Computational Analyses of the Electoral College: Campaigning Is Hard But Approximately Manageable

AAAI'21

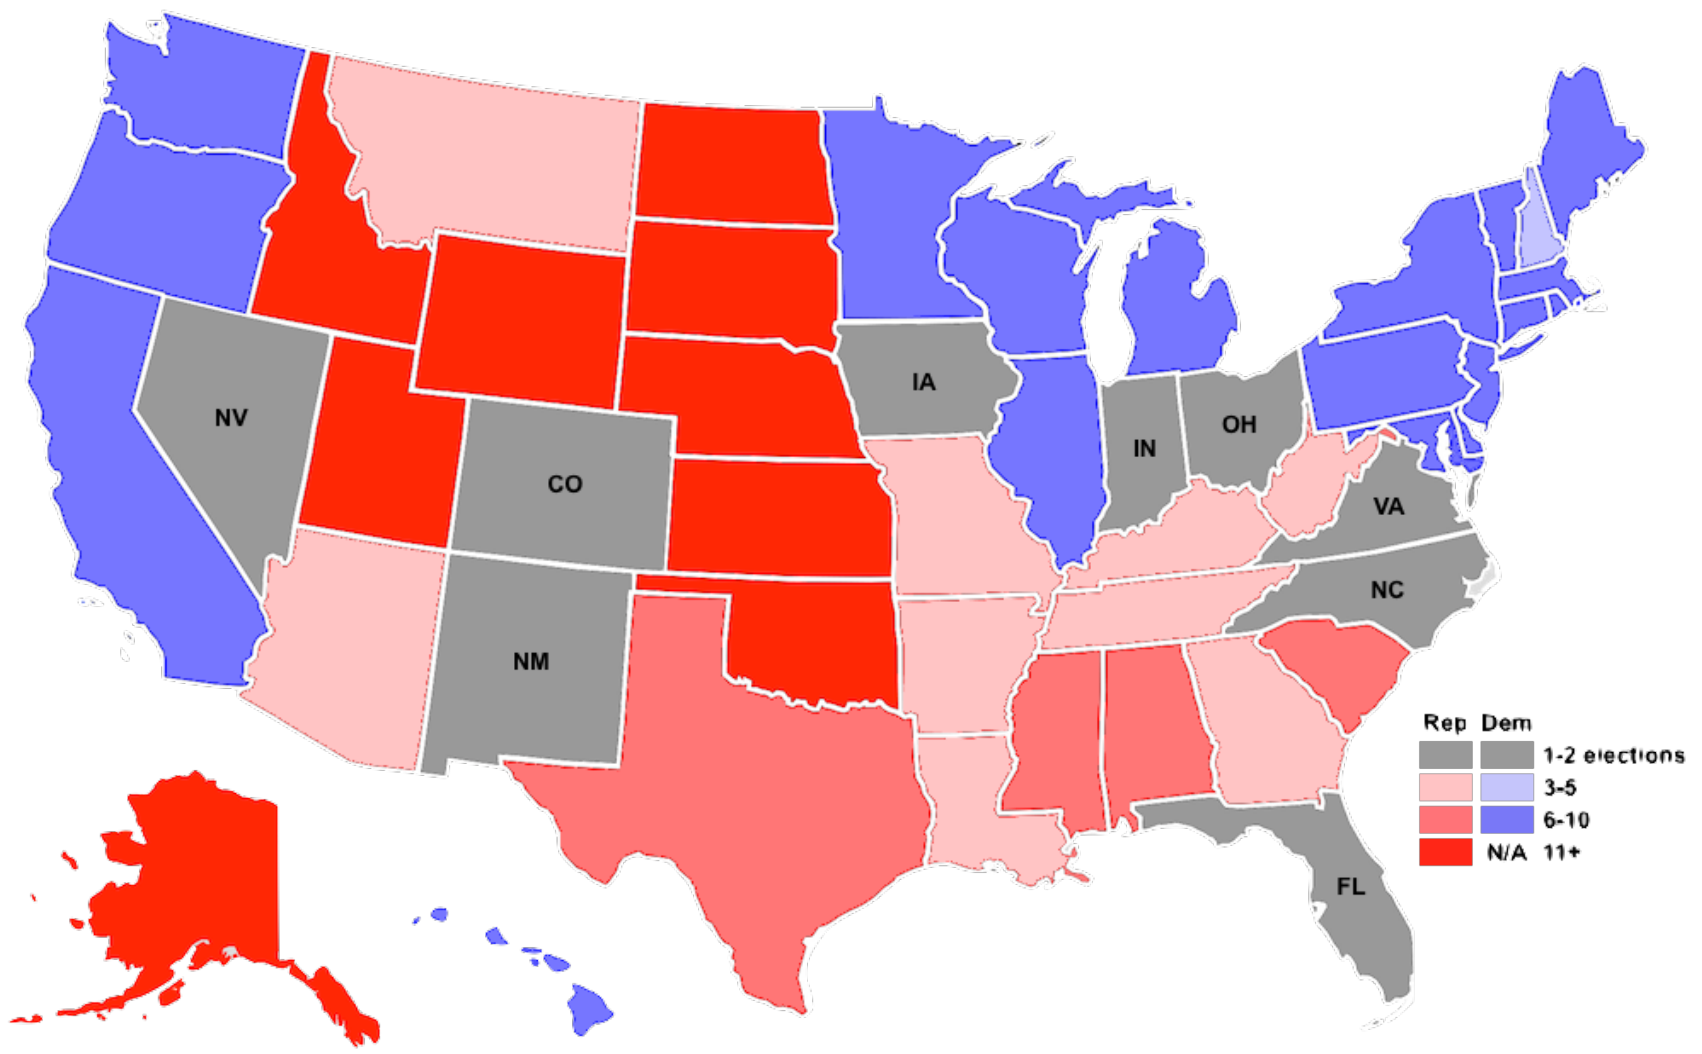

Sina Dehghani[1], Hamed Saleh[2], Saeed Seddighin[3], Shanghua Teng[4]

1 Institute for Research in Fundamental Sciences
2 University of Maryland
3 TTIC
4 University of Southern California

Colonel Blotto Game

Colonel Blotto Game

- Two colonels A and B are playing a game.
- Colonels A and B have m and n troops respectively.
- There are k battlefields.
- Colonels distribute their troops simultaneously across battlefields.
- The payoff of each battlefield is decided by winner-take-all policy.

Colonel Blotto Game

- Colonels A and B have m and n troops respectively.
- There are k battlefields.
- Pure strategies of each player:
    - A k-partitioning of the available troops.

Colonel Blotto Game

- Colonels A and B have m and n troops respectively.
- There are k battlefields.
- Randomized (mixed) strategies:
  - A probability distribution vector $\mathbf{X}$ over all feasible pure strategies.

Colonel Blotto Game

- Colonels A and B have m and n troops respectively.
- There are k battlefields.
- Constant-sum game:
  - The total payoff of both colonels is always constant (at each battlefield)
  - Maxmin strategies $\equiv$ Minmax strategies $\equiv$ Nash equilibria

Colonel Blotto Game

- Colonels A and B have m and n troops respectively.
- There are k battlefields.
- Applications:
    - Political Campaigns: U.S. presidential election
    - Marketing Campaigns: Apple vs Samsung
    - Sport Competitions

Colonel Blotto Game

- Colonels A and B have m and n troops respectively.
- There are k battlefields.
- Introduced by Borel and Ville (1921).
- Many attempts to solve the problem:
    - Continuous resources Roberson(2006).
    - Special cases Hart(2008).
- First polynomial solution Ahmadinejad et. al (2016).

Colonel Blotto Game

- Colonels A and B have m and n troops respectively.
- There are k battlefields.
- First polynomial solution Ahmadinejad et. al (2016).
- Linear Program to model the problem.
- Exponential number of variables and constraints.
- Ellipsoid method.

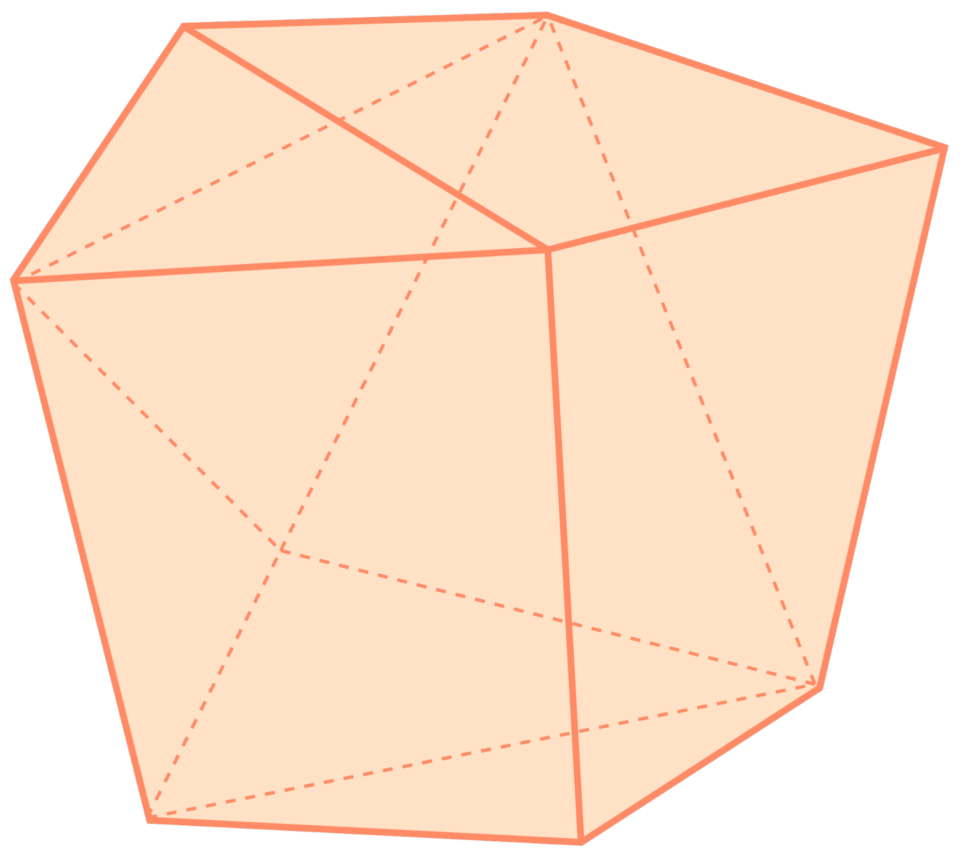

Colonel Blotto Game

- Colonels A and B have m and n troops respectively.
- There are k battlefields.
- First polynomial solution Ahmadinejad et. al (2016).
- Key idea: Reduce finding a maxmin strategy to finding a best response strategy.

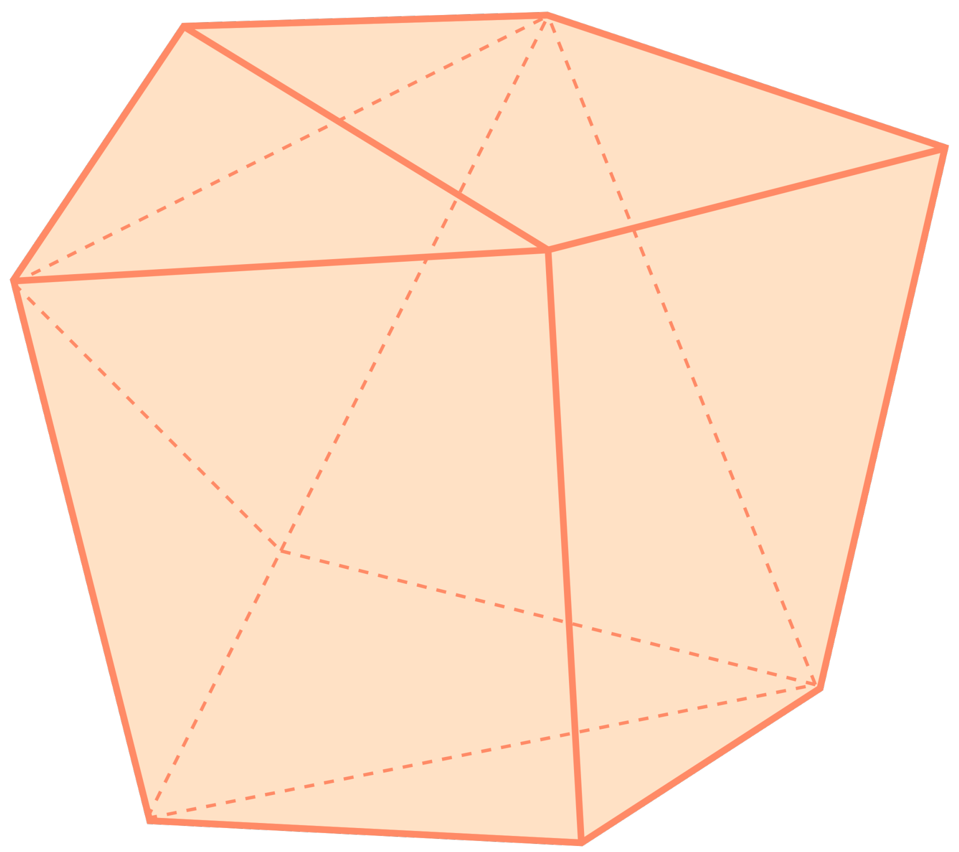

Colonel Blotto Game

- Colonels A and B have m and n troops respectively.
- There are k battlefields.
- Limitations of the original setting:
    - Troops are homogenous w.r.t. different battlegrounds.
        - Although we can assign weights $(\mu_1, \mu_2, \ldots, \mu_k)$ to battlefields, and it doesn't change the set of pure strategies.

Colonel Blotto Game

- Colonels A and B have m and n troops respectively.
- There are k battlefields.
- Limitations of the original setting:
    - All troops have the same strength.
    - Troops are homogenous w.r.t. different battlegrounds.
    - The payoff of each battleground is determined by winner-take-all policy.

Colonel Blotto Game

With Multifaceted Resources

Colonel Blotto Game

With Multifaceted Resources

- Colonels A and B have m and n troops respectively, and there are k battlefields.
- A $k \times (m+n)$ matrix W is given, where $w_{b,i}$ shows the strength of the i'th troop in battlefield b.
    - The strength of troops is additive: For colonel A, the total strength of the subset of troops $S \subseteq [m]$ assigned to battlefield b is equal to:

$$\sum_{i \in S} w_{b,i}$$

Colonel Blotto Game

With Multifaceted Resources

- Colonels A and B have m and n troops respectively, and there are k battlefields.
- A $k \times (m+n)$ matrix W is given.
- Two sets of utility functions $\{\mu_1^A, \mu_2^A, \ldots, \mu_k^A\}$ and $\{\mu_1^B, \mu_2^B, \ldots, \mu_k^B\}$ which determine the payoff in a battlefield based on the total strength of troops.
    - The utility functions are constant-sum, monotone, and non-negative.
    - The domain of utility functions is $\{0,1,\ldots,\max_f\}^2$, where $\max_f$ is an upper-bound on the total strength of the troops over all battlefields.

Colonel Blotto Game

With Multifaceted Resources

- Denote a pure strategy, which again is a k-partitioning of the available troops, by a vector X where X_b specifies the set of troops assigned to battlefield b.
- Define Y similarly for player B.
- The total payoff of each player for given pure strategies X and Y equals to:

$$\begin{cases} \mu^A(X,Y) = \sum_{b=1}^{k} \mu_b^A(w_b(X_b), w_b(Y_b)) \\ \mu^B(X,Y) = \sum_{b=1}^{k} \mu_b^B(w_b(X_b), w_b(Y_b)) \end{cases}$$

Colonel Blotto Game

With Multifaceted Resources

- Denote a mixed strategy for player A and B by $\mathbf{X}$ and $\mathbf{Y}$ respectively.
- The total payoff of each player for given mixed strategies $\mathbf{X}$ and $\mathbf{Y}$ equals to:

$$\begin{cases} \mu^A(\mathbf{X}, \mathbf{Y}) = \mathbb{E}_{X \sim \mathbf{X}, Y \sim \mathbf{Y}}[\mu^A(X, Y)] \\ \mu^B(\mathbf{X}, \mathbf{Y}) = \mathbb{E}_{X \sim \mathbf{X}, Y \sim \mathbf{Y}}[\mu^B(X, Y)] \end{cases}$$

Applications

U.S. Presidential Election

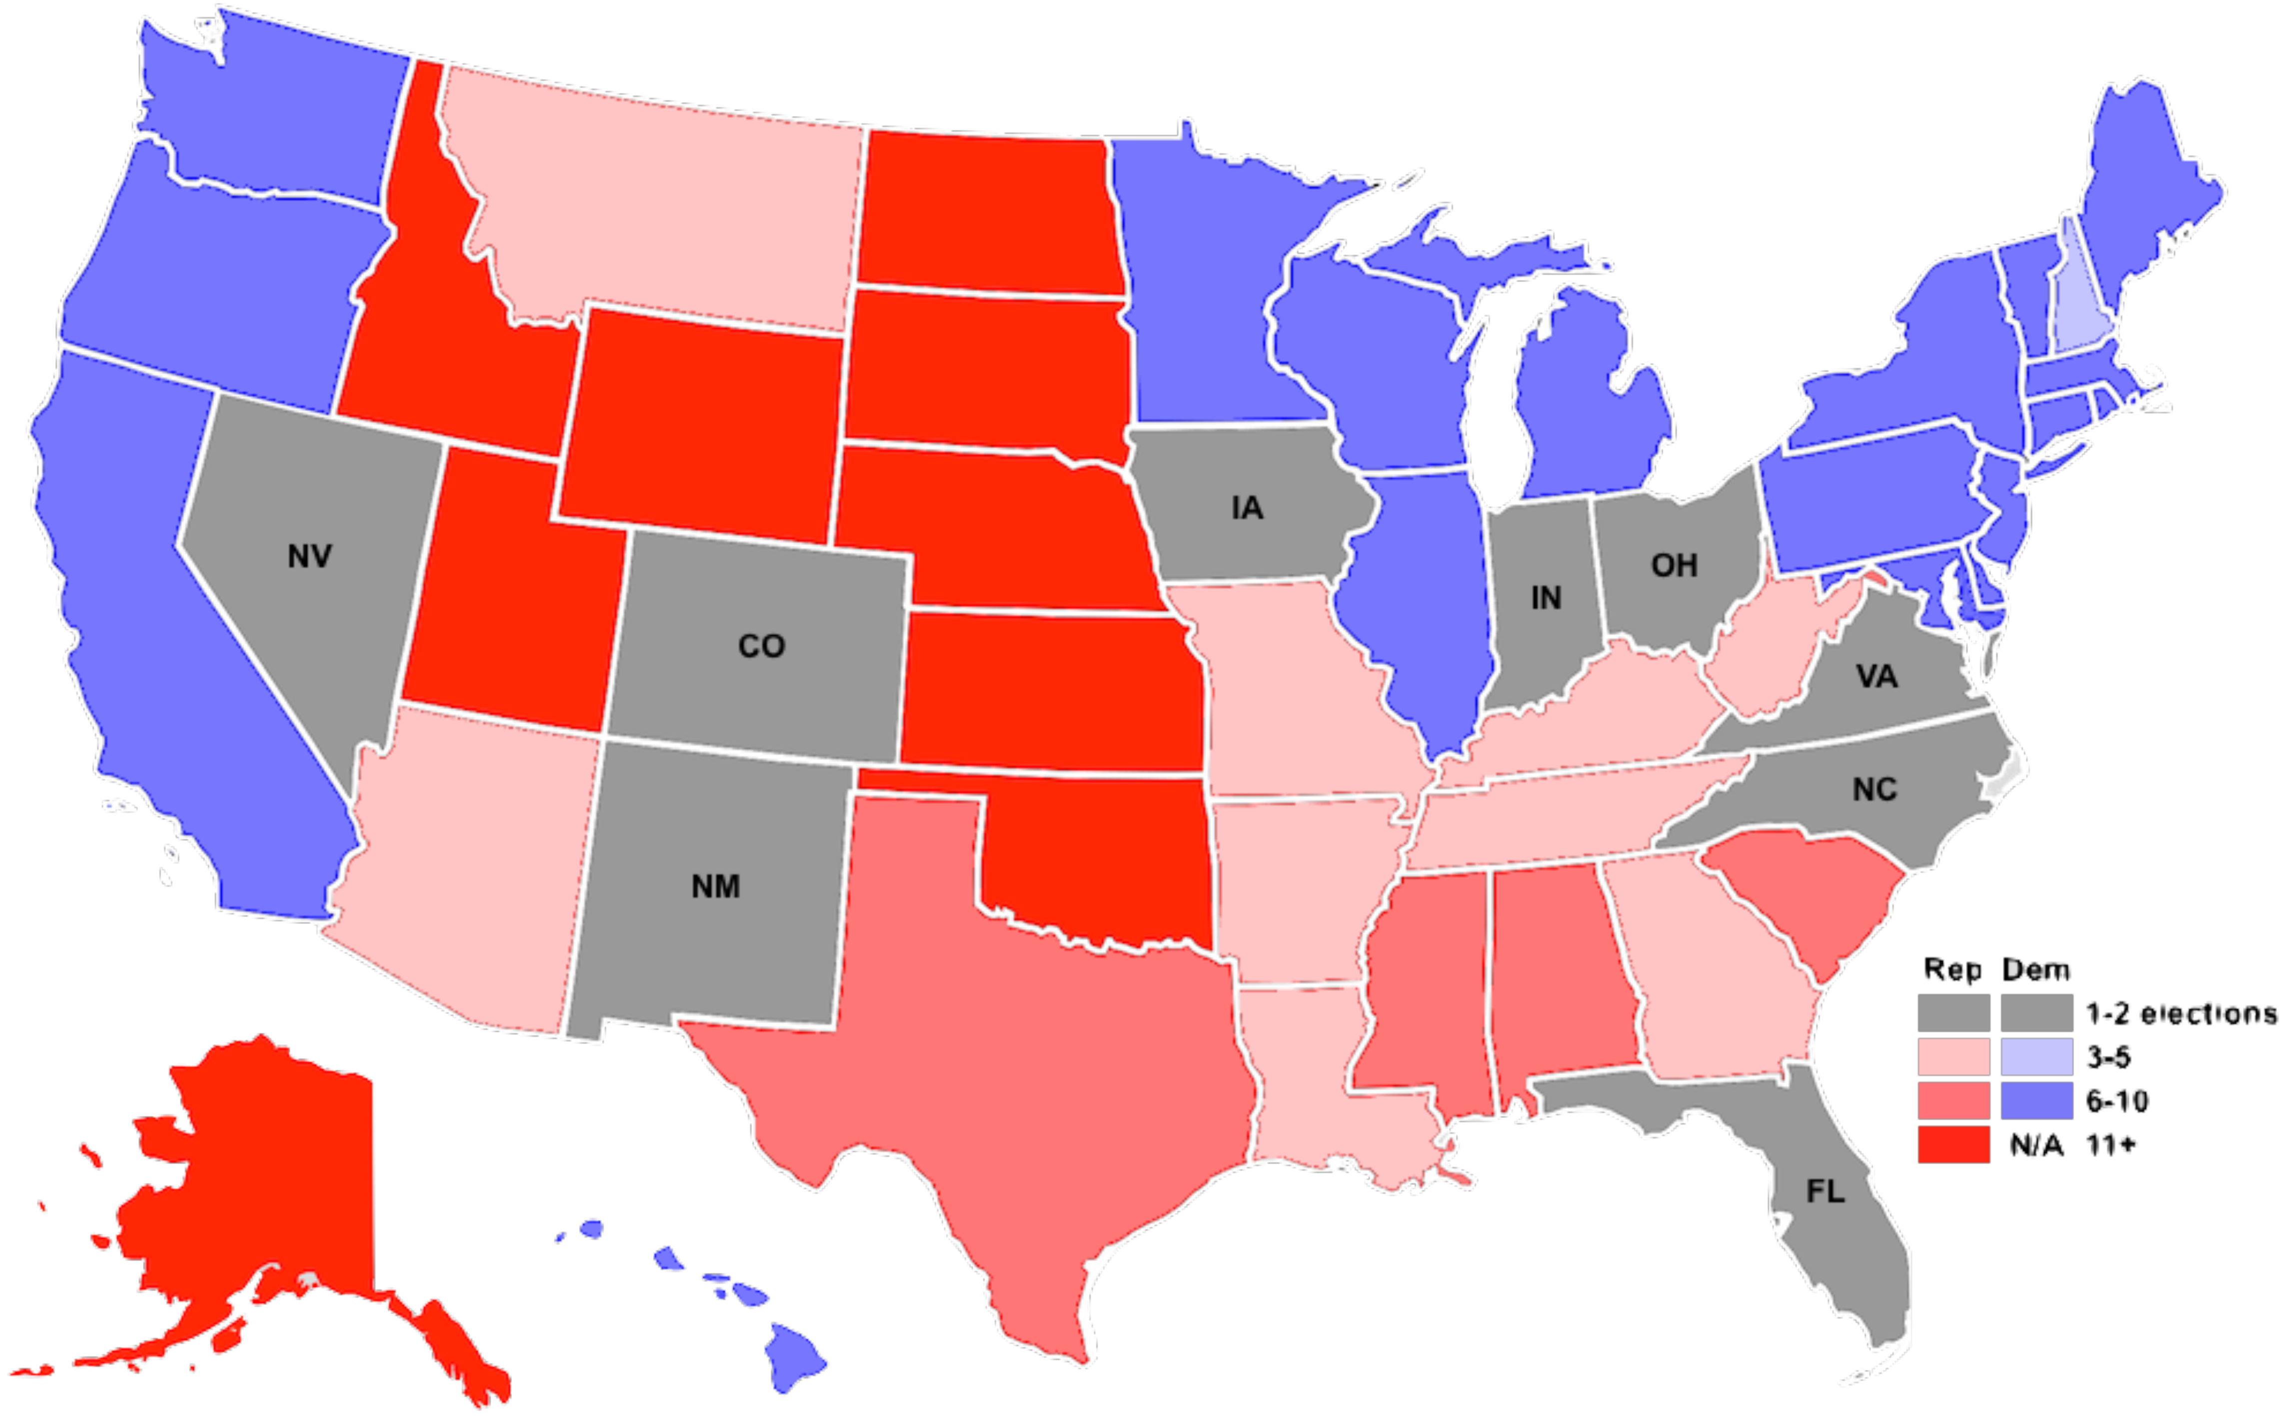

U.S. Presidential Election

- Swing states
- Maine and Nebraska
- Troops may include the following:
    - Money
    - Candidate’s time
    - On-the-ground staff
    - Campaign managers

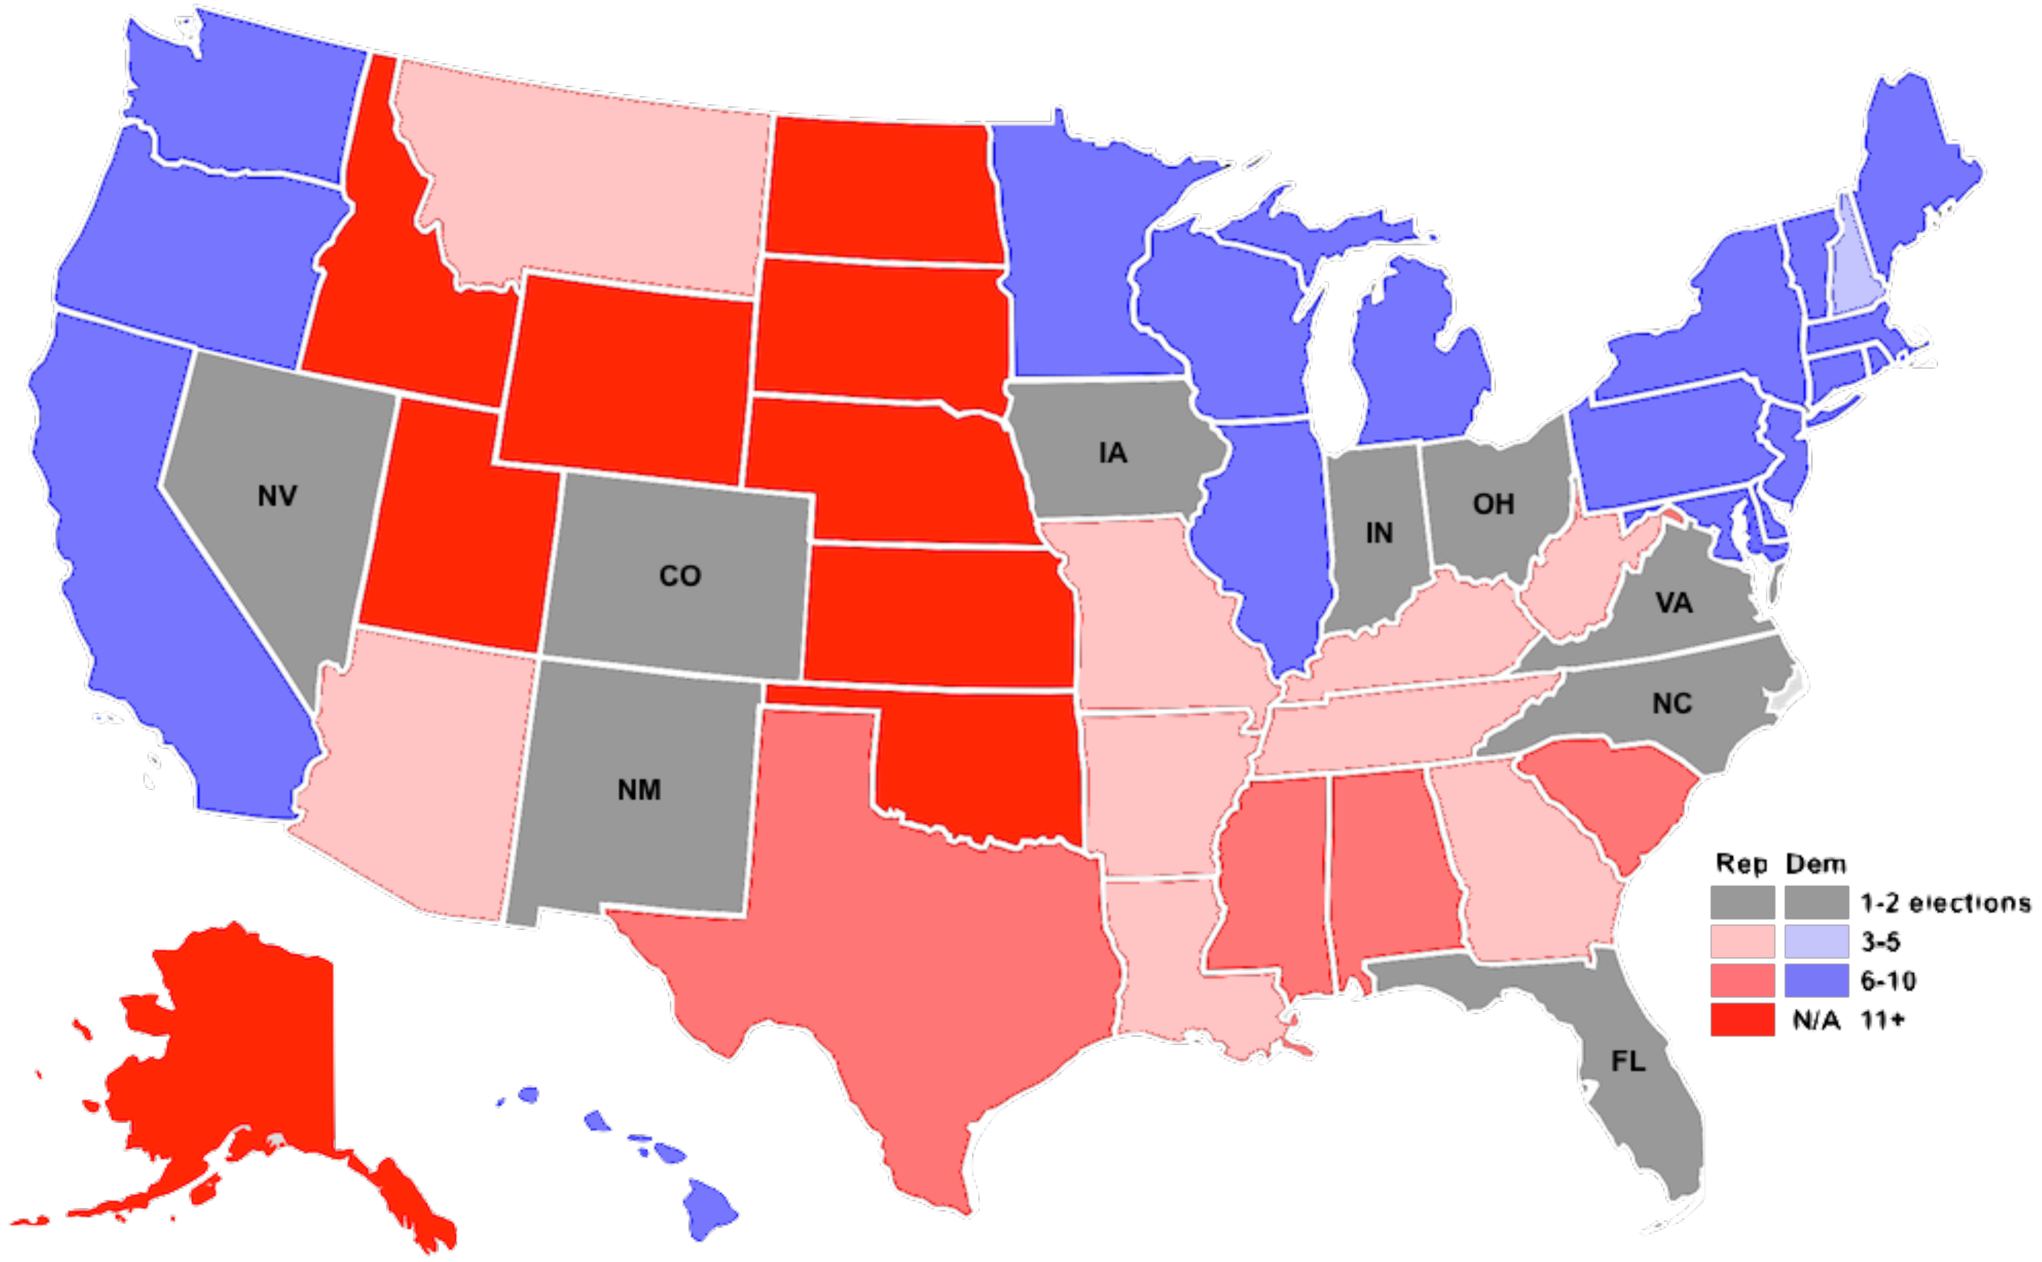

Tech companies competition

- Battlegrounds
    - Smartphone
    - Tablet
    - Laptop

Approximation Hardness

Hardness Result

- It is hard to approximate the best response strategy within $\sqrt{n}$ factor.
- Reduction from the Welfare Maximization for Single-minded Bidders problem.
- The approximation hardness of this problem is known by a reduction from Set Packing. Lehman et. al (2002) Sandholm (1999)

Welfare Maximization for Single-minded Bidders

- Allocation of a set of n indivisible items among m bidders.
- Each bidder i has a subset T_i of items which values $v_i(T_i)$.
- For a subset T', $v_i(T')$ equals:
    - $v_i(T_i)$ if $T_i \subseteq T'$
    - 0 otherwise
- Find an allocation which maximizes the total utility of bidders.

Welfare Maximization for Single-minded Bidders

- An example of reduction to an instance of Colonel Blotto.

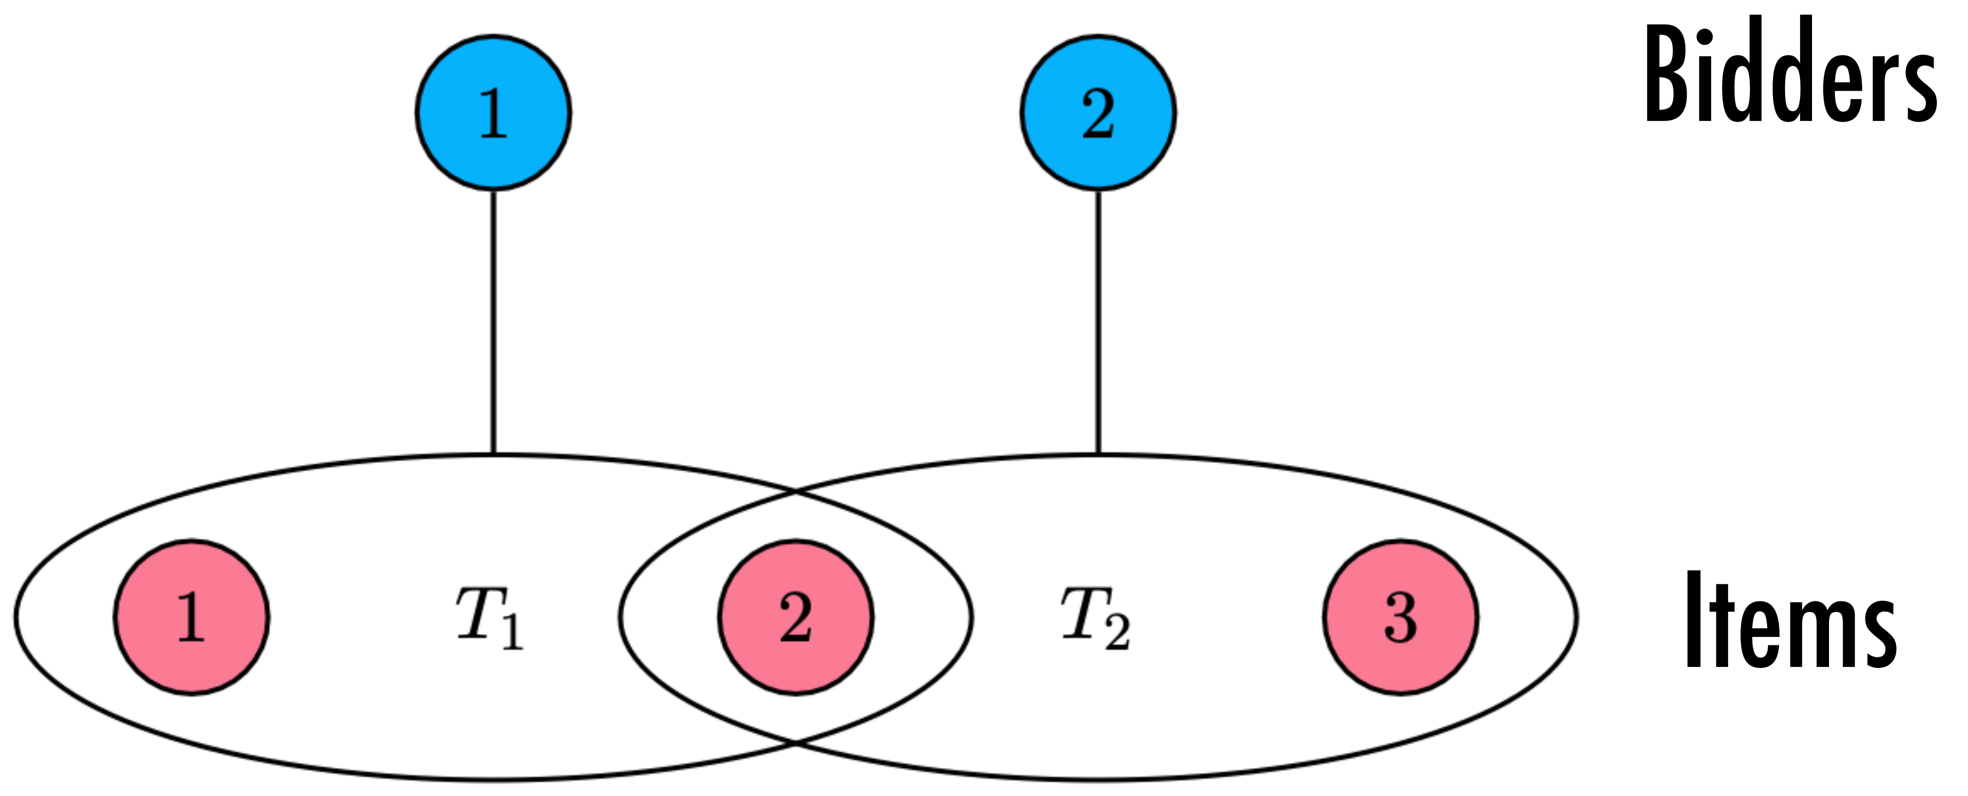

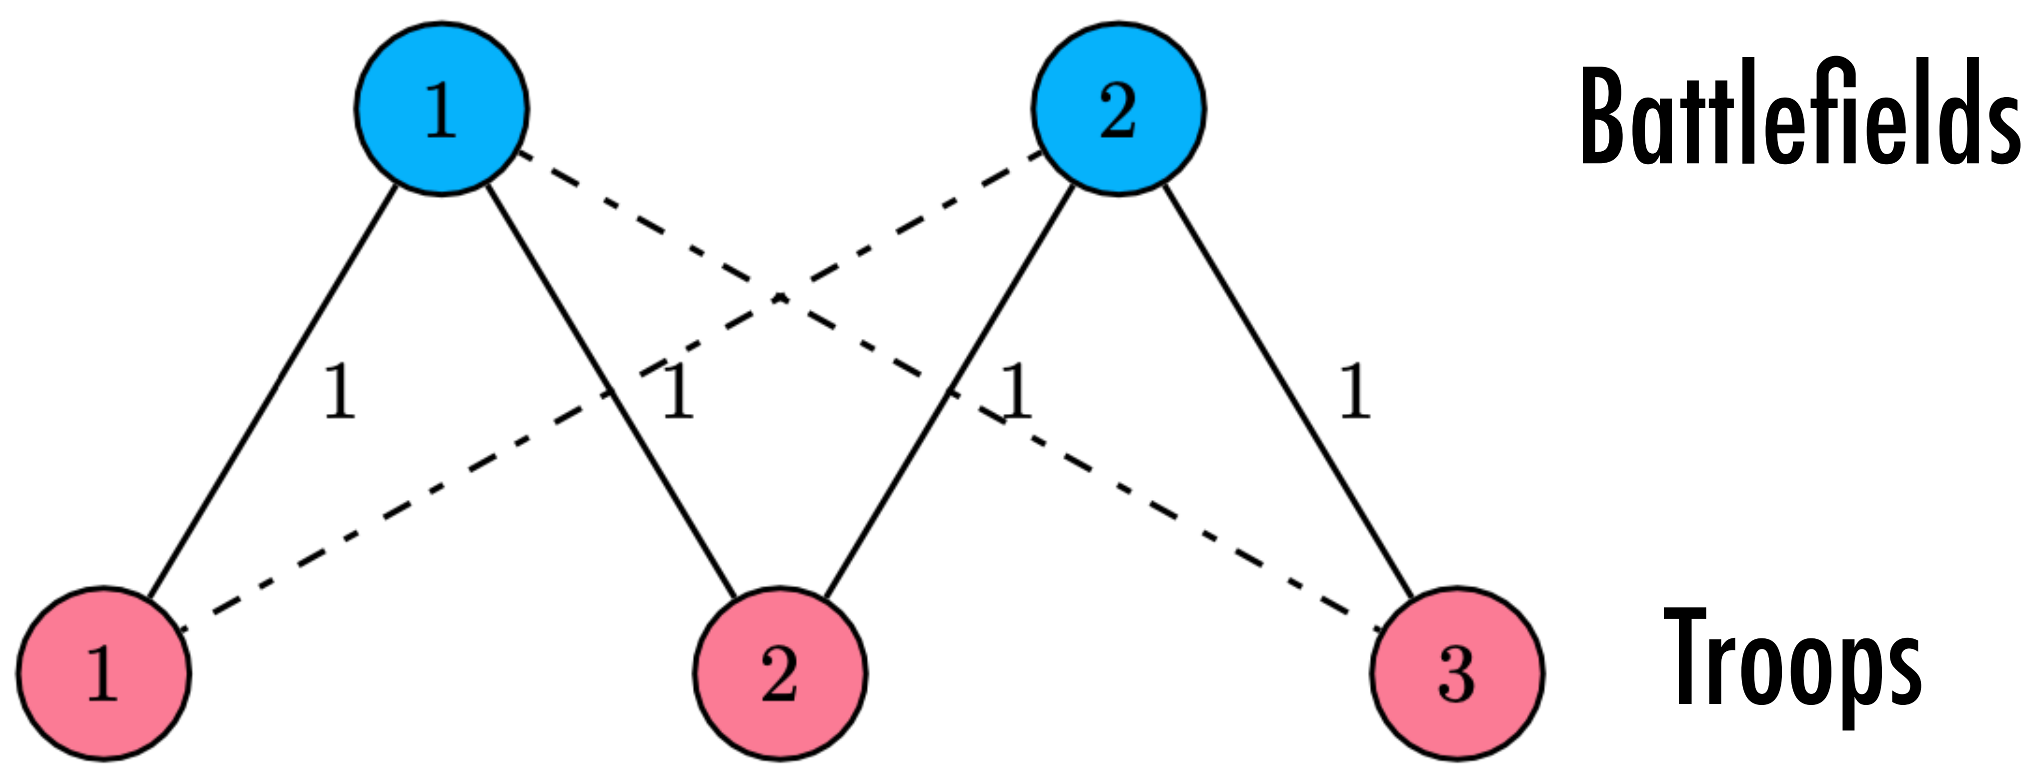

Hardness Result

Theorem. Unless NP = P, there is no polynomial-time algorithm that can always find an $O\left(\sqrt{min(m,n)}\right)$-approximate best response in the multi-faceted Colonel Blotto game.

From

Approximate **Best Response**

to approximate **Maxmin** strategies

Bicriteria Approximation

Multiplicative

- A strategy $\mathbf{Y}$ is an (α, β)-approximate best response strategy to a strategy $\mathbf{X}$ of opponent if:
  - $\mathbf{Y}$ is allowed to use up to α copies of each troop.
  - The payoff is at least $1/\beta$ fraction of the optimal best response against $\mathbf{X}$.

Bicriteria Approximation

Additive

- A strategy $\mathbf{X}$ is an (α, δ)-approximate maxmin strategy if:
  - $\mathbf{X}$ is allowed to use up to α copies of each troop.
  - Let u be the $\mathbf{X}$'s minimum utility against opponent's strategies.
  - Let u^* be the optimal maxmin strategy's minimum utility against an opponent who is allowed to use up to α copies of each troop.
  - $u^* - u \leq \delta$

Bicriteria Approximation

Additive

- A strategy $\mathbf{X}$ is an (α, δ)-approximate maxmin strategy if:
  - $\mathbf{X}$ is allowed to use up to α copies of each troop.
  - Let u be the $\mathbf{X}$'s minimum utility against opponent's strategies.
  - Let u^* be the optimal maxmin strategy's minimum utility against an opponent who is allowed to use up to α copies of each troop.
  - $u^* - u \leq \delta$
- W.l.o.g. Assumption: $\mu^A(\mathbf{X}, \mathbf{Y}) = 1 - \mu^B(\mathbf{X}, \mathbf{Y})$

(α, β)-approximate Best Response

- Given an exact best response oracle the solution of Ahmadinejad et. al (2016) finds a maxmin strategy.
- It leverages the ellipsoid method to find a maxmin strategy.

(α, β)-approximate Best Response

- Given an exact best response oracle the solution of Ahmadinejad et. al (2016) finds a maxmin strategy.
- We don't have access to such oracle here.
- However, as we show later, we construct an (α,β)-approximate best response oracle.

(α, β)-approximate Best Response

- Given an exact best response oracle the solution of Ahmadinejad et. al (2016) finds a maxmin strategy.
- We construct an (α,β)-approximate best response oracle.
- We obtain $(\alpha, 2 - \frac{2}{\beta})$-approximate maxmin strategies using an (α,β)-approximate best response oracle.

Reduction from approximate minmax to approximate best response

- The following LP models the problem:
  - A mixed strategy $\hat{x}$ denotes a point in $k \,.\, \max_f$ dimensions.
  - Each dimension (s^A, b) shows the probability of putting troops with total strength s^A in battlefield b.

$$
\begin{aligned}
\max. \quad & U & & \\
\text{s.t.} \quad & \hat{x} \in S(\mathbf{A}) & & \text{Membership constraints} \\
& \mu^{\mathbf{A}}(\hat{x}, \hat{y}) \geq U, \quad & \forall \hat{y} \in S(\mathbf{B}) \quad & \text{Payoff constraints}
\end{aligned}
$$

Reduction from approximate minmax to approximate best response

- The following LP models the problem:
  - $S(A)$ denotes the set of all feasible strategies for player A.
  - $S(B)$ denotes the set of all feasible strategies for player B.

$$\begin{aligned} \max. \quad & U & & \text{Membership constraints} \\ \text{s.t.} \quad & \hat{x} \in S(\mathbf{A}) & & \\ & \mu^{\mathbf{A}}(\hat{x}, \hat{y}) \geq U, & \forall \hat{y} \in S(\mathbf{B}) & \quad \text{Payoff constraints} \end{aligned}$$

Membership constraints

- We are given a convex polytope Z whose vertices are the pure strategies of the game.
- We wish to find a hyperplane which separates a given point $\hat{x}$ from Z.

$$\begin{aligned} \max. \quad & 0 \\ \text{s.t.} \quad & a_0 + \sum_{i=1}^{d} a_i \hat{x}_i \geq 0 \\ & a_0 + \sum_{i=1}^{d} a_i \hat{z}_i < 0, \qquad \forall \hat{z} \in Z \end{aligned}$$

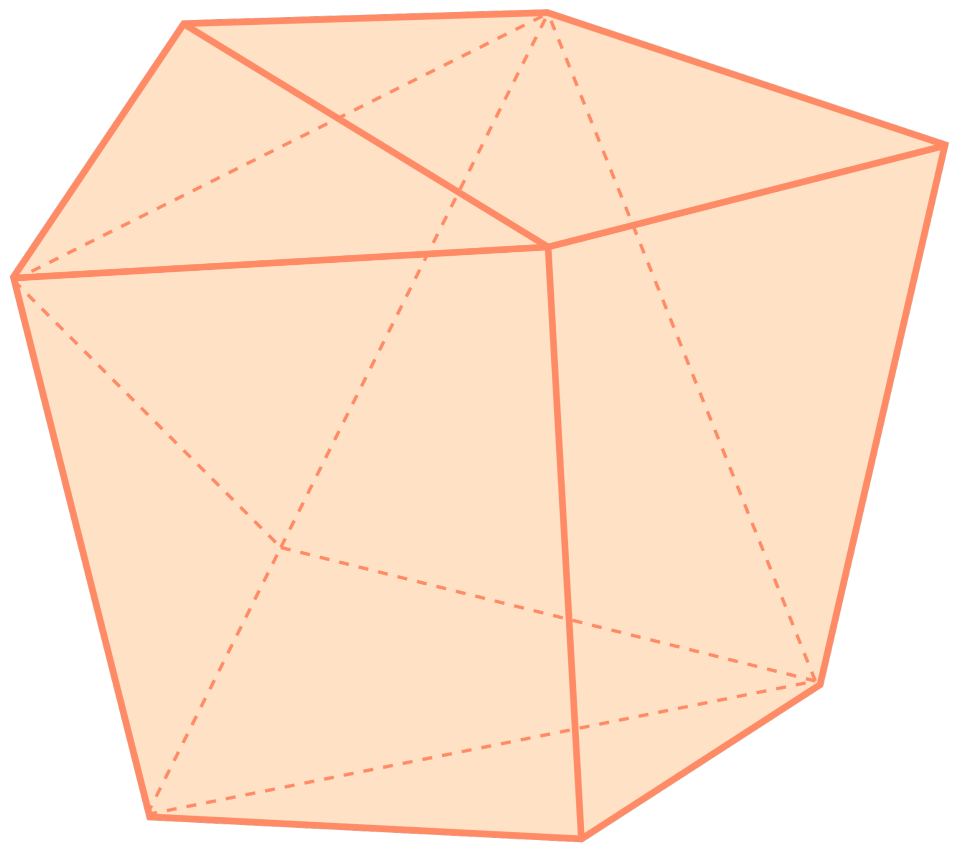

The set of feasible strategies $S(A)$, specified by polytope Z.

Membership constraints

- We wish to find a hyperplane which separates a given point $\hat{x}$ from Z.
- Point $\hat{x}$ is inside Z Iff no such hyperplane exists.

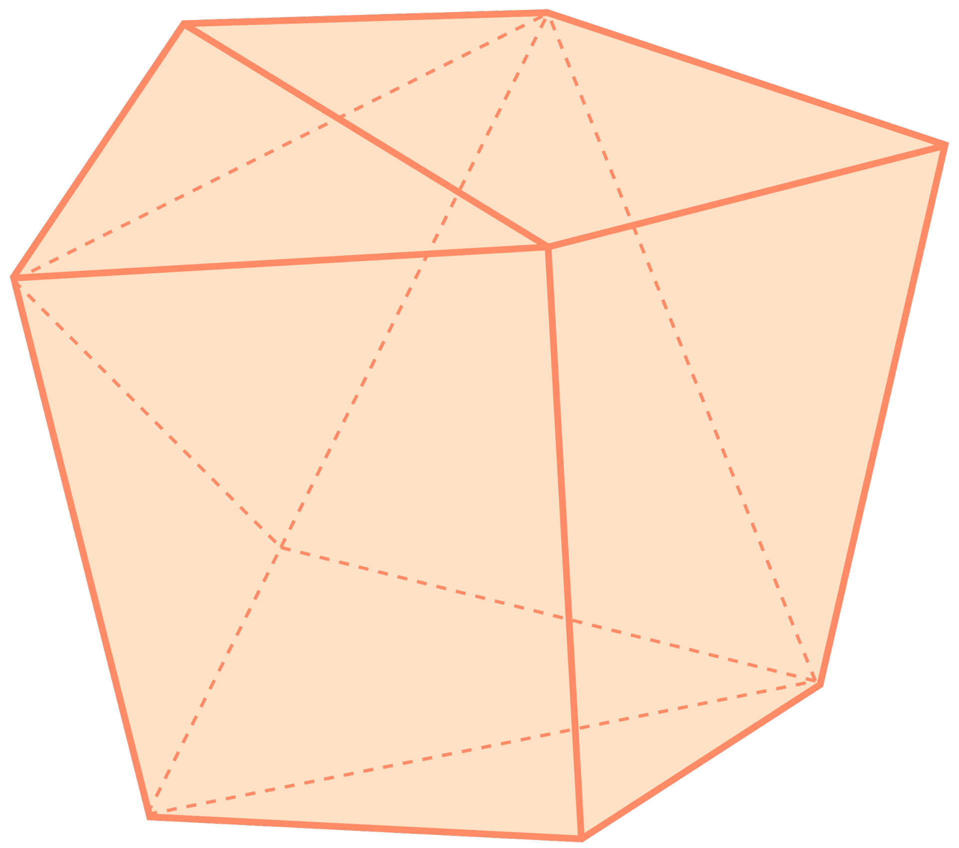

The set of feasible strategies $S(A)$, specified by polytope Z.

$$
\begin{aligned}
\max. \quad & 0 \\
\text{s.t.} \quad & a_0 + \sum_{i=1}^{d} a_i \hat{x}_i \geq 0 \\
& a_0 + \sum_{i=1}^{d} a_i \hat{z}_i < 0, \qquad \forall \hat{z} \in Z
\end{aligned}
$$

Membership constraints

- We wish to find a hyperplane which separates a given point $\hat{x}$ from Z.
- The hyperplane is formulated by $\{a_0, a_1, \ldots, a_d\}$.

$$
\begin{aligned}
\max. \quad & 0 \\
\text{s.t.} \quad & a_0 + \sum_{i=1}^{d} a_i \hat{x}_i \geq 0 \\
& a_0 + \sum_{i=1}^{d} a_i \hat{z}_i < 0, \qquad \forall \hat{z} \in Z
\end{aligned}
$$

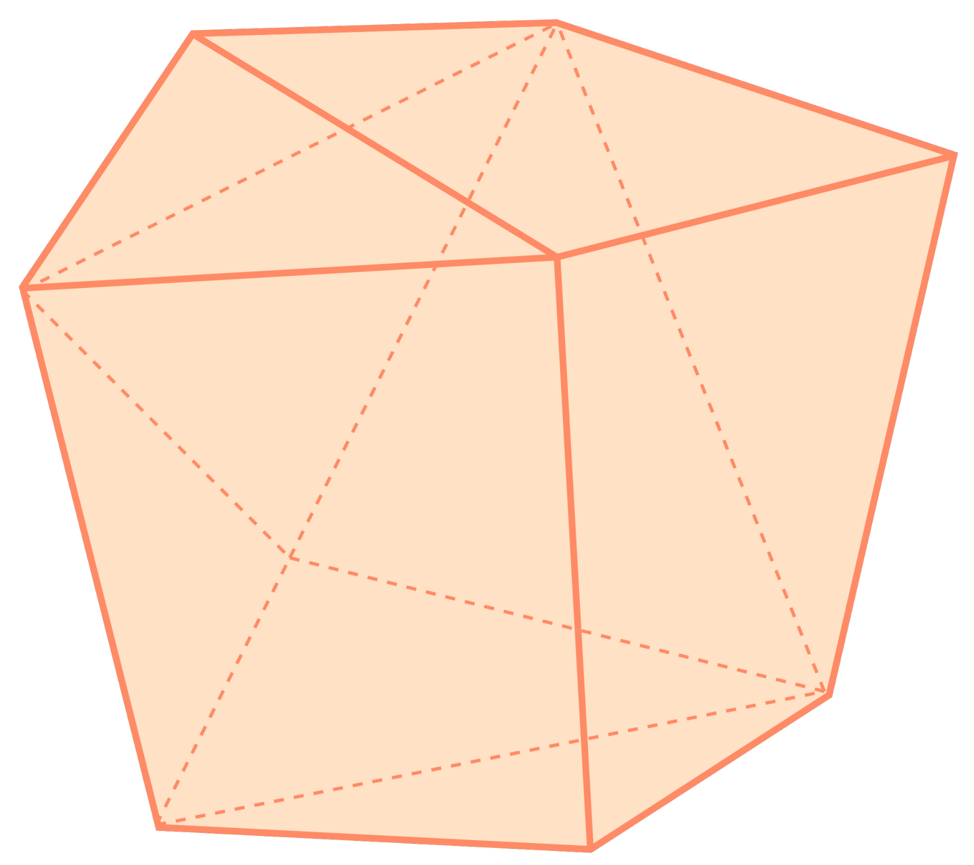

The set of feasible strategies $S(A)$, specified by polytope Z.

Membership constraints

- We wish to find a hyperplane which separates a given point $\hat{x}$ from Z.
- We can simplify the second set of constraints by only considering $\hat{z}_{\text{max}}(a)$, the vertex which maximizes the summation.

$$\begin{aligned} \max. \quad & 0 \\ \text{s.t.} \quad & a_0 + \sum_{i=1}^{d} a_i \hat{x}_i \geq 0 \\ & a_0 + \sum_{i=1}^{d} a_i \hat{z}_i < 0, \qquad \forall \hat{z} \in Z \end{aligned}$$

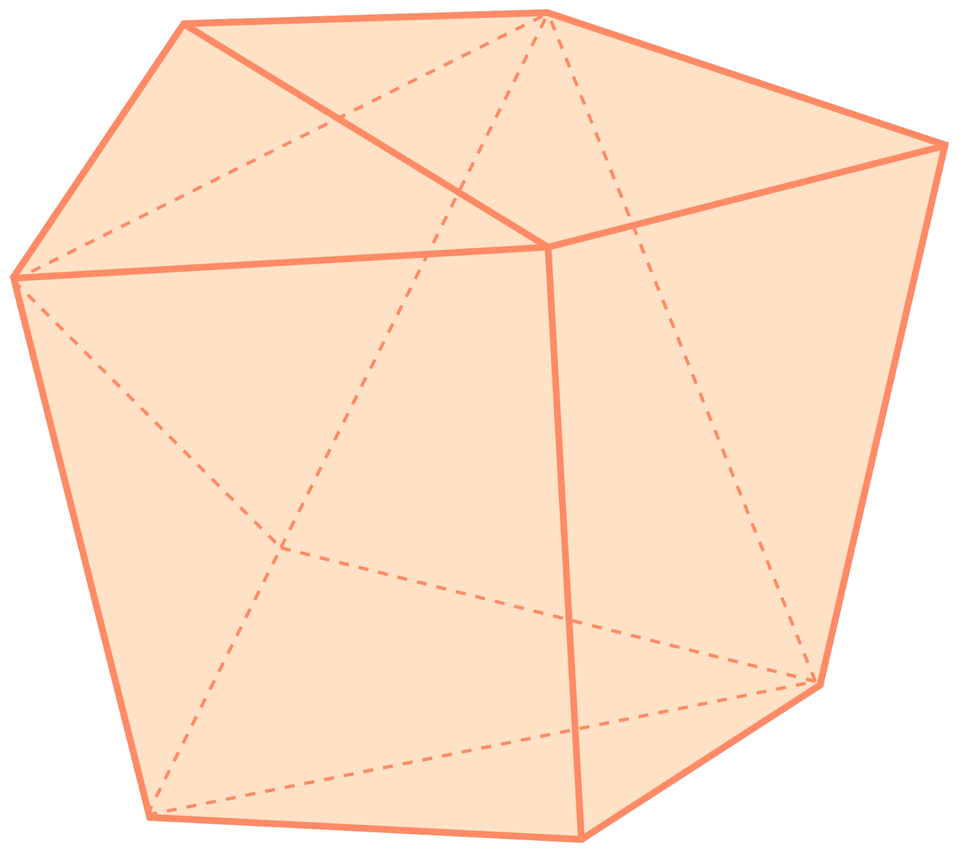

The set of feasible strategies $S(A)$, specified by polytope Z.

Membership constraints

- We can simplify the second set of constraints by only considering $\hat{z}_{\max}(a)$, the vertex which maximizes the summation.
- It is possible to find $\hat{z}_{\max}(a)$ in polynomial time if we have access to an exact best response oracle.

$$\begin{aligned} \max.\quad & 0 \\ \text{s.t.}\quad & a_0 + \sum_{i=1}^{d} a_i \hat{x}_i \geq 0 \\ & a_0 + \sum_{i=1}^{d} a_i \hat{z}_i < 0, \qquad \forall \hat{z} \in Z \end{aligned}$$

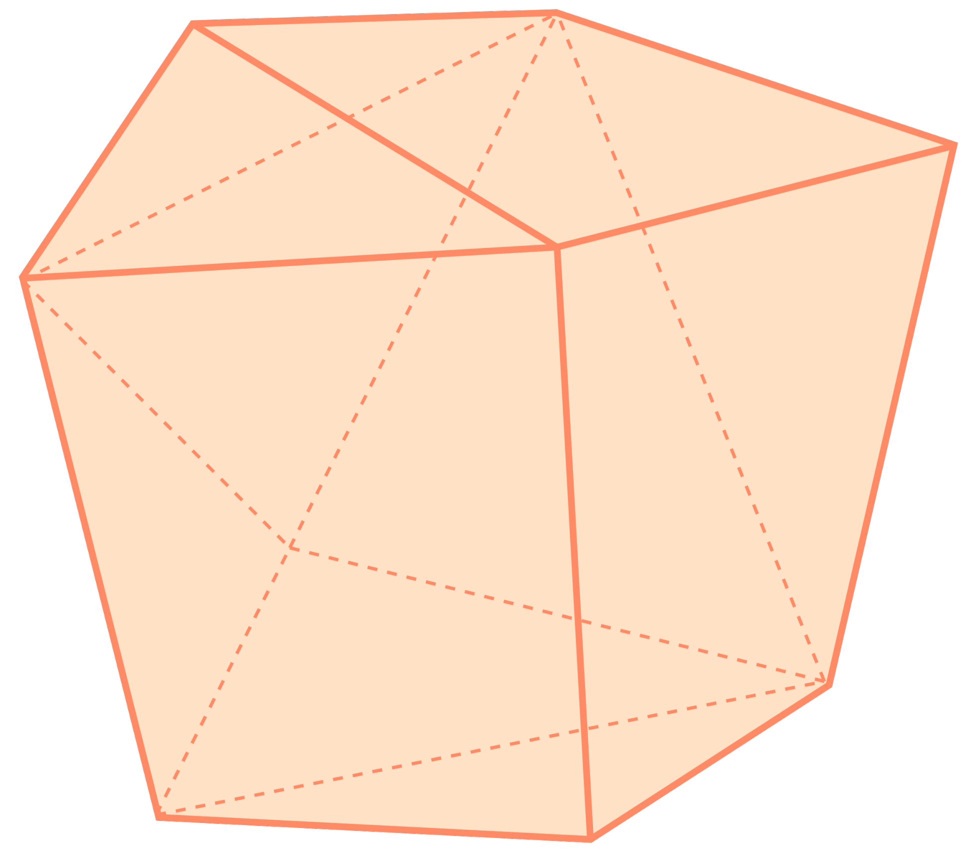

The set of feasible strategies $S(A)$, specified by polytope Z.

Membership constraints

- We can simplify the second set of constraints by only considering $\hat{z}_{\max}(a)$, the vertex which maximizes the summation.
- Instead of $\hat{z}_{\max}(a)$, we find $\hat{z}^*(a)$:
  - A feasible strategy if we have α copies of each troop
  - $$\sum_{i=1}^{d} a_i \hat{z}^*(a)_i \geq \frac{1}{\beta} \sum_{i=1}^{d} a_i \hat{z}_{\max}(a)_i$$

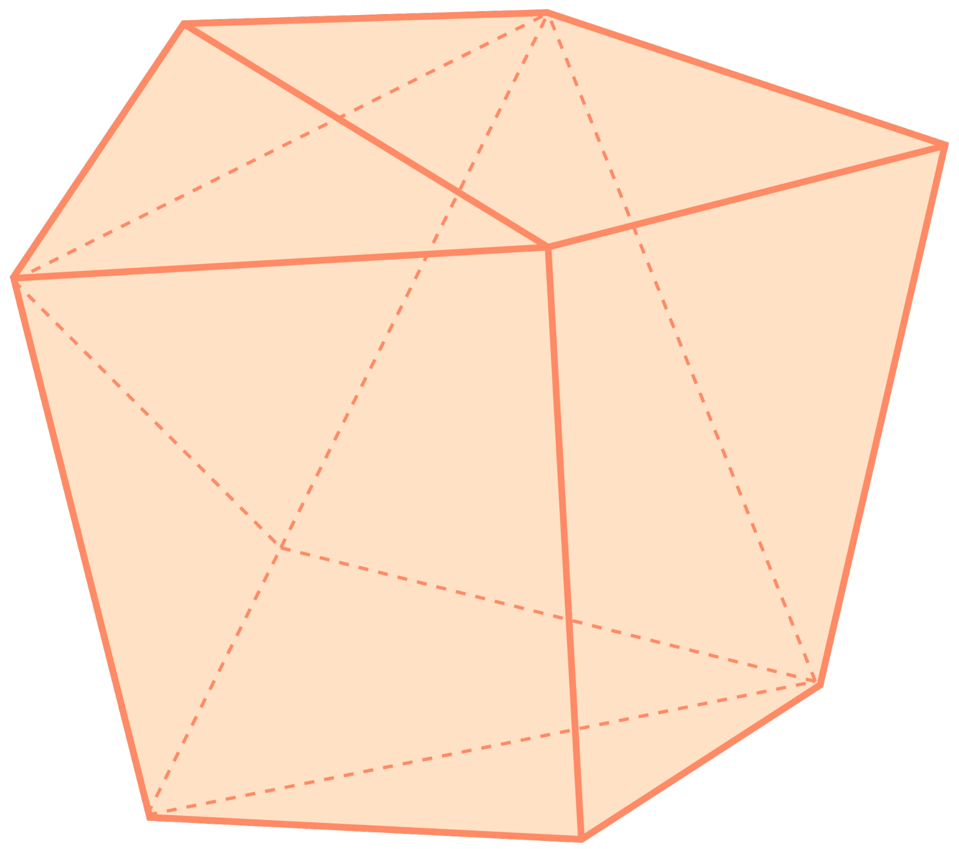

The set of feasible strategies $S(A)$, specified by polytope Z.

Membership constraints

- We can simplify the second set of constraints by only considering $\hat{z}_{\max}(a)$, the vertex which maximizes the summation.
- Instead of $\hat{z}_{\max}(a)$, we find $\hat{z}^*(a)$.
- We define an instance of (α,β)-approximate best response oracle as following:
    - The utility of a strategy $\hat{z}$ equals: $\sum_{i=1}^{d} a_i \hat{z}_i$
- Let $\hat{z}^*(a)$ be the best response strategy returned by the oracle.

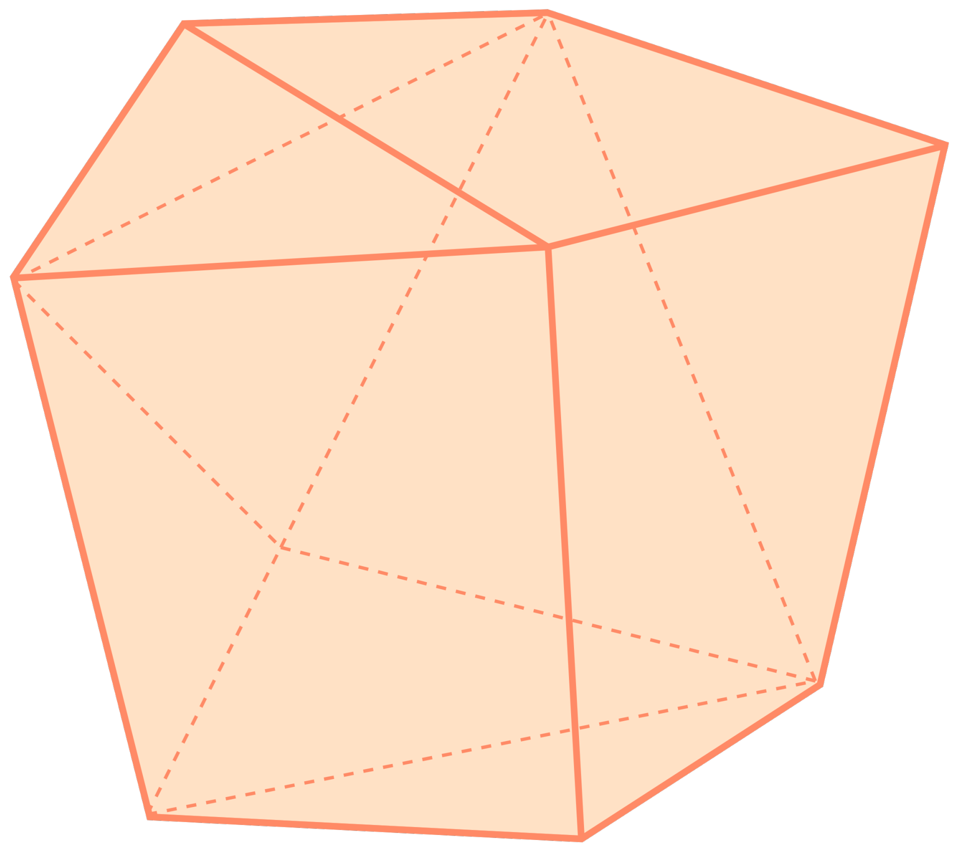

The set of feasible strategies $S(A)$, specified by polytope Z.

Membership constraints

- We try to solve our LP using Ellipsoid method and $\hat{z}^*(a)$: the oracle only checks if the current hyperplane satisfies $\hat{z}^*(a)$.
- Let $S'(A)$ denote the set of points that this algorithm admits.

$$
\begin{aligned}
\max. \quad & 0 \\
\text{s.t.} \quad & a_0 + \sum_{i=1}^{d} a_i \hat{x}_i \geq 0 \\
& a_0 + \sum_{i=1}^{d} a_i \hat{z}_i < 0, \qquad \forall \hat{z} \in Z^\alpha
\end{aligned}
$$

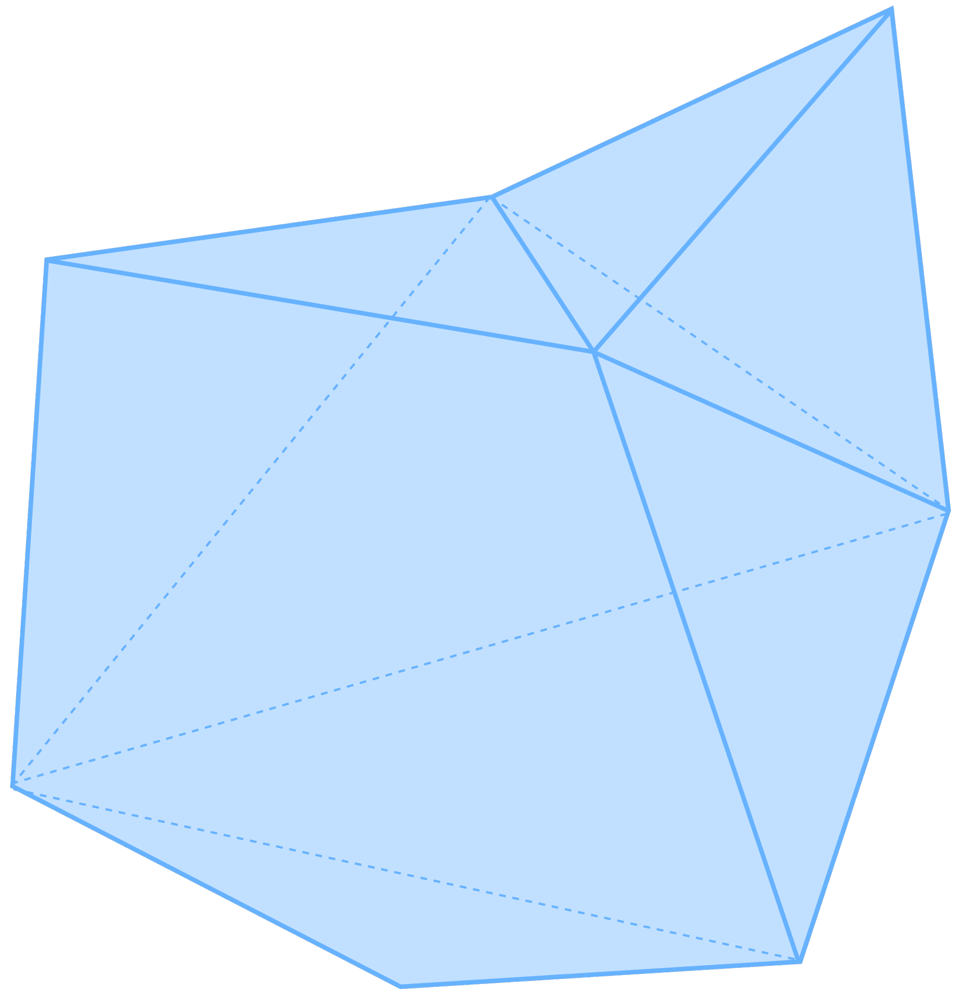

The set of strategies $S'(A)$, which are admitted by our algorithm.

Membership constraints

- Let $S'(A)$ denote the set of points that this algorithm admits.
- $S'(A)$ is not necessarily convex.

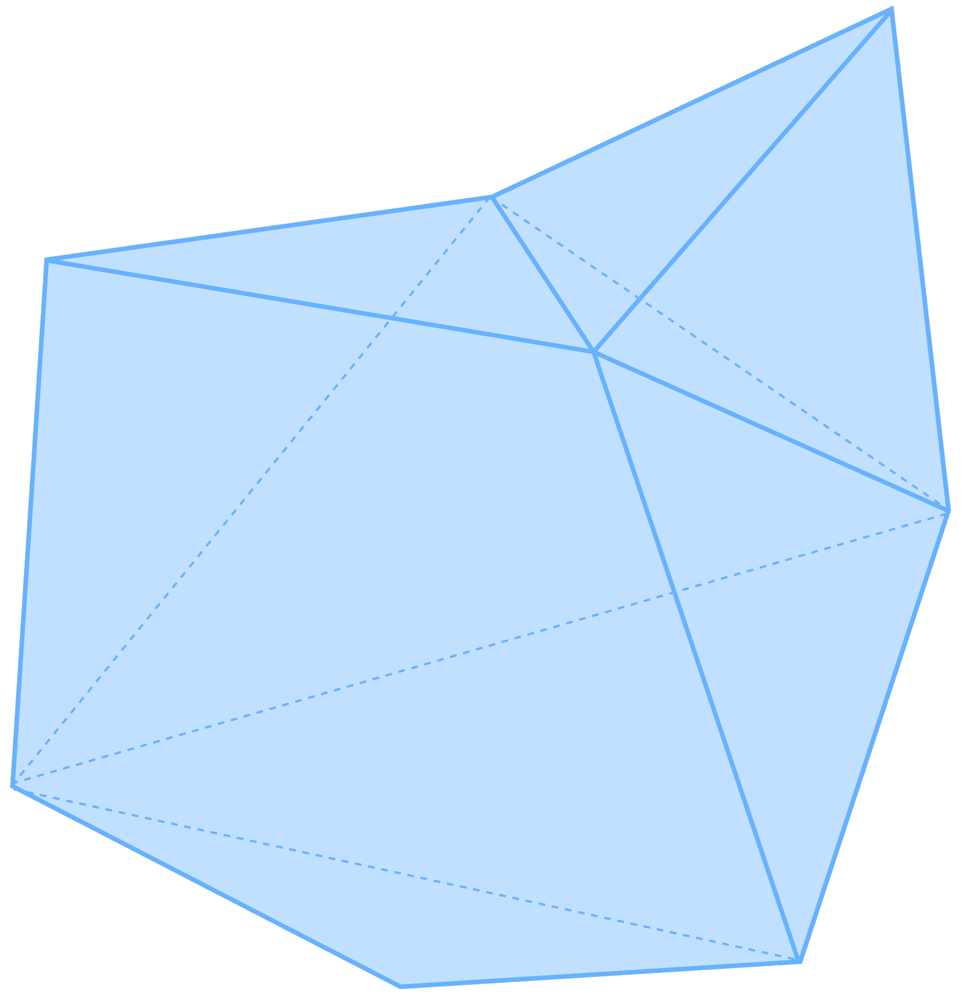

The set of strategies $S'(A)$, which are admitted by our algorithm.

$$
\begin{aligned}
\max. \quad & 0 \\
\text{s.t.} \quad & a_0 + \sum_{i=1}^{d} a_i \hat{x}_i \geq 0 \\
& a_0 + \sum_{i=1}^{d} a_i \hat{z}_i < 0, \qquad \forall \hat{z} \in Z^{\alpha}
\end{aligned}
$$

Membership constraints

- Let $S'(A)$ denote the set of points that this algorithm admits.
- But we have the following properties for any $\hat{x}' \in S'(A)$:
  - $\hat{x}'$ is a feasible strategy if we allow α copies of each troop.
  - If $\hat{x} \in S(A)$, then $\frac{\hat{x}}{\beta} \in S'(A)$.

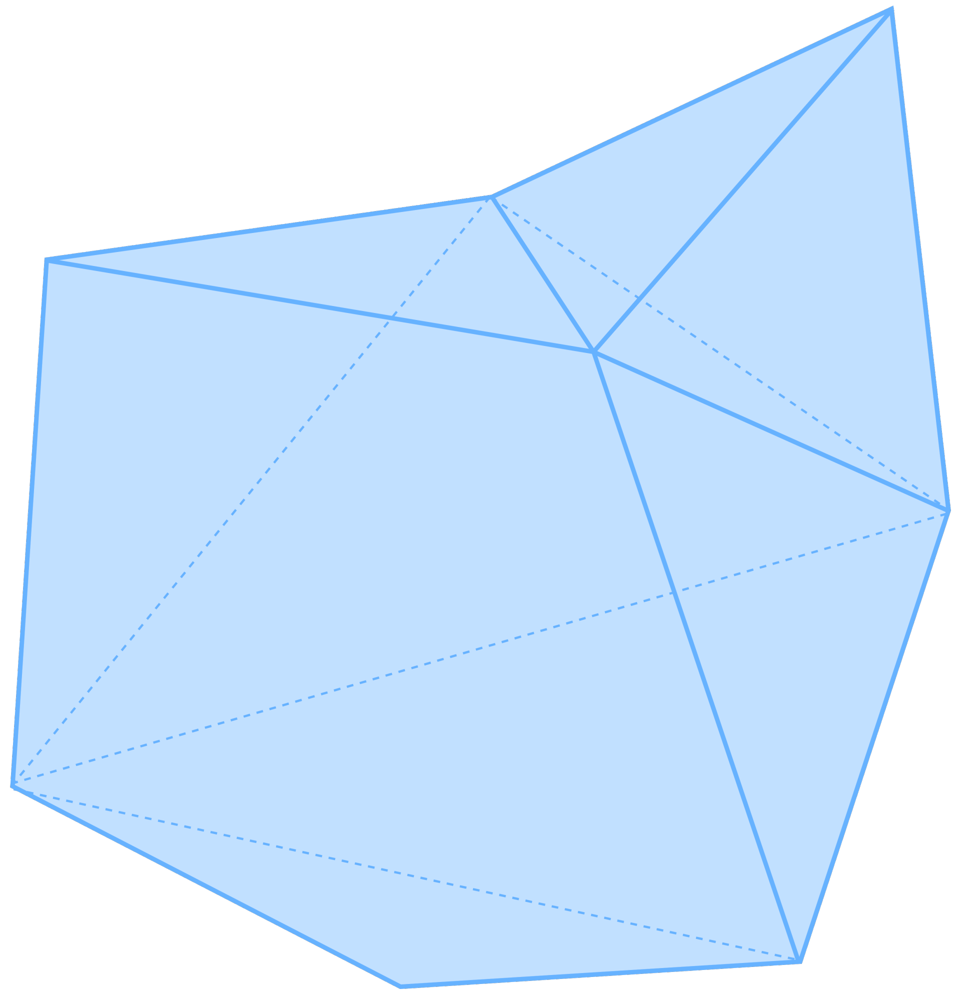

The set of strategies $S'(A)$, which are admitted by our algorithm.

Membership constraints

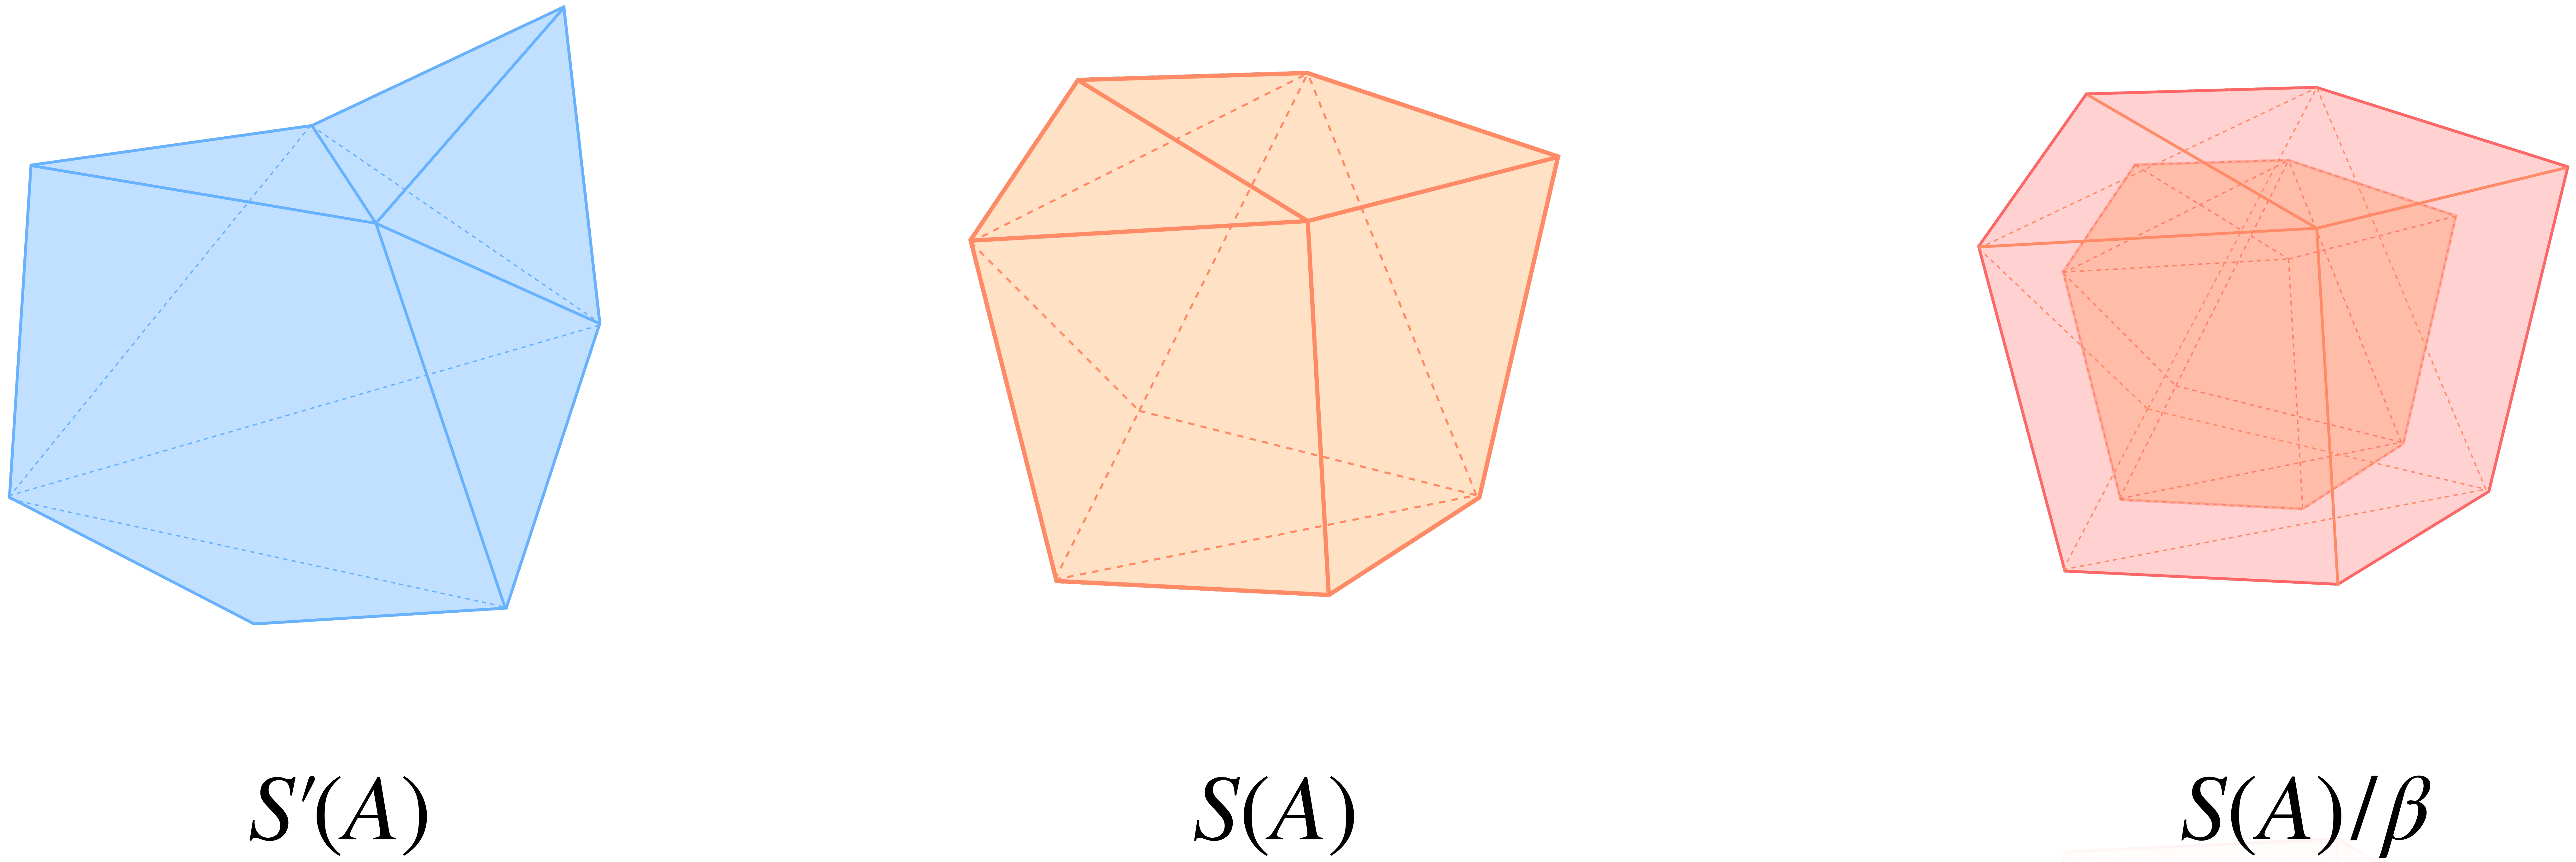

Payoff constraints

- In order to use (α,β)-approximate best response oracle for our payoff constraints, we need to reformulate it as a minmax LP:

$$\begin{aligned} \min. \quad & U \\ \text{s.t.} \quad & \hat{x} \in S(\mathbf{A}) \\ & \mu^{\mathbf{B}}(\hat{x}, B^{\alpha,\beta}(\hat{x})) \leq U \end{aligned}$$

Payoff constraints

- We show that by using approximate best response in the last constraint, the total loss in approximation is bounded by $2 - 2/\beta$.

$$
\begin{aligned}
\min. \quad & U \\
\text{s.t.} \quad & \hat{x} \in S(\mathbf{A}) \\
& \mu^{\mathbf{B}}(\hat{x}, B^{\alpha,\beta}(\hat{x})) \leq U
\end{aligned}
$$

Reduction from approximate minmax to approximate best response

Theorem. Given a polynomial time algorithm that finds an (α,β)-approximate best-response for the generalized Colonel Blotto game, one can find an $(\alpha, 2-\frac{2}{\beta})$-approximate minmax solution for the game in polynomial time.

Approximate Best Response

Heterogenous troops w.r.t battlegrounds

Theorem. For any $\epsilon > 0$, There exists a polynomial-time algorithm which obtains an $\left(O\left(\frac{\ln 1/\epsilon}{\epsilon}\right), 2\epsilon\right)$-maxmin strategy for the generalized Colonel Blotto game in the heterogenous setting.

Heterogenous troops w.r.t battlegrounds

Theorem. For any $\epsilon > 0$, There exists a polynomial-time algorithm which obtains an $\left(O\left(\frac{\ln 1/\epsilon}{\epsilon}\right), 2\epsilon\right)$-maxmin strategy for the generalized Colonel Blotto game in the heterogenous setting.

- Obtained by plugging an $\left(O\left(\frac{\ln 1/\epsilon}{\epsilon}\right), \frac{1}{1-\epsilon}\right)$-best response into the reduction.

Heterogenous troops w.r.t battlegrounds

Theorem. For any $\epsilon > 0$, There exists a polynomial-time algorithm which obtains an $\left(O\left(\frac{\ln 1/\epsilon}{\epsilon}\right), 2\epsilon\right)$-maxmin strategy for the generalized Colonel Blotto game in the heterogenous setting.

- The number of copies of each troop we need is 1 in expectation.
- But in the worst case we may require $o\left(\frac{\ln 1/\epsilon}{\epsilon}\right)$ copies of each troop.

Improved solutions for **Homogenous** battlegrounds

Homogenous troops w.r.t battlegrounds

- Search space dimensions reduces from $k \cdot (\max_f + 1)$ to $(\max_f + 1)$.
- We can represent the best response with a vector p of probability coefficients with length $(\max_f + 1)$.
- Reduce to Prize-collecting Knapsack problem.

Prize-collecting Knapsack

- A set of bag types $\mathscr{B} = \{1,2,\ldots,|\mathscr{B}|\}$ is given.
- Each bag type i has size v_i and prize p_i.
- Unlimited copies of each bag is available.
- A set of items $\mathscr{N} = \{1,2,\ldots,|\mathscr{N}|\}$ is given each with size a_i.
- We gain profit of p_i whenever we fill a bag of type i by a subset of items with total size of at least v_i.

Prize-collecting Knapsack

- We obtain a $(1 + \epsilon, 1)$-approximation of prize-collecting knapsack using dynamic programming.
- Key ideas:
    - Discretize the size of items by rounding to the nearest $(1 + \epsilon)^k$ value.
        - $O(\log(\text{max}_f))$ different sizes.

Prize-collecting Knapsack

- We obtain a $(1 + \epsilon, 1)$-approximation of prize-collecting knapsack using dynamic programming.
- Key ideas:
    - Divide items into three groups based on their size w.r.t each bag type i:
        - Large items ($\mathscr{L}$): $v_i < a_j$
        - Regular items ($\mathscr{R}$): $\epsilon v_i \leq a_j \leq v_i$
        - Small items ($\mathscr{S}$): $a_j < \epsilon v_i$

Prize-collecting Knapsack DP

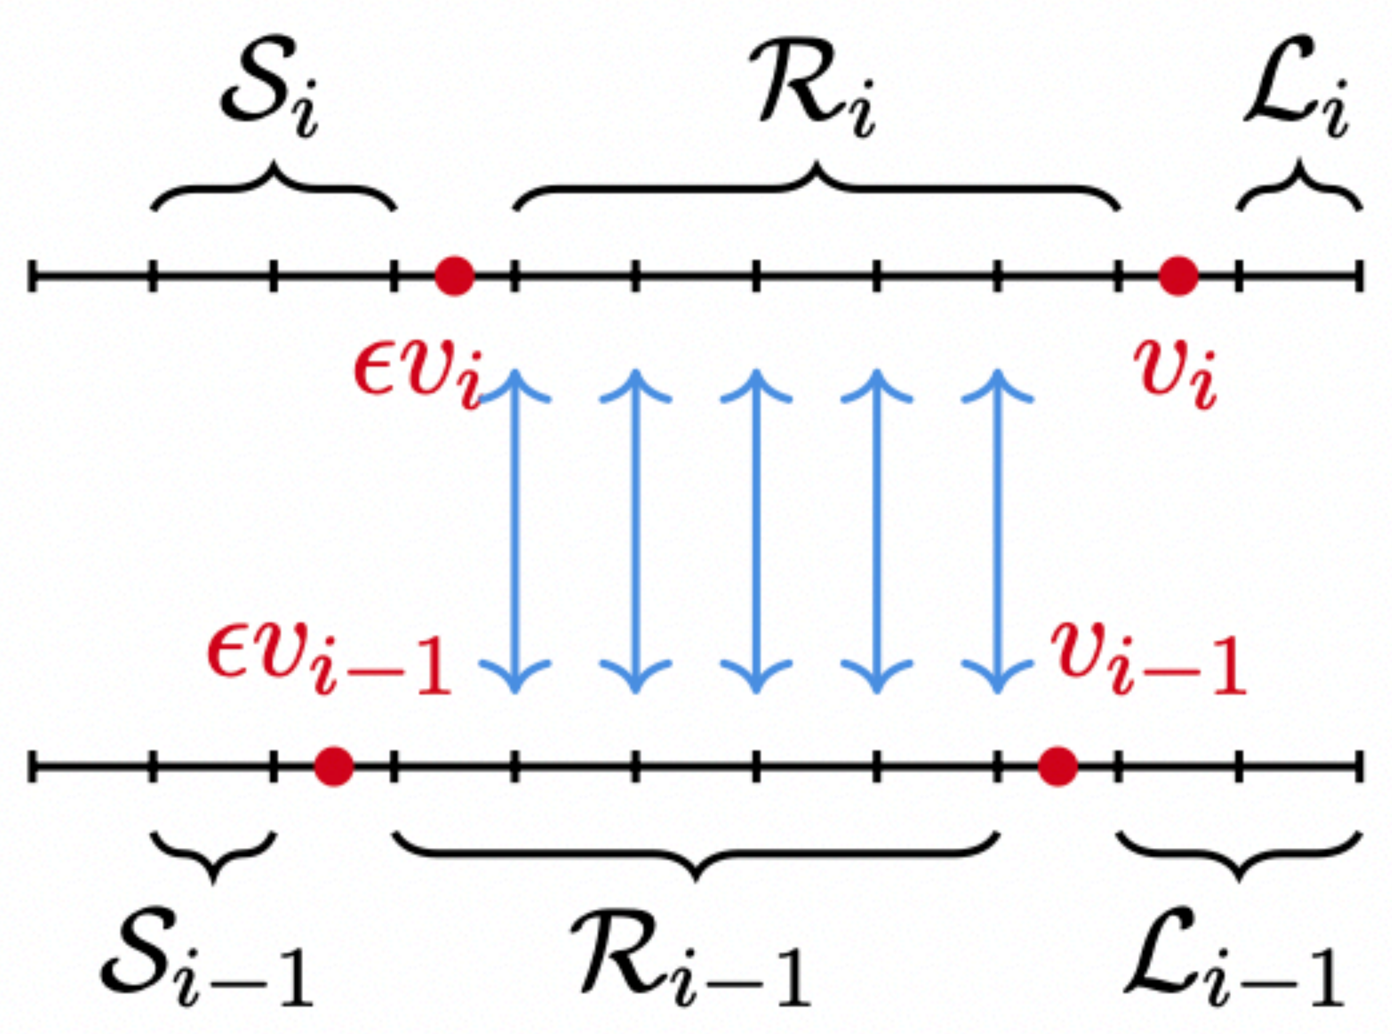

$\epsilon v_i \leq v_{i-1}$

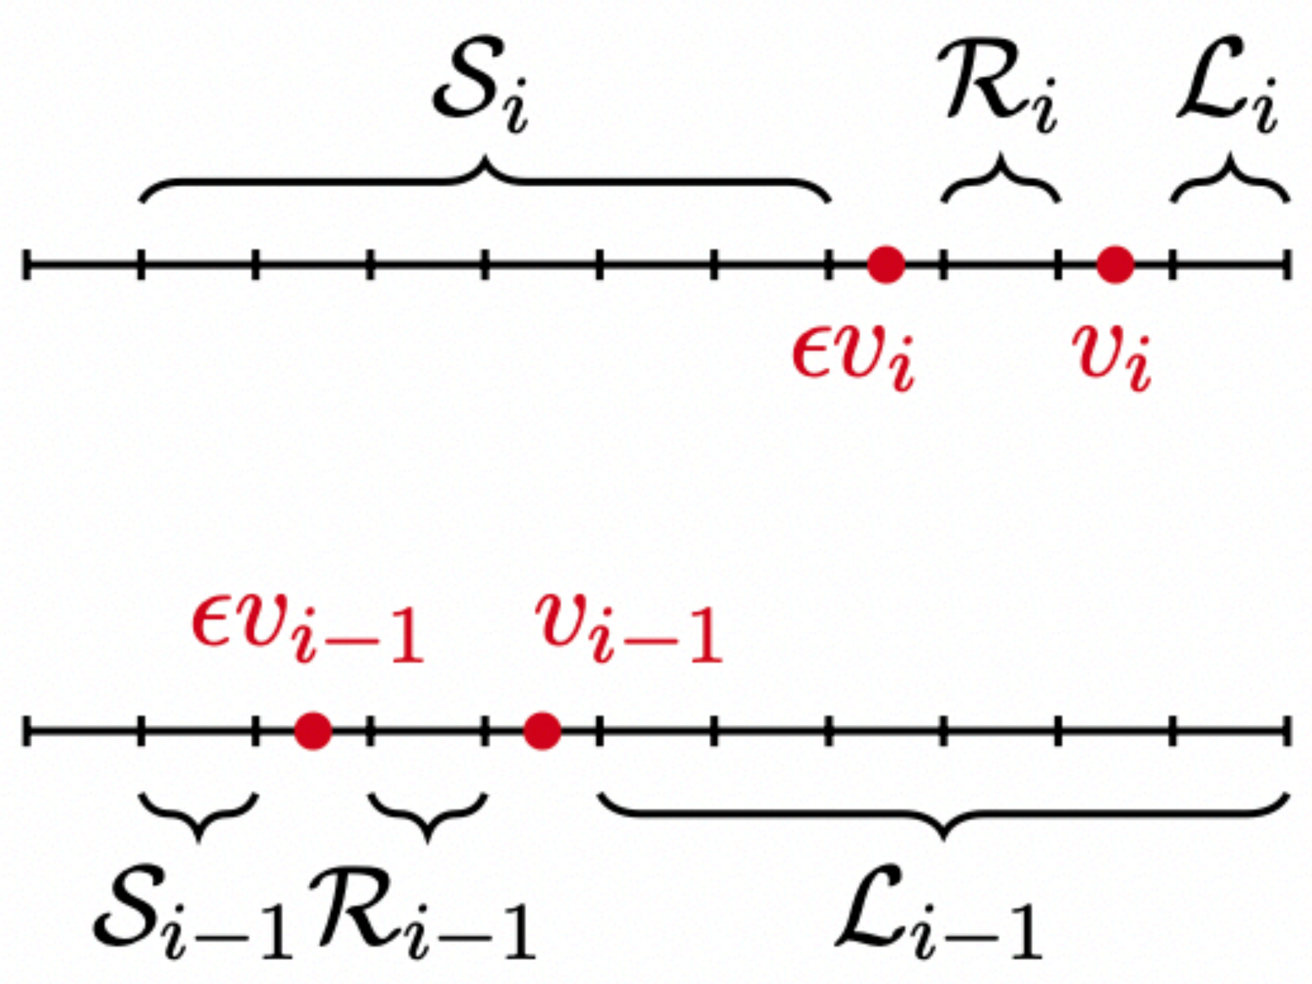

$\epsilon v_i > v_{i-1}$

Reduction to Prize-collecting Knapsack

- Randomly permuting the battlegrounds of an optimal solution preserves optimality.

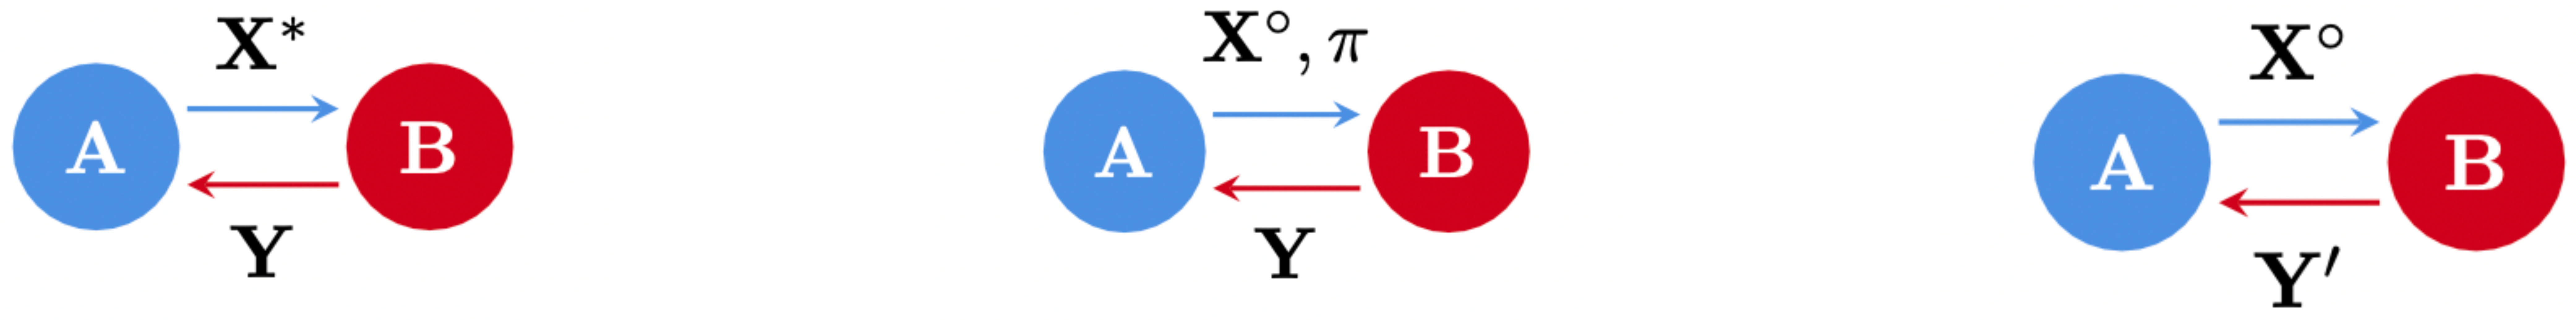

Homogenous troops w.r.t battlegrounds

Theorem. We can approximate the maxmin strategy of the generalized Colonel Blotto game within a bi-criteria approximation factor of $(1+\varepsilon, 0)$ in the homogenous setting in polynomial time.

Beyond
Zero-sum and Linearity

Are all assumptions necessary?

- A more generalized version of problem covering a broad range of multi-battlefield two player games.
- What happens if we eliminate each assumption of our current formulation?
  - Linearity of utilities
  - Zero-sum payoffs

Removing Linearity Constraint

Theorem. The problem of computing an equilibrium in non-linear battlefield-wide zero-sum two-player-multi-battlefield games is PPAD-hard.

Removing Zero-Sum Constraint

Theorem. The problem of computing an equilibrium in linear non-zero- sum two-player-multi-battlefield games is PPAD-hard.

Thank You!