

Colonel Blotto with Battlefield Games

Salam Afouni¹, Jakub Černý¹, Chun Kai Ling², Christian Kroer¹

¹Columbia University, USA

²National University of Singapore, Singapore

sa4316@columbia.edu, jakub.cerny@columbia.edu, chunkail@nus.edu.sg, christian.kroer@columbia.edu

Abstract

We study a class of two-player zero-sum Colonel Blotto games in which, after allocating soldiers across battlefields, players engage in (possibly distinct) normal-form games on each battlefield. Per-battlefield payoffs are parameterized by the soldier allocations. This generalizes the classical Blotto setting, where outcomes depend only on relative soldier allocations. We consider both discrete and continuous allocation models and examine two types of aggregate objectives: linear aggregation and worst-case battlefield value. For each setting, we analyze the existence and computability of Nash equilibrium. The general problem is not convex-concave, which limits the applicability of standard convex optimization techniques. However, we show that in several settings it is possible to reformulate the strategy space in a way where convex-concave structure is recovered. We evaluate the proposed methods on synthetic and real-world instances inspired by security applications, suggesting that our approaches scale well in practice.

Code — github.com/CoffeeAndConvexity/ColonelBlottoWithBattlefieldGames

Extended version — arxiv.org/abs/2511.06518

1 Introduction

Colonel Blotto games are a classic game class used to model competitive resource allocation (Borel and Ville 1991; Roberson 2006; Kovenock and Roberson 2012b). Introduced by Émile Borel in the 1920s, Blotto games describe situations where players simultaneously allocate soldiers (resources) between battlefields, after which each battlefield is “won” by the player who allocated more soldiers to it. In the classic setting, each battlefield features a winner-takes-all property, creating an interesting game theoretic conundrum where both players would like to either just barely win or badly lose every battlefield. Blotto games and their extensions have been studied in many disciplines including computer and social sciences, with applications to politics (Che and Gale 2008; Laslier and Picard 2002; Myerson 1993), warfare (Gross and Wagner 1950; Shubik and Weber 1981), and behavioral science (Chowdhury, Kovenock, and Sheremeta 2013; Arad and Rubinstein 2012).

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In this paper, we consider a variant of a two-player Colonel Blotto game played over n heterogeneous battlefields that features two levels of play. In our setting, each battlefield $i \in [n]$ also comprises a zero-sum “subgame” G_i , the structure and payoffs of which depends on the soldiers each player allocated to i . This captures many real-world settings where players follow-up with individual battlefield-level strategies. For instance, a political party first allocates money between state elections, after which, each individual state decides how they should spend that money. Likewise in warfare, the tactics at a battlefield level play a huge part in the success of the overarching conflict; these battlefield level tactics are themselves games played at a lower level, the outcomes of which depend greatly on the number of soldiers assigned. In more abstract terms, players first simultaneously determine a soldier assignment over battlefields (unobserved by the opponent), then play, in *every* battlefield, a subgame whose payoff matrix is parametrized by the player’s soldier allocation on that battlefield. We refer to this game as the “two-level Blotto game”. By convention, Player 1 (resp. Player 2) seeks to minimize (resp. maximize) the payoff. Accordingly, we often refer to Player 1 (resp. Player 2) as the minimizing (resp. maximizing) player.

We analyze three dimensions of the proposed two-level Blotto games. First, we study *discrete versus continuous* soldier types; the former implies that individual soldiers cannot be subdivided, though players are allowed to adopt randomization. Second, we study two ways of *aggregating* payoffs from individual battlefields – either with sum or min aggregators. Both these axes have been extensively explored in the Blotto literature (see, e.g. Roberson (2011); Washburn (2011), or Vu (2020)). Third, we study two *strategic settings*: the two-sided case, where both players allocate soldiers across battlefields and the simpler one-sided case, in which only the maximizing player is allowed soldier allocations. The one-sided setting is motivated by security applications, where only the “defending” player has soldiers to allocate, and provides stronger guarantees of equilibrium existence and computational tractability.

Our contributions are as follows. First, for every combination of soldier type (discrete vs. continuous), aggregator (sum vs. min), and strategic setting (two-sided vs. one-sided), we address the following: (i) does a Nash equilibrium exist, and (ii) are the max-min and min-max strategies well-

Setting		NE	Mx/Mn exist.
Discrete	Two-sided sum	✓	✓
	Two-sided min	✗	✓
	One-sided min	✓	✓
Continuous	Two-sided sum	✗	✓*
	One-sided sum	✗	✓*
	One-sided sum (L)	✓	✓
	Two-sided min	✗	✓*
	One-sided min	✓*	✓*
	One-sided min (L)	✓	✓

Table 1: Existence of NE and max-min strategies in different settings of discrete and continuous Blotto games. * holds only when the utility u is continuous in the soldiers allocation. (L) means that battlefield utilities are linear in the maximizing player’s soldier allocation.

defined and attained? ¹ Our results are summarized in Table 1. Second, for each of the cases where an equilibrium exists, we provide algorithms for solving them. In most cases, these problems reduce to linear programs, though in one special case the problem reduces to a quasiconcave problems. Third, we empirically evaluate these algorithms on synthetic data, demonstrating that these algorithms are usable.

Related work. Due to space constraints, we only provide a small sampling of relevant work. Computationally, Ahmadinejad et al. (2019) provided an LP based approach for separable battlefields (similar to one of our settings) to solve Blotto games but using the ellipsoid method; their method extends to continuous *Lotto* games (Hart 2008; Dziubiński 2013), which admit infinite action spaces, though they do not have subgames. Behnezhad et al. (2023) avoid the impractical ellipsoid algorithm and propose a practical LP formulation based on layered graphs, and is very closely related to one of our approaches. Kvasov (2007) and Roberson and Kvasov (2012) study a non-constant (zero) sum variant of Blotto games, while Hortala-Vallve and Llorente-Saguer (2012) study pure strategy equilibria in such variants. Kovenock and Roberson (2012a) study a team variant where a single player plays against a coalition of 2 other players while Boix-Adserà, Edelman, and Jayanti (2020) study more general multiplayer settings. Stephenson (2024) study payoffs which depend on a function of allocated soldiers per player. Powell (2009) and Hausken (2012) consider sequential variants of Blotto games.

2 Preliminaries

Let $n > 0$ be the number of battlefields, and m^1, m^2 be the total number of soldiers available to Players 1 and 2. Let \mathcal{Z}^j denote the set of all possible soldier assignments z^j for player $j \in \{1, 2\}$, where $z^j = (z_1^j, \dots, z_n^j)$ is a vector speci-

¹A positive answer to both questions implies a valid minimax theorem in that regime. Moreover, any existence result in the two-sided model implies the one-sided case by fixing the minimizer’s allocations to zero.

fying soldier distributions across the n battlefields, subject to the budget constraint $\sum_{i=1}^n z_i^j = m^j$. When soldiers are discrete, we let $z^j = k^j \in \mathbb{Z}_{\geq 0}^n$, and when soldiers are continuous, $z^j = \sigma^j \in \mathbb{R}_{\geq 0}^n$. We denote by \mathcal{K}^j and Σ^j the sets of all possible soldier assignments for player j in the discrete and continuous case, respectively. Each subgame G_i is a finite normal-form zero-sum game defined as a tuple $(\mathcal{A}_i^1, \mathcal{A}_i^2, u_i)$, where \mathcal{A}_i^j denotes the action space of player $j \in \{1, 2\}$ in battlefield i , and $u_i : \mathcal{A}_i^1 \times \mathcal{A}_i^2 \times \mathcal{Z}^1 \times \mathcal{Z}^2 \rightarrow \mathbb{R}$ denotes the payoff in that battlefield that Player 2 (resp. Player 1) seeks to maximize (resp. minimize). For each battlefield $i \in [n]$, the payoff u_i is a function of both the number of soldiers players allocated to that battlefield and the actions they play in the subgame G_i . We denote by $\alpha_i^j \in \mathcal{A}_i^j$ the action player j plays on battlefield i . For battlefield i and any allocation profile (z_i^1, z_i^2) , we let $v_i^*(z_i^1, z_i^2)$ be the equilibrium value of subgame G_i when player j allocates z_i^j soldiers to it. To compute the overall utility in the two-level Blotto game, we consider two aggregation functions: 1) the sum of the expected utilities across battlefields $U = \sum_{i \in [n]} \mathbb{E}[u_i(\cdot)]$, and 2) the minimum of the expected utilities in all battlefields $U = \min_{i \in [n]} \mathbb{E}[u_i(\cdot)]$.

Across the variants of the two-level Blotto game we analyze, players choose mixed strategies over their subgame actions. The use of randomization at the allocation stage depends on whether the model is discrete or continuous. In the discrete case, the allocating players (the maximizing player in the one-sided case and both players in the two-sided case) randomize over integer allocation vectors. In the continuous case, players could also mix over allocations; yet we show that in some settings, equilibrium can be guaranteed even when players commit to a deterministic (pure) allocation. This motivates our focus on pure strategies at the allocation stage, which offer simpler representation and tractable computation. The various strategy representations are summarized in the appendix. Finally, a strategy pair $(x^*, y^*) \in X \times Y$ is a Nash equilibrium (NE) if $U(x^*, y) \leq U(x^*, y^*) \leq U(x, y^*)$ for all $x \in X, y \in Y$, where X and Y are the strategy spaces of Player 1 and 2, respectively.

Minimax theorems play a central role in our analysis, particularly in establishing equilibrium existence. In fact, the existence of a minimax theorem for a setting is equivalent to equilibrium existence (Proposition 1). We rely on Sion’s minimax theorem (Sion 1958) in several of our proofs and also examine its applicability across various settings. In the setting where only Player 2 allocates a continuum of soldiers across battlefields and individual subgame payoffs are aggregated via the minimum operator, we apply a variant known as the Kneser-Fan minimax theorem (Sion 1958).

Proposition 1. *Let $u : X \times Y \rightarrow \mathbb{R}$ be an arbitrary function on arbitrary sets X, Y . Then $\max_{x \in X} \inf_{y \in Y} u(x, y) = \min_{y \in Y} \sup_{x \in X} u(x, y)$ if and only if there exists a NE in a two-player zero-sum game with action spaces X and Y and utility function u .*

Theorem 1 (Sion’s minimax theorem). *Let X and Y be nonempty convex and compact subsets of two linear topolog-*

ical spaces, and $f : X \times Y \rightarrow \mathbb{R}$ be a function that is upper semicontinuous and quasiconcave in the first variable and lower semicontinuous and quasiconvex in the second. Then $\min_{y \in Y} \max_{x \in X} f(x, y) = \max_{x \in X} \min_{y \in Y} f(x, y)$.

Definition 1 (Concavelike, Convexlike, and Concave-convexlike). Consider two spaces X and Y , and let f be a function f on $X \times Y$. We say f is concavelike in X if for every $x_1, x_2 \in X$ and $0 \leq t \leq 1$, there is an $x \in X$ such that $tf(x_1, y) + (1-t)f(x_2, y) \leq f(x, y) \forall y \in Y$, and f is convexlike in Y if for every $y_1, y_2 \in Y$ and $0 \leq t \leq 1$, there is a $y \in Y$ such that $tf(x, y_1) + (1-t)f(x, y_2) \geq f(x, y) \forall x \in X$. Finally, f is called concave-convexlike if it is concavelike in X and convexlike in Y .

Theorem 2 (Kneser-Fan minimax theorem). Let X be compact, Y any space, and f a function on $X \times Y$ that is concave-convexlike. If $f(x, y)$ is upper semicontinuous in x for each y , then $\sup_{x \in X} \inf_{y \in Y} f = \inf_{y \in Y} \sup_{x \in X} f$.

3 Colonel Blotto with Discrete Soldiers

In the discrete two-level Blotto game, players allocate a finite number of indivisible soldiers across battlefields and randomize over both soldier assignments and subgame strategies. Let $k^j = (k_1^j, \dots, k_n^j) \in \mathcal{K}^j$ be player j 's discrete allocation vector. Then the payoff in battlefield i is given by $u_i(\alpha_i^1, \alpha_i^2, k_i^1, k_i^2)$. In our two-level formulation, each player j adopts a two-level strategy composed of 1) a probability distribution γ^j over soldier assignments, and 2) for each battlefield i and each soldier count k_i^j , a mixed strategy $\delta_{i, k_i^j}^j$ over actions $\alpha_i^j \in \mathcal{A}_i^j$ in subgame G_i . The action set in each battlefield is independent of the number of soldiers allocated to that battlefield (this may be relaxed easily). Since the number of possible assignments is finite, each subgame can be viewed as a Bayesian game in which player j has types $\{0, \dots, m^j\}$, each occurring with probability equal to the chance that the player assigns that number of soldiers to the battlefield under γ^j .

3.1 Sum Aggregator

The overall utility under the sum aggregator is given by

$$U(\delta, \gamma) = \sum_{i \in [n]} \sum_{k^1 \in \mathcal{K}^1} \sum_{k^2 \in \mathcal{K}^2} \gamma^1(k^1) \gamma^2(k^2) \times \sum_{\alpha_i^1 \in \mathcal{A}_i^1} \sum_{\alpha_i^2 \in \mathcal{A}_i^2} \delta_{i, k_i^1}^1(\alpha_i^1) \delta_{i, k_i^2}^2(\alpha_i^2) u_i(\alpha_i^1, \alpha_i^2, k_i^1, k_i^2).$$

In this setting, we can formulate the two-level Blotto game as a linear program. For each battlefield i , and given a fixed distribution over k_i^1 and k_i^2 , the optimal subgame strategy corresponds to solving a Bayesian extensive-form game where nature randomizes over the $(m^1 + 1) \cdot (m^2 + 1)$ possible types. We model this randomization in two sequential phases – first determining Player 1's probabilities, then Player 2's – thereby inducing a product distribution over the joint types. By translating the type assignment into two stages, we can assign the first phase to Player 1 and the second phase to Player 2 while retaining perfect recall.

More precisely, for player $j \in \{1, 2\}$, we define a flow polytope Γ^j with two types of variables. First, the variable $x_{i, k_i^j, \alpha_i^j}^j$ denotes the probability that player j allocates k_i^j soldiers to battlefield i and plays action α_i^j in subgame G_i . Second, the flow variable $h_{i, a, b}^j$ denotes the flow through battlefield i for player j : it represents the transition from a remaining soldiers to b remaining soldiers after allocating $a - b$ soldiers to battlefield i . The polytope Γ^j is defined by the following constraints:

$$\Gamma^j = \left\{ h_{i, a, b}^j, x_{i, k_i^j, \alpha_i^j}^j \in [0, 1] \text{ such that:} \right. \\ h_{0, m^j, m^j}^j = 1, \quad h_{n-1, 0, 0}^j = 1, \\ h_{0, a, b}^j = 0 \quad \forall a, b \in [0, m^j], (a, b) \neq (m^j, m^j), \\ \sum_{a=c}^{m^j} h_{i-1, a, c}^j = \sum_{b=0}^c h_{i, c, b}^j \quad \forall i \in [1, n-1], c \in [0, m^j], \\ \sum_{\alpha_i^j \in \mathcal{A}_i^j} x_{i, k_i^j, \alpha_i^j}^j = \sum_{k_i^j \leq r \leq m^j} h_{i, r, r-k_i^j}^j \quad \forall i \in [1, n-1], \\ \left. k_i^j \in [0, m^j] \right\}.$$

By encoding strategies with a flow polytope, we obtain a polynomial-size representation of strategies that is payoff equivalent to the original strategy space $\Delta(\mathcal{K}^j) \times \prod_{i \in [n]} \Delta(\mathcal{A}_i^j)$. This flow polytope is a generalization of the layered graph approach of Behnezhad et al. (2023), with the addition of subgames.

Theorem 3 (Kuhn's theorem under two-sided sum aggregator). Consider a discrete two-level Blotto game with sum aggregator, and suppose $\gamma^j \in \Delta(\mathcal{K}^j)$, $\delta_i^j \in \Delta(\mathcal{A}_i^j)$, $j \in \{1, 2\}$. Then the overall payoff is bilinear in the strategies x^1 and x^2 . Specifically, $U = \sum_i \sum_{\alpha_i^1} \sum_{\alpha_i^2} \sum_{k_i^1} \sum_{k_i^2} u_i(\alpha_i^1, \alpha_i^2, k_i^1, k_i^2) \cdot x_{i, k_i^1, \alpha_i^1}^1 \cdot x_{i, k_i^2, \alpha_i^2}^2$.

Theorem 3 implies that rather than optimizing over $\Delta(\mathcal{K}^j) \times \prod_{i \in [n]} \Delta(\mathcal{A}_i^j)$, we can optimize over a flow polytope Γ^j . Computing a NE may be expressed as a bilinear saddle point problem $\max_{x^2 \in \Gamma^2} \min_{x^1 \in \Gamma^1} U$. Since Γ^1 and Γ^2 are compact and convex sets, Sion's minimax theorem holds. Then, by Proposition 1, we have:

Corollary 1. Under the sum-aggregator, the discrete two-level Blotto game admits a NE.

We propose two methods to solve the bilinear saddle point problem. The first uses a common trick of taking the dual of the inner optimization problem, converting the min-max problem into a single minimization problem that can be solved by a linear program (Behnezhad et al. 2023). For brevity, details of the LP are deferred to the appendix. The second method is based on online learning and self-play. In the classic setting of zero-sum matrix games, each player's strategy space is the probability simplex, and it is known that the recommendations given by the online learners converge

on average to a NE (Hart and Mas-Colell 2000; Roughgarden 2010). This is readily adapted to imperfect information extensive-form games by defining regret minimizers over the *treeplex* (Zinkevich et al. 2007). In our setting, the regret minimizers are over the flow polytopes Γ^1 and Γ^2 , which may be done using scaled extensions (Farina et al. 2019) or kernel-based approaches (Farina et al. 2022; Takimoto and Warmuth 2003). For brevity, we defer details to the appendix, but remark that the regret minimizers require $\mathcal{O}(n \cdot \max_{i,j} \{m^j(m^j + |\mathcal{A}_i^j|\})\})$ space (i.e., *linear* in the size of the flow polytope), $\mathcal{O}(n \cdot \max_{i,j} \{m^j(m^j + |\mathcal{A}_i^j|\}) + \max_i \{m_i^1 m_i^2 |\mathcal{A}_i^1| |\mathcal{A}_i^2|\})$ time-per-iteration and incurs total regret at a rate of $\mathcal{O}(\sqrt{T})$, depending on the implementation details. Our methods readily extend to settings where each subgame is a perfect recall extensive-form game. This is done by reformulating the strategy space in each G_i by the treeplex (Von Stengel 1996). Our results here extend trivially to the one-sided case.

Proposition 2. *Under the sum aggregator, a NE of the discrete two-level Blotto game is polynomial-time computable.*

Remark 1. *Note that all results extend directly to the one-sided setting, in which the minimizing player does not participate in the soldier allocation.*

3.2 Min Aggregator

When the overall utility is defined as the minimum across all battlefields, the aggregate utility is given by

$$U(\delta, \gamma) = \min_{i \in [n]} \sum_{k^1 \in \mathcal{K}^1} \sum_{k^2 \in \mathcal{K}^2} \gamma^1(k^1) \gamma^2(k^2) \times \left(\sum_{\alpha_i^1 \in \mathcal{A}_i^1} \sum_{\alpha_i^2 \in \mathcal{A}_i^2} \delta_{i,k^1}^1(\alpha_i^1) \delta_{i,k^2}^2(\alpha_i^2) u_i(\alpha_i^1, \alpha_i^2, k_i^1, k_i^2) \right).$$

Unlike with the sum aggregator, when both players allocate soldiers to battlefields, the two-level Blotto game may not admit a NE even when battlefield utilities are continuous.

Theorem 4. *Consider a two-sided discrete two-level Blotto game under the min aggregator. Even if all battlefield utilities are continuous, a NE may not exist. Nevertheless, both the max-min and min-max values are well-defined and are attained.*

This motivates our study of the one-sided model, in which only the maximizing player allocates soldiers to battlefields, then both players engage in independent zero-sum subgames. This variant captures common security scenarios where for example, only the defender allocates resources, and guarantees equilibrium existence while admitting efficient computation. In this setting, the total utility is

$$U(\delta, \gamma) = \min_{i \in [n]} \sum_{k^2 \in \mathcal{K}^2} \gamma^2(k^2) \times \left(\sum_{\alpha_i^1 \in \mathcal{A}_i^1} \sum_{\alpha_i^2 \in \mathcal{A}_i^2} \delta_i^1(\alpha_i^1) \delta_{i,k^2}^2(\alpha_i^2) u_i(\alpha_i^1, \alpha_i^2, k_i^2) \right).$$

This formulation admits a tractable NE.

Theorem 5 (Kuhn’s theorem under one-sided min aggregator). *Consider a one-sided discrete two-level Blotto game with min aggregator, and suppose $\gamma^2 \in \Delta(\mathcal{K}^2)$, $\delta_i^j \in \Delta(\mathcal{A}_i^j)$, $j \in \{1, 2\}$. Then the overall payoff is convex and bilinear in p and x^2 , where $p \in \Delta(n)$ is a convexification variable and x^2 is the treeplex strategy for Player 2. Specifically, $U = \min_p \sum_i \sum_{\alpha_i^2} \sum_{k_i^2} p_i \mathbb{E}_{\alpha_i^1 \sim \delta_i^1} [u_i(\alpha_i^1, \alpha_i^2, k_i^2)] x_{i,k_i^2,\alpha_i^2}^2$.*

Using the bilinear formulation from Theorem 5 we can now express the min-max problem as $\min_{\delta^1} \max_{x^2} \min_p \sum_i \sum_{\alpha_i^2} \sum_{k_i^2} p_i x_{i,k_i^2,\alpha_i^2}^2 \mathbb{E}[u_i(\alpha_i^1, \alpha_i^2, k_i^2)]$. Exploiting the bilinear structure of the payoff, we invoke Sion’s minimax theorem to interchange the inner maximization and minimization operators. By then combining the two minimizations, we define a sequence-form strategy over a two-level game for Player 1, in which he first chooses a battlefield and then plays an action in it, using $y^1(i, \alpha_i^1)$ to denote the product $p_i \cdot \delta_i^1(\alpha_i^1)$. Under this reformulation, we recover the game’s max-min structure, which yields the following result.

Theorem 6. *Consider a one-sided discrete two-level Blotto game with min aggregator. Then the minimax theorem holds. Specifically, $\min_{\delta^1} \max_{\delta^2, \gamma^2} \min_i \mathbb{E}[u_i(\alpha_i^1, \alpha_i^2, k_i^2)] = \max_{\delta^2, \gamma^2} \min_{\delta^1} \min_i \mathbb{E}[u_i(\alpha_i^1, \alpha_i^2, k_i^2)]$.*

The flow polytope representation of Player 2’s strategy introduces a set of constraints captured by the polytope Γ^2 . Analogously, the sequence-form product representation of Player 1’s strategy induces the constraints defined over the following polytope

$$\mathcal{P} = \left\{ \begin{array}{l} y^1(\emptyset) = 1, \\ \sum_{i \in [n]} y^1(i) = y^1(\emptyset), \\ \sum_{\alpha_i^1 \in \mathcal{A}_i^1} y^1(i, \alpha_i^1) = y^1(i). \end{array} \right. \left. \begin{array}{l} y^1(i) \geq 0 \\ y^1(i, \alpha_i^1) \geq 0 \end{array} \right\}$$

Therefore, a NE can be computed by solving a LP over the polytopes Γ^2 and \mathcal{P} , obtained by dualizing either player’s best-response program. Since both Γ^2 and \mathcal{P} have polynomial-size representations, we have:

Corollary 2. *In the one-sided model, a NE of the two-level Blotto game with min aggregator exists and can be computed in polynomial time.*

4 Colonel Blotto with Continuous Soldiers

In the *continuous two-level Blotto game*, we focus on equilibria where players choose deterministic (pure) soldier allocations and may randomize over subgame strategies. Unlike the discrete setting where equilibrium existence requires randomization at the resource allocation level, we show that when soldier allocations are continuous, deterministic allocations suffice to guarantee the existence of equilibrium in some cases. This yields several practical benefits: such equilibria admit exact representation, are computationally tractable, and, in applications such as security, can be easier to deploy when randomization at the allocation level is impractical or undesirable. We view this setting as a natural starting point for analyzing randomized allocation strategies, whose equilibria may lack finite representations and

are more difficult to characterize. We leave the analysis of such strategies to future work.

In this setting, players allocate any real-valued fraction of soldiers to each battlefield. Let $\sigma^j = (\sigma_1^j, \dots, \sigma_n^j) \in \Sigma^j$ be player j 's continuous allocation vector. Then the payoff in battlefield i is $u_i(\alpha_i^1, \alpha_i^2, \sigma_i^1, \sigma_i^2)$. Considering both the sum and min aggregators, we show that, in the two-sided setting, a partially pure Nash equilibrium (PPNE) – i.e. one in which players use pure soldier allocation strategies and mixed subgame strategies – may fail to exist under either aggregator, even when the utility functions in all battlefields are continuous in the allocation variables.

Theorem 7. *Consider a two-sided continuous two-level Blotto game under the sum or min aggregator. Even if battlefield utilities are continuous in σ , a PPNE may not exist.*

We proceed to analyze the one-sided setting. We show that under the sum aggregator, a PPNE may still fail to exist, whereas under the min aggregator, a PPNE is guaranteed to exist. We explain this divergence at the end of the section.

4.1 Sum Aggregator

In the continuous two-sided model, each player j deterministically selects a soldier allocation $\sigma^j \in \Sigma^j$, and employs a mixed strategy $\delta_{i, \sigma_i^j}^j$ over actions $\alpha_i^j \in \mathcal{A}_i^j$ in subgame G_i . We can write the overall in the two-sided model as: $U(\delta, \sigma) = \sum_i \left(\sum_{\alpha_i^1} \sum_{\alpha_i^2} \delta_{i, \sigma_i^1}^1(\alpha_i^1) \delta_{i, \sigma_i^2}^2(\alpha_i^2) u_i(\alpha_i^1, \alpha_i^2, \sigma_i^1, \sigma_i^2) \right)$.

The total payoff is non-convex non-concave in the minimizing and maximizing player's strategies, respectively. This remains true even in the one-sided allocation model. Consequently, a PPNE may still fail to exist, even when each u_i is continuous in the soldier allocation variables.

Proposition 3. *Even in the one-sided case, a continuous two-level Blotto with continuous battlefield utilities may not admit a PPNE under the sum aggregator. The max-min and min-max values are attained but do not coincide.*

Although equilibrium existence cannot be guaranteed in the one-sided model under general continuous utilities, we show that if each battlefield payoff is linear in the maximizing player's allocation, namely $u_i(\sigma_i^2) = c_i \sigma_i^2$, $c_i > 0$, a PPNE is guaranteed to exist. In fact, with linear battlefield utilities, the game can be cast as a bilinear saddle point problem by absorbing the soldier allocation variable σ^2 and the subgame distribution over actions δ^2 into a single sequence-form variable $y_{\sigma^2}^2$ via the treplex construction. The product variable $y_{\sigma^2}^2$ is defined over a polytope \mathcal{Q} that is represented by the following constraints:

$$\mathcal{Q} = \left\{ \begin{array}{l} y_{\sigma^2}^2(i) \geq 0 \\ y_{\sigma^2}^2(i, \alpha_i^2) \geq 0 \end{array} \middle| \begin{array}{l} y^2(\emptyset) = m^2, \\ \sum_{i \in [n]} y_{\sigma^2}^2(i) = y^2(\emptyset), \\ \sum_{\alpha_i^2 \in \mathcal{A}_i^2} y_{\sigma^2}^2(i, \alpha_i^2) = y_{\sigma^2}^2(i). \end{array} \right\}$$

The resulting payoff is bilinear in the variables $y_{\sigma^2}^2$ and δ^1 , which allows us to invoke the minimax theorem and guarantee the existence of a PPNE.

Theorem 8. *Consider a one-sided continuous two-level Blotto game with sum aggregator, and suppose $\sigma^2 \in \Sigma^2$, $\delta_i^j \in \Delta(\mathcal{A}_i^j)$, $j \in \{1, 2\}$. Assume that in each battlefield, the utility is linear in the maximizing player's allocation, i.e. $u_i = c_i \cdot \sigma_i^2$ for some constant $c_i > 0$. Then the minimax theorem applies.*

It follows from Theorem 8 that the continuous two-level Blotto game with sum aggregator and linear utilities admits an equilibrium, which can be computed by solving a linear program defined over the polytope \mathcal{Q} . Since \mathcal{Q} has a polynomial size representation, we have:

Corollary 3. *When battlefield utilities are linear in the maximizing player's allocation, a PPNE of the one-sided continuous two-level Blotto game with sum aggregator exists and can be computed in polynomial time.*

4.2 Min aggregator

In the two-sided model, when the overall payoff is defined as the minimum of the expected utilities across battlefields, it can be written as $U(\delta, \sigma) = \min_{i \in [n]} \left(\sum_{\alpha_i^1} \sum_{\alpha_i^2} \delta_{i, \sigma_i^1}^1(\alpha_i^1) \delta_{i, \sigma_i^2}^2(\alpha_i^2) u_i(\alpha_i^1, \alpha_i^2, \sigma_i^1, \sigma_i^2) \right)$.

Since a PPNE may fail to exist in the two-sided setting (Theorem 7), we focus on the one-sided model. We show that equilibrium non-existence arises only if the utility function in every battlefield i is discontinuous in σ_i^2 .

Theorem 9. *Consider a one-sided continuous two-level Blotto game with min aggregator. Then a PPNE may not exist when battlefield utilities are discontinuous.*

However, when the payoff in each battlefield is a continuous function of the soldier allocation, the continuous two-level Blotto game with min aggregator admits a saddle point in the one-sided formulation: the max-min and min-max values coincide and are attained, so a PPNE exists. We establish this by invoking the Kneser-Fan minimax theorem. Suppose that only Player 2 allocates soldiers to battlefields, and both players employ mixed strategies to engage in the battlefield subgames. Then, the min-max formulation of the continuous two-level Blotto game is:

$$\begin{aligned} & \min_{\delta^1} \max_{\delta^2, \sigma^2} \min_{i \in [n]} \mathbb{E} [u_i(\alpha_i^1, \alpha_i^2, \sigma_i^2)] \\ & = \min_{\delta^1} \max_{\delta^2, \sigma^2} \min_{i \in [n]} \sum_{\alpha_i^1} \sum_{\alpha_i^2} \delta_i^1(\alpha_i^1) \delta_{i, \sigma_i^2}^2(\alpha_i^2) u_i(\alpha_i^1, \alpha_i^2, \sigma_i^2). \end{aligned}$$

Ideally, we would like to invoke Sion's minimax theorem to interchange the outer minimization and maximization, thereby recovering the equivalent max-min formulation of the game and establishing the minimax identity. However, the inner objective function cannot be made simultaneously quasiconvex in the minimizing player's strategy and quasiconcave in the maximizing player's strategy under any re-ordering of the operators and/or regrouping of the variables. Consequently, Sion's conditions are not met. Therefore, we appeal to the Kneser-Fan minimax theorem (Theorem 2). Taking the dual of the inner minimization problem over $i \in [n]$, allows us to obtain the following result:

Theorem 10. Consider a one-sided continuous two-level Blotto game with continuous battlefield utilities. Then the minimax theorem holds. Specifically, $\min_{\delta^1} \max_{\delta^2, \sigma^2} \min_i \mathbb{E}[u_i(\alpha_i^1, \alpha_i^2, \sigma_i^2)] = \max_{\delta^2, \sigma^2} \min_{\delta^1} \min_i \mathbb{E}[u_i(\alpha_i^1, \alpha_i^2, \sigma_i^2)]$.

Remark 2. This approach does not extend to any of the other settings we analyzed in the continuous domain, which is consistent with the non-existence results established in Theorem 7 and Proposition 3. In particular, (1) in the one-sided continuous setting with sum aggregator, the payoff fails to be concavelike in the maximizing player's strategy, which immediately extends to the sum-aggregated two-sided setting; whereas (2) in the two-sided continuous setting with min aggregator, convexlikeness of the payoff in the minimizing player's strategy cannot be guaranteed.

For the special case where battlefield utilities are linear in σ^2 , we obtain a sequence-form representation of the strategies of both players defined over the polytopes \mathcal{P} and \mathcal{Q} . Hence, we have the following theorem:

Theorem 11. Consider a one-sided continuous two-level Blotto game under the min aggregator, and suppose $\sigma^2 \in \Sigma^2$, $\delta_i^j \in \Delta(\mathcal{A}_i^j)$, $j \in \{1, 2\}$. Assume that in every battlefield, the utility is linear in the maximizing player's allocation, i.e. $u_i = c_i \cdot \sigma_i^2$ for some constant $c_i > 0$. Then we can compute a PPNE of the game by solving a linear program.

Corollary 4. When battlefield utilities are linear in the soldier allocation of the maximizing player, the one-sided continuous two-level Blotto game with min aggregator admits a PPNE that can be computed in polynomial time.

Computing the max-min strategy In this paragraph, we focus on computing the max-min strategy for Player 2 in the one-sided continuous two-level Blotto game with min aggregator. We show that it can be reformulated as a maximization over the minimum Nash value across battlefields and hence it can be computed using a subgradient ascent algorithm.

Proposition 4. Consider a one-sided continuous two-level Blotto game with min aggregator, and suppose $\sigma^2 \in \Sigma^2$, $\delta_i^j \in \Delta(\mathcal{A}_i^j)$, $j \in \{1, 2\}$. Then the maxmin problem can be formulated as a maximization problem over the minimum Nash value across battlefields. Specifically, $\max_{\delta^2, \sigma^2} \min_{\delta^1} \min_i \mathbb{E}[u_i(\alpha_i^1, \alpha_i^2, \sigma_i^2)] = \max_{\sigma^2} \min_{i \in [n]} v_i^*(\sigma_i^2)$.

Let $V(\sigma^2) := \min_i v_i^*(\sigma_i^2)$ denote the objective function of this maximization problem. We show that this function is quasiconcave. Notice that, contrary to the sum, the min aggregator preserves quasiconcavity. Hence, to achieve overall quasiconcavity of $V(\sigma^2)$, it suffices to have that the individual terms $v_i^*(\sigma)$ are quasiconcave.

Lemma 1. If for each battlefield $i \in [n]$, $u_i(\alpha_i^1, \alpha_i^2, \sigma_i^2)$ is increasing in σ_i^2 for all $\alpha_i^1 \in \mathcal{A}_i^1$ and $\alpha_i^2 \in \mathcal{A}_i^2$, then the aggregate Nash value function $V(\sigma^2)$ is quasiconcave.

Since $V(\sigma^2)$ is quasiconcave on the compact convex set of soldier assignments Σ^2 , then finding the optimal $\sigma^{2*} =$

$\arg \max_{\sigma^2 \in \Sigma^2} V(\sigma^2)$ reduces to maximizing a quasiconcave function over a convex set. Moreover, under the assumption that each battlefield utility function $u_i(\sigma^2)$ is Lipschitz continuous on the soldiers assignments simplex Σ^2 , not only does a global maximizer σ^{2*} exist, but one can actually compute it using a simple projected subgradient-ascent (PSA) method. Specifically, at each iteration t , we select a subgradient $g_t \in \partial V(\sigma_t^2)$, a step size $\eta_t > 0$ and perform the update $\sigma_{t+1}^2 = P_{\Sigma^2}(\sigma_t^2 + \eta_t g_t)$, where P_{Σ^2} denotes the projection onto Player 2's soldiers' simplex. For completeness, the full pseudocode is given in the appendix. Next, we derive a closed-form description of the subdifferential $\partial V(\sigma_t^2)$, which enables efficient computation of each g_t .

Proposition 5. Consider the two-player zero-sum subgame G_i on battlefield i . Suppose that for every pure profile $(\alpha_i^1, \alpha_i^2) \in \mathcal{A}_i^1 \times \mathcal{A}_i^2$, the payoff $u_i(\alpha_i^1, \alpha_i^2, \sigma_i^2)$ is Lipschitz continuous in σ_i^2 . Let $(\delta_i^{1*}, \delta_i^{2*})$ be any NE of G_i at a given σ_i^2 . Then the vector $g_i = \sum_{\alpha_i^1 \in \mathcal{A}_i^1} \sum_{\alpha_i^2 \in \mathcal{A}_i^2} \delta_i^{1*}(\alpha_i^1) \delta_i^{2*}(\alpha_i^2) z_{\alpha_i^1, \alpha_i^2}$, $z_{\alpha_i^1, \alpha_i^2} \in \partial_{\sigma_i^2} u_i(\alpha_i^1, \alpha_i^2, \sigma_i^2)$, is a subgradient of $v_i^*(\sigma^2)$.

Corollary 5. Under the assumptions of Proposition 5, V is Lipschitz on Σ^2 and its subdifferentials $\partial V(\sigma^2)$ lie in $\text{Conv}\{\partial v_i^*(\sigma^2) : i \in \arg \min_{j \in [n]} v_j^*(\sigma^2)\}$.

Finally, to obtain the optimal soldier allocation σ^{2*} for Player 2, we run the PSA algorithm. Each iteration t requires a subgradient $g_t \in \partial V(\sigma_t^2)$. By Corollary 5, it suffices to pick any active battlefield $i^* \in \arg \min_i v_i^*(\sigma_t^2)$ and use its subgradient. Proposition 5 then provides a concrete "Nash subgradient" for battlefield i^* : choose for each action pair (α_i^1, α_i^2) a subgradient $z_{\alpha_i^1, \alpha_i^2} \in \partial_{\sigma_i^2} u_{i^*}(\alpha_i^1, \alpha_i^2, \sigma_i^2)$, and set $g_t = \sum_{\alpha_i^1 \in \mathcal{A}_i^1} \sum_{\alpha_i^2 \in \mathcal{A}_i^2} \delta_i^{1*}(\alpha_i^1) \delta_i^{2*}(\alpha_i^2) z_{\alpha_i^1, \alpha_i^2} \in \partial v_{i^*}^*(\sigma_t^2)$. This g_t is exactly the update direction used in the PSA algorithm.

Convergence Guarantees By Lemma 1 and Corollary 5, our aggregate Nash-value function V is quasiconcave and Lipschitz continuous on the soldiers simplex. Hence, the projected subgradient-ascent iterates $\{\sigma_t^2\}$ satisfy the general basic-inequality conditions (H1)-(H2) of (Hu, Li, and Yu 2020), which guarantee global convergence of any sequence with a suitably chosen step size rule. Moreover, they establish in Theorems 3.3-3.5 that, under an additional *weak sharp minima of Hölderian order* assumption and upper-bounded noise, such a sequence converges at a linear rate when the Hölder exponent equals 1. In our maximization context, this translates into a *weak sharp maxima* condition:

Definition 2 (Weak Sharp Maxima). There exist constants $\epsilon > 0$, $p \in (0, 1]$ such that $\forall \sigma^2 \in \Sigma^2$, $V(\sigma^{2*}) - V(\sigma^2) \geq \epsilon \|\sigma^2 - \sigma^{2*}\|^p$, where σ^{2*} is the unique maximizer of V .

When battlefield utilities are affine functions of σ^2 , the induced Nash value $v_i^*(\sigma_i^2)$ remains an affine function of σ_i^2 . It follows that the aggregate function $V = \min_i v_i^*$ is concave on the soldiers simplex Σ^2 . Applying the subgradient inequality at the maximizer $\sigma^{2*} \in \arg \max_{\Sigma^2} V$ then yields the following sharp-max error bound for V .

Proposition 6. If every battlefield utility is affine in Player 2's soldier allocation, i.e. $u_i(\sigma^2) = c_i \sigma_i^2 + d_i$ with $c_i >$




Figure 1: Convergence of the PSA algorithm in the one-sided min-aggregated setting with affine (left), quadratic (middle), and security-inspired (right) battlefield utilities.

0, then V satisfies the weak sharp-maxima condition with exponent $p = 1$.

It follows immediately from Theorem 3.3 of (Hu, Li, and Yu 2020) that the subgradient ascent iterates converge linearly to the optimal σ^{2*} .

5 Empirical Evaluation

Discrete two-sided with sum-aggregator We first evaluate in the most basic setting of discrete soldiers under the sum aggregator (Section 3.1). We consider (i) the LP-based approach similar to Behnezhad et al. (2023) and (ii) our approach based on online learning. The $i \in [n]$ battlefield has value i , such that winning (resp. losing) it gives $+i$ (resp. $-i$) reward. For consistency, we re-normalize values of battlefields such that they sum to 1. The i -th battlefield is won by player 1 with probability $k_i^1 / (k_i^1 + k_i^2)$, with a random winner selected if no soldiers are allocated. For each subgame, players decide whether to *double* their stakes. If just one player doubles, payoffs and losses for that battlefield are doubled; if both players double, they are quadrupled.

We report running times for both methods. For the former, we use the default configuration for Gurobi (Gurobi Optimization, LLC 2024). For the latter, we adopt the Regret Matching-Plus (RM+) algorithm (Tammelin 2014) and construct Γ^j using scaled extensions (Farina et al. 2019). We declare convergence when the saddle-point gap is less than 0.002. More details on experimental setup are reported in the appendix. The results are shown in Table 2. We observe that when the game is small, the LP finds the NE easily. However, as the size of the game grows, LPs slow down dramatically and online learning approaches begin to shine. In fact, for larger instances Gurobi’s sometimes fails to converge because of numerical issues or simply going out-of-memory.² In contrast, methods based on online learning scale better.

Continuous one-sided with min-aggregator We evaluate the PSA algorithm on randomly generated and security-

²In our experiments, it is unclear if Barrier or Primal/Dual simplexes are superior. For larger games, it would appear as though simplex performs better in practice, while barrier tends to lead to non-convergence.

n	m^1	m^2	$ \Gamma^1 , \Gamma^2 $	LP [s]	RM+ [s]
30	100	50	85, 49	91	0.02
35	125	70	281k, 92k	$4.8 \cdot 10^3$	121
40	150	100	1M, 262k	$5.4 \cdot 10^3$	525
50	200	100	12.6M, 2.05M	NA ($> 2.6 \cdot 10^4$)	$9.5 \cdot 10^3$

Table 2: Runtime for discrete soldiers with sum aggregators.

inspired battlefield utilities. For random instances, we consider (i) affine utilities $u_i(\sigma^2) = c_i \sigma^2 + d_i$, and (ii) quadratic utilities $u_i(\sigma^2) = b_i (\sigma^2)^2 + c_i \sigma^2 + d_i$, with $b_i, c_i, d_i \sim \text{Uniform}[0, 100]$. Player 2 is given 20 soldiers to allocate across 5 battlefields. We vary players’ action spaces in $\{20, 30, 40\}$ and generate 10 independent instances per game size. In the first 2 plots of Figure 1, we show the aggregated normalized value \bar{V}^t and its standard error $\text{SE}(t) = \sqrt{\frac{s_t}{\ell}}$, $s_t = \frac{1}{\ell-1} \sum_{k=1}^{\ell} (V_k^t - \bar{V}^t)^2$, across $\ell = 10$ instances, where V_k^t denotes the normalized objective value on instance k at iteration t . Under both utility functions, all three sizes reach near-optimality with very low dispersion. Next, we consider a real-world scenario inspired by security applications. We model a one-sided min-aggregated continuous two-level Blotto game where Player 2 allocates 10 soldiers across three battlefield, each featuring a two-player zero-sum security subgame where Player 2 is the defender’s and Player 1 the attacker. We follow the framework of (Kreuer et al. 2025). More details on the experimental setup are provided in the appendix. The rightmost plot of Figure 1 confirms the algorithm’s robustness in this practical setting.

6 Conclusion

In this paper, we introduced and analyzed a two-level variant of the Colonel Blotto game, where players allocate resources across battlefields then engage in parametrized subgames. We studied multiple modeling dimensions – soldier types, payoff aggregation methods, and strategic settings – and established equilibrium existence, provided solution algorithms, and validated their practicality through experiments.

Acknowledgments

This research was supported by the Office of Naval Research awards N00014-22-1-2530 and N00014-23-1-2374, the National Science Foundation awards IIS-2147361 and IIS-2238960, the Ministry of Education, Singapore, under the Academic Research Fund Tier 1 (FY2025) and by the National University of Singapore, under the Start-Up Grant Scheme. We thank Noah Krever for help with setting up the security-inspired subgames.

References

- Ahmadinejad, A.; Dehghani, S.; Hajiaghayi, M.; Lucier, B.; Mahini, H.; and Seddighin, S. 2019. From duels to battlefields: Computing equilibria of blotto and other games. *Mathematics of Operations Research*, 44(4): 1304–1325.
- Arad, A.; and Rubinstein, A. 2012. Multi-dimensional iterative reasoning in action: The case of the Colonel Blotto game. *Journal of Economic Behavior & Organization*, 84(2): 571–585.
- Behnezhad, S.; Dehghani, S.; Derakhshan, M.; Hajiaghayi, M.; and Seddighin, S. 2023. Fast and simple solutions of Blotto games. *Operations Research*, 71(2): 506–516.
- Boix-Adserà, E.; Edelman, B. L.; and Jayanti, S. 2020. The multiplayer colonel blotto game. In *Proceedings of the 21st ACM Conference on Economics and Computation*, 47–48.
- Borel, E.; and Ville, J. 1991. Application de la théorie des probabilités aux jeux de hasard, original edition by Gauthier-Villars, Paris, 1938; reprinted at the end of *Théorie mathématique du bridgée la portée de tous*, by E. Borel & A. Chéron, Editions Jacques Gabay, Paris.
- Che, Y.-K.; and Gale, I. L. 2008. Caps on political lobbying. In *40 Years of Research on Rent Seeking 2*, 337–345. Springer.
- Chowdhury, S. M.; Kovenock, D.; and Sheremeta, R. M. 2013. An experimental investigation of Colonel Blotto games. *Economic Theory*, 52: 833–861.
- Dziubiński, M. 2013. Non-symmetric discrete General Lotto games. *International Journal of Game Theory*, 42: 801–833.
- Farina, G.; Lee, C.-W.; Luo, H.; and Kroer, C. 2022. Kernelized multiplicative weights for 0/1-polyhedral games: Bridging the gap between learning in extensive-form and normal-form games. In *International Conference on Machine Learning*, 6337–6357. PMLR.
- Farina, G.; Ling, C. K.; Fang, F.; and Sandholm, T. 2019. Efficient regret minimization algorithm for extensive-form correlated equilibrium. *Advances in Neural Information Processing Systems*, 32.
- Gross, O.; and Wagner, R. 1950. A continuous Colonel Blotto game. RAND Corporation. *Research Memorandum*, 408.
- Gurobi Optimization, LLC. 2024. *Gurobi Optimizer Reference Manual*. <https://www.gurobi.com>.
- Hart, S. 2008. Discrete colonel blotto and general lotto games. *International Journal of Game Theory*, 36(3-4): 441–460.
- Hart, S.; and Mas-Colell, A. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5): 1127–1150.
- Hausken, K. 2012. On the impossibility of deterrence in sequential colonel blotto games. *International Game Theory Review*, 14(02): 1250011.
- Hortala-Vallve, R.; and Llorente-Saguer, A. 2012. Pure strategy Nash equilibria in non-zero sum colonel Blotto games. *International Journal of Game Theory*, 41: 331–343.
- Hu, Y.; Li, J.; and Yu, C. K. W. 2020. Convergence rates of subgradient methods for quasi-convex optimization problems. *Computational Optimization and Applications*, 77(1): 183–212.
- Kovenock, D.; and Roberson, B. 2012a. Coalitional Colonel Blotto games with application to the economics of alliances. *Journal of Public Economic Theory*, 14(4): 653–676.
- Kovenock, D.; and Roberson, B. 2012b. Conflicts with multiple battlefields.
- Krever, N.; Černý, J.; Blanchard, M.; and Kroer, C. 2025. GUARD: Constructing Realistic Two-Player Matrix and Security Games for Benchmarking Game-Theoretic Algorithms. arXiv:2505.14547.
- Kvasov, D. 2007. Contests with limited resources. *Journal of Economic Theory*, 136(1): 738–748.
- Laslier, J.-F.; and Picard, N. 2002. Distributive politics and electoral competition. *Journal of Economic Theory*, 103(1): 106–130.
- Myerson, R. B. 1993. Incentives to cultivate favored minorities under alternative electoral systems. *American Political Science Review*, 87(4): 856–869.
- Powell, R. 2009. Sequential, nonzero-sum “Blotto”: Allocating defensive resources prior to attack. *Games and Economic Behavior*, 67(2): 611–615.
- Roberson, B. 2006. The colonel blotto game. *Economic Theory*, 29(1): 1–24.
- Roberson, B. 2011. *Allocation Games*. John Wiley, Ltd.
- Roberson, B.; and Kvasov, D. 2012. The non-constant-sum Colonel Blotto game. *Economic Theory*, 51: 397–433.
- Roughgarden, T. 2010. Algorithmic game theory. *Communications of the ACM*, 53(7): 78–86.
- Shubik, M.; and Weber, R. J. 1981. Systems defense games: Colonel Blotto, command and control. *Naval Research Logistics Quarterly*, 28(2): 281–287.
- Sion, M. 1958. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1): 171 – 176.
- Stephenson, D. 2024. Multi-battle contests over complementary battlefields. *Review of Economic Design*, 1–16.
- Takimoto, E.; and Warmuth, M. K. 2003. Path kernels and multiplicative updates. *Journal of Machine Learning Research*, 4(Oct): 773–818.
- Tammelin, O. 2014. Solving large imperfect information games using CFR+. *arXiv preprint arXiv:1407.5042*.
- Von Stengel, B. 1996. Efficient computation of behavior strategies. *Games and Economic Behavior*, 14(2): 220–246.

Vu, D. Q. 2020. *Models and solutions of strategic resource allocation problems: Approximate equilibrium and online learning in blotto games*. Ph.D. thesis, Sorbonne Universites, UPMC University of Paris 6.

Washburn, A. R. 2011. *TPZS Applications: Blotto Games*. John Wiley, Ltd.

Zinkevich, M.; Johanson, M.; Bowling, M.; and Piccione, C. 2007. Regret minimization in games with incomplete information. *Advances in neural information processing systems*, 20.