Stanisław Kaźmierowski

Supervised by dr hab. Marcin Dziubiński

SOLVING SUCCINCT GAMES



University of Warsaw
Faculty of Mathematics, Informatics and Mechanics
Warsaw, 2025

Abstract

In this thesis, we study the algorithmic aspects of games with succinct representation—formal models of strategic interactions whose compact descriptions encode exponentially large strategy spaces. Our main objective is to determine which subclasses of succinct games admit efficient equilibrium computation and how structural properties of the model affect computational tractability.

We consider four classes of games with succinct representation, each of which can be interpreted as a conflict with multiple battlefields. As a reference point, we take the two-player constant-sum Colonel Blotto games, in which each player distributes a fixed amount of resources across several battlefields, and the player allocating more to a battlefield wins it. The payoffs in the Colonel Blotto game are bilinear, i.e., a player's payoff is the sum of payoffs from the individual battlefields, and the game is constant-sum. As shown by Ahmadinejad et al. [2019] and Behnezhad et al. [2023], these properties allow for computation of Nash equilibria in these games in a time that is polynomial with respect to the model parameters, i.e., the number of battlefields and the players' resources. We treat the Colonel Blotto games as a benchmark for studying how deviations from the two-player setting, bilinearity, or the constant-sum property affect equilibrium computation.

The first class of models we study, conflicts with non-bilinear aggregation functions, departs from the classical Blotto formulation by violating bilinearity while preserving the constant-sum property. The second, the Colonel Blotto game with costs, introduces global cost terms that break both bilinearity and the zero-sum property. The third model, the Colonel Blotto game with flexible tie-breaking, departs from the constant-sum assumption while preserving bilinear payoffs. Finally, the attack and interception game on networks extends the two-player Blotto framework to a setting with multiple players and network-defined strategy spaces, where feasible strategies of the attacker correspond to all simple paths in the network.

Our results characterize how these structural departures affect the feasibility of equilibrium computation. We prove that computing a best response in non-bilinear conflicts cannot be done in polynomial time unless $\mathsf{P} = \mathsf{NP}$. Nonetheless, we show that equilibria can still be computed efficiently in practice, using a combination of symmetrization, clash-matrix evaluation, and the Double Oracle Algorithm. For the Colonel Blotto game with costs, we demonstrate that each instance is strategically equivalent to a zero-sum Colonel Blotto game with one additional battlefield, enabling polynomial-time equilibrium computation. For the flexible tie-breaking variant, we provide a partial characterization of the equilibrium strategies and identify an equilibrium that remains robust under perturbations of the tie-breaking rule. Finally, for attack and interception games on networks, we show that Nash equilibria can be computed in polynomial time by exploiting the structure of the underlying graph.

Our results have been published in the form of four articles:

(1) S. Kaźmierowski and M. Dziubiński. Efficient Method for Finding Optimal Strategies in Chopstick Auctions with Uniform Objects Values. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '24, pages

992–1000, Richland, SC, 2024. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9798400704864

(2) S. Kaźmierowski. Equilibria of the Colonel Blotto Games with Costs. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(13):13969–13976, Apr. 2025. doi: 10.1609/aaai.v39i13.33528. URL https://ojs.aaai.org/index.php/AAAI/article/view/33528

(3) C. Ewerhart and S. Kaźmierowski. An Equilibrium Analysis of the Arad-Rubinstein Game. *Journal of Economic Behavior & Organization*, 226:106696, 2024. ISSN 0167-2681. doi: https://doi.org/10.1016/j.jebo.2024.106696. URL https://www.sciencedirect.com/science/article/pii/S0167268124003020

(4) S. Kaźmierowski and M. Dziubiński. Computation of Nash Equilibria of Attack and Defense Games on Networks. In A. Deligkas and A. Filos-Ratsikas, editors, *Algorithmic Game Theory*, pages 3–21, Cham, 2023. Springer Nature Switzerland. ISBN 978-3-031-43254-5

Keywords: game theory, algorithmic game theory, Nash equilibrium, optimal strategy, constant-sum games, conflicts with multiple battlefields, succinct games, Colonel Blotto game

Streszczenie

W rozprawie przedstawiamy wyniki naszych badań nad algorytmicznymi aspektami gier o zwięzłej reprezentacji — modeli strategicznych, w których liczba strategii jest wykładnicza względem opisu modelu. Naszym celem jest ustalenie, dla których podklas gier ze zwięzłą reprezentacją możliwe jest efektywne obliczanie równowag Nasha oraz jak cechy strukturalne modeli wpływają na złożoność obliczeniową wyznaczania równowag.

Analizujemy cztery klasy gier o zwięzłej reprezentacji, które można interpretować jako konflikty z wieloma polami bitwy. Punktem odniesienia jest gra o sumie stałej Pułkownik Blotto, w której dwóch graczy rozdziela ograniczone zasoby pomiędzy pola bitwy, a każde pole wygrywa ten, kto przydzieli do niego więcej zasobów. W tej grze wypłaty są dwuliniowe — całkowita wypłata jest sumą wyników z poszczególnych pól oraz jest to gra o sumie stałej. Jak pokazali Ahmadinejad et al. [2019] i Behnezhad et al. [2023], własności te umożliwiają obliczanie równowag Nasha w czasie wielomianowym względem parametrów modelu, czyli liczby pól bitew oraz liczby zasobów obydwu graczy. Grę Pułkownik Blotto traktujemy jako punkt odniesienia do analizy, jak odstępstwa od dwóch graczy, dwuliniowości, lub własności stałej sumy wypłat wpływają na złożoność obliczania równowag.

Pierwsza z badanych klas, konflikty z niedwuliniowymi funkcjami agregacji, zachowuje własność sumy stałej, ale narusza dwuliniowość. W drugiej, grze Pułkownik Blotto z kosztami, koszty łamią zarówno dwuliniowość, jak i własność stałej sumy. Trzeci model, gra Pułkownik Blotto z elastycznym rozstrzyganiem remisów, odchodzi od założenia sumy stałej, zachowując jednak dwuliniowość wypłat. Ostatni z analizowanych modeli — gra ataku i przechwytywania na sieciach — rozszerza klasyczny schemat gry Pułkownik Blotto do środowiska z wieloma graczami, w którym zbiór strategii gracza atakującego to zbiór wszystkich prostych ścieżek w sieci.

Nasze wyniki pokazują, jak te różnice strukturalne wpływają na obliczeniową złożoność gier. W przypadku konfliktów z niedwuliniowymi funkcjami agregacji dowodzimy, że obliczenie najlepszej odpowiedzi nie jest możliwe w czasie wielomianowym, chyba że $\mathsf{P} = \mathsf{NP}$. Mimo to proponujemy metodę umożliwiającą efektywne w praktyce obliczanie równowag dzięki połączeniu symetryzacji, algorytmu macierzy starć i algorytmu dwóch orakli. W grze z kosztami pokazaliśmy, że każda instancja rozważanego modelu jest strategicznie równoważna klasycznej grze Pułkownik Blotto o sumie stałej z jednym dodatkowym polem bitwy, co pozwala na obliczanie równowag w czasie wielomianowym. W wariancie z elastycznym rozstrzyganiem remisów scharakteryzowaliśmy zbiór strategii równowagowych i zidentyfikowaliśmy równowagę odporną na zmiany reguły rozstrzygania remisów. Dla gier ataku i przechwytywania na sieciach pokazaliśmy, że dzięki wykorzystaniu struktury sieci równowagi Nasha można obliczać w czasie wielomianowym.

Nasze wyniki zostały opublikowane jako cztery osobne artykuły:

(1) S. Kaźmierowski and M. Dziubiński. Efficient Method for Finding Optimal Strategies in Chopstick Auctions with Uniform Objects Values. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '24, pages

992–1000, Richland, SC, 2024. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9798400704864

(2) S. Kaźmierowski. Equilibria of the Colonel Blotto Games with Costs. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(13):13969–13976, Apr. 2025. doi: 10.1609/aaai.v39i13.33528. URL https://ojs.aaai.org/index.php/AAAI/article/view/33528

(3) C. Ewerhart and S. Kaźmierowski. An Equilibrium Analysis of the Arad-Rubinstein Game. *Journal of Economic Behavior & Organization*, 226:106696, 2024. ISSN 0167-2681. doi: https://doi.org/10.1016/j.jebo.2024.106696. URL https://www.sciencedirect.com/science/article/pii/S0167268124003020

(4) S. Kaźmierowski and M. Dziubiński. Computation of Nash Equilibria of Attack and Defense Games on Networks. In A. Deligkas and A. Filos-Ratsikas, editors, *Algorithmic Game Theory*, pages 3–21, Cham, 2023. Springer Nature Switzerland. ISBN 978-3-031-43254-5

Słowa kluczowe: teoria gier, algorytmiczna teoria gier, równowaga Nasha, strategia optymalna, gra o sumie stałej, konflikty z wieloma polami bitew, gry o zwięzłej reprezentacji, gra Pułkownik Blotto

Acknowledgments (Podziękowania)

First and foremost, I am very grateful to my supervisor, dr hab. Marcin Dziubiński. His kindness, wisdom, and commitment have been absolutely extraordinary. Marcin, it is a privilege to have you as my mentor and my friend.

I consider myself very lucky — the list of all people who have contributed to my academic and personal growth over the past four years is far too long to include here in full. I am deeply grateful to everyone who has supported, inspired, challenged, or encouraged me along the way.

To my friends, thank you for always being there for me. Your support has meant more than I can express.

To my family, thank you for being the source of strength and stability throughout this journey. Your love and confidence in me have carried me through every stage of this work.

To my colleagues, thank you for the stimulating conversations, thoughtful guidance, and the countless laughs we have shared. Working alongside you has been rewarding and joyful.

To my coauthors, thank you for your ideas and commitment. I have learned a great deal from each joint project, and I am grateful for the opportunity to work together.

* * *

My research has been supported by the Foundation for Polish Science (FNP) and the National Science Centre (NCN) with the following grants:

- NCN OPUS grant no. 2018/29/B/ST6/00174 (*Conflicts with Multiple Battlefields and Discrete Resources*),
- NCN SONATA BIS grant no. 2021/42/E/HS4/00196 (*Motivations: Multidisciplinary Approach*),
- ERC Starting grant no. 101076570 (*Proportional Algorithms for Democratic Decision*),
- ERC Consolidator grant no. 101002854 (*Pragmatics of Multiwinner Voting: Algorithms and Preference Data Analysis*).

Contents

1 Introduction **2**
- 1.1 Succinct Games . . . 3
- 1.2 Equilibrium Computation . . . 4
- 1.3 Literature Review . . . 5
- 1.4 Our Contribution . . . 8
  - 1.4.1 Symmetric Conflicts with Multiple Battlefields: Beyond Bilinear Objectives . . . 9
  - 1.4.2 Attack and Interception Games in Networks . . . 9
  - 1.4.3 Equilibria of the Colonel Blotto Game with Flexible Tie-breaking . . . 10
  - 1.4.4 Equilibria of the Colonel Blotto Game with Costs . . . 10

2 Symmetric Conflicts with Multiple Battlefields: Beyond Bilinear Objectives **11**
- 2.1 Motivating Examples . . . 11
- 2.2 Model . . . 12
- 2.3 Opening Discussion . . . 14
  - 2.3.1 Challenges . . . 15
  - 2.3.2 Our Contribution . . . 17
- 2.4 Symmetrized Conflicts with Multiple Battlefields . . . 17
  - 2.4.1 Definition . . . 17
  - 2.4.2 Equilibrium Analysis . . . 18
  - 2.4.3 The Number of Symmetrized Strategies . . . 19
- 2.5 Computation of Payoffs in the Symmetrized Models . . . 21
  - 2.5.1 Clash Matrix . . . 21
  - 2.5.2 Recursive Formula . . . 23
  - 2.5.3 Computation of Payoffs in Symmetrized Strategy Space . . . 25
- 2.6 Computation of Nash Equilibria . . . 26
  - 2.6.1 Oracle . . . 27
  - 2.6.2 Heuristic for Computing Best Response . . . 27
  - 2.6.3 Symmetrized Payoffs Computation . . . 28
  - 2.6.4 Equilibrium Computation . . . 29
- 2.7 Conclusions . . . 30

3 Attack and Interception Games in Networks **31**
3.1 Motivating Examples . . . 31
3.2 Model . . . 33
3.3 Opening Discussion . . . 35
3.3.1 Overview of Existing Results . . . 35
3.3.2 Counterexample . . . 36
3.3.3 Our Contribution . . . 37
3.4 Preliminary Results . . . 37
3.5 Network Structure . . . 39
3.5.1 Linkers . . . 39
3.5.2 Network Reduction . . . 39
3.5.3 Ascending Networks . . . 42
3.6 Equilibrium Analysis . . . 43
3.6.1 Fully Targeted Equilibrium . . . 44
3.6.2 General Case . . . 50
3.7 Equilibrium Computation . . . 52
3.8 Conclusions . . . 55

4 Equilibria of the Colonel Blotto Game with Flexible Tie-breaking **56**
4.1 Motivating Examples . . . 57
4.2 Model . . . 58
4.2.1 Examples . . . 59
4.3 Equilibrium of the Colonel Blotto Game with Flexible Tie-breaking . . . 61
4.3.1 A canonical equilibrium . . . 61
4.3.2 Discussion . . . 62
4.3.3 Illustration . . . 63
4.3.4 Proof of Proposition 4.3 . . . 63
4.4 Understanding the Equilibrium Set . . . 64
4.4.1 Pure strategies that arise in some equilibrium . . . 64
4.4.2 Proof of Proposition 4.6 . . . 65
4.4.3 Strategies that are never “good” . . . 65
4.5 Refinements . . . 67
4.6 Further Discussion . . . 68
4.6.1 Equilibria in pure strategies . . . 68
4.6.2 Payoff-nonequivalent equilibria . . . 69
4.6.3 The robustness of uniform marginals . . . 70
4.7 Conclusions . . . 71

5 Equilibria of the Colonel Blotto Games with Costs **72**
5.1 Motivating Examples . . . 73
5.2 Model . . . 74
5.3 Analysis . . . 75

5.3.1 Properties of Equilibria . . . 76
5.4 Equilibrium Computation . . . 77
5.4.1 Colonel Blotto Game with Sunk Costs . . . 77
5.4.2 Expressing $\mathcal{B}_0$ as $\mathcal{B}$ . . . 78
5.4.3 Computational Complexity . . . 79
5.5 Experimental Results . . . 80
5.6 Conclusions . . . 85

6 Concluding Remarks **86**

Bibliography **88**

Chapter 1

Introduction

Game theory provides a mathematical framework for analyzing strategic interactions among rational agents. In many settings, the underlying strategy spaces are extremely large, yet their structure allows a compact description of the game. Instead of enumerating all strategies and payoffs explicitly, one can represent them implicitly using algebraic, combinatorial, or algorithmic rules. Such games are called *succinct games*—the formal description can be polynomial in natural parameters even though the corresponding normal form may be exponentially large.

Solving succinct games is of interest to the research community for both theoretical and practical reasons. On the theoretical side, succinct representations expose a concrete gap between the compactness of a model and the computational cost of reasoning about it: tasks that are simple in the explicit normal form may become intractable, once the same game is given succinctly. For example, evaluating a best response—polynomial-time when strategies are listed explicitly—can require exponential time for succinct encodings. On the practical side, succinct games model real decision problems where explicit enumeration is infeasible, such as resource-allocation contests, electoral campaigning, cybersecurity and interception in networks. Many decision-support systems in economics, AI, and operations research build on theoretical results for such models, so progress on their algorithmic analysis directly improves tools for large-scale, strategic decision making.

At its core, this thesis explores the algorithmic frontier of succinct games: when and how can these compactly represented games be solved efficiently, and where do computational barriers arise? While general results often point to intractability, important subclasses admit efficient methods once their structural properties are carefully exploited. The chapters that follow develop this perspective through a series of case studies centered on conflicts with multiple battlefields. These include both Colonel Blotto–like games and attack and interception games in networks, which can be viewed as conflicts where defender nodes play the role of battlefields. Across these formulations, this thesis combines hardness results that clarify the limits of tractability with algorithmic and structural insights that make equilibrium computation feasible in practice.

1.1 Succinct Games

Different types of strategic games with compactly given strategy spaces, such as the Colonel Blotto game Borel [1921] or symmetric games Nash [1951], have been considered since the early beginnings of game theory. Papadimitriou and Roughgarden [2008] proposed a formal definition of *succinct games* expressed as a computational problem. In their framework, a game is said to be succinctly representable if its payoffs can be computed by a polynomial-time algorithm given the players' actions as input. This definition captures the idea that the normal form of a game may be exponentially large, yet the underlying rules of interaction admit a compact description.

While this computational definition formalizes succinctness rigorously, many strategically relevant models fall outside its finite framework yet share the same algorithmic essence. One of the limits of the definition of succinct games proposed by Papadimitriou and Roughgarden [2008] is that it only encapsulates finite games, i.e., games where strategy sets of all players are finite. Hence, in this thesis, we adopt a broader use of the term. We consider a succinct game to be any strategic game where at least one player's strategy set is represented implicitly, rather than by explicit enumeration. For example, in attack and interception games in networks, the attacker's pure strategies correspond to all simple paths in a graph originating at the attacker's node. The number of such paths can be exponential or even factorial in the number of nodes, making explicit representation infeasible. Because the defenders' strategy spaces in this model are continuous, the game falls outside the formal framework of Papadimitriou and Roughgarden [2008], yet it has a similar algorithmic character: compact input size, combined with a large induced attacker's strategy space.

We move on to discuss several important classes of games that fall under the umbrella of succinct representations. Those are:

- Conflicts with Multiple Battlefields with Discrete Resources, such as the discrete Colonel Blotto game, first introduced by Borel [1921], where players distribute limited resources over multiple battlefields, generating exponentially many allocations.
- Network interdiction games: In its classic formulation, the problem involves two players: an evader, who sends the traffic through the network, and an interceptor, whose objective is to stop the traffic.
- Symmetric Games Nash [1951]: Games where all players are identical and indistinguishable, i.e., have the same strategy sets and payoff functions.
- Anonymous Games Blonski [2000]: A generalization of symmetric games where each player's payoff function depends symmetrically on the actions of the other players, but may differ across individuals. These games are common in large population settings, such as traffic models or market competition.
- Graphical Games Kakade et al. [2003], Kearns et al. [2001]: A compact representation of multi-player interactions where each player's payoff depends only on a subset of other

players, defined by a graph. This reduces the computational time complexity from exponential in the number of players to polynomial in the graph structure.

- Polymatrix Games Eaves [1973], Howson [1972]: A further extension of graphical games where interactions between players are represented as pairwise matrix games, enabling efficient decomposition of the overall game into smaller subproblems.
- Hypergraphical Games Papadimitriou and Roughgarden [2008]: A generalization of both graphical and polymatrix games, where each player's payoff depends on multi-way interactions encoded by a hypergraph. This class captures more complex strategic dependencies found in multi-agent collaboration and competition.
- Congestion Games Rosenthal [1973]: Games in which players compete for shared resources, such as network traffic routing, where the utility of a chosen strategy depends on how many other players select the same option. These games are widely used in transportation and distributed computing.
- Local Effect Games Leyton-Brown and Tennenholtz [2003]: A variation of congestion games where the influence of other players' choices is restricted to a local neighborhood, further reducing computational complexity.
- Scheduling Games Fotakis et al. [2002]: Games where players compete for limited processing time on shared machines, modeling job scheduling in distributed systems.
- Sparse Games Papadimitriou and Roughgarden [2008]: Games where most player interactions have zero influence, i.e., some utility values are given explicitly, and the rest are assumed to be zero.

These diverse classes of games demonstrate the breadth of succinct representations, each introducing new computational and theoretical challenges.

1.2 Equilibrium Computation

In this section, we discuss the complexity of computing a Nash equilibrium in finite games.

Nash Equilibrium. The Nash equilibrium is the central solution concept in non-cooperative game theory, where no player can improve their payoff by unilaterally deviating from their strategy. Computing Nash equilibria in general finite games is known to be PPAD-complete, meaning that it is as hard as any problem in the complexity class PPAD (Polynomial Parity Arguments on Directed graphs). This classification, established through seminal results by Daskalakis et al. [2009], implies that there is no known polynomial-time algorithm for computing a Nash equilibrium in general games.

Two-Player Zero-Sum Games. A fundamental class of games with well-defined equilibrium properties is that of two-player zero-sum games, in which the sum of players' payoffs is constant—typically normalized to zero—so that any gain of one player corresponds exactly to the loss of the other. These games have a strong theoretical foundation, as established by von Neumann's minimax theorem von Neumann [1953], which guarantees that both players have optimal strategies that maximize their worst-case expected payoff. Unlike general-sum games, where multiple Nash equilibria may yield different payoffs, every pair of optimal strategies in a zero-sum game constitutes a Nash equilibrium, and all equilibria yield the same payoff for both players. The non-negative equilibrium payoff to one of the players is called the game value, representing the amount that player can guarantee regardless of the opponent's strategy.

An optimal strategy for a player in a two-player zero-sum game can be computed by solving a linear program, as shown by Dantzig [1951]. The formulation consists of a number of variables equal to the number of strategies available to the player and a number of constraints equal to the number of strategies available to the opponent. Since linear programming can be solved in polynomial time with respect to the number of variables and constraints, optimal strategies in standard representations of these games can be computed in polynomial time with respect to the number of strategies of both players.

Equilibrium Computation in Succinct Games. When dealing with succinct games, the number of pure strategies available to each player can be exponentially large in terms of the game's description parameters. This makes a direct application of standard linear programming methods impractical, as explicitly formulating and solving the linear program would require exponential time. To address this challenge, it becomes essential to design efficient computational methods that exploit the structure of the game to avoid explicitly enumerating all strategies, thereby enabling scalable equilibrium computation.

1.3 Literature Review

In this section, we provide a detailed overview of the literature concerning equilibrium computation in conflicts with multiple battlefields and interdiction games in networks.

Colonel Blotto Game. The Colonel Blotto game was first introduced by Borel [1921]. In its original formulation, two colonels are tasked with simultaneously assigning their limited troops (discrete or continuous) across the set of battlefields. Each colonel wins all the battlefields where they assigns strictly more resources than the opponent, and the remaining battlefields are tied. In the original formulation of the model, all battlefields are equally valuable, and the payoff to a colonel is the difference between the number of won and lost battlefields. Although first applications of the Colonel Blotto game were military Blackett [1954], Weber and Shubik [1981], the model has been since applied to analyzing competitions in various scenarios, such as politics Laslier and Picard [2002], Munger et al. [2005], Washburn [2013], Dehghani et al. [2021], network security Guan et al. [2018, 2019], Chia and Chuang [2011], and marketing Chowdhury

et al. [2009], Kovenock and Roberson [2012], Jiao and Xu [2020]. A simple description of this game, combined with a rich structure and many possible applications, resulted in ongoing interest among researchers in multiple areas in solutions of the Colonel Blotto game. Examples of these fields are operations research Blackett [1954], Beale and Heselden [1962], Bell and Cover [1980], Weber and Shubik [1981], economics Hart [2008], Roberson [2006], Roberson and Kvasov [2010], Kovenock and Roberson [2021], and computer science Behnezhad et al. [2018b], Ahmadinejad et al. [2019], Vu and Loiseau [2021], Behnezhad et al. [2023].

One commonly studied simplification of the model replaces the discrete allocation of troops with continuous resource division, allowing players to split their forces across battlefields.

The first stream of research concerning solutions of the Colonel Blotto game aims to obtain (at least a partial) characterization. Although in most formulations the game is zero-sum, the exponential number of strategies with respect to the model parameters makes the problem of characterization of optimal strategies hard, and most results were obtained under restrictive assumptions that narrow their possible applications. Borel and Ville [1938] solved the continuous variant of game for three battlefields and the same number of resources of both players. Their results were extended by Gross and Wagner [1950], who obtained a general solution under the symmetric budget of both players. Roberson [2006] extended the knowledge on continuous variants of the game, providing solutions for the general case both in terms of the number of battlefields and asymmetric budgets of players. These results were further generalized by Kovenock and Roberson [2021], who solved the variant allowing for different valuations across battlefields. As for the discrete variant of the problem, Hart [2008] provided partial characterization of solutions of the budget-symmetric model as well as for some of the budget-asymmetric cases for a general number of battlefields. To this day, the general characterization of the equilibrium set in Colonel Blotto games with discrete strategy spaces remains unknown.

Difficulties in characterizing solutions of the more general variants of the Colonel Blotto game (e.g. different values across battlefields or different values of battlefields across players) resulted in increased attention regarding the problem of computing or approximating solutions of such models. Beale and Heselden [1962] introduced a method of approximating solutions of the Colonel Blotto game, based on a new game involving the mean numbers of forces allocated to each battlefield in the original game. Penn [1971]) introduced a generalized Langrange-multipliers method, allowing for computing solutions of a broader class of games, including the Colonel Blotto game. Another implementation for computing the Colonel Blotto game solutions, based on fictitious play, was proposed by Eisen and Le Mat [1968]. None of the aforementioned methods guaranteed a polynomial runtime with respect to the number of resources of both players and the number of battlefields.

Recently, Ahmadinejad et al. [2019] proposed a first method of computing solutions for the Colonel Blotto game in polynomial time with respect to the model parameters. The proposed method utilizes the idea of a separation oracle and was proved applicable to a broader class of games satisfying the bilinearity of payoffs, i.e., the payoff to a player is the sum of payoffs across single battlefields. Although considered a breakthrough because of the guarantees of polynomial runtime, the method proposed by Ahmadinejad et al. [2019] uses the ellipsoid method, which makes it insufficient for larger model parameters. Behnezhad et al.

[2023] addressed this problem, providing the first polynomial-size LP formulation of the problem, resulting in a simpler and significantly faster algorithm. Furthermore, they showed that their LP representation is asymptotically tight. Their method, based on the idea of representing a mixed strategy of a player as a flow in the corresponding layered graph, is applicable to every zero-sum Colonel Blotto game satisfying the bilinearity of payoffs and the zero-sum property at every battlefield, and it is considered the state-of-the-art algorithm for computing solutions of the Colonel Blotto game.

While the Colonel Blotto game captures additive contests across multiple dimensions, many real-world conflicts exhibit non-additive aggregation or network constraints. In the next part of this section, we discuss how these features have been explored in non-bilinear and network-based models.

Conflicts with Multiple Battlefields and Non-bilinear Objective. Conflicts with multiple battlefields under non-bilinear objectives arise when the payoffs depend not on the additive sum of independently evaluated battlefields, as in the Colonel Blotto game, but rather on a more complex aggregation of outcomes. Examples include majoritarian rules, threshold objectives, or contest success functions that depart from linear additivity. These models are motivated by settings where what matters is not the cumulative value of battlefields, but rather the probability of surpassing an opponent on a critical fraction of them, or where nonlinear effects shape the ultimate payoff.

The first stream of research on such conflicts has focused on continuous formulations. Szentes and Rosenthal [2003b,a] analyzed chopstick auctions, in which players divide resources across multiple objects and winning probabilities depend on non-additive aggregation rules. For the three-object case, they provided explicit equilibrium constructions, and they extended their analysis to more general multi-object settings. Ewerhart [2022] studied further continuous models with nonlinear payoffs and showed that equilibrium strategies can exhibit fractal structures, highlighting the mathematical richness of such games.

Discrete formulations have also received attention. Behnezhad et al. [2018b] considered a variant of the Colonel Blotto game in which players aim to guarantee a minimum probability of achieving a target payoff, a problem that departs from bilinearity and requires new algorithmic tools. Carvalho et al. [2022] proposed the general framework of integer programming games, which naturally encompasses multi-battlefield conflicts with nonlinear objectives and connects them to combinatorial optimization.

Overall, the study of conflicts with multiple battlefields and non-bilinear objectives has revealed both the expressive power of these models and the mathematical challenges they entail. Compared to bilinear Blotto-type contests, the lack of bilinearity of payoffs makes characterization and computation significantly harder, and existing results are often confined to special cases or rely on strong assumptions. Nevertheless, this literature demonstrates that such models capture important strategic scenarios and motivate the development of new methods for understanding succinct games beyond the bilinear paradigm.

Interdiction Games in Networks. Network interdiction games model strategic interactions in which an evader seeks to traverse a network while an interdictor aims to hinder or prevent the evader's success. Depending on the formulation, the evader's goal may vary: minimizing the length of the path from a source to a sink Golden [1978], avoiding detection by devices placed on edges Washburn and Wood [1995], or maximizing the amount of flow sent between a source and a sink Fulkerson and Harding [1977], Ghare et al. [1971], McMasters and Mustin [1970]. Each of these models captures different operational aspects of the evader–interdictor interaction while retaining the combinatorial difficulty inherent in reasoning about paths and flows in networks.

The applications of network interdiction problems are broad and well-documented. They include smuggling prevention Dell [2015], Mirzaei et al. [2021], detection of nuclear materials Morton et al. [2007], and disruption of enemy troop movements McMasters and Mustin [1970], Ghare et al. [1971]. These examples highlight the central role of interdiction games in both security and defense contexts, where limited defensive resources must be allocated to counter adversaries who adapt their strategies to the network structure.

Over time, interdiction games have become a rich field of study at the intersection of operations research, game theory, and network optimization. A recent contribution is the attack and interception game in networks, introduced by Bloch et al. [2023], which explicitly formulates the interaction between an attacker moving through the network and defenders allocating interception probabilities at nodes. Their work established equilibrium existence and uniqueness under natural structural assumptions, and highlighted open challenges in characterizing targeted nodes in general networks. For comprehensive treatments of the broader interdiction literature, excellent surveys are provided by Collado and Papp [2010] and Smith and Song [2020]. These works synthesize decades of research, showing how interdiction games have evolved from early military applications to modern contexts such as supply-chain security and counterterrorism.

1.4 Our Contribution

The dissertation is based on the following research papers:

(1) S. Kaźmierowski and M. Dziubiński. Computation of Nash Equilibria of Attack and Defense Games on Networks. In A. Deligkas and A. Filos-Ratsikas, editors, *Algorithmic Game Theory*, pages 3–21, Cham, 2023. Springer Nature Switzerland. ISBN 978-3-031-43254-5

(2) S. Kaźmierowski and M. Dziubiński. Efficient Method for Finding Optimal Strategies in Chopstick Auctions with Uniform Objects Values. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '24, pages 992–1000, Richland, SC, 2024. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9798400704864

(3) C. Ewerhart and S. Kaźmierowski. An Equilibrium Analysis of the Arad-Rubinstein Game. *Journal of Economic Behavior & Organization*, 226:106696, 2024. ISSN

0167-2681. doi: https://doi.org/10.1016/j.jebo.2024.106696. URL https://www.sciencedirect.com/science/article/pii/S0167268124003020

(4) S. Kaźmierowski. Equilibria of the Colonel Blotto Games with Costs. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(13):13969–13976, Apr. 2025. doi: 10.1609/aaai.v39i13.33528. URL https://ojs.aaai.org/index.php/AAAI/article/view/33528

I contributed to all technical results presented in this dissertation. Each of the four subsequent chapters is based on one of our research papers, possibly with extended or refined results. Below, we provide a brief overview of these contributions.

1.4.1 Symmetric Conflicts with Multiple Battlefields: Beyond Bilinear Objectives

This chapter presents the results of Kaźmierowski and Dziubiński [2024]. We analyze a class of battlefield-symmetric conflicts with multiple battlefields under non-linear aggregation functions, specifically the majoritarian and more-than-opponent functions, in the setting where all battlefield values are uniform. Unlike in the classical Colonel Blotto game, these aggregation rules yield payoff functions that are not bilinear, which invalidates the known polynomial-time solution methods for the Colonel Blotto games.

We show that computing a best response to a mixed strategy in this setting cannot be done in polynomial time unless $\mathsf{P} = \mathsf{NP}$, in sharp contrast with bilinear Blotto games, where best responses can be computed in polynomial time. To address this difficulty, we extend the symmetrization approach, proposed by Hart [2008], to these conflicts, thereby reducing the effective strategy space by an exponential factor. We further introduce the clash matrix algorithm, a novel method that allows for computing payoffs from symmetrized strategy profiles in polynomial time. By combining these ideas with the Double Oracle Algorithm, supported by a heuristic tailored to the monotonicity properties of best responses, we design a method that efficiently computes optimal strategies of the considered games. Experimental evaluation demonstrates that our algorithm outperforms standard approaches on a range of instances by three orders of magnitude.

1.4.2 Attack and Interception Games in Networks

This chapter presents the results of Kaźmierowski and Dziubiński [2023]. In this work, we study attack and interception games in networks, introduced by Bloch et al. [2023]. In these games, the attacker selects a path in a graph while the defenders, each controlling a unique node in the network, allocate interception resources across the nodes. The attacker's strategy set consists of all simple paths in the network originating at the attacker's node, which can be exponential or even factorial in the number of nodes, making the game a natural example of a succinct representation.

We show that, despite this possibly exponential blowup of the strategy space, Nash equilibria can be computed in polynomial time for a broad class of cost functions. The key to our approach is the development of efficient algorithms that exploit the structural properties of networks. In particular, we show that the set of attack paths used by the attacker in equilibrium is unique, polynomial-size, and can be found in polynomial time.

This chapter demonstrates that at least some types of succinctly represented games in networks can be solved efficiently when their structure is carefully exploited.

1.4.3 Equilibria of the Colonel Blotto Game with Flexible Tie-breaking

This chapter presents the results of Ewerhart and Kaźmierowski [2024]. Motivated by the example of Arad and Rubinstein [2012], we study Colonel Blotto games with flexible tie-breaking rules. We develop a constructive approach to characterize Nash equilibria in this setting. Our analysis reveals that the equilibrium set is very large: every pure strategy that does not allocate an excessively large share of resources to a single battlefield is used with positive probability in some mixed-strategy equilibrium. We further examine criteria for equilibrium selection and show that neither the elimination of dominated strategies nor robustness with respect to the tie-breaking rule provides a meaningful refinement.

1.4.4 Equilibria of the Colonel Blotto Game with Costs

This chapter presents the results of Kaźmierowski [2025]. We consider a generalized variant of the Colonel Blotto game that incorporates explicit costs associated both with obtaining and deploying resources across battlefields. In contrast to the classical model, which assumes a strict use-it-or-lose-it budget, this cost-augmented formulation captures the strategic importance of resource acquisition and deployment costs.

We show that every instance of the Colonel Blotto game with costs is strategically equivalent to an instance of the zero-sum Colonel Blotto game with one additional battlefield. This reduction allows us to compute Nash equilibria of the cost-augmented game in polynomial time with respect to the natural game parameters, the number of battlefields and the number of resources available to the players.

This chapter extends the scope of Blotto-type models by incorporating costs, while preserving algorithmic tractability through a polynomial-time reduction to the classical zero-sum case.

Chapter 2

Symmetric Conflicts with Multiple Battlefields: Beyond Bilinear Objectives

Chapter Overview

In this chapter, we investigate the problem of computing Nash equilibria in symmetric conflicts with multiple battlefields under non-bilinear objectives such as majoritarian and chopstick aggregation functions. While the general complexity of equilibrium computation remains open, we establish that computing a best response to a given mixed strategy is NP-hard in these settings—suggesting fundamental difficulty. Despite this, we show that by focusing on symmetrized strategies and leveraging structural properties of the game, we can substantially reduce the computational burden.

The rest of the chapter is organized as follows. In Section 2.1, we provide real-life motivations of our study. The model is introduced in Section 2.2. Section 2.3 provides a discussion on the challenges and outlines our contribution. In Section 2.4 we define the symmetrized version of the model. Computation of payoffs from symmetrized strategy profiles and the clash matrix algorithm are given in Section 2.5. In Section 2.6, we present our method of computation of Nash equilibria of the model in question, and we provide experimental evaluation of our approach.

2.1 Motivating Examples

In competitive resource allocation scenarios, agents are often faced with the challenge of simultaneously engaging in multiple local contests while constrained by global budgets. A classical example of this phenomenon arises in military planning, where opposing commanders must allocate troops across a set of strategic locations or battlefields. Victory in each individual battle is determined by which side commits more resources, but overall success depends on how those local victories aggregate into a global outcome. In some settings, each battlefield contributes linearly to the final payoff—this is the case in the well-known Colonel Blotto game, where

the utility is simply the difference between the number of won and lost battlefields. In other settings, however, the aggregation may follow a non-linear, often threshold-based rule.

For instance, in majoritarian democratic systems with uniform voting districts, political candidates compete over a set of identical constituencies, each awarding a point to the candidate who allocates more campaign effort. Each candidate has a limited number of campaign days to distribute. The candidate who wins a strict majority of constituencies wins the election overall. This is a textbook example of a non-linear payoff: winning just over half the districts results in full victory, while one district less results in complete defeat, regardless of the margins.

A similar structure appears in high-stakes legal disputes argued before *multi-judge panels*. Two legal teams must allocate their limited time and effort across a panel of judges. Each judge is persuaded by the side that dedicates more effective argumentation, and the side that convinces a majority of the judges wins the case. Hence, the outcome is determined by securing the support of a majority of the judges.

Another example arises in security audit competitions between cybersecurity firms. Consider two firms competing for a government contract, evaluated through a series of independent, standardized audits. Each firm must distribute a discrete number of skilled testers across audit stations. A station is won by the firm allocating more testing resources to it. The overall contract goes to the firm that wins more audit stations than its competitor. Crucially, the margin of victory within each audit does not matter—only the total count of wins. This scenario exemplifies chopstick-style aggregation, where the payoff depends solely on outperforming the opponent in a greater number of symmetric contests.

These examples all reflect conflicts where resource allocation decisions must be made globally, yet payoffs are determined by local head-to-head comparisons and aggregated non-linearly. Whether in courts, elections, or technical evaluations, such environments highlight the practical significance of studying conflicts with multiple battlefields under non-bilinear objectives.

2.2 Model

Conflicts with multiple battlefields model a competition between two players, denoted by A and B. Each player has a number of discrete resources, e.g., military units or coins, denoted by D^{A} and D^{B}, respectively. Players compete on the set $[n] = \{1, 2, \dots n\}$ of n battlefields. We assume that $n \geq 2$. Each battlefield $i \in [n]$ has a value of 1. The player's strategy is to distribute her resources across battlefields. The set of strategies of player $P \in \{\mathsf{A}, \mathsf{B}\}$ is:

$$S(D^P, n) = \left\{ \boldsymbol{s} \in \mathbb{N}^n : \sum_{i=1}^{n} s_i = D^P \right\}. \tag{2.1}$$

When the parameters are clear from the context, we use $S^P \equiv S(D^P, n)$ to denote the set op pure strategies of player $P \in \{\mathsf{A}, \mathsf{B}\}$. A pair of strategies $(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}}) \in S^{\mathsf{A}} \times S^{\mathsf{B}}$ is called a *strategy profile*.

Each battlefield is won by the player assigning strictly more resources to it, or tied otherwise.

Formally, the payoff from a single battlefield $i \in [n]$ is defined as

$$u_i^{\mathsf{A}}\left(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}}\right) = -u_i^{\mathsf{B}}\left(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}}\right) = \operatorname{sign}\left(s_i^{\mathsf{A}} - s_i^{\mathsf{B}}\right). \tag{2.2}$$

Applying (2.2) to a strategy profile $(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}})$, we get the battlefield outcomes vector

$$\boldsymbol{u}^{\mathsf{A}}\left(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}}\right) = -\boldsymbol{u}^{\mathsf{B}}\left(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}}\right) = \left(u_1^{\mathsf{A}}\left(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}}\right), \ldots, u_n^{\mathsf{A}}\left(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}}\right)\right). \tag{2.3}$$

Values of vectors $\boldsymbol{u}^{\mathsf{A}}$ and $\boldsymbol{u}^{\mathsf{B}}$ are in $\{-1, 0, 1\}^n$.

We proceed to define payoffs to the players. We call a function $f : \{-1, 0, 1\}^n \to \mathbb{R}$ an *aggregation function* if for every vector $\boldsymbol{u} \in \{-1, 0, 1\}^n$, f is invariant under the permutations of $\boldsymbol{u}$, i.e., for every permutation $\boldsymbol{u}'$ of vector $\boldsymbol{u}$, $f(\boldsymbol{u}) = f(\boldsymbol{u}')$. Given an aggregation function f and a strategy profile $\left(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}}\right)$, payoffs to players satisfy:

$$\pi_f^{\mathsf{A}}\left(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}}\right) = -\pi_f^{\mathsf{B}}\left(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}}\right) = f\left(\boldsymbol{u}^{\mathsf{A}}\left(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}}\right)\right). \tag{2.4}$$

Definition 2.1 (Conflict with Multiple Battlefields)**.** For a given quadruple $(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f)$, a strategic game with strategy sets defined by Equation (2.1) and payoffs defined by Equation (2.4) is denoted with $\mathcal{M}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f)$ and called the conflict with multiple battlefields with aggregation function f.

The players are allowed to make randomized choices. A *mixed strategy* of player $P \in \{\mathsf{A}, \mathsf{B}\}$ is a probability distribution on S^P. Given a non-empty set X, let $\Delta(X)$ denote the set of all probability distributions on X. The expected payoff to player $P \in \{\mathsf{A}, \mathsf{B}\}$ from a pair of mixed strategies $(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}}) \in \Delta(S^{\mathsf{A}}) \times \Delta(S^{\mathsf{B}})$ is equal to

$$\Pi_f^P\left(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}}\right) = \sum_{(\boldsymbol{x}, \boldsymbol{z}) \in S^{\mathsf{A}} \times S^{\mathsf{B}}} \xi_{\boldsymbol{x}}^{\mathsf{A}} \xi_{\boldsymbol{z}}^{\mathsf{B}} \pi_f^P(\boldsymbol{x}, \boldsymbol{z}),$$

where given the probability distribution ξ^P, $\xi_{\boldsymbol{x}}^P$ is the probability associated with pure strategy $\boldsymbol{x}$. We assume that the players make their decisions "simultaneously", i.e., each player chooses her strategy without observing the choice of the opponent. We are interested in mixed strategy Nash equilibria (MSNE) of the game, i.e., mixed strategy profiles $(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}}) \in \Delta(S^{\mathsf{A}}) \times \Delta(S^{\mathsf{B}})$ such that no player can improve her expected payoff by changing her strategy unilaterally, i.e., for each $P \in \{\mathsf{A}, \mathsf{B}\}$ and all $\boldsymbol{\zeta} \in \Delta\left(S^P\right)$,

$$\Pi_f^P\left(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}}\right) \geq \Pi_f^P\left(\boldsymbol{\zeta}, \boldsymbol{\xi}^{-P}\right),$$

where $-P$ denotes the player other than P and $(\boldsymbol{\zeta}, \boldsymbol{\xi}^{-P})$ is the strategy profile obtained from $(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}})$ by replacing $\boldsymbol{\xi}^P$ with $\boldsymbol{\zeta}$.

We conclude this section with three examples of symmetric conflicts with multiple battlefields considered in the literature. The first example presents the Colonel Blotto game, where each player's utility is the difference between the number of won and lost battlefields.

Example 2.2 (Colonel Blotto Game)**.** Consider an aggregation function

$$f_{blotto}(\boldsymbol{u}) = \sum_{i=1}^{n} u_i. \tag{2.5}$$

Game $\mathcal{M}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f_{blotto})$ is called the Colonel Blotto game.

The second example presents the Majoritarian Auction, where the utility of a player is 1, if the player won strictly more than half of the battlefields, -1, if the player's opponent won strictly more than half of the battlefields, and 0 otherwise.

Example 2.3 (Majoritarian Auction)**.** Consider an aggregation function

$$f_{maj}(\boldsymbol{u}) = \left[\sum_{i=1}^{n} [u_i = 1] > n/2\right] - \left[\sum_{i=1}^{n} [u_i = \text{-}1] > n/2\right], \tag{2.6}$$

where, given a condition φ, $[\varphi]$ is the Iverson's bracket taking value 1 if φ is satisfied and value 0 otherwise. Game $\mathcal{M}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f_{maj})$ is called the Majoritarian Auction.

The third example presents the Chopstick Auction, where the utility of a player is 1, if the player won strictly more battlefields than her opponent, -1 if the player's opponent won strictly more battlefields than the player, and 0 otherwise.

Example 2.4 (Chopstick Auction)**.** Consider an aggregation function

$$f_{chop}(\boldsymbol{u}) = \operatorname{sign}\left(\sum_{i=1}^{n} u_i\right), \tag{2.7}$$

Game $\mathcal{M}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f_{chop})$ is called the Chopstick Auction.

The aggregation functions defined by Equations (2.6) and (2.7) are two extensions of the chopstick auction defined by Szentes and Rosenthal [2003b,a], where two players compete with an objective of winning two out of three battlefields. In general, many more aggregation functions generate nontrivial games. Throughout this chapter, we focus on the problem of computing Nash equilibria of Majoritarian and Chopstick Auctions.

2.3 Opening Discussion

This section outlines the key difficulties in computing equilibria for symmetric multi-battlefield conflicts with non-bilinear aggregation functions, including Chopstick and Majoritarian Auctions, and presents our contribution toward resolving them.

2.3.1 Challenges

Since every symmetric conflict with multiple battlefields is a finite, two-player, zero-sum game, every profile of optimal strategies of both players is a Nash equilibrium, as shown by von Neumann [1953]. The problem of computing an optimal strategy of a player can be expressed as a linear program (LP) Ferber [1959], Raghavan [1994] and requires polynomial time with respect to the number of strategies of both players. However, a player with D^P resources to distribute over n battlefields has

$$\binom{n + D^P - 1}{n - 1} \tag{2.8}$$

pure strategies available. To formalize the discussion on the number of available pure strategies, throughout this chapter, we assume that when increasing the number of battlefields, the number of resources of each player increases as a linear function, such that both coefficients are strictly positive numbers. Formally, for a number of battlefields $n \geq 2$, we assume that the number of resources of player $P \in \{\mathsf{A}, \mathsf{B}\}$ is given by a linear function $g^P(n) = a^P n + k^P$, such that $a^P, k^P \in \mathbb{N}_+$. Hence, given player $P \in \{\mathsf{A}, \mathsf{B}\}$, a function $g^P(n) = a^P n + k^P$, and a number of battlefields $n \geq 2$, the number of pure strategies of player P is:

$$\begin{aligned}
|S(g^P(n), n)| &= \binom{n + g^P(n) - 1}{n - 1} \\
&= \binom{(a^P + 1)n + k^P - 1}{n - 1} \\
&= \frac{a^P n + k^P - 1}{n} \binom{(a^P + 1)n + k^P - 1}{n} \\
&\geq \binom{2n}{n} \\
&\geq 2^n.
\end{aligned} \tag{2.9}$$

The exponential number of pure strategies with respect to the model parameters makes the standard LP-solver approach inapplicable in the case of conflicts with multiple battlefields. In the case of bilinear aggregation functions, e.g., the Colonel Blotto game, optimal strategies can be computed in polynomial time, as shown by Ahmadinejad et al. [2019], Behnezhad et al. [2023], by exploiting the game structure. As for the Majoritarian and Chopstick Auctions, there are no known methods for computing optimal strategies in polynomial time with respect to the model parameters.

Although the complexity of computing Nash equilibria of Chopstick and Majoritarian Auctions remains unknown, we show that finding a pure strategy best response to a given mixed strategy is a computationally hard problem. It is in stark contrast to the Colonel Blotto game, where a pure strategy best response can be computed in polynomial time with respect to the model parameters. We prove the result using a reduction from the max coverage problem, following the idea of Behnezhad et al. [2018a], where computing a pure strategy best response to a mixed strategy, that maximizes the probability of obtaining a payoff above a given threshold in the Colonel Blotto game, is shown to be computationally hard.

The hardness result of finding the best response does not necessarily imply a hardness result for the computation of a Nash equilibrium of a game, however, it is considered an insightful indicator of how computationally hard the game is to solve. See Xu [2016] for results about a class of games where the hardness of computing a best response implies the hardness of computing a Nash equilibrium of the game.

Assume a mixed strategy $\boldsymbol{\xi}^{\mathsf{B}}$ of player B, given as a list of pure strategies in its support, and their associated probabilities. The best response problem for a given aggregation function f, denoted by BR_f, is to find a pure strategy $\boldsymbol{s}^{\mathsf{A}}$ of player A that maximizes the $\Pi_f^{\mathsf{A}}(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}})$.

Proposition 2.5. *For aggregation functions f given by (2.6) and (2.7), there is no polynomial time algorithm to solve BR_f unless $\mathsf{P} = \mathsf{NP}$.*

Proof. To prove this hardness we provide a reduction from the max-coverage problem to BR_f. A number k and a collection of sets $S = \{S_1, S_2, \ldots, S_m\}$ are given. The maximum coverage problem is to find a collection $S' \subseteq S$, such that $|S'| \leq k$ and the number of covered elements (i.e., $|\bigcup_{S_i \in S'} S_i|$) is maximized.

Let $E = \bigcup_{S_i \in S} S_i$ denote the set of all elements in the given max-coverage instance. Assume, that the number of sets in S in which a given element appears is the same for every element in E and denote this number by t. If that is not the case, a polynomial number (with respect to $|E| \cdot |S|$) of additional singleton sets can be added to collection S, without interfering with the max-coverage complexity, or the resulting collection S' as no singleton is more desired than any other sets containing the only element of the considered singleton.

Consider a game with aggregation function f defined by (2.6), $2 \cdot |S|$ battlefields, where the player A has $|S| + k$ resources and player B has $(|S| + k + 1) \cdot (|S| - t)$ resources. The first $|S|$ of $2 \cdot |S|$ battlefields correspond to the sets in S. For every element $e \in E$, we denote a corresponding pure strategy of player B, denoted by $s^{\mathsf{B},e}$, by assigning $|S| + k + 1$ resources to each battlefield corresponding to a set that does not contain e and assigning 0 resources to all other battlefields. Assume that player B uses a mixed strategy defined by choosing uniformly from the set of strategies $\bigcup_{e \in E} s^{\mathsf{B},e}$.

Assume that the best response of player A, that maximizes her expected payoff, is given. Note, that as player B always assigns either 0 or $|S| + k + 1$ resources to a single battlefield, it is sufficient for the player A to assign either 0 or 1 resources on every battlefield, as the bid of at least $|S| + k + 1$ is not feasible due to the number of resources $D^{\mathsf{A}} = |S| + k$.

We claim that the best response of player A assigns $|S|$ resources to the remaining $|S|$ battlefields (one to each), which are always assigned 0 by player B. This guarantees that player A never loses, as she is guaranteed to win at least half of the battlefields. The remaining k resources are assigned between the first $|S|$ battlefields, yielding a max-coverage of size k from sets in $|S|$, where each battlefield corresponding to the set in the considered max-coverage will be assigned exactly one resource.

The proof can be easily adapted for the aggregation function (2.7). Instead of assigning additional $|S|$ battlefields that are assigned more than 0 by player B, we only add $|S| - t$ additional battlefields, and we reduce the number of resources of player A from $|S| + k$ to $|S| - t + k$. □

2.3.2 Our Contribution

Having described the main challenges of the computation of Nash equilibria of symmetric conflicts with multiple battlefields with non-bilinear aggregation functions, we proceed to present our contribution.

First, following the idea of Hart [2008], we prove that when computing a Nash equilibrium of any game in question, we can restrict attention to symmetrized strategies. A symmetrized pure strategy is a mixed strategy of the original model, defined by a partition of the number of resources into n non-negative integers, where every of $n!$ permutations of such a partition is assigned a probability of $1/n!$. Although the number of available partitions is still exponential with respect to the game parameters, it is also exponentially smaller than the number of pure strategies. Hence, restricting attention to symmetrized strategies allows for a reduction in the number of strategies by an exponential factor.

Second, we address the problem of computing the payoffs in the symmetrized strategy space. Given a profile of symmetrized strategies, a straightforward method of computing a payoff to a player from such a profile requires considering up to $n!$ pure strategy profiles – one for each permutation. We address this challenge by proposing a *Clash Matrix algorithm* that allows for computing such payoffs in polynomial time, $\mathcal{O}(n^7)$, with respect to the number of battlefields, n. The idea of clash matrix approach was initially introduced by Kaźmierowski [2021] for the chopstick auctions. In this thesis, we present the generalization of this method to all conflicts with multiple battlefields fitting Definition 2.1. We provide an experimental comparison of our method with the CPU and GPU-based naïve methods, and show that our approach scales significantly better.

Third, we address the problem of computing the entire payoff matrix in order to compute a Nash equilibrium of the game in question by applying the Double Oracle Algorithm (DOA) McMahan et al. [2003]. Our experimental evaluation shows that DOA scales significantly better than any method requiring calculation of the entire payoff matrix of a game, such as the LP-solver method.

2.4 Symmetrized Conflicts with Multiple Battlefields

In this section, we introduce symmetrized conflicts with multiple battlefields.

2.4.1 Definition

To address the problem of large strategy sets, we follow the idea of Hart [2008] and consider a symmetrized variant of the model in question, defined as follows. Let P_n denote the set of all permutations over $[n]$. Given a strategy $\boldsymbol{s} = (s_1, \ldots, s_n) \in S^P$ of player P let $\sigma(\boldsymbol{s})$ be the mixed strategy which for each permutation $p \in P_n$ over $[n]$ chooses, with probability $1/n!$, an assignment $p(\boldsymbol{s}) = (s_{p(1)}, \ldots, s_{p(n)})$ of resources. Let

$$\sigma\left(S^P\right) = \left\{\sigma(\boldsymbol{s}) : \boldsymbol{s} \in S^P\right\}, \tag{2.10}$$

denote the set of all symmetric strategies of player $P \in \{\mathsf{A}, \mathsf{B}\}$.

Definition 2.6 (Symmetrized Conflict with Multiple Battlefields)**.** For a given quadruple $(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f)$, a strategic game with strategy sets defined by Equation (2.10) and payoffs defined by Equation (2.4) is denoted with $\mathcal{SM}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f)$ and called the *symmetrized conflict with multiple battlefields* with aggregation function f.

The elements of $\sigma\left(S^P\right)$ are called *symmetric strategies*. Since symmetric strategies are order-invariant, when referring to a symmetric strategy $\sigma(\boldsymbol{s})$, we assume that $s_1 \geq \ldots \geq s_n$. Both models (defined in Definitions 2.1 and 2.6) describe finite games and therefore, by the Nash [1951] theorem, each of these games yields a Nash equilibrium in mixed strategies.

Given a mixed strategy $\boldsymbol{\xi} \in \Delta\left(S^P\right)$ of player P in the original game, let $\sigma(\boldsymbol{\xi})$ be the mixed strategy that for each permutation $p \in P_n$ and each $\boldsymbol{s} \in S^P$ chooses, with probability $\xi_{\boldsymbol{s}}/n!$, assignment $(s_{p(1)}, \ldots, s_{p(n)})$. Note that for any mixed strategy $\boldsymbol{\zeta} \in \Delta(\sigma\left(S^P\right))$ in the symmetrized game there exists a mixed strategy $\boldsymbol{\xi}$ in the original game such that $\boldsymbol{\zeta} = \sigma(\boldsymbol{\xi})$. We call the mixed strategies of the symmetrized game *symmetric mixed strategies*.

2.4.2 Equilibrium Analysis

The following proposition implies that any NE of any game that matches Definition 2.6 is also an NE of a corresponding game described by Definition 2.1.

Proposition 2.7. *Fix a quadruple $(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f)$. If a mixed strategy profile $\left(\sigma\left(\boldsymbol{\xi}^{\mathsf{A}}\right), \sigma\left(\boldsymbol{\xi}^{\mathsf{B}}\right)\right)$ is a Nash equilibrium of $\mathcal{SM}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f)$, then it is also a Nash equilibrium of $\mathcal{M}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f)$.*

Proof. First, observe that for any player $P \in \{\mathsf{A}, \mathsf{B}\}$, any aggregation function f, any two mixed strategies $\boldsymbol{\xi} \in \Delta\left(S^P\right)$, of P and $\boldsymbol{\zeta} \in \Delta\left(S^{-P}\right)$ of the other player and any permutation $q \in P_n$, it holds, that:

$$\begin{aligned}
\Pi_f^P(\sigma(\boldsymbol{\xi}), q(\boldsymbol{\zeta})) &= \sum_{p \in P_n} \sum_{\boldsymbol{x} \in S^P} \sum_{\boldsymbol{z} \in S^{-P}} \frac{\xi_{\boldsymbol{x}}}{n!} \zeta_{\boldsymbol{z}} \pi_f^P\left(p\left(\boldsymbol{x}\right), q\left(\boldsymbol{z}\right)\right) \\
&= \sum_{p \in P_n} \sum_{\boldsymbol{x} \in S^P} \sum_{\boldsymbol{z} \in S^{-P}} \frac{\xi_{\boldsymbol{x}}}{n!} \zeta_{\boldsymbol{z}} \pi_f^P\left(q^{-1}\left(p(\boldsymbol{x})\right), \boldsymbol{z}\right) \\
&= \sum_{p \in P_n} \sum_{\boldsymbol{x} \in S^P} \sum_{\boldsymbol{z} \in S^{-P}} \frac{\xi_{\boldsymbol{x}}}{n!} \zeta_{\boldsymbol{z}} \pi_f^P\left(p\left(\boldsymbol{x}\right), \boldsymbol{z}\right) \\
&= \Pi_f^P(\sigma(\boldsymbol{\xi}), \boldsymbol{\zeta}).
\end{aligned} \tag{2.11}$$

Take any mixed strategy profile $(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}})$ of the game in question and suppose that $\left(\sigma\left(\boldsymbol{\xi}^{\mathsf{A}}\right), \sigma\left(\boldsymbol{\xi}^{\mathsf{B}}\right)\right)$ is a Nash equilibrium of $\mathcal{SM}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f)$. Assume, to the contrary, that $\left(\sigma\left(\boldsymbol{\xi}^{\mathsf{A}}\right), \sigma\left(\boldsymbol{\xi}^{\mathsf{B}}\right)\right)$ is not a Nash equilibrium of $\mathcal{M}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f)$. Hence, there exists a player $P \in \{\mathsf{A}, \mathsf{B}\}$ and a mixed strategy $\boldsymbol{\zeta} \in \Delta\left(S^P\right)$ of P such that

$$\Pi_f^P\left(\boldsymbol{\zeta}, \sigma\left(\boldsymbol{\xi}^{-P}\right)\right) > \Pi_f^P\left(\sigma\left(\boldsymbol{\xi}^{\mathsf{A}}\right), \sigma\left(\boldsymbol{\xi}^{\mathsf{B}}\right)\right).$$

By Equation (2.11), it holds that:

$$\Pi_f^P\left(\boldsymbol{\zeta}, \sigma\left(\boldsymbol{\xi}^{-P}\right)\right) = \Pi_f^P\left(\sigma\left(\boldsymbol{\zeta}\right), \sigma\left(\boldsymbol{\xi}^{-P}\right)\right).$$

Hence, it holds that

$$\Pi_f^P\left(\sigma\left(\boldsymbol{\zeta}\right), \sigma\left(\boldsymbol{\xi}^{-P}\right)\right) > \Pi_f^P\left(\sigma\left(\boldsymbol{\xi}^{\mathsf{A}}\right), \sigma\left(\boldsymbol{\xi}^{\mathsf{B}}\right)\right).$$

Since $\sigma(\boldsymbol{\zeta})$ is a symmetric strategy, that is, $\sigma(\boldsymbol{\zeta}) \in \Delta(\sigma(S^P))$, this contradicts the assumption that the strategy profile $\left(\sigma\left(\boldsymbol{\xi}^{\mathsf{A}}\right), \sigma\left(\boldsymbol{\xi}^{\mathsf{B}}\right)\right)$ is a Nash Equilibrium of $\mathcal{SM}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f)$. Thus, the claim of the proposition is satisfied. □

2.4.3 The Number of Symmetrized Strategies

As before, to formalize the discussion on the size of strategy sets, we assume that for a given number of battlefields $n \geq 2$, the number of resources of a player P is given by a linear function $g^P(n) = a^P n + k^P$, where $a^P, k^P \in \mathbb{N}_+$. Hence, when increasing the number of battlefields in the model, the increase in the number of resources of a player P is defined by the function g^P.

We move on to present two observations concerning the number of pure symmetrized strategies when the number of resources is a linear function of the number of battlefields. The first observation states that the number of pure symmetrized strategies is asymptotically exponential in the number of battlefields.

Observation 2.8. The number of symmetrized strategies of a player P partitioning $g^P(n) = a^P + k^P$ resources into n integers, asymptotically satisfies:

$$\left|\sigma\left(S\left(g^P\left(n\right), n\right)\right)\right| \in \left[\mathcal{O}\left(\frac{1}{n}\exp\left(\pi\sqrt{\frac{2n}{3}}\right)\right), \mathcal{O}\left(\frac{1}{n}\exp\left(\pi\sqrt{\frac{2a^P n}{3}}\right)\right)\right]. \tag{2.12}$$

Proof. The number $p(n)$ of different partitions of n into the sum of non-negative integers has asymptotic growth of:

$$p\left(n\right) \sim \frac{1}{4n\sqrt{3}} \cdot \exp\left(\pi\sqrt{\frac{2n}{3}}\right), \tag{2.13}$$

as shown by Hardy and Ramanujan [1918]. The number of partitions of $g^P(n)$ into a sum of n non-negative integers is larger than the number of partitions $p(n)$, of n, and smaller than the number of partitions $p(g^P(n))$ of $g^P(n)$. Hence, the number of symmetrized strategies of a player with $g^P(n)$ resources to distribute over n battlefields asymptotically satisfies:

$$\left|\sigma\left(S\left(g^P\left(n\right), n\right)\right)\right| \in \left[\mathcal{O}\left(\frac{1}{n}\exp\left(\pi\sqrt{\frac{2n}{3}}\right)\right), \mathcal{O}\left(\frac{1}{n}\exp\left(\pi\sqrt{\frac{2a^P n}{3}}\right)\right)\right]. \tag{2.14}$$

□

The second observation states that although the number of pure symmetrized strategies is exponential in n, the reduction factor, i.e., the number of pure strategies divided by the number of pure symmetrized strategies, is also exponential in n.

Observation 2.9 (Reduction by an exponential factor)**.** The reduction factor between the strategy set S^P and the symmetrized strategy set $\sigma(S^P)$ of a player assigning $g^P(n) = a^P + k^P$ resources over n battlefields is exponential in n.

Proof. From Inequality 2.9 it follows, that the number of pure strategies of player P satisfies

$$\left|S\left(g^P\left(n\right), n\right)\right| \geq 2^n. \tag{2.15}$$

On the other hand, from Observation 2.8 follows, that there exists a constant C, such that for sufficiently large n's it holds that

$$\left|S\left(g^P\left(n\right), n\right)\right| < C \cdot \frac{1}{n} \exp\left(\pi\sqrt{\frac{2a^P n}{3}}\right). \tag{2.16}$$

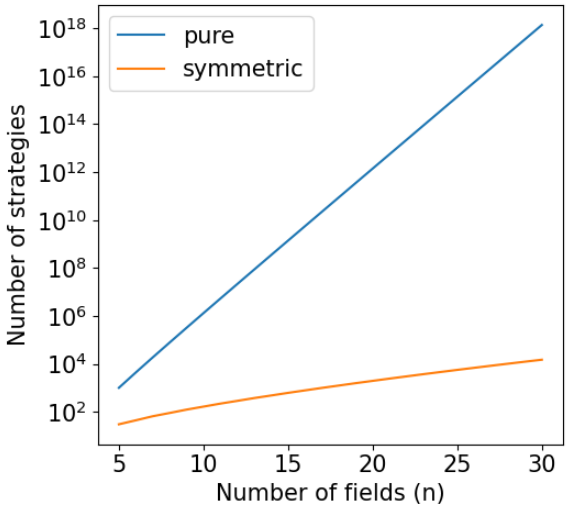
Hence, for sufficiently large n's it holds that

$$\frac{\left|S\left(g^P\left(n\right), n\right)\right|}{\left|\sigma\left(S\left(g^P\left(n\right), n\right)\right)\right|} > \frac{1}{C} \cdot n \cdot \exp\left(n\log 2 - \pi\sqrt{\frac{2a^P n}{3}}\right) \tag{2.17}$$

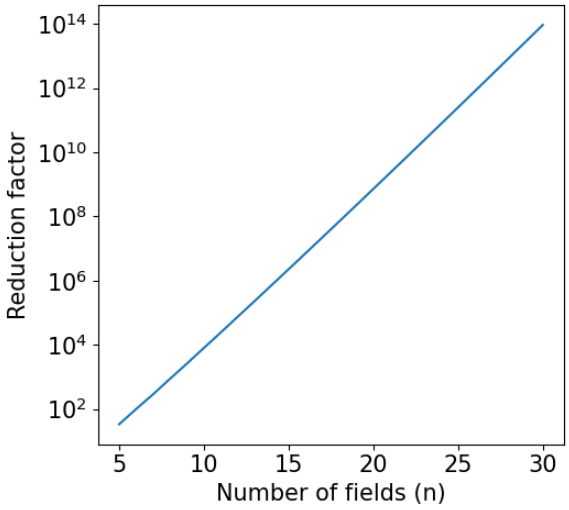
□

Figure 1(a) shows the comparison of the number of pure strategies and the number of symmetric strategies, when the number of resources is a linear function of the number of battlefields (here, the number of battlefields increased by 5). The vertical axis shows the number of strategies using a logarithmic scale.

Figure 1(b) shows the ratio of the number of pure strategies to the number of symmetric strategies (the reduction factor) for data in Figure 1(a). As stated by Observation 2.9, although the number of symmetric strategies is still exponential, the reduction is by an exponential factor with respect to the parameters of the model.



(a) The number of pure and symmetric strategies for $n+5$ resources.



(b) The reduction factor, i.e., the number of pure strategies divided by the number of pure symmetrized strategies.

Figure 1: Comparison of number of pure and pure symmetric strategies

2.5 Computation of Payoffs in the Symmetrized Models

In this section, we present an efficient method for computing payoffs in the symmetrized models. The payoff to player $P \in \{\mathsf{A}, \mathsf{B}\}$ under a strategy profile $\left(\sigma\left(\boldsymbol{s}^{\mathsf{A}}\right), \sigma\left(\boldsymbol{s}^{\mathsf{B}}\right)\right) \in \sigma\left(S^{\mathsf{A}}\right) \times \sigma\left(S^{\mathsf{B}}\right)$ is given by:

$$\Pi_f^P\left(\sigma\left(\boldsymbol{s}^{\mathsf{A}}\right), \sigma\left(\boldsymbol{s}^{\mathsf{B}}\right)\right) = \frac{1}{(n!)^2} \sum_{q \in P_n} \sum_{p \in P_n} \pi_f^P\left(q\left(\boldsymbol{s}^{\mathsf{A}}\right), p\left(\boldsymbol{s}^{\mathsf{B}}\right)\right) = \frac{1}{n!} \sum_{p \in P_n} \pi_f^P\left(\boldsymbol{s}^{\mathsf{A}}, p\left(\boldsymbol{s}^{\mathsf{B}}\right)\right). \quad (2.18)$$

Hence, to calculate the payoff under a symmetric strategy profile $\left(\sigma\left(\boldsymbol{s}^{\mathsf{A}}\right), \sigma\left(\boldsymbol{s}^{\mathsf{B}}\right)\right)$ in a naïve way, one needs to calculate the payoff of the original model for each permutation. We address this issue by providing an algorithm that computes payoffs under symmetric strategy profiles in polynomial time with respect to the number of battlefields, n.

2.5.1 Clash Matrix

By a clash matrix of a pair of symmetric strategies $\left(\sigma\left(\boldsymbol{s}^{\mathsf{A}}\right), \sigma\left(\boldsymbol{s}^{\mathsf{B}}\right)\right)$ we mean a matrix of dimension $n \times n$ defined as follows:

$$\boldsymbol{M}_{i,j}^{\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}}} = \operatorname{sign}\left(s_i^{\mathsf{A}} - s_j^{\mathsf{B}}\right).$$

That is, the matrix cell at coordinates (i, j) stores information (difference sign) of the result of clashing the i-th resource in player A's vector $\boldsymbol{s}^{\mathsf{A}}$ with the j-th resource in player B's vector $\boldsymbol{s}^{\mathsf{B}}$.

Example 2.10 (Clash matrix)**.** Consider a symmetrizes strategy profile $(\sigma(3,1,0),\sigma(2,2,0))$. The clash matrix of this profile is:

$$\boldsymbol{M}^{(3,1,0),(2,2,0)} = \begin{array}{c} \\ 3 \\ 1 \\ 0 \end{array} \begin{array}{c} \begin{array}{ccc} 2 & 2 & 0 \end{array} \\ \begin{pmatrix} 1 & 1 & 1 \\ -1 & -1 & 1 \\ -1 & -1 & 0 \end{pmatrix} \end{array}$$

Note that each permutation in Equation (2.18) can be used to select a subset of elements in matrix $\boldsymbol{M}^{\boldsymbol{s}^{\mathsf{A}},\boldsymbol{s}^{\mathsf{B}}}$ in such a way that exactly one element is selected from each row and each column. Formally:

$$\Pi_f^{\mathsf{A}}\left(\sigma\left(\boldsymbol{s}^{\mathsf{A}}\right),\sigma\left(\boldsymbol{s}^{\mathsf{B}}\right)\right) = \frac{1}{n!}\sum_{p\in P_n}\pi_f^{\mathsf{A}}\left(\boldsymbol{s}^{\mathsf{A}},p\left(\boldsymbol{s}^{\mathsf{B}}\right)\right) = \frac{1}{n!}\sum_{p\in P_n} f\left(\boldsymbol{M}^{\boldsymbol{s}^{\mathsf{A}},\boldsymbol{s}^{\mathsf{B}}}_{1,p(1)},\ldots,\boldsymbol{M}^{\boldsymbol{s}^{\mathsf{A}},\boldsymbol{s}^{\mathsf{B}}}_{n,p(n)}\right) \tag{2.19}$$

We can think of the problem of computing a single payoff under a symmetric strategy profile as follows. A matrix $\boldsymbol{M}^{\boldsymbol{s}^{\mathsf{A}},\boldsymbol{s}^{\mathsf{B}}}$ is a chessboard of dimension $n\times n$, divided into three distinct areas:

$$\begin{aligned} W^{\boldsymbol{s}^{\mathsf{A}},\boldsymbol{s}^{\mathsf{B}}} &= \{(i,j)\in[n]\times[n] : \boldsymbol{M}^{\boldsymbol{s}^{\mathsf{A}},\boldsymbol{s}^{\mathsf{B}}}_{i,j} = \ \ 1\},\\ T^{\boldsymbol{s}^{\mathsf{A}},\boldsymbol{s}^{\mathsf{B}}} &= \{(i,j)\in[n]\times[n] : \boldsymbol{M}^{\boldsymbol{s}^{\mathsf{A}},\boldsymbol{s}^{\mathsf{B}}}_{i,j} = \ \ 0\},\\ L^{\boldsymbol{s}^{\mathsf{A}},\boldsymbol{s}^{\mathsf{B}}} &= \{(i,j)\in[n]\times[n] : \boldsymbol{M}^{\boldsymbol{s}^{\mathsf{A}},\boldsymbol{s}^{\mathsf{B}}}_{i,j} = -1\}. \end{aligned}$$

Each permutation $p\in P_n$ corresponds to distributing n rooks on the mentioned chessboard of dimension $n\times n$ in such a way that no two rooks attack each other, i.e., each row contains exactly one rook and the rook in the i'th row is located in $p(i)$'th column. For $D\in\{W,T,L\}$, let $K_D(p\mid\boldsymbol{M})$ denote the number of rooks distributed in area D of clash matrix $\boldsymbol{M}$ by permutation p. Since for every battlefields outcome vector $\boldsymbol{u}\in\{-1,0,1,\}^n$, every aggregation function $f:\{-1,0,1,\}^n\longrightarrow\mathbb{R}$ is invariant under the permutations of $\boldsymbol{u}$, f can be defined as $f(\boldsymbol{u})=f_n(k_W,k_L)$, where k_W and k_L describe the number of 1s and -1s in vector $\boldsymbol{u}$, respectively. Hence:

$$\pi_f^{\mathsf{A}}\left(\boldsymbol{s}^{\mathsf{A}},p\left(\boldsymbol{s}^{\mathsf{B}}\right)\right) = f_n(K_W(p\,|\,\boldsymbol{M}),K_L(p\,|\,\boldsymbol{M}))\,. \tag{2.20}$$

Let $h(k_W,k_L\mid\boldsymbol{M})$ be the number of permutations $p\in P_n$ that place exactly k_W rooks in area W and k_L rooks in area L, for a given clash matrix $\boldsymbol{M}$, i.e.

$$h(k_W,k_L\,|\,\boldsymbol{M}) = \sum_{p\in P_n}[K_W(p\,|\,\boldsymbol{M})=k_W\wedge K_L(p\,|\,\boldsymbol{M})=k_L]. \tag{2.21}$$

Using function h and Equation (2.20), the payoff given by Equation (2.19) can be restated as:

$$\Pi_f^{\mathsf{A}}\left(\sigma\left(\boldsymbol{s}^{\mathsf{A}}\right),\sigma\left(\boldsymbol{s}^{\mathsf{B}}\right)\right) = \frac{1}{n!}\sum_{k_W=0}^{n}\sum_{k_L=0}^{n} h(k_W,k_L)\,f_n(k_W,k_L). \tag{2.22}$$

2.5.2 Recursive Formula

In this section, we derive a recursive formula for computing the values of the h function, defined in Equation (2.21). Note that due to the non-increasing ordering of the strategy vectors of both players, each column and row of a clash matrix is a weakly monotonic sequence of length n. Because of that, all clash matrices share a very particular structure. Area W occupies the upper part of the matrix, potentially reaching lower and lower in each successive column, and area T consists of rectangles, perhaps touching one another at the corners but never sharing a side.

To derive a recursive formula for function h, we "cut off" certain parts of the matrix (submatrix) with a certain number of distributed rooks and thus reduce the size of the considered matrix (submatrix).

Example 2.11 (Successive cut-offs of the clash matrix)**.** An example of the clash matrix with the successive "cut-offs", marked by their numbers (violet). Areas are colored as follows: yellow (L), blue (T), and red (W). An example of a strategy pair that yields the considered clash matrix is $s_A = (8, 8, 6, 5, 4, 2, 1, 1, 0)$, $s_B = (8, 8, 8, 7, 5, 3, 3, 1, 0)$.

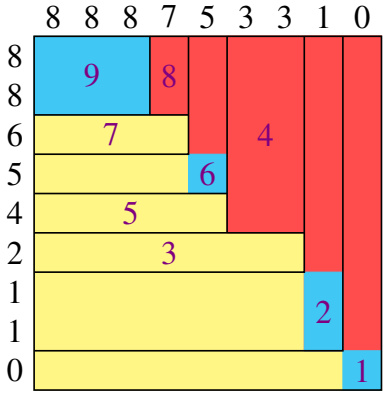


Figure 2: Clash matrix with successive cut-offs

We use the following notation.

$\boldsymbol{M}_{|(i,j)}$ – the submatrix of $\boldsymbol{M}$ consisting of the intersection of the first i rows and first j columns of $\boldsymbol{M}$.

$H(i, j, m, k_W, k_L \mid \boldsymbol{M})$ – the number of ways in which m rooks can be distributed in the intersection of the first i rows and j columns of the clash matrix $\boldsymbol{M}$, such that there are exactly k_W rooks in area W and exactly k_L rooks are in area L.

$R(i, j, t) = \binom{j}{t}\binom{i}{t}t!$ – the number of ways to distribute t rooks in a uniform area with i rows and j columns.

From the definition of H it follows that

$$h\left(k_W, k_L \mid \boldsymbol{M}\right) = H\left(n, n, n, k_W, k_L \mid \boldsymbol{M}\right).$$

We begin the description of the recursive formula with the description of the base case, that is, when the entire clash matrix is entirely within one of the three areas under consideration:

$$H_0(i, j, m, k_W, k_L \mid \boldsymbol{M}') = \begin{cases} R(i,j,m), & \text{if } k_W = m \text{ and } k_L = 0 \text{ and } \boldsymbol{M}' \subseteq W, \\ R(i,j,m), & \text{if } k_L = m \text{ and } k_W = 0 \text{ and } \boldsymbol{M}' \subseteq L, \\ R(i,j,m), & \text{if } k_L = k_W = 0 \text{ and } \boldsymbol{M}' \subseteq T, \\ 0, & \text{otherwise.} \end{cases} \tag{2.23}$$

The recursive formula is:

$$\begin{aligned} H\left(i, j, m, k_W, k_L \mid \boldsymbol{M}\right) = \sum_{\substack{r_1, r_2, r_3 \geq 0 \\ r_1 + r_2 + r_3 \leq m}} & H(i', j', m - r_s, k_W - r_3, k_L - r_1 \mid \boldsymbol{M}_{|(i',j')}) \\ & \cdot R(i - i', j' - (m - r_1 - r_2 - r_3), r_1) \\ & \cdot R(i - i' - r_1, j - j', r_2) \\ & \cdot R(i' - (m - r_1 - r_2 - r_3), j - j' - r_2, r_3). \end{aligned} \tag{2.24}$$

Proposition 2.12. *Fix $m \in \mathbb{N}$, such that $0 \leq m < \min\{i, j\}$. The recursive formula for $H\left(i, j, m, k_W, k_L \mid \boldsymbol{M}\right)$ satisfies Equation (2.24), with the base case defined by Equation (2.23).*

Proof. Consider a division of m rooks into four groups, as indicated below in Figure 3.

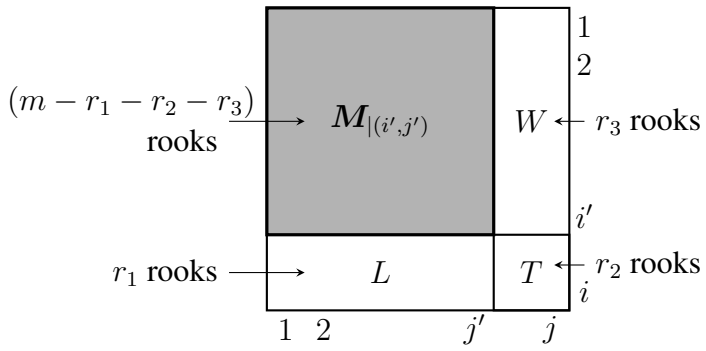


Figure 3: Cutting-off procedure

1. We place $(m - r_1 - r_2 - r_3)$ rooks in area $\boldsymbol{M}_{|(i',j')}$.
2. In the rectangle lying in L (where $(m - r_1 - r_2 - r_3)$ columns have already been excluded in 1), we place r_1 rooks in $i - i'$ available rows and $j' - (m - r_1 - r_2 - r_3)$ available columns.

3. In the rectangle lying in T (where r_1 rows have already been excluded in 2), we place r_2 rooks in $i - i' - r_1$ available rows and $j - j'$ available columns.

4. In the rectangle lying in W (where $(m - r_1 - r_2 - r_3)$ rows and r_2 columns have already been excluded in 1 and 3) we place r_3 rooks in $i' - (m - r_1 - r_2 - r_3)$ available rows and $j - j' - r_2$ available columns.

5. As r_3 rooks were placed in W and r_1 rooks were placed in L, we subtract those values from k_W and k_L respectively in the recursive call of H in point 1.

The recursive formula sums over every feasible division of rooks into four groups. □

2.5.3 Computation of Payoffs in Symmetrized Strategy Space

We proceed to prove that values of H function, defined by Equations (2.23) and (2.24), can be computed in polynomial time with respect to n, using dynamic programming. We work under the following assumption.

Assumption 1. Addition, subtraction, multiplication, and division take $\mathcal{O}(1)$.

First, we prove the following lemma, stating that the values of the H functions can be computed in polynomial time with respect to the size of the clash matrix.

Lemma 2.13. *Fix a clash matrix $\boldsymbol{M}$ of size $n \times n$. The two dimensional array $values[][]$ of size $(n+1) \times (n+1)$, such that $values[k_W][k_L] = h(k_W, k_L|\boldsymbol{M})$ can be computed in pessimistic (worst case) time $\mathcal{O}(n^7)$.*

Proof. During the recursion procedure on a clash matrix $\boldsymbol{M}$ of size $n \times n$, at most $2n - 1$ different matrices (including matrix $\boldsymbol{M}$) are considered, as at each recursive step at least one dimension of the considered matrix is decreased by at least 1, and the matrix of size 1×1 is guaranteed to be in only one of the three considered regions.

Let $\mathbb{M} = (\boldsymbol{M}^0, \boldsymbol{M}^1, \ldots, \boldsymbol{M})$ be the sequence of the at most $2n - 1$ matrices considered in the recursive formula given by Equation (2.24)

Fix a matrix $\boldsymbol{M}^d \in \mathbb{M}$ of size i, j.

Induction base: If $\boldsymbol{M}^d = \boldsymbol{M}^0$, then for each pair of k_W, k_L, such that $k_W + k_L \leq n$, the value of function H_0 defined by Equation (2.23) can be computed in $\mathcal{O}(n)$, as it requires computing a constant number of factorials of numbers not exceeding n. Since the number of considered pairs (k_W, k_L) is smaller than n^2, the overall pesymistic time is $\mathcal{O}(n^3)$.

Induction step: Assume that for all smaller matrices, all of the values of the H functions are computed in time of $\mathcal{O}(n^7)$. Fix k_W, k_L, such that $k_W + k_L \leq n$ and $m \leq \min\{i, j\}$. Equation (2.24) states that to compute $H(i, j, m, k_K, k_L, | \boldsymbol{M}^d)$ we need to consider every triple $r_1, r_2, r_3 \geq 0$ such that $r_1 + r_2 + r_3 \leq m$. There are fewer than $m^3 < n^3$ such triples. For each triple, the values of the H function for the $\boldsymbol{M}^{d-1}$ matrix are already computed. The single call of the R function takes time of $\mathcal{O}(n)$, and for a fixed triple r_1, r_2, r_3, the number of calls of

the R function is constant. Hence, for a fixed pair of k_W, k_L values, the number of arithmetic operations is at most $\mathcal{O}(n^4)$. Since there are no more then n^2 such pairs, the time of computing values of the function H for matrix $\boldsymbol{M}^d$ for every pair k_W, k_L is at most $O(n^6)$. From the induction assumption, the overall time is still at most $O(n^7)$. □

As a consequence of Lemma 2.13, the payoffs from symmetric strategy profiles can be computed in polynomial time with respect to the number of battlefields, n.

Proposition 2.14. *Impose Assumption 1. Fix symmetrized conflict with multiple battlefields $\mathcal{SM}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, f)$, and strategy profile $(\sigma(s^{\mathsf{A}}), \sigma(s^{\mathsf{B}}))$. The value $\pi_f^A(\sigma(s^{\mathsf{A}}), \sigma(s^{\mathsf{B}}))$ can be computed in time $\mathcal{O}(n^7)$.*

Proof. It follows from Lemma 2.13. Once the values of the h function are computed, Equation (2.22) can be computed in $\mathcal{O}(n^2)$. □

2.6 Computation of Nash Equilibria

In the previous section, we presented an algorithm for computing single payoffs from symmetric strategy profiles in polynomial time with respect to the model parameters. Sizes of the sets of symmetric strategies of both players, although much smaller, are still exponential with respect to the model parameters. Because of that, methods for solving zero-sum games, which require calculating the whole payoff matrix, such as LP-solvers, are still inefficient for large parameter values. For that reason, we use the Double Oracle Algorithm (DOA), which can be used for computing a Nash Equilibrium of zero-sum games, possibly without calculating the entire payoff matrix. Double Oracle Algorithm correctness was proven by McMahan et al. [2003]. The Double Oracle Algorithm is presented in Algorithm 1.

Algorithm 1 The Double Oracle Algorithm for Symmetrized Conflicts with Multiple Battlefields

Input: Pure strategy profile $(\sigma(\boldsymbol{s}^{\mathsf{A}}), \sigma(\boldsymbol{s}^{\mathsf{B}}))$ and an aggregation function f.
Output: NE of the zero-sum game.

1: $X^{\mathsf{A}} \leftarrow \sigma(\boldsymbol{s}^{\mathsf{A}})$.
2: $X^{\mathsf{B}} \leftarrow \sigma(\boldsymbol{s}^{\mathsf{B}})$.
3: **while** not convergence **do**
4: $\quad (\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}}) \leftarrow \text{CoreLP}(X^{\mathsf{A}}, X^{\mathsf{B}})$.
5: $\quad X^{\mathsf{A}} \leftarrow X^{\mathsf{A}} \cup \{best_response(\boldsymbol{\xi}^{\mathsf{B}})\}$.
6: $\quad X^{\mathsf{B}} \leftarrow X^{\mathsf{B}} \cup \{best_response(\boldsymbol{\xi}^{\mathsf{A}})\}$.
7: **end while**
8: **return** $(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}})$.

2.6.1 Oracle

When using DOA, one has to use an *oracle* that, for a given mixed strategy of a player P, returns a best-response pure strategy of the opponent. As already shown in Proposition 2.5, computing a best-response to a given mixed strategy of the opponent can not be done in polynomial time unless $\mathsf{P} = \mathsf{NP}$. Given a mixed strategy $\boldsymbol{\xi}^P$ of player P, we consider all strategy profiles $(\boldsymbol{\xi}^P, \boldsymbol{s}^{\text{-}P})$, where $\boldsymbol{s}^{\text{-}P}$ is a pure strategy of player $\text{-}P$, and find a pure strategy that maximizes the payoff to player $\text{-}P$. To compute these payoffs, we use a dynamic implementation of the recursive formula given by Equations (2.23) and (2.24). Given two sets of pure strategies, X^{A} and X^{B}, of players A and B, respectively, by CoreLP$(X^{\mathsf{A}}, X^{\mathsf{B}})$ we mean a strategy profile that is NE of the zero-sum game restricted to the pure strategies in X^{A} and X^{B}, computed by an LP-solver.

As we can choose the initial strategies of both players arbitrarily, we decided to start with the most even assignment between all the battlefields for each player.

2.6.2 Heuristic for Computing Best Response

In this section we propose a simple but effective heuristic that allows us to avoid unnecessary calculations for a subset of strategies. We first prove a simple observation regarding the pure strategies that are the best responses to a mixed strategy of the opponent. By a *best response* of player P to a mixed strategy $\boldsymbol{\xi}^{\text{-}P}$ of the opponent we mean a pure strategy $\boldsymbol{s}^P_{best}$ that maximizes the payoff of P against $\boldsymbol{\xi}^{\text{-}P}$. For a given mixed strategy $\boldsymbol{\xi}^P$, we define $\text{max_assign}(\boldsymbol{\xi}^P) \in N$, of player P as

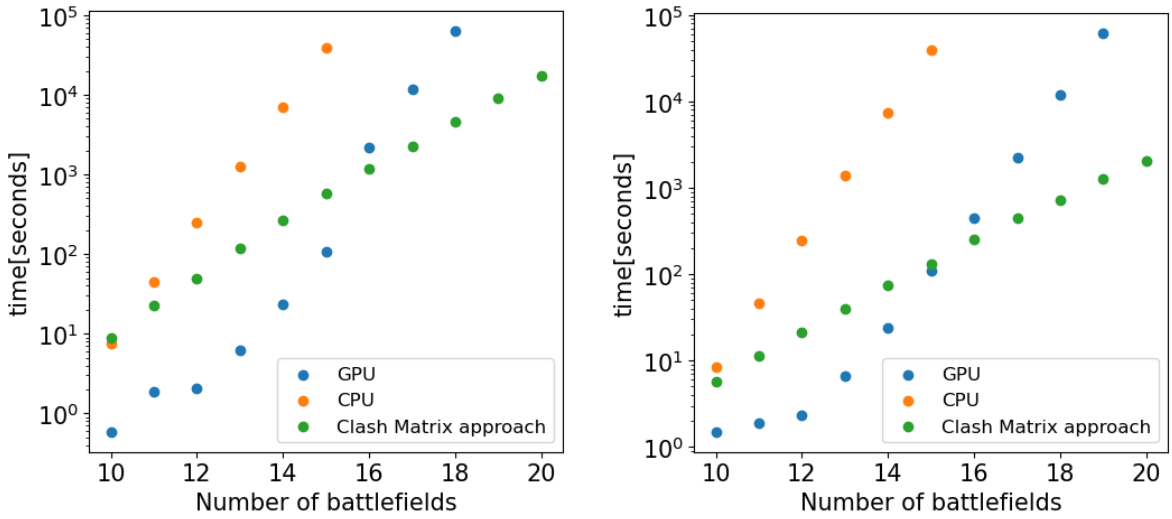
$$\max_{\substack{s \in S^P \\ \boldsymbol{\xi}^P(s) > 0}} \max_{i \in \{1,\ldots,n\}} s_i, \tag{2.25}$$

i.e., the largest number of resources that player P assigns to a single battlefield with positive probability when playing the strategy $\boldsymbol{\xi}^P$.

Proposition 2.15. *Let f be an aggregation function that is non-decreasing in its first argument and non-increasing in its second argument. If $|D^{\mathsf{A}} - D^{\mathsf{B}}| \leq n$, then for any mixed strategy $\boldsymbol{\xi}^P$ of player P, there exists a best response pure strategy $\boldsymbol{s}^{\text{-}P}_{best}$ of player $\text{-}P$, that satisfies*

$$\text{max_assign}(\boldsymbol{s}^{\text{-}P}_{best}) \leq \text{max_assign}(\boldsymbol{\xi}^P) + 1. \tag{2.26}$$

Proof. Take any pure strategy $\boldsymbol{s}^{\text{-}P}$ that maximizes payoff of player $\text{-}P$ against a given mixed strategy $\boldsymbol{\xi}^P$. If this strategy does not satisfy (2.26) then we can reduce the number of resources at battlefields where the number of resources exceeds $(\text{max_assign}(\boldsymbol{\xi}^P) + 1)$ to $(\text{max_assign}(\boldsymbol{\xi}^P) + 1)$, as those battlefields are still always won by $\text{-}P$, and increase the number of resources at the battlefields where the current assignment is below $(\text{max_assign}(\boldsymbol{\xi}^P) + 1)$. As the aggregation function is monotonic non-decreasing in its first argument and monotonic non-increasing in its second argument, the payoff of player $\text{-}P$ does not decrease after such a redistribution. □



(a) More than opponent

(b) Majoritarian

Figure 4: Run times of computing symmetrized payoff matrices with respect to the number of battlefields, n. The number of resources of each player is $n + 5$.

By Proposition 2.15, when computing a best response to a given mixed strategy $\boldsymbol{\xi}^P$, we can restrict attention to the pure strategies of the opponent for which the $\mathrm{max_assign}$ value exceeds the $\mathrm{max_assign}$ value of the mixed strategy ξ^C by at most one. This significantly reduces the number of necessary calculations and therefore the time required to compute the NE of the game in question.

We move on to present the results of experimental evaluation of the proposed methods.

2.6.3 Symmetrized Payoffs Computation

We implemented a dynamic program using the clash matrix approach for computing payoff matrices of the symmetrized games with the aggregation functions given by Equation (2.6) and Equation (2.7). The implementation in using $C{+}{+}$'s *vector* class from $C{+}{+}$'s standard library. We compared the acquired run times of the algorithm with CPU and GPU-based approaches that calculate payoffs based on (2.18) for finding payoff matrices of the symmetrized game. Experiments were run on a computational cluster *Entropy* with a central processing unit (CPU) Intel Xeon E5-2640 v4 (10 cores) and the GPU unit Titan V. Four cores of the processor were used while conducting each of the experiments. Every reported time result is an average from 10 runs of the program. The program we used to obtain payoff matrices with CPU and GPU is proposed and described in Litwin [2020]. Figure 4 shows the comparison of the times required by all three methods for computing the entire payoff matrices of both games. The number of battlefields is shown on the horizontal axis. The number of resources of each player is a linear function of the number of battlefields (the number of battlefields increased by 5). The vertical axis shows the run times of the programs using a logarithmic scale. The time limit for each

experiment was set to 10^5 seconds.

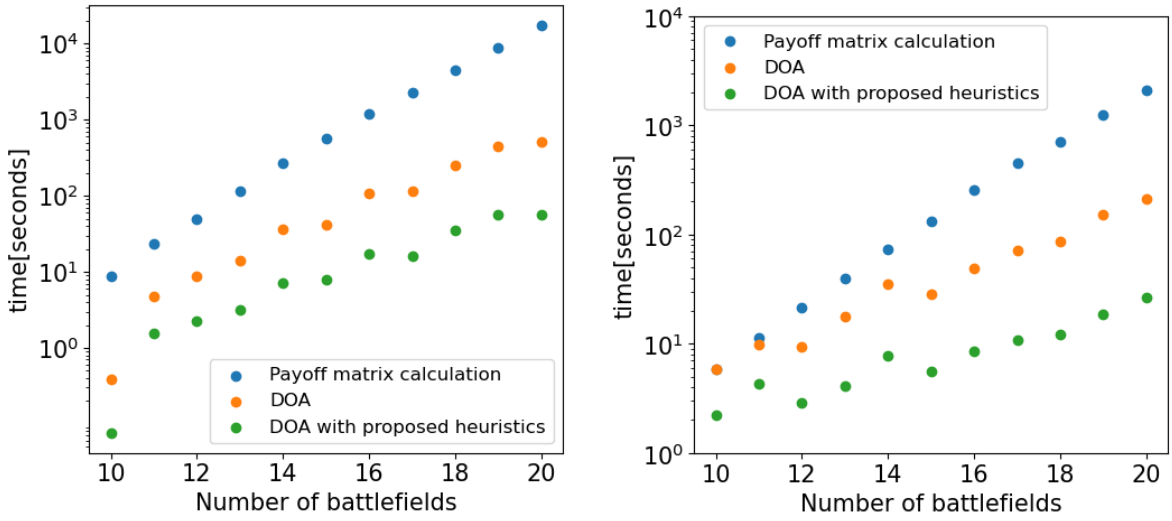
When considering the *more than opponent* (Figure 4(a)) game, the maximal value of parameters that GPU based approach was able to calculate the payoff matrix for was 23 resources to distribute over 18 battlefields. The achieved speedup of the clash matrix for the same parameters was 18 times. For the *majoritarian* (Figure 4(b)) game, the maximal value of parameters that GPU based approach was able to calculate the payoff matrix for was 24 resources to distribute over 19 battlefields. The achieved speedup of the clash matrix for the same parameters was 48 times. The speedup for 23 resources to distribute over 18 battlefields was 16 times (similar for both aggregation functions).

2.6.4 Equilibrium Computation

We proceed to report the experimental results on computing Nash equilibria of the games in question using the Double Oracle Algorithm. We use the time required for computing the entire payoff matrix as an optimistic time bound of every method that requires the entire payoff matrix to solve the game (e.g. the standard LP-solver approach). All of the programs use a symmetrized model and calculate single payoffs using the clash matrix.

Figure 5 presents a comparison of the time results of the calculation of the entire payoff matrix of the games in question, with the two versions of DOA (with and without the proposed heuristic) that compute a Nash Equilibrium of the given game. The number of battlefields is shown on the horizontal axis. The number of resources of each player is once again a linear function of the number of battlefields (the number of battlefields increased by 5). We use the state-of-the-art *Gurobi* Gurobi Optimization, LLC [2022] LP-solver for computing the CoreLP function used in Double Oracle Algorithm.

For the *more than opponent* game (Figure 5(a)), speed-up achieved for the game where both players have 25 resources to distribute over 20 battlefields by the DOA algorithm when compared to the LP-solver is 34 times. In contrast, when using the DOA algorithm with the heuristic that we propose, one gets a speed-up of more than 300 times. When considering the *more than opponent* game (Figure 5(b)), the speed-up achieved for the game where both players have 25 resources to distribute over 20 battlefields by the DOA algorithm when compared to the LP-solver is 10 times. In contrast, when using the DOA algorithm with the heuristic that we propose, one gets a speed-up of more than 80 times. This shows that the proposed heuristic, although simple, is very effective in both models. When comparing the LP-solver-based approach, where payoff matrices are computed using a GPU, and the clash matrix-based DOA approach, using the proposed heuristic, the achieved speed-up for the game where both players have 25 resources to distribute over 20 battlefields would be more than 3000 times for both models. Unfortunately, the run time of the GPU-based method for finding the payoff matrix takes too long for us to know the exact speedups that could be achieved.



(a) More than opponent

(b) Majoritarian

Figure 5: Run times of different algorithms

2.7 Conclusions

In this chapter, we examined symmetric conflicts with multiple battlefields under non-bilinear aggregation functions such as the majoritarian and chopstick rules. We established that computing a pure best response in these settings cannot be achieved in polynomial time unless $\mathsf{P} = \mathsf{NP}$, demonstrating a sharp complexity contrast with bilinear Blotto-like games.

To address this difficulty, we introduced a symmetrized formulation of the model, which reduces the effective strategy space by an exponential factor. We further developed the *clash matrix* algorithm, enabling the computation of payoffs from symmetrized strategy profiles in polynomial time. Combined with the Double Oracle Algorithm and a monotonicity-based heuristic for oracle selection, this framework allows for efficient computation of Nash equilibria despite the exponential growth of the underlying strategy sets.

Chapter 3

Attack and Interception Games in Networks

Chapter Overview

In this chapter, we study the computation of Nash equilibria in attack and interception games in networks, introduced by Bloch et al. [2023]. In these games, an attacker selects a path through the network while multiple defenders allocate interception efforts across their respective nodes. Although the attacker's strategy space—consisting of all simple paths in the network—may grow exponentially with the number of nodes, we show that equilibrium computation remains tractable for a broad class of cost functions. By exploiting the structural properties of networks and the interplay between attacker and defender strategies, we develop a polynomial-time algorithm for finding Nash equilibria in this setting.

The rest of this chapter is organized as follows. In Section 3.1, we provide real-life motivations for our study of the attack and interception games in networks. The model is introduced in Section 3.2. Section 3.3 provides a discussion on the already established results of Bloch et al. [2023]. In Section 3.4, we give preliminary results that are crucial for our analysis. Section 3.5 concerns the equilibria properties following from the network structure. In Section 3.6, we present how to apply the results concerning the network structure to find the set of targeted nodes and, in consequence, to compute a Nash equilibrium of the game in question. Finally, in Section 3.7, we prove that for a broad class of cost functions, our method runs in polynomial time with respect to the number of nodes in the network.

3.1 Motivating Examples

In the world of international security and intelligence, adversaries often seek to infiltrate high-value institutions—such as government agencies, research centers, or corporate boards—by strategically maneuvering through layers of access and trust. For example, a foreign agent aiming to manipulate a policy decision or steal classified information might begin by approaching

a loosely connected staff member, gradually moving through higher levels of the organizational hierarchy to reach the intended target. Each intermediary, from assistants to department heads, serves as a potential gatekeeper: they may scrutinize requests or interactions, but doing so requires effort, attention, and resources. Only the final target—say, a senior executive or agency director—is directly harmed if the infiltration is successful.

A similar structure emerges in the context of transnational smuggling. Here, a trafficker attempts to deliver contraband—such as drugs, weapons, or endangered goods—into a lucrative market by selecting a route across multiple national borders. Each country along the path can choose its own level of enforcement, such as inspection intensity or border patrol strength, but these measures come at a cost. While every country has an incentive to monitor traffic, only the final destination suffers the full impact of a successful smuggling attempt—be it social harm, crime, or market destabilization—making defensive investments a matter of strategic self-interest rather than global coordination.

An analogous pattern appears in cyber operations targeting secure systems. To breach a protected server or data vault, a hacker may need to traverse a multi-layered digital infrastructure, bypassing firewalls, user authentication, and access-control mechanisms. Each subsystem in the attack path can independently invest in intrusion detection or security protocols. However, intermediate systems themselves are not the objective—they serve only as stepping stones toward the core asset. If the final node is reached undetected, the organization suffers a breach, but if the attacker is intercepted along the way, the damage is prevented.

Despite differences in domain—spanning physical, geopolitical, and digital environments—these scenarios share a common strategic structure:

- The attacker selects a **specific target** within a network and chooses a **fixed path** to reach it.
- **Only the final target** is harmed if the attack succeeds; intermediaries are affected only through their defensive efforts.
- Each node acts **independently**, choosing its own level of protection based on localized trade-offs between risk and cost.
- Movement through the network is **constrained by edges**, requiring the attacker to traverse adjacent nodes en route to the target.

A formal model capturing these types of interactions exists in the literature: the *attack and interception game in networks*, introduced by Bloch et al. [2023]. This game-theoretic framework encapsulates the strategic tension between a single attacker choosing a path to a target node and a set of defenders who independently invest in protection against interception. Given its relevance to a wide range of practical applications—from espionage and cyberattacks to smuggling and secure communication systems—it is highly desirable to be able to compute solutions of this model efficiently.

One of the central challenges in solving attack and interception games arises from the *exponential* size of the attacker's strategy space. Specifically, the attacker can choose any simple path originating from a given source node, and the number of such paths can grow factorially with the number of nodes in the network (e.g., $n!$ in the case of a clique). This explosion of possible strategies places the model within the class of *succinct games*, where the representation

of the game is compact, but the underlying strategy space is large.

In this chapter, we study the computation of Nash equilibria in attack and interception games in networks. We show that for a broad and natural class of cost functions, equilibrium strategies can be computed in *polynomial time* with respect to the number of nodes in the network.

3.2 Model

The attack and interception game in a network, introduced by Bloch et al. [2023], is played by $n+1$ individuals, where $n \geq 1$. The first player is called the *attacker*, and referred to as the player 0. The remaining n players are called *defenders*. The set of defenders is denoted by $[n] = \{1, 2, \ldots, n\}$. The attacker and the defenders are connected in a network modeled by an exogenously given undirected graph, $G = \langle V, E \rangle$, where $V = [n] \cup \{0\}$ is the set of nodes and $E \subseteq \binom{V}{2}$ is the set of edges ($\binom{V}{2}$ denotes the set of all 2-element subsets of V). We assume that graph G is connected, meaning that every target is reachable from the attacker by some simple path p (i.e. a path where every node appears at most once). Throughout this chapter, when using the term path, we mean a simple path. Given a graph G, we will use $E(G)$ to denote the set of edges in G and $V(G)$ to denote the set of nodes in G.

A strategy of the attacker is a path of attack, which is any path in the network G that originates at the node 0. The terminating node of the path of attack represents the target of the attack, while the rest of the path represents the route leading to the targeted node. Each defender, anticipating a possible attack, chooses an interception probability, i.e., the probability of intercepting an attack. Intercepting the attack means stopping the attacker, regardless of whether the defender is the target or simply lies on an attack path. The interception probability of defender j, $x_j \in [0, 1]$, comes at a cost, $c_j(x_j)$. Defenders choose their interception probabilities independently and simultaneously, while the attacker chooses an attack path p. Formally, the strategy set of the attacker is the set of all paths in G originating at 0, denoted by $P(G)$, while the strategy set of every defender $j \in [n]$ is the interval $[0, 1]$. A tuple of players strategies, $(p, \boldsymbol{x}) \in P(G) \times \mathbb{R}^n$ is called a strategy profile.

For any attack path $p \in P(G)$, let $t(p)$ denote the terminal node of path p, i.e., the targeted node. Given an attack path $p \in P(G)$ and $j \in p$, the set of predecessors of j in p is $Pred(p, j) = \{k \in p : \ k \text{ lies on path } p \text{ between 0 and } j\ \}$. Fix a strategy profile $(p, \boldsymbol{x}) \in P(G) \times \mathbb{R}^n$ and a node j on path p. The probability that the attack along p reaches node j is

$$\alpha_j(p, \boldsymbol{x}_{-j}) = \prod_{k \in Pred(p,j)} (1 - x_k),$$

while the probability that the attack on the targeted node, $t(p)$ is successful is

$$\beta_{t(p)}(p, \boldsymbol{x}) = \alpha_{t(p)}(p, \boldsymbol{x}_{-t(p)}) \cdot (1 - x_{t(p)}).$$

For each defender $j \in [n]$, there are two strictly positive values, b_j and d_j. The value $b_j > 0$ describes the defender's strategic importance to the attacker, i.e., the utility of the attacker

from a successful attack on node i. Hence, the payoff to the attacker under a strategy profile $(p, \boldsymbol{x}) \in P(G) \times \mathbb{R}^n$ is

$$U(p, \boldsymbol{x}) = \beta_{t(p)}(p, \boldsymbol{x}) b_{t(p)}.$$

On the other hand, the value $d_j > 0$ represents the defender's loss upon a successful attack. A defender i loses his value d_j if and only if i is the targeted node and the attack is successful. Hence, the payoff to the defender j under a strategy profile $(p, \boldsymbol{x}) \in P(G) \times \mathbb{R}^n$ is

$$V_j(p, \boldsymbol{x}) = \begin{cases} \beta_j(p, \boldsymbol{x})(-d_j) - c_j(x_j), \text{ if } j = t(p), \\ -c_j(x_j), \text{ otherwise.} \end{cases} \tag{3.1}$$

Definition 3.1 (Attack and Interception Game in a Network Bloch et al. [2023])**.** Quadruple $(G, \boldsymbol{b}, \boldsymbol{d}, \boldsymbol{c})$ defines an attack and interception game in a network with network G, set of players $V(G)$, defenders' cost functions, $\boldsymbol{c}$, attacker's evaluations, $\boldsymbol{b}$, and defenders' evaluations $\boldsymbol{d}$.

We allow the attacker to use mixed strategies, choosing a probability distribution $\boldsymbol{\pi}$ over all paths in $P(G)$. Let $\Delta(P(G))$ denote the set of all probability distributions over $P(G)$. The expected payoff of the attacker from a mixed strategy profile $(\boldsymbol{\pi}, \boldsymbol{x}) \in \Delta(P(G)) \times [0, 1]^n$ is

$$U(\boldsymbol{\pi}, \boldsymbol{x}) = \sum_{p \in P(G)} \boldsymbol{\pi}(p) \beta_{t(p)}(p, \boldsymbol{x}) b_{t(p)}, \tag{3.2}$$

while the expected payoff to defender $j \in [n]$ is

$$V_j(\boldsymbol{\pi}, \boldsymbol{x}) = \sum_{\substack{p \in P(G) \\ t(p) = j}} \boldsymbol{\pi}(p) \alpha_j(p, \boldsymbol{x}_{-j})(1 - x_j)(-d_j) - c_j(x_j). \tag{3.3}$$

We are interested in computing (mixed strategy) Nash equilibria (NE) of attack and interception games in a network defined by Definition 3.1. A strategy profile $(\boldsymbol{\pi}^*, \boldsymbol{x}^*)$ is a NE if and only if for every mixed strategy $\boldsymbol{\pi} \in \Delta(P(G))$ of the attacker, $U(\boldsymbol{\pi}^*, \boldsymbol{x}^*) \geq U(\boldsymbol{\pi}, \boldsymbol{x}^*)$, and for every defender $j \in [n]$ and every strategy $x_j \in [0, 1]$, $V_j(\boldsymbol{\pi}^*, (x_j, \boldsymbol{x}^*_{-j})) \leq V_j(\boldsymbol{\pi}^*, \boldsymbol{x}^*)$.

We move on to present the assumptions that we impose in our analysis. The first assumption states that no two defenders have the same strategic importance to the attacker.

Assumption 2. For any two defenders j and k, $b_j \neq b_k$.

Without loss of generality, we assume that the defenders are numbered in increasing order with respect to their strategic importance to the attacker, that is, $j < k \implies b_j < b_k$. The second assumption characterizes a class of cost functions we allow in the model, that is, that every cost function is increasing, strictly convex, with derivative at 0 equal to 0 and derivative at 1 smaller than the node's loss upon a successful attack.

Assumption 3. Every cost function $c_j(x_j)$ of defender $j \in [n]$ is increasing, differentiable, and strictly convex. Moreover, $c'_j(1) > d_j$ and $c'_j(0) = 0$.

3.3 Opening Discussion

In this section, we discuss the already known results concerning the attack and interception games in networks and identify their shortcomings. Moreover, we present our contribution and give the outline of the rest of the chapter.

3.3.1 Overview of Existing Results

The game of attack and interception in a network was introduced by Bloch et al. [2023]. Their results are obtained under three simplifying assumptions:

1. For every defender, the strategic importance of his node for the attacker is equal to the defender's loss upon a successful attack, i.e., $\forall_{i\in[n]} b_i = d_i$,
2. the values are normalized to be smaller than 1, i.e., $\forall_{i\in[n]} b_i < 1$,
3. the cost function of every defender is half of the square of the interception probability, i.e., $\forall_{i\in[n]} c_i(x_i) = \frac{x_i^2}{2}$.

In their formulation, the authors disregard neutral nodes—those that are never attacked in any equilibrium. Because such nodes maximize their payoff by setting their interception probability to zero, their only strategic role is to maintain connectivity within the network. Furthermore, an attack path is uniquely characterized by the sequence of targeted (i.e., non-neutral) nodes it traverses, and two paths that include the same sequence of targeted nodes are treated as identical. The results of Bloch et al. [2023] concerning the Nash equilibria of the attack and interception game in a network are:

1. **Theorem 1** The game of attack and interception in networks always admits a Nash equilibrium in mixed strategies.
2. **Theorem 2** Given Assumption 2, the game of attack and interception admits a unique Nash equilibrium.

The authors prove that in every Nash equilibrium, every targeted node in the network is targeted through a unique *equilibrium attack path*, and provide the set of first-order conditions on the unique Nash equilibrium of the game in question. Finally, they present a procedure for finding the set of targeted nodes when the network in question is a linear graph. As for the general network structure, the authors left the question of identifying the set of targeted nodes open. In consequence, up to this point, there is no known method of computing mixed strategy Nash equilibria of the game in question.

Although the simplifying assumptions may seem quite restrictive, the authors claim, informally, that their results hold for any game in question, such, that

1. no two defenders have the same strategic importance for the attacker (Assumption 2),
2. every cost function c_i is strictly increasing and strictly convex.

3.3.2 Counterexample

In this section, we illustrate with a counterexample that the results of Bloch et al. [2023] do not extend in line with their informal claim. In more detail, we present an instance of the game in question, where Theorem 2 of Bloch et al. [2023] does not hold, although no two defenders have the same strategic importance for the attacker, and every cost function is strictly increasing and strictly convex.

Example 3.2. Consider an attack and interception game in a network, Γ, such that

- $n = 2$,
- $E(G) = \{\{0, 1\}, \{1, 2\}\}$,
- $\boldsymbol{b} = \boldsymbol{d} = (3, 10)$,
- $c_1(x_1) = x_1^2$,
- $c_2(x_2) = 10x_2^2$.

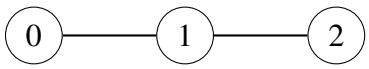


Figure 6: Network considered in Example 3.2

We claim, that both strategy profiles $(p^1 = (0, 1), x_1 = 1, x_2 = 0)$ and $(p^2 = (0, 1, 2), x_1 = 0, x_2 = \frac{1}{2})$ are Nash equilibria of Γ, contradicting the informal claim of Bloch et al. [2023] about the extension of their Theorem 2 to the general setting.

Proof. Consider strategy profile $(p^1 = (0, 1), x_1 = 1, x_2 = 0)$. The attacker can not increase her payoff by deviating, as the payoff from strategy profile $(p^2 = (0, 1, 2), x_1 = 1, x_2 = 0)$ to the attacker is also 0. Defender 2 maximizes his payoff by choosing the interception probability $x_2 = 0$. Consider defender 1. The payoff function to 1 if

$$V_1(p^1, x_1, x_2 = 0) = -(1 - x_1) \cdot 3 - x_1^2,$$

with derivative

$$\frac{\partial V_1(p^1, x_1, x_2 = 0)}{\partial x_1} = 3 - 2 \cdot x_1.$$

Since $\frac{\partial V_1(p^1, x_1, x_2 = 0)}{\partial x_1} > 0$ for $x_1 \in [0, 1]$, hence, $x_1 = 1$ is the unique maximizer of defender 1's payoff given the strategy profile $(p^1, x_2 = 0)$ of the remaining players.

Consider strategy profile $(p^2 = (0, 1, 2), x_1 = 0, x_2 = \frac{1}{2})$. The payoff to the attacker is

$$U(p^2 = (0, 1, 2), x_1 = 0, x_2 = \frac{1}{2}) = (1 - \frac{1}{2}) \cdot 10 = 5.$$

The attacker can not increase her payoff by deviating, as the payoff from the strategy profile $(p^1 = (0,1), x_1 = 0, x_2 = \frac{1}{2})$ is 3. Defender 1 maximizes his payoff by choosing the interception probability $x_1 = 0$, as he is not targeted. Consider defender 2. The payoff to the defender 2 is

$$V_2(p^2, x_1 = 0, x_2) = -(1 - x_2) \cdot 10 - 10 \cdot x_2^2,$$

with derivative

$$\frac{\partial V_2(p^2, x_1 = 0, x_2)}{\partial x_2} = 10 - 20 \cdot x_2.$$

Since $V_2(p^2, x_1 = 0, x_2)$ is a strictly concave function, and

$$\frac{\partial V_2(p^2, x_1 = 0, x_2)}{\partial x_2} \left(\frac{1}{2}\right) = 0,$$

$x_2 = \frac{1}{2}$ is the unique maximizer of the defender's 2 payoff given the strategy profile $(p^2, x_1 = 0)$ of the remaining players. □

3.3.3 Our Contribution

We move on to present the contribution of this chapter. We introduce an algorithm for computing a Nash equilibrium of the attack and interception games in networks in the generalized setting, that is, for any game satisfying Assumptions 2 and 3. Moreover, we provide an analysis of the pessimistic time complexity of our algorithm. In more detail, we characterize a necessary condition on the cost functions for the proposed algorithm to run in polynomial time. We show that for any homogeneous quadratic cost functions satisfying Assumption 3, the pessimistic time complexity of our algorithm is $\mathcal{O}(n^3)$. As a part of our analysis leading to the mentioned algorithm, we formally prove Theorem 2 of Bloch et al. [2023] for the class of games satisfying Assumptions 2 and 3.

3.4 Preliminary Results

In this section, we present three key observations that follow from Assumptions 2 and 3 and are central to our subsequent analysis. For a given strategy profile $(\boldsymbol{\pi}, \boldsymbol{x})$, let $\gamma_j(\boldsymbol{\pi}, \boldsymbol{x}_{-j})$ denote the probability of a successful attack targeted at node j under strategy profile $(\boldsymbol{\pi}, \boldsymbol{x})$, i.e.,

$$\gamma_j(\boldsymbol{\pi}, \boldsymbol{x}_{-j}) = \sum_{\substack{p \in P(G) \\ t(p) = j}} \boldsymbol{\pi}(p) \alpha_j(p, \boldsymbol{x}_{-j}).$$

We proceed to provide a preliminary result concerning the players' behavior in any Nash equilibrium of the game in question. In more detail, we prove that in any Nash equilibrium: no defender chooses an interception probability of 1; a defender's interception probability is positive if and only if it is targeted with positive probability; and the payoff to the attacker is always strictly positive.

Lemma 3.3. *Impose Assumptions 2 and 3. Fix game* Γ*. Fix strategy profile* $(\boldsymbol{\pi}^*, \boldsymbol{x}^*)$ *that is a Nash equilibrium of* Γ*. Fix defender* $j \in [n]$*. It holds that*

(i) $x_j^* < 1$,

(ii) $x_j^* > 0 \iff \gamma_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*) > 0$,

(iii) $U(\boldsymbol{\pi}^*, \boldsymbol{x}^*) > 0$.

Proof. We prove each part in turn:
Point (i) $x_j^* < 1$:

Consider the defender j's payoff function under opponents' strategy profile $(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*)$:

$$V_j(\boldsymbol{\pi}^*, \boldsymbol{x}^*) = \gamma_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*)(1 - x_j)(-d_j) - c_j(x_j).$$

The derivative with respect to x_j is:

$$\frac{\partial V_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*, x_j)}{\partial x_j} = \gamma_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*) \cdot d_j - c_j'(x_j).$$

Since $c_j(x_j)$ is strictly convex and increasing (by Assumption 3), the payoff function V_j is strictly concave in x_j. Moreover,

$$\frac{\partial V_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*, x_j)}{\partial x_j}(0) = \gamma_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*) \cdot d_j - c_j'(0) \geq 0,$$
$$\frac{\partial V_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*, x_j)}{\partial x_j}(1) = \gamma_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*) \cdot d_j - c_j'(1) \leq d_j - c_j'(1).$$

By Assumption 3, the second term is strictly negative and, in consequence, the maximizing argument x_j^* satisfies $x_j^* < 1$.
Point (ii) $x_j^* > 0 \implies \gamma_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*) > 0$:

Assume $x_j^* > 0$. Since the strategy profile $(\boldsymbol{\pi}^*, \boldsymbol{x}^*)$ is a Nash equilibrium, the defender j cannot increase his payoff by deviating to strategy $x_j = 0$. Hence, it has to hold that node j is successfully attacked with positive probablity under $(\boldsymbol{\pi}^*, \boldsymbol{x}^*)$, i.e., $\gamma_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*) > 0$.
Point (ii) $x_j^* > 0 \impliedby \gamma_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*) > 0$:

Assume $\gamma_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*) > 0$. It holds that

$$\frac{\partial V_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*, x_j)}{\partial x_j}(0) = \gamma_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*) \cdot d_j - c_j'(0) > 0,$$
$$\frac{\partial V_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*, x_j)}{\partial x_j}(1) = \gamma_j(\boldsymbol{\pi}^*, \boldsymbol{x}_{-j}^*) \cdot d_j - c_j'(1) \leq d_j - c_j'(1).$$

By Assumption 3, the first term is strictly positive, and the second term is strictly negative. Hence, the maximizing argument $x_j^* > 0$.
Point (iii) $U(\boldsymbol{\pi}^*, \boldsymbol{x}^*) > 0$:

By Point (i), every strategy $p \in P(G)$ results in a strictly positive payoff to the attacker under the strategy profile $(\boldsymbol{x}^*)$ of defenders.

□

3.5 Network Structure

In this section, we present our results concerning equilibria properties following from the network structure. In more detail, in Section 3.5.1 we characterize a subset of nodes that are never attacked in any Nash equilibrium of the game. We refer to these nodes as *linkers*, as their primary strategic importance lies in providing connectivity within the network. Building on that observation, we introduce a procedure of network reduction in Section 3.5.2. Network reduction allows for removing nodes that are never attacked in any Nash equilibrium, while guaranteeing a correspondence between every Nash equilibrium of the original game and the game on the reduced network. In Section 3.5.3, building on the notion of network reduction, we present a procedure that allows for obtaining a network without any linkers, while maintaining a correspondence of the original game and the game restricted to the resulting network.

3.5.1 Linkers

In this section, we characterize a type of node that is never attacked in any Nash equilibrium of the game in question. For a given game Γ, let $D(\Gamma) \equiv D$ denote the set of nodes that are attacked in any Nash equilibrium of Γ. We call elements of D *targeted nodes*, while referring to the nodes in $[n] \setminus D$ as *neutral nodes*.

To identify which nodes can never be targets in equilibrium, we define a structural class of nodes that we call *linkers*. Let us call a node m a *linker* if it is not directly connected to the attacker and all of its neighbors' valuations, b_j, are greater than b_m. Formally, node m is a linker in a network G if, and only if, $\{0, m\} \notin E(G)$ and $(\{m, j\} \in E(G) \implies b_j > b_m)$.

Proposition 3.4. *Under Assumptions 2 and 3 every linker is a neutral node.*

Proof. Impose Assumptions 2 and 3. Fix game Γ. Fix a Nash equilibrium $(\boldsymbol{\pi}^*, \boldsymbol{x}^*)$ of Γ. Assume, to the contrary, that there exists a linker $m \in [n]$ and an attack path $p \in P(G)$ with $t(p) = m$ such that $\boldsymbol{\pi}(p) > 0$, that is, m is targeted with positive probability. Let $p = (0, p_i, \ldots, k(m), m)$, i.e., node $k(m)$ is the direct predecessor of node m on path p. Notice that the probability $\beta_m(m, p)$ of the successful attack on node m through path p satisfies

$$\beta_m(p, \boldsymbol{x}^*) = (1 - x_m)\beta_{k(m)}(p', \boldsymbol{x}^*) \leq \beta_{k(m)}(p', \boldsymbol{x}^*),$$

where $p' = (0, p_i, \ldots, k(m))$ is the path p stripped from the last node, m. From Lemma 3.3 it follows, that $\beta_m(p, \boldsymbol{x}_{-m}) > 0$, as no node chooses the interception probability of 1 in any Nash equilibrium. As $b_m < b_{k(m)}$, by the definition of a linker, the strategy $p' = (0, p_i, \ldots, k(m))$ yields a strictly greater payoff to the attacker. Hence, p is not in the support of the attacker's equilibrium strategy $\boldsymbol{\pi}^*$, which contradicts the initial assumption that the node m is targeted with positive probability under $\boldsymbol{\pi}^*$. □

3.5.2 Network Reduction

Having defined the linker nodes, we proceed to define an operation of network reduction. The idea is straightforward—we want to be able to remove neutral nodes from the network, while

retaining the correspondence between equilibrium sets of the two games. To do so, we define the network reduction in the following way—a node is removed from the network, but every pair of its neighbors that were not neighbors are connected with an edge.

Definition 3.5 (Network Reduction)**.** For a given game network G and a node $m \in V(G)$, we define a network $G \setminus m$, such that:

- the set of nodes of $G \setminus m$ is $V(G \setminus m) = V(G) \setminus \{m\}$,
- the set of edges is $E(G \setminus m) = E(G) \setminus \{\{i, m\} : i \in V(G)\} \cup \{\{i, k\} : i \neq k \wedge \{i, m\} \in E(G) \wedge \{m, k\} \in E(G)\}$.

We call the resulting network $G \setminus m$ network G reduced by the node m.

We move on to prove that reducing a network by a neutral node retains a correspondence between the equilibrium sets of two games. First, given a network G and a node $m \in V(G) \setminus \{0\}$, we formally define a correspondence between the attack paths in any network G and a reduced network $G \setminus m$. Let $h_m^G : P(G) \to P(G \setminus m)$ be a function such that, for a given attack path $p \in P(G)$,

$$h_m^G(p) = \begin{cases} p, & \text{if } m \notin p, \\ p \setminus \{m\}, & \text{otherwise.} \end{cases} \tag{3.4}$$

Function h_m^G defines a 1 to 1 correspondence between the attack paths in network G and the attack paths in network $G \setminus m$. Moreover, the function h_m^G, defined for the set of the pure strategies of the attacker, extends naturally to a function on mixed strategies $H_m^G : \Delta(P(G)) \to \Delta(P(G \setminus m))$ such that, for every probability distribution $\boldsymbol{\pi} \in \Delta(P(G))$ over the set of paths in G, and every path $p' \in P(G \setminus m)$,

$$H_m^G(\boldsymbol{\pi})(p') = \sum_{\substack{p \in P(G) \\ h_m^G(p) = p'}} \boldsymbol{\pi}(p).$$

For a given node m, the function H_m^G defines a correspondence between the mixed strategies of the attacker in the original network and the network reduced by m. This allows us to prove the following proposition concerning the sets of equilibria of two games when the node m is a neutral node.

Lemma 3.6. *Impose Assumptions 2 and 3. Fix game $\Gamma = (G, \boldsymbol{b}, \boldsymbol{d}, \boldsymbol{c})$. If node $m \in [n]$ is a neutral node then a strategy profile $(\boldsymbol{\pi}^*, \boldsymbol{x})$ is a Nash Equilibrium of the game Γ if and only if the strategy profile $(H_m^G(\boldsymbol{\pi}^*), \boldsymbol{x}_{-m}^*)$ is a NE of $\Gamma \setminus m = (G \setminus m, \boldsymbol{b}_{-m}, \boldsymbol{d}_{-m}, \boldsymbol{c}_{-m})$.*

Proof. Since node m is neutral, Lemma 3.3 implies that $x_m^* = 0$. Therefore, for every node $j \in [n] \setminus \{m\}$, and every attack path $p \in P(G)$, such that j is the terminating node $t(p)$ of p, it holds that

$$\alpha_j(p, \boldsymbol{x}_{-j}) = \alpha_j(h_m^G(p), \boldsymbol{x}_{-j}). \tag{3.5}$$

That is, removing a neutral node does not change the probability of the attack reaching any other node along the path. Hence, every attack path $p \in P(G)$ in game Γ, such that the terminating node $t(p)$ is not m, yields the same payoff to the attacker as the corresponding attack path $h_m^G(p)$ in game $\Gamma \setminus m$. Since m is a neutral node, no deviation to an attack path terminating in m is beneficial to the attacker, and, in consequence, a strategy $\boldsymbol{\pi}^*$ is a best response to strategy profile $\boldsymbol{x}^*$ in game Γ if and only if $H_m^G(\boldsymbol{\pi}^*)$ is a best response to strategy profile $\boldsymbol{x}^*_{-m}$ in game $\Gamma \setminus m$.

Consider any defender $j \in [n] \setminus \{m\}$. By Equation (3.5), the payoff to j under strategy profile $(\boldsymbol{\pi}^*, \boldsymbol{x}^*_{-j})$ in Γ is the same function of j's interception probability x_j as the payoff to j under strategy profile $(H_m^G(\boldsymbol{\pi}^*), \boldsymbol{x}^*_{-\{j,m\}})$ in game $\Gamma \setminus m$. Therefore, defender j's best response remains unchanged under network reduction, preserving the equilibrium condition. □

Next, we extend the reduction procedure to any sequence of neutral nodes. In more detail, we prove that for every game in question in network G, there exists a correspondence between the equilibrium sets of the game on the original network and the corresponding game on the network obtained via a series of reductions of neutral nodes. For a given $j \in \{1, 2, \ldots, k-1\}$, let $\boldsymbol{s^j}$ be the vector of first j elements of $\boldsymbol{s}$, and let $G \setminus \boldsymbol{s}^j$ denote the network G after reducing in by the set of nodes in vector $\boldsymbol{s}^j$ in the order of their appearance in vector $\boldsymbol{s}^j$.

Proposition 3.7. *Impose Assumptions 2 and 3. Fix game $\Gamma = (G, \boldsymbol{b}, \boldsymbol{d}, \boldsymbol{c})$. Let $\boldsymbol{s} = (s_i, s_2, \ldots, s_k)$ be any vector of neutral nodes in game Γ, where every neutral node appears at most once. Let S be the set of nodes appearing in $\boldsymbol{s}$. Function*

$$H_{\boldsymbol{s}}^G = H_{s_k}^{G \setminus \boldsymbol{s}^{k-1}} \circ H_{s_{k-1}}^{G \setminus \boldsymbol{s}^{k-2}} \circ \ldots \circ H_{s_2}^{G \setminus \boldsymbol{s}^1} \circ H_{s_1}^G, \tag{3.6}$$

defines a correspondence between the Nash equilibria of games Γ and $\Gamma' = (G \setminus \boldsymbol{s}, \boldsymbol{b}_{-S}, \boldsymbol{d}_{-S}, \boldsymbol{c}_{-S})$, i.e., a strategy profile $(\boldsymbol{\pi}^*, \boldsymbol{x}^*)$ is a NE of Γ if and only if the strategy profile $(H_{\boldsymbol{s}}^G(\boldsymbol{\pi}^*), \boldsymbol{x}^*_{-S})$ is a NE of Γ'.

Proof. We prove the claim using induction over the number of elements of vector $\boldsymbol{s}$.
Induction base For any single-element vector, the claim follows directly from Lemma 3.6.
Induction step Assume that the claim holds for every vector of length at most $k \geq 1$. Let $\boldsymbol{s} = (s_1, s_2, \ldots, s_{k+1})$ be a sequence of distinct neutral nodes. Define $s' = (s_2, s_3, \ldots, s_{k+1})$, and S' be the set of all elements of vector $\boldsymbol{s'}$. By the inductive hypothesis, the reduction of G by $\boldsymbol{s}'$ preserves equilibrium structure. It remains to reduce the resulting network by s_1, which is neutral in $G \setminus \boldsymbol{s}'$ by Lemma 3.6. Thus, by transitivity, the full reduction by $\boldsymbol{s}$ preserves the correspondence between the Nash equilibrium sets of two games. □

We move on to present a short result, characterizing the necessary and sufficient conditions for an edge to exist in a reduced network.

Lemma 3.8. *Fix network G and a sequence of nodes $\boldsymbol{s} = (s_1, \ldots, s_k)$ where each node $s_i \in [n]$ appears at most once. Let S denote the set of nodes appearing in the sequence $\boldsymbol{s}$. Consider the reduced network $G' = G \setminus \boldsymbol{s}$. For any $\{i, j\} \in V(G')$, $\{i, j\} \in E(G')$ if and only if there exists a path p from i to j in network G, such that the only appearing nodes on path p are elements of vector $\boldsymbol{s}$.*

Proof. Let p denote a path from i to j in network G , such that every internal node on p, i.e., any point on p different from i and j, is in sequence $\boldsymbol{s}$. By the definition of the reduction, $H^G_{\boldsymbol{s}}(p) = \{i, j\}$. Hence, the existence of path p in the network G implies existence of the edge $\{i, j\} \in E(G')$.

Assume that $\{i, j\} \in E(G')$. That means that there exists a path p in network G such that $(H^G_{\boldsymbol{s}})^{(-1)}(\{i, j\}) = p$. Moreover, path p does not contain any node outside of S other than i and j, as any such node would not be reduced. Hence, the existence of the edge $\{i, j\} \in E(G')$ implies the existence of a path p in network G, such that the only appearing nodes on path p are elements of vector $\boldsymbol{s}$. □

Lemma 3.8 has a crucial consequence for our further analysis. Since the condition characterized in Lemma 3.8 does not depend on the ordering of the nodes in the sequence $\boldsymbol{s}$, the particular ordering of the sequence $\boldsymbol{s}$ does not influence the resulting network. This is stated formally by the following corollary.

Corollary 3.9. *The outcome of the network reduction is invariant under reordering of the sequence* $\boldsymbol{s}$.

3.5.3 Ascending Networks

Using network reduction, we define an iterative procedure that allows us to obtain a network without linkers, while retaining a correspondence between the Nash equilibria sets of the original and the reduced game. First, we define a procedure that iteratively removes linkers from a given graph. The procedure is given by Algorithm 2.

Algorithm 2 Computing a corresponding network without linkers.

Input: network G
Output: corresponding network G' with no linkers

for $i \in \{1, 2, \ldots, n\}$ **do**
  if i is a linker in G **then**
    $G = G \setminus i$
  end if
end for
return G

We proceed to describe the properties of any network obtained as a result of Algorithm 2. First, we prove that in the resulting network, there are no linkers, as stated by the following lemma.

Lemma 3.10. *Take any network G and the network G' that results from using the Algorithm 2 on G. There are no linkers in G'.*

Proof. Assume otherwise, i.e., there exists a node $m \in V(G')$ that is a linker. When the iteration of the main loop went over node m, m was not removed, hence, it was not a linker at

that point. Therefore, at that iteration, there was at least one neighbor j of m in G, such that $j < m$. Since node j was of a smaller index, it could not have been removed in any iteration of the algorithm after node m had been considered. In consequence, edge $\{j, m\}$ exists in the resulting network, giving a contradiction. □

Take any network G, such that there are no linker nodes in G. Since there are no linker nodes in G, every node $i \in [n]$ has at least one neighbor of a smaller index (it might be node 0), and in consequence, every node can be reached from the attacker by a path of ascending indices. Because of that, we call every network that doesn't have any linkers an *ascending* network. The following proposition characterizes a property of Nash equilibria of every game of attack and interception on an ascending network.

Proposition 3.11. *Impose Assumptions 2 and 3. Fix game $\Gamma = (G, \boldsymbol{b}, \boldsymbol{d}, \boldsymbol{c})$, where network G is an ascending network. Take any strategy profile $(\boldsymbol{\pi}^*, \boldsymbol{x}^*)$ that is a Nash equilibrium of Γ. The set of nodes targeted with positive probability under the attacker's strategy $\boldsymbol{\pi}^*$ is $\{k, k+1, \ldots, n\}$ for some $k \in [n]$.*

Proof. Let $M = \{m, m+1, \ldots, m+l\} \subseteq [n]$, where $m + l < n$, be a subset of nodes target with positive probability under $\boldsymbol{\pi}^*$, such that m is the lowest-index node targeted with a positive probability under $\boldsymbol{\pi}^*$. Assume node $m + l + 1$ is not targeted under $\boldsymbol{\pi}^*$. If that is the case, from Lemma 3.3 it follows that node $m + l + 1$ maximizes its payoff uniquely with the zero-effort strategy, i.e., $x^*_{m+l+1} = 0$. Let $j \in [n]$ be the lowest-index neighbor of $m+l+1$. Since $m+l+1$ is not a linker, it holds that $j < m+l+1$ and in consequence, $b_j < b_{m+l+1}$. Consider two cases.

1. $j \in M$. Then, for any path $p^j = (0, p^j_1, \ldots, j) \in P(G)$ used in $\boldsymbol{\pi}^*$ with positive probability to target node j, deviation to $p' = (0, p^j_1, \ldots, j, m + l + 1)$ is strictly beneficial for the attacker, as it has the same probability of successful attack, which by Lemma 3.3 is strictly positive, and results in a strictly higher reward.

2. $j < m$. Take any attack path of ascending indices from 0 to j, $p^j = (0, p^j_1, \ldots, j) \in P(G)$. None of the nodes on p^j is attacked with positive probability, and in consequence, all nodes on p^j choose the interception probability equal to 0. Therefore, by deviating to attack path $(0, p^j_1, \ldots, j, m + 1) \in P(G)$, the attacker receives a strictly higher payoff than from attacking node m, as the attack is successful with probability of 1, and yields strictly higher reward.

Thus, in any Nash equilibrium of an ascending network, the attacker only targets a consecutive set of the highest-index defenders. □

3.6 Equilibrium Analysis

In this section, we present our analysis of the properties of Nash equilibria of the game in question. First, in Section 3.6.1, we present properties of every fully targeted Nash equilibrium, that is a Nash equilibrium where every node is targeted, and provide necessary and sufficient

conditions for a game to admit a fully targeted Nash equilibrium. Next, in Section 3.6.2 we present results concerning the general case.

3.6.1 Fully Targeted Equilibrium

In this section, we present results concerning Nash equilibria of the attack and interception game in a network, where every defender is targeted with positive probability in equilibrium. We call every Nash equilibrium where every node $j \in [n]$ is attacked with a strictly positive probability a *fully targeted Nash equilibrium*. The following lemma characterizes the properties of attack paths used in every fully targeted Nash equilibrium.

Lemma 3.12. *Impose Assumptions, 2 and 3. Fix game Γ. Let $(\boldsymbol{\pi}^*, \boldsymbol{x}^*) \in \Delta(P(G)) \times [0,1]^n$ be a fully targeted Nash equilibrium of Γ. Fix node $j \in [n]$. Let $k(j)$ be the lowest index neighbor of j. Node j is targeted through a unique path $p^j \in P(G)$, satisfying*

1. *if $k(j) = 0$, $p^j = (0, j)$.*
2. *if $k(j) > 0$, $p^j = (0, p_1^{k(j)}, \ldots, k(j), j)$, where $p^{k(j)} = (0, p_1^{k(j)}, \ldots, k(j))$ is the equilibrium attack path to node $k(j)$.*

Proof. We proceed to prove the following claims.

Claim 1 Network G is an ascending network.

Assume otherwise. Then, there exists at least one linker node in G. By Proposition 3.4, this node is not attacked in any equilibrium, which contradicts that $(\boldsymbol{\pi}^*, \boldsymbol{x}^*)$ is a fully-targeted Nash equilibrium.

Claim 2 For each node $i \in [n]$, it holds that $x_i^* \in (0, 1)$.

This follows directly from Lemma 3.3.

Claim 3 The attacker payoff, $U^* = U(\boldsymbol{\pi}^*, \boldsymbol{x}^*)$ under the strategy profile $(\boldsymbol{\pi}^*, \boldsymbol{x}^*)$ is strictly positive.

This follows directly from Lemma 3.3.

Claim 4 For every node $j \in [n]$, it holds that $U^* = \beta_j^{best}(\boldsymbol{\pi}^*, \boldsymbol{x}^*) \cdot b_j$, where

$$\beta_j^{best}(\boldsymbol{\pi}^*, \boldsymbol{x}^*) = (1 - x_j^*) \cdot \max_{\substack{p \in P(G) \\ t(p)=j}} \alpha_j(p, \boldsymbol{x}_{-j}^*).$$

This follows from the attacker's indifference over the set of pure strategies used in equilibrium.

Claim 5 If $j < k$, than $\beta_j^{best} > \beta_k^{best}$.

This follows directly from Claim 4.

Induction base As every node $i \in [n]$ is targeted under the attacker's strategy $\boldsymbol{\pi}^*$, from Lemma 3.3 it follows, that $x_i^* > 0$ for every node i. Hence, every node $j \in [n]$ that is a neighbor of 0 is uniquely attacked through a single-edge path $(0, j)$, as going through any other node would strictly decrease the probability of a successful attack. Since network G is an ascending network, node 1 is a neighbor of node 0.

Induction step Assume Lemma 3.12 holds for nodes $\{1, 2, \ldots, j\}$, where $1 \leq j < n$. Consider node $j+1$. In equilibrium, node $j+1$ needs to be reached by one of the attack paths resulting in the highest attack probability. From Claim 5, we know that every such path goes through node $k(j)$, which is the lowest index neighbor of j. By Claim 1, $k(j) < j$. From the induction assumption, we get that this path is unique. □

We call the unique attack path to node j, characterized by Lemma 3.12, an *equilibrium attack path* to node j. From Lemma 3.12 it follows, that the set of equilibrium attack paths $\{p^j\}_{j\in[n]}$ constitutes a tree. We call such a tree an *equilibrium attack tree*, and for a given ascending network G, we denote it $T(G)$. As stated by the following definition, we formally define the equilibrium attack tree for any ascending network.

Definition 3.13 (Equilibrium Attack Tree)**.** For any ascending network G, we call a tree such that every node $j \in [n]$ is connected with an edge to its lowest neighbor $k(j)$, an equilibrium attack tree.

As a consequence of Lemma 3.12, a mixed strategy of the attacker in every fully-targeted Nash equilibrium can be expressed as a vector $\boldsymbol{q}^*$, where q_j^* is the probability of choosing the equilibrium attack path p^j. Hence, from now on, we refer to the attacker's strategy in any fully targeted Nash equilibrium with $\boldsymbol{q}^*$ to stress that it is in fact a distribution on n equilibrium attack paths. The following proposition is an extension of the Equations (4-8) proposed by Bloch et al. [2023], that characterizes the first-order condition on every strategy profile $(\boldsymbol{q}^*, \boldsymbol{x}^*)$ that is a fully targeted Nash equilibrium.

Proposition 3.14. *Impose Assumptions 2 and 3. Fix game Γ. A fully targeted Nash equilibrium $(\boldsymbol{q}^*, \boldsymbol{x}^*) \in \Delta(P(T(G))) \times [0,1]^n$ of Γ satisfies the following conditions*

$$x_j^* = 1 - \frac{U^*}{b_j}, \qquad \text{if } j \in D_0, \tag{3.7}$$

$$x_j^* = 1 - \frac{b_{k(j)}}{b_j}, \qquad \text{if } j \in D \setminus D_0, \tag{3.8}$$

$$q_j^* = \frac{c_j' \left(1 - \frac{U^*}{b_j}\right)}{d_j}, \qquad \text{if } j \in D_0, \tag{3.9}$$

$$q_j^* = \frac{b_{k(j)} \cdot c_j' \left(1 - \frac{b_{k(j)}}{b_j}\right)}{U^* \cdot d_j}, \text{ if } j \in D \setminus D_0, \tag{3.10}$$

$$\sum_j q_j^* = 1, \tag{3.11}$$

where $k(j)$ is the lowest-index neighbor of j, U^* is the attacker's payoff under the strategy profile $(\boldsymbol{q}^*, \boldsymbol{x}^*)$, and $D_0 \subseteq [n]$ is the set of neighbors of node 0.

Proof. Equations (3.7) and (3.8) are obtained from the equations guaranteeing that the attacker is indifferent among the targets in the support, that is:

$$b_j(1 - x_j^*) = U^*, \quad \text{for } j \in D_0,$$
$$b_j(1 - x_j^*) = b_{k(j)}, \text{ for } j \notin D_0.$$

Equations (3.9) and (3.10) follow from the fact that every defender needs to best respond in any Nash equilibrium. Let $\alpha_j(p^j, \boldsymbol{x}^*_{-j})$ be the probability of the attack reaching node j using the equilibrium attack path p^j under the profile of defenders' interception probabilities $\boldsymbol{x}^*_{-j}$. The derivative of defender's j payoff under strategy profile $(\boldsymbol{q}^*, \boldsymbol{x}^*)$ is

$$\frac{\partial V_j(q, x_1, \ldots, x_n)}{\partial x_j} = \alpha_j(p^j, \boldsymbol{x}^*_{-j}) x_j q_j^* d_j - c_j'(x_j). \tag{3.12}$$

The function $V_j(x_j)$ is strictly concave, therefore, it is only increasing in an interval where $\alpha_j x_j q_j^* d_j \geq c_j'(x_j)$. It follows from Lemma 3.3, that 0 is in this interval while 1 is not, hence

$$c_j'(x_j^*) = \alpha_j(p^j, \boldsymbol{x}^*_{-j}) q_j^* d_j. \tag{3.13}$$

In any Nash equilibrium, the attacker is indifferent over the strategies in the support, i.e., for every node $j \in [n]$ it holds that

$$U^* = b_j \alpha_j(p^j, \boldsymbol{x}^*_{-j})(1 - x_j^*).$$

After transforming this equation, we get

$$\alpha_j(p^j, \boldsymbol{x}^*_{-j}) = \frac{U^*}{b_j(1 - x_j^*)},$$

and after substituting $\alpha_j(p^j, \boldsymbol{x}^*_{-j})$ in Equation (3.13),

$$c_j'(x_j^*) = \frac{U^* q_j^* d_j}{b_j(1 - x_j^*)}.$$

In consequence,

$$q_j^* = \frac{c_j'(x_j^*) b_j(1 - x_j^*)}{U^* \cdot d_j}. \tag{3.14}$$

Equations (3.7) and (3.8) can be restated as

$$1 - x_j^* = \frac{U^*}{b_j} \text{ , if } j \in D_0, \tag{3.15}$$

$$1 - x_j^* = \frac{b_{k(j)}}{b_j}, \text{ if } j \in D \setminus D_0, . \tag{3.16}$$

After substituting x_j^* from Equation (3.7) or (3.8) and $(1 - x_j^*)$ from Equation (3.15) or (3.16) in Equation (3.14), we get Equations (3.9) and (3.10) respectively. Equation (3.11) states that the probabilities in any mixed strategy of the attacker sum up to 1. □

We move on to prove that if the game in question admits a fully targeted Nash equilibrium, then it is the unique Nash equilibrium of the game.

Proposition 3.15. *Impose Assumptions 2 and 3. Fix game* Γ. *Let* $(\boldsymbol{q}^*, \boldsymbol{x}^*) \in \Delta(P(G)) \times [0,1]^n$ *be a fully targeted Nash equilibrium of* Γ. *Strategy profile* $(\boldsymbol{q}^*, \boldsymbol{x}^*)$ *is the unique Nash equilibrium of* Γ.

Proof. First, we prove that $(\boldsymbol{q}^*, \boldsymbol{x}^*)$ is a unique fully targeted Nash equilibrium. We do so by proving that the set of first-order conditions presented in Proposition 3.14 has at most one solution. Consider function F of the attacker's payoff U such that

$$F(U) = \sum_{i \in [n]} q_i^*(U),$$

where $q_i^*(U)$ is given by either Equation (3.9) or (3.10). Since every component of the presented sum is strictly decreasing in U, so is the function F. Hence, there exists at most one value U^* of the attacker's payoff, such that $F(U^*) = 1$. Given a unique value of the attacker's equilibrium payoff U^*, Equations (3.7), (3.8), (3.9) and (3.10) admit unique solutions.

We proceed to prove that the unique fully targeted Nash equilibrium $(\boldsymbol{q}^*, \boldsymbol{x}^*)$ is the unique Nash equilibrium of game Γ. Assume otherwise. Fix strategy profile $(\hat{\boldsymbol{\pi}}, \hat{\boldsymbol{x}}) \neq (\boldsymbol{q}^*, \boldsymbol{x}^*)$ that is a Nash equilibrium of Γ. Let $\hat{U}$ denote the payoff to attacker under strategy profile $(\hat{\boldsymbol{\pi}}, \hat{\boldsymbol{x}})$, and $D \subseteq [n]$ denote the set of nodes attacked with strictly positive probability under the same strategy profile.

Claim 1 Network G of game Γ is an ascending network.

This follows directly from the fact, that $(\boldsymbol{q}^*, \boldsymbol{x}^*)$ is a fully targeted equilibrium of Γ.

Claim 2 $D = \{m, m+1, \ldots, n\}$ for some $m > 1$.

Since $(\hat{\boldsymbol{\pi}}, \hat{\boldsymbol{x}}) \neq (\boldsymbol{q}^*, \boldsymbol{x}^*)$ and $(\boldsymbol{q}^*, \boldsymbol{x}^*)$ is the unique fully targeted Nash equilibrium, the set of D is a strict subset of $[n]$. As the network G is ascending, the claim follows from Proposition 3.11.

Claim 3 $\hat{U} \geq b_{m-1} > U^*$.

By Lemma 3.3, $x_1^* > 0$, and, in consequence, $U^* < b_1$. By Claim 2, node $m-1$ is not attacked under the attacker's strategy $\hat{\boldsymbol{\pi}}$. By Claim 1, $m-1$ is reachable from node 0 by an ascending path $p \in P(G)$. By Claim 2, no node on path p is attacked with positive probability, and in consequence $\alpha_{m-1}(p, \hat{\boldsymbol{x}}_{-(m-1)}) = 1$. Since the attacker cannot beneficially deviate to strategy p, it holds that $\hat{U} \geq b_{m-1}$.

Claim 4 For every node $i \in [n]$, $\alpha_j^{best}(\hat{\boldsymbol{x}}_{-\boldsymbol{j}}) \geq \alpha_j(p^j, \boldsymbol{x}^*_{-j})$, where

$$\alpha_j^{best}(\hat{\boldsymbol{x}}_{-\boldsymbol{j}}) = \max_{\substack{p \in P(G) \\ t(p)=j}} \alpha_j(p, \hat{\boldsymbol{x}}_{-j}),$$

and p^j is the equilibrium attack path to node j.

Take any node $j \in [n]$ and its lowest index neighbor $k(j)$. If $k(j) < m$, $\alpha_j^{best}(\hat{\boldsymbol{x}}_{-j}) = 1$, as there exists an ascending path from 0 to node j, on which no node is attacked under attacker's strategy $\hat{\boldsymbol{\pi}}$.

Assume $k(j) \geq m$. Fix any path $p = (0, p_1, \ldots, k(j)) \in P(G)$ that is in the support of $\hat{\boldsymbol{\pi}}$, such that the terminating node $t(p)$ is $k(j)$. It holds that

$$\alpha_j^{best}(\hat{\boldsymbol{x}}_{-j}) \geq (1 - \hat{x}_{k(j)})\alpha_{k(j)}(p, \hat{\boldsymbol{x}}_{-k(j)}) = \beta_{k(j)}(p, \hat{\boldsymbol{x}}).$$

Moreover, since path p is in the support of $\hat{\boldsymbol{\pi}}$,

$$\beta_{k(j)}(p, \hat{\boldsymbol{x}}) \cdot b_{k(j)} = \hat{U}.$$

From Proposition 3.12, the equilibrium attack path to node j is the equilibrium attack path to node $k(j)$ extended by the edge $(k(j), j)$. Therefore,

$$\beta_{k(j)}(p^{k(j)}, \boldsymbol{x}^*) \cdot b_{k(j)} = \alpha_j(p^j, \boldsymbol{x}^*_{-j}) \cdot b_{k(j)} = U^*.$$

From Claim 3, we get

$$\alpha_j(p^j, \boldsymbol{x}^*_{-j}) < \beta_{k(j)}(p, \hat{\boldsymbol{x}}),$$

which concludes the argument for Claim 4.

Claim 5 For every node $j \in [n]$ it holds that $x_j^* \geq \hat{x}_j$.

Fix node j.

If $j < m$, $\hat{x}_j = 0$, as node j is not attacked under the strategy profile $\hat{\boldsymbol{\pi}}$ and it is maximizing its payoff by choosing the interception probability of 0. Thus, the claim holds.

Assume $j \geq m$. Fix the lowest index neighbor $k(j)$ of j. If $k(j) = 0$, then the equilibrium attack path p^j to j is $(0, j)$, and it holds, that

$$\begin{aligned} U^* &= (1 - x_j^*) \cdot b_j, \text{ and,} \\ \hat{U} &= (1 - \hat{x}_j) \cdot b_j. \end{aligned}$$

From Claim 3 follows that

$$(1 - x_j^*) \cdot b_j < (1 - \hat{x}_j) \cdot b_j,$$

and in consequence

$$x_j^* > \hat{x}_j.$$

Assume $0 < k(j) < m$. Then, $\alpha_j^{best}(\hat{\boldsymbol{x}}_{-j}) = 1$. From Equation (3.8), it holds that

$$b_{k(j)} = (1 - x_j^*) \cdot b_j.$$

Moreover, since $\alpha_j^{best}(\hat{\boldsymbol{x}}_{-j}) = 1$, and node $k(j)$ is not attacked under the strategy profile $\hat{\boldsymbol{\pi}}$, it holds that

$$b_{k(j)} \leq \hat{U} = (1 - \hat{x}_j) \cdot b_j,$$

meaning, that

$$(1 - x_j^*) \cdot b_j \leq (1 - \hat{x}_j) \cdot b_j,$$

and in consequence

$$x_j^* \geq \hat{x}_j.$$

Assume $k(j) \geq m$. Again, from Equation (3.8), it holds that

$$b_{k(j)} = (1 - x_j^*) \cdot b_j. \tag{3.17}$$

Since $k(j)$ is the lowest index neighbor of j, and $k(j) \geq m$, every neighbor of j is attacked with positive probability under the strategy $\hat{\boldsymbol{\pi}}$. From the indifference of the payoffs in a Nash equilibrium, for every neighbor i of j in network G, it holds that

$$\hat{U} = \alpha_i^{best}(\hat{\boldsymbol{x}_{-i}})(1 - \hat{x}_i) \cdot b_i.$$

As a consequence, the higher the node index (value), the lower the probability of reaching that node successfully. Therefore, any path maximizing $\alpha_j^{best}(\hat{\boldsymbol{x}_{-j}})$ has to end with an edge $(k(j), j)$. Hence, since nodes j and $k(j)$ are targeted under $\hat{\boldsymbol{\pi}}$, it holds that

$$\hat{U} = \alpha_{k(j)}^{best}(\hat{\boldsymbol{x}}_{-k(j)})(1 - \hat{x}_{k(j)})b_{k(j)} = \alpha_{k(j)}^{best}(\hat{\boldsymbol{x}}_{-k(j)})(1 - \hat{x}_{k(j)}) \cdot (1 - \hat{x}_j)b_j,$$

and in consequence

$$b_{k(j)} = (1 - \hat{x}_j)b_j,$$

which together with Equation (3.17) gives

$$x_j^* = \hat{x}_j.$$

concluding the argument for Claim 5.

Claim 6 For every node $i \in [n]$, it holds that

$$\hat{q}_j(\hat{\boldsymbol{\pi}}) = \sum_{\substack{p \in P(G) \\ t(p)=j}} \boldsymbol{\pi}(p) \leq q_j^*. \tag{3.18}$$

Fix a node j. If $j < m$, the claim holds. Assume $j \geq m$. Consider the defender j's payoff function under opponents' strategy profile $(\hat{\boldsymbol{\pi}}, \hat{\boldsymbol{x}}_{-j})$:

$$V_j(\hat{\boldsymbol{\pi}}, (\hat{\boldsymbol{x}}_{-j}, x_j)) = \alpha_j^{best}(\hat{\boldsymbol{x}}_{-j})(1 - x_j)(-d_j) - c_j(x_j).$$

The derivative with respect to x_j is:

$$\frac{\partial V_j(\hat{\boldsymbol{\pi}}, \hat{\boldsymbol{x}}_{-j}, x_j)}{\partial x_j} = \alpha_j^{best}(\hat{\boldsymbol{x}}_{-j}) \cdot d_j - c_j'(x_j).$$

Since $c_j(x_j)$ is strictly convex and increasing (by Assumption 3), the payoff function V_j is strictly concave in x_j. By Lemma 3.3, the unique maximizer $\hat{x}_j$ is inside the interval $[0, 1]$. Therefore,

$$\alpha_j^{best}(\hat{\boldsymbol{x}}_{-j}) \cdot d_j - c_j'(\hat{x}_j) = 0,$$

and, in consequence,

$$\hat{x}_j = (c_j')^{(-1)}(\hat{q}_j \cdot \alpha_j^{best}(\hat{\boldsymbol{x}}_{-j}) \cdot d_j).$$

Since c_j is a strictly convex function, its derivative, c_j', is strictly increasing. Therefore, $(c_j')^{(-1)}$ is also strictly increasing, and as a consequence of Claim 5, it holds that

$$q_j^* \cdot \alpha_j(p^j, \boldsymbol{x}_{-j}^*) \cdot d_j \geq \hat{q}_j \cdot \alpha_j^{best}(\hat{\boldsymbol{x}}_{-j}) \cdot d_j.$$

Claim 4 concludes the argument for Claim 6.
Since the Inequality (3.18) is strict for at least one node, we get a contradiction. □

We move on to characterize a necessary and sufficient condition for the existence of a fully targeted Nash equilibrium. Intuitively, $F(U)$, used in the following proposition, represents the total probability mass that would be assigned to all targets under attacker payoff U. Thus, a fully targeted equilibrium exists if this total can reach 1 before the attacker's payoff exceeds the lowest target value.

Proposition 3.16. *Impose Assumptions 2 and 3. Fix game Γ with an ascending network G. Consider function F of the attacker's payoff U such that*

$$F(U) = \sum_{i \in [n]} q_i^*(U),$$

where q_i^* is given by either Equation (3.9) or (3.10). Game Γ admits a fully targeted Nash equilibrium if and only if $F(b_1) < 1$.

Proof. Note that F is a strictly decreasing function of U. If node 1 is targeted, from Lemma 3.3 it follows, that the equilibrium payoff to the attacker, U^*, is strictly smaller than b_1, as $x_1^* > 0$. On the other hand, if U^* is the attacker's equilibrium payoff, it has to hold that $F(U^*) = 1$. Hence, there exists a U^* such that $F(U^*) = 1$, and $U^* < b_1$ if and only if $F(b_1) < 1$. □

We move on to characterize a condition on the attacker's equilibrium payoff in the fully targeted equilibrium.

Proposition 3.17. *Impose Assumptions 2 and 3. Fix game Γ in a network G. Let $(\boldsymbol{q}^*, \boldsymbol{x}^*)$ be a fully targeted Nash equilibrium of Γ. Let $D_0 \subseteq [n]$ be the set of neighbors of the attacker node, 0. The equilibrium payoff U^* to the attacker satisfies*

$$\sum_{i \in D_0} \frac{c_i'(1 - \frac{U^*}{b_i})}{d_i} + \sum_{i \in D \setminus D_0} \frac{b_{k(i)} \cdot c_i'(1 - \frac{b_{k(i)}}{b_i})}{U^* \cdot d_i} = 1. \tag{3.19}$$

Proof. This follows directly from Proposition 3.14. □

3.6.2 General Case

Finally, we proceed to present the results concerning the general case. First, building on already obtained results, we prove that for any game Γ, the set of nodes targeted under a Nash equilibrium of the game is independent of the particular Nash equilibrium. This is formalized by the following proposition.

Proposition 3.18. *Impose Assumptions 2 and 3. Fix game Γ played in network G. The set of targeted nodes is invariant over the set of Nash equilibria of Γ.*

Proof. First, we prove the claim for ascending networks. Assume G is an ascending network. Let $j \in [n]$ denote the lowest-index node attacked with strictly positive probability under any Nash equilibrium of Γ.
Induction base: If $j = 1$, from Proposition 3.11 it follows, that there exists a fully targeted Nash equilibrium of Γ. By Proposition 3.15, this Nash equilibrium is unique.
Induction step: If $j > 1$, 1 is a neutral node. From Lemma 3.6 it follows, that we can reduce the network G by node 1, while maintaining the correspondence given by Equation (3.4) between the Nash equilibrium sets of the original game and the game $\Gamma \setminus 1$ in network $G \setminus 1$.

We proceed to prove the claim for the general case. Let G be any network. By Proposition 3.7, we can compute the corresponding ascending network G' using Algorithm 2, and consider game Γ' on the network G'. □

Following Proposition 3.18, for a given game Γ, we denote the set of nodes attacked with positive probability under any Nash equilibrium of Γ by $D(\Gamma) \equiv D$. We move on to describe the high-level idea of a procedure for computing a Nash equilibrium of the game in question, given by Algorithm 3.

Algorithm 3 Computing a Nash equilibrium.

Input: Game $\Gamma = (G, \boldsymbol{b}, \boldsymbol{d}, \boldsymbol{c})$.
Output: Nash equilibrium of Γ.

Compute ascending Network G' using Algorithm 2.
Take the lowest-index defender $i \in V(G')$.
while i is not targeted **do**
    Reduce G' by i.
    Set i to the lowest-index defender in network G'.
end while
Compute the equilibrium attack tree of T' of G'.
Compute the equilibrium payoff U^* to the attacker solving Equation (3.19).
Fix game $\Gamma' = (G', \boldsymbol{b}_{V(G')}, \boldsymbol{d}_{V(G')}, \boldsymbol{c}_{V(G')}$.
Compute the solution $(\boldsymbol{q}^*, \boldsymbol{x}^*_{V(G')})$ of Γ' in network G' (Proposition 3.14).
for each edge $e \in E(T')$ **do**
    Retrieve a path p_e in the network G, such that $H^G_D(p_e) = e$.
end for
for each equilibrium attack path $p^j \in P(T')$ **do**
    Compute the corresponding attack path $p'^j \in P(G)$, such that for each edge $e \in p^j$, replace e with p_e.
end for
Strategy profile $(\boldsymbol{q}^*, (\boldsymbol{x}^*_{V(G')}, \boldsymbol{0}_{-V(G')}))$ under mapping $p^j \leftrightarrow p'^j$ is a Nash equilibrium of Γ.

The following proposition states the correctness of Algorithm 3.

Proposition 3.19. *Impose Assumptions 2 and 3. Fix game* Γ. *Algorithm 3 computes a Nash equilibrium of the game* Γ.

Proof. The following arguments prove the statement of the proposition..
Line 1 From Proposition 3.10 it follows, that the resulting network G' is ascending. From Lemma 3.7 it follows that the correspondence between the original game Γ, and game Γ', that is, the game Γ restricted to the network G', remains.
Lines 2-6 By Proposition 3.16, we know the necessary and sufficient condition for a game to admit a fully targeted Nash equilibrium. At every iteration of the while loop, the network remains ascending, as every node either keeps its lowest index node, or gets a direct edge connecting him to the attacker, if his lowest-index neighbor was the removed node i.. Therefore, from Proposition 3.11 it follows that if the lowest-index node is not targeted, it is in fact a neutral node, and, in consequence, the network can be reduced by the said node.
Lines 7-10 By Theorem 1 of Bloch et al. [2023], the while loop is guaranteed to terminate with at least one node other than the attacker remaining in the network G'. By Proposition 3.16, after the while loop terminates, the game in question, restricted to the network G' admits a fully targeted Nash equilibrium. Lemma 3.12 gives the unique characterization of the edges in the equilibrium attack tree, while Proposition 3.17 gives the equation that needs to be solved in order to get the attacker's payoff. Proposition 3.14 defines the unique Nash equilibrium strategy profile as a function of the attacker's equilibrium payoff U^*.
Lines 11-13 By Lemma 3.8, for every edge $e = \{i, j\} \in E(G')$, there exists a path p_e in the original network G, such that p does not contain any targeted node other than i and j. Hence, for each edge $e \in E(G')$ we can retrieve such a path.
Line 17 By Lemma 3.3, every neutral node $i \in [n] \setminus D$ exerts zero-effort in every Nash equilibrium, i.e., $x_i^* = 0$.

□

3.7 Equilibrium Computation

In this section, we analyze the computational complexity of Algorithm 3, which computes a Nash equilibrium for a given attack-and-defense game in networks. We present assumptions, data structures, and a breakdown of the worst-case (pessimistic) runtime of each component of the algorithm. We work on the following data structures. The network is represented by an adjacency matrix $M[][]$, which is a two-dimensional array of size $(n+1) \times (n+1)$, such that

$$M[i][j] = \begin{cases} 1, \text{ if } \{i, j\} \in E(G), \\ 0, \text{ otherwise.} \end{cases}$$

The values $\boldsymbol{b}$ and $\boldsymbol{d}$ are kept in two double arrays of size n, called $b[]$ and $d[]$ respectively. Last but not least, for the sake of presentation, we operate in a lazy deletion scheme, that is, we keep a boolean array $is_active[]$ of size n. If a node is removed from a network, we change its status to false. Hence, any node can be removed from the network in constant time $\mathcal{O}(1)$. We work under the following assumptions.

Assumption 4. For every defender $i \in [n]$ and every interception probability $x_i \in [0, 1]$, the cost of choosing an interception probability x_i, $c_i(x_i)$ and it's derivative $der_i(x_i)$ can be computed in constant time $\mathcal{O}(1)$.

Assumption 5. Addition, subtraction, multiplication, and division take constant time $\mathcal{O}(1)$.

To analyze the runtime of Algorithm 3, we begin by establishing the cost of fundamental operations, starting with the removal of nodes from the network data structure.

Lemma 3.20. *Impose Assumptions 4 and 5. The reduction of any network G by any node $i \in [n]$ can be computed in time $\mathcal{O}(n^2)$.*

Proof. Note that 'removal' in the assumed data structure (lazy deletion) refers to marking a node as inactive; the actual network structure remains unchanged, and inactive nodes are ignored in subsequent steps. Hence, removing node i from the network takes constant time $\mathcal{O}(1)$. Getting a list of all neighbors of a node takes linear time $\mathcal{O}(n)$ — it can be achieved by traversing the i'th row of the adjacency matrix, where for each neighbor we check whether he is active before adding him to the resulting list. Adding an edge takes $\mathcal{O}(1)$, since we only need to overwrite two cells of the adjacency matrix. However, since we have to add an edge for every pair of node i neighbors, the final pessimistic runtime is $\mathcal{O}(n^2)$. □

We proceed to identify the total cost of identifying and removing all neutral nodes in the network. This process is central to the ascending network reduction phase of the algorithm.

Lemma 3.21. *Impose Assumptions 4 and 5. Fix ascending network G. The equilibrium attack tree $T(G)$ can be computed in time $\mathcal{O}(n^2)$.*

Proof. For every active node i, we need to traverse the corresponding row of the adjacency matrix in order to find the lowest-index active neighbor of i. Hence, doing so for every node takes at most $\mathcal{O}(n^2)$ constant-time operations. □

After the reduction step, the algorithm reconstructs the unique attack paths corresponding to each targeted node. The following result characterizes a bound on the time required for a path recovery step.

Lemma 3.22. *Impose Assumptions 4 and 5. Fix game Γ in a network G. Let D denote the set of targeted nodes in Γ, and $\boldsymbol{s}$ denote any ordered sequence of the nodes in $[n] \setminus D$. Let G' denote the network resulting from reducing the network G by the sequence $\boldsymbol{s}$. Fix an edge $e = \{i, j\} \in E(G')$. A corresponding path p_e in network G, such that $H^G_{\boldsymbol{s}}(p_e) = e$, can be computed in time $\mathcal{O}(n^2)$.*

Proof. Since p_e must only pass through neutral nodes, we deactivate all other targeted nodes temporarily before searching. Formally, retrieving a path p_e can be done by setting all nodes in $D \setminus \{i, j\}$ to temporarily inactive, and then running the BFS algorithm that ignores the nodes that are not active in order to find a path between i and j in G. The pessimistic runtime of the BFS algorithm operating on the adjacency matrix is $\mathcal{O}(n^2)$. □

Combining these three components — node removal, reduction of the network, and path recovery — we now establish the overall pessimistic runtime of Algorithm 3.

Proposition 3.23. *Impose Assumptions 2, 3, 4 and 5. Fix game* Γ. *A Nash equilibrium of* Γ *can be computed using Algorithm 3 in time*

$$\mathcal{O}(n^3 + \Phi),$$

where Φ *denotes the time required to solve Equation 3.19 for the attacker's payoff* U^*.

Proof. We present the arguments for each part of Algorithm 3.
Line 1 Algorithm 2 iterates exactly once over every node $i \in [n]$. Lemma 3.20 states, that each iteration takes at most $\mathcal{O}(n^2)$, hence, the pessimistic runtime is $\mathcal{O}(n^3)$.
Lines 2-6 The while loop iterates over each node at most once. To check whether a node is targeted, we need to compute the equilibrium attack tree, which takes at most $\mathcal{O}(n^2)$ time (Lemma 3.21), and check whether the condition presented in Proposition 3.16 holds, which takes at most $\mathcal{O}(n)$ time from Assumptions 4 and 5. Since in every loop at most one reduction is performed, the pessimistic runtime is in fact $\mathcal{O}(n^3)$.
Line 7 As stated by Lemma 3.21, computing the equilibrium attack tree takes at most $\mathcal{O}(n^2)$ time.
Line 8 By the proposition statement, solving Equation (3.19) requires time Φ.
Line 9 For each targeted node i, two corresponding values q_i^* and x_i^* are computed. From Assumptions 4 and 5, computing single value takes $\mathcal{O}(1)$, and in consequence, the pessimistic runtime is $\mathcal{O}(n)$.
Lines 10-12 The for loop iterates over the set of at most n edges spanning the equilibrium attack tree. Since for a single edge the pessimistic runtime is $\mathcal{O}(n^2)$, the pessimistic runtime of the entire loop is $\mathcal{O}(n^3)$.
Lines 12-15 The for loop iterates over the set of at most n paths equilibrium attack paths. Each equilibrium attack path has at most n edges. For each edge e, the corresponding path p_e has at most $n + 1$ vertices. Hence, the pessimistic runtime is not greater than $\mathcal{O}(n^3)$ □

In the general case, the overall runtime of Algorithm 3 depends on the computational effort required to solve Equation (3.19)), which characterizes the attacker's equilibrium payoff. However, when all defender cost functions are quadratic, this equation reduces to a closed-form quadratic equation in U^*. Although computing the full equilibrium still involves network reduction and path reconstruction, this refinement yields a total pessimistic runtime of $\mathcal{O}(n^3)$.

Proposition 3.24. *Impose Assumptions 2, 3, 4, and 5. If for every defender* $i \in [n]$, *the cost function* c_i *is a quadratic function, i.e.,* $c_i(x_i) = a_i \cdot x_i^2$, *Algorithm 3 runs in pessimistic time of* $\mathcal{O}(n^3)$.

Proof. Assume $c_i(x) = a_i x^2$, so that $c_i'(x) = 2a_i x$. Substituting into Equation (3.19), we obtain:

$$\sum_{i \in D_0} \left(\frac{2a_i}{d_i} - \frac{2a_i U^*}{b_i d_i} \right) + \sum_{i \in D \setminus D_0} \frac{2a_i b_{k(i)} \left(1 - \frac{b_{k(i)}}{b_i}\right)}{U^* d_i} = 1, \tag{3.20}$$

where D_0 is the set of nodes directly connected to the attacker and $k(i)$ is the predecessor of node i in the equilibrium attack tree. This expression simplifies to:

$$A - BU^* + \frac{C}{U^*} = 1, \tag{3.21}$$

where the constants are defined as:

$$A = \sum_{i \in D_0} \frac{2a_i}{d_i}, \quad B = \sum_{i \in D_0} \frac{2a_i}{b_i d_i}, \quad C = \sum_{i \in D \setminus D_0} \frac{2a_i b_{k(i)} \left(1 - \frac{b_{k(i)}}{b_i}\right)}{d_i}. \tag{3.22}$$

Multiplying both sides by U^* and rearranging terms, we obtain the quadratic equation:

$$B(U^*)^2 - (A-1)U^* - C = 0, \tag{3.23}$$

which can be solved in constant time using the quadratic formula.

Once U^* is known, the remaining steps of Algorithm 2—namely, network reduction and equilibrium path reconstruction—run in pessimistic time $\mathcal{O}(n^3)$ (as shown in Proposition 3.23). Therefore, the overall runtime of Algorithm 3 under quadratic cost functions is $\mathcal{O}(n^3)$.

□

3.8 Conclusions

In this chapter, we analyzed attack and interception games on networks, where an attacker selects a path through the network and defenders allocate interception resources across their respective nodes. Although the attacker's strategy space—the set of all simple paths—can grow exponentially with the number of nodes, we showed that Nash equilibria can be computed in polynomial time for a broad class of cost functions.

Our approach exploits structural properties of networks to reduce the computational burden. In particular, we demonstrated that the set of attack paths used in equilibrium is unique, of polynomial size, and can be determined efficiently. These results establish that the exponential expansion of the attacker's strategy space does not necessarily translate into intractability, provided that the network structure is leveraged appropriately.

This chapter thus highlights how combinatorial properties of networked environments can render equilibrium computation tractable in otherwise succinctly represented games. By combining structural insight with algorithmic design, we extend the scope of efficiently solvable models within algorithmic game theory.

Chapter 4

Equilibria of the Colonel Blotto Game with Flexible Tie-breaking

Chapter Overview

This chapter presents our analysis concerning the equilibria of the Colonel Blotto game with discrete strategies and a flexible tie-breaking rule. We allow tied outcomes on individual battlefields to yield an arbitrary share of the battlefield's value, which generalizes the classic constant-sum formulation. This modification captures a variety of real-world strategic environments where ties may be inefficient (e.g., military stalemates) or efficient (e.g., simultaneous scientific breakthroughs). Our focus is on understanding how such flexibility in the tie-breaking rule affects the existence, structure, and selection of equilibrium strategies.

We begin by showing that canonical equilibrium constructions from the standard model—with players randomizing uniformly over symmetric allocations—extend to the generalized setting. In particular, equilibrium strategies with uniform unidimensional marginal distributions remain valid regardless of how ties are valued. We prove that the likelihood of ties remains constant across bid levels used in equilibrium, which preserves the indifference conditions necessary for equilibrium play.

Next, we demonstrate that the equilibrium set is excessively large: any pure strategy that assigns no more than twice the fair share to each battlefield appears in the support of some equilibrium strategy. We further show that these strategies cannot be eliminated by standard refinement tools such as the iterated removal of weakly dominated strategies—especially when the ties are inefficient. This highlights a limitation of equilibrium refinements in discrete Blotto games and suggests that payoff structure alone cannot guide equilibrium selection.

Finally, we discuss robustness: while uniform marginal strategies persist as equilibria under small perturbations in tie efficiency, other equilibrium types (e.g., asymmetric or parity-based) may break down or emerge depending on how valuable ties are. The results provide insight into equilibrium multiplicity and selection in a widely studied class of games, and they point to the importance of understanding payoff externalities even in formally symmetric environments.

The remainder of the chapter is structured as follows. In Section 4.1, we provide real-life

motivations of our study. The model is introduced in Section 4.2. In Section 4.3, we present an approach for constructing Nash equilibria within the considered class of games. Section 4.4 concerns the equilibrium set. Section 4.5 presents the extensions of our main results, that are further discussed in Section 4.6. We conclude in Section 4.7.

4.1 Motivating Examples

In contemporary military conflicts, opposing forces must make strategic decisions about how to distribute limited units—such as troops, equipment, and air support—across multiple geographically distinct frontlines. Each region holds similar strategic importance, whether due to geography, political symbolism, or access to key infrastructure. However, a frontline contested by evenly matched forces often devolves into a prolonged stalemate. Such deadlocks typically prolong conflict, drain resources, and lead to political and humanitarian consequences, all without yielding decisive outcomes. Military planners must therefore consider not only how to distribute limited forces across simultaneous engagements, but also how the possibility of mutual obstruction affects the overall efficiency of their strategic deployment. The key questions become: How can forces be allocated to avoid protracted standoffs? What are the opportunity costs of spreading forces too thinly or concentrating them too narrowly? In this setting, evenly matched contests may be inherently inefficient and represent a loss of strategic momentum.

A similar dynamic appears in the realm of intellectual property litigation, where firms often find themselves simultaneously involved in multiple patent disputes of comparable importance. Legal resources—such as attorneys, expert witnesses, and administrative support—must be strategically allocated across these cases to achieve favorable outcomes. However, when competing firms commit similar resources to the same claim, the likely result is an unresolved legal conflict that is both lengthy and expensive. These unresolved disputes not only tie up legal budgets but also delay product launches and create uncertainty in the marketplace. In this context, ties do not offer any constructive resolution but instead amplify inefficiency. Legal strategists must therefore assess the risk of unproductive deadlock and consider how to allocate effort to either dominate certain cases or strategically concede others in order to maximize overall advantage.

By contrast, in the competitive world of scientific research, simultaneous success by rival teams can often lead to more efficient outcomes. Institutions racing to solve a set of high-impact scientific problems—such as developing new vaccines or advancing climate technologies—must allocate their limited research resources across equally valuable challenges. When two teams independently reach a breakthrough around the same time, the outcome may benefit the broader scientific community: validation of results becomes more credible, media attention increases, and policy or funding support is amplified. In these cases, ties can generate spillover benefits that exceed the value of exclusive discovery. Research strategists must therefore consider how to allocate resources not only to maximize the probability of success, but also to benefit from the reputational and collaborative advantages that may arise when progress occurs in parallel. Unlike adversarial settings, ties in this context can enhance efficiency rather than

diminish it.

The strategic settings described above share a common structure: two players compete across multiple equally valuable contests, each allocating limited and indivisible resources in pursuit of a global objective. However, while the formal setup may be identical, the consequences of ties differ substantially across environments. In military or legal contexts, ties often result in costly inefficiencies, such as prolonged conflict or unresolved litigation, leading strategic agents to avoid evenly matched allocations. In contrast, scientific competition presents a case where ties may lead to highly efficient outcomes, with simultaneous success generating reputational gains and collaborative spillovers. These observations suggest that the efficiency of ties—whether they reduce or enhance total utility—plays a central role in shaping optimal strategies. In this chapter, we formalize this idea by studying a variant of the Colonel Blotto game Borel [1921] that accommodates flexible tie-breaking rules. The objective is to understand how equilibrium behavior changes when the strategic value of a tie is allowed to vary, and to characterize the resulting equilibrium set. In doing so, the analysis highlights both the robustness and the limits of known equilibrium constructions under different assumptions about the efficiency of ties.

4.2 Model

Two players, denoted by A and B, each allocate a total of $D \geq 1$ resources over $n \geq 2$ battlefields.[1] Units of the resource are not divisible. Hence, a *pure strategy* of player $P \in \{\mathsf{A}, \mathsf{B}\}$ is a vector

$$s^P = \begin{pmatrix} s_1^P \\ \vdots \\ s_i^P \end{pmatrix},$$

such that $s_i^P \in \{0, 1, \ldots, N\}$ for every $i \in \{1, \ldots, n\}$, and

$$\sum_{i=1}^{n} s_i^P = D.$$

The set of strategies for players A and B is identical and denoted by S. As usual, we refer to the opponent of player P by $-P$. For a given pure strategy profile $(s^P, s^{-P}) \in S \times S$, the payoff of player $P \in \{\mathsf{A}, \mathsf{B}\}$ in the Colonel Blotto game is defined by

$$\pi^P\left(\boldsymbol{s}^P, \boldsymbol{s}^{-P}\right) = \sum_{i=1}^{n} \left(\left[s_i^P > s_i^{-P}\right] + \frac{\alpha}{2} \cdot \left[s_i^P = s_i^{-P}\right] \right),$$

where, given a condition φ, $[\varphi]$ is the Iverson's bracket taking value 1 if φ is satisfied and value 0 otherwise, and α is a parameter. The departure from the standard model is the introduction

[1] In the excluded cases where $D = 0$ or $n = 1$, players have only one strategy. However, note that we allow for $n = 2$, which is not necessarily trivial in our model.

of flexible tie-breaking, represented by α. We call the two-player game with strategy sets and payoffs defined as above the *Colonel Blotto game* $\mathcal{B}_\alpha \equiv \mathcal{B}_\alpha(D, n)$.[2]

4.2.1 Examples

Below, we recall two examples of finite Colonel Blotto games that have been considered in the literature.

Example 4.1 (Hart, 2008). In the standard version of the Colonel Blotto game, $\mathcal{B}_1(D, n)$, player P's payoff is defined as the number of battlefields won plus half of the tied battlefields. Although this is formally a different definition then the one used by Hart [2008], where the payoff to the player is the difference between the number of won and lost battlefields, the two definitions result in strategically equivalent games.

$$\pi^P\left(\boldsymbol{s}^P, \boldsymbol{s}^{-P}\right) = \sum_{i=1}^{n}\left(\left[s_i^P > s_i^{-P}\right] + \frac{1}{2}\cdot\left[s_i^P = s_i^{-P}\right]\right).$$

The constant-sum setup underlying Example 4.1 has indeed been prevalent in the literature. Borel [1921, p. 100] considered a particular case where $D = 7$ and $n = 3$, a solution of which has been offered by Hart [2008]. The setup in Example 4.1 has been tested as a symmetric control by Avrahami et al. [2014] for $D \in \{16, 24\}$ and $n = 8$.

Example 4.2 (Arad and Rubinstein, 2012). Two colonels are asked to distribute a total of $D = 120$ units of the resource over a total of $n = 6$ battlefields. The payoff of a player is the number of battlefields that she assigned strictly more resources to than her opponent. Thus, the game is $\mathcal{B}_0(120, 6)$, the common set of strategies is

$$S = \{s^P \in \{0, 1, \ldots, 120\}^6 : \sum_{i=1}^{6} s_i^P = 120\},$$

and the payoff of player i is defined by

$$\pi^P\left(\boldsymbol{s}^P, \boldsymbol{s}^{-P}\right) = \sum_{i=1}^{6}\left[s_i^P > s_i^{-P}\right]. \tag{4.1}$$

[2]The number of pure strategies in $\mathcal{B}_\alpha(N, K)$ is $|S| = \binom{N+K-1}{K-1}$. In the corresponding Colonel Lotto game, i.e., taking account of symmetries between battlefields, the number of pure strategies is given by the number of partitions of D into exactly n *weakly* positive parts, or equivalently, by $p(N + K, K)$, where $p(D, n)$ denotes the number of partitions of D into exactly n *strictly* positive parts. The recursive relationship

$$p(D, n) = p(N - 1, K - 1) + p(N - K, K)$$

with initial conditions $p(N, 1) = 1$ and $p(D, n) = 0$ if $n > N$ allows computing this number in specific examples [Gupta, 1970].

Abstracting from the fact that Example 4.1 keeps D and n flexible while Example 4.2 does not, the main difference between the two examples lies in the tie-breaking rule applied. In contrast to Example 4.1, the game in Example 4.2 fails to be constant-sum, as the sum of payoffs of both players depends on probabilities of ties occurring on individual battlefields. This complicates the equilibrium analysis but also makes the game more interesting. Intuitively, setting $\alpha = 0$ provides an additional incentive to outguess the opponent. For instance, in an experiment with $\alpha = 1$, subjects might perceive the pure strategy $(20, 20, 20, 20, 20, 20)'$ as a focal point that leads to a fair division.[3] With $\alpha = 0$, however, payoffs at this focal point are zero for both players, i.e., there is a strong incentive to engage in at least *some* additional reasoning. As this consideration is mentioned by Arad and Rubinstein [2012] early in their work, it might have contributed to their decision to depart from the standard tie-breaking rule.

For expositional reasons, we will initially work under a set of simplifying assumptions. The implications of dropping these assumptions will be discussed in the extensions section. Our first assumption concerns the parity of the number of battlefields.

Assumption 6. The number of battlefields, n, is even.

Assumption 6 simplifies the analysis but can often be dropped at the cost of additional arguments. Our second assumption concerns the relationship between D and n.

Assumption 7. The number of resources, D, is divisible by n.

Thus, in the main part of the analysis, we will assume that the number of resources D is a multiple of n. It follows from Assumption 2 that

$$m = \frac{D}{n}$$

is an integer. For an experimental subject, this means that the uniform allocation that assigns m units of the resource to each of the battlefields is a possibility. The last assumption concerns the efficiency parameter α.

Assumption 8. The efficiency parameter α satisfies $\alpha \in [0, 2]$.

Assumption 8 may be considered natural but nevertheless restricts the efficiency parameter in two ways. The restriction $\alpha \geq 0$ says that ties cannot cause inefficiencies beyond the complete loss of the battlefield value. For $\alpha < 0$, players would have a very strong incentive to avoid ties, which may lead to asymmetric equilibria. Conversely, for $\alpha > 2$, ties would be "superefficient," so that pure-strategy equilibria become natural. Such possibilities will be further discussed in the extensions section.

Note that Assumptions 6 through 8 hold in the Arad-Rubinstein game. Indeed, in Example 4.2, $n = 6$ is even, $m = \frac{120}{6} = 20$ is an integer, and $\alpha = 0 \in [0, 2]$.

Given a finite non-empty set X, let $\Delta(X)$ denote the set of all probability distributions on X. We are interested in *mixed-strategy Nash equilibria* of the game, i.e., mixed strategy

[3] Here and below, we represent the strategy as a row vector via transposition.

profiles $\boldsymbol{\sigma} = (\boldsymbol{\sigma}^P, \boldsymbol{\sigma}^{-P}) \in \Delta(S) \times \Delta(S)$ such that no player can improve her expected payoff by unilaterally changing her mixed strategy. By a *symmetric Nash equilibrium strategy*, we mean any mixed strategy $\boldsymbol{\sigma}^P$ such that $\boldsymbol{\sigma} = (\boldsymbol{\sigma}^P, \boldsymbol{\sigma}^P)$ is a mixed-strategy Nash equilibrium.

A crucial property of Colonel Blotto games is that payoffs are functions of the respective marginal distributions at each battlefield. Fix mixed strategy $\boldsymbol{\sigma}^P$ of player P. Let $\boldsymbol{\sigma}_i^P$ be a marginal distribution induced by strategy $\boldsymbol{\sigma}^P$ on battlefield i, i.e.,

$$\mathbf{Pr}(\boldsymbol{\sigma}_i^P = k) = \sum_{\boldsymbol{x} \in S} \boldsymbol{\sigma}_{\boldsymbol{x}} \cdot [x_i = k],$$

where $\boldsymbol{\sigma}_{\boldsymbol{x}}$ is the probability associated with pure strategy $\boldsymbol{x}$. It holds that

$$\begin{aligned}
\pi^{\mathsf{A}}\left(\boldsymbol{\sigma}^{\mathsf{A}}, \boldsymbol{\sigma}^{\mathsf{B}}\right) &= \sum_{(\boldsymbol{x},\boldsymbol{z}) \in S \times S} \boldsymbol{\sigma}_{\boldsymbol{x}}^{\mathsf{A}} \boldsymbol{\sigma}_{\boldsymbol{z}}^{\mathsf{B}} \pi^P(\boldsymbol{x}, \boldsymbol{z}) \\
&= \sum_{(\boldsymbol{x},\boldsymbol{z}) \in S \times S} \boldsymbol{\sigma}_{\boldsymbol{x}}^{\mathsf{A}} \boldsymbol{\sigma}_{\boldsymbol{z}}^{\mathsf{B}} \sum_{i=1}^{n} \left([x_i > z_i] + \frac{\alpha}{2} [x_i = z_i] \right) \\
&= \sum_{i=1}^{n} \sum_{(\boldsymbol{x},\boldsymbol{z}) \in S \times S} \boldsymbol{\sigma}_{\boldsymbol{x}}^{\mathsf{A}} \boldsymbol{\sigma}_{\boldsymbol{z}}^{\mathsf{B}} \left([x_i > z_i] + \frac{\alpha}{2} [x_i = z_i] \right) \\
&= \sum_{i=1}^{n} \left(\mathbf{Pr}\left(\boldsymbol{\sigma}_i^{\mathsf{A}} > \boldsymbol{\sigma}_i^{\mathsf{B}}\right) + \frac{\alpha}{2} \mathbf{Pr}\left(\boldsymbol{\sigma}_i^{\mathsf{A}} = \boldsymbol{\sigma}_i^{\mathsf{B}}\right) \right).
\end{aligned}$$

Given a mixed strategy $\boldsymbol{\sigma}^P$ and a battlefield $n \in \{1, \ldots, K\}$, let $\boldsymbol{\sigma}_i^P$ denote the marginal distribution of $\boldsymbol{\sigma}^P$ at battlefield i. Following Hart [2008], we denote the uniform distribution on natural numbers from 0 to $2m$ by U^m. A mixed strategy $\boldsymbol{\sigma}^P$ such that $\boldsymbol{\sigma}_i^P = U^m$ for every battlefield i and some $m \geq 1$ will be said to induce *uniform marginals*.

4.3 Equilibrium of the Colonel Blotto Game with Flexible Tie-breaking

In this section, we present a simple approach that allows constructing an equilibrium in the Colonel Blotto game with flexible tie-breaking. As mentioned before, this will lead us to an equilibrium in the Arad-Rubinstein game as well.

4.3.1 A canonical equilibrium

The following result characterizes one particular Nash equilibrium in $\mathcal{B}_\alpha(D, n)$.

Proposition 4.3. *Impose Assumptions 6 through 8. Then, a symmetric equilibrium strategy of* $\mathcal{B}_\alpha(D,n)$ *is given by uniform randomization over the set of pure strategies*

$$S_0 = \left\{ \begin{pmatrix} 0 \\ 2m \\ \vdots \\ 0 \\ 2m \end{pmatrix}, \begin{pmatrix} 1 \\ 2m-1 \\ \vdots \\ 1 \\ 2m-1 \end{pmatrix}, \ldots, \begin{pmatrix} 2m \\ 0 \\ \vdots \\ 2m \\ 0 \end{pmatrix} \right\}.$$

In the resulting equilibrium, players' expected payoffs amount to $\pi^* = K \cdot \frac{m+\frac{\alpha}{2}}{2m+1}$.

Proof. See Section 4.3.4. □

As can be seen, players partition the set of battlefields into pairs. This is possible, of course, because the number of battlefields has been assumed even via Assumption 6. To each pair of battlefields, a constant number of $2m$ resources is allocated, which is feasible due to Assumption 7. Moreover, the split among the two battlefields in each pair is uniformly distributed, with perfect correlation across pairs. We note that players' strategies induce uniform marginals. Indeed, for any battlefield i, every number of resources in $\{0,1,2,\ldots,2m\}$ is assigned to battlefield i with the same probability of $\frac{1}{2m+1}$.

4.3.2 Discussion

Hart [2008] has shown that, under Assumptions 6 and 7, strategies inducing uniform marginals form an equilibrium in the constant-sum game $\mathcal{B}_1(D,n)$. The point to note is that this remains the case even with flexible tie-breaking. The reason why the equilibrium property does not break down with more flexible tie-breaking is that, given uniform marginals, *the likelihood of getting tied is constant across all bid levels that are used in equilibrium with positive probability*. As a result, the indifference in the standard setup with $\alpha = 1$ is not affected by the modification of the payoff functions at ties.

One might wonder if, with the modified tie-breaking in place, players would not have an incentive to bid higher than $2m$ on some of the battlefields. Such an incentive might arise for $\alpha < 1$, because the tie-breaking is inefficient in that case. However, overbidding cannot raise a player's payoff. The reason is that any additional unit of the resource required to assign more than $2m$ on some battlefield needs to be taken from some other battlefield, where this lowers the probability of winning by $\frac{1}{2m+1}$. Indeed, every bid level in $\{0,1,\ldots,2m\}$ is used by the opponent on every battlefield with the same probability of $\frac{1}{2m+1}$. On the other hand, the increase in payoff from bidding $2m+1$ instead of $2m$ in a battlefield is $\frac{1}{2m+1} \times (1-\alpha) \leq \frac{1}{2m+1}$. Therefore, given Assumption 8, or more precisely given that $\alpha \geq 0$, the deviation never yields a strictly higher payoff.[4]

[4] As the discussion shows, the conclusion of Proposition 4.3 remains technically true for $\alpha > 2$. However, as mentioned before, there is no reason to hide one's strategy for $\alpha \geq 2$, i.e., symmetric pure-strategy equilibria may be more plausible in that case. Cf. Section 4.6.

4.3.3 Illustration

We illuminate Proposition 4.3 with an example.

Corollary 4.4 (Equilibrium of the Arad-Rubinstein game). *The following strategy is a symmetric Nash equilibrium strategy of $\mathcal{B}_0(120, 6)$. Both players individually and independently randomize uniformly over the set*

$$S_0 = \left\{ \begin{pmatrix} 0 \\ 40 \\ 0 \\ 40 \\ 0 \\ 40 \end{pmatrix}, \begin{pmatrix} 1 \\ 39 \\ 1 \\ 39 \\ 1 \\ 39 \end{pmatrix}, \ldots, \begin{pmatrix} 40 \\ 0 \\ 40 \\ 0 \\ 40 \\ 0 \end{pmatrix} \right\}.$$

The equilibrium payoff is $\pi^* = \frac{120}{41} \approx 2.927$.

Proof. Immediate from Proposition 4.3. □

Even though the equilibrium characterized by Proposition 4.3 has a canonical structure (e.g., it has uniform marginals, and is symmetric with respect to permutations of the battlefield pairings and within battlefield pairings), it is not the only equilibrium, as we show in Section 4.4.1.

4.3.4 Proof of Proposition 4.3

We start with an auxiliary result.

Lemma 4.5. *Impose Assumptions 7 and 8. Let $\boldsymbol{\sigma} = (\boldsymbol{\sigma}^{\mathsf{A}}, \boldsymbol{\sigma}^{\mathsf{B}})$ be a mixed strategy profile in $\mathcal{B}_\alpha(D, n)$ such that $\boldsymbol{\sigma}^P$ induces uniform marginals for $P \in \{\mathsf{A}, \mathsf{B}\}$. Then, $\boldsymbol{\sigma}$ is a Nash equilibrium.*

Proof. Suppose that both players use the assumed mixed strategies that induce uniform marginals. Then, the probability of a tie on any given battlefield is $\frac{1}{2m+1}$. Therefore, the expected payoff to player P under the strategy profile $\boldsymbol{\sigma} = (\boldsymbol{\sigma}^{\mathsf{A}}, \boldsymbol{\sigma}^{\mathsf{B}})$ is

$$\pi^P(\boldsymbol{\sigma}) = \frac{K}{2} + K \cdot \frac{1}{2m+1} \cdot \frac{\alpha - 1}{2} = \frac{K \cdot (2N + \alpha K)}{4N + 2K}.$$

Suppose that player $P \in \{\mathsf{A}, \mathsf{B}\}$ deviates to a pure strategy $s^P \in S$. The mixed strategy $\boldsymbol{\sigma}^{-P}$ allocates every number of resources in $\{0, 1, 2, \ldots, 2m\}$ on every battlefield i with the same probability of $\frac{1}{2m+1}$. If player P assigns strictly more than $2m$ resources to a battlefield, her payoff from that battlefield is 1. If player i assigns weakly less than $2m$ to battlefield i, then she overwhelms the opponent on that battlefield with probability $\frac{s_i^P}{2m+1}$ and achieves a tie with probability $\frac{1}{2m+1}$. Hence, player P's expected payoff from battlefield i is

$$\pi_i^P\left(s^P, \boldsymbol{\sigma}^{-P}\right) = \min\{1, \frac{s_i^P}{2m+1} + \frac{\alpha}{2} \cdot \frac{1}{2m+1}\}.$$

It follows that

$$\begin{aligned}\pi^P\left(s^P, \boldsymbol{\sigma}^{-P}\right) &= \sum_{i=1}^{n} \pi_i^P\left(s^P, \boldsymbol{\sigma}^{-P}\right) \\ &\leq \sum_{i=1}^{n}\left(\frac{s_i^P}{2m+1} + \frac{\alpha}{2}\cdot\frac{1}{2m+1}\right) \\ &= \frac{D}{2m+1} + \frac{n\cdot\alpha}{2\cdot(2m+1)} \\ &= \frac{n\cdot(2D+\alpha n)}{4D+2n}.\end{aligned} \tag{4.2}$$

As both P and s^P were arbitrary, no player can raise her payoff by deviating from $\boldsymbol{\sigma} = (\boldsymbol{\sigma}^{\mathsf{A}}, \boldsymbol{\sigma}^{\mathsf{B}})$. Hence, $\boldsymbol{\sigma}$ is a Nash equilibrium, which completes the proof of the lemma. □

Proof of Proposition 4.3. The mixed strategy $\boldsymbol{\sigma}^P$, which is well-defined given Assumption 6, assigns every number of resources $s_i^P \in \{0, 1, \ldots, 2m\}$ to every battlefield i with the same probability $\frac{1}{2m+1}$. Hence, $\boldsymbol{\sigma}^P$ induces uniform marginals. It therefore follows from Lemma 4.5 above that $\boldsymbol{\sigma}^P$ is a symmetric equilibrium strategy. □

4.4 Understanding the Equilibrium Set

The equilibrium set turns out to be very large. To make this point, we study the support of arbitrary randomized equilibrium strategies. We first identify pure strategies that are chosen with positive probability in some mixed-strategy Nash equilibrium (Subsection 4.4.1). Then, we outline the proof (Subsection 4.4.2). Finally, we turn to strategies that are never "good" (Subsection 4.4.3).

4.4.1 Pure strategies that arise in some equilibrium

The following result gives an indication of the size of the equilibrium set.

Proposition 4.6. *Impose Assumptions 6 through 8. Then, every pure strategy s^P such that $s_i^P \leq \frac{2D}{n} = 2m$ for every battlefield i is used with positive probability in some randomized equilibrium strategy of $\mathcal{B}_\alpha(D, n)$.*

Proof. See Subsection 4.4.2. □

Thus, every pure strategy that does not allocate an excessive number of resources to an individual battlefield is used in some mixed-strategy Nash equilibrium.

In the Arad-Rubinstein game $\mathcal{B}_0(120, 6)$, every pure strategy that assigns at most 40 soldiers to each of the battlefields is part of some mixed strategy Nash equilibrium. This observation illustrates a drawback of the Nash equilibrium concept for the analysis of Colonel Blotto games.

The set of equilibrium predictions is simply very large. We will come back to this issue in the next section.

Proposition 4.6 relates to observations made by Tukey [1949] saying that "there are good strategies in which a given player either (i) sends out no units, (ii) sends out more than half of some kind of unit, or (iii) sends units to more than half of the available sites.

4.4.2 Proof of Proposition 4.6

To understand why Proposition 4.6 is true, suppose that both players uniformly randomize over the set of pure strategies

$$S_1 = \left\{ \begin{pmatrix} s_1 \\ 2m - s_1 \\ \vdots \\ s_L \\ 2m - s_L \end{pmatrix} : s_1, \ldots, s_L \in \{0, 1, \ldots, 2m\} \right\}.$$

We denote the mixed strategy that randomizes uniformly over S_1 as $\hat{\boldsymbol{\sigma}}$. Clearly, $\hat{\boldsymbol{\sigma}}$ induces uniform marginals. Lemma 4.5 implies, that both players are actually using an equilibrium strategy.

Let s be any pure strategy such that $s_i \leq 2m$ for all $i \in \{1, \ldots, n\}$. Then, it suffices to replace the two pure strategies

$$\begin{pmatrix} s_1 \\ 2m - s_1 \\ \vdots \\ s_{n-1} \\ 2m - s_{n-1} \end{pmatrix}, \begin{pmatrix} 2m - s_2 \\ s_2 \\ \vdots \\ 2m - s_i \\ s_i \end{pmatrix}$$

in the support of the mixed strategy $\hat{\boldsymbol{\sigma}}$ by

$$\begin{pmatrix} s_1 \\ s_2 \\ \vdots \\ s_{n-1} \\ s_n \end{pmatrix}, \begin{pmatrix} 2m - s_2 \\ 2m - s_1 \\ \vdots \\ 2m - s_n \\ 2m - s_{n-1} \end{pmatrix},$$

respectively. The marginals do not change, and hence, we have found an equilibrium strategy in which s is played with positive probability. This concludes the argument.

4.4.3 Strategies that are never "good"

In analogy to Proposition 4.6, one may ask what type of pure strategy is *never* used in equilibrium. To address this question, we identify the maximum loss of efficiency that is feasible with

modified tie-breaking. Then, we derive conditions on pure strategies that make them render an expected payoff too low to correspond to the maximum efficiency loss. Proceeding along these lines, we show that any strategy that focuses on too few battlefields will never be part of any mixed-strategy Nash equilibrium.

Proposition 4.7. *Impose Assumptions 6 and 7, and let $\alpha \in [0, 1)$. Then, any pure strategy that allocates the resource to strictly less than*

$$n^* = \frac{Dn}{2D+n}\left(1 - \frac{\alpha}{2-\alpha}\right)$$

battlefields is never part of any equilibrium of $\mathcal{B}_\alpha(D, n)$.[5]

Proof. Consider an arbitrary pure strategy s^P that allocates a positive number of units of the resource to $n_+ \in \{1, \dots, K\}$ battlefields and zero units to the remaining battlefields. Then, player i's expected payoff of s^P against an arbitrary pure strategy s^{-P} satisfies the relationship

$$\begin{aligned} \pi^P\left(s^P, s^{-P}\right) &\leq n_+ \cdot 1 + (n - n_+) \cdot \frac{\alpha}{2} \\ &= \frac{\alpha}{2} \cdot n + \frac{2-\alpha}{2} \cdot n_+. \end{aligned} \tag{4.3}$$

Consider now a deviation by player i to some mixed strategy $\boldsymbol{\sigma}^P$ that induces uniform marginals. Proposition 4.3 provides an example of such a mixed strategy. The following analysis is analogous to the proof of Lemma 4.5. If player $-P$ assigns strictly more than $2m$ units of the resource to battlefield i, then P's payoff from that battlefield is zero. If, however, player $-P$ assigns weakly less than $2m$ to battlefield i, then player P overwhelms her opponent on that battlefield with probability $\frac{2m - s_i^{-P}}{2m+1}$ and achieves a tie with probability $\frac{1}{2m+1}$. Hence, player P's expected payoff from battlefield i is

$$\pi_i^P\left(\boldsymbol{\sigma}^P, s^{-P}\right) = \max\{0, \frac{2m - s_i^{-P}}{2m+1} + \frac{\alpha}{2} \cdot \frac{1}{2m+1}\}.$$

It follows that

$$\begin{aligned} \pi^P\left(\boldsymbol{\sigma}^P, s^{-P}\right) &= \sum_{i=1}^{n} \pi_i^P\left(\boldsymbol{\sigma}^P, s^{-P}\right) \\ &\geq \sum_{i=1}^{n} \left(\frac{2m - s_i^{-P}}{2m+1} + \frac{\alpha}{2} \cdot \frac{1}{2m+1}\right) \\ &= \frac{Dn}{2D+n} + \frac{\alpha}{2} \cdot \frac{n^2}{2D+n}. \end{aligned} \tag{4.4}$$

[5] As α becomes larger, the conclusion of Proposition 4.7 weakens. For instance, the conclusion becomes void in the limit case as $\alpha \to 1$.

A straightforward calculation shows that the right-hand side of equation (4.4) strictly exceeds the right-hand side of equation (4.3) if and only if

$$n_+ < n^* \equiv \frac{2}{2-\alpha} \cdot \frac{Dn(1-\alpha)}{2D+n}.$$

nIn particular, in that case, $\pi^P\left(\boldsymbol{\sigma}^P, s^{-P}\right) > \pi^P\left(s^P, s^{-P}\right)$, for any $s^{-P} \in S$. Thus, s^P is never a best response if $n_+ < K^*$. This proves the proposition. □

How strong is the conclusion of Proposition 4.7 in specific games? For the Arad-Rubinstein game,

$$K^* = \frac{120 \cdot 6}{2 \cdot 120 + 6} = 2.9268.$$

This means that a pure strategy that assigns a positive number of resources to less than three battlefields is never used in equilibrium.

Corollary 4.8. *In the Arad-Rubinstein game, any pure strategy that allocates the resource to just one or two battlefields is never part of any Nash equilibrium.*

Proof. See the text above. □

Thus, pure strategies such as $(120, 0, 0, 0, 0, 0)^T$ and $(80, 40, 0, 0, 0, 0)^T$ are never "good" strategies in the example. Regrettably, this leaves a gap when compared to the conclusion of Proposition 4.6. Specifically, we do not know for the Arad-Rubinstein game if strategies such as $(60, 30, 30, 0, 0, 0)'$ that assign strictly positive resources over at least three battlefields and more than 40 soldiers to at least one of the battlefields are "good."[6]

4.5 Refinements

Given that the Arad-Rubinstein game is a static game, it is natural to check if dominance relationships between strategies might help to narrow down the set of equilibria (e.g., [Kohlberg and Mertens, 1986]).

[6]In the standard model, however, the converse of Proposition 4.6 holds by Dziubiński [2017, Cor. 2]. Thus, an arbitrary pure strategy s^P is used with positive probability in some equilibrium of $\mathcal{B}_1(D, n)$ if and only if $s_i^P \leq \frac{2D}{n}$ for every battlefield i. For the reader's convenience, we offer a direct proof. Consider a mixed-strategy equilibrium $\boldsymbol{\sigma}^* = (\boldsymbol{\sigma}^{A,*}, \boldsymbol{\sigma}^{B,*})$ in $\mathcal{B}_1(D, n)$. To provoke a contradiction, suppose that there is some player $P \in \{\mathsf{A}, \mathsf{B}\}$ and some pure strategy s^P in the support of $\boldsymbol{\sigma}^{P,*}$ such that $s_i^P > 2m$ for some battlefield i. As equilibrium strategies in two-person constant-sum game are interchangeable [Osborne and Rubinstein, 1994, p. 23], $\boldsymbol{\sigma}^{i,*}$ is a best response also to the mixed strategy $\boldsymbol{\sigma}^{-P} = \boldsymbol{\sigma}^P$ identified in Proposition 4.3. Moreover, since the bid vector s^P is chosen with positive probability in the mixed strategy $\boldsymbol{\sigma}^{P,*}$, the pure strategy s^P is likewise a best response to $\boldsymbol{\sigma}^{-P}$. But from $s_i^P > 2m$ for some battlefield i, inequality (4.2) in the proof of Lemma 4.5 is strict, so that that

$$\pi^P\left(s^P, \boldsymbol{\sigma}^{-P}\right) < \frac{n \cdot (2D + \alpha n)}{4D + 2n} = \pi^P\left(\boldsymbol{\sigma}^P, \boldsymbol{\sigma}^{-P}\right).$$

Thus, s^P is not a best response to $\boldsymbol{\sigma}^{-P}$ after all. The contradiction proves the assertion.

Recall that a pure strategy $s^P \in S$ for player $P \in \{\mathsf{A}, \mathsf{B}\}$ is *weakly dominated* by another pure strategy $\hat{s}^P \in S$ if (i) $\pi^P(s^P, s^{-P}) \leq \pi^P(\hat{s}^P, s^{-P})$ for every pure strategy $s^{-P} \in S$ of the opponent, and (ii) there exists at least one pure strategy $s^{-P} \in S$ for the opponent such that $\pi^P(s^P, s^{-P}) < \pi^P(\hat{s}^P, s^{-P})$. Thus, a pure strategy is weakly dominated by another pure strategy if it never yields a greater payoff than the other strategy, but a strictly lower payoff than the other strategy against at least one pure strategy of the opponent.

It turns out that the elimination of dominated strategies by pure strategies is entirely ineffective for small α.

Proposition 4.9. *Suppose that $\alpha < \frac{2}{K}$. Then, no pure strategy in $\mathcal{B}_\alpha(N, K)$ is weakly dominated by any other pure strategy.*

Proof. By contradiction. Suppose that s is a pure strategy that is weakly dominated by another pure strategy $\hat{s}$. In $\mathcal{B}_\alpha(N, K)$, the diagonal entries of the payoff matrix correspond to an outcome with n ties and are, therefore, equal to $n \cdot \frac{\alpha}{2}$. In particular, this is the payoff of $\hat{s}$ against itself. However, as $\hat{s}$ is necessarily different from s, the strategy s bids strictly higher than $\hat{s}$ on at least one battlefield. Therefore, the payoff of s against $\hat{s}$ is at least one. Under the assumption made, this is strictly higher than $n \cdot \frac{\alpha}{2}$. The contradiction shows that no pure strategy can be weakly dominated by any other pure strategy. □

The condition in the proposition is satisfied, in particular, in the Arad-Rubinstein game. There, given that identical choices of pure strategies by the two players lead to ties in all battlefields, $\alpha = 0$ implies that the diagonal entries of the payoff matrix are all zero. In contrast, all of the off-diagonal entries of the payoff matrix are positive because at least one battlefield is won by each player if strategies differ. Therefore, no pure strategy is weakly dominated by any other pure strategy if $\alpha = 0$. This idea of the proof generalizes in a straightforward way to positive but sufficiently small α.

Corollary 4.10. *In the Arad-Rubinstein game, there are no strategies that are weakly dominated by a pure strategy.*

Proof. In this case, $\alpha = 0$. The claim is, therefore, immediate from Proposition 4.9. □

4.6 Further Discussion

We first discuss the possibility of pure strategy Nash equilibria (Subsection 4.6.1), then payoff-nonequivalent equilibria (Subsection 4.6.2), finally the robustness of uniform marginals in the standard model (Subsection 4.6.3).

4.6.1 Equilibria in pure strategies

Pure-strategy Nash equilibria, both symmetric and asymmetric, are feasible if parameters are outside of the usual range. For instance, if the payoffs from ties are sufficiently high (i.e., if α is close to or exceeds 2), then the Colonel Blotto game transforms into a coordination game with multiple *symmetric* pure-strategy equilibria.

Proposition 4.11. *If $\alpha \geq \frac{2\cdot(K-1)}{K}$, then every pure strategy $s^P \in S$ is a symmetric Nash equilibrium strategy in $\mathcal{B}_\alpha(D,n)$.*

Proof. Take a symmetric profile $s = (s^P, s^{-P})$ in pure strategies, where $s^P = s^{-P}$. The payoff to both players from s is equal to $n \cdot \frac{\alpha}{2}$. Suppose first that $\alpha \leq 2$. Then, as a deviating player i gives up at least one battlefield, her payoff is bounded from above by $n-1$. But $n-1 \leq n \cdot \frac{\alpha}{2}$ holds by assumption, so that the deviation is not profitable. Suppose next that $\alpha > 2$. Then, changing the matching bid on any battlefield strictly lowers the payoff from that battlefield. Hence, a deviation is not be profitable either. This completes the proof. □

The pathological outcome indicated by Proposition 4.11 in cases where α is excessively large might provide a rationale for our Assumption 8.

Asymmetric pure-strategy equilibria may emerge when the number of resources D is small relative to the number of battlefields n (and Assumption 8 is in place). Specifically, if $2D \leq n$, adversaries can easily avoid any conflict by dividing the set of battlefields between them.

4.6.2 Payoff-nonequivalent equilibria

In the constant-sum model, all equilibria yield the same payoff. This is not the case with flexible tie-breaking, however. To understand why, let U_{O}^m and U_{E}^m denote uniform marginals on the odd and even numbers between 0 and $2m$, respectively. As shown by Hart [2008, Thm. 7], any pair of strategies that induce marginals in the convex hull of $\{U_{\mathrm{O}}^m, U_{\mathrm{E}}^m\}$ is an equilibrium in $\mathcal{B}_1(D,n)$. In particular, any strategy profile $\boldsymbol{\sigma} = (\boldsymbol{\sigma}^P, \boldsymbol{\sigma}^{-P})$, where $\boldsymbol{\sigma}_i^P = U_{\mathrm{O}}^m$ and $\boldsymbol{\sigma}_i^{-P} = U_{\mathrm{E}}^m$ for every battlefield i, is an *asymmetric* equilibrium in $\mathcal{B}_1(D,n)$. But any such profile remains an equilibrium for $\alpha \in [0,1]$, because inefficient tie-breaking makes deviations less attractive. Moreover, given the absence of ties, equilibrium payoffs are $n/2$ for both players, which is different from the expression in Proposition 1. Similarly, every mixed strategy $\boldsymbol{\sigma}^P$, where $\boldsymbol{\sigma}_i^P = U_{\mathrm{O}}^m$ or $\boldsymbol{\sigma}_i^P = U_{\mathrm{E}}^m$ for every battlefield i is a *symmetric* equilibrium strategy in $\mathcal{B}_1(D,n)$. Since deviations make ties less likely, these profiles remain equilibria for $\alpha \in [1,2]$. We ilustrate those results with the following Propositions.

Proposition 4.12. *Impose Assumptions 6 and 7 and let $\alpha \in [0,1]$. Every strategy profile $(\boldsymbol{\sigma}^P, \boldsymbol{\sigma}^{-P})$, where $\boldsymbol{\sigma}_k^P = U_O^m$ and $\boldsymbol{\sigma}_k^{-P} = U_E^m$ for every battlefield i is a Nash equilibrium of the Colonel Blotto game $\mathcal{B}_\alpha(D,n)$.*

Proof. Impose Assumptions 6 and 7 and consider a mixed strategy profile $\boldsymbol{\sigma}$, where $\boldsymbol{\sigma}_i^P = U_{\mathrm{O}}^m$ and $\boldsymbol{\sigma}_i^{-P} = U_{\mathrm{E}}^m$ for every battlefield i. The expected payoffs of both players are the same as in the Colonel Blotto game ($\frac{n}{2}$) as the considered mixed strategy profile results in no ties across all of the battlefields. On the other hand, the payoff of player C from playing against $\boldsymbol{\sigma}^{-P}$ is bounded by $\frac{n}{2}$ in the Colonel Blotto game $\mathcal{B}_1(D,n)$ (Hart, 2008), hence this bound holds for payoffs in $\mathcal{B}_\alpha(D,n)$, for all $\alpha \in [0,1]$, as the payoff in the $\mathcal{B}_\alpha(D,n)$ game are never greater than the payoffs in $\mathcal{B}_1(D,n)$ game for the same strategy profiles. As the expected payoffs of both players are already equal to $\frac{n}{2}$, none of the players can increase her expected payoff by deviating from $(\boldsymbol{\sigma}^P, \boldsymbol{\sigma}^{-P})$. □

Proposition 4.13. *Impose Assumptions 6 and 7 and let* $\alpha \in [1, 2]$. *Every strategy profile* $(\boldsymbol{\sigma}^P, \boldsymbol{\sigma}^{-P})$, *where*

1. $\boldsymbol{\sigma}_i^P = U_O^m$ *for every player* P *and every battlefield* i,
2. $\boldsymbol{\sigma}_i^P = U_E^m$ *for every player* P *and every battlefield* i,

is a Nash equilibrium of the Colonel Blotto game $\mathcal{B}_\alpha(D, n)$.

Proof. Impose Assumptions 6 and 7. Consider any two mixed strategies $\boldsymbol{\sigma}^P$, $\boldsymbol{\sigma}^{-P}$, such that $\boldsymbol{\sigma}_i^P = \boldsymbol{\sigma}_i^{-P} = U_{\mathrm{O}}^m$ for every battlefield i. Strategy profile $(\boldsymbol{\sigma}^P, \boldsymbol{\sigma}^{-P})$ results in payoff of

$$\frac{n}{2} + n \cdot \frac{\alpha - 1}{2} \cdot \frac{1}{m}. \tag{4.5}$$

The payoff is calculated as follows – first, we calculate the payoff in $\mathcal{B}_1(D, n)$, which is $\frac{n}{2}$. Then the different tie-breaking rule is included. As the probability of a tie on every single battlefield is equal to $\frac{1}{m}$ and there are n battlefields, this results in $n \cdot \frac{\alpha-1}{2} \cdot \frac{1}{m}$. No pure strategy s^{-P} of player $-P$ can result in a higher payoff to player $-P$ as the sum of probabilities of ties on single battlefields is already maximized as well as the payoff in $\mathcal{B}_1(D, n)$. Hence, player $-P$ has no incentive to deviate. The same holds for player P. The proof when $\boldsymbol{\sigma}_i^P = U_{\mathrm{E}}^m$ for every player P and every battlefield i is analogous. □

4.6.3 The robustness of uniform marginals

This section develops a simple refinement concept for the standard model. For this, let $\boldsymbol{\sigma}_{\text{even}}$ and $\boldsymbol{\sigma}_{\text{odd}}$ denote mixed strategies with marginals uniform on the even and odd integers within $\{0, \ldots, 2m\}$, respectively. As shown by Hart [2008] under Assumptions 6 and 7, such strategies exist. Moreover, *any* convex combination of such strategies is an equilibrium strategy in $B_1(D, n)$.

Proposition 4.14. *The only type of equilibrium in* $B_1(D, n)$ *that remains an equilibrium for all values of* α *in a neighborhood of* $\alpha = 1$ *has uniform marginals.*

Proof. The probability of a tie at any given battlefield is zero in $(\boldsymbol{\sigma}_{\text{even}}, \boldsymbol{\sigma}_{\text{odd}})$, and $(\boldsymbol{\sigma}_{\text{odd}}, \boldsymbol{\sigma}_{\text{even}})$. Similarly, the probability of a tie is $\frac{m+1}{2m+1}$ in $(\boldsymbol{\sigma}_{\text{even}}, \boldsymbol{\sigma}_{\text{even}})$, and $\frac{m}{2m+1}$ in $(\boldsymbol{\sigma}_{\text{odd}}, \boldsymbol{\sigma}_{\text{odd}})$. If player $-i$'s marginal is not uniform and $\alpha < 1$ ($\alpha > 1$), then player i has the incentive to deviate to the parity that is used less (more) often by $-i$. □

The proposition above suggests a simple form of equilibrium selection in finite Colonel Blotto games with standard tie-breaking. Symmetric equilibria $(\boldsymbol{\sigma}_{\text{even}}, \boldsymbol{\sigma}_{\text{even}})$ and $(\boldsymbol{\sigma}_{\text{odd}}, \boldsymbol{\sigma}_{\text{odd}})$ break down for $\alpha < 1$, because players seek to avoid the efficiency loss. Similarly, asymmetric equilibria such as $(\boldsymbol{\sigma}_{\text{even}}, \boldsymbol{\sigma}_{\text{odd}})$ break down for $\alpha > 1$ because players wish to realize the efficiency gain from matching the respective other's strategy. These consideration of robustness might explain why, in the standard model, any equilibrium with uniform marginals, i.e., without parity considerations, is intuitively more appealing than any of the other combinations.

4.7 Conclusions

In this chapter, we examined the Colonel Blotto game with flexible tie-breaking, extending the classical zero-sum formulation to settings where ties may be efficient or inefficient depending on a parameter α. From a computational perspective, this generalization provides insight into how small perturbations of the payoff structure influence both the existence and the form of equilibria in multi-battlefield conflicts.

Our analysis shows that equilibria with uniform marginals remain valid for arbitrary values of α, demonstrating a striking robustness of equilibrium structure beyond the constant-sum case. In particular, we established that the only equilibria of $B_1(D, n)$ that persist for all values of α in a neighborhood of $\alpha = 1$ are those with uniform marginals. This property identifies a class of equilibria that are not only conceptually central but also computationally stable—unchanged under small modifications of the tie-breaking rule. From an algorithmic standpoint, this robustness suggests that uniform-marginal equilibria form a natural focal point for equilibrium computation and verification.

At the same time, we showed that the equilibrium set in this model can be exponentially large, containing a wide range of mixed strategies that satisfy equilibrium conditions. This combinatorial richness illustrates the potential difficulty of algorithmic equilibrium enumeration and highlights the importance of structural restrictions or selection mechanisms. Moreover, we showed that standard dominance-based elimination techniques fail to reduce the strategy space when tie efficiency is low.

Overall, our results clarify which properties of Colonel Blotto equilibria are artifacts of the zero-sum assumption and which persist under a more general, non-constant-sum formulation.

Chapter 5

Equilibria of the Colonel Blotto Games with Costs

Chapter Overview

In this chapter, we study computation of Nash equilibria in a generalized variant of the Colonel Blotto game, called the *Colonel Blotto game with costs*, which captures both global costs of obtaining the resources as well as battlefield-specific costs of assigning resources to the given battlefield. Although the game in question is not zero-sum, using the notion of strategic equivalence, we show that the Colonel Blotto game with costs satisfies the interchangeability of Nash equilibria. In consequence, the notion of *equilibrium strategy* of a player holds, and every pair of equilibrium strategies is a Nash equilibrium. Furthermore, we prove that every instance of the Colonel Blotto game with costs is strategically equivalent to a corresponding variant of the standard Colonel Blotto game where players have fixed budgets. By applying the results of Behnezhad et al. [2023] about solution computation in the standard Colonel Blotto games with fixed budgets, we show that the problem of computation of Nash equilibria of the Colonel Blotto game with costs can be expressed using a linear program with at most $\Theta(D^2 \cdot n)$ constraints and $\Theta(D^2 \cdot n)$ variables, where n is the number of battlefields and D is the maximum number of resources that any player is allowed to assign. Lastly, we present our experimental results and discuss their consistency with theoretical results for the continuous resources setting of Roberson and Kvasov [2010].

The rest of this chapter is organized as follows. In Section 5.1 we provide motivation for studying the problem of equilibria computation of the Colonel Blotto game with costs. In Section 5.2, we formally define the model in question. The theoretical analysis of the equilibria is provided in Section 5.3, and Section 5.4 presents the high-level idea of computing the equilibria of the game in question. We report and discuss the experimental results in Section 5.5 and conclude in Section 5.6.

5.1 Motivating Examples

In the competitive landscape of the global electric vehicle (EV) market, Tesla and Volkswagen are two major players making strategic decisions about increasing and allocating production capacity across multiple key markets. Each market has unique characteristics, such as production costs, market potential, and regulations, that influence both the potential return on investment and the costs associated with expanding production capacity. Furthermore, the costs of increasing production capacity in any given market may differ among companies due to factors like existing infrastructure, supply chain efficiency, and local partnerships. This creates a dual challenge: first, deciding how much to invest in expanding overall production capacity, and second, determining how to allocate this capacity across different markets to maximize competitive advantage. Strategic planners must carefully weigh these considerations to decide how to globally increase and distribute production capacity. The central questions are: How should a company balance its investments in expanding capacity with the costs and benefits of allocating it to specific markets? What is the optimal strategy for assigning the increased capacity across markets with varying cost structures? Finally, what are the total costs that each company incurs in pursuing this global expansion?

Similarly, in the competitive pharmaceutical industry, companies like Pfizer and Johnson & Johnson are engaged in a strategic race to develop new drugs. They must decide how much to invest in expanding their R&D capacity—such as hiring scientists and expanding lab facilities—-and how to allocate these resources to different research areas such as oncology, immunology, and neurology. Each area has its own set of challenges, including market potential, scientific difficulty, and regulatory requirements, which affect both the cost and the potential return on investment. In addition, the cost of conducting research in these areas can vary between companies due to factors such as existing expertise, access to specialized equipment, and previous partnerships with research institutions. The key strategic questions are: How should a company increase its R&D capacity? How should it allocate this capacity across various research areas to balance risk, cost, and reward? What are the overall costs for each company in this R&D competition?

Strategic decisions regarding the aforementioned scenarios share similar characteristics. First, they involve allocating resources across multiple contests, where strategic planners aim to maximize the sum of utilities across all contests. Second, both obtaining and assigning resources come at a cost. A strategic planner aiming to maximize the outcome of her decisions needs to consider both the anticipated outcomes of contests as well as the total costs incurred. The Colonel Blotto game [Borel, 1921] is a well-known model that captures the strategic problem of allocating limited resources across multiple contests. However, in its classic formulation, the Colonel Blotto game imposes a sunk-cost assumption, i.e., the budgets of competing parties are fixed, and there are no costs related to either obtaining resources or assigning them to specific battlefields. Since the sunk-cost assumption does not capture the scenarios mentioned above, we study the problem of computing equilibrium strategies in a generalized variant of the Colonel Blotto game, called the Colonel Blotto game with costs.

5.2 Model

This section introduces the Colonel Blotto game with discrete resources and costs, a generalization of the classic Colonel Blotto game that captures the strategic costs of assigning resources over a set of battlefields. The Colonel Blotto game with costs models competition between two players denoted by A and B. Each player has a number of available discrete resources, e.g. military units or coins, denoted by D^{A} and D^{B}, respectively. Players compete on the set $[n] = \{1, 2, \dots n\}$ of n battlefields, where $n \geq 2$. A strategy of player $P \in \{\mathsf{A}, \mathsf{B}\}$ is an assignment of at most D^P resources across the set of battlefields. Formally, the set of strategies of player $P \in \{\mathsf{A}, \mathsf{B}\}$ is

$$H(D^P, n) = \left\{ \boldsymbol{s} \in \mathbb{N}^n : \sum_{i=1}^{n} s_i \leq D^P \right\}. \tag{5.1}$$

Elements of $H(D^P, n)$ are partial assignments, as the players are not required to use all of their available resources. When the parameters are clear from the context, the set of partial assignments is denoted H^P.

Given a strategy profile $(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}})$, the payoffs to the players from battlefield $i \in [n]$ are defined as

$$\begin{aligned} u_i^{\mathsf{A}}\left(s_i^{\mathsf{A}}, s_i^{\mathsf{B}}\right) &= \quad v_i\left(s_i^{\mathsf{A}}, s_i^{\mathsf{B}}\right) - c_i^{\mathsf{A}}\left(s_i^{\mathsf{A}}\right), \\ u_i^{\mathsf{B}}\left(s_i^{\mathsf{A}}, s_i^{\mathsf{B}}\right) &= -v_i\left(s_i^{\mathsf{A}}, s_i^{\mathsf{B}}\right) - c_i^{\mathsf{B}}\left(s_i^{\mathsf{B}}\right), \end{aligned}$$

where $v_i : \mathbb{N} \times \mathbb{N} \longrightarrow \mathbb{R}$ is a valuation function associated with the battlefield i, and $c_i^{\mathsf{A}}, c_i^{\mathsf{B}} : \mathbb{N} \longrightarrow \mathbb{R}$ are non-decreasing *assignment cost functions* that capture the costs of assigning a given number of resources to the given battlefield by each of the players.

The payoff to player $P \in \{\mathsf{A}, \mathsf{B}\}$ from strategy profile $(\boldsymbol{s}^P, \boldsymbol{s}^{-P})$ is

$$\pi_{\$}^P\left(\boldsymbol{s}^P, \boldsymbol{s}^{-P}\right) = \sum_{i=1}^{n} u_i^P\left(s_i^P, s_i^{-P}\right) - g^P\left(\sum_{i=1}^{n} s_i^P\right), \tag{5.2}$$

where g^P is a non-decreasing *obtainment cost* function that describes the cost of obtaining the total number of resources assigned across all the battlefields.

Definition 5.1 (Colonel Blotto game with costs)**.** For a given tuple $(D^{\mathsf{A}}, D^{\mathsf{B}}, n, \boldsymbol{v}, \boldsymbol{c}^{\mathsf{A}}, \boldsymbol{c}^{\mathsf{B}}, g^{\mathsf{A}}, g^{\mathsf{B}})$, a strategic game with strategy sets defined by (5.1) and payoffs defined by (5.2) is denoted with $\mathcal{B}_{\$}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, \boldsymbol{v}, \boldsymbol{c}^{\mathsf{A}}, \boldsymbol{c}^{\mathsf{B}}, g^{\mathsf{A}}, g^{\mathsf{B}}) \equiv \mathcal{B}_{\$}$ and called the **Colonel Blotto game with costs**.

The players are allowed to make randomized choices. A *mixed strategy* of player $P \in \{\mathsf{A}, \mathsf{B}\}$ is a probability distribution on H^P. Given a non-empty set X, let $\Delta(X)$ denote the set of all probability distributions on X. The expected payoff to player $P \in \{\mathsf{A}, \mathsf{B}\}$ from a pair of mixed strategies $(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}}) \in \Delta(H^{\mathsf{A}}) \times \Delta(H^{\mathsf{B}})$ is equal to

$$\Pi_f^P\left(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}}\right) = \sum_{(\boldsymbol{x}, \boldsymbol{z}) \in H^{\mathsf{A}} \times H^{\mathsf{B}}} \xi_{\boldsymbol{x}}^{\mathsf{A}} \xi_{\boldsymbol{z}}^{\mathsf{B}} \pi_{\$}^P(\boldsymbol{x}, \boldsymbol{z}),$$

where given the probability distribution ξ^P, $\xi^P_{\boldsymbol{x}}$ is the probability associated with pure strategy $\boldsymbol{x}$. We assume that the players make their decisions "simultaneously", i.e., each player chooses her strategy without observing the choice of the opponent. We are interested in mixed strategy Nash equilibria (MSNE) of the game, i.e., mixed strategy profiles $(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}}) \in \Delta(H^{\mathsf{A}}) \times \Delta(H^{\mathsf{B}})$ such that no player can improve her expected payoff by changing her strategy unilaterally. Formally, for each $P \in \{\mathsf{A}, \mathsf{B}\}$ and all $\boldsymbol{\zeta} \in \Delta\left(S^P\right)$,

$$\Pi_f^P\left(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}}\right) \geq \Pi_f^P\left(\boldsymbol{\zeta}, \boldsymbol{\xi}^{-P}\right),$$

where $-P$ denotes the player other than P and $(\boldsymbol{\zeta}, \boldsymbol{\xi}^{-P})$ is the strategy profile obtained from $(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}})$ by replacing $\boldsymbol{\xi}^P$ with $\boldsymbol{\zeta}$.

5.3 Analysis

In this section, we present one of two key ideas of this chapter, that is to introduce an artificial zero-sum game, $\mathcal{B}_0$, that, given the same set of parameters, has the same set of Nash equilibria as the Colonel Blotto game with costs, $\mathcal{B}_\$$. The relation between the two games provides insight into the properties of equilibria of the Colonel Blotto game with costs.

For the earlier introduced set of parameters, let $\mathcal{B}_0(D^{\mathsf{A}}, D^{\mathsf{B}}, n, \boldsymbol{v}, \boldsymbol{c}^{\mathsf{A}}, \boldsymbol{c}^{\mathsf{B}}, g^{\mathsf{A}}, g^{\mathsf{B}}) \equiv \mathcal{B}_0$ denote a two-player strategic game with strategy sets of both players defined by (5.1). The payoff to player $P \in \{\mathsf{A}, \mathsf{B}\}$ in the game $\mathcal{B}_0$ under strategy profile $(\boldsymbol{s}^P, \boldsymbol{s}^{-P})$ is

$$\pi_0^P\left(\boldsymbol{s}^P, \boldsymbol{s}^{-P}\right) = \pi_\$^P\left(\boldsymbol{s}^P, \boldsymbol{s}^{-P}\right) + \sum_{i=1}^{n} c_i^{-P}\left(s_i^{-P}\right) + g^{-P}\left(\sum_{i=1}^{n} s_i^{-P}\right).$$

Game $\mathcal{B}_0$ is obtained from the corresponding game $\mathcal{B}_\$$ by adding to the payoff of each player the sum of costs (both general and battlefield-specific) carried by her opponent, resulting in the zero-sum property. The following proposition characterizes the relation between the Nash equilibria of $\mathcal{B}_\$$ and $\mathcal{B}_0$ given the same set of parameters.

Proposition 5.2. *For the same set of parameters $(D^{\mathsf{A}}, D^{\mathsf{B}}, n, \boldsymbol{v}, \boldsymbol{c}^{\mathsf{A}}, \boldsymbol{c}^{\mathsf{B}}, g^{\mathsf{A}}, g^{\mathsf{B}})$, games $\mathcal{B}_\$$ and $\mathcal{B}_0$ have the same set of Nash equilibria.*

Proof. Fix player $P \in \{\mathsf{A}, \mathsf{B}\}$, two strategies $\boldsymbol{s}^P, \boldsymbol{t}^P \in H^P$ and a strategy $s^{-P} \in H^{-P}$ of player $-P$. It holds, that

$$\pi_0^P\left(\boldsymbol{s}^P, \boldsymbol{s}^{-P}\right) - \pi_0^P\left(\boldsymbol{t}^P, \boldsymbol{s}^{-P}\right) = \pi_\$^P\left(\boldsymbol{s}^P, \boldsymbol{s}^{-P}\right) - \pi_\$^P\left(\boldsymbol{t}^P, \boldsymbol{s}^{-P}\right).$$

Hence, the two games are strategically equivalent in the sense that for every player $P \in \{\mathsf{A}, \mathsf{B}\}$, every two strategies $\boldsymbol{s}^P, \boldsymbol{t}^P \in H^P$ and every strategy of the opponent $s^{-P} \in H^{-P}$, it holds, that

$$\pi_0^P\left(\boldsymbol{s}^P, \boldsymbol{s}^{-P}\right) \geq \pi_0^P\left(\boldsymbol{t}^P, \boldsymbol{s}^{-P}\right) \iff \pi_\$^P\left(\boldsymbol{s}^P, \boldsymbol{s}^{-P}\right) \geq \pi_\$^P\left(\boldsymbol{t}^P, \boldsymbol{s}^{-P}\right).$$

As a consequence of Lemma 1 of Moulin and Vial [1978], the two games have the same set of Nash equilibria. □

5.3.1 Properties of Equilibria

This section provides a discussion of the properties of Nash equilibria of the Colonel Blotto game with costs. First, we give two corollaries that follow directly from Proposition 5.2. A strategic two-player game satisfies the *interchangeability of Nash equilibria* [Nash, 1951, Naumov and Nicholls, 2013] if for every two (possibly mixed) strategy profiles $(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}}), (\boldsymbol{\sigma}^{\mathsf{A}}, \boldsymbol{\sigma}^{\mathsf{B}}) \in \Delta(H^{\mathsf{A}}) \times \Delta(H^{\mathsf{B}})$ that are Nash equilibria of game $\mathcal{B}_{\$}$, it holds that strategy profiles $(\boldsymbol{\xi}^{\mathsf{A}}, \boldsymbol{\sigma}^{\mathsf{B}})$ and $(\boldsymbol{\sigma}^{\mathsf{A}}, \boldsymbol{\xi}^{\mathsf{B}})$ are also Nash equilibria of game $\mathcal{B}_{\$}$. As every finite two-player zero-sum game satisfies the interchangeability of Nash equilibria, the following corollary holds.

Corollary 5.3. *The Colonel Blotto game with costs satisfies the interchangeability of Nash equilibria.*

Note that one of the consequences of the interchangeability of Nash equilibria in a two-player finite game is that every player has a set of equilibrium strategies, and every pair of equilibrium strategies is a Nash equilibrium of the game. In the case of zero-sum games, equilibria strategies are called optimal strategies. As the Colonel Blotto game with costs is not zero-sum, we use the term *equilibrium strategy*. Since the set of optimal strategies in every finite two-player zero-sum game is convex, the following corollary is also a direct consequence of Proposition 5.2. Formally, this is a direct consequence of the the work of Chin et al. [1974], stating that for any two-player finite game the interchangeability and convexity of the set of equilibrium points are synonymous.

Corollary 5.4. *For each player P the set of equilibrium strategies of player P is convex in every Colonel Blotto game with costs.*

We illustrate the consequences of these corollaries in the following example.

Example 5.5. Consider an instance $\mathcal{B}^*_{\$}$ of the Colonel Blotto game with costs, where

- the players compete on two battlefields,
- each player can assign at most 2 resources,
- each battlefield is won by the player assigning strictly more resources to it (and tied otherwise),
- obtaining a unit of resources yields a cost of 1 for each of the players,
- assigning resources across battlefields yields no costs.

Formally, $\mathcal{B}^*_{\$} \equiv \mathcal{B}_{\$}(2, 2, 2, \boldsymbol{v} = (\mathbf{sign}), \mathbf{0}, \mathbf{0}, g^{\mathsf{A}}, g^{\mathsf{B}})$, where

$$g^{\mathsf{A}}(\boldsymbol{s}) = g^{\mathsf{B}}(\boldsymbol{s}) = s_1 + s_2.$$

The strategy sets of both players are

$$H^{\mathsf{A}} = H^{\mathsf{B}} = \{(0,0), (0,1), (1,0), (1,1), (0,2), (2,0)\},$$

and it can be easily calculated, that every strategy in

$$H^* = \{(0,0),(0,1),(1,0),(1,1)\},$$

is an equilibrium strategy of $\mathcal{B}^*_\$$, while the remaining pure strategies $(0,2),(2,0)$ are never used with positive probability in any equilibrium strategy. A consequence of Corollary 5.4 is that every mixed strategy with support restricted to any subset of H^* is also an equilibrium strategy of $\mathcal{B}^*_\$$. Figure 7 present the payoff matrix of game $\mathcal{B}^*_\$$. Parts of the payoff matrix corresponding to the pure strategies outside of H^* are shadowed, to focus attention on the strategies in H^*. If the opponent's equilibrium strategy is fixed, the player cannot change her payoff by deviating

	(0,0)	(0,1)	(1,0)	(1,1)	(0,2)	(2,0)
$(0,0)$	0, 0	-1, 0	-1, 0	-2, 0	-1, -1	-1, -1
(0,1)	0, -1	-1, -1	-1, -1	-2, -1	-1, -2	-2, -1
(1,0)	0, -1	-1, -1	-1, -1	-2, -1	-2, -1	-1, -2
(1,1)	0, -2	-1, -2	-1, -2	-2, -2	-2, -2	-2, -2
(0,2)	-1, -1	-1, -2	-2, -1	-2, -2	-2, -2	-2, -2
(2,0)	-1, -1	-2, -1	-1, -2	-2, -2	-2, -2	-2, -2

Figure 7: The payoff matrix of game $\mathcal{B}^*_\$$.

from one strategy in H^* to another strategy in H^* (this must be the case as every strategy in H^* is an equilibrium strategy). However, by changing her strategy, the player can influence her opponent's equilibrium payoff.

5.4 Equilibrium Computation

In this section, we explain how every instance of $\mathcal{B}_0$ can be expressed as a variant of the Colonel Blotto game with fixed budgets under a well-defined strategy mapping. As a consequence, the problem of computing an equilibrium strategy of a player in the Colonel Blotto game with costs can be expressed as a linear program with a polynomial (in the number of battlefields and resources of both players) number of both constraints and variables using the BDDHS[1] linear program of Behnezhad et al. [2023]).

5.4.1 Colonel Blotto Game with Sunk Costs

In the classic Colonel Blotto game, a strategy of all player $P \in \{\mathsf{A},\mathsf{B}\}$ is an assignment of all of her available D^P resources over the set of the $[n]$ battlefields. Hence, the strategy set of a player P is

$$S(D^P,n) = \left\{ \boldsymbol{s} \in \mathbb{N}^n : \sum_{i=1}^{n} s_i = D^P \right\}, \tag{5.3}$$

[1]The name of the linear program comes from the first letters of authors surnames.

which constitutes a subset of the player's strategy set $H(D^P, n)$ of the Colonel Blotto game with costs. Elements of set $S(D^P, n)$ are called *full assignments*, as players are required to assign all of their available resources. When the parameters are clear from the context, the set of full assignments is denoted S^P.

Definition 5.6. The Colonel Blotto game with costs, $\mathcal{B}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, \boldsymbol{v}, \mathbf{0}, \mathbf{0}, \mathbf{0}, \mathbf{0})$, where all costs functions are trivial (equal to $\mathbf{0}$) and strategy sets of both players are restricted to full assignments (5.3) is called the **Colonel Blotto game with sunk costs** and denoted $\mathcal{B}(D^{\mathsf{A}}, D^{\mathsf{B}}, n, \boldsymbol{v}) \equiv \mathcal{B}$.

5.4.2 Expressing $\mathcal{B}_0$ as $\mathcal{B}$

For a given game $\mathcal{B}_0(D^{\mathsf{A}}, D^{\mathsf{B}}, n, \boldsymbol{v}, \boldsymbol{c}^{\mathsf{A}}, \boldsymbol{c}^{\mathsf{B}}, g^{\mathsf{A}}, g^{\mathsf{B}})$, consider the *corresponding* Colonel Blotto game with sunk costs $\mathcal{B}(D^{\mathsf{A}}, D^{\mathsf{B}}, n+1, \hat{\boldsymbol{v}})$, with valuation functions $\hat{v}_i : \mathbb{N} \times \mathbb{N} \longrightarrow \mathbb{R}$ defined as

$$\hat{v}_i(s_i^{\mathsf{A}}, s_i^{\mathsf{B}}) = \begin{cases} v_i\left(s_i^{\mathsf{A}}, s_i^{\mathsf{B}}\right) - c_i^{\mathsf{A}}\left(s_i^{\mathsf{A}}\right) + c_i^{\mathsf{B}}\left(s_i^{\mathsf{B}}\right), & \text{if } i \in \{1, 2, \ldots, n\}, \\ -g^{\mathsf{A}}\left(D^{\mathsf{A}} - s_{n+1}^{\mathsf{A}}\right) + g^{\mathsf{B}}\left(D^{\mathsf{B}} - s_{n+1}^{\mathsf{B}}\right), & \text{if } i = n+1. \end{cases}$$

The following result states that the Colonel Blotto game with costs and the corresponding Colonel Blotto game with sunk costs are strategically equivalent.

Proposition 5.7. *For a given a tuple of parameters, the game $\mathcal{B}_0(D^{\mathsf{A}}, D^{\mathsf{B}}, n, \boldsymbol{v}, \boldsymbol{c}^{\mathsf{A}}, \boldsymbol{c}^{\mathsf{B}}, g^{\mathsf{A}}, g^{\mathsf{B}})$ and the corresponding Colonel Blotto game with sunk costs $\mathcal{B}(D^{\mathsf{A}}, D^{\mathsf{B}}, n+1, \hat{\boldsymbol{v}})$ are the same game under the bijective strategy mapping $\boldsymbol{m} : H(D^P, n) \longrightarrow S(D^P, n+1)$, such that*

$$\boldsymbol{m}(\boldsymbol{s}^P) = \left(s_1^P, s_2^P, \ldots, s_n^P, D^P - \sum_{i=1}^{n} s_i^P\right). \tag{5.4}$$

Proof. For every strategy profile $(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}})$ it holds that

$$\begin{aligned}
&\pi_0^{\mathsf{A}}\left(\boldsymbol{s}^{\mathsf{A}}, \boldsymbol{s}^{\mathsf{B}}\right) \\
&= \sum_{i=1}^{n}\left(v_i\left(s_i^{\mathsf{A}}, s_i^{\mathsf{B}}\right) - c_i^{\mathsf{A}}\left(s_i^{\mathsf{A}}\right) + c_i^{\mathsf{B}}\left(s_i^{\mathsf{B}}\right)\right) - g^{\mathsf{A}}\left(\sum_{i=1}^{n} s_i^{\mathsf{A}}\right) + g^{\mathsf{B}}\left(\sum_{i=1}^{n} s_i^{\mathsf{B}}\right) \\
&= \sum_{i=1}^{n}\left(v_i\left(s_i^{\mathsf{A}}, s_i^{\mathsf{B}}\right) - c_i^{\mathsf{A}}\left(s_i^{\mathsf{A}}\right) + c_i^{\mathsf{B}}\left(s_i^{\mathsf{B}}\right)\right) - g^{\mathsf{A}}\left(D^{\mathsf{A}} - m_{n+1}\left(\boldsymbol{s}^{\mathsf{A}}\right)\right) + g^{\mathsf{B}}\left(D^{\mathsf{B}} - m_{n+1}\left(\boldsymbol{s}^{\mathsf{B}}\right)\right) \\
&= \sum_{i=1}^{n+1} \hat{v}_i\left(m_i\left(\boldsymbol{s}^{\mathsf{A}}\right), m_i\left(\boldsymbol{s}^{\mathsf{B}}\right)\right) \\
&= \pi^{\mathsf{A}}\left(\boldsymbol{m}\left(\boldsymbol{s}^{\mathsf{A}}\right), \boldsymbol{m}\left(\boldsymbol{s}^{\mathsf{B}}\right)\right),
\end{aligned}$$

where $\pi^{\mathsf{A}}(\boldsymbol{m}(\boldsymbol{s}^{\mathsf{A}}), \boldsymbol{m}\left(\boldsymbol{s}^{\mathsf{B}}\right))$ is the payoff to player A in the Colonel Blotto game with sunk costs $\mathcal{B}(D^{\mathsf{A}}, D^{\mathsf{B}}, n+1, \hat{\boldsymbol{v}})$ under strategy profile $(\boldsymbol{m}(\boldsymbol{s}^{\mathsf{A}}), \boldsymbol{m}\left(\boldsymbol{s}^{\mathsf{B}}\right)) \in S(D^{\mathsf{A}}, n+1) \times S(D^{\mathsf{B}}, n+1)$.

The payoff equality in both games holds for player B as well (the reason for using player A is purely technical, as it determines the sign of v_i functions). Because of that, the two games are equivalent under the considered strategy mapping. □

We conclude this subsection with the following observation, regarding the relation between the strategies of the Colonel Blotto game with costs and the strategies of the corresponding Colonel Blotto game with sunk costs.

Observation 5.8. Consider the Colonel Blotto game with costs, $\mathcal{B}_\$(D^\mathsf{A}, D^\mathsf{B}, n, \boldsymbol{v}, \boldsymbol{c}^\mathsf{A}, \boldsymbol{c}^\mathsf{B}, g^\mathsf{A}, g^\mathsf{B}) \equiv \mathcal{B}_\$$, and the corresponding Colonel Blotto game with sunk costs, $\mathcal{B}(D^\mathsf{A}, D^\mathsf{B}, n+1, \hat{\boldsymbol{v}}) \equiv \mathcal{B}$. Fix Player P and her pure strategy $\boldsymbol{s}^P \in S(D^P, n+1)$ in $\mathcal{B}$. The total number of resources obtained by player P according to the corresponding pure strategy $\boldsymbol{m}^{-1}(\boldsymbol{s}^P) \in S(D^P, n)$ of $\mathcal{B}_\$$ is exactly $D^P - s^P_{n+1}$.

5.4.3 Computational Complexity

As a consequence of Propositions 5.2 and 5.7, the problem of computing an equilibrium strategy of the Colonel Blotto game with costs can be restated as a problem of computing an optimal strategy of the corresponding Colonel Blotto game with sunk costs with one additional battlefield and the same numbers of resources of both players.

Proposition 5.9. *The problem of computing a player equilibrium strategy of the Colonel Blotto game with costs $\mathcal{B}_\$(D^\mathsf{A}, D^\mathsf{B}, n, \boldsymbol{v}, \boldsymbol{c}^\mathsf{A}, \boldsymbol{c}^\mathsf{B}, g^\mathsf{A}, g^\mathsf{B})$ can be formulated as a linear program with at most $\Theta((D^{\max})^2 \cdot n)$ constraints and $\Theta((D^{\max})^2 \cdot n)$ variables, where $D^{\max} = \max\{D^\mathsf{A}, D^\mathsf{B}\}$.*

Proof. Consider an instance of the Colonel Blotto game with costs, $\mathcal{B}_\$(D^\mathsf{A}, D^\mathsf{B}, n, \boldsymbol{v}, \boldsymbol{c}^\mathsf{A}, \boldsymbol{c}^\mathsf{B}, g^\mathsf{A}, g^\mathsf{B})$. By Proposition 5.2, computing an equilibrium strategy of $\mathcal{B}_\$$ is equivalent to computing optimal strategy of the corresponding game $\mathcal{B}_0(D^\mathsf{A}, D^\mathsf{B}, n, \boldsymbol{v}, \boldsymbol{c}^\mathsf{A}, \boldsymbol{c}^\mathsf{B}, g^\mathsf{A}, g^\mathsf{B})$. By Proposition 5.7, this is equivalent to computing an optimal strategy of the corresponding Colonel Blotto game with sunk costs, $\mathcal{B}(D^\mathsf{A}, D^\mathsf{B}, n+1, \hat{\boldsymbol{v}})$. By Theorem 1 [Behnezhad et al., 2023], an optimal strategy of the Colonel Blotto game with sunk costs, $\mathcal{B}(D^\mathsf{A}, D^\mathsf{B}, n+1, \hat{\boldsymbol{v}})$ can be expressed as a linear program with at most $\Theta((D^{\max})^2 \cdot (n+1))$ constraints and $\Theta((D^{\max})^2 \cdot (n+1))$ variables, where $D^{\max} = \max\{D^\mathsf{A}, D^\mathsf{B}\}$. Since $\Theta((D^{\max})^2 \cdot (n+1)) = \Theta((D^{\max})^2 \cdot n)$, this concludes the proof. □

As every pair of equilibrium strategies is a Nash equilibrium of the Colonel Blotto game with costs (Corollary 5.3), the following theorem is a direct consequence of Proposition 5.9.

Theorem 5.10. *A Nash equilibrium of the Colonel Blotto game with costs can be computed in polynomial time with respect to the number of battlefields and the number of resources of both players.*

We conclude this section with the high-level idea of our implementation of the program computing an equilibrium strategy of the Colonel Blotto game with costs. In the description of our implementations we refer to the linear program BDDHS proposed by Behnezhad et

al. (2023), which, given an instance of the Colonel Blotto game with sunk costs, $\mathcal{B}$, defines a set of variables and constraints capturing any mixed strategy of a player. For clarity and conciseness, we only refer to the variables we operate on directly:

- u^{B} - a variable of the program describing the bound on the players' B maximal payoff in $\mathcal{B}_0(n, D^{\mathsf{A}}, D^{\mathsf{B}}, \boldsymbol{v})$;
- $P_i(t)$ - a variable describing the probability of player A assigning exactly t resources to battlefield i.

Algorithm 4 describes the idea of our implementation. Formally, Algorithm 4 returns a

Algorithm 4 Computing an equilibrium strategy of player A of the Colonel Blotto game with costs.

Input: A set of parameters $(D^{\mathsf{A}}, D^{\mathsf{B}}, n, \boldsymbol{v}, \boldsymbol{c}^{\mathsf{A}}, \boldsymbol{c}^{\mathsf{B}}, g^{\mathsf{A}}, g^{\mathsf{B}})$.
Output: Linear program describing an equilibrium strategy of player A.

1: LP = BDDHH($\mathcal{B}(D^{\mathsf{A}}, D^{\mathsf{B}}, n+1, \hat{\boldsymbol{v}})$)
   # Payoff function vector $\hat{\boldsymbol{v}}$ is defined as in
   # Equation (5.4).
2: LP.minimize(u^{B})
3: **return** LP

linear program describing an optimal strategy of the Colonel Blotto game with sunk costs $\mathcal{B}(D^{\mathsf{A}}, D^{\mathsf{B}}, n+1, \hat{\boldsymbol{v}})$. However, the inverse m^{-1} of strategy mapping m defined in Equation (5.4) defines the exact relation between any optimal strategy described by LP and the corresponding equilibrium strategy of the Colonel Blotto game with costs. As stated by Observation 5.8, for a given number of resources $t \in \{0, 1, \ldots, D^{\mathsf{A}}\}$, the variable $P_{n+1}(t)$ in the resulting LP program describes the probability that player A obtained exactly $(D^{\mathsf{A}} - t)$ resources in the equilibrium strategy.

5.5 Experimental Results

In this section, we report on the carried-out experiments, focusing on the amount of resources players obtain in equilibrium. First, we provide a high-level description of the implemented program, used for finding the lower (upper) bound on the number of resources obtained in equilibrium. Second, we report and discuss the obtained results. Lastly, we report the running times of our program.

In order to investigate the number of resources the players obtain in equilibrium, we introduced a new variable, D, which describes the expected number of resources obtained by player A. Algorithm 5 describes the high-level idea of computing the maximal (minimal) number of resources obtained by a player in equilibrium.

Algorithm 5 Computing the maximal (minimal) expected number of resources in any equilibrium strategy of player A of the Colonel Blotto game with costs.

Input: A set of parameters $(D^{\mathsf{A}}, D^{\mathsf{B}}, n, \boldsymbol{v}, \boldsymbol{c}^{\mathsf{A}}, \boldsymbol{c}^{\mathsf{B}}, g^{\mathsf{A}}, g^{\mathsf{B}})$.
Output: Maximal (minimal) number of resources obtained by player A in equilibrium.

1: LP = BDDHH($\mathcal{B}(D^{\mathsf{A}}, D^{\mathsf{B}}, n+1, \hat{\boldsymbol{v}})$)
   # Payoff function vector $\hat{\boldsymbol{v}}$ is defined as in
   # Equation (5.4).
2: LP.minimize(u^{B})
3: LP.fixVariableValue(u^{B})
   # LP is restricted to equilibrium strategies.
4: LP.addVariable(D)
5: LP.addConstraint(D $= \sum_{t=0}^{D^{\mathsf{A}}} P_{n+1}(t) \cdot (D^{\mathsf{A}} - t)$)
   # D describes the expected number of resources
   # obtained by player A.
6: LP.maximize(D)
7: **return** D

We proceed to report our results. Specifically, we investigate the relation between the costs of obtaining resources and the number of resources players obtain in equilibrium of the Colonel Blotto game with costs, using the implementation of Algorithm 5. We consider models with linear costs, $\mathcal{B}_{\$}$, such that

$$D^{\mathsf{A}} = D^{\mathsf{B}}, \text{ and}$$
$$\pi^P(\boldsymbol{s}^P, \boldsymbol{s}^{-P}) = \sum_{i=1}^{n} \operatorname{sign}(s_i^P - s_i^{-P}) - c_0 \cdot \sum_{i=1}^{n} s_i^P,$$

for some unit costs coefficient c_0. The following hypothesis captures the structure of the preliminary results of our experiments.

Hypothesis 1. All of the results concerning the equilibrium number of resources, denoted by $D(D^{\mathsf{A}}, n, c)$, satisfy

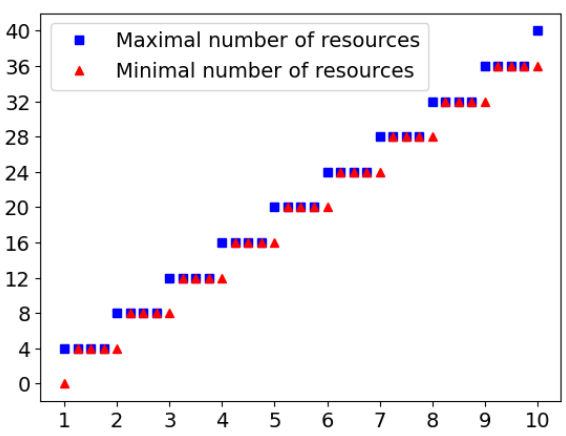
$$D\left(D^{\mathsf{A}}, n, c_0\right) = D^{\mathsf{A}}, \text{ if } D^{\mathsf{A}} \leq n \cdot \left(\left\lfloor \frac{1}{c_0} \right\rfloor - 1\right),$$
$$D\left(D^{\mathsf{A}}, n, c_0\right) = \min\left\{D^{\mathsf{A}}, n \cdot \left\lfloor \frac{1}{c_0} \right\rfloor\right\}, \text{ if } \left\lfloor \frac{1}{c_0} \right\rfloor < \frac{1}{c_0},$$
$$D\left(D^{\mathsf{A}}, n, c_0\right) \in \left(n \cdot \left(\frac{1}{c_0} - 1\right), \min\left\{n \cdot \frac{1}{c_0}, D^{\mathsf{A}}\right\}\right),$$

otherwise, i.e., when

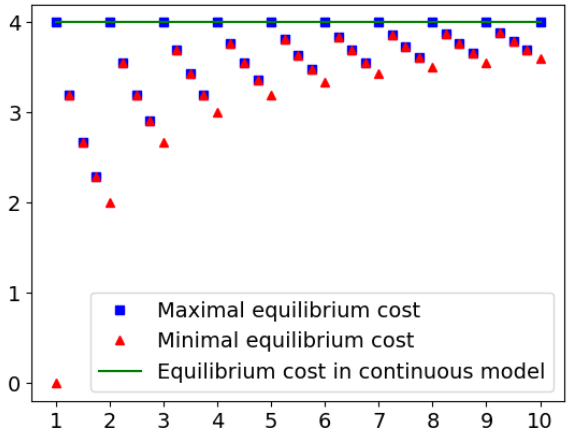
$$D^{\mathsf{A}} > n \cdot \left(\left\lfloor \frac{1}{c_0} \right\rfloor - 1\right) \text{ and } \left\lfloor \frac{1}{c_0} \right\rfloor = \frac{1}{c_0}.$$

	n	D_A	D_B	c_0^{-1}
min	3	1	1	1
max	5	50	50	10
interval	1	1	1	0.25

Figure 8: The range of parameters considered in the experiments aimed to test Hypothesis 1. The min (max) row corresponds to the lower (upper) bound on the considered parameter, and the interval row describes the difference between the two consecutive values of the parameter.



(a) The y-axis is the number of resources used in equilibrium and the x-axis is the inverse c_0^{-1} of the unit costs, c_0. The considered number of available resources of both players is $D^{\mathsf{A}} = 40$ and players compete on $n = 4$ battlefields.



(b) The y-axis is the cost carried in an equilibrium and the x-axis is the inverse c_0^{-1} of the unit costs, c_0.

Figure 9: The number of resources obtained in an equilibrium.

The preliminary results suggest the following structure. The only case when the number of resources used in equilibrium is not unique (i.e., the lower and the upper bounds are different), is when the inverse of unit costs, c_0^{-1}, is a natural number and the number of resources, D^{A} is sufficiently large. Since the set of equilibrium strategies is convex, for such cases every number of resources between the lower and upper bound is also obtained in some equilibrium strategy. In order to test Hypothesis 1, we ran experiments operating on the inverse of the unit costs as a parameter. Figure 8 presents the considered range of parameters. The results fit Hypothesis 1. Figure 9(a) illustrates the maximal and minimal number of resources used by a player in equilibrium as a function of the inverse of the unit costs, c_0^{-1}, for $c_0^{-1} \in \{1, \frac{5}{4}, \frac{6}{4}, \ldots, 10\}$. As illustrated in Figure 9(a), the number of resources obtained in equilibrium increases when the unit cost decreases. The intuition behind it is that the same total carried cost allows players to obtain more resources as the unit cost decreases.

We look at the total costs carried by a player in equilibrium and check the idea of con-

vergence to the continuous model. Figure 9(b) presents the maximal and minimal equilibrium expenditure for the same experimental data that are presented in Figure 9(a). Figure 9(b) also includes a constant line presenting the unique equilibrium expenditure of the corresponding continuous model of Roberson and Kvasov [2010].

We move on to report on the running times of our implementation of Algorithm 4. The BDDHS linear program used in Algorithm 4 was implemented as presented by Behnezhad et al. [2023]. Our implementation was done in $C{+}{+}$, using the state-of-the-art *Gurobi* [Gurobi Optimization, LLC, 2022] LP-optimizer. We investigated the program running time for two settings. In the first setting, we considered linear obtainment costs with unit costs of a resource equal to 0.05 and no assignment costs. Hence, the payoff $\hat{\pi}$ to player P is

$$\hat{\pi}^P(\boldsymbol{s}^P, \boldsymbol{s}^{-P}) = \sum_{i=1}^{n} \text{sign}(s_i^P - s_i^{-P}) - \frac{1}{20} \cdot \sum_{i=1}^{n} s_i^P.$$

In the second setting, we considered squared assignment costs with a coefficient of 0.01 and no obtainment costs. Hence, the payoff $\bar{\pi}$ to player P is

$$\bar{\pi}^P(\boldsymbol{s}^P, \boldsymbol{s}^{-P}) = \sum_{i=1}^{n} \text{sign}(s_i^P - s_i^{-P}) - \frac{1}{100} \cdot \sum_{i=1}^{n} (s_i^P)^2.$$

Figure 10 presents the obtained running times. The code was ran on a machine with an Apple M3 Pro processor unit and an 18-GB memory.

n	D^{A}	D^{B}	Time ($\hat{\pi}$)	Time ($\bar{\pi}$)
10	30	30	0.534s	0.486s
10	30	40	1.159s	1.041s
10	30	50	2.114s	1.629s
10	40	40	1.913s	1.462s
10	40	50	4.027s	2.360s
10	50	50	3.964s	4.110s
20	30	30	1.296s	1.615s
20	30	40	4.742s	3.680s
20	30	50	8.070s	6.739s
20	40	40	5.626s	4.711s
20	40	50	8.294s	9.511s
20	50	50	10.287s	10.968s
30	30	30	3.354s	2.180s
30	30	40	7.881s	6.656s
30	30	50	15.177s	16.376s
30	40	40	9.117s	6.828s
30	40	50	13.832s	18.069s
30	50	50	15.182s	19.988s

Figure 10: The averages of 10 running times of our implementation of Algorithm 4. The fourth column ($\hat{\pi}$) presents the running times for models with linear obtainment costs, while the fifth column ($\bar{\pi}$) presents the running times for models with squared assignment costs.

5.6 Conclusions

In this chapter, we studied the Colonel Blotto game with costs. We demonstrated that the equilibrium strategies of this generalized game can be computed in polynomial time with respect to the game parameters. Furthermore, we provided a comprehensive method for computing these equilibrium strategies. The significance of our contribution is twofold. Firstly, the descriptiveness of the studied model, which allows for capturing multiple real-life scenarios, makes our method for computing equilibrium strategies an good candidate for integration into decision-support systems. Secondly, the structured nature of the experimental results offers a promising avenue for testing hypotheses regarding the properties of equilibria, which may aid in obtaining further theoretical insights into the model.

Chapter 6

Concluding Remarks

In this thesis, we presented our study of four succinctly representable games. Although all the analyzed models can be interpreted as *conflicts with multiple battlefields*, they differ substantially in both their strategic structure and computational properties.

A natural benchmark for comparison of our results is the two-player constant-sum Colonel Blotto game. In this model, each player distributes a fixed number of resources across multiple battlefields, and the payoffs to players are bilinear, i.e., they are given by the sum of utilities across individual battlefields. As showed by Ahmadinejad et al. [2019], Behnezhad et al. [2023], Nash equilibria of the Colonel Blotto game can be computed in polynomial time with respect to the model parameters.

The four models studied in this thesis can be viewed as systematic departures from this benchmark. The conflicts with non-bilinear aggregation functions preserve the constant-sum property but violate bilinearity. The Colonel Blotto game with Costs and the Colonel Blotto game with Flexible Tie-breaking depart from the classical model by breaking the zero-sum property. Moreover, the global costs in the Colonel Blotto game with Costs also break bilinearity, as the payoff to a player is no longer the sum of single-battlefield utilities. Finally, the attack and interception game on networks departs from the two-player Blotto framework in two fundamental ways: it involves multiple defenders, resulting in a multiplayer interaction, and it replaces the independent battlefields with a network structure that defines feasible strategies and payoffs.

Our results characterize how these structural departures affect the feasibility of equilibrium computation.

In conflicts with non-bilinear aggregation functions, we proved that computing a best response to a mixed strategy cannot be done in polynomial time unless $\mathsf{P} = \mathsf{NP}$, demonstrating that the loss of bilinearity introduces a fundamental computational barrier. Despite this hardness, we developed an efficient algorithm for computing equilibria. The proposed method combines a symmetrization procedure, which reduces the effective strategy space exponentially, with the clash-matrix algorithm for efficient payoff computation and a Double Oracle framework supported by a monotonicity-based heuristic.

In the Colonel Blotto game with Costs, we showed that every instance is strategically equiv-

alent to a zero-sum Colonel Blotto game with one additional battlefield. This equivalence result allows the use of polynomial-time algorithms for the classical Blotto model to compute equilibria in the cost-augmented variant.

In the Colonel Blotto game with Flexible Tie-breaking, we constructed and analyzed the set of Nash equilibria, showing that it is extensive and highly non-unique. In particular, every pure strategy that does not assign excessive resources to individual battlefields is played with positive probability in some mixed-strategy equilibrium. Moreover, we showed that for small values of the efficiency parameter, no pure strategy is dominated by another pure strategy. Finally, we proved that the only equilibrium that remains robust to tie-breaking perturbations is the one with uniform marginals.

In attack and interception games on networks, we showed that equilibrium computation remains tractable even though the attacker's pure strategies correspond to all simple paths in the network, a set that can grow exponentially with the number of nodes. By exploiting the network topology, we proved that the set of nodes targeted in equilibrium is unique and can be determined in polynomial time for a broad class of cost functions. This result demonstrates that when the structure of feasible strategies is defined by a network, its topological constraints can be leveraged to achieve efficient equilibrium computation in an otherwise exponentially large strategy space.

Bibliography

A. Ahmadinejad, S. Dehghani, M. Hajiaghayi, B. Lucier, H. Mahini, and S. Seddighin. From duels to battlefields: Computing equilibria of Blotto and other games. *Mathematics of Operations Research*, 44(4):1304–1325, 2019.

A. Arad and A. Rubinstein. Multi-dimensional iterative reasoning in action: The case of the colonel blotto game. *Journal of Economic Behavior & Organization*, 84:571–585, 11 2012. doi: 10.1016/j.jebo.2012.09.004.

J. Avrahami, Y. Kareev, P. M. Todd, and B. Silverman. Allocation of resources in asymmetric competitions: How do the weak maintain a chance of winning? *Journal of Economic Psychology*, 42:161–174, 2014.

E. Beale and G. Heselden. An approximate method of solving Blotto games. *Naval Research Logistics*, 9(2):65–79, 1962.

S. Behnezhad, A. Blum, M. Derakhshan, M. Hajiaghayi, M. Mahdian, C. Papadimitriou, R. Rivest, S. Seddighin, and P. Stark. From battlefields to elections: Winning strategies of Blotto and auditing games. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 2291–2310, 2018a.

S. Behnezhad, A. Blum, M. Derakhshan, MT. HajiAghayi, M. Mahdian, C. H. Papadimitriou, R. L. Rivest, S. Seddighin, and P. B. Stark. *From Battlefields to Elections: Winning Strategies of Blotto and Auditing Games*, pages 2291–2310. 2018b. doi: 10.1137/1.9781611975031.148.

S. Behnezhad, S. Dehghani, M. Derakhshan, M. Hajiaghayi, and S. Seddighin. Fast and simple solutions of blotto games. *Operations Research*, 71(2):506–516, 2023. doi: 10.1287/opre.2022.2261.

R. M. Bell and T. M. Cover. Competitive optimality of logarithmic investment. *Mathematics of Operations Research*, 5(2):161–166, 1980. ISSN 0364765X, 15265471. URL http://www.jstor.org/stable/3689147.

D. W. Blackett. Some blotto games. *Naval Research Logistics Quarterly*, 1(1):55–60, March 1954. doi: 10.1002/nav.3800010109. URL https://ideas.repec.org/a/wly/navlog/v1y1954i1p55-60.html.

F. Bloch, K. Chatterjee, and B. Dutta. Attack and interception in networks. *Theoretical Economics*, 18(4):1511–1546, 2023. doi: https://doi.org/10.3982/TE5122. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/TE5122.

M. Blonski. Characterization of pure strategy equilibria in finite anonymous games. *Journal of Mathematical Economics*, 34(2):225–233, 2000. ISSN 0304-4068. doi: https://doi.org/10.1016/S0304-4068(99)00008-7. URL https://www.sciencedirect.com/science/article/pii/S0304406899000087.

É. Borel. La Théorie du Jeu et les Équations Intégrales à Noyau Symétrique. *Comptes Rendus de l'Académie des Sciences*, 173:1304–1308, 1921. Translated by Savage L. J., The theory of play and integral equations with skew symmetric kernels. Econometrica 21 (1953) 97–100.

É. Borel and J. Ville. *Applications de la théorie des probabilités aux jeux de hasard*. J. Gabay, Paris, 1991 edition, 1938.

M. Carvalho, A. Lodi, and J. Pedroso. Computing equilibria for integer programming games. *European Journal of Operational Research*, 303(3):1057–1070, 2022. ISSN 0377-2217.

P. H. Chia and J. Chuang. Colonel Blotto in the phishing war. In *Proceedings of the Second International Conference on Decision and Game Theory for Security*, GameSec'11, pages 201–218, 2011.

H. Chin, T. Parthasarathy, and T. Raghavan. Structure of equilibria in n-person non-cooperative games. *International Journal of Game Theory - INT J GAME THEORY*, 3:1–19, 03 1974. doi: 10.1007/BF01766215.

S. Chowdhury, D. Kovenock, and R. Sheremeta. An experimental investigation of colonel blotto games. *Economic Theory*, 52, 01 2009. doi: 10.1007/s00199-011-0670-2.

R. A. Collado and D. Papp. Network interdiction – models , applications , unexplored directions. *Tech. Rep. RRR 4-2012, Rutgers University*, 2010.

G. Dantzig. A proof of the equivalence of the programming problem and the game problem. In T. Koopmans, editor, *Activity Analysis of Production and Allocation*, 1951. URL https://link.springer.com/chapter/10.1007/978-1-349-58802-2_13.

C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a nash equilibrium. *Commun. ACM*, 52(2):89–97, February 2009. ISSN 0001-0782. doi: 10.1145/1461928.1461951. URL https://doi.org/10.1145/1461928.1461951.

S. Dehghani, H. Saleh, S. Seddighin, and S-H. Teng. Computational analyses of the electoral college: Campaigning is hard but approximately manageable. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6):5294–5302, May 2021. doi: 10.1609/aaai.v35i6.16668. URL https://ojs.aaai.org/index.php/AAAI/article/view/16668.

M. Dell. Trafficking networks and the mexican drug war. *American Economic Review*, 105(6): 1738–79, June 2015.

M. Dziubiński. The spectrum of equilibria for the Colonel Blotto and the Colonel Lotto games. In V. Bilò and Michele Flammini, editors, *Algorithmic Game Theory*, pages 292–306, Cham, 2017. Springer International Publishing. ISBN 978-3-319-66700-3.

B. C. Eaves. Polymatrix games with joint constraints. *SIAM J. Appl. Math.*, 24(3):418–423, May 1973. ISSN 0036-1399. doi: 10.1137/0124043. URL https://doi.org/10.1137/0124043.

D. Eisen and M. Le Mat. Defense models, XIII: BLOTTO, a constrained matrix game solver, 1968.

C. Ewerhart. A "fractal" solution to the chopstick auction. *Economic Theory*, 74(4):1025–1041, 2022.

C. Ewerhart and S. Kaźmierowski. An Equilibrium Analysis of the Arad-Rubinstein Game. *Journal of Economic Behavior & Organization*, 226:106696, 2024. ISSN 0167-2681. doi: https://doi.org/10.1016/j.jebo.2024.106696. URL https://www.sciencedirect.com/science/article/pii/S0167268124003020.

R. Ferber. Book review: Games and decisions: Introduction and critical survey. *Journal of Marketing*, 23(4):469–470, 1959. doi: 10.1177/002224295902300423. URL https://doi.org/10.1177/002224295902300423.

D. Fotakis, S. C. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. G. Spirakis. The structure and complexity of nash equilibria for a selfish routing game. In *Proceedings of the 29th International Colloquium on Automata, Languages and Programming*, ICALP '02, page 123–134, Berlin, Heidelberg, 2002. Springer-Verlag. ISBN 3540438645.

D. R. Fulkerson and G. C. Harding. Maximizing the minimum source-sink path subject to a budget constraint. *Math. Program.*, 13(1):116–118, dec 1977. ISSN 0025-5610.

P. M. Ghare, D. C. Montgomery, and W. C. Turner. Optimal interdiction policy for a flow network. *Naval Research Logistics Quarterly*, 18(1):37–45, 1971.

B. Golden. A problem in network interdiction. *Naval Research Logistics Quarterly*, 25(4): 711–713, 1978.

O. A. Gross and R. A. Wagner. A continuous colonel blotto game. 1950. URL https://api.semanticscholar.org/CorpusID:107940045.

S. Guan, J. Wang, C. Jiang, Z. Han, Y. Ren, and A. Benslimane. Colonel blotto game aided attack-defense analysis in real-world networks. In *2018 IEEE Global Communications Conference (GLOBECOM)*, pages 1–6, 2018. doi: 10.1109/GLOCOM.2018.8647886.

S. Guan, J. Wang, H. Yao, C. Jiang, Z. Han, and Y. Ren. Colonel blotto games in network systems: Models, strategies, and applications. *IEEE Transactions on Network Science and Engineering*, 7:1–1, 03 2019. doi: 10.1109/TNSE.2019.2904530.

H. Gupta. Partitions–a survey. *Journal of Res. of Nat. Bur. Standards-B Math. Sciences B*, 74: 1–29, 1970.

Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2022.

G. H . Hardy and S. Ramanujan. Asymptotic formulae in combinatory analysis. *Proc. London Math. Soc*, 17(2):75–115, 1918.

S. Hart. Discrete Colonel Blotto and General Lotto games. *International Journal of Game Theory*, 36:441–460, 02 2008.

J. T. Howson. Equilibria of polymatrix games. *Management Science*, 18(5):312–318, 1972. ISSN 00251909, 15265501. URL http://www.jstor.org/stable/2634798.

Q. Jiao and J. Xu. Competition for networked agents in the lottery blotto game. *Economics Letters*, 197:109644, 2020. ISSN 0165-1765. doi: https://doi.org/10.1016/j.econlet.2020.109644. URL https://www.sciencedirect.com/science/article/pii/S0165176520304043.

S. Kakade, M. Kearns, J. Langford, and L. Ortiz. Correlated equilibria in graphical games. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, EC '03, page 42–47, New York, NY, USA, 2003. Association for Computing Machinery. ISBN 158113679X. doi: 10.1145/779928.779934. URL https://doi.org/10.1145/779928.779934.

S. Kaźmierowski. Equilibria of the Colonel Blotto Games with Costs. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(13):13969–13976, Apr. 2025. doi: 10.1609/aaai.v39i13.33528. URL https://ojs.aaai.org/index.php/AAAI/article/view/33528.

S. Kaźmierowski and M. Dziubiński. Computation of Nash Equilibria of Attack and Defense Games on Networks. In A. Deligkas and A. Filos-Ratsikas, editors, *Algorithmic Game Theory*, pages 3–21, Cham, 2023. Springer Nature Switzerland. ISBN 978-3-031-43254-5.

S. Kaźmierowski and M. Dziubiński. Efficient Method for Finding Optimal Strategies in Chopstick Auctions with Uniform Objects Values. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '24, pages 992–1000, Richland, SC, 2024. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9798400704864.

S. Kaźmierowski. Seeking ϵ-equilibrias in chopsticks auctions using multiplicative weights update algorithm. Master's thesis, University of Warsaw, 2021.

M. Kearns, M. L. Littman, and S. Singh. Graphical models for game theory. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, UAI'01, page 253–260, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. ISBN 1558608001.

E. Kohlberg and J-F. Mertens. On the strategic stability of equilibria. *Econometrica*, 54(5): 1003–1037, 1986.

D. Kovenock and B. Roberson. Coalitional colonel blotto games with application to the economics of alliances. *Journal of Public Economic Theory*, 14(4):653–676, 2012. doi: https://doi.org/10.1111/j.1467-9779.2012.01556.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9779.2012.01556.x.

D. Kovenock and B. Roberson. Generalizations of the General Lotto and Colonel Blotto games. *Economic Theory*, 71(3):997–1032, 2021.

J-F. Laslier and N. Picard. Distributive politics and electoral competition. *Journal of Economic Theory*, 103(1):106–130, 2002. ISSN 0022-0531. doi: https://doi.org/10.1006/jeth.2000.2775. URL https://www.sciencedirect.com/science/article/pii/S0022053100927753.

K. Leyton-Brown and M. Tennenholtz. Local-effect games. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, IJCAI'03, page 772–777, San Francisco, CA, USA, 2003. Morgan Kaufmann Publishers Inc.

J. Litwin. Equilibria of chopstick auctions with discrete resources. Master's thesis, University of Warsaw, 2020.

H. B. McMahan, G. J. Gordon, and A. Blum. Planning in the presence of cost functions controlled by an adversary. In T. Fawcett and N. Mishra, editors, *ICML*, pages 536–543. AAAI Press, 2003. ISBN 1-57735-189-4.

A. W. McMasters and T. M. Mustin. Optimal interdiction of a supply network. *Naval Research Logistics Quarterly*, 17(3):261–268, 1970.

M. Mirzaei, S. M. J. Mirzapour Al-e-hashem, and M. Akbarpour Shirazi. A maximum-flow network interdiction problem in an uncertain environment under information asymmetry condition: Application to smuggling goods. *Computers I& Industrial Engineering*, 162:107708, 2021. ISSN 0360-8352.

D. P. Morton, F. Pan, and K. J. Saeger. Models for nuclear smuggling interdiction. *IIE Transactions*, 39(1):3–14, 2007.

H. Moulin and J-P. Vial. Strategically zero-sum games: the class of games whose completely mixed equilibria connot be improved upon. LIDAM Reprints CORE 359, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE), January 1978. URL https://ideas.repec.org/p/cor/louvrp/359.html.

M. Munger, J. Merolla, and M. Tofias. In play: A commentary on strategies in the 2004 u.s. presidential election. *Public Choice*, 123:19–37, 02 2005. doi: 10.1007/s11127-005-3209-x.

J. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951. ISSN 0003486X, 19398980. URL http://www.jstor.org/stable/1969529.

P. Naumov and B. Nicholls. On interchangeability of nash equilibria in multi-player strategic games. *Synthese*, 190:57–78, 2013. ISSN 00397857, 15730964. URL http://www.jstor.org/stable/24019896.

M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

C. H. Papadimitriou and T. Roughgarden. Computing correlated equilibria in multi-player games. *J. ACM*, 55(3), August 2008. ISSN 0004-5411. doi: 10.1145/1379759.1379762. URL https://doi.org/10.1145/1379759.1379762.

A. Penn. A generalized lagrange-multiplier method for constrained matrix games. *Operations Research*, 19(4):933–945, 1971.

T. Raghavan. Zero-sum two-person games. In R. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 2, chapter 20, pages 735–768. Elsevier, 1 edition, 1994.

B. Roberson. The Colonel Blotto game. *Economic Theory*, 29(1):1–24, 2006.

B. Roberson and D. Kvasov. The Non-Constant-Sum Colonel Blotto Game. Purdue University Economics Working Papers 1252, Purdue University, Department of Economics, August 2010. URL https://ideas.repec.org/p/pur/prukra/1252.html.

R. W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *Int. J. Game Theory*, 2(1):65–67, December 1973. ISSN 0020-7276. doi: 10.1007/BF01737559. URL https://doi.org/10.1007/BF01737559.

J. C. Smith and Y. Song. A survey of network interdiction models and algorithms. *European Journal of Operational Research*, 283(3):797–811, 2020. ISSN 0377-2217.

B. Szentes and R. Rosenthal. Beyond chopsticks: Symmetric equilibria in majority auction games. *Games and Economic Behavior*, 45(2):278–295, 2003a. ISSN 0899-8256. Special Issue in Honor of Robert Rosenthal.

B. Szentes and R. Rosenthal. Three-object two-bidder simultaneous auctions: Chopsticks and tetrahedra. *Games and Economic Behavior*, 44:114–133, 07 2003b.

J. W. Tukey. A problem of strategy. *Econometrica*, 17(1):73, 1949.

J. von Neumann. A certain zero-sum two-person game equivalent to the optimal assignment problem. In *Contributions to the Theory of Games (AM-28), Volume II*, pages 5–12. Princeton University Press, 1953.

D. Q. Vu and P. Loiseau. Colonel blotto games with favoritism: Competitions with pre-allocations and asymmetric effectiveness. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, EC '21, page 862–863, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450385541. doi: 10.1145/3465456.3467527. URL https://doi.org/10.1145/3465456.3467527.

A. Washburn. OR Forum—Blotto Politics. *Operations Research*, 61(3):532–543, June 2013. doi: 10.1287/opre.1120.1142. URL https://ideas.repec.org/a/inm/oropre/v61y2013i3p532-543.html.

A. Washburn and K. Wood. Two-person zero-sum games for network interdiction. *Operations Research*, 43(2):243–251, 1995.

R. Weber and M. Shubik. Systems defense games: Colonel blotto, command and control. *Naval Research Logistics Quarterly*, 28:281–287, 1981.

H. Xu. The mysteries of security games: Equilibrium computation becomes combinatorial algorithm design. *Proceedings of the 2016 ACM Conference on Economics and Computation*, 2016. URL https://api.semanticscholar.org/CorpusID:707610.