{x}}_T^{(r)}) - F_\mu(\widehat{\mathbf{x}}_T^{(r)})|, \end{aligned}$$

using the fact that $\hat{F}_T(\mathbf{x}^*) \le \hat{F}_T(\widehat{\mathbf{x}}_T^{(r)})$.

For any fixed $\mathbf{x}_i$, the pointwise criterion error satisfies

$$|F_\mu(\mathbf{x}_i) - \hat{F}_T(\mathbf{x}_i)| \;\le \sum_{A:\, 1\le|A|\le r} |f_A(\mathbf{x}_{i,A})|\cdot|\hat{\phi}_A^{(T)} - \bar{\phi}_A(\mu_A)| \;\le \sum_{A:\, 1\le|A|\le r} C_{|A|}|\hat{\phi}_A^{(T)} - \bar{\phi}_A(\mu_A)|.$$

The right-hand side does not depend on $\mathbf{x}_i$ (the data-dependent errors $\hat{\phi}_A^{(T)} - \bar{\phi}_A$ are scalars), so this bound holds uniformly over $\mathbf{x}_i \in \mathcal{X}_i(b_i)$ without a covering argument. In particular it holds at $\mathbf{x}_i = \mathbf{x}^*$ and at the data-dependent $\mathbf{x}_i = \widehat{\mathbf{x}}_T^{(r)}$.

Blocks with zero range width are estimated exactly and may be omitted from the concentration argument. Taking expectations and using the Hoeffding L^1 bound for each remaining block: since $\phi_A(\mathbf{Y}_A^{(m)})$ is bounded in an interval of width $R_{|A|}$,

$$\mathbb{E}\big[|\hat{\phi}_A^{(T)} - \bar{\phi}_A(\mu_A)|\big] \;\le\; \int_0^\infty \Pr\big[|\hat{\phi}_A^{(T)} - \bar{\phi}_A| \ge \epsilon\big]\, d\epsilon \;\le\; \int_0^\infty 2\, e^{-2T\epsilon^2/R_{|A|}^2}\, d\epsilon \;=\; R_{|A|}\sqrt{\frac{\pi}{2T}}.$$

Summing and grouping by $|A| = j$:

$$\mathrm{Regret}_T \;\le\; 2 \sum_{A:\, 1\le|A|\le r} C_{|A|}\, R_{|A|}\sqrt{\frac{\pi}{2T}} \;=\; \sqrt{\frac{2\pi}{T}} \sum_{j=1}^{r} \binom{K}{j} C_j R_j.$$

Writing $N_{K,r} := \sum_{j=1}^{r} \binom{K}{j}$, the same argument and a union bound give, with probability at least $1-\delta$,

$$\mathrm{reg}_T(\mu; \widehat{\mathbf{x}}_T^{(r)}) \le \sqrt{\frac{2\log(2N_{K,r}/\delta)}{T}} \sum_{j=1}^{r} \binom{K}{j} C_j R_j.$$

For multilinear $f_A(\mathbf{x}_{i,A}) = \prod_{k\in A} x_{ik}$ with $\sum_k x_{ik} = b_i$, AM-GM gives $|\prod_{k\in A} x_{ik}| \le (b_i/j)^j$. Hence

$$\mathrm{Regret}_T \le \sqrt{\frac{2\pi}{T}} \sum_{j=1}^{r} \binom{K}{j} \left(\frac{b_i}{j}\right)^j R_j = O(K^r/\sqrt{T})$$

for fixed r, b_i, and bounded R_j. □

Remark 7 (Separation of statistical and computational complexity)**.** The proof shows that the $O(K^r/\sqrt{T})$ rate requires only scalar Hoeffding concentration on each of the $N_{K,r} = O(K^r)$ scalar statistics $\hat{\phi}_A^{(T)}$; no joint distribution estimation is needed. The rate depends on the number of retained scalar block statistics because the separability $g_A = f_A \cdot \phi_A$ makes the opponent's contribution to each local kernel a single real number. For multilinear $f_A(\mathbf{x}_{i,A}) = \prod_{k\in A} x_{ik}$, finding $\widehat{\mathbf{x}}_T^{(r)}$ amounts to degree-r polynomial optimization over a simplex, which can already be NP-hard for $r = 2$ when K varies [Motzkin and Straus, 1965]; statistical learnability of the state therefore need not make the induced optimization tractable.

Corollary 8 (Bandit-feedback regret for exact states)**.** *For $K \ge 2$, under linearly modulated payoffs $g_k(t, y) = t\phi_k(y)$ with an exact first-order opponent state, let $\bar{R}_\phi := \max_k \sup_{[0,M]} \phi_k - \min_k \inf_{[0,M]} \phi_k$ be the width of a common interval containing all modulator values. If $b_i\bar{R}_\phi > 0$, the reward $b_i\phi_k(Y_k)$ of every vertex $b_i\mathbf{e}_k$ lies in an interval of width $b_i\bar{R}_\phi$. Run EXP3 [Auer et al., 2002, Corollary 3.2] on the K vertices, with rewards affinely rescaled to $[0,1]$, exploration $\gamma = \min\{1, \sqrt{K\log K/((e-1)T)}\}$, and bandit feedback (only the realized reward of the played vertex is observed). Let $\mathfrak{R}_T$ be cumulative regret against the best fixed allocation. Then*

$$\mathbb{E}[\mathfrak{R}_T] \le 2\sqrt{e-1}\; b_i\bar{R}_\phi\, \sqrt{TK\log K},$$

i.e., per-round regret $O(b_i\bar{R}_\phi\sqrt{K\log K/T})$. The criterion is linear in $\mathbf{x}_i$, so the best fixed vertex is optimal on $\mathcal{X}_i(b_i)$, and the bound is the EXP3 guarantee scaled by the reward width. Unlike the full-information rate of Theorem 11, which depends only on the individual widths R_k, the bandit rate depends on the common width $\bar{R}_\phi \ge \max_k R_k$, because arms with different offsets must be compared from their realized rewards; it also has the usual $\sqrt{K}$ factor of K-action bandit problems. If $b_i\bar{R}_\phi = 0$, all vertices have the same constant reward and regret is zero, so rescaling is unnecessary. For $K = 1$, there is only one feasible allocation and regret is also zero.

F PPAD-hardness of equilibrium opponent modeling

Proof of Lemma 1. Let $x \in \Delta_K$ and define the dummy mass $m := \sum_{a=2}^{K-1} x_a$. If $m > 0$, move all dummy mass to coordinate 0: $x_0' := x_0 + m$, $x_1' := x_1$, $x_a' := 0$ for $a \ge 2$. The linear objective gains $+\tau m$ (each dummy coordinate scored $-\tau$; battlefield 0 scored 0). The quadratic regularizer changes by

$$-\tfrac{\lambda}{2}\big((x_0+m)^2 - x_0^2 - \textstyle\sum_{a\ge2} x_a^2\big) \ge -\tfrac{\lambda}{2}(2x_0m + m^2) \ge -\tfrac{3\lambda m}{2}.$$

Hence the net payoff gain is at least $(\tau - 3\lambda/2)m > 0$, so no maximizer uses dummy battlefields.

After eliminating dummies, set $p := x_1$ and $1 - p := x_0$. The reduced problem is

$$\max_{p\in[0,1]} \Big\{ pS_1 - \tfrac{\lambda}{2}p^2 - \tfrac{\lambda}{2}(1-p)^2 \Big\}.$$

The derivative is $S_1 + \lambda(1 - 2p)$, yielding unconstrained maximizer $p^* = \frac{1}{2} + \frac{S_1}{2\lambda}$. Clipping to $[0, 1]$ gives the stated $x_1^* = \mathrm{clip}(\frac{1}{2} + \frac{S_1}{2\lambda}, 0, 1)$.

For the approximation bound: the objective is λ-strongly concave on Δ_K. For any mixed response with mean $\bar{x}$ and expected payoff within β of the maximum, Jensen's inequality implies the value at $\bar{x}$ is also within β of the maximum; strong concavity then gives $\frac{\lambda}{2}\|\bar{x} - x^*\|_2^2 \le \beta$, hence $|\bar{x}_1 - x_1^*| \le \|\bar{x} - x^*\|_2 \le \sqrt{2\beta/\lambda}$. □

Full proof of Theorem 6. We analyze the baseline tolerance $\beta := \lambda\varepsilon_0^2/128$; every γ-BNE with $0 < \gamma \le \beta$ is also a β-BNE, so every output at target tolerance γ is the mean profile of a β-BNE. We give a polynomial-time many-one reduction from the problem of computing an ε_0-threshold equilibrium of a directed graph with bounded in-degree (PPAD-hard for every fixed $\varepsilon_0 < 1/6$, even when every in- and out-degree is at most two, by Deligkas et al., 2024, Theorem 4.23) to the regularized diagonal BNE computation problem (Definition 6).

Step 1. Source problem. The source instance is a directed graph $G = (V, E)$ with $|V| = n$, maximum in-degree at most two, and in-neighborhoods $N_v \subseteq V$. Set $D := 2$. A vector $z = (z_v)_{v\in V} \in [0, 1]^n$ is an ε_0-threshold equilibrium if for every $v \in V$:

$$\sum_{u\in N_v} z_u > \tfrac{1}{2} + \varepsilon_0 \;\Rightarrow\; z_v \le \varepsilon_0, \qquad \sum_{u\in N_v} z_u < \tfrac{1}{2} - \varepsilon_0 \;\Rightarrow\; z_v \ge 1 - \varepsilon_0.$$

Step 2. Exact-budget allocation instance construction. Set $T_1 = T_2 = V$ with $\Pr(t_1 = v, t_2 = u) = 1/n^2$ for all $v, u \in V$, so $F_{-i}(u \mid v) = 1/n$ for all u, v. Decode threshold variables as $z_v := p_{1,v,1}$, the mean allocation on battlefield 1 of player 1 type v in the output. Battlefields are indexed $\{0, 1, \ldots, K-1\}$ with $K \ge 3$; battlefields 0 and 1 are active gadget battlefields; battlefields $a \ge 2$ are dummy battlefields penalized by $\tau := 4\lambda$. Set

$$A := \frac{4\lambda}{\varepsilon_0}, \qquad \beta = \frac{\lambda\varepsilon_0^2}{32D^2}, \qquad e := \sqrt{\frac{2\beta}{\lambda}} = \frac{\varepsilon_0}{4D},$$

so that $De = \varepsilon_0/4$.

Step 3. Player 2 copy gadget. Set player-2 coefficients: $d^2_{v,0} := 0$, $d^2_{v,1} := -\lambda$, $d^2_{v,a} := -\tau$ for $a \ge 2$; $\theta^2_{v,u,1} := 2\lambda n\, \mathbf{1}_{u=v}$; all other $\theta^2_{v,u,a} := 0$. For player 2 type v, the interim score for battlefield a is $q^2_{v,a} = d^2_{v,a} + (1/n)\sum_u \theta^2_{v,u,a}\, p_{1,u,a}$. Computing each battlefield: $q^2_{v,a} = -\tau$ for $a \ge 2$; $q^2_{v,0} = 0$; $q^2_{v,1} = -\lambda + 2\lambda p_{1,v,1} = \lambda(2z_v - 1)$. The score advantage of battlefield 1 over battlefield 0 is $S_v^2 := \lambda(2z_v - 1)$. Since $\tau = 4\lambda > 3\lambda/2$, Lemma 1 applies with $S_1 := S_v^2$: the unique exact response is $x^*_{2,v,1} = \mathrm{clip}(\frac{1}{2} + \frac{S_v^2}{2\lambda}, 0, 1) = \mathrm{clip}(z_v, 0, 1) = z_v$ (since $z_v \in [0, 1]$). For any β-best response with mean $p_{2,v,1}$, the approximation bound gives $|p_{2,v,1} - z_v| \le e$.

Step 4. Player 1 threshold gadget. Set player-1 coefficients: $d^1_{v,0} := 0$, $d^1_{v,1} := A/2$, $d^1_{v,a} := -\tau$ for $a \ge 2$; $\theta^1_{v,u,1} := -An\, \mathbf{1}_{u\in N_v}$; all other $\theta^1_{v,u,a} := 0$. For player 1 type v, the interim score for battlefield 1 is

$$q^1_{v,1} = \frac{A}{2} + \frac{1}{n}\sum_{u\in V}(-An)\, \mathbf{1}_{u\in N_v}\, p_{2,u,1} = \frac{A}{2} - A\sum_{u\in N_v} p_{2,u,1}.$$

The score advantage of battlefield 1 over battlefield 0 ($q^1_{v,0} = 0$) is $S_v^1 := A\big(\frac{1}{2} - \sum_{u\in N_v} p_{2,u,1}\big)$. By Lemma 1, the exact response is $x^*_{1,v,1} = \mathrm{clip}(\frac{1}{2} + \frac{S_v^1}{2\lambda}, 0, 1)$, and in a β-BNE, $|z_v - x^*_{1,v,1}| \le e$.

Step 5. Dummy-battlefield elimination. For both players and every type, dummy battlefields $a \geq 2$ have interim score $-\tau = -4\lambda < -3\lambda/2$. By Lemma 1 (applied separately for each player-type), every exact best response places zero mass on dummy battlefields. The above analysis of the exact responses for player 1 and player 2 therefore holds.

Step 6. BNE decodes to threshold equilibrium. From Step 3: $|p_{2,u,1} - z_u| \leq e$ for every $u \in V$. Summing over N_v ($|N_v| \leq D$):

$$\Big|\sum_{u\in N_v} p_{2,u,1} - \sum_{u\in N_v} z_u\Big| \;\leq\; |N_v|\cdot e \;\leq\; De \;=\; \frac{\varepsilon_0}{4}.$$

Case $\sum_{u\in N_v} z_u > \frac{1}{2} + \varepsilon_0$. Then $\sum_{u\in N_v} p_{2,u,1} > \frac{1}{2} + \varepsilon_0 - \frac{\varepsilon_0}{4} = \frac{1}{2} + \frac{3\varepsilon_0}{4}$. The score advantage $S_v^1 < A\big(\frac{1}{2} - \frac{1}{2} - \frac{3\varepsilon_0}{4}\big) = -\frac{3A\varepsilon_0}{4} = -3\lambda$. Then $\frac{1}{2} + \frac{S_v^1}{2\lambda} < \frac{1}{2} - \frac{3}{2} < 0$, so $x^*_{1,v,1} = 0$, and $z_v = p_{1,v,1} \leq x^*_{1,v,1} + e = e = \varepsilon_0/(4D) \leq \varepsilon_0$.

Case $\sum_{u\in N_v} z_u < \frac{1}{2} - \varepsilon_0$. Then $\sum_{u\in N_v} p_{2,u,1} < \frac{1}{2} - \frac{3\varepsilon_0}{4}$. The score advantage $S_v^1 > 3\lambda$, so $\frac{1}{2} + \frac{S_v^1}{2\lambda} > \frac{1}{2} + \frac{3}{2} > 1$, giving $x^*_{1,v,1} = 1$, and $z_v \geq x^*_{1,v,1} - e = 1 - e \geq 1 - \varepsilon_0$.

In both cases z_v satisfies the threshold equilibrium condition. Hence $z = (z_v)_{v\in V}$ is an ε_0-threshold equilibrium of the source graph.

Step 7. First-order exactness throughout the reduction. For player i type t_i, any own action $\mathbf{x}_i \in \Delta_K$, and opponent type u, the opponent-dependent part of the payoff slice is

$$\mathbf{y} \mapsto \sum_{a=0}^{K-1} \theta^i_{t_i,u,a}\, x_{i,a}\, y_a.$$

This is a linear function of $\mathbf{y}$ with each coordinate y_a entering independently, so it lies in the closed coordinate-additive span $\mathcal{A}_u^{(1)}$. The remaining payoff terms are constant in $\mathbf{y}$, so the full slice also lies in this span. By Theorem 1 applied type by type (Corollary 2), the typed first-order opponent state is exact for every realized type throughout the reduction.

Step 8. Polynomial bit complexity. The constructed exact-budget allocation instance has n^2 type pairs with prior $1/n^2$ (bit-size $O(\log n)$), K battlefields (fixed constant), and player-specific coefficients $A = 4\lambda/\varepsilon_0$, $\tau = 4\lambda$, $\theta^2_{v,u,1} = 2\lambda n\, \mathbf{1}_{u=v}$, and $\theta^1_{v,u,1} = -An\, \mathbf{1}_{u\in N_v}$ (rational, bit-size $O(\log n + \mathrm{bits}(\lambda))$ for fixed ε_0). The baseline $\beta = \lambda\varepsilon_0^2/(32D^2)$ is rational with polynomial bits; the target tolerance $\gamma \leq \beta$ has polynomial bit length by assumption. The output requires $O(nK)$ mean allocations in $[0,1]$. The reduction is polynomial-time. □

F.1 Learning in the regularized class

Proposition 9 (Payoff estimation from type-labeled samples)**.** *Define interim scores* $q^i_{t_i,a} := d^i_{t_i,a} + \sum_s F_{-i}(s \mid t_i)\theta^i_{t_i,s,a} \int y_a\, \pi_{-i}(d\mathbf{y} \mid s)$ *and, from* T *i.i.d. type-labeled samples* $(S^m, \mathbf{Y}^m)$ *of the typed posterior law, estimators* $\widehat{q}^i_{t_i,a} := T^{-1}\sum_{m=1}^T Z^m_a$ *with* $Z^m_a := d^i_{t_i,a} + \theta^i_{t_i,S^m,a} Y^m_a$. *Since* $Y^m_a \in [0,1]$ *(unit simplex),* Z^m_a *lies in an interval of width at most* $R_{t_i} := \max_a\big(\max_s(\theta^i_{t_i,s,a})^+ + \max_s(\theta^i_{t_i,s,a})^-\big) \leq 2\max_a \|\theta^i_{t_i,\cdot,a}\|_\infty$. *With probability at least* $1-\delta$, $\sup_{\mathbf{x}_i\in\Delta_K} |\widehat{U}_i - U_i| \leq R_{t_i}\sqrt{\log(2K/\delta)/(2T)}$, *and the plug-in maximizer has regret at most* $2R_{t_i}\sqrt{\log(2K/\delta)/(2T)}$.

Proof. The criterion is $U_i(\mathbf{x}_i) = \sum_a q^i_{t_i,a} x_{i,a} - \frac{\lambda}{2}\|\mathbf{x}_i\|_2^2$, and $\widehat{U}_i$ replaces $q^i_{t_i,a}$ by $\widehat{q}^i_{t_i,a}$, an average of i.i.d. variables with mean $q^i_{t_i,a}$. Since $\mathbf{x}_i \in \Delta_K$, $\sup_{\mathbf{x}_i} |\widehat{U}_i - U_i| = \sup_{\mathbf{x}_i} |\sum_a (\widehat{q}^i_{t_i,a} - q^i_{t_i,a}) x_{i,a}| \leq \max_a |\widehat{q}^i_{t_i,a} - q^i_{t_i,a}|$. Hoeffding's inequality for variables in an interval of width R_{t_i} and a union bound over the K battlefields give the first claim; the empirical-maximization decomposition gives the second. □

For any fixed opponent profile in the reduction instances, each coordinate's type-pair coefficients have a common sign, and $R_{t_i} \leq An = 4\lambda n/\varepsilon_0$. Apply Proposition 9 with failure probability $\delta/(2n)$

at each player-type. Regret is at most $\beta = \lambda\varepsilon_0^2/(32D^2)$ simultaneously over all player-types if each posterior supplies

$$T \geq 32768D^4n^2\varepsilon_0^{-6}\log(4nK/\delta)$$

i.i.d. type-labeled samples. For fixed ε_0, D, K, this is polynomial in n; the factor λ cancels. Thus the hard instances have polynomial statistical sample complexity. The general bound above retains its dependence on R_{t_i}.

F.2 Fixed battlefield and type counts

Proposition 10 (Exact BNE for fixed battlefield and type counts)**.** *For an instance $\mathcal{I}$ of Definition 6 with bit-size $|\mathcal{I}|$, an exact BNE in pure strategies exists and can be computed in time $2^{K(|T_1|+|T_2|)}\,\mathrm{poly}(|\mathcal{I}|)$.*

Proof. For player i and type t, the interim objective is $\sum_a q_{i,t,a}x_{i,t,a} - \frac{\lambda}{2}\|\mathbf{x}_{i,t}\|_2^2$, with interim scores $q_{i,t,a} := d^i_{t,a} + \sum_s F_{-i}(s \mid t)\theta^i_{t,s,a}x_{-i,s,a}$ that are affine in the opponent's allocations. The best-response map is single-valued by strong concavity and continuous by Berge's maximum theorem on the product of the types' simplices, so Brouwer's theorem gives a pure-strategy fixed point, which is an exact BNE. Because the constraints are linear and the objective is concave, $\mathbf{x}_{i,t}$ is a best response iff there are $\nu_{i,t} \in \mathbb{R}$ and $\sigma_{i,t,a} \geq 0$ with

$$q_{i,t,a} - \lambda x_{i,t,a} - \nu_{i,t} + \sigma_{i,t,a} = 0, \qquad x_{i,t,a} \geq 0, \qquad \sigma_{i,t,a}x_{i,t,a} = 0, \qquad \textstyle\sum_a x_{i,t,a} = 1,$$

a mixed complementarity system with K complementarity pairs per player-type. Fix, for every player-type, a set $S_{i,t}$ of battlefields, and impose $\sigma_{i,t,a} = 0$ for $a \in S_{i,t}$ and $x_{i,t,a} = 0$ for $a \notin S_{i,t}$. The remaining conditions are linear equalities and inequalities in $(\mathbf{x}, \boldsymbol{\nu}, \boldsymbol{\sigma})$ with rational coefficients of bit-size $\mathrm{poly}(|\mathcal{I}|)$, so linear programming decides feasibility and returns a rational solution in time $\mathrm{poly}(|\mathcal{I}|)$. A BNE exists, so some branch is feasible, and every feasible point of a branch satisfies all KKT systems simultaneously and is therefore an exact BNE. There are $2^{K(|T_1|+|T_2|)}$ branches. □

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction state the paper's contributions: an exact ambiguity identity and span criterion for observable opponent states, the first-order and higher-order characterization in exact-budget Blotto, the separation of representation error from estimation error for state-based rules, and PPAD-hardness of approximate BNE in a regularized diagonal exact-budget allocation subclass. Table 1 records the structure each result uses. These claims match the theorems proved in the body and appendix; no empirical claim is made.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The conclusion has a Limitations paragraph, and the scope of the hardness result is stated in Remark 3. The limitations stated there are: (i) the ambiguity identity and its necessity direction assume continuous payoff slices on compact action spaces, and discontinuous payoffs are covered through the upper bound and direct integration only; (ii) the typed results use finite opponent type sets T_{-i}, and estimation with type-dependent kernels assumes type-labeled samples; (iii) argmax-only preservation is not characterized; (iv) the setting is a one-shot interim decision problem, not a sequential or nonstationary one; (v) the PPAD lower bound applies to the regularized diagonal subclass, not to winner-take-all, Tullock, or other contest-success Blotto payoffs. Two further restrictions are stated with the corresponding results: (vi) our higher-order hierarchy construction uses $K \geq 2(r+1)$ battlefields (Theorem 4); and (vii) the marginal identification and minimality results require affine-determining battlefield kernel families (Appendix C), an assumption beyond those needed for the main exactness theorem.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be

used reliably to provide closed captions for online lectures because it fails to handle technical jargon.

- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: Every theorem, proposition, and lemma states its assumptions explicitly. The main text gives short proofs or proof sketches and points to the appendices for the longer arguments. Complete proofs appear in the appendices: Theorem 1, the levels of exactness (Corollary 1), and the example of a state that discards the type label in Appendix B; the first-order representation results in Appendix B.4–C; the hierarchy witnesses and the majoritarian example in Appendix D; the learning rates (including the winner-take-all rate, Proposition 8), the marginal-state regret floor, and minimax lower bounds in Appendix E; and the PPAD reduction, including Lemma 1, Theorem 6, and the fixed-size enumeration (Proposition 10), in Appendix F. The criterion floor of Theorem 5 is proved in Section 5.1.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [N/A]

   Justification: The paper is a theory paper and does not include experimental results that support the main claims.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may

be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.

- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [N/A]

   Justification: The paper does not include experiments requiring data or code for reproducing the main claims.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [N/A]

   Justification: The paper does not include experiments, training, or test settings.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [N/A]

   Justification: The paper does not include experiments or empirical performance claims.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
   - If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [N/A]

   Justification: The paper does not include experiments requiring compute-resource reporting.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: This paper studies abstract mathematical game theory. It involves no human subjects, personal data collection, or deployed system, and it conforms with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

    Justification: This paper develops mathematical theory for resource-allocation games and is not tied to a specific application or deployment.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The paper releases no models, scraped datasets, or other high-risk assets, so no misuse-specific safeguards are required.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [N/A]

    Justification: This paper uses no external datasets, software packages, or third-party assets beyond standard mathematical references, which are cited in the bibliography.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not introduce new datasets, models, benchmarks, or software assets as part of its main contribution.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: This paper involves no crowdsourcing and no research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: This paper involves no human subjects research and therefore requires no IRB approval.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: LLMs were not used as a component of the research methodology.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.