Hardness and Totality of Problems Related to Nash Equilibrium

by

Koosha Samieefar

B.Sc., Amirkabir University of Technology, 2016
M.Sc., Amirkabir University of Technology, 2018

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Computer Science

We acknowledge and respect the Lək̓ʷəŋən (Songhees and Xʷsepsəm/Esquimalt) Peoples on whose territory the university stands, and the Lək̓ʷəŋən and W̱SÁNEĆ Peoples whose historical relationships with the land continue to this day.

Hardness and Totality of Problems Related to Nash Equilibrium

by

Koosha Samieefar

B.Sc., Amirkabir University of Technology, 2016
M.Sc., Amirkabir University of Technology, 2018

Supervisory Committee

Dr. Bruce M. Kapron, Supervisor
(Department of Computer Science)

Dr. Sajin Koroth, Departmental Member
(Department of Computer Science)

Dr. David Scoones, Outside Member
(Department of Economics)

ABSTRACT

We investigate the computational complexity of various equilibrium problems and their applications—a central topic in algorithmic game theory. While the existence of solutions for classical equilibrium notions, such as generalized Nash equilibria and multi-leader–follower games with convex constraints, is well established via foundational fixed-point theorems, the complexity of computing such equilibria has remained largely unresolved, especially in settings involving uncertainty.

In particular, we study the computational complexity of approximating both Nash and generalized Nash equilibria in more complex formats of games under general convexity assumptions, encompassing broad settings with both differentiable and non-differentiable utility functions. To be specific, building on and extending the computational framework for Kakutani's fixed-point theorem, we establish PPAD-completeness results for several classes of variational inequality problems and extend the results to multi-leader–follower games and resilient equilibria, as well as some other relevant definitions in games with or without uncertainty. Moreover, we propose a family of natural yet specific non-convex constraints—termed strategic constraints—and investigate their computational complexity and a range of potential applications, such as strategic fair division.

Contents

List of Tables

List of Figures

ACKNOWLEDGEMENTS

I would like to thank:

My Supervisor for being always helpful and generous with his time and energy, and helping me to discover the beauty of theoretical subjects in computer science. I'm so grateful to have had the opportunity to work under his supervision.

University of Victoria I will forever be grateful for the opportunity that was given to me to study in a wonderful environment.

The Supervisory committee for accepting the supervisory role and devoting their precious time to reviewing my research.

And my family for being supportive in my entire life unconditionally.

No one who achieves success does so without acknowledging the help of others.
Alfred North Whitehead

DEDICATION

To my supervisor, whose guidance, encouragement, and support shaped not only this research but also my growth as a human being. Your patience and the opportunities you gave me helped me to become a better person. Your encouragement gave me the courage to continue, and your belief in me carried me through moments of doubt.

To my family and my partner, whose strength gave me hope when times were difficult. Despite the distance, the situation in my country, and personal hardships, you were always the reason I continued.

This thesis reflects not only academic effort but also the lessons of life I gained along the way.

Chapter 1

Introduction

1.1 Game Theory, Computational Complexity and Equilibrium Problems

Game theory is the mathematical study of strategic decision-making in competitive or cooperative situations, where an agent's choices depend on the choices of others. It includes various game formats and extensions, each offering unique insights into interactions. These formats emerged to address the need for modeling complex strategic interactions, whether competitive or cooperative, across diverse fields such as economics, political science, engineering, biology, and artificial intelligence. A central pursuit within game theory is the study of the computational complexity of equilibrium problems, aiming to understand the algorithmic feasibility of predicting or computing stable outcomes in strategic settings. In this area, we investigate whether equilibria—such as Nash equilibria and their restrictions or relaxations—can be computed efficiently, and what complexity-theoretic barriers arise across different game formats.

1.1.1 Equilibrium Problems in Game Theory

The notion of equilibrium has long been central to the study of game theory, economic behavior, and multi-agent systems. Among various equilibrium concepts, the Nash equilibrium remains one of the most studied and celebrated in the non-cooperative setting. Introduced in the seminal work of Nash, the equilibrium captures the idea of mutual best responses in a non-cooperative setting: no player can improve their utility by unilaterally deviating from their strategy [144, 145]. Moreover, Nash [144, 145] proved that for all well-defined finite games, there exists at least one Nash equilibrium. The computational complexity of finding a Nash equilibrium has been extensively studied [45, 59, 70, 138, 165], and shown to lie within the class TFNP of total NP search problems. The key result of [45] shows that computing a mixed Nash equilibrium for two-player games is PPAD-complete, meaning it is among the hardest problems in PPAD, a subclass of TFNP that includes problems such as Sperner's Lemma and Brouwer's fixed-point theorem. In the case that we have more than two players, the existence of a rational Nash equilibrium (even without extra constraints) is not guaranteed. It is known that there are 3-player games that can only produce irrational solutions. Discussion of the complexity in such cases requires the use of classes such as FIXP [70] and for Nash equilibrium with additional conditions $\exists\mathbb{R}$ [31]. To tackle this problem, the notion of approximation where a rational solution is guaranteed to exist was proposed [59]. A variety of approximation algorithms are known for approximating a Nash equilibrium [49, 56, 60, 130].

Beyond the basic Nash equilibrium problem, we can consider equilibria with additional constraints. In this direction, there are two well-known approaches. The goal of the first approach is to find a Nash equilibrium that has certain properties (e.g., achieving a specified social welfare) [52, 89, 135]. In this setting, unlike the situation with unconstrained Nash equilibrium, a solution is not guaranteed [52, 89, 135]. For example, determining the existence of an equilibrium achieving some standard fairness notions, such as utilitarian or egalitarian social welfare, is NP-hard [52]. In the second approach, the goal is to find a generalized solution—one in which no player would deviate, assuming all players choose strategies that satisfy the given constraint [4, 62, 72, 161]. Unlike in Nash equilibrium, where the players do not have an incentive to deviate to all alternative strategies, in a generalized equilibrium, the players do not have an incentive to deviate to only those that lie within the constraint imposed by other players. This extension was introduced by Debreu, who also proved the existence of a generalized equilibrium under certain assumptions [62]. Many researchers from different areas have worked on this notion, which in turn explains why it has developed with a number of names in the literature including pseudo-games, social equilibrium, and common (coupled) constrained equilibrium [4, 72, 73, 170, 185]. Rosen also introduced a restricted formulation, known as "concave games", in which the constraints imposed on the players are the same, and provided existence and uniqueness results under some assumptions, such as convexity of the constrained strategies.

Multi-leader-follower games are a class of games in which multiple agents, referred to as *leaders* and *followers*, interact strategically to achieve their respective objectives. Leaders and followers often have conflicting objectives or interests, in particular aiming to maximize their own profit or minimize their loss. This setup resembles a bi-level program, where the leaders engage in competitive, non-cooperative Nash games at the upper level, making their decisions while considering the followers' responses. After the leaders have made their choices, the followers then engage in a parametric, non-cooperative game at the lower level, with the strategies of the leaders treated as external parameters. The concept of multi-leader-follower games has a variety of applications that arise from situations where there are multiple oligopoly firms operating in the market [46, 105, 106, 128, 150]. Oligopoly markets are markets dominated by a small number of suppliers. The simplest form of the multi-leader-follower game is a Stackelberg game [18, 172, 180] in which one leader and multiple followers react to the leader's strategies. These games have applications in various fields, including economics, engineering, and multi-agent systems [47, 104, 149]. While the multi-leader-follower problem provides a sound mathematical framework with a well-defined solution concept and practical applications, its elevated level of complexity and technical intricacies make it computationally challenging or intractable. Specifically, it resembles an equilibrium problem in a more complex form of Debreu-Rosen style games, requiring each leader to solve a non-convex mathematical program with equilibrium constraints [20, 150, 179]. This formulation faces two significant issues: a potential absence of an equilibrium solution due to non-convexity, and computational intractability. In response to these challenges, [106, 150] proposed a careful analysis and selection of "remedial models" aimed at deriving sensible equilibrium solutions.

Strong Nash equilibrium *(SNE)* is another concept in game theory that restricts the traditional notion of Nash Equilibrium. In SNE, not only does each player have no incentive to deviate from their strategy given the strategies of others unilaterally, but this equilibrium also withstands deviations by any possible coalition of players [10]. In other words, even if a group of players collaboratively tries to deviate from the equilibrium, they cannot collectively gain from such deviations. This concept had faced criticism for being excessively "strong", particularly in environments permitting unlimited private communication (for more information, see [29]). Furthermore, a strong Nash equilibrium is weakly Pareto-efficient (no one can be made better off without making someone else worse off), a fact which may be verified by considering a deviation of the "grand coalition" of all players. In conclusion, the existence of such an equilibrium is unlikely (even if mixed strategies are allowed) [29], except under strict conditions and assumptions (extensively discussed in [6, 111, 146, 156]). In contrast, for *coalition-proof Nash equilibrium (CPNE)*, limitations on private communication are imposed [29]. CPNE emphasizes stability against coordinated actions by any subset of players, regardless of the coalition's size while SNE goes further by requiring that not only should deviations by any coalition be unsuccessful, but any subset of players (including the entire group) should also be unable to deviate successfully. In contrast to the NP-completeness of deciding whether an SNE exists [52, 83, 85], the problem of finding an SNE is in smoothed polynomial time [35, 36, 84][1]. *Resilient Nash equilibrium*, introduced in [1], is a more restricted fundamental notion in applications such as secret sharing and multi-party computation [1, 2, 95, 96, 125]. We may also restrict ourselves to coalitions of size at most t. This limitation appears reasonable in practice, as forming and coordinating large coalitions can be challenging. This fact in turn motivated the definition of *t-strong Nash equilibrium* and *t-resilient Nash equilibrium* [1, 6]. Informally, no member of a coalition of size up to t can do better, even if the whole coalition defects. Apart from the technical interest of these variations, numerous papers explore their practical applications [1, 2, 30, 81].

Games with uncertain data, also known as imperfect information games, add another layer to traditional games which is unpredictability, as players often lack complete information about their opponents' strategies or the game's state. This uncertainty necessitates strategic decision-making under ambiguity. This uncertainty can arise from various sources, such as hidden information, randomness, or incomplete knowledge of the game's mechanics. Assuming the Bayesian hypothesis on probability distributions, Harsanyi [99, 100, 101] investigated games featuring incomplete information, where players lack full knowledge of key game parameters. By assuming that all players share common knowledge of these probability distributions, the game can ultimately be re-conceptualized as one with complete information, referred to as the Bayes equivalent of the original game. Furthermore, stochastic optimization techniques are shown to be suitable for addressing Stackelberg games involving uncertain data, enabling probabilistic decision-making. In contrast to traditional games, these models incorporate randomness, allowing decision-makers to account for uncertain

[1]These results concern the problem of finding (mixed strategy and exact solution of) SNE in bi-matrix games, which is a special case. Specifically, they show that if there is a mixed strategy, the payoffs restricted to the actions in the support must satisfy strict geometric conditions.

parameters such as resource availability. Detailed insights into stochastic Stackelberg and multi-leader-follower games and their applications in numerous fields are available in [33, 50, 66, 129, 171, 132].

In addition to models involving strategy distributions conditioned on player types, recent years have seen a rise in the adoption of distribution-free models based on worst-case scenarios [3, 50, 107, 108, 113, 123]. In these models, each player's decision-making process is guided by the principles of robust optimization (see also [21, 22, 23]). Robust optimization operates under the assumption that uncertain data lie within an uncertainty set, seeking solutions that account for the worst-case outcomes in terms of objective function values and/or constraints. For example, Hu et al. [107] use the robust optimization technique for multi-leader-follower games under uncertainty assumptions. Additionally, [189] explores the problem of multi-agent reinforcement learning (MARL) in scenarios where agents have imperfect knowledge of the model and formulates it as a robust Markov game to find policies that are robust to model uncertainty. A recent paper also applies theoretical developments to explore the impact of uncertainty on robust game solutions [55][2]. *Robust Nash equilibrium* is a solution concept that accounts for uncertainty in the parameters of a game where each player's strategy remains optimal even when faced with uncertainty about the strategies of other players, or about the underlying conditions of the game. This concept ensures the stability of strategies even in the presence of unpredictability or incomplete information. Hayashi et al. explored the problem of robust Nash equilibria in bi-matrix games and presented some existence results under specific assumptions on the uncertainty sets [103]. Aghassi and Bertsimas [3] studied robust Nash equilibria in n-player games with bounded polyhedral uncertainty sets, where each player solves a linear programming problem, and also proposed a method for finding such equilibria. The discussed models of [3, 103] focus on linear objective functions in each player's problem. In contrast, more recent work by [147] investigates a more general NE with uncertain data, where each player solves an optimization problem with a nonlinear objective function. Moreover, under mild assumptions on the uncertainty sets, the authors presented results regarding the existence, and uniqueness of robust Nash equilibria in this setting. Hu and Fukushima [107] further extended their work in [106] under the uncertainty assumption by using the robust optimization technique

[2]Specifically, the paper investigates the sensitivity of equilibria to different levels of uncertainty and establishes a relationship between the robust optimization approach and approximate Nash equilibrium.

for multi-leader-follower games with uncertainty. Specifically, they introduced the concept *robust L/F-equilibrium* and established its existence and uniqueness for a specific class of multi-leader-follower games with a special structure.

1.2 Our Contribution and Organization

This thesis investigates the computational complexity of a variety of equilibrium problems in game theory, particularly under structural constraints. We focus on both classical and generalized notions of equilibrium and study their algorithmic properties using tools from convex analysis, variational inequalities, and fixed-point theorems. Our contributions span both hardness results (NP-completeness and PPAD-completeness) and algorithmic techniques (including metaheuristic approaches). The thesis is organized as follows:

1.2.1 Chapter 2: Preliminary Definitions and Existence Results

The second chapter includes preliminary definitions and foundational concepts from the literature, covering all the above-mentioned equilibrium problems—including Nash equilibrium, generalized Nash equilibrium, and strong and resilient Nash variants—as well as key analytical tools such as variational inequalities and fixed-point theorems that are central to our research.

1.2.2 Chapter 3: Convex Optimization, Variational Inequalities

Recently, the computational complexity of general convex optimization problems has been getting noticeable attention [5, 28, 76, 79, 120, 187]. The existence of a generalized Nash equilibrium in Debreu-Rosen style games under convexity assumptions is proven by Kakutani's fixed-point theorem [62, 161]. Problems for which the existence of a solution follows from Kakutani's fixed-point theorem have not been extensively explored from a complexity standpoint until recently. Specifically, considering the application of generalized equilibrium problems defined by Debreu [62] and Rosen [161] in economics and other areas (see [72, 73]), the complexity of these problems has not received much attention until recently [28, 78, 79, 120, 122, 152].

In Chapter 3, we focus on the computational complexity of variational inequality problems and establish several PPAD-completeness results using tools from computational convex geometry. These results will form the basis for our analysis in Chapter 4, where we study the complexity of a broad class of equilibrium problems. A *Variational inequality problem (VI)* involves finding a point in a closed convex set where a function satisfies a specific inequality for all points in the set. More generalized forms, such as *quasi-variational inequalities (QVIs)* and *generalized quasi-variational inequalities (GQVIs)*, have also been proposed [43, 75]. QVIs extend VIs by introducing set-valued mappings and feasible sets that depend on external variables, while GQVIs further generalize by allowing the function itself to be a correspondence, increasing complexity compared to VIs and QVIs. Commencing with the formal definition of Variational inequalities proposed by Bensoussan and Lion ([24, 25, 26, 27]), researchers have engaged in an exploration of algorithmic solutions, focusing on conditions governing convergence. Various numerical and mathematical techniques, including fixed-point methods, penalty methods, and projection methods, can be used to find solutions to variational techniques [19, 74, 88, 150, 159]. There are also research papers which focus on examining whether a solution to a problem exists [32, 43, 126, 137, 151, 186].

Variational inequality problems also find application in diverse fields, including optimization theory, economics, and engineering, serving as a modeling and solution tool for a broad spectrum of real-world problems [88, 102, 139, 150]. This is especially true in cases where the aim is to find a solution subject to a particular inequality condition that is defined by a set of functions or operators. For example, under some assumptions such as differentiability, Debreu-Rosen style games (see [97, 150]) can be expressed as variational inequalities, aiding their analysis through variational techniques, and offering a unified framework for various non-cooperative games[3].

The study of the computational aspects of these variational inequalities has begun to advance, with recent work (including ours) contributing to a growing body of results [28, 120, 119]. Starting from [43], an existence result for the GQVI was proved using the Eilenberg-Montgomery fixed point theorem (see [68]) and Kakutani's fixed point under different conditions. In [153], a problem known as Kakutani was introduced, with a brief overview of its inclusion in the complexity class PPAD. The primary challenge in formulating Kakutani as a computational problem lies

[3]Multiple other motivating reasons exist for using variational inequalities in game theory problems [181, 150]. For example, they naturally centralize the role of dynamics in understanding the behavior of interacting agents [181]. Also, several papers study properties of the Nash equilibrium in network games by using variational inequalities [136, 142, 154].

in the limitations of conventional approaches to explicitly and succinctly represent convex sets such as convex polytopes. A more suitable computational formulation of the Kakutani problem has recently been introduced by leveraging the computational convex geometry methods introduced in [93]. These methods represent convex sets using separation oracles. This formulation was subsequently applied in [152], which considers the computational complexity of finding approximate equilibrium solutions for Rosen-style games.

As discussed in [152], simpler representations of convex sets, such as using convex hulls or polytopes, are too restrictive and often not useful for capturing the practical applications of Kakutani fixed points in game theory. On the theoretical side, our goal is to establish PPAD membership for the most general forms of these equilibrium problems. A key consideration here is how convex constraints are represented, since the choice of representation can substantially affect both the applicability of fixed-point theorems and the computational formulation. While convex hull or polytope descriptions might seem simpler, using them would require replacing standard tools such as Berge's maximum theorem and Kakutani's fixed point theorem with more specialized techniques, and thus would not necessarily simplify our proofs[4]. Moreover, many practical applications—particularly in economics—naturally lead to the need for more general definitions of convex sets, such as in the Walrasian equilibrium model of [114], which are more naturally handled via oracle-based representations.

The recent paper [28] establishes PPAD-completeness of (QVIs), limited to certain constraints, using an approach different from [152] and also establishes PPAD-completeness of equilibria with limited constraints and smooth loss (utility) functions. However, we consider a more generalized version of the variational inequality problem (GQVI) and also games with loss (or utility) functions that are not necessarily differentiable, such as those represented by linear arithmetic circuits [76]. Moreover, with the notion of weak separation oracles, our approach is more general even for the sort of constraints considered in [28]. We apply the methods of [152] to establish the PPAD-completeness of a more restricted class of variational inequalities.

1.2.3 Chapter 4: Equilibrium With Convex Constraints

In Chapter 4, we focus on the computational complexity of several equilibrium problems. We extend our investigation to encompass a broader class of equilibrium prob-

[4] For more information, see the definitions of Kakutani's fixed point and Theorem E.1 in [152].

lems, including multi-leader-follower games, resilient Nash equilibria, and related notions under reasonable assumptions. Building on the remedial solutions proposed for L/F-equilibrium in multi-leader-follower games [106, 150], we propose a more general definition that unifies remedial models. We then carefully establish the PPAD-completeness of the corresponding computational formulation. This chapter also includes a computational treatment of robust equilibrium. This is a game model that extends standard notions by introducing uncertain data. We consider this model in different classes of games, including multi-leader-follower games. Doing this requires the development of more general results on variational inequalities, as well as addressing significant technical challenges needed to deal with this broader approach. Finally, we note that our approach provides an alternate route to proving the PPAD-completeness of Rosen-style games (as in [152]) and even Debreu-style games (see also [28, 79]).

1.2.4 Chapter 5: Disjointness, Farness as Non-Convex Constraints

Informally, our main focus in this chapter is on equilibrium problems where players are trying to avoid “overlap" between their strategies. In the extreme case, players are prohibited from using strategies used by other players, i.e., the intersection of their supports is empty. We call this problem *disjoint Nash* and show that deciding whether a bi-matrix game has such an equilibrium, even under polynomial additive approximation, is NP-complete. We also consider a stronger form of disjoint Nash, which we call *partition Nash*, in which the players' supports must also partition the strategy space[5].

We will also consider a generalization of disjoint Nash by allowing overlap but requiring that mixed strategies be at least 2δ statistically far (i.e., with respect to $L1$ distance), obtaining a problem that we call *δ-far Nash.* It is natural to consider whether there is a threshold approximation error below which the problem δ-far Nash becomes total (has at least one guaranteed solution). While we do not have an answer to this question, we provide a bound in which a solution is guaranteed by the existence of a Nash equilibrium. Moreover, by viewing the mixed strategies of players as a representative of the strategies of a large number of agents, we explain that the statistical difference between the mixed strategies of the players can provide

[5] In the next chapter, we will apply this notion in the context of strategic resource allocation.

useful information for the analysis of public behavior, and we apply this perspective to income tax games as a motivating example.

As mentioned, there are two simple variations of Nash equilibrium. By constraining a player's best response, we consider the generalization of Nash's notion and consider disjointness and related constraints in this setting. For example, in the generalized version of a (mixed) disjoint Nash equilibrium (which we call a generalized disjoint equilibrium), a player in this equilibrium (not necessarily a Nash equilibrium) cannot deviate to a mixed strategy whose support has a nonempty intersection with those of the other players. This is unlike a disjoint Nash equilibrium, where a player may consider all possible strategies. While much of the work on generalized equilibria focuses on pure strategies in general games, both Charnes [44] and McKinsey [134] provide arguments for the reasonableness of considering constraints on mixed strategies.

Furthermore, work discussed in Chapter 3 and the existing literature on finding generalized equilibria focus on convex constraints where solutions are guaranteed to exist by well-known mathematical facts [57, 62, 152, 161]. However, most problems we consider do not fall into this setting. In the generalized setting, at least one of the problems we propose, namely, approximate generalized disjoint equilibrium, is tractable and has at least one (trivial) solution. However, we also show in this setting that one selfish player can degrade the social welfare of all other players. With this drawback in mind, we direct our attention to another type of generalized equilibrium (which we call *restricted equilibrium*) by excluding the possibility of minor probabilities.

1.2.5 Chapter 6: A Meta Heuristic Approach For Strategic Resource Allocation

Strategic fair division is a branch of fair division in which players (agents) may act uncooperatively to maximize their utility. In particular, the participants may hide their true preferences, rather than play sincerely according to their true preferences. One branch of strategic fair division is related to game theory and studies the equilibria in games created by fair division algorithms [40, 37, 177]. We can cite *envy-freeness* [86] as one of the well-known fairness constraints in strategic fair resource allocation problems. The other branch is related to mechanism design and aims to find truthful mechanisms for fair division problems. Several notions are defined in the literature with respect to designing mechanisms that incentivize players to reveal their true preferences. In a *truthful* mechanism, strategic behavior cannot help any player achieve a better payoff. We do not address truthful mechanisms in this research.

We will propose a simple approach to model item allocation problems with players who have strategic behavior by using a more constrained notion of the Nash equilibrium problem. In particular, we argue that we can use the concept of Nash equilibrium with certain constraints to generate a fair allocation. For example, one of the constraints for such an equilibrium ensures that no player can use an item chosen by other players – which is an essential condition for fair division problems with indivisible items. To ensure a fair allocation, we add more constraints to ensure that all of the players can achieve a certain minimum payoff[6]. We introduce a problem that we call (*strategic arrangement problem*) that aims to find a constrained Nash equilibrium in standard normal-form games. In other words, we show that a solution to a strategic arrangement problem could imply a fair item allocation in a strategic setting. While most papers on fair division are focused on one specific fairness constraint (or multiple constraints separately), in this research, we consider alternative fairness criteria that include a variety of constraints, each of which is quite natural[7].

We begin with a general fair division problem of indivisible items (*fair item allocation*) without strategic behavior and transform it into a pseudo-game (*simple arrangement problem*). We show that even when we do not include strategic behavior, the problem of whether there exists a fair allocation remains NP-complete. Next, for the purpose of modeling strategic behavior, we expand this problem into a standard normal-form game (strategic arrangement problem) where the goal becomes finding a constrained Nash equilibrium. The quasi-polynomial algorithm provided in [130] is known to be close to optimal assuming the exponential time hypothesis [166]. This quasi-polynomial algorithm iterates blindly through certain mixed strategies, some of which correspond to our constraints, and eventually will find an approximate Nash equilibrium. However, we show that this method for finding our constrained Nash equilibrium has some disadvantages. This is because a solution with all the constraints that we desire is not guaranteed to exist.

To the best of our knowledge, limited progress has been made on local search-adopted techniques for the computation of unconstrained Nash equilibria with mixed

[6]Several works address the relationship between Nash equilibrium and social welfare functions for various problems [52, 39, 40, 157, 133].

[7]These fairness criteria are separately discussed in Chapter 5, and we consider them as fairness constraints (all enforced together) in this chapter.

strategies [42, 69]. Recently, human-based meta-heuristic algorithms are receiving considerable attention for optimization purposes [8, 14, 34, 63, 178, 182]. For the purpose of finding a constrained Nash equilibrium, we will employ a human-based meta-heuristic algorithm to search the potential strategy space faster based on the structure of the game. A gently modified version of the bus transportation algorithm proposed in [34] will help us perform a faster search for finding a desirable equilibrium.

Chapter 2

Preliminary Definitions and Existence Results

This Chapter is devoted to presenting fundamental definitions and basic facts used in the thesis.

2.1 Elementary Mathematical Definitions

Throughout this thesis, we will use ℓ_2 distance unless otherwise stated. We also use $\mathbb{1}$ to denote the all ones vector. To prevent confusion, we use $\prec$ and $\succ$ for comparison between vectors as opposed to $>$ and $<$ for scalars.

2.1.1 Convex Sets and Functions

Definition 2.1.1. *A set $C \subseteq \mathbb{R}^n$ is called convex if, for any two points $x_1, x_2 \in C$ and for any $\alpha \in [0, 1]$, the following holds:*

$$\alpha x_1 + (1 - \alpha)x_2 \in C.$$

Definition 2.1.2. *A vector-valued function $f : \mathbb{R}^n \to \mathbb{R}^m$ is convex if, for any $x_1, x_2 \in \mathbb{R}^n$ and $\alpha \in [0, 1]$, the following holds:*

$$f(\alpha x_1 + (1 - \alpha)x_2) \preceq \alpha f(x_1) + (1 - \alpha)f(x_2),$$

where $\preceq$ denotes element-wise inequality.

Definition 2.1.3. *A function $f : \mathbb{R}^n \to \mathbb{R}$ is said to be* convex with respect to the variable x_i *if, for all fixed values of the other variables x_{-i}, and for all $x_{i1}, x_{i2} \in \mathbb{R}$ and $\alpha \in [0, 1]$, the following holds:*

$$f\left(\alpha x_{i1} + (1-\alpha)x_{i2}, x_{-i}\right) \leq \alpha f\left(x_{i1}, x_{-i}\right) + (1-\alpha) f\left(x_{i2}, x_{-i}\right),$$

where $x_{-i} \in \mathbb{R}^{n-1}$ denotes the vector of all variables other than x_i, held fixed.

Remark 2.1.4. The definition of general convex games (which we will define later) assumes this condition on the loss function. We can similarly define a similar notion for $f : \mathbb{R}^n \to \mathbb{R}^m$.

Definition 2.1.5. *A function $f_i : S_1 \times \cdots \times S_k \to \mathbb{R}$ is t-*multi-convex *($t \leq k$) if, for every subset $J \subseteq [t]$ with $|J| \leq t$ and every fixed $\mathbf{s}_{-J} \in \prod_{j \notin J} S_j$, the restricted function $f(\cdot, \mathbf{s}_{-J}) : \prod_{j \in J} S_j \to \mathbb{R}$ is convex. That is, for all $s_J, s'_J \in \prod_{j \in J} S_j$ and $\lambda \in [0, 1]$,*

$$f(\lambda s_J + (1-\lambda)s'_J, s_{-J}) \leq \lambda f(\mathbf{s}_J, \mathbf{s}_{-J}) + (1-\lambda) f(s'_J, s_{-J}).$$

2.1.2 Correspondences and Semi-Continuity

Definition 2.1.6. *For sets X, Y, a* correspondence *(also called a* set-valued map *or a* a point-to-set map*) from X to Y is a function $\mathcal{R} : X \to 2^Y$. Sometimes we will denote such a mapping by $\mathcal{R} : X \rightrightarrows Y$.*

Definition 2.1.7. *The correspondence $\mathcal{R} : X \rightrightarrows Y$ is called* convex-valued *if for all $x \in X$, $\mathcal{R}(x)$ is a convex set.*

Definition 2.1.8. *The correspondence $\mathcal{R}$ is* upper semicontinuous *at $x \in X$ if for every open neighborhood V of $\mathcal{R}(x)$ there is an open neighborhood U of x such that for every $y \in U$, $\mathcal{R}(y) \subseteq V$.*

Definition 2.1.9. *The correspondence $\mathcal{R}$ is* lower semicontinuous *at x if for any open set V such that $\mathcal{R}(x) \cap V$ is nonempty, there exists a neighborhood U of x such that $\mathcal{R}(y) \cap V$ is nonempty for all $y \in U$.*

2.1.3 Distance, Projections and Hausdorff Metrics

Definition 2.1.10. *Let X be a convex, non-empty, and compact set.*

- *We let* $\mathrm{d}(x, z)$ *denote the Euclidean distance between two points* $x, z \in \mathbb{R}^m$.
- *The projection map of a point* x *to the set* X *is* $\Pi_X(x) = \arg\inf_{z \in X} \ \mathrm{d}(x, z)$.
- *The closed* ϵ*-parallel body of* X *is defined as* $\overline{\mathrm{B}}(X, \epsilon) := \bigcup_{x \in X}\{z \in \mathcal{M} : \mathrm{d}(x, z) \leq \epsilon\}$. *Note that* $\mathcal{M}$ *is the universe. The inner closed* ϵ*-parallel body of* X *is* $\overline{\mathrm{B}}(X, -\epsilon) = \{x \in X : \overline{\mathrm{B}}(x, \epsilon) \subseteq X\}$. *The elements of* $\overline{\mathrm{B}}(X, -\epsilon)$ *can be viewed as the points "deep inside of* X*", while* $\overline{\mathrm{B}}(X, \epsilon)$ *as the points that are "almost inside of* X*".*

Definition 2.1.11. *We also need to define Hausdorff semi-continuity notions:*

- *We define the set-point distance of a point* x *from a set* X *to be* $\mathrm{dist}(x, X) := \inf_{z \in X} \ \mathrm{d}(x, z)$.
- *Let* X *and* Y *be two non-empty sets. We define their Hausdorff distance* $\mathrm{d_H}(X, Y)$ *for two sets* X *and* Y *to be:*

$$\max\left\{\sup_{x \in X} \mathrm{dist}(x, Y), \sup_{y \in Y} \mathrm{dist}(X, y)\right\} = \inf\{\epsilon \geq 0 : X \subseteq \overline{\mathrm{B}}(Y, \epsilon) \wedge y \subseteq \overline{\mathrm{B}}(X, \epsilon)\}$$

- *A correspondence* $\mathcal{F} : X \rightrightarrows Y$ *is called (Hausdorff) upper semi-continuous at a point* $x_0 \in X$ *if and only if for every* $\epsilon > 0$ *there exists a neighborhood* N *of* x_0 *such that for all* $x \in N(x_0)$, $\mathcal{F}(x) \subseteq \overline{\mathrm{B}}(\mathcal{F}(x_0), \epsilon)$.
- *A correspondence* $\mathcal{F} : X \rightrightarrows Y$ *is called (Hausdorff) lower semi-continuous at a point* $x_0 \in X$ *if and only if for every* $\epsilon > 0$ *there is a neighborhood* N *of* x_0 *such that* $\mathcal{F}(x_0) \subseteq \overline{\mathrm{B}}(\mathcal{F}(x), \epsilon)$ *for all* $x \in N(x_0)$.
- *A correspondence* $\mathcal{F} : X \rightrightarrows Y$ *is called L-Hausdorff Lipschitz continuous if there exists a positive constant* $L \in \mathbb{R}^+$ *such that, for all* x_1 *and* x_2 *in* $X, \mathrm{d_H}(\mathcal{F}(x_1), \mathcal{F}(x_2)) \leq L \cdot \mathrm{d}(x_1, x_2)$.
- *A correspondence* $\mathcal{F}$ *is said to be* $(\eta, \sqrt{m}, L)$*-well-conditioned if it is L-Hausdorff Lipschitz, and for every* x*, its output is a closed, compact, and convex subset of* $\mathbb{R}^m$ *that contains a ball of radius* η *for some fixed* $\eta > 0$. *This definition will be used as a basic standard assumption for correspondences we focus on in this thesis when weak separation oracles are considered, similar to the approach of [152].*

Remark 2.1.12. If the correspondence $\mathcal{F}(x)$ is compact for every x and possesses global Hausdorff Lipschitz properties, then the conditions for both upper and lower semi-continuity are inherently satisfied.

2.2 General Games and Normal Form Games

Definition 2.2.1 (General Game)**.** *A general game consists of:*

- *A finite set of k* players*, indexed by $i \in \{1, \ldots, k\}$,*
- *For each player i, a (possibly infinite) set S_i of* pure strategies*,*
- *For each player i, a* utility function *(also called* payoff *function) $u_i : S_1 \times \cdots \times S_k \to \mathbb{R}$, which assigns a real-valued payoff to each strategy profile*[1]*.*

The tuple $(\mathcal{G}, \mathcal{S}, \mathcal{U})$ where $\mathcal{S} = \{S_1, \ldots, S_k\}$ and $\mathcal{U} = \{u_1, \ldots, u_k\}$, defines the game. A strategy profile *is an element $(s_1, \ldots, s_k) \in S_1 \times \cdots \times S_k$, and player i's payoff under this profile is $u_i(s_1, \ldots, s_k)$.*

Remark 2.2.2. In the minimization setting, we may equivalently consider *loss function $\theta_i : S_1 \times \cdots \times S_k \to \mathbb{R}$* for each player i, representing the loss (rather than utility) associated with each strategy profile. For various definitions, we adhere to their classical definitions, which are framed as either maximization or minimization problems, to maintain consistency with the original terminology.

Definition 2.2.3 (Normal-Form Game)**.** Normal-form games *are a special case of general games in which:*

- *Each player has a finite* pure strategy set *S_i,*
- *The strategy sets are often indexed uniformly as $S_i = \{1, \ldots, n\}$ for all i (when the sets are of equal size)*[2]*,*
- *The utility functions $u_i : S_1 \times \cdots \times S_k \to \mathbb{R}$ are represented as explicit tables or multidimensional arrays (also called* payoff matrices*).*

[1] They are typically represented by polynomially computable circuits.

[2] A more general formulation provides different strategy sets for each player. Later, we show that there is no essential difference between such a formulation and the one used in this research when we focus on the notion of exact (i.e., not approximate) Nash equilibrium (see Lemma 5.2.5).

Remark 2.2.4. For the case of two players, the game is called a *bimatrix game*, and each player's utility can be represented by a matrix $A_1, A_2 \in \mathbb{R}^{n\times n}$.

In many situations, it is useful to allow players to randomize over pure strategies. This leads to the following:

Definition 2.2.5. *Let $\triangle_n$ be the the probability simplex in n-dimension,*

$$\triangle_n = \left\{ x \in \mathbb{R}^n \mid x_i \geq 0, \forall i, \sum_{i=1}^{n} x_i = 1 \right\}.$$

A mixed strategy *is an element of $\triangle_n$, corresponding to a distribution on strategies. If x is a mixed strategy,* Supp(x) *(*support *of x) denotes the set of elements of S to which x assigns positive probability. We can extend definition of utilities u_i to mixed strategy profiles by defining $u_i(\mathbf{x})$ as $\mathbb{E}_{\mathbf{s}\leftarrow\mathbf{x}}[u_i(\mathbf{s})]$.*

Remark 2.2.6. Note that we may view a pure strategy s as a mixed strategy assigning probability 1 to s. Recall that a pure strategy profile is a k-dimensional vector $\mathbf{s} = (s_1, \ldots, s_k)$ of strategies, where s_i is the pure strategy played by player i. We can define a *mixed strategy profile* $\mathbf{x} = (x_1, \ldots, x_k)$ similarly. In a bi-matrix game, for a mixed strategy profile $\mathbf{z} = (x, y)$ we have:

$$u_1(x,y) = \sum_{1\leq i\leq n} \sum_{1\leq j\leq n} \alpha_{ij} u_1(i,j) \text{ and } u_2(x,y) = \sum_{1\leq i\leq n} \sum_{1\leq j\leq n} \alpha_{ij} u_2(i,j),$$

where $\alpha_{ij} = x_i y_j$.

2.3 Nash Equilibrium

Definition 2.3.1. *A pure strategy s, is a* best response *for player i to a strategy profile $\boldsymbol{s}^*$ if for all $s' \in S_i$ we have $u_i\left(s, \boldsymbol{s}^*_{-i}\right) \geq u_i\left(s', \boldsymbol{s}^*_{-i}\right)$. We can also define a* mixed best response *strategy to a mixed strategy profile in the same manner.*

Definition 2.3.2. *A* Nash equilibrium *is a strategy profile $\boldsymbol{s} = (s_1^*, \ldots, s_k^*)$ such that for $1 \leq i \leq k$ and*

$$u_i(\boldsymbol{s}^*) \geq u_i(s_i, \boldsymbol{s}^*_{-i}), \qquad \forall s_i \in S_i \tag{2.1}$$

where $\boldsymbol{s}^*_{-i}$ includes all strategies except player i.

Definition 2.3.3. *An* ϵ-approximate Nash equilibrium *is a strategy profile* $\boldsymbol{s} = (s_1^*, \ldots, s_k^*)$ *such that for* $1 \leq i \leq k$ *and*

$$u_i(\boldsymbol{s}^*) + \epsilon \geq u_i(s_i, \boldsymbol{s}_{-i}^*), \qquad \forall s_i \in S_i \tag{2.2}$$

where $\boldsymbol{s}_{-i}^*$ *includes all strategies except player* i*. We also define* ϵ-best response *similarly.*

Definition 2.3.4. *We can similarly define the equivalent minimization version for the loss functions* $\Theta = (\theta_1, \cdots, \theta_k)$*:*

$$\theta_i(\boldsymbol{s}^*) \leq \theta_i(s_i, \boldsymbol{s}_{-i}^*) + \epsilon, \qquad \forall s_i \in S_i \tag{2.3}$$

where $\boldsymbol{s}_{-i}^*$ *includes all strategies except player* i*.*

Remark 2.3.5. We can assume that each S_i is a bounded subset of $\mathbb{R}^{n_i}$ and each player i controls $dim(S_i) = n_i$ variables of the strategy profile. Furthermore, $\sum_{i=1}^k n_i = n$.

Remark 2.3.6. We note that there are equivalent versions of the definitions of best response and Nash equilibrium in which the restriction to pure strategies is relaxed, i.e., players compare the payoff using a particular mixed strategy against all alternative mixed strategies.

2.4 Special Games and Definitions

Definition 2.4.1. *A* strictly dominated strategy *is a strategy that always delivers a lower payoff compared to all alternative strategies, regardless of what strategy the opponent chooses.*

Remark 2.4.2. A strictly dominated strategy cannot participate in a Nash equilibrium. However, in an approximate Nash equilibrium, a strictly dominated strategy can be played with a small probability based on each player's utility matrix and the approximation error.

Definition 2.4.3. *The (additive)* social welfare *of a strategy profile* $\boldsymbol{x}$ *is the sum* $\sum_{i=1}^k u_i(\boldsymbol{x})$ *of all players' payoffs.*

Definition 2.4.4. *A* fully mixed strategy *is a mixed strategy in which all pure strategies in S are assigned with a positive (greater than zero) probability. This concept has been extensively studied, which explains why there are various definitions [112, 117, 118, 138, 155, 158].*

Throughout the thesis, we will refer to well-known restrictions, including *zero-sum* and *symmetric* games:

Definition 2.4.5. *A 2-player game is a* zero-sum game *if $u_1(s_1, s_2) + u_2(s_1, s_2) = 0$ for all $s_1, s_2 \in S$. This means that one player's payoff is equal to the other player's loss on any given play of the game.*

Definition 2.4.6. *The game $\mathcal{G}$ is* symmetric *if both players have the same strategies and $u_1(s_1, s_2) = u_2(s_2, s_1)$. In addition, we call x^* a* symmetric equilibrium *if (x^*, x^*) is a Nash equilibrium.*

2.5 Beyond Nash Equilibrium

In this section, we present a range of equilibrium problems, including generalized Nash equilibrium, L/F-equilibrium in multi-leader-follower games, and resilient Nash equilibrium.

2.5.1 Generalized Nash Equilibrium

A natural way to limit the notion of a player's best response is to suppose that player i, when considering responses to $\mathbf{s}^*_{-i}$, is constrained to those lying in some subset of S_i, determined by $\mathbf{s}^*_{-i}$. This constraint might be given as a correspondence $\mathcal{R} : S_{-i} \rightrightarrows S_i$, where S_{-i} denotes $S_1 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_k$. We then have the following conditions:

$$u_i(s^*_i, \mathbf{s}^*_{-i}) + \epsilon \geq u_i(s_i, \mathbf{s}^*_{-i}), \qquad \forall s_i \in \mathcal{R}_i(\mathbf{s}^*_{-i}) \qquad (**)$$

Given this notion of best response, Debreu generalized the definition of Nash equilibrium:

Definition 2.5.1. *A* generalized (social) equilibrium *is a (pure) strategy profile $\boldsymbol{s} = (s^*_1, \ldots, s^*_n)$ which satisfies $s^*_i \in \mathcal{R}_i(s^*_{-i})$ and (**) for $1 \leq i \leq k$.*

Remark 2.5.2 (Alternative Representations)**.** We can consider the minimization version by considering loss or regret functions. Also, following the approach in some applications in generalized (Nash) equilibria and multi-leader-follower games [72, 150], we may assume that each player i's constraints are represented by the non-empty set given by:

$$\mathcal{R}_i\left(\mathbf{x}_{-i}\right) \equiv \left\{x_i \in S_i \mid g_i\left(x_i, \mathbf{x}_{-i}\right) \leq 0,\ h_i\left(x_i\right) \leq 0\right\},$$

where $x_i \in S_i$ and $g_i : \Pi_{j=1}^{k} S_j \to \mathbb{R}^{m_i}$ and $h_i : S_{\mathrm{i}} \to \mathbb{R}^{l_i}$ are continuously differentiable, h_i is convex, and $g_i\left(\cdot, \mathbf{x}_{-i}\right)$ is convex for any fixed $\mathbf{x}_{-i}$. Note that m_i and l_i are some given constants.

Remark 2.5.3. We can similarly consider this definition in normal-form games with mixed strategies.

2.5.2 L/F-Equilibrium in Multi-Leader-Follower Games

We now extend the equilibrium concept to multi-leader–follower games. An L/F equilibrium is a solution concept in multi-leader–follower games in which the leaders' and followers' strategies are optimal given the other followers' and leaders' responses and, no individual player, whether a leader or a follower, can improve their utility (or minimize their *loss* or *regret* functions) by unilaterally altering their current strategy. Stackelberg games, which are a special case of multi-leader-follower games with only one leader, may be formulated as mathematical programs with equilibrium constraints (MPEC). In this case, the followers' problems are replaced by a constraint given by their optimality conditions. In a broader context, an MPEC is an optimization problem that encompasses two sets of variables, namely decision variables and response variables [77, 131, 148]. The framework commonly used to represent the multi-leader-follower game is referred to as the equilibrium problem with equilibrium constraints (EPEC). An EPEC [67, 80, 109, 110, 175, 176] is essentially an equilibrium problem composed of multiple parametric MPECs, each of which incorporates other players' strategies as parameters.

For simplicity, we only consider games that have two leaders, labeled by I and II, and k followers, labeled by $i \in [k]$ with strategies S_i. We denote the *domain* of strategies of leaders I and II by X^{I} and X^{II} respectively. The leaders' loss functions are denoted by $\phi_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}, y)$ and $\phi_{\mathrm{II}}(x_{\mathrm{I}}, x_{\mathrm{II}}, y)$. The notation implies that the loss of each leader is determined by both its own strategies and those of the opposing leader, as well as the strategies of the followers represented as the vector y. The followers also respond to the leaders' strategies as follows. For each follower $i = 1, \ldots, k$, let

$\theta_i(x_{\mathrm{I}}, x_{\mathrm{II}}, y)$ and $\mathcal{R}_i(x_{\mathrm{I}}, x_{\mathrm{II}}, y_{-i})$ denote follower i's loss function and available constrained strategy set respectively. This strategy set depends on the pair of strategies $(x_{\mathrm{I}}, x_{\mathrm{II}}) \in X^{\mathrm{I}} \times X^{\mathrm{II}}$. For each such pair $(x_{\mathrm{I}}, x_{\mathrm{II}})$, the followers' problem is modeled by a Debreu game (generalized Nash equilibrium problem) parameterized by the leaders' strategies. Let $Y_{sol}(x_{\mathrm{I}}, x_{\mathrm{II}})$ denote the set of such solutions (not necessarily a singleton). Each element $\bar{y} \in Y_{sol}(x_{\mathrm{I}}, x_{\mathrm{II}}) \subseteq \Pi_{i=1}^{k} S_i$ is a tuple $(\bar{y}_i)_{i=1}^{k}$ where for each follower i, $\bar{y}_i$ is a solution of problem 2.4. The tuple $(x_{\mathrm{I}}, x_{\mathrm{II}}, \bar{y}_{-i})$ is external to the minimization program 2.4, and y_i is the primary variable that must be computed.

$$\begin{aligned} &\text{Min}\, \theta_i\left(x_{\mathrm{I}}, x_{\mathrm{II}}, \bar{y}_{-i}, y_i\right) \\ &\text{s.t } y_i \in \mathcal{R}_i\left(x_{\mathrm{I}}, x_{\mathrm{II}}, \bar{y}_{-i}\right), \end{aligned} \tag{2.4}$$

We now define the concept of equilibrium in multi-leader-follower games. A pair $(x^*_{\mathrm{I}}, x^*_{\mathrm{II}}) \in X^{\mathrm{I}} \times X^{\mathrm{II}}$ is called a *L/F-(generalized Nash) equilibrium*, if there exists $(y^*_{\mathrm{I}}, y^*_{\mathrm{II}})$ such that $(x^*_{\mathrm{I}}, y^*_{\mathrm{I}})$ is an optimal solution of leader I's problem, which tries to find a pair $(x_{\mathrm{I}}, y_{\mathrm{I}})$ to the following:

$$\begin{aligned} &\text{Min}\, \phi_{\mathrm{I}}\left(x_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{\mathrm{I}}\right) \\ &\text{s.t } x_{\mathrm{I}} \in X^{\mathrm{I}} \\ &\text{and } y_{\mathrm{I}} \in Y_{sol}\left(x_{\mathrm{I}}, x^*_{\mathrm{II}}\right) \end{aligned} \tag{2.5}$$

and $(x^*_{\mathrm{II}}, y^*_{\mathrm{II}})$ is an optimal solution for leader II trying to find a pair $(x_{\mathrm{II}}, y_{\mathrm{II}})$ to the following problem:

$$\begin{aligned} &\text{Min}\, \phi_{\mathrm{II}}\left(x^*_{\mathrm{I}}, x_{\mathrm{II}}, y_{\mathrm{II}}\right) \\ &\text{s.t } x_{\mathrm{II}} \in X^{\mathrm{II}} \\ &\text{and } \quad y_{\mathrm{II}} \in Y_{sol}\left(x^*_{\mathrm{I}}, x_{\mathrm{II}}\right) \end{aligned} \tag{2.6}$$

In the given definition, the equilibrium strategies of the followers, represented as y^*_{I} and y^*_{II}, belong to the same set $Y_{sol}(x^*_{\mathrm{I}}, x^*_{\mathrm{II}})$. However, they are not required to be the same. This flexibility arises because these strategies are based on different anticipations by Leader I and Leader II of how the followers collectively respond to the pair of strategies $(x^*_{\mathrm{I}}, x^*_{\mathrm{II}})$. One could introduce another variant of this problem by enforcing $y^*_{\mathrm{I}} = y^*_{\mathrm{II}}$. However, even if $Y_{sol}(x_{\mathrm{I}}, x_{\mathrm{II}})$ contains only one unique response for all pairs of strategies $(x_{\mathrm{I}}, x_{\mathrm{II}})$, a L/F-equilibrium may not exist [150].

Proposition 2.5.4. *Assume that in a multi-leader-follower game , we have k followers with continuously differentiable loss functions* $\Theta = (\theta_1, \ldots, \theta_k)$*, and strategies*

$\mathcal{S} = (S_1, \ldots, S_k)$ and constraints $\mathcal{R} = (\mathcal{R}_1, \ldots, \mathcal{R}_k)$. In addition, we have 2 leaders with strategies $\mathcal{X} = (X^{\mathrm{I}}, X^{\mathrm{II}})$ and utilities $\Phi = (\phi_{\mathrm{I}}, \phi_{\mathrm{II}})$. The problem of finding a L/F-equilibrium of this game can be transformed into an EPEC problem.

Proof. A pair $(x^*_{\mathrm{I}}, x^*_{\mathrm{II}}) \in X^{\mathrm{I}} \times X^{\mathrm{II}}$ is a *L/F-equilibrium*, if there exists $(y^*_{\mathrm{I}}, y^*_{\mathrm{II}})$ such that $(x^*_{\mathrm{I}}, y^*_{\mathrm{I}})$ is an optimal solution of leader I's problem, which tries to find $(x_{\mathrm{I}}, y_{\mathrm{I}})$ such that[3]:

$$\begin{aligned}
&\text{Min}\, \phi_{\mathrm{I}}\left(x_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{\mathrm{I}}\right) \\
&\text{s.t } x_{\mathrm{I}} \in X^{\mathrm{I}} \\
&\text{and } y_{\mathrm{I}} \text{ solves } QVI(Y_{sol}(x_{\mathrm{I}}, x^*_{\mathrm{II}}), F_{\mathrm{I}}(x_{\mathrm{I}}, x^*_{\mathrm{II}}, \cdot))
\end{aligned}$$

where F_{I} is defined as follows:

$$F_{\mathrm{I}}\left(x_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{\mathrm{I}}\right) := \begin{pmatrix} \partial_{y_{(1,\mathrm{I})}} \theta_1 \left(x_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{(1,\mathrm{I})}, y^*_{(-1,\mathrm{I})}\right) \\ \vdots \\ \partial_{y_{(k,\mathrm{I})}} \theta_k \left(x_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{(k,\mathrm{I})}, y^*_{(-k,\mathrm{I})}\right) \end{pmatrix}$$

For clarity, all exogenous variables are marked by $*$. The notation $y^*_{(-k,\mathrm{I})}$ refers to all followers' anticipated strategies except k for leader I. A similar transformation can be used for the second leader.

□

2.5.3 Equilibria and Coalitions

Considerable effort in game theory has been devoted to understanding deviations by coalitions of players, tracing back to the foundational work of Aumann [10]. Here we review some of the relevant notions. Consider a k-player game $(\mathcal{G}, \mathcal{U}, \mathcal{S})$. Let $\mathcal{J}$ be the set of subsets of $\{1, \ldots, k\}$, $J \in \mathcal{J}$ denote a *coalition* and $S_J \equiv \prod_{j \in J} S_j$ *the possible strategy set* for the coalition J. strong Nash equilibrium is particularly useful in areas such as the study of voting systems, in which a certain level of stability and resistance to coordinated deviations is required. Somewhat confusingly, the concept of a strong Nash equilibrium is unrelated to that of a weak Nash equilibrium, which is a well-known notion in game theory.

Definition 2.5.5 ([10]). *In a game with k-players, a strategy profile $\boldsymbol{s}^*$ is considered a* strong Nash equilibrium *if, for any coalition $J \subseteq \mathcal{J}$ and any deviation $\boldsymbol{s}'_J \in \Pi_{j \in J} S_j$*

[3] See Definition 2.7.1

by the players in the coalition J , there exists a player $j \in J$ such that the following condition holds:

$$u_j(\boldsymbol{s}^*) \geq u_j\left(\boldsymbol{s}'_J, \boldsymbol{s}^*_{-J}\right)$$

where $-J$ denote the complement of J in $[k]$ (which is equal to $[k] - J$).

Remark 2.5.6. It is equivalent to consider a strong Nash equilibrium to be a strategy profile in which for any $J \in \mathcal{J}$, there exists no player $j \in J$ and an alternative strategy $\mathbf{s}'_J$, such that the condition $u_j(\mathbf{s}^*) < u_j\left(\mathbf{s}'_J, \mathbf{s}^*_{-J}\right)$ holds.

Definition 2.5.7 ([6]). *A t-*strong Nash equilibrium *is a Nash equilibrium in which no coalition of size t or less of players can deviate so that all its members strictly benefit.*

The idea of a coalition-proof Nash equilibrium is relevant in situations where players can talk about their strategies but can't make binding commitments (see [29]). In simple terms, it emphasizes immunity to deviations (alternative strategies) that are self-enforcing.

Definition 2.5.8 ([29]). *For each $\boldsymbol{s}^0_{-J} \in S_{-J}$, we use $\mathcal{G}/\boldsymbol{s}^0_{-J}$ to denote the game induced on coalition J by the strategies s^0_{-J} for coalition $-J$, i.e., for all $j \in J$ and $\boldsymbol{s}_J \in S_J$ we define the utilities ($\bar{u}_j : S_J \to \mathbb{R}$) of the members of coalition J to be $u_j(\boldsymbol{s}_J, \boldsymbol{s}^0_{-J})$. Next,* self-enforceability *and* coalition-proofness *are defined by mutual recursion as follows:*

1. *In a single-player game $s^* \in S_1$ is a coalition-proof Nash equilibrium if and only if $\boldsymbol{s}^*$ maximizes $u_1(\boldsymbol{s})$.*

2. *Let $k > 1$ and assume that a coalition-proof Nash equilibrium has been defined for games with fewer than k players. Then,*

   - *For any game $\mathcal{G}$ with k players, $\boldsymbol{s}^* \in S$ is self-enforcing if, for all $J \in \mathcal{J}$, $\boldsymbol{s}^*_J$ is a coalition-proof Nash equilibrium in the game $\mathcal{G}/\boldsymbol{s}^*_{-J}$.*

   - *For any game $\mathcal{G}$ with k players, $\boldsymbol{s}^* \in S$ is a coalition-proof Nash equilibrium if it is self-enforcing and if there does not exist another self-enforcing strategy vector $\boldsymbol{s} \in S$ such that $u_j(\boldsymbol{s}^*) < u_j(\boldsymbol{s})$ for all $j \in [k]$.*

Next, we define a more restricted equilibrium notion called t-resilient Nash equilibrium, introduced in [1], which is fundamental to applications such as secret sharing and multi-party computation [1, 2, 95, 96, 125].

Definition 2.5.9 ([1]). *Given a nonempty set $J \subseteq [k], \boldsymbol{s}^*_J \in S_J$ is a group best response for J to $\boldsymbol{s}^*_{-J} \in S_{-J}$ if, for all $\boldsymbol{s}'_J \in S_J$ and all $j \in J$, we have:*

$$u_j\left(\boldsymbol{s}^*\right) \geq u_j\left(\boldsymbol{s}'_J, \boldsymbol{s}^*_{-J}\right).$$

*A strategy profile $\boldsymbol{s}^* \in \mathcal{S}$ is a t-*resilient Nash equilibrium *if, $\forall J \subseteq [k]$ with $|J| \leq t, s^*_J$ is a group best response for J to $\boldsymbol{s}^*_{-J}$*[4]*.*

Remark 2.5.10. The t-resilient Nash equilibrium problem is the most restricted of these three solution concepts, as every t-resilient Nash equilibrium is also a t-strong Nash equilibrium. In other words, in a t-strong Nash, every coalition deviation fails because at least one member does not strictly gain, whereas in a t-resilient Nash, every coalition deviation fails because no member can strictly gain. Every strong Nash equilibrium is also inherently a coalition-proof Nash equilibrium. This follows by definition; the set of possible deviations in the latter is a subset of that in the former. 1-resilient and 1-strong Nash equilibria are equivalent to a Nash equilibrium. Furthermore, a coalition-proof Nash equilibrium for two players is equivalent to a two-player Nash equilibrium. The connection between Nash and coalition-proof Nash equilibrium is less straightforward for more than two players. The paper [29] presents an example of a three-player game lacking coalition-proof Nash equilibria. To our knowledge, no other definitive inclusion relationship between these two classes has been established for multiplayer games.

2.5.4 Pareto Efficiency (Maximization Version)

The following definitions are provided for the sake of completeness.

Definition 2.5.11. *A strategy profile s^* is* strictly(strong) Pareto optimal *if there is no alternative strategy profile s such that there exists at least one player i such that for all $j \in [k] - \{i\}$:*

$$u_j(s^*) \leq u_j(s) \text{ and } u_i(s^*) < u_i(s)$$

Definition 2.5.12. *A strategy profile s^* is* weakly Pareto optimal *if there is no strategy profile s such that for all $j \in [k]$:*

$$u_j(s^*) < u_j(s)$$

[4]We call a strategy *strongly resilient* if it is t-resilient for all $t \leq k - 1$.

Remark 2.5.13. Every strong Nash equilibrium is weakly Pareto optimal. This condition does not necessarily hold in a coalition-proof Nash equilibrium. A strong Nash equilibrium with strong Pareto optimality is called *super Nash equilibrium*[163].

Remark 2.5.14. Informally, a strong Pareto optimal strategy is a strategy profile such that any alternative strategy will make at least one player worse off. A weak Pareto optimal is a strategy profile such that any alternative strategy will make at least one player worse off, but may not make any player worse off.

2.6 Games With Uncertainty

We now consider the above-mentioned equilibrium problems under uncertainty within a distribution-free model. Specifically, we consider the notion of Robust Nash equilibrium in different game formats.

2.6.1 Robust Nash Equilibrium for General Games

Robust Nash equilibrium, introduced by [3, 103], is one approach to addressing uncertainty in general games. This concept was initially developed for simpler games and was subsequently extended in [147] to k-player non-cooperative games with nonlinear loss functions. To address uncertainty in this setting, a specific *uncertainty parameter* of the form $u_i \in U^i \subset \mathbb{R}^{l_i}$ is introduced for each player i, and it becomes part of the player's loss function, allowing a representation of the uncertainty. In particular, the loss function may now be expressed as $\theta_i : S_i \times S_{-i} \times U^i \to \mathbb{R}$. Player i does not know the exact value of parameter u_i, only that it must belong to a given nonempty set $U^i \subseteq \mathbb{R}^{l_i}$. Finally, player i tries to find x_i solving the following optimization problem with parameter $u_i \in U^i$ and x^*_{-i} as a given external parameter:

$$\begin{aligned} &\text{Min}\, \theta_i \left(x_i, x^*_{-i}, u_i\right) \\ &\text{s.t } x_i \in S_i \end{aligned} \tag{2.7}$$

By the robust optimization paradigm, we assume that each player i tries to minimize their worst-case loss function. Under this assumption, each player i considers the worst-case loss function $\tilde{f}_i : S_i \times S_{-i} \to \mathbb{R}$, defined as $\tilde{f}_i\,(x_i, x_{-i}) :=$

$\sup \{\theta_i (x_i, x_{-i}, u_i) \mid u_i \in U^i\}$, and tries to solve the following problem:

$$\begin{aligned} &\text{Min } \tilde{f}_i \left(x_i, x_{-i}^*\right) \\ &\text{s.t } x_i \in S_i \end{aligned} \tag{2.8}$$

This indeed is a Nash equilibrium problem with complete information. We now define the resulting notion of equilibrium with uncertainty. The definition can be extended to the approximate version, as well as to generalized forms of Nash equilibrium.

Definition 2.6.1. *A strategy profile x^* is called a* robust Nash equilibrium *of the game of the optimization program (2.7) if x^* is a (Nash) solution of (2.8).*

2.6.2 Robust Equilibrium for Multi-Leader-Follower Games

Next, we consider extending the solution concept of robust Nash equilibrium to multi-leader-follower games. We consider only two types of uncertainty caused by observation errors of the leaders towards other leaders and also the followers similar to [107]. In a multi-leader-follower with two leaders I and II, leader I tries to find x_{I} and y_{I} solving the following optimization problem:

$$\begin{aligned} &\text{Min } \phi_{\mathrm{I}} (x_{\mathrm{I}}, x_{\mathrm{II}}^*, y_{\mathrm{I}}, u_{\mathrm{I}}) \\ &\text{s.t } x_{\mathrm{I}} \in X^{\mathrm{I}} \\ &\text{and } y_{\mathrm{I}} \in Y_{sol} (x_{\mathrm{I}}, x_{\mathrm{II}}^*) \end{aligned} \tag{2.9}$$

where the loss function $\phi_{\mathrm{I}} : X^{\mathrm{I}} \times X^{\mathrm{II}} \times \Pi_{i=1}^{k} S_i \times U^i \to \mathbb{R}$ has an uncertainty parameter $u_{\mathrm{I}} \in \mathbb{R}^{l_{\mathrm{I}}}$. As described above, although leader I does not know the exact value of parameter u_i, the leader knows that it must belong to a given nonempty set U^{I}. We can similarly write the optimization problem for leader II:

$$\begin{aligned} &\text{Min } \phi_{\mathrm{II}} (x_{\mathrm{I}}^*, x_{\mathrm{II}}, y_{\mathrm{II}}, u_{\mathrm{II}}) \\ &\text{s.t } x_{\mathrm{II}} \in X^{\mathrm{II}} \\ &\text{and } y_{\mathrm{II}} \in Y_{sol} (x_{\mathrm{I}}^*, x_{\mathrm{II}}) \end{aligned} \tag{2.10}$$

The optimization program (2.4) only examines the followers' problem solely from the perspective of the followers, while the concept of equilibrium in multi-leader-follower games reflects a hierarchical structure. From the perspective of the leaders, they cannot exactly anticipate the response of the followers due to uncertainty. In conclusion, leader I estimates that follower i will try to find $y_{(i,\mathrm{I})}$ which solves the following optimization problem:

$$\begin{aligned} &\operatorname{Min} \theta_i \left(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*, \bar{y}_{(\mathrm{I},-i)}, y_{(\mathrm{I},i)}, e_{\mathrm{I}}\right) \\ &\text{s.t } y_{(\mathrm{I},i)} \in \mathcal{R}_i \left(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*, \bar{y}_{(\mathrm{I},-i)}\right) \end{aligned} \tag{2.11}$$

where each element $\bar{y}_{\mathrm{I}} \in Y_{sol}(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*) \subseteq Y(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*) \subseteq \Pi_{i=1}^k S_i$ is a tuple $(\bar{y}_{(\mathrm{I},i)})_{i=1}^k$ such that for each follower i, $\bar{y}_{\mathrm{I}}$ is a solution of (2.11). The tuple $(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*, \bar{y}_{(\mathrm{I},-i)})$ is external to (2.11), and $y_{(\mathrm{I},i)}$ is the variable that must be computed. Here $\bar{y}_{(\mathrm{I},-i)}$ is the estimated strategy of leader I of all followers except i. Similarly, leader II estimates that follower i will try to find $y_{(i,\mathrm{II})}$ which solves:

$$\begin{aligned} &\operatorname{Min} \theta_i \left(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*, \bar{y}_{(\mathrm{II},-i)}, y_{(\mathrm{II},i)}, e_{\mathrm{II}}\right) \\ &\text{s.t } y_{(\mathrm{II},i)} \in \mathcal{R}_i \left(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*, \bar{y}_{(\mathrm{II},-i)}\right), \end{aligned} \tag{2.12}$$

We considered the fact that each leader's anticipation of the followers can be different. Also, the leaders do not have any information about the uncertainty parameters e_{I} and e_{II} but they know that they must belong to the sets $E^{\mathrm{I}} \subset \mathbb{R}^{k_{\mathrm{I}}}$ and $E^{\mathrm{II}} \subset \mathbb{R}^{k_{\mathrm{II}}}$ respectively[5]. To define robust L/F-Nash equilibrium, we reformulate the problem for leader I with uncertainty parameters e_{I} and u_{I}[6]:

$$\begin{aligned} &\operatorname{Min} \phi_{\mathrm{I}} \left(x_{\mathrm{I}}, x_{\mathrm{II}}^*, y_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}^*, e_{\mathrm{I}}), u_{\mathrm{I}}\right) \\ &\text{s.t } x_{\mathrm{I}} \in X^{\mathrm{I}} \end{aligned} \tag{2.13}$$

For leader II, we have:

$$\begin{aligned} &\operatorname{Min} \phi_{\mathrm{II}} \left(x_{\mathrm{I}}^*, x_{\mathrm{II}}, y_{\mathrm{II}}(x_{\mathrm{I}}^*, x_{\mathrm{II}}, e_{\mathrm{II}}), u_{\mathrm{II}}\right) \\ &\text{s.t } x_{\mathrm{II}} \in X^{\mathrm{II}} \end{aligned} \tag{2.14}$$

where $y_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}^*, e_{\mathrm{I}})$ and $y_{\mathrm{II}}(x_{\mathrm{I}}^*, x_{\mathrm{II}}, e_{\mathrm{II}})$ represent one of solution of Equation (2.11) and (2.12) respectively and belong to Y_{sol}. Utilizing the robust optimization, we define a new game $\tilde{\mathcal{G}}$ with worst-case loss functions $\tilde{\psi}_{\mathrm{I}} : X^{\mathrm{I}} \times X^{\mathrm{II}} \to \mathbb{R}$ and $\tilde{\psi}_{\mathrm{II}} : X^{\mathrm{I}} \times X^{\mathrm{II}} \to \mathbb{R}$

[5] Recall that $\mathcal{R}_i \left(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*, \bar{y}_{(\mathrm{I},-i)}\right)$ is a nonempty, closed and convex correspondence. In [107, 147], they assume that $\mathcal{R}$ is a set and is independent from $y_{(\mathrm{I},-i)}$ and $y_{(\mathrm{II},-i)}$.

[6] This reformulation is well-known in the literature, see [107, 108]. Informally, remedial solutions try to characterize Y_{sol} in a computationally friendly manner.

as follows:

$$\begin{aligned}\tilde{\psi}_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}^*\right) &:= \sup\left\{\phi_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}^*, y_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}^*, e_{\mathrm{I}}), u_{\mathrm{I}}\right) \mid u_{\mathrm{I}} \in U^{\mathrm{I}}, e_{\mathrm{I}} \in E^{\mathrm{I}}\right\} \\ \tilde{\psi}_{\mathrm{II}}\left(x_{\mathrm{I}}^*, x_{\mathrm{II}}\right) &:= \sup\left\{\phi_{\mathrm{II}}\left(x_{\mathrm{I}}^*, x_{\mathrm{II}}, y_{\mathrm{II}}(x_{\mathrm{I}}^*, x_{\mathrm{II}}, e_{\mathrm{II}}), u_{\mathrm{II}}\right) \mid u_{\mathrm{II}} \in U^{\mathrm{II}}, e_{\mathrm{II}} \in E^{\mathrm{II}}\right\}\end{aligned} \tag{2.15}$$

Definition 2.6.2. *Assume that we have a multi-leader-follower game $\mathcal{G}$ with k followers with the format (*2.13*) and (*2.14*) with loss functions ϕ_{I} and ϕ_{II}. A strategy profile $(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*)$ is called a* robust L/F-(Nash) equilibrium *of $\mathcal{G}$ if $(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*)$ is a L/F-equilibrium of $\tilde{\mathcal{G}}$ with loss functions $\tilde{\psi}_{\mathrm{I}}$ and $\tilde{\psi}_{\mathrm{II}}$ defined in (2.15).*

2.7 Variational Inequality Problems

Variational inequalities offer a general framework for constrained optimization and equilibrium problems. They also naturally capture equilibrium conditions in game-theoretic settings, particularly with continuous strategy spaces.

Definition 2.7.1. *Given a function $F : \mathbb{R}^m \to \mathbb{R}^m$ and a correspondence $\mathcal{R} : \mathbb{R}^m \rightrightarrows \mathbb{R}^m$, an approximate solution to the* quasi-variational inequality *$QVI\,(\mathcal{R}, F)$ is a vector $x^* \in \mathcal{R}(x^*)$ such that:*

$$\left(y - x^*\right)^T F\left(x^*\right) + \epsilon \geq 0, \quad \forall y \in \mathcal{R}\left(x^*\right)$$

Definition 2.7.2. *A* variational inequality *is a special case of the preceding definition when $\mathcal{R}(x)$ is independent of x, say, $\mathcal{R}(x) = \mathcal{R}$ for all x. An approximate solution to the variational inequality $VI\ (\mathcal{R}, F)$ is a vector $x^* \in \mathcal{R}$ such that:*

$$(y - x^*)^T F(x^*) + \epsilon \geq 0, \quad \forall y \in \mathcal{R}$$

Generalized quasi-variational inequalities (GQVI) extend the notion of quasi-variational inequality (QVI). GQVI is a unification of QVI and another notion known as generalized VI [43].

Definition 2.7.3. *Given correspondences $\mathcal{F} : \mathbb{R}^m \rightrightarrows \mathbb{R}^m$ and $\mathcal{R} : \mathbb{R}^m \rightrightarrows \mathbb{R}^m$, an approximate solution to the* generalized quasi-variational inequality *$GQVI\ (\mathcal{R}, \mathcal{F})$ consists of two vectors $x^* \in \mathcal{R}(x^*)$ and $w^* \in \mathcal{F}(x^*)$ such that:*

$$\left(y - x^*\right)^T w^* + \epsilon \geq 0, \quad \forall y \in \mathcal{R}\left(x^*\right)$$

2.7.1 Properties of Quasi Variational Inequalities

The following propositions will play an important part in establishing a relationship between quasi-variational inequalities and Debreu-Rosen style games (see Proposition 2.7.6).

Proposition 2.7.4. *Let x^* be a solution to the following optimization problem where $f : \mathbb{R}^m \to \mathbb{R}$ is continuously differentiable and $\mathcal{R}$ is closed and convex correspondence:*

$$\begin{aligned} &\text{Min } \ f(x) \\ &s.t \ \ x \in \mathcal{R}(x) \end{aligned} \tag{2.16}$$

Then x^* is a solution of the quasi-variational inequality problem ($F = \nabla f$ where ∇ refers to the gradient operator):

$$(y - x^*)^T F(x^*) \geq 0, \quad \forall y \in \mathcal{R}(x^*) \tag{2.17}$$

Proof. Let $\Psi(t) = f(x^* + t(y - x^*))$, for $t \in [0, 1]$. Since $\Psi(t)$ gets its minimum at $t = 0$, $0 \leq \Psi'(0) = (y - x^*)^T \nabla f(x^*)$ which means x^* is a solution of Equation 2.17. □

Proposition 2.7.5. *Suppose that $f : \mathbb{R}^m \to \mathbb{R}$ is a differentiable convex function and x^* is a solution to* QVI$(\nabla f, \mathcal{R})$*, then x^* is a solution to the minimization problem (2.16).*

Proof. By the fact that f is convex:

$$f(y) \geq f(x^*) + (y - x^*)^T \nabla f(x^*), \quad \forall y \in \mathcal{R}(x^*)$$

But $(y - x^*)^T \nabla f(x^*) \geq 0$, since x^* is a solution to QVI$(\nabla f, \mathcal{R})$. Therefore, we can conclude that:

$$f(y) \geq f(x^*), \quad \forall y \in \mathcal{R}(x^*)$$

that is, x^* is a solution (minimum point) for Equation 2.16. □

Proposition 2.7.6. *Suppose that we have a game with concave and continuously differentiable utilities $\mathcal{U} = (u_1, \ldots, u_k)$, strategies $\mathcal{S} = (S_1, \ldots, S_k)$ and constraints $\mathcal{R} = (\mathcal{R}_1, \ldots, \mathcal{R}_k)$ which are closed convex subsets of $\mathcal{S} = (S_1, \ldots, S_k)$. The problem* *of finding an approximate generalized equilibrium of this game can be transformed into finding an approximate solution to a QVI problem*[7]*.*

Proof. We inspect the exact version of the problem without considering the approximation. This version is easier to establish and is also a well-known result in the literature [97, 150]. A more complete proof is available in [119]. Define $\theta_i = -u_i$ for each $i \in [k]$ to denote the *loss* function of each player i. $\theta_i(\cdot, \mathbf{x}_{-i})$ is convex in $\mathbf{x}_{-i}$. Finding a generalized equilibrium in this game is equivalent to finding a solution to the following optimization problem, where the goal is finding $\mathbf{x}^* = (x_i^*, \mathbf{x}_{-i}^*)$ such that $\forall i \in [k]$:

$$\begin{aligned} &\text{Min}\, \theta_i(x_i, \mathbf{x}_{-i}^*) \\ &\text{s.t } x_i \in \mathcal{R}_i(\mathbf{x}_{-i}^*) \end{aligned} \tag{2.18}$$

Define $\mathcal{R}(\mathbf{x}) \equiv \prod_{i=1}^{k} \mathcal{R}_i\left(\mathbf{x}_{-i}\right)$ and $F(\mathbf{x}) = \left(\nabla_{x_i}\theta_i(\mathbf{x})\right)_{i=1}^{k} \in \mathbb{R}^k$. Based on Proposition 2.7.4 and 2.7.5, we can see that x^* is a generalized equilibrium if and only if $\mathbf{x}^* \in \mathcal{R}\left(\mathbf{x}^*\right)$ and it satisfies:

$$\left(\mathbf{y} - \mathbf{x}^*\right)^T F\left(\mathbf{x}^*\right) \geq 0, \quad \forall \mathbf{y} \in \mathcal{R}(\mathbf{x}^*) \tag{2.19}$$

□

Corollary 2.7.7. *Suppose that we have a game with the same conditions on* $\mathcal{U} = (u_1, \ldots, u_k)$*, and strategies* $\mathcal{S} = (S_1, \ldots, S_k)$*. The problem of finding an approximate Nash equilibrium of this game can be transformed into finding an approximate solution to a VI problem.*

2.8 Existence Results and Fixed Point Theorems

For the sake of completeness, we review some well-known existence results related to problems that we study.

2.8.1 Kakutani's Fixed Point Theorem

Kakutani's Theorem generalizes Brouwer's fixed point theorem concerning functions [115] to correspondences.

[7]This proposition is a simple extension of Harker's transformation ([97]) for the approximate case. The proof is available in the Appendix.

Theorem 2.8.1 ([115]). *Suppose S is a non-empty, compact, and convex subset of $\mathbb{R}^n$ and $\mathcal{R} : S \rightrightarrows S$ is a correspondence with the following properties, then $\mathcal{R}$ has a fixed point, i.e., a point $s \in S$ such that $s \in \mathcal{R}(s)$.*

1. *$\mathcal{R}$ is upper semi-continuous on S;*

2. *$\mathcal{R}(s)$ is non-empty and convex for all $s \in S$.*

Definition 2.8.2. *We say that $\mathcal{R}$ has a* closed graph *if $\{(x, y) \mid y \in \mathcal{R}(x)\}$ is a closed subset of $X \times Y$ in the product topology.*

By the *closed graph theorem*, since S is compact, the first condition in Theorem 2.8.1 may be replaced by $\mathcal{R}$ having a closed graph.

2.8.2 Existence of GQVI

The existence of a solution of a generalized quasi-variational inequality is guaranteed by the Eilenberg-Montgomery fixed-point theorem [43].

Proposition 2.8.3 ([43]). *Let $\mathcal{F}$ and $\mathcal{R}$ be correspondences from $\mathbb{R}^n$ into itself. Suppose that there exists a nonempty compact convex set C such that:*

1. *$\mathcal{R}(C) \subseteq C$;*

2. *$\mathcal{F}$ is a nonempty contractible compact valued upper semicontinuous correspondence on C;*

3. *$\mathcal{R}$ is a nonempty continuous convex valued correspondence on C.*

Then there exists a solution to $GQVI(\mathcal{R}, \mathcal{F})$.

Remark 2.8.4. Under the following additional assumption, Kakutani's fixed point theorem may be applied, instead of Eilenberg-Montgomery:

4. *$\mathcal{F}$ is convex valued mapping on C*

Proof. By assumption, $\mathcal{R}(x) \times \mathcal{F}(x)$ is nonempty, compact, and convex because C, $\mathcal{R}(x)$, and $\mathcal{F}(x)$ are assumed to be compact and convex. For $(x, w) \in \mathcal{R}(x) \times \mathcal{F}(x)$, define the function $\phi(y, x, w)$ as $\phi(y, x, w) = -(y - x)^T w$. This function is continuous in (y, x, w), and the concavity of ϕ in $y \in C$ follows from the definition. Now, for each $(x, w) \in \mathcal{R}(x) \times \mathcal{F}(x)$, we also define the correspondence $\Pi(x, w)$ as:

$$\Pi(x, w) = \arg \max_{y \in \mathcal{R}(x)} \phi(y, x, w).$$

By the Maximum Theorem, since $\phi(y, x, w)$ is continuous and concave in y, $\Pi(x, w)$ is nonempty, compact, convex, and upper semi-continuous. Now, define the correspondence $\Psi : \mathcal{R}(x) \times \mathcal{F}(x) \to \mathcal{R}(x) \times \mathcal{F}(x)$ as:

$$\Psi(x, w) = (\Pi(x, w), \mathcal{F}(x)) .$$

The correspondence Ψ is upper semicontinuous, and by the properties of Π and $\mathcal{F}$, Ψ is nonempty, compact-valued, and convex-valued. Now by Kakutani's fixed-point theorem, we can guarantee the existence of a fixed point $(x^*, w^*) \in \mathcal{R}(x^*) \times \mathcal{F}(x^*)$ such that $(x^*, w^*) \in \Psi(x^*, w^*)$. This means that:

$$x^* \in \Pi(x^*, w^*) \quad \text{and} \quad w^* \in \mathcal{F}(x^*).$$

From the definition of $\Pi(x^*, w^*)$, we have:

$$\phi(y, x^*, w^*) \leq \phi(x^*, x^*, w^*) \quad \forall y \in \mathcal{R}(x^*),$$

which is equivalent to:

$$\langle w^*, y - x^* \rangle \geq 0 \quad \forall y \in \mathcal{R}(x^*).$$

Therefore, $x^* \in \mathcal{R}(x^*)$ and $w^* \in \mathcal{F}(x^*)$, giving a solution for $GQVI(\mathcal{R}, \mathcal{F})$.

□

2.8.3 Existence Result for Generalized Nash Equilibrium

A proof of the existence of an equilibrium for bi-matrix games and mixed strategies was provided by Nash in his celebrated papers [144, 145]. Nash's result follows in a fairly direct way from Kakutani's or Brouwer's fixed-point theorem. Debreu's theorem (Theorem 2.8.5) specifies necessary conditions that guarantee the existence of a social equilibrium.

Theorem 2.8.5 ([62])**.** *Consider a (Debreu) game represented by* $\mathcal{U} = (u_1, \ldots, u_k)$ *and* $\mathcal{S} = (S_1, \ldots, S_k)$ *with the constraints imposed on the players* $\mathcal{R} = (\mathcal{R}_1, \ldots, \mathcal{R}_k)$*,*

$1 \le i \le k$, *and suppose that:*

1. u_i *is continuous (for all* $i \in [k]$*);*

2. *For any* $\lambda \in [0,1]$ $s_i, s_i' \in S_i$*, we have* $\boldsymbol{s}_{-i} \in S_{-i}$, $u_i(\lambda s_i + (1-\lambda)s_i', \boldsymbol{s}_{-i}) \ge \min\{u_i(s_i, \boldsymbol{s}_{-i}), u_i(s_i', \boldsymbol{s}_{-i})\}$ *(concavity with respect to i-th variable);*

3. $\mathcal{R}_i$ *is upper and lower semicontinuous for all* $i \in [k]$.

4. *For every* $\boldsymbol{s}_{-i} \in S_{-i}$, $\mathcal{R}_i(\boldsymbol{s}_{-i})$ *is convex and non-empty (convex-valued and non-empty).*

Then, the game specified by $\mathcal{U}$, $\mathcal{S}$, $\mathcal{R}$ *admits a social (generalized) equilibrium.*

2.8.4 Rosen's Contribution and Concave Games

In [161], Rosen gives an alternate approach to specifying generalized equilibrium, and an existence result that guarantees an equilibrium under potentially weaker conditions. Rosen's result applies in a setting where there is a single *commonly coupled constraint*, i.e., a relation $\mathcal{R} \subseteq S_1 \times \cdots \times S_k$ which constrains the best response for all players. In particular, (**) becomes;

$$u_i(\mathbf{s}^*) + \epsilon \ge u_i(s_i, \mathbf{s}^*_{-i}), \qquad \forall s_i \text{ s.t. } \mathcal{R}(s_i, \mathbf{s}^*_{-i}) \tag{$\dagger$}$$

An equilibrium in this case is a strategy profile $s^* \in \mathcal{R}$ that satisfies (†).

Theorem 2.8.1 can be used to guarantee the existence of an equilibrium for these problems, but the required continuity conditions on $\mathcal{R}_i$ are often too strong to be satisfied in many settings. Rosen defines $\phi : \mathcal{R} \times \mathcal{R} \to \mathbb{R}$ by: $\phi(\mathbf{s}, \mathbf{t}) = \sum_{i \in [k]} u_i(\mathbf{t}_i, \mathbf{s}_{-i})$, and the correspondence $\Gamma : \mathcal{R} \rightrightarrows \mathcal{R}$ by:

$$\Gamma(\mathbf{s}) = \{\mathbf{t} \mid \phi(\mathbf{s}, \mathbf{t}) = \max_{\mathbf{r}} \phi(\mathbf{s}, \mathbf{r})\}$$

Theorem 2.8.6 ([161]). *Assume that for* $i \in [n]$, $u_i(\boldsymbol{s})$ *is continuous at every* $\boldsymbol{s} \in \mathcal{S}$ *and for each fixed value of* $\boldsymbol{s}_{-i} \in S_{-i}$, $u_i(s_i, \boldsymbol{s}_{-i})$ *is concave in* s_i*, and that* $\mathcal{R}$ *is convex and compact. Then:*

1. Γ *admits a fixed point* $\boldsymbol{s}^* \in \Gamma(\boldsymbol{s}^*)$.

2. *Any fixed point* $\boldsymbol{s}^*$ *of* Γ*, is an equilibrium in the sense of* (†)

Note that $S_{-i} = (S_1 \times \ldots S_{i-1}, S_{i+1} \ldots S_k)$. It also should be noted that Theorem 2.8.6 guarantees the existence of an equilibrium for games with coupled constraints under reasonable assumptions. Because of their close connection with variational inequalities [71, 97], equilibria obtained in this way have been termed *variational equilibria*, and their properties have been extensively studied (see Proposition 2.7.6). We present the proof only for the sake of completeness and comparison, as one of the techniques used in the proof will be used later.

Proof. First note that ϕ is continuous at every $(\mathbf{s}, \mathbf{t}) \in \mathcal{R} \times \mathcal{R}$, and that for $\mathbf{s} \in \mathcal{R}$, $\phi(\mathbf{s}, \mathbf{t})$ is concave in $\mathbf{t}$. By assumption, $\mathcal{R}$ is compact and it is clear that for any $\mathbf{s} \in \mathcal{R}$, $\Gamma(\mathbf{s})$ is nonempty and convex. So, $\mathcal{R}$ and Γ satisfy the assumptions of Theorem 2.8.1 (Kakutani's fixed point) and there exists $\mathbf{s}^* \in \mathcal{R}$ such that $\mathbf{s}^* \in \Gamma(\mathbf{s}^*)$, which means that:

$$\phi(\mathbf{s}^*, \mathbf{s}^*) = \max_{\mathbf{r} \in \mathcal{R}} \phi(\mathbf{s}^*, \mathbf{r}). \tag{$\dagger\dagger$}$$

The fixed point $\mathbf{s}^*$ must satisfy (†). If not, for some $i \in [n]$ and s_i such that $\mathcal{R}(s_i, \mathbf{s}^*_{-i})$, $u_i(s_i, \mathbf{s}^*_{-i}) > u_i(s^*_i, \mathbf{s}^*_{-i})$. But then, for $\overline{s} = (s_i, \mathbf{s}^*_{-i})$, $\phi(\mathbf{s}^*, \overline{s}) > \phi(\mathbf{s}^*, \mathbf{s}^*)$, contradicting (††). □

2.8.5 Existence Result for Remedial L/F Equilibrium

Theorem 5 (implied by Theorem 2) in [150] provides the required conditions for the existence of a remedial L/F equilibrium[8], and we prove a stronger result, which is inclusion in PPAD.

Theorem 2.8.7 ([150]). *Let X^{I} and X^{II} be nonempty, bounded polyhedra and assume the following conditions:*

1. *For each $(x_{\mathrm{I}}, x_{\mathrm{II}}) \in X^{\mathrm{I}} \times X^{\mathrm{II}}$, the functions $\phi_{\mathrm{I}}(\cdot, x_{\mathrm{II}}, \cdot)$ and $\phi_{\mathrm{II}}(x_{\mathrm{I}}, \cdot, \cdot)$ are convex and continuously differentiable;*

2. *For all $i = 1, \ldots, k$, the function $c_i\left(x_{\mathrm{I}}, x_{\mathrm{II}}, y_{(-i,\mathrm{I})}\right)$ is affine, and the graphs of the four set-valued maps $W_{(i,\mathrm{I})}, V_{(i,\mathrm{I})}, W_{(i,\mathrm{II})}$, and $V_{(i,\mathrm{II})}$ are polyhedra;*

3. *For each $(x_{\mathrm{I}}, x_{\mathrm{II}}) \in X^{\mathrm{I}} \times X^{\mathrm{II}}, Z^{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}})$ and $Z^{\mathrm{II}}(x_{\mathrm{I}}, x_{\mathrm{II}})$ are nonempty;*

4. *$Z^{\mathrm{I}}\left(X^{\mathrm{I}}, X^{\mathrm{II}}\right)$ and $Z^{\mathrm{II}}\left(X^{\mathrm{I}}, X^{\mathrm{II}}\right)$ are bounded.*

Then there exists a remedial L/F equilibrium.

[8] For the formal definition, see Section 4.1.1

2.9 Computational Complexity Definitions

2.9.1 Total Search Problems, Reductions and PPAD

In computational complexity theory, total search problems are those in which a solution is guaranteed to exist for every valid input. Total search problems are central to this thesis because they encompass a wide range of equilibrium problems we study for which the existence of a solution is guaranteed by the application of fixed-point theorems.

Definition 2.9.1. *TFNP (Total Functional Nondeterministic Polynomial Time) is the class of total search problems with solutions that are poly-time verifiable. Formally, given a poly-balanced poly-time relation $R(x, y)$, the associated* NP search problem *is the (partial) multi-valued function $Q(x) = \{y \mid R(x, y)\}$. The problem is* total *if $Q(x)$ is nonempty for all x. The class* FNP *consists of all NP search problems, while TFNP consists of all total NP search problems. A search problem Q is in* FP *if there is a poly-time function Q' such that for all x, $Q'(x) \in Q(x)$. The* decision problem *associated with R is to determine, given x, whether there is some y such that $R(x, y)$. NP is the class of all such decision problems.*

Definition 2.9.2. *We can extend the notion of many-one reduction to total search problems as follows: For two total search problems R and S, we say $R \leq_m S$ if there exist poly-time computable functions f, g such that for all x, y if $(f(x), y) \in S$ then $(x, g(y)) \in R$.*

Remark 2.9.3. This form of reduction between search problems is equivalent to a Cook reduction with one call to the oracle.

The subclass PPAD of TFNP plays a central role in the complexity of the equilibrium problems in game theory and is defined as follows:

Definition 2.9.4. *The* end-of-the-line (EOTL) *problem is defined as follows: Given a directed graph $\mathcal{G}$, represented by two poly-sized circuits that return the predecessor and the successor of a node (represented in binary), where each vertex has at most one predecessor and successor and a vertex u in $\mathcal{G}$ with no predecessor, find another vertex $v \neq u$ with no predecessor or no successor. A total search problem is in* PPAD *if it is many-one reducible to EOTL.*

2.9.2 Approximation Schemes

When considering the computational complexity of additive approximation, we need to take into account the representation of the approximation term ϵ. In particular, we consider the following forms of approximation [169]:

Definition 2.9.5.

- *By* EXACT *we indicate an exact solution to a problem (i.e.,* $\epsilon = 0$*).*
- *For exponential additive approximation or* EXP-approx*, the input* ϵ *is encoded as a (dyadic) rational number.*
- *For polynomial additive approximation or* POLY-approx *the input* $\epsilon = \frac{1}{k}$ *is represented (in unary) as* 1^k*.*

2.9.3 Linear Arithmetic Circuits

Linear arithmetic circuits are a restriction of general arithmetic circuits that do not allow general multiplication gates [70, 76]. Such circuits can efficiently approximate any polynomially computable function; in particular, they can approximate functions represented by well-behaved general circuits. Linear arithmetic circuits have a variety of useful properties, such as Lipschitzness (see Lemma 2.9.8). Given the set of gates in a linear arithmetic circuit, it is differentiable if and only if it represents a linear (affine) function. In general, the presence of min and max gates makes these circuits non-differentiable. However, as noted in [152], automatic differentiation can compute a vector belonging to subgradients, removing the need for a dedicated circuit [17]. Given these properties and their effectiveness in facilitating the analysis[9], they are utilized throughout this research.

Definition 2.9.6. *A* linear arithmetic circuit C *is a circuit represented as a directed acyclic graph with nodes labeled either as input nodes or as output nodes or as gate nodes with one of the following possible gates* $\{+, -, \min, \max, \times\zeta\}$*, where the* $\times\zeta$ *gate refers to the multiplication by a constant. Also, rational constant gates are allowed. We use* $\text{size}(C)$ *to refer to the number of nodes of* C*, also including the constants.*

Next, we restate Theorem E.1 of [76] in which the functions that can be computed by arithmetic circuits can be approximated by linear arithmetic circuits with minimal

[9] Discussion is relegated to chapter 3.

error. An arithmetic circuit f is well-behaved if, on any directed path that leads to output, there are at most $log(size(f))$ true multiplication gates. A true multiplication gate is one where both inputs are non-constant nodes of the circuit. Note that in linear arithmetic circuits, we do not have multiplication gates.

Theorem 2.9.7 ([76]). *Given a well-behaved arithmetic circuit $f : [0,1]^n \to \mathbb{R}^d$, a purported Lipschitz constant $L > 0$, and a precision parameter $\epsilon > 0$, in polynomial time in* $\text{size}(f), \log L$ *and* $\log(1/\epsilon)$*, we can construct a linear arithmetic circuit $F : [0,1]^n \to \mathbb{R}^d$ such that for any $x \in [0,1]^n$ it holds that:*

- $\|f(x) - F(x)\|_\infty \leq \epsilon$, or
- given x, we can efficiently compute $y \in [0,1]^n$ such that:

$$\|f(x) - f(y)\|_\infty > L\|x - y\|_\infty.$$

Here "efficiently" means in polynomial time in $\text{size}(x), \text{size}(f), \log L$ and $\log(1/\epsilon)$.

The following lemma from (Lemma A.1 in [76]) states that any linear arithmetic circuit f mapping $\mathbb{R}^n$ to $\mathbb{R}^m$ is Lipschitz-continuous with respect to the ℓ_∞-norm. The Lipschitz constant for the circuit is bounded by $2^{\text{size}(f)^2}$, where $\text{size}(f)$ is the number of nodes in the circuit. This means that the output of the circuit changes at a controlled rate, depending on its size.

Lemma 2.9.8 ([76]). *Any linear arithmetic circuit $F : \mathbb{R}^m \to \mathbb{R}^m$ is $2^{size(F)^2}$-Lipschitz-continuous (with respect to the ℓ_∞-norm) over $\mathbb{R}^m$.*

Chapter 3

Variational Inequalities

3.1 Organization

The goal of this chapter is to introduce computational versions of the variational inequality problems we defined in the previous chapter. We progressively develop computational formulations of variational inequality problems, culminating in establishing their computational complexity based on Kakutani's fixed-point theorem. Specifically, we formally provide computational frameworks for approximating various variational inequality problems and establishing their PPAD-completeness by leveraging computational geometry techniques (see [93]), carefully adapting the technical contributions of [152]. Our foundational results serve as the basis for analyzing the complexity of more advanced equilibrium problems in Chapter 4. Introducing a very general framework for the variational inequality framework lets us address diverse equilibrium problems within a single, coherent approach, rather than handling each equilibrium problem with separate, problem-specific reasoning. For example, in the literature, the existence and the computational complexity of Nash equilibria are typically derived via the Brouwer fixed-point theorem for functions or the linear complementarity problem (LCP), whereas generalized Nash equilibrium would require dealing with correspondences and Kakutani's fixed-point.

We begin by defining separation oracles, which are standard tools for the computational treatment of convex sets and correspondences. We then introduce computational versions of several fundamental problems in convex optimization—namely, constrained optimization and approximate projection (see also [93])—which serve as essential building blocks for our computational definitions of GQVI, QVI, and VI and also for the final reduction to a computational version of Kakutani's fixed-point theorem which was proven to be PPAD-complete [152]. We further introduce a strictly more general problem, MGQVI (see Section 3.3), which is required for our analysis of t-resilient Nash equilibria, and prove its PPAD-completeness as well.

3.2 Complexity of Variational Inequalities

So far, we have motivated the applicability of variational inequalities by the fundamental relationship between variational inequalities and the equilibrium problems (see Proposition 2.7.6 and 2.7.7). As highlighted by existence results derived from fixed-point theorems (such as Proposition 2.8.3), the convexity of the domain is essential. An efficient representation of the convex domain is also crucial for computational tractability. This requirement naturally motivates the use of separation oracles, which employ separating hyperplanes as certificates to efficiently determine whether a given point belongs to a convex set.

For simplicity and to avoid some cumbersome definitions, we initially investigate strong separation oracles[1]. Next, the computational versions and statements of results are proposed. For the sake of completeness, we provide a short review of the results of [152] and include a direct comparison to our work. Finally, we provide proof for PPAD-completeness. Without loss of generality, we restrict our attention to the metric space $(\mathbb{R}^m, \ell_2)$ where the inputs and the outputs are restricted to a well-bounded compact box denoted by $\mathbb{R}^{m*}$[2].

3.2.1 Using Strong Separation Oracles

Given the discussions above, and the introduction, for addressing the limitations of Kakutani's fixed-point definition[3], the paper [152] assumes an implicit representation of convex sets using a polynomially-sized circuit that computes (weak/strong) separation oracles. While this formulation has been proven useful in many game-theoretic applications, it introduces some technical challenges [152]. We follow the assumption that the sets are bounded by a box that is a subset of $\mathbb{R}^m$ and denote this by $\mathbb{R}^{m*}$.

[1] Discussion of several computational problems for weak separation oracles is relegated to the Appendix B.2. One major problem is that a disparity arises as an unavoidable curse of weak separation oracles. For more information, see Appendix B in [152].

[2] We can extend our results for different norms and any well-bounded box similar to the approach of [152]

[3] See Section 1.2.2 and Theorem E.1 in [152].

A STRONG SEPARATION ORACLE (VIA A CIRCUIT $SO_{\mathcal{R}(x)}$)

Input: A vector $z \in \mathbb{Q}^m \cap \mathbb{R}^{m*}$.

Output: $(a, b) \in \mathbb{Q}^m \times \mathbb{Q}$ such that the threshold $b \in [0, 1] \cap \mathbb{Q}$ denotes the membership of z in $\mathcal{R}(x)$[4]:

- If $z \in \mathcal{R}(x)$ then $b > \frac{1}{2}$ and the vector $a \in \mathbb{Q}^m$ will be $\perp$. In other words, a is meaningful only when $b \leq \frac{1}{2}$
- $b \leq \frac{1}{2}$ and vector a , defines a separating hyperplane $\mathcal{H}(a, z) := \{y \in \mathbb{R}^{m*} : \langle a, y - z \rangle = 0\}$ between the vector z and the set $\mathcal{R}(x)$ such that $\langle a, y - z \rangle \leq 0$ for every $y \in \mathcal{R}(x)$.

The following problems can be solved using the sub-gradient ellipsoid central cut method [152]. We assume that if the function is not provided in an explicit form, we have oracle access to a subgradient or a suitable generalized derivative, enabling optimization algorithms to proceed[5].

STRONG CONSTRAINED CONVEX OPTIMIZATION PROBLEM

Input: A zeroth and first order oracle[6] for the convex function $F : \mathbb{R}^{m*} \to \mathbb{R}$, two rational numbers $\delta, \epsilon > 0$ and a strong separation oracle $\mathrm{SO}_{\mathcal{R}}$ for a non-empty closed convex set $\mathcal{R} \subseteq \mathbb{R}^{m*}$.

Output: One of the following cases:

- (Violation of non-emptiness of points deep inside $\mathcal{R}$): A failure symbol $\perp$ with a polynomial-sized certificate that $\overline{\mathrm{B}}(\mathcal{R}, -\delta) = \emptyset$.
- (Approximate minimization): A vector $z \in \mathbb{Q}^m \cap \mathcal{R}$, with $F(z) \leq \min_{y \in \mathcal{R}} F(y) + \epsilon$.

We also define the strong approximate version of the projection problem, denoted by $\widetilde{\Pi}^{\epsilon}_{\mathcal{X}}(x)$, which is an instance of a strong constrained optimization problem.

[4] This threshold is considered a rational number for the sake of consistency [152].
[5] This is indeed a general and necessary assumption. For more information, see [61, 79, 152].
[6] A zeroth-order oracle returns the function value at a queried point, while a first-order oracle provides both the function value and a sub-gradient or gradient at that point.

STRONG APPROXIMATE PROJECTION PROBLEM $\widetilde{\Pi}^{\epsilon}_{\mathcal{X}}(x)$

Input: A rational number $\epsilon > 0$ and a strong separation oracle $\mathrm{SO}_{\mathcal{X}}$ for a non-empty closed convex set $\mathcal{X} \subseteq \mathbb{R}^{m*}$ and a vector x that belongs to $\mathbb{Q}^m \cap \mathbb{R}^{m*}$.

Output: One of the following cases:

- (Violation of non-emptiness): A failure symbol $\perp$ and a polynomial-sized certificate that $\overline{\mathrm{B}}(\mathcal{X}, -\epsilon) = \emptyset$,
- (Approximate Projection): A vector $z \in \mathbb{Q}^m \cap \mathcal{X}$, such that:

$$\|z - x\|_2^2 \leq \min_{y \in \mathcal{X}} \|x - y\|_2^2 + \epsilon$$

3.2.2 Computational Definition for Generalized Variational Inequalities (GQVI)

In this section, we define the computational version of the generalized quasi-variational inequality problem (GQVI) under convexity assumptions and establish its PPAD-completeness.

The computational complexity of solving variational inequalities is influenced by the problem's dimension, as well as the properties and representations of sets and mappings. We follow the assumption that the sets are bounded by a box that is a subset of $\mathbb{R}^m$ and denote this by $\mathbb{R}^{m*}$. For Lipschitzness of correspondences, we use a more relaxed (algorithmic) version of the concept for correspondences, which utilizes strong approximate projection [152][7].

[7]Details are relegated to Appendix A.2

$GQVI(\mathcal{F}, \mathcal{R})$ With Strong Separation Oracles

---

Input: We receive as input all the following:

- An arithmetic circuit $C_{\mathcal{R}}$ which represents a strong separation oracle for a convex-valued $L_{\mathcal{R}}$-Hausdorff Lipschitz correspondence $\mathcal{R} : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$,
- An arithmetic circuit $C_{\mathcal{F}}$ which represents a strong separation oracle for a convex-valued and $L_{\mathcal{F}}$-Hausdorff Lipschitz correspondence $\mathcal{F} : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$,
- An accuracy parameter β.

Output: One of the following cases:

- (Violation of non-emptiness): A vector $x \in \mathbb{R}^{m*}$ such that $\overline{\mathrm{B}}(\mathcal{R}(x), -\epsilon) = \emptyset$ or $x \in \mathbb{R}^{m*}$ such that $\overline{\mathrm{B}}(\mathcal{F}(x), -\epsilon) = \emptyset$,
- (Violation of $L_{\mathcal{R}}$-Hausdorff Lipschitzness of $\mathcal{R}$): Four vectors $p, q, z, w \in \mathbb{R}^{m*}$ and a constant $\epsilon > 0$ such that $w = \widetilde{\Pi}^{\epsilon}_{\mathcal{R}(q)}(q)$ and $z = \widetilde{\Pi}^{\epsilon}_{\mathcal{R}(p)}(w)$ but $\|z - w\| > L_{\mathcal{R}}\|p - q\| + 3\epsilon$.
- (Violation of $L_{\mathcal{F}}$-Hausdorff Lipschitzness of $\mathcal{F}$) Four vectors $p, q, z, w \in \mathbb{R}^{m*}$ and a constant $\epsilon > 0$ such that $w = \widetilde{\Pi}^{\epsilon}_{\mathcal{F}(q)}(q)$ and $z = \widetilde{\Pi}^{\epsilon}_{\mathcal{F}(p)}(w)$ but $\|z - w\| > L_{\mathcal{F}}\|p - q\| + 3\epsilon$,
- Three vectors (x, x^*, w^*) with $\|x - x^*\|_2^2 \leq \beta$ such that $x^* \in \mathcal{R}(x)$ and $w^* \in \mathcal{F}(x)$ and $(y - x)^T w^* + \beta \geq 0, \quad \forall y \in \mathcal{R}(x)$

Remark 3.2.1. The given approximate solution does not satisfy the property $x^* \in \mathcal{R}(x^*)$. However, we show that this assumption is not needed for the applications we investigate in game theory. We can also impose reasonable assumptions on this problem (GQVI) such as $\mathcal{R}$ being symmetric (i.e, if $y \in \mathcal{R}(x)$, then $x \in \mathcal{R}(y)$), however, this might detract from the generality of the problem. We can consider different problem statements considering similar or more restricting assumptions such as considering the notion of strong convexity (see Section A.3) and assumptions of [28].

Theorem 3.2.2. *The above-mentioned generalized quasi-variational inequality problem $GQVI(\mathcal{F}, \mathcal{R})$ for $\mathrm{SO}_{\mathcal{R}}$ and $\mathrm{SO}_{\mathcal{F}}$ is PPAD-complete.*

Proof Sketch. The complete proof of inclusion in PPAD is relegated to Section 3.2.4. Our approach to establishing PPAD membership for variational inequality problems leverages the computational version of Kakutani's fixed point theorem and the robust version of Berge's maximum theorem of [152]. Our approach, while similar to that of [152], has key adaptations[8]. Specifically, we utilize the proof of the existence of Generalized Quasi-Variational Inequalities (GQVI) provided in [43] instead of relying on Rosen's existence theorem [161]. Considering the technical difficulty of the framework proposed in [152], a careful and rigorous analysis is essential to address these adaptations[9]. Note that the paper [43] uses the Eilenberg-Montgomery fixed-point theorem to establish the existence of GQVI, while we instead use Kakutani's fixed-point theorem (see Proposition 2.8.3). PPAD-hardness follows directly from the PPAD-hardness of problems such as Nash equilibrium (see Corollary 2.7.7), which are easily formulated using the simpler framework of variational inequalities (VI). □

3.2.3 Special Cases: Quasi Variational Inequalities (QVI) and Variational Inequalities (VI)

Below, we define a computational variant of QVI that takes as input a strong separation oracle and a circuit representing a function, unlike GQVIs that take two correspondences.

[8]A direct and complete comparison is provided in Section 3.2.4

[9]In addition to using separation oracles rather than conventional approaches such as using polytopes, this work addresses several challenges, such as carefully handling Lipschitz continuity syntactically.

$QVI(F,\mathcal{R})$ With a strong Separation Oracle

Input: We receive as input all the following:

- A circuit $C_{\mathcal{R}}$ which represents a strong separation oracle for a convex valued and $L_{\mathcal{R}}$-Hausdorff Lipschitz correspondence $\mathcal{R} : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$,
- A circuit C_F which represents a L_F-Lipschitz function $F : \mathbb{R}^{m*} \to \mathbb{R}^{m*}$,
- An accuracy parameter β.

Output: One of the following cases:

- (Violation of non-emptiness): A vector $x \in \mathbb{R}^{m*}$ such that $\overline{\mathrm{B}}(\mathcal{R}(x), -\epsilon) = \emptyset$,
- (Violation of $L_{\mathcal{R}}$-Hausdorff Lipschitzness of $\mathcal{R}$): Four vectors $p, q, z, w \in \mathbb{R}^{m*}$ and a constant $\epsilon > 0$ such that $w = \widetilde{\Pi}^{\epsilon}_{\mathcal{R}(q)}(q)$ and $z = \widetilde{\Pi}^{\epsilon}_{\mathcal{R}(p)}(w)$ but $\|z - w\| > L_{\mathcal{R}}\|p - q\| + 3\epsilon$,
- (Violation of L_F-Lipschitzness of F): Two vectors $p, q \in \mathbb{R}^{m*}$ such that $\|F(p) - F(q)\| > L_F\|p - q\|$,
- Two vectors x^* and x with $\|x - x^*\| \leq \beta$ such that $x^* \in \mathcal{R}(x)$ and $(y - x)^T F(x^*) + \beta \geq 0, \quad \forall y \in \mathcal{R}(x)$

Proposition 3.2.3. *The quasi-variational inequality (QVI) problem for a strong separation oracle* $\mathrm{SO}_{\mathcal{R}}$ *is PPAD-complete.*

Proof. Inclusion in PPAD follows by Theorem 3.2.2. PPAD-hardness of this problem is implied by the hardness of VI, which in turn is implied by the PPAD-hardness of Nash equilibrium.

□

We next introduce a computational version of VI. For this problem, the input consists of a strong separation oracle for a set and a circuit representing a function.

$VI(F)$ WITH A STRONG SEPARATION ORACLE

Input: We receive as input all the following:

- An arithmetic circuit $C_{\mathcal{R}}$ which represents a strong separation oracle for a non-empty, compact, and convex set $\mathcal{R}$,
- A arithmetic circuit C_F which represents a L_F-Lipschitz function $F : \mathbb{R}^{m*} \to \mathbb{R}^{m*}$,
- An accuracy parameter β.

Output: One of the following cases:

- (Violation of non-emptiness): A symbol $\perp$ with a polynomial-sized witness that certifies $\overline{\mathrm{B}}(\mathcal{R}, -\epsilon) = \emptyset$,
- (Violation of L_F-Lipschitzness of F): Two vectors $p, q \in \mathbb{R}^{m*}$ such that $\|F(p) - F(q)\| > L_F\|p - q\|$,
- One vector $x^* \in \mathcal{R}$ such that $(y - x^*)^T F(x^*) + \beta \geq 0, \quad \forall y \in \mathcal{R}$

Theorem 3.2.4. *The variational inequality problem (VI) for a strong separation oracle* $\mathrm{SO}_{\mathcal{R}}$ *is PPAD-complete.*

Proof. Inclusion in PPAD simply is implied by Theorem 3.2.2. PPAD-hardness of this problem is implied by the hardness of approximate Nash equilibrium. Finding a mixed (a well-known concept in game theory) Nash equilibrium in two-player games is a special case that is shown to be PPAD-hard [45] where maximizing (or equivalently minimizing the loss) the expected payoff is the optimization goal. The rest of the proof is similar to Corollary 2.7.7 for differentiable utility functions. □

Remark 3.2.5. To satisfy properties such as Lipschitz continuity syntactically, we may assume inputs are provided as linear arithmetic circuits. As discussed in [76], we do not lose representation power (see the previous chapter).

3.2.4 Proof of Inclusion in PPAD of GQVI

This section is organized as follows. First, we present additional definitions and theorems from [152] that we use directly, indirectly, or with modifications. Specifically, for strong separation oracles, in addition to the definition of the strong constrained convex optimization and the strong approximate projection problem, we provide the computational definition of Kakutani's fixed point, finding an approximate generalized equilibrium in concave games and their PPAD-completeness statement. Also, we provide the robust version of Berge's maximum theorem, which provides a more robust Lipschitz continuity for the maximum of a Lipschitz function and is essential in the transformation of the generalized equilibrium problem to Kakutani's fixed point. Finally, we state and prove the PPAD-completeness result.

Computational Kakutani's Fixed Point Problem

We define the computational version of Kakutani's fixed point problem and restate the results of [152]. The definition of the computational problem of finding a Kakutani fixed point is given as follows:

KAKUTANI WITH A STRONG SEPARATION ORACLE (VIA $C_{\mathcal{R}(x)}$)

Input: A circuit $C_{\mathcal{R}}$ that represents strong separation oracle for a L-Hausdorff Lipschitz correspondence $\mathcal{R} : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$ and an accuracy parameter α.

Output: One of the following cases:

- (Violation of non-emptiness): A vector $x \in \mathbb{R}^{m*}$ such that $\overline{\mathrm{B}}(\mathcal{R}(x), -\epsilon) = \emptyset$.
- (Violation of L-Hausdorff Lipschitzness) Four vectors $p, q, z, w \in \mathbb{R}^{m*}$ and a constant $\epsilon > 0$ such that $w = \widetilde{\Pi}^{\epsilon}_{\mathcal{R}(q)}(q)$ and $z = \widetilde{\Pi}^{\epsilon}_{\mathcal{R}(p)}(w)$ but $\|z - w\| > L\|p - q\| + 3\epsilon$.
- Vectors $x, z \in \mathbb{R}^{m*}$ such that $\|x - z\| \leq \alpha$ and $z \in \mathcal{R}(x) \Leftrightarrow d(x, \mathcal{R}(x)) \leq \alpha$.

Theorem 3.2.6 ([152])**.** *The Kakutani fixed-point problem for a strong separation oracle is PPAD-complete.*

Remark 3.2.7. The convexity of $\mathcal{R}$ (i.e., $\mathcal{R}$ is a convex-valued correspondence) is adequate for ensuring the uniqueness of the nearest point (for more information, see [152]).

Using Linear Arithmetic Circuits

As mentioned, linear arithmetic circuits are proven to have useful properties such as Lipschitzness, and they are general enough as we do not lose representation power (see the discussion in Section 2.9.3). In general, when we are restricted to queries to a zeroth and first-order oracle or when inspecting a circuit or Turing machine description, evaluating consistency between function values and their gradients poses computational challenges[10]. The proofs in [152] utilize instances where this consistency can be assured syntactically by using linear arithmetic circuits.

For a computational version of Kakutani's fixed-point, an essential requirement that we need to impose on these separation oracles is that the output, which is a separating hyperplane, should have polynomial bit complexity with respect to the given input parameters[11]. The absence of this constraint enables the construction of a malicious oracle that consistently returns separating hyperplanes with exponentially large bit complexity, making algorithms such as the Ellipsoid method ineffective regardless of optimization efforts. Similar to [152], we will utilize these circuits for defining versions of the problems, considering these benefits.

Robust Berge's Maximum Theorem

The paper [152] leverages the following quantified version of Berge's theorem for general convex functions.

Definition 3.2.8. *A real-valued correspondence $\mathcal{F}$ on Euclidean space satisfies Holder's condition, or is (q, p)-Holder continuous (where $q > 0$ and $p > 1$), when there are real positive constants κ and c such that $d_H(\mathcal{F}(y_1), \mathcal{F}(y_2)) \leq \kappa \|(y_1) - (y_2)\|_p^q + c$.*

Theorem 3.2.9 ([152])**.** *(Robust Berge Maximum Theorem). Let $A \subseteq \mathbb{R}^n$ and $B \subseteq \mathbb{R}^m$. Consider a continuous function $f : A \times B \to \mathbb{R}$ that is μ-strongly concave $\forall a \in A$, L-Lipschitz in $A \times B$ and also a L'-Hausdorff Lipschitz, non-empty, convex-set, compact-valued correspondence $g : A \rightrightarrows B$. Define $f^*(a) = \max_{b \in g(a)} f(a, b)$ and $g^*(a) = \arg\max_{b \in g(a)} f(a, b)$. Then, we observe f^* is continuous and g^* is upper*

[10] Consistency of function values and its gradients" means that the gradients (or sub-gradients) provided correspond correctly to the function values at given points — in other words, the gradient information accurately reflects the local behavior of the function. For more information, see Appendix A in [152]

[11] The bit complexity of a separating hyperplane is the number of bits needed to represent its normal vector a and threshold b as rational numbers, which we require to be polynomial in the input size.

semi-continuous and single-valued, i.e., continuous. Furthermore, f^* and g^* are Lipschitz and $\left(L' + 2\sqrt{\frac{4}{\mu}}\sqrt{(L + L \cdot L')}\right) - (1/2, p)$*-Holder continuous respectively (for sufficiently small differences).*

Strongly Concave Games

The paper [152] establishes PPAD-completeness of an instance of games introduced in the celebrated work of Rosen [161] (also called *concave games* [57, 152, 161]).

Theorem 3.2.10 ([152]). *The Strongly Concave Games Problem (defined below) for a strong separation oracle is PPAD-complete.*

Proof Sketch. The complete proof is available in paper [152] (Theorem E.1). The difference (compared to Rosen's original proof which is also available in Section 2.8.4) lies in how ϕ is defined:

$$\phi(\mathbf{x}, \mathbf{y}) = \sum_{i \in [n]} u_i(y_i, \mathbf{x}_{-i}) - \gamma \cdot \|\mathbf{y}\|_2^2$$

For this proof, adding $-\gamma \cdot \|\mathbf{y}\|_2^2$ ensures that ϕ stays a 2γ-strongly concave function in y. This property is one of the necessary assumptions for the robust Berge's maximum theorem.

Utilizing the properties of $\phi(x, y)$, the proof of [152] constructs a correspondence Γ, proves the non-emptiness and existence of a separation oracle, and passes this correspondence as an input to Kakutani's fixed point. Specifically, the separation oracle for $\Gamma(\mathbf{x}) = \{\mathbf{y} \mid \phi(\mathbf{s}, \mathbf{y}) \geq \max_{\mathbf{r} \in S} \phi(\mathbf{x}, \mathbf{r}) - \epsilon\}$ is constructed, and Lipschitzness is proven by utilizing the robust version of Berge's maximum theorem. This is done in polynomial time by using the constrained convex optimization problem (given a separation oracle that represents the convex set S, which represents the set of strategies that satisfy a commonly coupled constraint). The reduction concludes by proving a solution provided by Kakutani's fixed point given Γ as the input can form an approximate solution to the generalized Nash equilibrium problem in concave games using some techniques that are similar to Rosen's proof. We will use this approach later in the proof of inclusion in PPAD of variational inequality problems. However, we utilize the proof of the existence of variational inequalities provided by [43] instead of the techniques provided by Rosen, which require a meticulous analysis of approximation errors. □

STRONGLY CONCAVE GAMES PROBLEM WITH SO

Input: We receive as input all the following:

- k arithmetic circuits representing the utility functions $(u)_{i=1}^k$ for all k players,
- A Lipschitzness parameter L, a strong concavity parameter μ, and accuracy parameter ϵ,
- $\mathcal{S} = \Pi_{i=1}^k S_i$ that is a convex set called *the strategy domain* where S_i represent the strategy domain for each player i,
- An arithmetic circuit representing a strong separation oracle for the common constraint $\mathcal{R}$ that is a non-empty, convex, and compact subset of $\mathcal{S}$,
- Accuracy parameter ϵ.

Question: We output one of the following:

- (Violation of non-emptiness): A certificate that $\overline{\mathrm{B}}(\mathcal{R}, -\epsilon) = \emptyset$.
- (Violation of Lipschitz continuity): A certification that there exist at least two vectors $x, y \in \mathcal{S}$ and an index $i \in [n]$ such that $|u_i(x) - u_i(y)| > L\|x - y\|$.
- (Violation of strong concavity): An index $i \in [n]$, three vectors $x_i, y_i \in S_i, \mathbf{x}_{-i} \in S_{-i} = \Pi_{j=1,j\neq i}^k S_j$ and a number $\mu \in [0, 1]$ such that:

$$u_i\left(\lambda x_i + (1-\lambda)y_i, \mathbf{x}_{-i}\right) < \lambda u_i\left(x_i, \mathbf{x}_{-i}\right) + (1-\lambda)u_i\left(y_i, \mathbf{x}_{-i}\right) + \frac{\lambda(1-\lambda)}{2}\mu \cdot \|(x_i, \mathbf{x}_{-i}) - (y_i, \mathbf{x}_{-i})\|_2^2$$

- An ϵ-approximate generalized equilibrium in the sense of (†).

A Direct Comparison to [152]

Given the discussion above, we are now ready to compare our proof of inclusion in PPAD to the approach of [152]:

We will use the definition of the strong approximate projection problem from Section 3.2 with no modification in the computational definition of the problems we provide. For this research, especially inclusion in PPAD for resilient Nash equilibrium, we require a different version of the strong constrained convex optimization problem (we define two different versions for two main applications we study in this paper, which we call “modified strong constrained convex optimization problem") with different inputs and outputs. As mentioned, the strong constrained convex optimization problem can be solved using the subgradient ellipsoid central cut algorithm. However, for the modified versions we define, for every single problem, namely, multi-leader-followers and resilient Nash equilibrium, we need to modify the sub-gradient ellipsoid central cut algorithm to make it appropriate for different inputs and outputs (especially for violation of emptiness exceptions)[12]. For multi-leader-follower games, the modifications are easier, while for resilient Nash equilibrium, we need to do more modifications, such as a different version of Berge's maximum theorem.

Finally, we can directly apply Kakutani's fixed point in the proof of inclusion in PPAD for GQVIs. However, the construction of the separation oracle for the final correspondence is more complex because we apply the techniques of [152] to the reduction provided by [43], and the existence of such a separation oracle needs careful use of properties of convex sets and their separation oracles. Moreover, for a resilient Nash equilibrium, providing suitable inputs for Kakutani's fixed point problem and bounding the approximation error for the desired t-resilient Nash equilibrium is quite tedious, given the simple format of the inputs of the t-resilient Nash equilibrium problem and the exponentially many constraints in t (which is a constant) that need to be satisfied. In conclusion, a significant change in the format of the inputs will be required, and we provide a different version of GQVI that handles more complex inputs that are a combination of vectors and matrices[13].

Proof of Inclusion of GQVI in PPAD

Finally, we are now ready to proceed to the proof of Theorem 3.2.2.

[12]Appendix B.3 provides all the details.

[13]For more information, see the definition of MGQVI (Section 3.3).

Proof. Without loss of generality, we can consider $[-1,1]^m$ instead of $\mathbb{R}^{m*}$ (similar to [152]). The proof is organized as follows. First, we define the following function and correspondence[14]:

$$\Phi(y,x,w) = -(y-x)^T w - \gamma(\|y\|_2^2)$$

$$\Pi(x,w) = \{y \in \mathcal{R}(x) \mid \Phi(y,x,w) > max_{y\in\mathcal{R}(x)}\Phi(y,x,w) - \epsilon\}$$

We will show that for constant κ', the correspondence $\Psi(x,w) = (\Pi(x,w), \mathcal{F}(x))$ is κ'-Hausdorff Lipschitz continuous in (x,w) utilizing Berge's maximum theorem. Then, we construct a strong separation oracle for Ψ. Finally, we show that an approximate Kakutani's fixed point of this correspondence will provide an approximate solution to the given GQVI.

To show that Ψ is κ'-Lipschitz continuous, we show that $d_H(\Psi(x_1,w_1),\Psi(x_2,w_2)) \leq \kappa'\|(x_1,w_1)-(x_2,w_2)\|_p^q + c$ for some constants q,p and c. First, similar to Theorem 3.2.10, to leverage Berge's maximum theorem, we need to ensure that for some μ, μ-strong concavity holds for Φ. Here $\Phi(y,x,w)$ is (2γ)-strongly concave in y. This follows by the fact that $-(y-x)^T w$ is 0-strongly concave function of y and also $-\gamma(\|y\|_2^2)$ is 2γ-strongly concave function of y by Lemma A.1.2. In addition, Φ is G-Lipschitz continuous where $G = m + 2\gamma m$ for any pair (x,w).

The procedure to show that $\Pi(x,w)$ is Lipschitz continuous in x and w is similar to Theorem E1 of [152]. Let $H(x,w) = \{y \in \mathcal{R}(x) \mid \Phi(y,x,w) = max_{y\in\mathcal{R}(x)}\Phi(y,x,w)\}$. By optimality KKT conditions for maximization of a concave function with respect to the constraint set $\mathcal{R}(x)$ for all $y \in \mathcal{R}(x)$ we have:

$$\partial\Phi(y,x,w)^T(y^* - y) \geq 0 \text{ where } y^* = argmax_{y\in\mathcal{R}(x)}\Phi(y,x,w)$$

Furthermore, (2γ)-strong-concavity of $\Phi(\cdot,x,w)$ results the following inequality:

$$\Phi\left(y^*,x,w\right) - \Phi(y,x,w) \geq \partial\Phi\left(y^*,x,w\right)^\top\left(y^*-y\right) + \gamma\left\|y-y^*\right\|_2^2$$

Finally, by combining the previous inequalities, we have:

$$\Phi\left(y^*,x,w\right) - \Phi(y,x,w) \geq \gamma\left\|y-y^*\right\|_2^2$$

[14] We will deal with the constants later.

In conclusion, for $y^* \in H(x, w)$ and $y \in \Pi(x, w)$, we have the following inequality:

$$\gamma \|y - y^*\|_2^2 \leq \epsilon \text{ or equivalently } \|y - y^*\|_2 \leq \sqrt{\frac{\epsilon}{\gamma}} \tag{3.1}$$

Recall that we showed $\Phi(\cdot, x, w)$ is strongly concave. Now, we can apply Theorem 3.2.9 (Berge's Maximum Theorem) to $f((x, w), y) = \Phi(y, x, w)$ and $g((w, x)) = \mathcal{R}(x)$, where g is a non-empty, compact, convex-valued correspondence.

By Theorem 3.2.9, there exists a constant $\kappa = \left(L_{\mathcal{R}} + 2\sqrt{\frac{2}{\gamma}}\sqrt{(G + G \cdot L_{\mathcal{R}})}\right)$ such that H is κ-Lipschitz (and also Holder) continuous and also is single-valued. In conclusion, we have the following:

$$d_H(H(x_1, w_1), H(x_2, w_2)) \leq \kappa \|(x_1, w_1) - (x_2, w_2)\|_2^{\frac{1}{2}} \tag{3.2}$$

Combining Equations 3.1 and 3.2 results in:

$$\begin{aligned} d_H(\Pi(x_1, w_1), \Pi(x_2, w_2)) &\leq d_H(H(x_1, w_1), H(x_2, w_2)) \\ &+ d_H(\Pi(x_1, w_1), H(x_1, w_1)) \\ &+ d_H(\Pi(x_2, w_2), H(x_2, w_2)) \end{aligned} \tag{3.3}$$

In conclusion, we have the following:

$$d_H(\Pi(x_1, w_1), \Pi(x_2, w_2)) \leq \kappa \|(x_1, w_1) - (x_2, w_2)\|_2^{\frac{1}{2}} + 2\sqrt{\frac{\epsilon}{\gamma}} \tag{3.4}$$

Since $\mathcal{F}$ is $L_{\mathcal{F}}$-Hausdorff Lipschitz, by the definition of Ψ we can deduce that there exists a constant $\kappa' = \max(\kappa, L_{\mathcal{F}})$:

$$d_H(\Psi(x_1, w_1), \Psi(x_2, w_2)) \leq \kappa' \|(x_1, w_1) - (x_2, w_2)\|_2 + 2\sqrt{\frac{\epsilon}{\gamma}} \tag{3.5}$$

This shows that Ψ is Hausdorff Lipschitz continuous. Next, to establish a reduction employing the computational variant of Kakutani's problem, it is essential to establish the non-emptiness of the correspondence $\Psi(x, w)$ (i.e, $\overline{\mathrm{B}}(\Psi(x, w), -\epsilon) \neq \emptyset$). We know that $\overline{\mathrm{B}}(\mathcal{R}(x), -\epsilon)$ and also $\overline{\mathrm{B}}(\mathcal{F}(x), -\epsilon)$ are nonempty otherwise, we output violation of emptiness. Let $y_x^* = argmax_{y \in \mathcal{R}(x)} \Phi(y, x, w)$. Then, we define a none-empty region $V_x = \overline{\mathrm{B}}(y_x^\star, +\frac{\epsilon}{G}) \cap \mathcal{R}(x) = \overline{\mathrm{B}}(y_x^\star, +\frac{\epsilon}{G})$ where G is the Lipschitz constant of Φ. Moreover, $\overline{\mathrm{B}}(V_x, \frac{-\epsilon}{G})$ is non-empty due to the fact that $\overline{\mathrm{B}}(\mathcal{R}(x), -\epsilon)$ is non-empty for any x. By

Lipschitzness of $\Phi(\cdot, x, w)$, we have:

$$\forall y \in V_x, \quad |\Phi(y, x, w) - \Phi(y_x^*, x, w)| \leq G\frac{\epsilon}{G} = \epsilon$$

In conclusion, $\Phi(V_x, x, w) \subseteq [\Phi(y_x^*, x, w) - \epsilon, \Phi(y_x^*, x, w)]$ is non-empty. Considering the definition of Π, y_x^* and V_x, we can conclude that $\Phi(V_x, x, w) \subseteq \Pi(x, w)$ (see the definition of Π). We also know that for any $x \in \mathcal{R}(x)$, $\overline{\mathrm{B}}(\mathcal{F}(x), -\epsilon)$ is non-empty, then, we can similarly find a non-empty region C_x such that $C_x \subseteq \overline{\mathrm{B}}(\mathcal{F}(x), -\epsilon)$. This is sufficient to show that for all x and w, the correspondence $\Psi(x, w)$ does not violate the non-emptiness given the non-emptiness guarantees of $\mathcal{R}(x)$ and $\mathcal{F}(x)$.

Now we proceed to constructing a strong separation oracle for $\Psi(x, w)$ leveraging a slightly modified version of the Strong Constrained Convex Optimization framework stated in the following:

MODIFIED STRONG CONSTRAINED CONVEX OPTIMIZATION PROBLEM

Input: A zeroth and first order oracle for the concave function $G : \mathbb{R}^{m*} \to \mathbb{R}$, two rational numbers $\delta, \epsilon > 0$ and a strong separation oracle $\mathrm{SO}_{\mathcal{R}}$ for a non-empty closed convex-valued correspondence $\mathcal{R} : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$ and one input x.

Output: One of the following cases:

- (Violation of non-emptiness) A failure symbol $\perp$ with a polynomial-sized witness that certifies that $\overline{\mathrm{B}}(\mathcal{R}(x), -\delta) = \emptyset$,
- (Approximate Maximization) A vector $z \in \mathbb{Q}^m \cap \mathcal{R}(x)$, such that $G(z) + \epsilon \geq \max_{y \in \mathcal{R}(x)} G(y)$ for all values.

The problem mentioned above still can be solved in polynomial time using the sub-gradient ellipsoid central cut method given the separation oracle as we only changed the input while we know that for any x, $\mathcal{R}(x)$ is a convex set as the correspondence $\mathcal{R}$ is convex-valued (for more information, see Appendix B.3 and also Appendix G in [152])[15]. We can compute a solution $y^* \in \mathcal{R}(x)$ using the sub-gradient ellipsoid central cut method such that $\Phi(y^*, x, w) \geq max_{y \in \mathcal{R}(x)} \Phi(y, x, w) - \epsilon$. Consider a

[15]Note that the modification in this problem is that it is a maximization problem.

strong separation oracle for the following set:

$$\overline{\Psi}_s(x,w) = \{(y,w) \in (\mathcal{R}(x), \mathcal{F}(x)) \mid -\Phi(y,x,w) \leq \sigma\}$$

where $\sigma = -\Phi(y^*,x,w)$. In other words, we are looking for $(y,w) \in (\mathcal{R}(x), \mathcal{F}(x))$ such that $(y-x)^T w + \gamma\|y\|_2^2 \geq (y^*-x)^T w + \gamma\|y^*\|_2^2$ given the strong separation oracles representing $\mathcal{F}$ and $\mathcal{R}$. This separation oracle exists due to Theorem H.3 in [152], which was also used for the proof of inclusion in PPAD of Walrasian equilibrium. Thus, using a careful adaptation of the techniques in [152], the existence of a strong separation oracle is achievable in polynomial time for Ψ.

Now, we can give the constructed separation oracle for Ψ as input to the computational Kakutani problem with accuracy parameter $\alpha = \frac{\epsilon'}{\kappa'}$ where $\epsilon' = \epsilon h$. The output of this Kakutani instance will be two points $(x,w) \in ([-1,1]^m, [-1,1]^m)$ and $z = (x^*, w^*) \in \Psi(x,w)$ where $\|(x,w) - (x^*,w^*)\| \leq \frac{\epsilon'}{\kappa'}$ and $d((x,w), \Psi(x,w)) \leq \frac{\epsilon'}{\kappa'}$. Considering the definition of Ψ, we can infer that $w^* \in \mathcal{F}(x)$ and $d(w, \mathcal{F}(x)) \leq \frac{\epsilon'}{\kappa'}$. In addition, $x^* \in \Pi(x,w)$ and $d(x, \Pi(x,w)) \leq \frac{\epsilon'}{\kappa'}$. By the definition of Π, for every $y \in \mathcal{R}(x)$, $\Phi(x^*,x,w) \geq \Phi(y,x,w) - \epsilon$. In conclusion:

$$(y-x)^T w + \gamma\|y\|_2^2 \geq (x^*-x)^T w + \gamma\|x^*\|_2^2 - \epsilon, \ \forall y \in \mathcal{R}(x)$$

So, recalling that $x, y, w \in [-1,1]^m$:

$$(y-x)^T w \geq \pm\frac{2\sqrt{m}\epsilon'}{\kappa'} \pm 4\gamma - \epsilon, \quad \forall y \in \mathcal{R}(x)$$

Finally, knowing that $\|(x,w) - (x^*,w^*)\|_2^2 \leq \frac{\epsilon'}{\kappa'}$ implies that $\|w - w^*\|_2^2 \leq \frac{\epsilon'}{\kappa'}$. We can easily conclude:

$$(y-x)^T w^* \geq \pm\frac{2\sqrt{m}\epsilon'}{\kappa'} \pm 4\gamma \pm \frac{\epsilon'}{\kappa'}\|y-x\|_2 - \epsilon, \quad \forall y \in \mathcal{R}(x)$$

Since w, w^*, x, y are limited to the box constraint $[-1,1]^m$. We can conclude that:

$$(y-x)^T w^* \geq \pm\frac{4\sqrt{m}\epsilon'}{\kappa'} \pm 4\gamma - \epsilon, \quad \forall y \in \mathcal{R}(x)$$

Considering appropriate numbers for ϵ, h (recall that $\epsilon' = \epsilon h$) and γ results the following for the desired approximation error β (note that $w^* \in \mathcal{F}(x)$):

$$(y-x)^T w^* + \beta \geq 0, \quad \forall y \in \mathcal{R}(x)$$

□

Remark 3.2.11. Note that we assumed that both $\mathcal{F}$ and $\mathcal{R}$ are convex-valued. The convexity condition on $\mathcal{F}$ is important for the construction of the separation oracle for Ψ, while the convexity of $\mathcal{R}$ is also necessary for the modified strong constrained optimization problem constraint. Indeed, Ψ must be a convex-valued otherwise membership of an element in Ψ cannot be assessed by using the strong separation oracle framework, which is necessary for the computational Kakutani's fixed point problem.

3.3 A More Generalized Form of GQVI

We need to introduce a more general version of variational inequalities appropriate for t-resilient Nash equilibrium and address the existence of a solution for them. Informally, this stems from the fact that for a t-resilient Nash equilibrium, we need to check exponentially many conditions in the constant t (see the definition of t-resilient Nash equilibrium) and impose appropriate constraints on the variational inequality problem to reflect the constraints of resilient Nash equilibrium. First, we extend the definition of multi-convexity to correspondences.

Definition 3.3.1 (Multi-Convexity of a Correspondence)**.** *Let $\mathcal{F} = (\mathcal{F}_1, \mathcal{F}_2, \ldots, \mathcal{F}_r)$ be a correspondence where each $\mathcal{F}_i(x)$ is a correspondence and for $x \in X$, $\mathcal{F}_i(x)$ is a set. We say that $\mathcal{F}$ is t-multi-convex if for any subset $J \subseteq [t]$ with $|J| \leq t$, and for any $y_1, y_2 \in F(x)$, the convex combination of y_1 and y_2 with respect to J satisfies:*

$$F_J(\lambda y_J^1 + (1-\lambda) y_J^2, x_{-J}) \supseteq \lambda \mathcal{F}_J(y_J^1, x_{-J}) + (1-\lambda) F_J(y_J^2, x_{-J}),$$

for all $\lambda \in [0,1]$. Similarly, in this definition, $\mathcal{F}_J(x) = \prod_{j \in J} \mathcal{F}_j(x)$, and x_{-J} denotes the components of x that are not indexed by J. This condition ensures that the correspondence $\mathcal{F}$ is convex in the components indexed by J, while keeping other components fixed.

To guarantee the existence of a solution, we will also need to have the following assumption:

Definition 3.3.2. *Sub-gradient Compatibility Let $\mathcal{F} = (\mathcal{F}_1, \mathcal{F}_2, \ldots, \mathcal{F}_r)$ be a correspondence where each $\mathcal{F}_i(x)$ is a correspondence and for $x \in X$, $\mathcal{F}_i(x)$ is a set. The correspondence $\mathcal{F}$* Sub-gradient Compatibility *has if there exists a convex function $g : S \to \mathbb{R}$, such that:*

$$\partial g(\mathbf{x}) \subseteq \bigcap_{i=1}^{k} \mathcal{F}_i(\mathbf{x}), \quad \forall \mathbf{x} \in S$$

For completeness and brevity, we only investigate the most complicated case. Computational MQVI and MVI can be defined similarly. The computational problem (MGQVI) with strong separation oracles is defined as follows [16]:

[16] To avoid repetition, we do not investigate this problem for weak separation oracles.

$MGQVI(\mathcal{F}, \mathcal{R})$ With strong Separation Oracles

---

Input: We receive as input all the following:

- A circuit $C_{\mathcal{R}}$ which represents a strong separation oracle for a $L_{\mathcal{R}}$-Hausdorff Lipschitz correspondence $\mathcal{R} : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$,
- A multi-convex correspondence $\mathcal{F} = (\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_r)$ with sub-gradient compatibility which has r circuits such that for each $i \in [r]$, $C_{\mathcal{F}_i}$ represents a strong separation oracle for a $L_{\mathcal{F}}$-Hausdorff Lipschitz and convex-valued correspondence $\mathcal{F}_i : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$,
- An accuracy parameter β.

Output: One of the following cases:

- (Violation of non-emptiness): A vector $x \in \mathbb{R}^{m*}$ such that $\overline{\mathrm{B}}(\mathcal{R}(x), -\epsilon) = \emptyset$ or an index i such that for some $x \in \mathbb{R}^{m*}$ we have that $\overline{\mathrm{B}}(\mathcal{F}_i(x), -\epsilon) = \emptyset$,
- (Violation of $L_{\mathcal{R}}$-Hausdorff Lipschitzness of $\mathcal{R}$): Four vectors $p, q, z, w \in \mathbb{R}^{m*}$ and a constant $\epsilon > 0$ such that $w = \widetilde{\Pi}^{\epsilon}_{\mathcal{R}(q)}(q)$ and $z = \widetilde{\Pi}^{\epsilon}_{\mathcal{R}(p)}(w)$ but $\|z - w\| > L_{\mathcal{R}}\|p - q\| + 3\epsilon$.
- (Violation of $L_{\mathcal{F}}$-Hausdorff Lipschitzness of $\mathcal{F}$): Four vectors $p, q, z, w \in \mathbb{R}^{m*}$ and a constant $\epsilon > 0$ and one index i such that $w = \widetilde{\Pi}^{\epsilon}_{\mathcal{F}_i(q)}(q)$ and $z = \widetilde{\Pi}^{\epsilon}_{\mathcal{F}_i(p)}(w)$ but $\|z - w\| > L_{\mathcal{F}}\|p - q\| + 3\epsilon$,
- Two tuples (x, w) and (x^*, w^*) with $\|x - x^*\| \leq \beta$ and $\|w - w^*\| \leq \beta$ such that $x^* \in \mathcal{R}(x)$ and $w^* \in \mathcal{F}(x)$ and $(y - x)^T w^* + \beta \mathbb{1} \geq 0, \quad \forall y \in \mathcal{R}(x)$ [17].

---

[17] Recall that $\mathbb{1}$ is all ones vector.

Our main result for MGQVI is the following:

Theorem 3.3.3. *The problem of finding a solution to MGQVI for strong separation oracles is in PPAD.*

Proof. We again consider $[-1,1]^m$ instead of $\mathbb{R}^{m*}$. First, we define the following[18]:

$$\Phi(y,x,w) = -(y-x)^\top w - \gamma(\|y\|_2^2)\mathbb{1}$$

$$\Pi(x,w) = \{y \in \mathcal{R}(x) \mid \forall i \in [r],\ (\Phi(y,x,w))_i > max_{y\in\mathcal{R}(x)}\,(\Phi(y,x,w))_i - \epsilon\}$$

Note that here w is a matrix and Φ is a vector-valued function. Next, we show that for a constant κ', the correspondence $\Psi(x,w) = (\Pi(x,w), \mathcal{F}(x))$ is κ'-Lipschitz continuous in (x,w). The next step will similarly (compared to proof of Theorem 3.2.2) be constructing a strong separation oracle for Ψ. Finally, we show that an approximate solution to Kakutani's fixed point of this correspondence will provide an approximate solution to the given MGQVI.

Remark 3.3.4. We use the Frobenius Norm for matrix distance defined in the following:

$$\|w_1 - w_2\|_M = \sqrt{\sum_{i=1}^{m}\sum_{j=1}^{r}(w_{1,ij} - w_{2,ij})^2}$$

Similar to proof of Theorem 3.2.2, we will need to bound the distance between pairs (x_1,w_1) and (x_2,w_2). Thus, we now define[19]:

$$\|(x_1,w_1) - (x_2,w_2)\|_X = \sqrt{\|x_1 - x_2\|_2^2 + \|w_1 - w_2\|_M^2}$$

To show that Ψ is κ'-Lipschitz continuous, we prove $d_H(\Psi(x_1,w_1),\Psi(x_2,w_2)) \leq \kappa'\|(x_1,w_1) - (x_2,w_2)\|_X^q + c$ for some constants q, and c.

Similar to the proof of the simper version (Theorem 3.2.2), $\Phi(y,x,w)$ is (2γ)-strongly concave (with respect to each coordinate) in y. In addition, each i-th coordinate of Φ is G-Lipschitz continuous where $G = m + 2\gamma m$ for all x and w. We also have access to the sub-gradients of $(\Phi)_i$ for each $i \in [r]$.

[18]Note that, unlike Theorem 3.2.10, Φ is a function that takes two vectors x and y (of size $m \times 1$), one matrix w (of size $m \times r$ which has r vectors given as columns of the matrix), and outputs one vector.

[19]We can extend the proof to use other norms such as the spectral norm.

Let $H(x,w) = \{y \in \mathcal{R}(x) \mid \forall i \in [r],\ (\Phi(y,x,w))_i = max_{y\in\mathcal{R}(x)}(\Phi(y,x,w))_i\}$. This means that $y \in H(x,w)$, must maximize all coordinates of Φ. We claim that such y must exists due the definition of multi-convexity of F and prove it later (see Lemma 3.3.6). By optimality KKT conditions for maximization of a concave function with respect to the constraint set $\mathcal{R}(x)$, for all $y \in \mathcal{R}(x)$, for all $i \in [r]$ we have:

$$\partial(\Phi(y,x,w))_i^T(y^* - y) \geq 0 \text{ where } y^* \in H(x,w)$$

In addition, (2γ)-strong-concavity of $(\Phi(\cdot,x,w))_i$ results the following inequality:

$$(\Phi(y^*,x,w))_i - (\Phi(y,x,w))_i \geq \partial\left((\Phi(y^*,x,w))_i\right)^\top (y^* - y) + \gamma \|y - y^*\|_2^2$$

Finally, by combining the previous inequalities we have:

$$(\Phi(y^*,x,w))_i - (\Phi(y,x,w))_i \geq \gamma \|y - y^*\|_2^2$$

In conclusion, for $y^* \in H(x,w)$ and $y \in \Pi(x,w)$, we have the following inequality:

$$\gamma \|y - y^*\|_2^2 \leq \epsilon \text{ or equivalently } \|y - y^*\|_2 \leq \sqrt{\frac{\epsilon}{\gamma}} \tag{3.6}$$

Let us define $H'(x,w) = \{y \in \mathcal{R}(x) \mid \Phi(y,x,w) = \text{argmaxall}_{y\in\mathcal{R}(x)}(\Phi(y,x,w))_i\}$ where argmaxall is an operator which tries to find a maximizer that maximizes all coordinates of Φ (see the modified Berge's maximum Theorem B.1.1). It is easy to verify that $H'(x,w) = H(x,w)$.

Now, we apply Theorem B.1.1 (relegated to the appendix) to $f((x,w),y) = \Phi(y,x,w)$ and $g((w,x)) = \mathcal{R}(x)$ which is a modified version of Berge's maximum theorem. By Theorem B.1.1, there exists a constant κ such that H' (g^* in Theorem B.1.1) is κ-Lipschitz continuous:

$$d_H(H'(x_1,w_1), H'(x_2,w_2)) \leq \kappa \|(x_1,w_1) - (x_2,w_2)\|_X \tag{3.7}$$

Recall that $H'(x,w) = H(x,w)$. Combining Equations 3.6 and 3.7 results in:

$$\begin{aligned} d_H(\Pi(x_1,w_1), \Pi(x_2,w_2)) \leq\ & d_H(H(x_1,w_1), H(x_2,w_2)) \\ & + d_H(\Pi(x_1,w_1), H(x_1,w_1)) \\ & + d_H(\Pi(x_2,w_2), H(x_2,w_2)) \end{aligned} \tag{3.8}$$

In conclusion, we have:

$$d_H(\Pi(x_1, w_1), \Pi(x_2, w_2)) \leq \kappa \|(x_1, w_1) - (x_2, w_2)\|_2^{\frac{1}{2}} + 2\sqrt{\frac{\epsilon}{\gamma}} \tag{3.9}$$

Since for each $i \in [r]$, $\mathcal{F}_i$ is $L_{\mathcal{F}}$-Hausdorff Lipschitz, by the definition of Ψ we can deduce there exists a constant κ':

$$d_H(\Psi(x_1, w_1), \Psi(x_2, w_2)) \leq \kappa' \|(x_1, w_1) - (x_2, w_2)\|_2^{\frac{1}{2}} + c', \tag{3.10}$$

showing that $\Psi(x, w)$ is κ'-Lipschitz continuous in (x, w).

To establish a reduction employing the computational variatnt of Kakutani, it is essential to establish the existence of a bounded-radius ball within the correspondence $\Psi(x, w)$. Let $y_x^* = \text{argmaxall}_{y \in \mathcal{R}(x)} \Phi(y, x, w)$ (recall that such y^* always exists due to Lemma 3.3.6). Similarly assume that $\overline{\mathrm{B}}(\mathcal{R}(x), -\epsilon)$ and also $\overline{\mathrm{B}}(\mathcal{F}(x), -\epsilon)$ are non-empty. Then, define a non-empty region $V_x = \overline{\mathrm{B}}(y_x^\star, +\frac{\epsilon}{G'}) \cap \mathcal{R}(x) = \overline{\mathrm{B}}(y_x^\star, +\frac{\epsilon}{G'})$ where $G' = r \cdot G$. Note that G is the Lipschitz constant of Φ_i for all i. By Lipschitzness of $\Phi(\cdot, x, w)$, we have:

$$\forall y \in V_x, \quad \|\Phi(y, x, w) - \Phi(y_x^*, x, w)\|_X \leq G' \frac{\epsilon}{G'} = \epsilon$$

In conclusion, $\Phi(V_x, x, w) \subseteq [\Phi(y_x^*, x, w) - \epsilon \mathbb{1}, \Phi(y_x^*, x, w)]$ is non-empty[20]. Then, we can show that $\Phi(V_x, x, w) \subseteq \Pi(x, w)$. As for any $x \in \mathcal{R}(x)$, $\overline{\mathrm{B}}(\mathcal{F}(x), -\epsilon)$ is non-empty, we can find a non-empty region C_x such that $C_x \subseteq \overline{\mathrm{B}}(\mathcal{F}(x), -\epsilon)$. This is sufficient to show that for all x and w, $\Psi(x, w)$ is non-empty. Now we are ready to construct a strong separation oracle for $\Psi(x, w)$ using a modified version of the Strong Constrained Convex Optimization framework stated in the following.

[20]We use the convention that a vector lies between two others when each of its coordinates lies strictly between the corresponding coordinates of those two vectors.

Modified Strong Constrained Convex Optimization Problem (2)

Input: A zeroth and first order oracle for the multi-concave function $G : \mathbb{R}^{m*} \to \mathbb{R}^{r*}$, two rational numbers $\delta, \epsilon > 0$ and a strong separation oracle $\mathrm{SO}_{\mathcal{R}}$ for a non-empty closed convex-valued correspondence $\mathcal{R} : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$ and one input x.

Output: One of the following cases:

- (Violation of non-emptiness): A failure symbol $\perp$ with a polynomial-sized witness that certifies that $\overline{\mathrm{B}}(\mathcal{R}(x), -\delta) = \emptyset$,
- (Approximate Maximization): A vector $z \in \mathbb{Q}^m \cap \mathcal{R}(x)$, such that $G(z) + \epsilon\mathbb{1} \geq \mathrm{maxall}_{y\in\mathcal{R}(x)}\, G(y)$ for all values.

This modified strong constrained convex optimization problem can be solved by a modified version of the sub-gradient ellipsoid central cut (see Appendix B.3). We can compute such a solution $y^* \in \mathcal{R}(x)$ such that $\Phi\left(y^*, x, w\right) \succeq max_{y\in\mathcal{R}(x)}\Phi\left(y, x, w\right) - \epsilon\mathbb{1}$. Thus, given the strong separation oracles representing $\mathcal{F}$ and $\mathcal{R}$, we can construct a strong separation oracle for Ψ by considering a strong separation oracle for the following set:

$$\overline{\Psi}_s(x, w) = \{(y, w) \in (\mathcal{R}(x), \mathcal{F}(x)) \mid -\Phi(y, x, w) \preceq \sigma\}$$

where $\sigma = -\Phi(y^*, x, w)$.

Now, we can give Ψ as input to the Kakutani problem presented in [152] with accuracy parameter α, however, we need to modify the input. This is due to the fact that our input is a separation oracle for Ψ which takes one vector x and one matrix w while Kakutani's problem takes a correspondence in the format of $\mathcal{F} : \mathbb{R}^{d*} \rightrightarrows \mathbb{R}^{d*}$. We need to consider $d = mr + m$ and encode the given matrix w with r-vectors to make it completely compatible with the conditions of the Kakutani's fixed point. In other words, we are aim to find a fixed point of the correspondence Ψ' where $\Psi'(x, w_1, \ldots, w_r) = \Psi(x, \sum_{i=1}^{r} e_i w_i^\top)$. In the following we show that using the X norm we defined does not change the outcome if we encode the correspondence given the above-mentioned scheme[21].

[21] Alternatively, we may consider and define a format of Kakutani's fixed point such that it works with the correspondences of the format $\mathcal{F} : \mathbb{R}^{a_1*} \times \mathbb{R}^{a_2\times a_3*} \rightrightarrows \mathbb{R}^{b_1*} \times \mathbb{R}^{b_2\times b_3*}$ However, this will require redoing the proofs of [152]

Now for achieving the desired approximation, we set $\alpha = \frac{\epsilon'}{\kappa'}$ where $\epsilon' = \epsilon h$. Given the output of Kakutani, we decode the output by reverting it to the matrix format. We get two points $z = (x, w) \in ([-1,1]^m, [-1,1]^{m\times r})$ and also $z^* = (x^*, w^*) \in \Psi(x, w)$ where $\|(x, \sum_{i=1}^r e_i w_i^\top) - (x^*, \sum_{i=1}^r e_i w_i^{*\top})\|_2 \leq \frac{\epsilon'}{\kappa'}$ resulting $\|(x,w) - (x^*, w^*)\|_X \leq \frac{\epsilon'}{\kappa'}$ $d((x,w), \Psi(x,w)) \leq \frac{\epsilon'}{\kappa'}$.

Thus, $w^* \in \mathcal{F}(x)$ and $d(w, \mathcal{F}(x)) \leq \frac{\epsilon'}{\kappa'}$. In addition, by the definition of Ψ, $x^* \in \Pi(x, w)$ and $d(x, \Pi(x,w)) \leq \frac{\epsilon'}{\kappa'}$. By the definition of Π, for every $y \in \mathcal{R}(x)$, $\Phi(x^*, x, w) \geq \Phi(y, x, w) - \epsilon\mathbb{1}$. In conclusion, we have the following:

$$(y-x)^T w + 2\gamma\|y\|_2^2\mathbb{1} \succeq (x^* - x)^T w + 2\gamma\|x^*\|_2^2\mathbb{1} - \epsilon\mathbb{1},\ \forall y \in \mathcal{R}(x)$$

So, recalling that $x, y, w \in [-1,1]^m$:

$$(y-x)^T w \succeq \pm\frac{2\sqrt{m}\epsilon'}{\kappa'}\mathbb{1} \pm 4\gamma\mathbb{1} - \epsilon\mathbb{1}, \quad \forall y \in \mathcal{R}(x)$$

Finally, by knowing the facts that $\|(x,w) - (x^*, w^*)\|_2^2 \leq \frac{\epsilon'}{\kappa'}$ and $w^* \in \mathcal{F}(x)$, we can simply have:

$$(y-x)^T w^* \succeq \pm\frac{2\sqrt{m}\epsilon'}{\kappa'}\mathbb{1} \pm 4\gamma\mathbb{1} \pm \frac{2\sqrt{m}\epsilon'}{\kappa'}\|y - x\|\mathbb{1} - \epsilon\mathbb{1},, \quad \forall y \in \mathcal{R}(x)$$

Considering the fact that $x, y, w \in [-1,1]^m$:

$$(y-x)^T w^* \succeq \pm\frac{2\sqrt{m}\epsilon'}{\kappa'}\mathbb{1} \pm 4\gamma\mathbb{1} \pm \frac{2\sqrt{m}\epsilon'}{\kappa'}\mathbb{1} - \epsilon\mathbb{1},, \quad \forall y \in \mathcal{R}(x)$$

Considering appropriate small numbers for ϵ, h and γ will imply the following for $x^* \in \mathcal{R}(x)$:

$$(y-x)^T w^* \succeq -\beta\mathbb{1}, \quad \forall y \in \mathcal{R}(x)$$

□

Remark 3.3.5. If we do not use X norm and for example use the spectral norm, we need to consider a constant θ for the error caused by the simple decoding. This error is caused by the fact that the norm that we defined for the input given as a matrix and a vector becomes different from $L2$ distance approximation guarantee we get from the computational Kakutani's fixed point theorem. This does not cause a problem as the matrix has constant number of columns and these different norms are related (for example, see Lemma 3.3.8). The proof can be adapted for other norms using an appropriate number h to cancel out the encoding error caused by encoding the matrices.

Below, we show that multi-convexity of $\mathcal{F}$ ensures a unique maximizer for Φ in all coordinates.

Lemma 3.3.6. *If the correspondence $\mathcal{F}$ defied above satisfies multi-convexity, sub-gradient compatibility conditions then we have a unique maximizer for the set* $H(x, w) = \left\{y \in \mathcal{R}(x) \mid \forall i \in [r],\, (\Phi(y, x, w))_i = \max_{y' \in \mathcal{R}(x)} (\Phi(y', x, w))_i\right\}$.

Proof. The function $\Phi(y, x, w)$ is strictly concave in y. Since $\mathcal{F}(x)$ satisfies multi-convexity and sub-gradient compatibility, the feasible set remains convex for any restricted components given the fact that $w \in \mathcal{F}(x)$. These two conditions ensure that the function $\Phi(y, x, w)$ has at least one maximizer that is common in all of the variational inequalities. As discussed, for each i, $\Phi_i(y, x, w)$ has a unique maximizer. Therefore, the set $H(x, w)$, which consists of the maximizers of $\Phi(y, x, w)$, contains a unique maximizer. □

Remark 3.3.7. Note that one may consider the a maximum solution for $(\Phi(y, x, w))_i$ would be $y^* = -\frac{w_i}{2\gamma}$ and infer that a unified maximizer must have the condition $w_i = w_j$ but this y^* must belong to $\mathcal{R}(x)$. Thus, when y is restricted to lie in a convex set $R(x)$ (i.e. under projection), the unique maximizer can be given by $y^* = \Pi_{R(x)}\left(-\frac{w_i}{2\gamma}\right)$. Even if the unconstrained maximizers $-\frac{w_i}{2\gamma}$ differ, their projections can coincide. Therefore, in the presence of domain restrictions $(x \in R(x))$, a common maximizer y^* does not necessarily imply that we must have $w_i = w_j$ for all i, j.

Lemma 3.3.8. *Suppose $x_1, x_2 \in [-1, 1]^m$ and $w_1, w_2 \in \mathbb{R}^{r \times m}$, where m_1 and m_2 encode the matrix w_1 and w_2, respectively. If $\|(x_1, w_1) - (x_2, w_2)\|_2 \le \sigma$, then the bound for the norm $\|(x_1, w_1) - (x_2, w_2)\|_X$ is given by $\|(x_1, w_1) - (x_2, w_2)\|_X \le \sigma$.*

Proof. Given that $\|(x_1, m_1) - (x_2, m_2)\|_2 \le \sigma$, where m_1 and m_2 are the vectorized forms of w_1 and w_2, we can express the norm $\|(x_1, w_1) - (x_2, w_2)\|_X$ as square root of the sum of the norms of their respective components. Utilizing the properties of the vectorized Frobenius norm, we find that $\|m_1 - m_2\|_2$ is equivalent to $\|(w_1 - w_2)\|_M$. □

Chapter 4

Equilibrium With Convex Constraints

A primary objective of algorithmic game theory is to categorize the complexity of key economic concepts. PPAD-completeness has become a significant unifying principle within this endeavor [58, 152]. In this chapter, applying the computational versions of the variational inequality problems, we demonstrate that several central problems in game theory are PPAD-complete under reasonable convexity assumptions. We first review well-known remedial solutions for leader/follower equilibria in multi-leader-follower games, both with and without uncertainty, and present a unified mathematical definition that encompasses these remedial approaches. Building on this, we analyze the computational complexity of appropriate computational formulations of these problems, establishing their PPAD-completeness using similar methodologies. Finally, we consider the computational complexity of resilient Nash equilibrium and related solution concepts under given specific continuity and concavity conditions (in particular, a more restricted condition that we call multi-concavity) assumptions. Specifically, we demonstrate that these problems can be formulated as a more generalized class of variational inequalities, which we refer to as MGQVI in Chapter 3.

4.1 Remedial Approaches for Multi-Leader-Follower Games

In addition to the complexity of the leaders' optimization problem, the non-convexity of the leader's constraints may mean that a L/F-equilibrium does not always exist[1]. The algebraic relaxation/restriction approach (*remedial solutions*) introduced in [107, 150] addresses this issue. These approaches may consider *Karush–Kuhn–Tucker conditions* (KKT) and introduce various relaxations/restrictions[2]. Several other approaches, such as using convex hull or mixed strategies, exist [150].

4.1.1 First Remedial Model of [150] for Multi-leader Follower Games

As discussed, we may assume follower i's constraints are represented by a non-empty convex correspondence.

$$\mathcal{R}_i\left(x_{\mathrm{I}}, x_{\mathrm{II}}, y_{-i}\right) \equiv \left\{y_i \in S_i \mid g_i\left(x_{\mathrm{I}}, x_{\mathrm{II}}, y\right) \leq 0,\ h_i\left(x_{\mathrm{I}}, x_{\mathrm{II}}, y_i\right) \leq 0\right\}$$

where $x_{\mathrm{I}}, x_{\mathrm{II}} \in X^{\mathrm{I}} \times X^{\mathrm{II}}$ are the leaders' strategies and $g_i : X^{\mathrm{I}} \times X^{\mathrm{II}} \times \Pi_{j=1}^{k} S_j \rightarrow \mathbb{R}^{m_i}$ and $h_i : X^{\mathrm{I}} \times X^{\mathrm{II}} \times S_i \rightarrow \mathbb{R}^{l_i}$ are convex and Lipschitz continuous. Note that m_i and l_i are some given constants.

By incorporating the KKT conditions of the followers' optimization problems with x^*_{II} as an exogenous variable (representing the second leader's optimal solution), leader I's optimization problem can be formulated as finding a solution $(x_{\mathrm{I}}, y_{\mathrm{I}}, \lambda_{(i,\mathrm{I})}, \mu_{(i,\mathrm{I})})$ to the following problem:

$$\begin{aligned}
&\operatorname{Min} \phi_{\mathrm{I}}\left(x_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{\mathrm{I}}\right) \\
&\text{s.t } x_{\mathrm{I}} \in X^{\mathrm{I}} \\
&\forall i \in [k],\ \partial_{y_{(i,\mathrm{I})}} \theta_i\left(x_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{\mathrm{I}}\right) + \sum_{j=1}^{m_i} \lambda^j_{(i,\mathrm{I})} \partial_{y_{(i,\mathrm{I})}} g^j_i\left(x_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{\mathrm{I}}\right) + \sum_{j=1}^{l_i} \mu^j_{(i,\mathrm{I})} \partial_{y_{(i,\mathrm{I})}} h^j_i(x_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{(i,\mathrm{I})}) = 0 \\
&\forall i \in [k],\ 0 \leq \lambda_{(i,\mathrm{I})} \perp g_i\left(x_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{\mathrm{I}}\right) \leq 0 \\
&\forall i \in [k],\ 0 \leq \mu_{(i,\mathrm{I})} \perp h_i(x_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{(i,\mathrm{I})}) \leq 0
\end{aligned}$$

In general, the constraints in the aforementioned problem lack convexity with respect to their variables. By considering a restriction that still encompasses a wide range of practical models, the remedial model of [150] addresses this issue. More

[1] See Example 4 in [150].

[2] Note that we only investigate the approximate version for the remedial models. There are a variety of hardness results for Stackelberg games (e.g. [51, 132]).

precisely, we will assume that the functions θ_i, g_i, and h_i take the following forms:

$$\theta_i\left(x_{\mathrm{I}}, x_{\mathrm{II}}, y\right) \equiv \frac{1}{2}\left(y_i\right)^T M_i y_i + \left(c_i\left(x_{\mathrm{I}}, x_{\mathrm{II}}, y_{-i}\right)\right)^T y_i + \psi_i\left(x_{\mathrm{I}}, x_{\mathrm{II}}, y_{-i}\right)$$
$$g_i\left(x_{\mathrm{I}}, x_{\mathrm{II}}, y\right) \equiv A_{(i,\mathrm{I})} x_{\mathrm{I}} + A_{(i,\mathrm{II})} x_{\mathrm{II}} + \sum_{j=1}^{k} B_{i,j} y_j \text{ and } h_i\left(x_{\mathrm{I}}, x_{\mathrm{II}}, y_i\right) \equiv C_{(i,\mathrm{I})} x_{\mathrm{I}} + C_{(i,\mathrm{II})} x_{\mathrm{II}} + D_i y_i \tag{4.1}$$

for some matrices $M_i \in \mathbb{R}^{n_i \times n_i}$, which are symmetric positive semidefinite, $A_{(i,\mathrm{I})} \in \mathbb{R}^{m_i \times n_{\mathrm{I}}}$, $A_{(i,\mathrm{II})} \in \mathbb{R}^{m_i \times n_{\mathrm{II}}}$, $B_{i,j} \in \mathbb{R}^{m_i \times n_j}$, $C_{(i,\mathrm{I})} \in \mathbb{R}^{l_i \times n_{\mathrm{I}}}$, $C_{(i,\mathrm{II})} \in \mathbb{R}^{l_i \times n_{\mathrm{II}}}$, $D_i \in \mathbb{R}^{l_i \times n_i}$, affine functions $c_i : \mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}+n_{-i}} \to \mathbb{R}^{n_i}$, and arbitrary real-valued functions $\psi_i : \mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}+n_{-i}} \to \mathbb{R}$.

Remark 4.1.1. Recall that n_{I}, n_{II} and n_i denote the variables that the leaders and the follower i control (out of $n_{\mathrm{I}}+n_{\mathrm{II}}+n$ where $n = \sum_{i=1}^{k} n_i$) respectively. Furthermore, n_{-i} denotes the variables that all other followers other than i control.

With these modifications, we formulate leader I's optimization problem as follows: With x_{II}^* as an exogenous variable , find a solution $(x_{\mathrm{I}}, y_{\mathrm{I}}, \lambda_{(i,\mathrm{I})}, \mu_{(i,\mathrm{I})}) \in X^{\mathrm{I}} \times X^{\mathrm{II}} \times \Pi_{j=1}^{k} S_j \times \mathbb{R}^{m_i+l_i}$ to the following:

$$\begin{aligned}
&\text{Min } \phi_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}^*, y_{\mathrm{I}}\right)\\
&\text{s.t } x_{\mathrm{I}} \in X^{\mathrm{I}}\\
&c_i\left(x_{\mathrm{I}}, x_{\mathrm{II}}^*, y_{(-i,\mathrm{I})}\right) + M_i y_{(i,\mathrm{I})} + \left(B_{i,i}\right)^T \lambda_{(i,\mathrm{I})} + D_i^T \mu_{(i,\mathrm{I})} = 0\\
&\forall i \in [k], \quad \left(\lambda_{(i,\mathrm{I})}\right)^T \left[A_{(i,\mathrm{I})} x_{\mathrm{I}} + A_{(i,\mathrm{I})} x_{\mathrm{II}}^* + \sum_{j=1}^{k} B_{i,j} y_{(j,\mathrm{I})}\right] = 0\\
&\forall i \in [k], \quad \left(\mu_{(i,\mathrm{I})}\right)^T \left[C_{(i,\mathrm{I})} x_{\mathrm{I}} + C_{(i,\mathrm{I})} x_{\mathrm{II}}^* + D_i y_{(i,\mathrm{I})}\right] = 0
\end{aligned} \tag{4.2}$$

Recall that for leader I, we denote an anticipated strategy of all k followers except i by $y_{(-i,\mathrm{I})}$. With the exception of the set X^{I}, which might have nonlinear (though convex) constraints and the orthogonality conditions, the remaining constraints in leader I's problem are linear. Informally, the approach described in [150] for handling these nonconvex constraints is described as follows: Each leader has access to limited information about the followers' reactions, characterized by specific "favorable" sets $Z^{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}})$ and $Z^{\mathrm{II}}(x_{\mathrm{I}}, x_{\mathrm{II}})$, respectively, and subject to which the leaders optimize their objective functions. In general, such partial information could be of one of two kinds, *restricted* or *relaxed* information, as follows[3]:

$$Z^{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}\right) \subseteq Y_{sol}\left(x_{\mathrm{I}}, x_{\mathrm{II}}\right) \text{ or } Y_{sol}\left(x_{\mathrm{I}}, x_{\mathrm{II}}\right) \subseteq Z^{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}\right) \tag{4.3}$$

A restricted response is an equilibrium response, whereas a relaxed one may not be. Similar classifications apply to the partial responses available to leader II. By separating their responses, one leader may have restricted follower responses while the other has relaxed ones. The reference [150] suggests substituting the two non-convex orthogonality conditions in leader I's constraints with the following conditions:

$$\left(x_{\mathrm{I}}, y_{\mathrm{I}}, \lambda_{(i,\mathrm{I})}\right) \in W_{(i,\mathrm{I})}\left(x_{\mathrm{II}}\right) \qquad \text{and} \qquad \left(x_{\mathrm{I}}, y_{(i,\mathrm{I})}, \mu_{(i,\mathrm{I})}\right) \in V_{(i,\mathrm{I})}\left(x_{\mathrm{II}}\right)$$

where $W_{(i,\mathrm{I})}\left(x_{\mathrm{II}}\right) \subseteq X^{\mathrm{I}} \times \Pi_{j=1}^{k} S_j \times \mathbb{R}^{m_i}$ and $V_{(i,\mathrm{I})}\left(x_{\mathrm{II}}\right) \subseteq X^{\mathrm{I}} \times S_i \times \mathbb{R}^{l_i}$ are appropriate polyhedral sets such that together they represent either a restriction or a relaxation of the complementarity constraints in $Y_{sol}(x_{\mathrm{I}}, x_{\mathrm{II}}^*)$. For all $i = 1, \ldots, k$, we define $Z^{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}\right)$ to be the set of all tuples $y_{\mathrm{I}} = \left(y_{(j,\mathrm{I})}\right)_{j=1}^{k}$ for which there exists $(\lambda_{\mathrm{I}}, \mu_{\mathrm{I}})$ that satisfies:

$$c_i\left(x_{\mathrm{I}}, x_{\mathrm{II}}, y_{(-i,\mathrm{I})}\right) + M_i y_{(i,\mathrm{I})} + \left(B_{i,i}\right)^T \lambda_{(i,\mathrm{I})} + \left(D_i\right)^T \mu_{(i,\mathrm{I})} = 0$$

$$0 \leq \lambda_{(i,\mathrm{I})}, \quad A_{i,\mathrm{I}} x_{\mathrm{I}} + A_{i,\mathrm{II}} x_{\mathrm{II}} + \sum_{j=1}^{k} B_{i,j} y_{(j,\mathrm{I})} \leq 0$$

$$0 \leq \mu_{(i,\mathrm{I})}, \quad C_{(i,\mathrm{I})} x_{\mathrm{I}} + C_{(i,\mathrm{II})} x_{\mathrm{II}} + D_i y_{(i,\mathrm{I})} \leq 0$$

We call the elements of $Z^{\mathrm{I}}\left(x_{\mathrm{I}}^*, x_{\mathrm{II}}\right)$ the followers' partial responses anticipated by (or available to) leader I and categorize such responses as restricted or relaxed based on Equation 4.3[4]. Finally, in terms of the partial responses, leader I, given x_{II}^* as an exogenous variable, must find a solution $(x_{\mathrm{I}}, y_{\mathrm{I}})$ to the following surrogate optimization problem:

$$\begin{aligned} &\text{Min } \phi_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}^*, y_{\mathrm{I}}\right) \\ &\text{s.t } x_{\mathrm{I}} \in X^{\mathrm{I}} \\ &\text{and} \qquad \left(x_{\mathrm{I}}, y_{\mathrm{I}}\right) \in \text{graph } Z^{\mathrm{I}}\left(\cdot, x_{\mathrm{II}}^*\right) \end{aligned} \tag{4.4}$$

For leader II, the optimization problem with the surrogate complementarity conditions can be defined by considering $Z^{\mathrm{II}}(x_{\mathrm{I}}^*, x_{\mathrm{II}})$, $W_{(i,\mathrm{II})}(x_{\mathrm{I}}^*)$ and $V_{(i,\mathrm{II})}(x_{\mathrm{I}}^*)$ where x_{I}^*

[3]Recall that Y_{sol} denotes the solution of the follower problem in Equation 2.11

[4]We can construct a strong separation oracle for this set given the polyhedral sets' separation oracles.

will be an exogenous variable:

$$\begin{aligned} &\text{Min}\,\phi_{\text{II}}\left(x_{\text{I}}^{*}, x_{\text{II}}, y_{\text{I}}\right) \\ &\text{s.t}\ x_{\text{II}} \in X^{\text{II}} \\ &\text{and} \qquad \left(x_{\text{II}}, y_{\text{II}}\right) \in \text{graph}\, Z^{\text{II}}\left(x_{\text{I}}^{*}, \cdot\right) \end{aligned} \tag{4.5}$$

Remark 4.1.2. Note that by combining the given constraints, we are able to write Equation 4.4 in terms of correspondences G_{I} and G_{II}. In particular, the optimization problem for leader I and II can be written as a generalized Nash equilibrium problem:

$$\begin{aligned} &\text{Min}\,\phi_{\text{I}}\left(x_{\text{I}}, x_{\text{II}}^{*}, y_{\text{I}}\right) \\ &\text{s.t}\ \left(x_{\text{I}}, y_{\text{I}}\right) \in G_{\text{I}}\left(x_{\text{II}}^{*}\right) \end{aligned} \tag{4.6}$$

And for leader II, we have the following:

$$\begin{aligned} &\text{Min}\,\phi_{\text{II}}\left(x_{\text{I}}^{*}, x_{\text{II}}, y_{\text{I}}\right) \\ &\text{s.t}\ \left(x_{\text{II}}, y_{\text{II}}\right) \in G_{\text{II}}\left(x_{\text{I}}^{*}\right) \end{aligned} \tag{4.7}$$

Definition 4.1.3. *A* remedial L/F-equilibrium *is a pair of leaders' strategies* $(x_{\text{I}}^{*}, x_{\text{II}}^{*})$ *such that there exists a pair* $(y_{\text{I}}^{*}, y_{\text{II}}^{*})$ *such that* $(x_{\text{I}}^{*}, y_{\text{I}}^{*})$ *and* $(x_{\text{II}}^{*}, y_{\text{II}}^{*})$ *constitute a solution of the two leaders' surrogate optimization problems 4.4 and 4.5, respectively.*

We define $Sol(4.4)$ and $Sol(4.5)$ to be the minimized value (of ϕ_{I} and ϕ_{II} given their respective exogenous variables) of a solution of the two leaders' surrogate optimization problems 4.4 and 4.5.

Before proceeding to proof of PPAD-completeness, given the format proposed in the Remark above (generalized equilibrium reformulation), we might assume that we can use Proposition 2.7.6 to transform the remedial problem to the QVI format under some conditions. However, this proposition assumes the differentiability of the loss functions (utilities). Since these are now represented by linear arithmetic circuits, which as mentioned above are not necessarily differentiable, we need to take a more general approach carefully applying the following proposition.

Proposition 4.1.4. *Suppose that we have a game with convex loss functions* $\phi = (\phi_1, \ldots, \phi_{\text{k}})$*, and strategies* $\mathcal{X} = (X^1, \ldots, X^{\text{k}})$ *represented by linear arithmetic circuits and constraints* $\mathcal{G} = (G_1, \ldots, G_{\text{k}})$ *that have separation oracles represented by linear arithmetic circuits which are closed convex subsets of* $\mathcal{X} = (X^1, \ldots, X^{\text{k}})$*. The* *problem of finding an approximate generalized (Nash) equilibrium of this game can be transformed into finding an approximate solution to a GQVI problem*[5].

Proof. For simplicity, we prove the proposition for $k = 2$ players. Given the definitions of Equations 4.6 and 4.7, we actually can consider only $\phi_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}})$ and $\phi_{\mathrm{II}}(x_{\mathrm{I}}, x_{\mathrm{II}})$ with two variables. Define the correspondence $\mathcal{F}(x) : \mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}} \to \mathcal{P}\left(\mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}}\right)$ to be:

$$\mathcal{F}(x) = \partial_{x_{\mathrm{I}}}\phi_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}) \times \partial_{x_{\mathrm{II}}}\phi_{\mathrm{II}}(x_{\mathrm{I}}, x_{\mathrm{II}})$$

Note that $\mathcal{F}(x)$ is a correspondence as the sub-differentials are not necessarily unique due to the definition. $\mathcal{P}(A)$ also denotes the power set of A. We formulate the following problem by concatenating the first-order optimality conditions of all leaders' problems where the goal is finding a vector $x = (x_{\mathrm{I}}, x_{\mathrm{II}}) \in \mathcal{X}$ (where $\mathcal{X} = X^{\mathrm{I}} \times X^{\mathrm{II}}$) such that:

$$\exists w \in \mathcal{F}(x), \quad (y - x)^\top w + \epsilon \geq 0, \quad \forall y \in \mathcal{G}(x)$$

Note that $\mathcal{G}(x) = (G_{\mathrm{I}}(x_{\mathrm{I}}), G_{\mathrm{II}}(x_{\mathrm{II}}))$. The functions ϕ_{I} and ϕ_{II} are linear arithmetic circuits and they are piece-wise linear convex functions. In a piecewise linear convex function, kink points (where segments meet) are points of non-differentiability, as the gradient changes abruptly. At these points, the function has a sub-differential, representing all possible slopes of adjacent segments. Despite this non-differentiability, convexity is preserved (for the subgradients), ensuring any local minimum, even at a kink point, is also a global minimum. The sub-differential at a kink point includes any convex combination of all possible slopes of the adjacent segments. Thus, a useful property of linear arithmetic circuits is that the sub-differentials at one point always form a convex set.

□

Computational Definition

We are now ready to define the computational version of finding an equilibrium and state the PPAD-completeness result in this setting. We could consider alternative definitions that omit Z^{I} and Z^{II} and consider strong separation oracles for the above-mentioned appropriate polyhedral sets as inputs. This is because all of the requirements for constructing separation oracles for Z^{I} and Z^{II} can be assured by the linearity

[5]This proposition is a slightly modified and generalized version of the transformation provided in [107].

of the loss functions of the followers and having the possibility of representing the approximate polyhedral sets by separation oracles. We chose this version for the sake of a simpler presentation.

REMEDIAL ([150]) L/F EQUILIBRIUM WITH STRONG SEPARATION ORACLES

Input: We receive as input all the following:

- Two linear arithmetic circuits representing the convex loss functions (ϕ_{I},ϕ_{II}) for two leaders,
- Two linear arithmetic circuits representing strong separation oracles for non-empty, convex, and compact sets X^{I} and X^{II} for the leaders I and II,
- Two linear arithmetic circuits representing strong separation oracles for Z^{I} and Z^{II} that represent restricted or relaxed responses of the followers for each leader respectively that are two non-empty, Hausdorff Lipschitz, convex-valued, and compact correspondences.
- An accuracy parameter β.

Output: One of the following cases:

- (Violation of non-emptiness): A certificate indicating at least one of the X^{I}, X^{II}, $Z^{\mathrm{I}}(x, \cdot)$ for some $x \in X^{\mathrm{I}}$ or $Z^{\mathrm{II}}(\cdot, x)$ for some $x \in X^{\mathrm{II}}$ is empty.
- (Violation of convexity) of any of the loss functions of the inputs,
- (Approximate minimization): Vectors $(x_{\mathrm{I}}^*, y_{\mathrm{I}}^*, x_{\mathrm{II}}^*, y_{\mathrm{II}}^*)$ having the following relationship[6]:
  - $\phi_{\mathrm{I}}(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*, y_{\mathrm{I}}^*) \leq \beta + \text{ Sol}(4.4)$
  - $\phi_{\mathrm{II}}(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*, y_{\mathrm{II}}^*) \leq \beta + \text{ Sol}(4.5)$

Remark 4.1.5. We can also consider circuits that are sums of monomials, for which Lipschitz continuity and computing the sub-gradient is straightforward. The case

[6]Here, the output vectors are needed to satisfy all feasibility constraints of (4.4) and (4.5).

(Violation of convexity) is meaningful as output whenever the form is explicitly given (e.g. sum of monomials or linear arithmetic circuits) otherwise convexity holds as a promise. The non-emptiness exceptions can be distinguished by a slightly generalized ellipsoid algorithm we discuss in Appendix B.3.

Theorem 4.1.6. *In a multi-leader-follower game, the problem of finding an approximate remedial L/F-equilibrium where the constraints of the followers are given by strong separation oracles is PPAD-complete.*

Proof Sketch. Consider the format given in Equations 4.6 and 4.7 and carefully apply Proposition 4.1.4, which allows us to transform the problem into a GQVI and to establish inclusion in PPAD. PPAD-hardness of this problem follows from the hardness of finding a mixed Nash equilibrium (see [45]) in a game with 2 leaders, where the loss functions represent the expected payoff of mixed strategies where the followers have only one strategy and no restrictions[7]. □

Proof of Theorem 4.1.6

Before proceeding to the complete proof, we will need several propositions and tools.

A Separation Oracle for Sub-gradients

We provide a separation oracle for the sub-gradients of linear arithmetic circuits, specifically for correspondence $\mathcal{F}$ in Proposition 4.1.4. According to the results of [17], we can efficiently compute only one vector belonging to the sub-gradients of a function. For a given point x, we first determine which segment it belongs to and obtain the corresponding slope as the sub-gradient. To check the validity of a proposed sub-gradient y, compare it to the slopes of the segments at the specified point x: if y matches one slope, it is valid; otherwise, it lies outside the range of slopes. If y is invalid, the separation oracle outputs a hyperplane that separates d from the valid sub-gradients.

However, we need to describe the convex hull of in polynomial time and space. This process is efficient and easy, as shown in Algorithm 1. In short, Algorithm 1 can efficiently and concisely represent all sub-gradients of a vector-valued linear arithmetic circuit. In each step, a set of sub-gradients is either a constraint or just a number and will be propagated for the calculation of the next step. This approach works

[7]The complete proof (also for Proposition 4.1.4) is available in Section 4.1.1.

because the structure of linear arithmetic circuits ensures that each sub-gradient set is manageable in size.

ALGORITHM 1: Separation Oracle for Sub-gradients in Linear Arithmetic Circuits

---

Input: Linear arithmetic circuit C computing function $f(x)$, input point x, and point y to check.

Output: "Yes" if y is a valid sub-gradient for $f(x)$; otherwise, return a separating hyperplane.

Function `Subgradient Oracle`(C, x, y)**:**

  Evaluate $f(x)$ and also calculate the set of possible sub-gradients $S(x)$ at x as follows: **Represent $S(\cdot)$ implicitly: create symbolic gate-level vector variables s_j and collect linear constraints (do not enumerate extreme points).**;

  for *each node in* C **do**

    if *the operation is* $+$ *or* $-$ **then**

      Apply standard rules for $+$ or $-$ and update $S(x)$.

    end

    else if *the operation is multiplication by a constant* c **then**

      Apply standard rules for multiplication and update $S(x)$.

    end

    else if *the operation is* max *or* min **then**

      if *two inputs to the gate,* g_1 *and* g_2*, satisfy* $g_1(x) = g_2(x)$ **then**

        Create scalar mixing variable λ_j with $0 \leq \lambda_j \leq 1$ and add linear constraint $s_j = \lambda_j s_{g_1} + (1 - \lambda_j) s_{g_2}$ to the collected constraints.;

      end

      else

        Choose the subgradient from the function that attains the maximum or minimum at x.

      end

    end

  end

  Propagate the set $S(x)$ backward through the circuit, maintaining convex sub-gradient structure.;

  Add equalities $s_{\text{output}} = y$ to the collected linear constraints and solve the resulting linear-feasibility LP (exact rationals); if feasible return "Valid", else extract LP infeasibility/Farkas certificate and return the resulting rational separating hyperplane.;

Remark 4.1.7. The algorithm efficiently propagates subgradient sets through linear arithmetic circuits by maintaining an implicit representation via linear programming constraints. At each max/min gate where both inputs are equal at the evaluation point, we introduce a scalar mixing variable $\lambda_j \in [0, 1]$ and enforce that the output subgradient is a convex combination $s_j = \lambda_j s_{g_1} + (1 - \lambda_j) s_{g_2}$, adding $d + 2$ constraints for d-dimensional subgradients. This approach avoids explicit convex hull computation, ensuring the constraint system grows polynomially with circuit size, and the final membership test reduces to a polynomial-time linear feasibility problem.

Bounding the Approximation Error

The following lemma can also be applied to remedial L/F-equilibrium, as remedial L/F-equilibrium can be written in the format of a generalized equilibrium problem. We will prove this Lemma in the next Chapter.

Lemma 4.1.8 (Generalized Penalty Lemma)**.** *Suppose that x is a generalized equilibrium in game $(S, \mathcal{R}, \Theta)$ with L-Lipschitz loss functions Θ[8]. Let x^* be a strategy profile that is β-statistically close to x and $x^* \in \mathcal{R}(x^*)$. Then, x^* is a $f(\beta)$-approximate generalized equilibrium of this game where $f(\beta) = L\beta$.*

Proof of Theorem 4.1.6

We are now ready to provide a formal proof for Theorem 4.1.6.

Proof. To establish inclusion in PPAD, we must consider solutions of the computational variant of the transformed GQVI from Proposition 4.1.4. For the computational formulation of GQVI, a solution consists of two tuples (x, w) and (x^*, w^*) with $\|(x, w) - (x^*, w^*)\|_2^2 \leq \beta$ such that $x^* \in \mathcal{G}(x)$ and $w^* \in \mathcal{F}(x)$ and $(y - x)^T w^* + \beta \geq 0$, for every $y \in \mathcal{G}(x)$. We take $x^* \in \mathcal{G}(x)$ as the final approximate solution for remedial L/F-equilibrium. x^* is β-statistically close to x, using Lemma 4.1.8 can imply that the strategy x^*, is a $f(\beta)$-approximate L/F-equilibrium. Note that for our proof, we do not need the condition $x^* \in \mathcal{G}(x^*)$ due to the definition of G_{I} and G_{II}.

□

4.1.2 Second Remedial Model of [106] for Multi Leader Follower Games

First, we review the mathematical definition of the second remedial approach. Similar to the first remedial approach, this remedial approach assumes differentiability and the existence of gradients. To avoid repetition, we do not inspect the computational

[8] With respect to each norm, we can have different constants.

version of this approach as we will inspect a more general version later in Section 4.3.1.

We assume that we have two leaders who try to solve the following optimization problems (4.8) and (4.9), respectively.

$$\begin{aligned} &\text{Min}\, \phi_{\text{I}}\left(x_{\text{I}}, x_{\text{II}}^{*}\right) + \left(x_{\text{I}}\right)^{\top} D_{\text{I}} y_{\text{I}} \\ &\text{s.t } g_{\text{I}}(x_{\text{I}}, x_{\text{II}}^{*}) \leq 0 \\ &\text{and } h_{\text{I}}(x_{\text{I}}) = 0 \end{aligned} \tag{4.8}$$

and for the second leader we have:

$$\begin{aligned} &\text{Min}\, \phi_{\text{II}}\left(x_{\text{I}}^{*}, x_{\text{II}}\right) + \left(x_{\text{II}}\right)^{\top} D_{\text{II}} y_{\text{II}} \\ &\text{s.t } g_{\text{II}}(x_{\text{I}}^{*}, x_{\text{II}}) \leq 0 \\ &\text{and } h_{\text{II}}(x_{\text{II}}) = 0 \end{aligned} \tag{4.9}$$

Also, all followers $i \in [k]$ tries to find $y = (y_1, \ldots, y_i, \ldots, y_k) \in \mathbb{R}^n$:

$$\begin{aligned} &\text{Min}\, \frac{1}{2}\left(y^{\top} B y\right) + c^{\top} y - \left(x_{\text{I}}\right)^{\top} D_{\text{I}} y - \left(x_{\text{II}}\right)^{\top} D_{\text{II}} y \\ &\text{s.t } Ay + a = 0. \end{aligned} \tag{4.10}$$

For each $\text{i} \in \mathcal{I} = \{\text{I}, \text{II}\}$, $\phi_{\text{i}} : \mathbb{R}^{n_{\text{I}}+n_{\text{II}}}$ is a twice continuously differentiable and convex function with respect to the variable x_{i}. Furthermore, $g_{\text{i}} : X^{\text{I}} \times X^{\text{II}} \to \mathbb{R}^{m_{\text{i}}}$ is a convex function and $h_{\text{i}} : X^{\text{i}} \to \mathbb{R}^{l_{\text{i}}}$ is an affine function. Matrix $B \in \mathbb{R}^{n \times n}$ is symmetric and positive definite matrix, matrix $D_{\text{i}} \in \mathbb{R}^{n_{\text{i}} \times n}$ for each $\text{i} \in \mathcal{I}$ and $A \in \mathbb{R}^{p \times n}$ is a matrix with full row rank. And, $c \in \mathbb{R}^n$ and $a \in \mathbb{R}^p$ are constant vectors.

Problem (4.10) is a strictly convex quadratic programming problem with equality constraints. Using the equivalent KKT system of linear equations which finds a pair $(y, \lambda) \in \mathbb{R}^{n \times p}$ satisfying the following, we will reformulate the problem:

$$\begin{aligned} By + c - \left(D_{\text{I}}\right)^{T} x_{\text{I}} - \left(D_{\text{II}}\right)^{T} x_{\text{II}} + A^{T}\lambda = 0, \\ Ay + a = 0. \end{aligned}$$

A has full row rank. The uniqueness of the solution ensure the convexity of $Y_{sol}(x_{\text{I}}^{*}, x_{\text{II}})$. We will later show that adding monotonicity condition on the loss functions of the leaders will ensure the uniqueness of L/F-equilibrium. Under the given assumptions, a KKT pair (y, λ) exists uniquely for each $(x_{\text{I}}, x_{\text{II}})$ and is denoted by

$(y(x_{\mathrm{I}}, x_{\mathrm{II}}), \lambda(x_{\mathrm{I}}, x_{\mathrm{II}}))$. A straightforward calculation can show that:

$$\begin{aligned}
y(x_{\mathrm{I}}, x_{\mathrm{II}}) = &- B^{-1}c - B^{-1}A^{T}\left(AB^{-1}A^{T}\right)^{-1}\left(a - AB^{-1}c\right) \\
&+ \left[B^{-1}(D_{\mathrm{I}})^{T} - B^{-1}A^{T}\left(AB^{-1}A^{T}\right)^{-1}AB^{-1}(D_{\mathrm{I}})^{T}\right]x_{\mathrm{I}} \\
&+ \left[B^{-1}(D_{\mathrm{II}})^{T} - B^{-1}A^{T}\left(AB^{-1}A^{T}\right)^{-1}AB^{-1}(D_{\mathrm{II}})^{T}\right]x_{\mathrm{II}}, \\
\lambda(x_{\mathrm{I}}, x_{\mathrm{II}}) = &\left(AB^{-1}A^{T}\right)^{-1}\left(a - AB^{-1}c\right) + \left(AB^{-1}A^{T}\right)^{-1}AB^{-1}(D_{\mathrm{I}})^{T}x_{\mathrm{I}} \\
&+ \left(AB^{-1}A^{T}\right)^{-1}AB^{-1}(D_{\mathrm{II}})^{T}x_{\mathrm{II}}
\end{aligned}$$

In conclusion, for a given vector $x = (x_{\mathrm{I}}, x_{\mathrm{II}}) \subset \mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}}$, we can substitute the unique optimal response $y(x_{\mathrm{I}}, x_{\mathrm{II}})$ for $y_{\mathrm{I}} = y_{\mathrm{II}}$ in the leaders' problems. Then, we will show that the above-mentioned multi-leader-follower game can be reformulated in the format of a generalized Nash equilibrium problem with the constraints $\mathcal{R}^{\mathrm{I}}$ and $\mathcal{R}^{\mathrm{II}}$ which are defined by $\mathcal{R}^{\mathrm{I}}(x_{\mathrm{II}}) = \{x_{\mathrm{I}} : g_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}) \leq 0, h_{\mathrm{I}}(x_{\mathrm{I}}) = 0\}$ and $\mathcal{R}^{\mathrm{II}}(x_{\mathrm{I}}) = \{x_{\mathrm{II}} : g_{\mathrm{II}}(x_{\mathrm{I}}, x_{\mathrm{II}}) \leq 0, h_{\mathrm{II}}(x_{\mathrm{II}}) = 0\}$. Finally, leader I aims to find x_{I} such that:

$$\begin{aligned}
&\mathrm{Min}\ \Psi_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}^{*}) \\
&\mathrm{s.t}\ g_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}^{*}) \leq 0 \\
&\mathrm{and}\ h_{\mathrm{I}}(x_{\mathrm{I}}) = 0
\end{aligned} \tag{4.11}$$

And for leader II:

$$\begin{aligned}
&\mathrm{Min}\ \Psi_{\mathrm{II}}(x_{\mathrm{I}}^{*}, x_{\mathrm{II}}) \\
&\mathrm{s.t}\ g_{\mathrm{II}}(x_{\mathrm{I}}^{*}, x_{\mathrm{II}}) \leq 0 \\
&\mathrm{and}\ h_{\mathrm{II}}(x_{\mathrm{II}}) = 0
\end{aligned} \tag{4.12}$$

where Ψ_{I} and Ψ_{II} are defined:

$$\begin{aligned}
\Psi_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}^{*}) &= \phi_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}^{*}) + (x_{\mathrm{I}})^{\top} D_{\mathrm{I}} r + (x_{\mathrm{I}})^{\top} D_{\mathrm{I}} C x_{\mathrm{I}} + (x_{\mathrm{I}})^{\top} D_{\mathrm{I}} H x_{\mathrm{I}} \\
\Psi_{\mathrm{II}}(x_{\mathrm{I}}^{*}, x_{\mathrm{II}}) &= \phi_{\mathrm{II}}(x_{\mathrm{I}}^{*}, x_{\mathrm{II}}) + (x_{\mathrm{II}})^{\top} D_{\mathrm{II}} r + (x_{\mathrm{II}})^{\top} D_{\mathrm{II}} G x_{\mathrm{II}} + (x_{\mathrm{II}})^{\top} D_{\mathrm{II}} H x_{\mathrm{II}}
\end{aligned}$$

where $G \in \mathbb{R}^{n \times n_{\mathrm{I}}}, H \in \mathbb{R}^{n \times n_{\mathrm{II}}}$, and $r \in \mathbb{R}^{n}$ are defined as follows:

$$\begin{aligned}
G &= B^{-1}(D_{\mathrm{I}})^{\top} - B^{-1}A^{\top}\left(AB^{-1}A^{\top}\right)^{-1}AB^{-1}(D_{\mathrm{I}})^{\top} \\
H &= B^{-1}(D_{\mathrm{II}})^{\top} - B^{-1}A^{\top}\left(AB^{-1}A^{\top}\right)^{-1}AB^{-1}(D_{\mathrm{II}})^{\top} \\
r &= -B^{-1}c - B^{-1}A^{\top}\left(AB^{-1}A^{\top}\right)^{-1}\left(a - AB^{-1}c\right)
\end{aligned}$$

Proposition 4.1.9. *The problem of finding an approximate solution to remedial L/F-equilibrium denoted in (4.11) and (4.12) can be transformed into finding an approximate solution to a QVI problem.*

Proof. Proof is similar to Proposition 2.7.6 with the exception that $F : \mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}} \to \mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}}$ is defined as follows:

$$F(x) := \begin{pmatrix} \nabla_{x_{\mathrm{I}}} \psi_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}) \\ \nabla_{x_{\mathrm{II}}} \psi_{\mathrm{II}}(x_{\mathrm{I}}, x_{\mathrm{II}}) \end{pmatrix} = \begin{pmatrix} \nabla_{x_{\mathrm{I}}} \phi_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}) + D_{\mathrm{I}} r + 2D_{\mathrm{I}} G x_{\mathrm{I}} + D_{\mathrm{I}} H x_{\mathrm{II}} \\ \nabla_{x_{\mathrm{II}}} \phi_{\mathrm{II}}(x_{\mathrm{I}}, x_{\mathrm{II}}) + D_{\mathrm{II}} r + D_{\mathrm{II}} G x_{\mathrm{I}} + 2D_{\mathrm{II}} H x_{\mathrm{II}} \end{pmatrix}.$$

□

The reference [106] considers two leaders and one follower in the setting that the functions $g_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}^*) = g_{\mathrm{I}}(x_{\mathrm{I}})$ and $g_{\mathrm{II}}(x_{\mathrm{I}}^*, x_{\mathrm{II}}) = g_{\mathrm{II}}(x_{\mathrm{II}})$.

This indicates that here $X^{\mathrm{I}} = \{x_{\mathrm{I}} : g_{\mathrm{I}}(x_{\mathrm{I}}) \le 0, h_{\mathrm{I}}(x_{\mathrm{I}}) = 0\}$ and also the set $X^{\mathrm{II}} = \{x_{\mathrm{II}} : g_{\mathrm{II}}(x_{\mathrm{II}}) \le 0, h_{\mathrm{II}}(x_{\mathrm{II}}) = 0\}$ are the strategy sets for the leaders (no mutual constraint). Next, we recite some existence and uniqueness results on L/F-equilibrium for the multi-leader-follower game in this setting (a simpler form will also hold for the Nash equilibrium problem). The following lemmas will help us establish the uniqueness of L/F-equilibrium under the given conditions[9].

Lemma 4.1.10. *Suppose that F is a function and $\mathcal{R}$ is a set*[10]*. If F is strictly monotone on $\mathcal{R}$, i.e., $\forall x, y \in \mathcal{R}, x \neq y$ we have $(F(x) - F(y))^\top (x - y) > 0$, and $VI(\mathcal{R}, F)$ has at least one solution, then the solution is unique. Moreover, if F is strongly monotone on $\mathcal{R}$, i.e., there exists $\mu > 0$ such that $\forall x, y \in \mathcal{R}$, $(F(x) - F(y))^\top (x - y) \ge \mu \|x - y\|^2$, then there exists a unique solution to the $VI(\mathcal{R}, F)$.*

Corollary 4.1.11. *If the function $F(x) = (\nabla_{x_{\mathrm{I}}} \phi_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}), \nabla_{x_{\mathrm{II}}} \phi_{\mathrm{II}}(x_{\mathrm{I}}, x_{\mathrm{II}}))$ is strictly monotone and (4.11) and (4.12) have at least one solution, then the L/F-equilibrium is unique.*

Corollary 4.1.12. *If the function $F(x) = (\nabla_{x_{\mathrm{I}}} \phi_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}), \nabla_{x_{\mathrm{II}}} \phi_{\mathrm{II}}(x_{\mathrm{I}}, x_{\mathrm{II}}))$ is strictly monotone, then a unique solution is guaranteed which means the existence of a solution is not a necessary condition for uniqueness.*

[9] For more information, see [98, 106].

[10] Note that we assume that the set $\mathcal{R}$ is a well-defined set serving as the domain of the well-defined function F.

4.2 Remedial Solution to Multi-Leader-Follower Games with Uncertainty [107]

To the best of our knowledge, the only remedial approach for multi-leader-follower games with uncertainty studied in the literature is the extension of the second approach of [106]. We only review the second approach as we will investigate the more general version later.

4.2.1 Some General Existence Results

For the sake of completeness, we inspect a general existence result for robust L/F-Nash equilibrium for the multi-leader-follower games with uncertainty under some assumptions [107]. To prevent confusion due to multiple equivalent definitions, we recite the definition of L/F-robust Nash equilibrium which we reformulated the problem for leader I with uncertainty parameters e_{I} and u_{I}(see 2.13 and 2.14)[11]:

$$\begin{aligned}&\mathrm{Min}\, \phi_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}^{*}, y_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}^{*}, e_{\mathrm{I}}), u_{\mathrm{I}}\right)\\ &\text{s.t } x_{\mathrm{I}} \in X^{\mathrm{I}}\end{aligned}$$

and for leader II:

$$\begin{aligned}&\mathrm{Min}\, \phi_{\mathrm{II}}\left(x_{\mathrm{I}}^{*}, x_{\mathrm{II}}, y_{\mathrm{II}}(x_{\mathrm{I}}^{*}, x_{\mathrm{II}}, e_{\mathrm{II}}), u_{\mathrm{II}}\right)\\ &\text{s.t } x_{\mathrm{II}} \in X^{\mathrm{II}}\end{aligned}$$

Also, recall that utilizing the robust optimization paradigm, we defined a new game $\tilde{\mathcal{G}}$ with worst loss functions $\tilde{\psi}_{\mathrm{I}} : X^{\mathrm{I}} \times X^{\mathrm{II}} \to \mathbb{R}$ and $\tilde{\psi}_{\mathrm{II}} : X^{\mathrm{I}} \times X^{\mathrm{II}} \to \mathbb{R}$ (see Equation 2.15):

$$\begin{aligned}\tilde{\psi}_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}^{*}\right) &:= \sup\left\{\phi_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}^{*}, y_{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}^{*}, e_{\mathrm{I}}), u_{\mathrm{I}}\right) \mid u_{\mathrm{I}} \in U^{\mathrm{I}}, e_{\mathrm{I}} \in E^{\mathrm{I}}\right\}\\ \tilde{\psi}_{\mathrm{II}}\left(x_{\mathrm{I}}^{*}, x_{\mathrm{II}}\right) &:= \sup\left\{\phi_{\mathrm{II}}\left(x_{\mathrm{I}}^{*}, x_{\mathrm{II}}, y_{\mathrm{II}}(x_{\mathrm{I}}^{*}, x_{\mathrm{II}}, e_{\mathrm{II}}), u_{\mathrm{II}}\right) \mid u_{\mathrm{II}} \in U^{\mathrm{II}}, e_{\mathrm{II}} \in E^{\mathrm{II}}\right\}\end{aligned}$$

Assumption 1. *The following assumption will guarantee the existence of a robust L/F-equilibrium (see the optimization programs (2.13) and (2.14)).*

- *The functions* $\phi_{\mathrm{I}} : \mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}} \times \mathbb{R}^{n} \times \mathbb{R}^{l_{\mathrm{I}}} \to \mathbb{R}$ *and* $\phi_{\mathrm{II}} : \mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}} \times \mathbb{R}^{n} \times \mathbb{R}^{l_{\mathrm{II}}} \to \mathbb{R}$ *are both continuous.*

[11] Note that our definition of remedial models will be more general compared to this definition which comes from the original definition of robust L/F-Nash equilibrium in the literature without considering the generalized constrained format (we may call this problem robust L/F-generalized Nash equilibrium.)

- *The functions* $y_{\mathrm{I}} : \mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}} \times \mathbb{R}^{k_{\mathrm{I}}} \to \mathbb{R}^n$ *and* $y_{\mathrm{II}} : \mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}} \times \mathbb{R}^{k_{\mathrm{II}}} \to \mathbb{R}^n$ *are also continuous.*

- *The uncertainty sets* $U^{\mathrm{I}} \subseteq \mathbb{R}^{l_{\mathrm{I}}}$, $U^{\mathrm{II}} \subseteq \mathbb{R}^{l_{\mathrm{II}}}$, $E^{\mathrm{I}} \subseteq \mathbb{R}^{k_{\mathrm{I}}}$ *and* $E^{\mathrm{II}} \subseteq \mathbb{R}^{k_{\mathrm{II}}}$ *are nonempty and compact.*

- *The strategy sets* X^{I} *and* X^{II} *are nonempty, compact and convex.*

- *For* $\mathrm{i} \in \{\mathrm{I}, \mathrm{II}\}$, *the functions* $\phi_{\mathrm{i}}\left(x_{\mathrm{i}}, x^*_{-\mathrm{i}}, y_{\mathrm{i}}(x_{\mathrm{i}}, x^*_{-\mathrm{i}}, e_{\mathrm{i}}), u_{\mathrm{i}}\right)$ *is convex for any fixed* u_{i}, $x^*_{-\mathrm{i}}$ *and* e_{i}.

Under Assumption 1, we have the following property for correspondences $\tilde{\psi}_{\mathrm{I}}$ and $\tilde{\psi}_{\mathrm{II}}$ defined in 2.15:

Proposition 4.2.1 ([107])**.** *For each leader* $\mathrm{i} \in \{\mathrm{I}, \mathrm{II}\}$, *under Assumption 1 we have that:*

- $\tilde{\psi}_{\mathrm{i}}(x)$ *is finite for any* $x \in X^{\mathrm{I}} \times X^{\mathrm{II}}$.

- $\tilde{\psi}_{\mathrm{i}} : \mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}} \to \mathbb{R}$ *is also continuous.*

- $\tilde{\psi}_{\mathrm{i}}(\cdot, x_{-\mathrm{i}})$ *is convex on* X^{i} *for any fixed* $x_{-\mathrm{i}}$

Theorem 4.2.2 ([107])**.** *If Assumption 1 holds, then the multi-leader-follower game with uncertainty comprising problems (2.13) and (2.14) has at least one robust L/F-Nash equilibrium.*

Note that a similar existence result holds for the basic Nash equilibrium problem with uncertainty (see Theorem 3.3 in [147]). Note that instead of viewing y_{I} and y_{II} as functions, we used the correspondences F_{I} and F_{II} and the constraints G_{I} and G_{II} in the first approach.

The Remedial Approach

Given above strict assumptions above, the paper [107] extended the results of [106] to cover uncertainty. For this approach, we assume that one can formulate a robust L/F equilibrium problem as follows:

$$\begin{aligned} &\operatorname{Min} \phi_{\mathrm{I}}\left(x_{\mathrm{I}}, x^*_{\mathrm{II}}, y, u_{\mathrm{I}}\right) := \omega_{\mathrm{I}}\left(x_{\mathrm{I}}, x^*_{\mathrm{II}}, , u_{\mathrm{II}}\right) + \pi_{\mathrm{I}}\left(x_{\mathrm{I}}, y\right) \\ &\text{ s.t } g_{\mathrm{I}}\left(x_{\mathrm{I}}, x^*_{\mathrm{II}}\right) \leq 0, \quad h_{\mathrm{I}}\left(x_{\mathrm{I}}\right) = 0. \end{aligned}$$

And for Leader II:

$$\begin{aligned}&\text{Min}\, \phi_{\text{II}}\left(x_{\text{I}}, x_{\text{II}}^{*}, y, u_{\text{II}}\right) := \omega_{\text{II}}\left(x_{\text{I}}, x_{\text{II}}^{*}, u_{\text{II}}\right) + \pi_{\text{II}}\left(x_{\text{II}}, y\right)\\ &\text{s.t } g_{\text{II}}\left(x_{\text{I}}^{*}, x_{\text{II}}\right) \leq 0, \quad h_{\text{II}}\left(x_{\text{II}}\right) = 0.\end{aligned}$$

And each follower $i \in [k]$ will try to find y_i given the strategies of the leaders and other followers $(\bar{y}_{-i})$:

$$\begin{aligned}&\text{Min}\, \theta_i(x_{\text{I}}^{*}, x_{\text{II}}^{*}, y) := \gamma_i(y) - \pi_{\text{I}}\left(x_{\text{I}}^{*}, y\right) - \pi_{\text{II}}\left(x_{\text{II}}^{*}, y\right)\\ &\text{s.t } y_i \in Y_i.\end{aligned}$$

where $y = (y_i, \bar{y}_{-i}) \in \mathcal{Y}$ and $\mathcal{Y} \subseteq Y_1 \times \cdots \times Y_k \subseteq Y_{sol}$.

The reference [107] assumes the following explicit representation for π_{I} and π_{II} for leaders[12]. Also, for each follower $i \in [k]$, γ_i and Y_i will have an explicit representation:

$$\begin{aligned}\omega_{\text{I}}\left(x_{\text{I}}, x_{\text{II}}, u_{\text{I}}\right) &:= \frac{1}{2}\left(x_{\text{I}}\right)^{\top} H_{\text{I}} x_{\text{I}} + \left(x_{\text{I}}\right)^{\top} N_{(\text{I},\text{II})} x_{\text{II}} + \left(x_{\text{I}}\right)^{\top} R_{\text{I}} u_{\text{I}},\\ \omega_{\text{II}}\left(x_{\text{I}}, x_{\text{II}}, u_{\text{II}}\right) &:= \frac{1}{2}\left(x_{\text{II}}\right)^{\top} H_{\text{II}} x_{\text{II}} + \left(x_{\text{II}}\right)^{\top} N_{(\text{II},\text{I})} x_{\text{I}} + \left(x_{\text{II}}\right)^{\top} R_{\text{II}} u_{\text{II}},\\ \pi_{\text{I}}\left(x_{\text{I}}, y\right) &:= \left(x_{\text{I}}\right)^{\top} D_{\text{I}} y,\\ \pi_{\text{II}}\left(x_{\text{II}}, y\right) &:= \left(x_{\text{II}}\right)^{\top} D_{\text{II}} y,\\ \forall i \in [k],\ : \gamma_i(y) &:= \frac{1}{2} y^{\top} B y + c^{\top} y,\\ \mathcal{Y} &:= \left\{ y = (y_1, \ldots, y_k) \in \mathbb{R}^n \mid Ay + a = 0 \right\},\end{aligned}$$

where for i $\in \{\text{I}, \text{II}\}$, $H_{\text{i}} \in \mathbb{R}^{n_{\text{i}} \times n_{\text{i}}}$ is symmetric, $D_{\text{i}} \in \mathbb{R}^{n_{\text{i}} \times n}, R_{\text{i}} \in \mathbb{R}^{n_{\text{i}} \times l_{\text{i}}}, N_{(\text{I},\text{II})} \in \mathbb{R}^{n_{\text{I}} \times n_{\text{II}}}$, and $c \in \mathbb{R}^n$. Matrix $B \in \mathbb{R}^{n \times n}$ is assumed to be symmetric and positive definite. Moreover, $A \in \mathbb{R}^{p \times n}, a \in \mathbb{R}^p$, and A has full row rank. Note that it is obvious that in the followers' problem, this approach assumes the loss function for all followers is the same. Also, recall that each leader i cannot exactly know the follower $j \in [k]$'s problem but only can anticipate it as follows:

$$\begin{aligned}&\text{Min}\, \theta_{(\text{i},j)}\left(x_{\text{I}}^{*}, x_{\text{I}}^{*}, y_{\text{i}}, v_{\text{I}}\right) := \tfrac{1}{2} y_{\text{i}}^{\top} B y_{\text{i}} + \left(c + v_{\text{i}}\right)^{\top} y_{\text{i}} - \pi_{\text{I}}\left(x_{\text{I}}^{*}, y_{\text{i}}\right) - \pi_{\text{II}}\left(x_{\text{II}}^{*}, y_{\text{i}}\right)\\ &\text{s.t } y_{\text{i}} \in Y_i.\end{aligned}$$

[12] See Section 4 in [107] for more details specifically for the case when we have more than two leaders. The formulation is slightly different for that case. For consistency, we did not go beyond two leaders.

In the remainder, we inspect a restricted version of the remedial approach and also the uniqueness of a robust L/F Nash equilibrium for this special class of multi-leader-follower games with uncertainty, where each leader $\mathrm{i} \in \{\mathrm{I}, \mathrm{II}\}$ is assumed to solve the following optimization problem:

$$\begin{array}{l} \text{Min } \frac{1}{2}\left(x_{\mathrm{i}}\right)^{\top} H_{\mathrm{i}} x_{\mathrm{i}} + \left(x_{\mathrm{i}}\right)^{\top} N_{(\mathrm{i},-\mathrm{i})} x_{-\mathrm{i}} + \left(x_{\mathrm{i}}\right)^{\top} R_{\mathrm{i}} u_{\mathrm{i}} + \left(x_{\mathrm{i}}\right)^{\top} D_{\mathrm{i}} y \\ \text{s.t } x_{\mathrm{i}} \in X^{\mathrm{i}}, \end{array} \tag{4.13}$$

where y is an optimal solution of the following follower's problem anticipated by leader i :

$$\begin{array}{l} \text{Min } \frac{1}{2} y^{\top} B y + \left(c + e_{\mathrm{i}}\right)^{\top} y - \left(x_{\mathrm{I}}\right)^{\top} D_{\mathrm{I}} y - \left(x_{\mathrm{II}}\right)^{\top} D_{\mathrm{II}} y \\ \text{s.t } A y + a = 0, \end{array} \tag{4.14}$$

Since the followers' problems estimated by the two leaders are both strictly convex quadratic programming problems with equality constraints, each of them is equivalent to finding a pair $(y, \lambda) \in \mathbb{R}^n \times \mathbb{R}^p$ satisfying the following KKT system of linear equations for leaders I and II respectively:

$$\begin{aligned} B y_{\mathrm{I}} + c + e_{\mathrm{I}} - \left(D_{\mathrm{I}}\right)^T x_{\mathrm{I}} - \left(D_{\mathrm{II}}\right)^T x_{\mathrm{II}} + A^T \lambda = 0, \\ A y_{\mathrm{I}} + a = 0. \end{aligned}$$

$$\begin{aligned} B y_{\mathrm{II}} + c + e_{\mathrm{II}} - \left(D_{\mathrm{I}}\right)^T x_{\mathrm{I}} - \left(D_{\mathrm{II}}\right)^T x_{\mathrm{II}} + A^T \lambda = 0, \\ A y_{\mathrm{II}} + a = 0. \end{aligned}$$

Under the assumptions where we know that A has full row rank, a KKT pair $(y_{\mathrm{I}}, \lambda_{\mathrm{I}})$ exists uniquely for $(x_{\mathrm{I}}, x_{\mathrm{II}}, e_{\mathrm{I}})$ and shown by $\left(y_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}, e_{\mathrm{I}}\right), \lambda\left(x_{\mathrm{I}}, x_{\mathrm{II}}, e_{\mathrm{I}}\right)\right)$. A similar pair for leader II also exists. A straightforward calculation can show that for $\mathrm{i} \in \{\mathrm{I}, \mathrm{II}\}$:

$$\begin{aligned} y_{\mathrm{i}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}, e_{\mathrm{i}}\right) = & - B^{-1}(c + e_{\mathrm{i}}) - B^{-1} A^T \left(A B^{-1} A^T\right)^{-1} \left(a - A B^{-1}(c + e_{\mathrm{i}})\right) \\ & + \left[B^{-1}\left(D_{\mathrm{I}}\right)^T - B^{-1} A^T \left(A B^{-1} A^T\right)^{-1} A B^{-1} \left(D_{\mathrm{I}}\right)^T\right] x_{\mathrm{I}} \\ & + \left[B^{-1}\left(D_{\mathrm{II}}\right)^T - B^{-1} A^T \left(A B^{-1} A^T\right)^{-1} A B^{-1} \left(D_{\mathrm{II}}\right)^T\right] x_{\mathrm{II}}, \\ \lambda_{\mathrm{i}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}, e_{\mathrm{i}}\right) = & \left(A B^{-1} A^T\right)^{-1} \left(a - A B^{-1}(c + e_{\mathrm{i}})\right) + \left(A B^{-1} A^T\right)^{-1} A B^{-1} \left(D_{\mathrm{I}}\right)^T x_{\mathrm{I}} \\ & + \left(A B^{-1} A^T\right)^{-1} A B^{-1} \left(D_{\mathrm{II}}\right)^T x_{\mathrm{II}}. \end{aligned}$$

Let $P = I - B^{-\frac{1}{2}} A^{\top} \left(A B^{-1} A^{\top}\right)^{-1} A B^{-\frac{1}{2}}$. By substituting each $y_{\mathrm{i}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}, e_{\mathrm{i}}\right)$ for y_{i} in the respective leader's problem, the objective function of each leader i can be reformulated as:

$$\begin{aligned}
&\psi_{\mathrm{i}}\left(x_{\mathrm{i}}, x_{-\mathrm{i}}, e_{\mathrm{i}}, u_{\mathrm{i}}\right) \\
&:= \phi_{\mathrm{i}}\left(x_{\mathrm{i}}, x_{-\mathrm{i}}, y_{\mathrm{i}}\left(x_{\mathrm{i}}, x_{-\mathrm{i}}, e_{\mathrm{i}}\right), u_{\mathrm{i}}\right) \\
&= \frac{1}{2}\left(x_{\mathrm{i}}\right)^{\top} H_{\mathrm{i}} x_{\mathrm{i}} + \left(x_{\mathrm{i}}\right)^{\top} D_{\mathrm{i}} Q_{\mathrm{i}} x_{\mathrm{i}} + \left(x_{\mathrm{i}}\right)^{\top} R_{\mathrm{i}} u_{\mathrm{i}} + \left(x_{\mathrm{i}}\right)^{\top} D_{\mathrm{i}} r \\
&+ \left(x_{\mathrm{i}}\right)^{\top}\left(D_{\mathrm{i}} Q_{-\mathrm{i}} + N_{(\mathrm{i},-\mathrm{i})}\right) x_{-\mathrm{i}} - \left(x_{\mathrm{i}}\right)^{\top} D_{\mathrm{i}} B^{-\frac{1}{2}} P B^{-\frac{1}{2}} e_{\mathrm{i}}.
\end{aligned}$$

Here, $Q_{\mathrm{I}} \in \mathbb{R}^{n \times n_{\mathrm{I}}}, Q_{\mathrm{II}} \in \mathbb{R}^{n \times n_{\mathrm{II}}}$, and $r \in \mathbb{R}^{n}$ are given by:

$$\begin{aligned}
Q_{\mathrm{I}} &= B^{-\frac{1}{2}} P B^{-\frac{1}{2}}\left(D_{\mathrm{I}}\right)^{T}, \\
Q_{\mathrm{II}} &= B^{-\frac{1}{2}} P B^{-\frac{1}{2}}\left(D_{\mathrm{II}}\right)^{T}, \\
r &= -B^{-\frac{1}{2}} P B^{-\frac{1}{2}} c - B^{-1} A^{T}\left(A B^{-1} A^{T}\right)^{-1} a.
\end{aligned}$$

With the function ψ_{i} for each leader i, we can transform the above multi-leader single-follower game with uncertainty to a Nash equilibrium problem with complete information using a robust optimization technique similarly:

$$\begin{aligned}
\operatorname*{Min}_{x_{\mathrm{i}}} \quad & \tilde{\psi}_{\mathrm{i}}\left(x_{\mathrm{i}}, x_{-\mathrm{i}}\right) \\
& \text{s.t } x_{\mathrm{i}} \in X^{\mathrm{i}}.
\end{aligned}$$

Here, $\tilde{\psi}_{\mathrm{i}} : \mathbb{R}^{n_{\mathrm{i}}} \times \mathbb{R}^{n_{-\mathrm{i}}} \rightarrow \mathbb{R}$ is defined by:

$$\begin{aligned}
\tilde{\psi}_{\mathrm{i}}\left(x_{\mathrm{i}}, x_{-\mathrm{i}}\right) := & \sup\left\{\psi_{\mathrm{i}}\left(x_{\mathrm{i}}, x_{-\mathrm{i}}, e_{\mathrm{i}}, u_{\mathrm{i}}\right) \mid u_{\mathrm{i}} \in U^{\mathrm{i}}, e_{\mathrm{i}} \in E^{\mathrm{i}}\right\} \\
= & \frac{1}{2}\left(x_{\mathrm{i}}\right)^{\top} H_{\mathrm{i}} x_{\mathrm{i}} + \left(x_{\mathrm{i}}\right)^{\top} D_{\mathrm{i}} Q_{\mathrm{i}} x_{\mathrm{i}} + \left(x_{\mathrm{i}}\right)^{\top} D_{\mathrm{i}} r \\
& + \left(x_{\mathrm{i}}\right)^{\top}\left(D_{\mathrm{i}} Q_{-\mathrm{i}} + N_{(\mathrm{i},-\mathrm{i})}\right) x_{-\mathrm{i}} + \phi_{\mathrm{i}}\left(x_{\mathrm{i}}\right)
\end{aligned}$$

where for each $\mathrm{i} \in \{\mathrm{I}, \mathrm{II}\}$, $\phi_{\mathrm{i}} : \mathbb{R}^{n_{\mathrm{i}}} \rightarrow \mathbb{R}$ is defined by:

$$\begin{aligned}
\phi_{\mathrm{i}}\left(x_{\mathrm{i}}\right) := & \sup\left\{\left(x^{\mathrm{i}}\right)^{\top} R_{\mathrm{i}} u_{\mathrm{i}} \mid u_{\mathrm{i}} \in U^{\mathrm{i}}\right\} \\
& + \sup\left\{-\left(x^{\mathrm{i}}\right)^{\top} D_{\mathrm{i}} B^{-\frac{1}{2}} P B^{-\frac{1}{2}} e_{\mathrm{i}} \,\middle|\, e_{\mathrm{i}} \in E^{\mathrm{i}}\right\}.
\end{aligned}$$

Theorem 4.2.3 ([107]). *Suppose that for each* $\mathrm{i} \in \{\mathrm{I}, \mathrm{II}\}$, *the strategy set* X^{i} *is nonempty, compact, and convex, the matrix* $H_{\mathrm{i}} \in \mathbb{R}^{n_{\mathrm{i}} \times n_{\mathrm{i}}}$ *is symmetric and positive semi-definite, and the uncertainty sets* U^{i} *and* E^{i} *are nonempty and compact.*

Then, the multi-leader-follower game with uncertainty comprising optimization program (4.13) has at least one robust L/F-equilibrium.

To establish uniqueness, let us define $N_{\mathrm{I}} = N_{(\mathrm{I,II})}$ and $N_{(\mathrm{II,I})}$ and:

$$\mathcal{J} := \begin{pmatrix} H_{\mathrm{I}} & N_{\mathrm{I}} \\ N_{\mathrm{II}} & H_{\mathrm{II}} \end{pmatrix} \tag{4.15}$$

Theorem 4.2.4 ([107])**.** *Suppose that matrix $\mathcal{J}$ is positive definite and the uncertainty sets U^{i} and E^{i} are nonempty and compact. Then the multi-leader-follower game with uncertainty comprising problems (4.13) and (4.14) has a unique robust L/F-equilibrium.*

4.3 Our General Remedial Approach

We introduce a general formulation that is inspired by both of these approaches and can encompass both remedial models[13]. First, we provide a mathematical definition for a general remedial approach dealing with multi-leader-follower games with uncertainty. Assume that we have two leaders and k followers. We first assume that we have the following optimization programs from the perspective of the leaders:

$$\begin{aligned}
&\operatorname{Min} \phi_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}^*, y, u_{\mathrm{I}}\right) := \omega_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}^*, u_{\mathrm{I}}\right) + \pi_{\mathrm{I}}\left(x_{\mathrm{I}}, y\right) \\
&\text{s.t } g_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}^*\right) \leq 0, \quad h_{\mathrm{I}}\left(x_{\mathrm{I}}\right) = 0. \\
&\operatorname{Min} \phi_{\mathrm{II}}\left(x_{\mathrm{I}}^*, x_{\mathrm{II}}, y, u_{\mathrm{II}}\right) := \omega_{\mathrm{II}}\left(x_{\mathrm{I}}^*, x_{\mathrm{II}}, u_{\mathrm{II}}\right) + \pi_{\mathrm{II}}\left(x_{\mathrm{II}}, y\right) \\
&\text{s.t } g_{\mathrm{II}}\left(x_{\mathrm{I}}^*, x_{\mathrm{II}}\right) \leq 0, \quad h_{\mathrm{II}}\left(x_{\mathrm{II}}\right) = 0.
\end{aligned} \tag{4.16}$$

where $u_{\mathrm{I}} \in U^{\mathrm{I}}$ and $u_{\mathrm{II}} \in U^{\mathrm{II}}$ are uncertainty parameters of the leaders. The uncertainty sets are convex and compact. Function $\omega_{\mathrm{I}} : \mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}} \to \mathbb{R}$ is convex for fixed x_{II} and u_{I} (similar assumption holds for ω_{II}). Also, for each $\mathrm{i} \in \{\mathrm{I}, \mathrm{II}\}$, we have the assumption that $g_{\mathrm{i}} : X^{\mathrm{I}} \times X^{\mathrm{II}} \to \mathbb{R}^{m_{\mathrm{i}}}$ is a convex function for fixed $x_{-\mathrm{i}}$ and $h_{\mathrm{i}} : X^{\mathrm{i}} \to \mathbb{R}^{l_{\mathrm{i}}}$ is an affine function.

From the perspective of the followers, each follower $i \in [k]$ will try to find y_i given the strategies of the leaders and other followers ($\bar{y}_{-i}$):

$$\begin{aligned}
&\operatorname{Min} \theta_i(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*, y) := \gamma_i(y) - \pi_{\mathrm{I}}\left(x_{\mathrm{I}}^*, y\right) - \pi_{\mathrm{II}}\left(x_{\mathrm{II}}^*, y\right) \\
&\text{s.t } g_i'(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*, y) \leq 0, \quad h_i'\left(y_i\right) = 0
\end{aligned}$$

[13]This obviously covers the case without uncertainty, in which the uncertainty sets are singletons.

where $y = (y_i, \bar{y}_{-i})$. For follower i, we consider γ_i, g'_i and h'_i to have an explicit representation[14]. Also, g'_i and h'_i for $i \in [k]$ have similar assumptions such as convexity.

Instead of $y_i \in \mathcal{R}_i(\bar{y}_{-i})$ as in traditional L/F-equilibrium, we take a similar approach combining the ideas of [106, 107, 150], restricting or relaxing followers' feasible solutions. We assume that a valid solution $\bar{y}_i$ for follower i belongs to a set Z_i that we call the *restricted or relaxed feasible set of follower* i: Assuming $Z = \Pi_{i=1}^{k} Z_i$, we can either consider $Z \subseteq Y_{sol}$ or $Y_{sol} \subseteq Z$[15].

As mentioned, each leader's anticipation of the strategies of their followers could be different. Also, the uncertainty parameter of the followers needs to be taken into account. In conclusion, leaders I and II anticipate that each follower will solve the following optimization problems, respectively:

$$\begin{aligned}
&\text{Min}\, \theta_{(\mathrm{I},i)}(x^*_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{\mathrm{I}}, e_{\mathrm{I}}) := \gamma_i(y_{\mathrm{I}}, e_{\mathrm{I}}) - \pi_{\mathrm{I}}\,(x^*_{\mathrm{I}}, y_{\mathrm{I}}) - \pi_{\mathrm{II}}\,(x^*_{\mathrm{II}}, y_{\mathrm{I}}) \\
&\text{s.t } y_{(\mathrm{I},i)} \in Z_i \\
&\text{Min}\, \theta_{(\mathrm{II},i)}(x^*_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{\mathrm{II}}, e_{\mathrm{II}}) := \gamma_i(y_{\mathrm{II}}, e_{\mathrm{II}}) - \pi_{\mathrm{I}}\,(x^*_{\mathrm{I}}, y_{\mathrm{II}}) - \pi_{\mathrm{II}}\,(x^*_{\mathrm{II}}, y_{\mathrm{II}}) \\
&\text{s.t } y_{(\mathrm{II},i)} \in Z_i
\end{aligned} \tag{4.17}$$

where $y_{\mathrm{I}} = (y_{(\mathrm{I},i)}, \bar{y}_{(\mathrm{I},-i)})$ and $y_{\mathrm{II}} = (y_{(\mathrm{II},i)}, \bar{y}_{(\mathrm{II},-i)})$. As observed, each remedial approach has a specific feasible set and relaxation method, aiming to simplify the followers' problem into a feasible solution set that is convex. In conclusion, a remedial version (such as approaches in [107, 150]) can be formulated in the format of a generalized (Nash) equilibrium problem efficiently (in polynomial time).

$$\begin{aligned}
&\text{Min}\, F_{\mathrm{I}}\,(x_{\mathrm{I}}, x^*_{\mathrm{II}}, y_{\mathrm{I}}, u_{\mathrm{I}}, e_{\mathrm{I}}) \\
&\text{s.t } \ (x_{\mathrm{I}}, y_{\mathrm{I}}) \in G_{\mathrm{I}}\,(x^*_{\mathrm{II}}, e_{\mathrm{I}})
\end{aligned} \tag{4.18}$$

and for leader II, we have:

$$\begin{aligned}
&\text{Min}\, F_{\mathrm{II}}\,(x^*_{\mathrm{I}}, x_{\mathrm{II}}, y_{\mathrm{I}}, u_{\mathrm{II}}, e_{\mathrm{II}}) \\
&\text{s.t } \ (x_{\mathrm{II}}, y_{\mathrm{II}}) \in G_{\mathrm{II}}\,(x^*_{\mathrm{I}}, e_{\mathrm{II}})
\end{aligned} \tag{4.19}$$

By using the robust optimization approach, one can construct the robust counterpart

[14] For more information, see Section 4.2.

[15] This is similar to the approach of [150] (see Section 4.1.1).

defining $\tilde{F}_{\mathrm{I}} : X^{\mathrm{I}} \times X^{\mathrm{II}} \rightrightarrows \mathbb{R}$ and $\tilde{F}_{\mathrm{II}} : X^{\mathrm{I}} \times X^{\mathrm{II}} \rightrightarrows \mathbb{R}$:

$$\begin{aligned}
&\text{Min } \tilde{F}_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}^{*}\right) := Max\left\{F_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}^{*}, y_{\mathrm{I}}, u_{\mathrm{I}}, e_{\mathrm{I}}\right) \mid u_{\mathrm{I}} \in U^{\mathrm{I}}, e_{\mathrm{I}} \in E^{\mathrm{I}}\right\} \\
&\text{s.t } \left(x_{\mathrm{I}}, y_{\mathrm{I}}\right) \in G_{\mathrm{I}}\left(x_{\mathrm{II}}^{*}, e_{\mathrm{I}}\right) \\
&\text{Min } \tilde{F}_{\mathrm{II}}\left(x_{\mathrm{I}}^{*}, x_{\mathrm{II}}\right) := Max\left\{F_{\mathrm{II}}\left(x_{\mathrm{I}}^{*}, x_{\mathrm{II}}, y_{\mathrm{II}}, u_{\mathrm{II}}, e_{\mathrm{II}}\right) \mid u_{\mathrm{II}} \in U^{\mathrm{II}}, e_{\mathrm{II}} \in E^{\mathrm{II}}\right\} \\
&\text{s.t } \left(x_{\mathrm{II}}, y_{\mathrm{II}}\right) \in G_{\mathrm{II}}\left(x_{\mathrm{I}}^{*}, e_{\mathrm{II}}\right)
\end{aligned} \tag{4.20}$$

We approximate Sup with Max[16]. The following proposition is a simple extension of [107], transforming the above-mentioned problem into GQVI[17].

Proposition 4.3.1. *In a multi-leader-follower game, the problem of finding a remedial L/F-equilibrium of the optimization programs in (4.20) with convex and compact uncertainty sets can be transformed into a GQVI problem.*

Proof. The proof is similar to Proposition 4.1.4. According to the definition of a robust Nash equilibrium problem, this problem can be viewed as a Nash equilibrium problem with complete information. The famous transformation of [97, 150] can transform an instance of generalized Nash equilibrium to a QVI problem (see Proposition 2.7.6). Define the correspondence $\tilde{\mathcal{D}}(x) : \mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}} \times \mathbb{R}^{l_{\mathrm{I}}+l_{\mathrm{II}}} \rightarrow \mathcal{P}\left(\mathbb{R}^{n_{\mathrm{I}}+n_{\mathrm{II}}}\right)$ to be ($\mathcal{P}(A)$ denotes the power set of A):

$$\tilde{\mathcal{D}}(x) = \partial_{x_{\mathrm{I}}} \tilde{F}_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}\right) \times \partial_{x_{\mathrm{II}}} \tilde{F}_{\mathrm{II}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}\right)$$

Note that $\tilde{\mathcal{D}}(x)$ is a correspondence as the sub-differentials are not necessarily unique. We can formulate the problem by concatenating the first-order optimality conditions of all leaders' problems, where the goal is finding a vector x such that:

$$\exists w \in \tilde{\mathcal{D}}\left(x\right), \quad \left(y - x\right)^{\top} w + \epsilon \geq 0, \quad \forall y \in \mathcal{G}(x)$$

where $\mathcal{G}(x) = (G_{\mathrm{I}}(x_{\mathrm{II}}, e_{\mathrm{I}}), G_{\mathrm{II}}(x_{\mathrm{I}}, e_{\mathrm{II}}))$. □

[16] If the set is compact and the function continuous, then $sup = max$. In our setting, the sets are closed due to compactness. A continuous function on a compact set always attains its maximum.

[17] They reduced their specific remedial equilibrium problem to GVI (see [43]), not GQVI, while we need this extension to possibly include the remedial approach of [150].

4.3.1 Computational Formulation

We now provide a computational definition for remedial solutions to robust L/F-equilibrium in multi-leader-follower games and establish their PPAD-completeness. We will show that if the loss functions and the constraints are given by linear arithmetic circuits, we can compute $\tilde{F}_{\text{I}}$ and $\tilde{F}_{\text{II}}$ in polynomial time[18].

TRANSFORMED REMEDIAL L/F ROBUST EQUILIBRIUM

Input: Transformed remedial game $\tilde{\mathcal{G}}$ as input with all the following:

- Two linear arithmetic circuits representing the convex loss functions ($\tilde{F}_{\text{I}}$,$\tilde{F}_{\text{II}}$) for two leaders,
- Two linear arithmetic circuits representing strong separation oracles for G_{I} and G_{II} that represents the *remedial constrained domain* of the strategies of the leaders I, II, and followers $i \in [k]$ that are two non-empty, convex, and compact correspondences,
- An accuracy parameter ϵ.

Output: One of the following cases:

- (Violation of non-emptiness): A certificate indicating at least one of the following cases is empty:
  - $G_{\text{I}}(x, e)$ for some $x \in X^{\text{II}}$ or $e \in E^{\text{I}}$
  - $G_{\text{II}}(x, e)$ for some $x \in X^{\text{I}}$ or $e \in E^{\text{II}}$
- (Violation of convexity): A certificate showing the loss functions are not convex.
- (Approximate minimization): Vectors $(x_{\text{I}}^*, y_{\text{I}}^*) \in G_{\text{I}}(x_{\text{II}}^*, e_{\text{I}})$ and $(x_{\text{II}}^*, y_{\text{II}}^*) \in G_{\text{II}}(x_{\text{I}}^*, e_{\text{II}})$ having the following relationship:
  - $\tilde{F}_{\text{I}}(x_{\text{I}}^*, x_{\text{II}}^*) \leq \epsilon + \underset{x_{\text{I}}}{\text{Min}}\ \tilde{F}_{\text{I}}(x_{\text{I}}, x_{\text{II}}^*)$
  - $\tilde{F}_{\text{II}}(x_{\text{I}}^*, x_{\text{II}}^*) \leq \epsilon + \underset{x_{\text{II}}}{\text{Min}}\ \tilde{F}_{\text{II}}(x_{\text{I}}^*, x_{\text{II}})$

---

[18] We can similarly provide a separation oracle for the sub-gradients of $\tilde{\mathcal{D}}(x)$ (see Algorithm 1).

Remark 4.3.2. Given the original game $\mathcal{G}$, the remedial game $\tilde{\mathcal{G}}$ can be constructed in polynomial time[19]. To ensure simplicity and clarity, we focused on the robust counterpart problem. The case (Violation of convexity) is meaningful as an output whenever the form of input is explicitly given; otherwise, convexity holds as a promise. Violation of non-emptiness is an exception that can be given due to the definition of the computational version of GQVI that has a direct relationship with Kakutani's problem and the ellipsoid algorithm. For this exception, the type of exception (first or second case) can be distinguished by a simple modification of the ellipsoid algorithm (see [119]).

Theorem 4.3.3. *Finding an approximate solution for the robust remedial L/F equilibrium (defined below) in a multi-leader-follower game with the above-mentioned assumptions for strong separation oracles is PPAD-complete.*

Proof Sketch. Inclusion in PPAD can be implied by careful adaptation of Proposition 4.3.1 as this problem can be transformed into a GQVI problem. The transformation can be done in polynomial time due to to construction of the GQVI instance in Proposition 4.3.1. The PPAD-hardness of this problem can be implied by the hardness of finding a mixed Nash equilibrium (see [45]) by which we can construct a game in which two leaders have loss functions that represent the expected loss of their mixed strategies. The followers have only one strategy and no restrictions and all the uncertainty sets are singletons.

□

4.3.2 Detailed Proof of Theorem 4.3.3

Before advancing to the proof, the following facts must be considered.

Constructing The Robust Counter-Part in Polynomial Time

It is noteworthy that the paper [107] did not consider the efficient computation of the constructed robust counterpart because it was focused only on the existence of solutions. In Equation 4.20, given strategies x_{I} and x^*_{II}, we must be able to calculate $\tilde{F}_{\mathrm{I}}\left(x_{\mathrm{I}}, x^*_{\mathrm{II}}\right)$ in polynomial time[20]. Finding an approximate solution for the minimum

[19] The remedial approaches in Section 4.2 and 4.1.1 can serve as good examples that provide two different transformations.

[20] Previously, we discussed remedial approaches indirectly find y_{I}

of a convex function is doable in polynomial time using simple and standard methods such as the ellipsoid algorithm. However, we would like to find a maximum solution of a convex function over a convex and compact domain. For a continuous function defined on a compact, closed, convex, and bounded domain, the Extreme Value Theorem guarantees that the function will attain both a maximum and a minimum value within that domain. Here, we only limit our attention to linear arithmetic circuits.

Lemma 4.3.4. *Let C be a linear arithmetic circuit of size s with g $\min / \max$ gates. Fix all inputs except two variables x, y, and let $D = [x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}]$ be a compact box with rational endpoints. Assume all circuit constants and the fixed inputs are rationals of polynomial bit-length B. Then $\max_{(x,y)\in D} C(x, y, \mathbf{z})$ can be computed in time polynomial in s and g (e.g., $O(g^2 \cdot s \cdot \text{poly}(B))$).*

Proof. Each $\min / \max$ gate induces an affine equality $f_1(x, y) = f_2(x, y)$, i.e., a line in the (x, y)-plane; the g such lines partition D into $O(g^2)$ polygonal regions. On any single region, the circuit simplifies to an affine function, whose maximum over that convex polygon is attained at a vertex. The total number of vertices is $O(g^2)$; evaluating C at one vertex takes $O(s)$ arithmetic operations. Hence, the algorithm that enumerates all vertices and evaluates C runs in $O(g^2 \cdot s)$ arithmetic steps. Under the stated bit-length assumptions, arithmetic on rationals can be done in time $\text{poly}(B)$, yielding the claimed polynomial-time bound. □

4.3.3 Adapting Proposition 4.3.1 and Establishing Proof of PPAD-completeness

We now proceed to apply Proposition 4.3.1, considering the computational restrictions and definitions. As mentioned, the linear arithmetic circuits that are piece-wise linear functions can perfectly approximate any well-behaved function. We also showed that, despite the non-differentiability, convexity is still preserved at kink points, and any local minimum is also a global minimum. This applies even if the minimum occurs at a kink point. Furthermore, we can similarly show that we can construct a strong separation oracle for $\tilde{\mathcal{D}}(x)$.

Finally, to establish inclusion in PPAD, we must consider solutions of the computational variant of the transformed GQVI from Proposition 4.3.1. The approach is similar to the proof of Theorem 4.1.6. For the final approximate solution, we apply Lemma 4.1.8 to the robust counterpart of the remedial L/F-equilibrium, as the remedial L/F-equilibrium can be written in the format of a generalized equilibrium problem.

4.4 Computational Complexity of Resilient Nash

We now proceed to the computational complexity of a total version of resilient Nash and related notions. Assume that we have a game $(\mathcal{G}, \mathcal{S}, \mathcal{U})$ with k players. When t is not a constant, neither t-resilient nor t-strong Nash (and also strong Nash when $t = k$ and coalition-proof Nash) are likely to be in PPAD according to their definitions (in particular, they are required to satisfy an exponential in t number of conditions). Therefore, we will assume that either the number of players in the game we discuss is constant or the parameter t is constant. Note that t-resilient Nash equilibrium is defined for $t \leq k-1$ [1].

Definition 4.4.1. *A game has the* t-coalition consistency property *if for every* $\mathbf{x} \in \mathcal{S}$ *and for all* $J_1, J_2 \subseteq [k]$ *with* $|J_1|, |J_2| \leq t$ *and* $J_1 \cap J_2 \neq \emptyset$, *there exists* $\mathbf{z} \in \mathcal{S}$ *such that* $\mathbf{z} \in \tilde{B}_{J_1}(\mathbf{x}_{-J_1}) \cap \tilde{B}_{J_2}(\mathbf{x}_{-J_2})$, *where for any coalition* $J \subseteq [k]$,

$$\tilde{B}_J(\mathbf{x}_{-J}) = \{s_J \in S_J \,|\, u_j(s_J, \mathbf{x}_{-J}) + \epsilon \geq u_j(s'_J, \mathbf{x}_{-J}),\ \forall s'_J \in S_J,\ \forall j \in J\}\,.$$

That is, $\tilde{B}_J(\mathbf{x}_{-J})$ *is the set of strategy profiles for coalition* J *that are* ϵ*-best responses to the strategies of players outside* J*, ensuring that no member of* J *has an incentive to change their strategy unilaterally.*

Remark 4.4.2. We do not require t-coalition consistency property if we either inspect the problem of 1-resilient Nash equilibrium or when we have k-multi-concavity as an assumption.

TOTAL t-RESILIENT NASH

Input: We receive as input all the following:

- A game $\mathcal{G}$ with t-coalition consistency property with k players represented by k linear arithmetic circuits that represent t-concave utility functions,
- The strategy sets $\mathcal{S}$,
- An accuracy parameter ϵ.

Output: One of the following cases:

- (Violation of t-multi-concavity): A subset $J \in \mathcal{J}$ of at most size t, one index j, three vectors[21] $\mathbf{x}_J, \mathbf{y}_J \in S_J$ and $\mathbf{x}_{-J} \in S_{-J}$ such that for some $\lambda \in (0,1)$:

$$u_j\left(\lambda \mathbf{x}_J + (1-\lambda)\mathbf{y}_J, \mathbf{x}_{-J}\right) < \lambda u_j\left(\mathbf{x}_J, \mathbf{x}_{-J}\right) + (1-\lambda)u_j\left(\mathbf{y}_J, \mathbf{x}_{-J}\right)$$

- (Violation of t-coalition consistency): One vector x and two sets $J_1, J_2 \subseteq [t]$ with $J_1 \cap J_2 \neq \emptyset$ indicating the emptiness of $\tilde{B}_{J_1}(\mathbf{x}_{-J_1}) \cap \tilde{B}_{J_2}(\mathbf{x}_{-J_2})$.
- One vector s^* which represents the strategy profile of all players that satisfies:

$$\left(\forall J \in \mathcal{J} \text{ s.t } |J| \leq t\right), \forall s'_J \in S_J, \forall j \in J : u_j\left(\mathbf{s}^*\right) + \epsilon \geq u_j\left(\mathbf{s}'_J, \mathbf{s}^*_{-J}\right)$$

Theorem 4.4.3. *The problem of finding a total t-resilient Nash equilibrium is PPAD-complete.*

Proof Sketch. We need to further generalize the techniques of [152] and also introduce a more general version of GQVIs appropriate for this application and address

[21] Recall that we declared that each player j will play from S_j and each player controls n_j variables. In computational game theory, the computational results relevant to strong and coalition-proof Nash equilibrium are usually restricted to mixed strategies and normal-form games. We kept the definitions closer to the generalized Nash equilibrium notion for consistency, which is also more general.

the existence of a solution for them. This stems from the fact that for a t-resilient Nash equilibrium with conditions such as multi-concavity, we need to involve exponentially many equations in the constant t (see the definition of t-resilient Nash equilibrium) and impose appropriate constraints on the variational inequality problem to reflect the constraints of resilient Nash equilibrium. Roughly, we show that a matrix form of variational inequalities (we call MGQVI) may still be solved using Kakutani's fixed point, by an appropriate transformation and careful consideration of the approximation errors.

Specifically, we need to make adjustments to the strong constrained optimization problem, the robust version of Berge's maximum theorem, and the sub-gradient ellipsoid central cut method. We provide a slightly generalized version of the sub-gradient ellipsoid central cut method that deals with vector-valued functions and also can provide meaningful emptiness exceptions required for the violation of emptiness, in Appendix B.3. The generalized version of the robust maximum Berge theorem we introduce can provide the Lipschitzness bound we require for the more generalized variational inequality problem that we introduce (Section B.1). The transformation from the t-resilient Nash equilibrium problem to our generalized version of variational inequalities is also more general compared to the famous transformation introduced by [97] (compare Corollary 2.7.7 and Lemma 4.4.7). Dealing with the non-differentiability of linear arithmetic circuits is similar to multi-leader-follower games[22]. For PPAD-hardness, we consider the problem of finding a mixed approximate Nash equilibrium (1-resilient Nash equilibrium) in which the expected payoff is concave. □

Corollary 4.4.4. *The problem of finding a t-strong Nash equilibrium ($t \leq k-1$ and t is a constant) and coalition-proof Nash equilibrium (with constant players) under multi-concavity and coalition consistency are PPAD-complete.*

Remark 4.4.5. The decision version of the t-resilient Nash equilibrium without the multi-concavity condition (even for normal-form games) is NP-hard in general. The details are available in Appendix B.5. This result can be derived by combining the established NP-hardness of computing an exact strong Nash equilibrium, as shown in [52], and using techniques similar to those provided in the next chapter for addressing the approximate version of the problem. This demonstrates that, although multi-

[22]We show the fact that the sub-gradient correspondence of a function represented by a linear arithmetic circuit is completely aligned with the generalization of the variational inequality problem we introduce (MGQVI). In other words, the sub-gradient correspondence satisfies the assumption required for transformation to MGQVI (see Lemma 4.4.10)

concavity is a reasonable assumption from the perspective of convex optimization and general games, it is quite restrictive.

4.4.1 Proof of Inclusion in PPAD

Finally, in the following lemma, we convert the problem of finding a t-resilient Nash equilibrium (for any constant t) to a series of variational equations as follows. Roughly speaking, the definition of t-resilient Nash equilibrium needs $\sum_{i=1}^{t}\binom{k}{i}$ conditions (immunity to deviations) to be checked, and through this lemma, we make sure that all of these conditions are satisfied by using a single variational inequality problem.

Remark 4.4.6. Note that the membership of a strategy profile in $\mathcal{S} = \Pi_{j\in[k]}S_j$ could be done in polynomial time. Therefore, a strong separation oracle for $\mathcal{S}$ exists.

Lemma 4.4.7. *Suppose that we have a game with multi-concave and continuously differentiable utilities $\mathcal{U} = (u_1, \ldots, u_k)$, and strategies $\mathcal{S} = (S_1, \ldots, S_k)$. The problem of finding an approximate t-resilient Nash equilibrium of this game can be reduced to an MVI (multi-variational inequality) problem.*

Proof. For simplicity, we only solve the case $t = 2$. Let $\theta_j = -u_j$ for each $j \in [k]$. Finding an approximate 2-resilient equilibrium in this game is equivalent to finding a solution to the following optimization problem, where the goal is finding $\mathbf{s}^*$ such that:

$$\begin{aligned} &\forall J \in \mathcal{J}, \forall j \in J, \quad \theta_j(\mathbf{s}^*) \leq \operatorname{Min} \theta_j(\mathbf{s}_J, \mathbf{s}^*_{-J}) + \epsilon \\ &\text{s.t } s_j \in S_J \end{aligned} \tag{4.21}$$

Define $F_i(x) = (\nabla_{x_{\pi(i,j)}} \theta_j(x))_{j=1}^{k} \in \mathbb{R}^k$ where the function π will help us indicate which one of the partial derivations should be considered. F is the matrix whose columns are F_i (for each i). And if $\pi(i,j) = 0$, $F(i,j) = 0$ will be 0.

$$\pi(i,j) = \begin{cases} i & \text{if } 1 \leq i \leq k \;\&\; i = j \\ b(i,j) & \text{if } k+1 \leq i \leq k+1+\binom{k}{2} \;\&\; b(i,j) \neq 0 \\ 0 & \text{o.w} \end{cases}$$

$$b(i,j) = \begin{cases} k - \operatorname{Mod}((i-k)-1, k-1) & \text{if } j = 1 + \lfloor ((i-k)-1)/(k-1) \rfloor, \\ 1 + \lfloor ((i-k)-1)/(k-1) \rfloor & \text{if } j = k - \operatorname{Mod}((i-k)-1, k-1), \\ 0 & \text{otherwise,} \end{cases}$$

We will show that $\mathbf{s}^*$ is a 2-resilient Nash equilibrium if and only if $\mathbf{s}^* \in \mathcal{S} = \Pi_{i\in[k]} S_i$ satisfies:

$$(\mathbf{y} - \mathbf{s}^*)^T F(\mathbf{s}^*) + \epsilon \mathbb{1} \geq 0, \quad \forall \mathbf{y} \in \mathcal{S} \tag{4.22}$$

Assume that we have an approximate solution for Equation 4.22. First, for each $J \in \mathcal{J}$ with one element (all $j \in [k]$), we have:

$$\theta_j(y_j, \mathbf{s}^*_{-j}) \geq \theta_j\left(s_j, \mathbf{s}^*_{-j}\right) + \left(y_j - s^*_j\right) \nabla_{y^*_j} \theta_j\left(s_j, \mathbf{s}^*_{-j}\right), \quad \forall y_j \in S_j$$

Similar to Proposition 2.7.6, $(y - \mathbf{s}^*)^T F(\mathbf{s}^*) + \epsilon \mathbb{1} \geq 0$, since $\mathbf{s}^*$ is a solution to MVI$(\nabla f, \mathcal{S})$. Then, we can conclude that s^* is a 1-resilient Nash equilibrium (which is also a Nash equilibrium):

$$\theta_j(y_j, \mathbf{s}^*_{-j}) + \epsilon \geq \theta_j(s_j, \mathbf{s}^*_{-j})$$

Next, for each coalition of size 2, for a pair (i, p) of players cannot favor themselves by forming a coalition and deviating from the equilibrium. By the fact that $\theta_j(\cdot, s_{-J})$ is 2-multi-convex for any $j \in [k]$, for any possible strategy $y_p \in S_p$ for player p, we have (note that either $j = p$ or $i = j$):

$$\theta_j(y_p, y_i, \mathbf{s}^*_{-\{i,p\}}) \geq \theta_j\left(\mathbf{s}^*_{\{i,p\}}, \mathbf{s}^*_{-\{i,p\}}\right) + (y_i - s^*_i) \nabla_{s^*_i} \theta_j\left(\mathbf{s}^*_{\{i,p\}}, \mathbf{s}^*_{-\{i,p\}}\right), \quad \forall y_i \in S_i$$

And also for any possible fixed strategy y_i for player i, we have:

$$\theta_j(y_p, y_i, \mathbf{s}^*_{-\{i,p\}}) \geq \theta_j\left(\mathbf{s}^*_{\{i,p\}}, \mathbf{s}^*_{-\{i,p\}}\right) + \left(y_p - s^*_p\right) \nabla_{s^*_p} \theta_j\left(\mathbf{s}^*_{\{i,p\}}, \mathbf{s}^*_{-\{i,p\}}\right), \quad \forall y_p \in S_p$$

Since $\mathbf{s}^*$ is a solution to MVI$(\nabla f, S)$, we have $(y - \mathbf{s}^*)^T F(\mathbf{s}^*) + \epsilon \mathbb{1} \geq 0$. This also means that the inequality $(y_i - s^*_i) \nabla_{s^*_i} \theta_j\left(\mathbf{s}^*_{\{i,p\}}, \mathbf{s}^*_{-\{i,p\}}\right) + \epsilon > 0$ and also inequality $\left(y_p - s^*_p\right) \nabla_{s^*_p} \theta_j\left(\mathbf{s}^*_{\{i,p\}}, \mathbf{s}^*_{-\{i,p\}}\right) + \epsilon > 0$ hold. Therefore, we can conclude that s^* is a 2-resilient Nash equilibrium as any two pairs (i, p) cannot form a coalition to get a better payoff:

$$\theta_j(\mathbf{s}^*_J, \mathbf{s}^*_{-J}) \leq \theta_j(\mathbf{y}_J, \mathbf{s}^*_{-J}) + \epsilon, \quad \forall y_J \in S_J \tag{4.23}$$

For the other direction, assume that x^* is a solution of Equation 4.22. We can use a similar approach used in Proposition 2.7.6 (corollary 2.7.7). □

Now we address the case where the functions are limited to linear arithmetic circuits that are not necessarily differentiable.

Proposition 4.4.8. *Suppose that we have a game with multi-concavity and given by linear arithmetic circuits $\mathcal{U} = (u_1, \ldots, u_k)$, and strategies $\mathcal{S} = (S_1, \ldots, S_k)$. The problem of finding an approximate t-resilient Nash equilibrium of this game can be reduced to an MGQVI (multi generalized-variational inequality) problem.*

Proof. The approach is similar to Theorem 4.1.6, which uses Lemma 4.1.8 on Proposition 4.1.4 to deal with the approximation error caused by the computational variant. Specifically, we need to apply Lemma 4.1.8 on Lemma 4.4.7, considering the following lemmas. □

Lemma 4.4.9. *Let $f = f_1, \ldots, f_r$ such that for each $i \in [r]$, f_i is a multi-convex function represented by a linear arithmetic circuit. Then, the correspondence $F(x) = (\partial f_1(x), \ldots, \partial f_r(x))$ is a multi-convex correspondence.*

Proof. Since each function f_i is multi-convex, we know that for any subset of variables $J \subseteq [m]$, fixing the remaining variables, the function remains convex in the selected variables. First, we show that the sub-gradient mapping $\partial f_i(x)$ of a multi-convex function satisfies the property that convex combinations of subgradients at two points also belong to the subgradient at their convex combination. That is, for any $y_J^1, y_J^2 \in F_J(x)$ and $\lambda \in [0,1]$, we have $\lambda \partial f_i(y_J^1, x_{-J}) + (1-\lambda)\partial f_i(y_J^2, x_{-J}) \subseteq \partial f_i(\lambda y_J^1 + (1-\lambda) y_J^2, x_{-J})$.

Given that f_i is constructed using linear arithmetic circuits (addition, scalar multiplication, max, and min), and these operations preserve multi-convexity, the subgradient mapping ∂f_i remains multi-convex. Next, we show $F(x) = (\partial f_1(x), \ldots, \partial f_n(x))$ preserves multi-convexity. Let $F = (F_1, F_2, \ldots, F_n)$ be the Cartesian product of the correspondences $F_i = \partial f_i$, and each f_i is a function represented by a multi-convex linear arithmetic circuit.

Let $F_J(x) = \prod_{i \in J} F_i(x)$ where $J \subseteq [n]$. For any subset $J \subseteq [n]$ and any two points $y_1, y_2 \in F_J(x)$, where y_1 and y_2 are convex combinations of points in the Cartesian product, we have:

$$\lambda F_J(y_1, x_{-J}) + (1-\lambda) F_J(y_2, x_{-J}) \subseteq F_J(\lambda y_1 + (1-\lambda) y_2, x_{-J}).$$

Since the Cartesian product of multi-convex correspondences satisfies the multi-convexity property for F_J any subset $J \in \mathcal{J}$, it follows that F is multi-convex as well.

□

Lemma 4.4.10. *Let f be a multi-convex function represented by an arithmetic circuit. Then ∂f is a multi-convex correspondence.*

Proof. We want to prove that for f, for any $J \subseteq \mathcal{J}$, y_J^1, y_J^2 and $\lambda \in [0,1]$, the sub-gradient correspondence satisfies:

$$\lambda \partial f(y_J^1, x_{-J}) + (1-\lambda)\partial f(y_J^2, x_{-J}) \subseteq \partial f(\lambda y_J^1 + (1-\lambda) y_J^2, x_{-J})$$

Since f is multi-convex, this means that it is convex in y_J for any fixed x_{-J}. This means that for any $g_J^1 \in \partial f(y_J^1, x_{-J})$ and $g_J^2 \in \partial f(y_J^2, x_{-J})$, we have the sub-gradient inequality:

$$f(y_J, x_{-J}) \geq f(y_J^1, x_{-J}) + \langle g_J^1, y_J - y_J^1 \rangle, \quad \forall y_J \in \mathbb{R}^{m|J|}$$

$$f(y_J, x_{-J}) \geq f(y_J^2, x_{-J}) + \langle g_J^2, y_J - y_J^2 \rangle, \quad \forall y_J \in \mathbb{R}^{m|J|}$$

Now, set $y_J^\lambda = \lambda y_J^1 + (1-\lambda) y_J^2$, and take a convex combination of the two inequalities:

$$f(y_J, x_{-J}) \geq \lambda f(y_J^1, x_{-J}) + (1-\lambda) f(y_J^2, x_{-J}) + \lambda \langle g_J^1, y_J - y_J^1 \rangle + (1-\lambda) \langle g_J^2, y_J - y_J^2 \rangle.$$

By convexity of f, we have $f(y_J^\lambda, x_{-J}) \leq \lambda f(y_J^1, x_{-J}) + (1-\lambda) f(y_J^2, x_{-J})$. Thus, a simple calculation gives:

$$f(y_J, x_{-J}) \geq f(y_J^\lambda, x_{-J}) + \lambda \langle g_J^1, y_J - y_J^1 \rangle + (1-\lambda) \langle g_J^2, y_J - y_J^2 \rangle.$$

Using the definition of sub-gradients, we rewrite the inequality to deduce that the convex combination $g_J = \lambda g_J^1 + (1-\lambda) g_J^2$ is also a sub-gradient:

$$f(y_J, x_{-J}) \geq f(y_J^\lambda, x_{-J}) + \langle \lambda g_J^1 + (1-\lambda) g_J^2, y_J - y_J^\lambda \rangle.$$

Finally, we have:

$$\lambda g_J^1 + (1-\lambda) g_J^2 \in \partial f(y_J^\lambda, x_{-J}).$$

□

Lemma 4.4.11. *Let a game with S_j as the strategy sets have utilities $u_j : S_j \to \mathbb{R}$, $j \in [k]$, that are t-multi-concave and Lipschitz continuous, and satisfy the t-coalition consistency property with $\epsilon \geq 0$. Define the MGQVI with $\mathcal{R}(\mathbf{x}) = \prod_j S_j$ and $\mathcal{F}_j(\mathbf{x})$ be defined according to Lemma 4.4.7 but considering the sub-gradients instead of gradient. Then, the t-coalition consistency property ensures the existence of a convex function $g : S \to \mathbb{R}$ (a point for $\epsilon = 0$) such that the subgradient compatibility condition:*

$$\partial g(\mathbf{x}) \subseteq \bigcap_{j=1}^{k} \mathcal{F}_j(\mathbf{x}), \quad \forall \mathbf{x} \in S$$

holds as a special case when the ϵ-best response sets $\tilde{B}_J(\mathbf{x}_{-J})$ for overlapping coalitions J_1, J_2 with $|J_1|, |J_2| \leq t$ and $J_1 \cap J_2 \neq \emptyset$ admit a common strategy $\mathbf{z} \in S$ that aligns with the minimizer of g. This allows the t-resilient Nash equilibrium problem to be reduced to solving the MGQVI.

4.4.2 Complexity of Robust and Resilient Nash Equilibrium

We now proceed to the computational complexity of the robust version of resilient Nash equilibrium. Assume that we have the game $\mathcal{G}$ with k players and given the uncertainty sets $U = (U_1, \ldots, U_k)$. Similarly, we can transform the game $\mathcal{G}$ to the robust counterpart $\tilde{\mathcal{G}}$ (see Equation (2.8)) and consider the following problem:

Total Robust and t-Resilient Nash

Input: We receive the robust counter part $\tilde{\mathcal{G}}$ as input all the following:

- k linear arithmetic circuits representing k convex loss functions,
- k linear arithmetic circuits representing strategy sets of S_i for $i \in [k]$,
- An accuracy parameter ϵ.

Output: One of the following cases:

- (Violation of t-multi-convexity): A subset $J \in \mathcal{J}$ of at most size t, one index j, three vectors $\mathbf{x}_J, \mathbf{y}_J \in S_J$ and $\mathbf{x}_{-J} \in S_{-J}$ such that for some $\alpha \in (0,1)$:

$$\tilde{\theta}_j\left(\alpha \mathbf{x}_J + (1-\alpha)\mathbf{y}_J, \mathbf{x}_{-J}\right) > \alpha \cdot \tilde{\theta}_j\left(\mathbf{x}_J, \mathbf{x}_{-J}\right) + (1-\alpha) \cdot \tilde{\theta}_j\left(\mathbf{y}_J, \mathbf{x}_{-J}\right)$$

- (Violation of t-coalition consistency): One vector x and two sets $J_1, J_2 \subseteq [t]$ with $J_1 \cap J_2 \neq \emptyset$ indicating the emptiness of $\tilde{B}_{J_1}(\mathbf{x}_{-J_1}) \cap \tilde{B}_{J_2}(\mathbf{x}_{-J_2})$.
- One vector s^* which represents the strategy profile of all players that satisfies:

$$\left(\forall J \in \mathcal{J} \text{ s.t } |J| \leq t\right), \forall s'_J \in S_J, \forall j \in J : \tilde{\theta}_j\left(\mathbf{s}^*\right) - \epsilon \leq \tilde{\theta}_j\left(\mathbf{s}'_J, \mathbf{s}^*_{-J}\right)$$

Theorem 4.4.12. *The problem of finding a total robust and t-resilient Nash equilibrium is PPAD-complete.*

Remark 4.4.13. Here, we defined the notion of t-coalition consistency for the robust counterpart, considering the minimization version. We can similarly extend this definition for multi-leader-follower games and introduce the concept of robust and resilient remedial L/F-equilibrium.

Chapter 5

Disjointness and Farness as Non-Convex Constraints

One of the foundational challenges in game theory is making the concept of mixed strategies more useful in practice [168]. There has been criticism of the concept of mixed strategies since the 1980s, since this concept may weaken the definition of Nash equilibrium [12, 164]. According to [53], mixed strategies are employed in games without a pure Nash equilibrium, but their primary purpose is to achieve noteworthy mathematical results. Another perspective on mixed strategies comes from behavioral game theory, which incorporates players' behavioral perspectives and psychological factors. While classical game theory assumes rationality and utility maximization, behavioral game theory incorporates social preferences and psychological factors, offering models that explain non-equilibrium behaviors [41, 48, 54, 141, 173]. One view treats them as beliefs about others' strategies rather than actions [9]; another sees them as representing populations of agents, where each agent plays a pure strategy and the resulting payoff then depends on the fraction of agents adopting each strategy, with the mixed strategy representing the distribution of pure strategies within the population [164]. In addition to all above-mentioned approaches, our approach (the constraints that we define) also considers strategies in mixed Nash equilibria as resources similar to [143] (for more information, see Section 5.1.2).

In this chapter, we mainly focus on the computational complexity of two forms of equilibrium with specific constraints, namely Nash and generalized equilibria with some form of constraints that are relevant to the strategies (specifically the support of the players) rather than their social welfare or payoff. Towards establishing hardness for these problems, we also consider some relevant problems involving equilibria with constraints. Informally stated, we show that approximate versions of disjoint Nash, partition Nash, and δ-far Nash and several related problems are NP-complete. Given the general landscape of Nash equilibria subject to constraints, these results are to be expected, although we show that their proofs require a significant extension of existing techniques.

As noted, existing work on generalized equilibria largely assumes convex constraints, where solutions are guaranteed [57, 62, 152, 161], and most problems we study lie outside this setting. We provide a variety of computational tractability and hardness results for generalized equilibria with strategic constraints. In contrast to the case of Nash equilibrium, the generalized setting is much more subtle, and the problems here vary considerably in their computational complexity. The results we obtain in this setting may represent a new direction in computational game theory (see also [57]).

5.1 Equilibria with Strategic Constraints

We now introduce the definitions and main theorems of problems related to equilibria with strategic constraints. For simplicity, we only investigate the computational aspects of two-player (bi-matrix) games, and we will assume that the utility matrices' entries are in the range $[0, 1]$ unless otherwise stated. The proofs can be easily extended to normal-form games with a constant many players. We consider the standard notion of *statistical distance* (*total variation distance*) between mixed strategies, viewed as probability distributions on pure strategies.

5.1.1 Problem Definitions

Definition 5.1.1. *Two (mixed) strategies $x \in \triangle_n$ and $y \in \triangle_n$ are δ-far if $||x-y||_1 \geq \delta$, where $||x-y||_1 = \sum_{i=1}^{n} |x_i - y_i|$.* δ-closeness *of strategies can be defined similarly.*

Remark 5.1.2. Note that for bi-matrix games, for the mixed strategies x and y for the row and the column player, we simply have $0 \leq ||x-y||_1 \leq 2$.

Definition 5.1.3. *A Nash equilibrium $\boldsymbol{x}^* = (x_1^*, \ldots, x_k^*)$ is called a* δ-far Nash equilibrium *if for all $i \neq j$, x_i^* and x_j^* are 2δ-far. The approximate version of this type of equilibrium will be denoted by (δ, ϵ)-far Nash equilibrium*[1].

[1]We normalized δ to be in the range $[0, 1]$ similarly to ϵ.

A special case of far Nash can also be introduced by prohibiting all players from using the strategies employed by all other players as follows:

Definition 5.1.4. *(Disjoint Nash): A strategy profile $\boldsymbol{x} = (x_1, \ldots, x_k)$ is a* disjoint Nash equilibrium *(or simply disjoint Nash) if it is a Nash equilibrium and $\bigcap_{i=1}^k P_i = \emptyset$ where $P_i = \text{Supp}(x_i)$*[2]*.*

We say that a mixed strategy profile $(x_1, \ldots, x_k)$ has a disjoint support profile *if for any player i and player j $(i \neq j)$ have disjoint supports.*

Remark 5.1.5. When $\delta = 0$, δ-far Nash is just the problem of finding a Nash equilibrium, and when $\delta = 1$, it is exactly a disjoint Nash equilibrium. The existence of a disjoint Nash is not necessarily guaranteed, e.g., consider a 2-player, 2-strategy game where $U_1(1,1) = U_2(1,1) = 1$ while for $(i,j) \neq (1,1)$, $U_1(i,j) = U_2(i,j) = 0$.

For the purpose of investigating one of the straightforward applications of equilibria with limited overlapping support, we define the following problem that can be derived by further constraining a disjoint Nash equilibrium. Analogously to Definition 2.3.3, we can define the approximate version of these problems (*ϵ-disjoint and ϵ-partition Nash*).

Definition 5.1.6. *(Partition Nash): A strategy profile $\boldsymbol{x} = (x_1, \ldots, x_k)$ is a* partition Nash equilibrium *if it is a disjoint Nash equilibrium and $\bigcup_{i=1}^k P_i = S$, where $P_i =$* $\text{Supp}(x_i)$*.*

We also consider generalized equilibria for this class of constraints. For example, in the generalized version of a disjoint Nash equilibrium (which we call a generalized disjoint equilibrium), a player cannot deviate to a mixed strategy whose support has a nonempty intersection with those of the other players.

Definition 5.1.7. *For bi-matrix games, a* generalized δ-far equilibrium *is a generalized equilibrium with common constraint:*

$$\mathcal{R}_\delta^n = \{(x, y) \in \Delta_n \times \Delta_n \mid ||x - y||_1 \geq 2\delta\}$$

while a generalized disjoint equilibrium *is a generalized equilibrium with common constraint:*

$$\mathcal{R}_{\text{disjoint}}^k = \{(x, y) \in \Delta_n \times \Delta_n \mid \langle x, y \rangle = 0\}$$

[2]Recall that $\text{Supp}(x_i)$ denotes the support of player i and has the strategies used by player i with non-zero probabilities. Recall that the strategies of all players are assumed to be in $S = [n]$.

where $\langle x, y\rangle$ denotes the inner product.

Remark 5.1.8. Assume that the column player plays y. For the row player (similarly for the column player), we can also consider the relations (see (**)):

$$\mathcal{R}_{\delta}^{n}(y) = \{x \in \Delta_n \mid ||x - y||_1 \geq 2\delta\}$$

$$\mathcal{R}_{\text{disjoint}}^{n}(y) = \{x \in \Delta_n \mid \langle x, y\rangle = 0\}$$

Remark 5.1.9. It is easy to observe that all games having exactly two pure strategies have a generalized disjoint equilibrium where both players play one pure strategy. The constrained strategy spaces that we have defined are not convex. As shown by Rosen [161] a solution is guaranteed for convex constraints. Recently, it has been proved that by using a computational version of Kakutani's fixed point theorem, finding an approximate equilibrium (a generalized equilibrium) for a game with convex constraints is PPAD-complete [152]. Our results show that a solution may be guaranteed when constraints are not convex.

In the generalized setting, at least one of the problems we propose, namely, approximate generalized disjoint equilibrium, is tractable (Theorem 5.2.3 (a)) and has at least one trivial solution. However, we show that one selfish player can degrade the social welfare of all other players. With this drawback in mind, we direct our attention to another type of generalized equilibrium by excluding the possibility of minor probabilities, which is defined as follows:

Definition 5.1.10. *A* (θ, δ)-restricted far equilibrium problem *is defined as follows: Given a bi-matrix game $\mathcal{G}$ find a generalized δ-far equilibrium such that all pure strategies played with a positive probability are played with probability greater than θ. We can relax this definition by using an additive approximation error ϵ and call it* $(\theta, \delta, \epsilon)$-restricted far equilibrium. *For $\delta = 1$, we call the problem* (θ, ϵ)-restricted disjoint equilibrium.

Other Related Problems

Analogously to notions of farness, one can define equilibria with closeness constraints, leading to the concept of a δ-close Nash equilibrium. This may also be viewed as a generalization of symmetric Nash equilibrium. Similar to δ-far Nash, the δ-close Nash problem is NP-hard to approximate. The problem of determining if a bi-matrix game

Table 5.1: Bach or Stravinsky and the opposite game

	Clara	
Ray	1,$\frac{1}{2}$	0,0
	0,0	$\frac{1}{2}$,1

	Clara	
Ray	0,0	1,$\frac{1}{2}$
	$\frac{1}{2}$,1	0,0

has more than one equilibrium is one of the well-known instances of Nash equilibria with constraints [52, 89, 135]. We observe that, given continuity considerations, every bi-matrix game has an infinite number of approximate Nash equilibria. As an approximate analog of this problem, one can consider the problem of discovering two equilibrium conditions with a given statistical distance between them [13, 64, 65]. In [13] (Theorem 1.4 and 4.1), it is shown that for a sufficiently small approximation error, an instance of this problem is as hard as finding a hidden clique of size $O(\log n)$ in the random graph $G(n, \frac{1}{2})$ with n vertices. Furthermore, under ETH (the exponential time hypothesis), a quasi-polynomial lower bound for a constant approximation error smaller than $\frac{1}{8}$ exists (see problem 3 in [64]). We show that this problem for bi-matrix games is NP-hard to approximate[3]. We also have another variant of Nash equilibrium for which only the exact problem was also investigated in [135]. This variant requires a Nash equilibrium such that no strategy can be played with minor probabilities.

5.1.2 Applications

To provide some motivation, we present three potential applications of our constraints, two of which concern the behavioral interpretation of equilibria.

5.1.3 Bach or Stravinsky

We have two simple examples that are based on the well-known "Bach or Stravinsky" game. In both games, there are two players with two strategies. In the first scenario, we have two players who must choose between two activities: a row player (Ray) and a column player (Clara). Ray prefers the first activity while Clara prefers the second one. In the first game, if they both choose the same activity, one will receive a payoff of $\frac{1}{2}$ and the other will get 1, based on their preference, and if they choose different

[3] A very obvious difference between this problem and far Nash equilibrium is that far Nash equilibrium focuses on the existence of one equilibrium in which the players' strategies are statistically far apart, while the other problem requires at least two Nash equilibria that are statistically far apart.

activities, they end up in a fight and both will receive 0. In the second game, the players get a payoff of 1 or $\frac{1}{2}$ (based on their preference) if they employ different strategies.

In the first game, there are two pure Nash equilibria, and both players play the same strategies. This game has a mixed Nash equilibrium $(\frac{2}{3}, \frac{1}{3})$ (for the row player) and $(\frac{1}{3}, \frac{2}{3})$ (for the column player) where $L1$ distance between the strategies is $\frac{2}{3}$. In addition, in the first game, we do not have a Nash equilibrium where two players choose different activities, indicating that they cannot reach a desirable payoff if both players exhibit uncooperative behavior. In the second game, we have two pure Nash equilibria in which both players play different strategies, while the unique mixed strategy equilibrium is similar to the previous game. Therefore, in this game, both players are inclined toward using different activities (strategies) based on their preferences. This simple example motivates the fact that the disjoint Nash could be of interest in the behavioral interpretation of the strategic choices[4].

5.1.4 Tax Cheats and Policy Design

Income tax games are well-known examples that incorporate two different interpretations of mixed strategies. In this section, we demonstrate how a far Nash equilibrium can provide useful information for policy designers. We illustrate this with two non-symmetric games modeling a simple tax auditing problem, shown in Table 5.2. Here, the auditor and the taxpayer have two possible strategies. The auditor's strategies are either auditing the taxpayer (strategy 1) or not auditing the taxpayer (strategy 2). The taxpayer can either pay their taxes truthfully (strategy 1) or cheat (strategy 2). These bi-matrix games have no pure Nash equilibrium. The first game's Nash equilibrium happens when the taxpayer pays their taxes truthfully with probability $\frac{2}{3}$ and cheats with probability $\frac{1}{3}$. The auditor in this equilibrium will audit the taxpayer with probability $\frac{3}{7}$ and will not audit with probability $\frac{4}{7}$. We can interpret this mixed strategy by assuming that the auditor is literally randomizing over the two strategies while the numbers that are related to the taxpayer indicate the proportion of taxpayers who are paying taxes and cheating, respectively. The goal of policy design could be designing a policy in which the taxpayer tries to deter cheating. In the second game with this new penalty of getting caught in place, the auditor will audit with

[4]Furthermore, in some situations, we might be only interested in finding a disjoint Nash equilibrium while other forms of equilibria exist.

Table 5.2: The income tax games

	Tax payer	
Auditor	2,0	4,-10
	4,0	0,4

	Tax payer	
Auditor	2,0	4,-20
	4,0	0,4

probability $\frac{1}{6}$ while the tax compliance rate will not change. In other words, in the Nash equilibrium of this game, the auditor will play $(\frac{1}{6}, \frac{5}{6})$, and the taxpayer will play $(\frac{2}{3}, \frac{1}{3})$. This shows that if we want to change tax compliance, we need to change the auditor's payoff. To get a better compliance rate, we can either increase the payoff of getting a cheater or lower the cost of auditing.

The policy designer who determines the penalties should consider some basic assumptions. For example, the penalty should not be extremely high. Auditing with a high probability also while the taxpayers are honest, is also not ideal. Taking these facts into account, the far Nash equilibrium notion could be of interest in designing the policy in simple tax auditing problems (and also similar applications). The existence of a far Nash equilibrium could imply the effectiveness of a policy since we want to lower the probability of playing the same strategies (based on the numbers that we assign to the strategies). Furthermore, the statistical distance between the strategies of the players in an equilibrium could imply how the players behave in an equilibrium[5].

5.1.5 Strategic Resource Allocation

Problems of resource allocation and fair division have received considerable research attention from a computational perspective, as they are a critical consideration in multi-agent systems. Fair division problems involve a group of agents where each has individual preferences for some collection of resources, typically represented by a utility function. The goal is the allocation of resources in a fair manner. A fundamental underlying question is how to define fairness criteria in the first place. Classical fairness definitions include competitive equilibrium [7] and social Nash welfare [116]. Many fairness notions for indivisible goods have been considered; however, they are typically hard to compute [38, 39, 40]. Furthermore, almost all of this research is limited to non-strategic settings.

[5]For games that have multiple mixed Nash equilibria, we could possibly measure the mean of the statistical distance of the mixed strategies of the players.

Strategic fair division introduces participants who may act uncooperatively in order to maximize their own utility. In particular, players may hide their true preferences. In the presence of strategic behavior, the formulation of appropriate fairness criteria is essential. One branch of fair division related to game theory considers equilibrium in games resulting from fair division algorithms [37, 40, 177]. An *envy-free* allocation is an allocation in which each player receives a share that is, in their eyes, at least as good as the share received by any other agent [86]. There are several different variations of envy-freeness [15, 16, 183]. We define a variation of envy-freeness in the following strategic allocation problem:

Definition 5.1.11. *The game $\mathcal{RV}$ for k players and n items is defined as follows: $\forall i \in [k]$ and $\forall t \in [n]$, $U_i(t, x_{-i}) = \alpha V_i(t) + \beta R_i(x_{-i})$, where α and β are constants. $V_i(t)$ indicates the value of item t for player i and $R_i(x_{-i})$ is a function that is computable in constant time which gives player i's fairness measure for the allocation of items $x_{-i} = (t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n)$ to the other players*[6].

Definition 5.1.12. *An allocation is an* (δ, ϵ)-item-wise envy-free *allocation if each player receives the items that are in the support of a ϵ-partition Nash equilibrium x^* of game $\mathcal{RV}$ with the following properties:*

- *Each player $i \in [k]$ has a support of size $\frac{n}{k}$.*
- *For each player $i \in [k]$, we have $||x_i^* - b||_1 \leq \delta$ where b is a given distribution that $\frac{n}{k}$ strategies are assigned with probability $\frac{k}{n}$.*

It is easy to observe that in a $(0, 0)$-*item-wise envy-free* allocation, no player would prefer to exchange any item allocated to another player for any of the items in their own allocation. Note that in an envy-free allocation, each player is only concerned with the value of the items. The following proposition follows from the hardness of partition Nash.

Proposition 5.1.13. *The decision problem of whether there exists* (δ, ϵ)-item-wise envy-free *allocation is NP-complete for some δ and ϵ poly-bounded in the size of the game.*

Remark 5.1.14. Another branch of fair division considers truthful resource allocation mechanisms, in which agents are incentivized to reveal their true valuations for resources. We do not address truthful mechanisms in this research.

[6] A more complex form of this problem is described in [167] and Chapter 6.

5.2 Main Results

A number of the problems we consider depend on parameters besides the additive approximation term. These parameters depend on n, which is the size of the strategy space S. We say that such a parameter α is *poly-bounded* if there is some $k > 0$ such that $\alpha(n) = O(\frac{1}{n^k})$. The theorems of this section admit a number of hardness results for other variations of Nash equilibrium, including additional information about the parameters. We have not stated them here in the interest of simplifying the presentation.

5.2.1 Main Theorems

For Nash equilibrium with constraints, we have the following results:

Theorem 5.2.1. *For bi-matrix games, the following decision problems are NP-complete under polynomial approximation:*

(a) Disjoint Nash

(b) Partition Nash

(c) δ-Far Nash equilibrium, for any poly-bounded δ

(d) δ-Close Nash equilibrium, for any poly-bounded δ

(e) More than one Nash equilibrium with a statistical difference of δ for some δ polynomial in $|S|$.

We may consider the question of whether there is a threshold approximation error below which δ-far Nash becomes a total problem. Despite the absence of a definitive answer to this question, we provide a bound within which a δ-far Nash equilibrium is guaranteed to exist in the following theorem:

Theorem 5.2.2. *For any game $\mathcal{G}$, there exists a $\delta(n) \geq \frac{1}{n}$ such that at least one $(\delta, 4\delta)$-far Nash equilibrium for $\mathcal{G}$ exists. Furthermore, this problem is in PPAD.*

The main complexity results related to generalized equilibrium versions of these problems are the following. These are with respect to bi-matrix games and polynomial approximation except (a).

Theorem 5.2.3. *(a) The problem of finding a generalized disjoint equilibrium under exponential approximation is computable in polynomial time.*

(b) There is a poly-bounded δ for which the problem of finding a δ-far generalized equilibrium is PPAD-hard.

(c) There is a poly-bounded δ for which the problem of finding a δ-far generalized equilibrium for games whose diagonally modified version does not admit a fully mixed Nash equilibrium (see Definition 5.2.9) is in PPAD.

(d) For some poly-bounded θ, the problem of deciding whether there exists a θ-restricted disjoint equilibrium is NP-complete.

(e) For some poly-bounded θ and any poly-bounded δ, the problem of deciding whether there exists a (θ, δ)-restricted far equilibrium is NP-hard.

Remark 5.2.4. The previous assumption (Theorem 5.2.3 (c)) can be easily checked by using part of a well-known *support enumeration* algorithm.

5.2.2 Technical Overview

Basic Techniques

Throughout this work, we use two standard techniques (technical lemmas) in multiple situations. For example, to derive hardness results for disjoint Nash through standard NP-hard reductions from well-known NP-hard problems, we need to ensure that the required game solutions have disjoint support. This is done by duplicating some of the strategies so that the players can only play strategies from their *associated* sets. We explain this technique with the following simple example:

Lemma 5.2.5 (Simple Duplication)**.** *The problem of finding a Nash equilibrium is poly-time reducible to finding a disjoint Nash equilibrium.*

Proof. Given a game $\mathcal{G}$ with S strategies, we create game $\mathcal{G}'$ with duplicates of S (S_1 and S_2). We prove that a Nash equilibrium in $\mathcal{G}$ will form a disjoint Nash equilibrium in $\mathcal{G}'$ in which the first and second player only select their strategies from S_1 and S_2 respectively.

In $\mathcal{G}'$, all players have the strategy space $S' = S_1 \cup S_2$ where $S_1 = [n]$ and $S_2 = \{n+1, \cdots, 2n\}$ are just copies of S. Let σ be a value that is strictly less than the minimum payoff for either player in $\mathcal{G}$. The new payoff matrix will force each player k to only play from set S_k. For any (i,j) from $S' \times S'$, the utility matrix for each player is defined as follows:

$$U_1'(i,j) = \begin{cases} U_1(i,j-n), & \text{if } i \in S_1 \ \& \ j \in S_2 \\ \sigma & \text{o.w} \end{cases}$$

$$U_2'(i,j) = \begin{cases} U_2(i,j-n), & \text{if } i \in S_1 \ \& \ j \in S_2 \\ \sigma & \text{o.w} \end{cases}$$

Suppose that (x^*, y^*) is a Nash equilibrium for $\mathcal{G}$. If each player k plays the same distribution on members of S_k that was played on S in (x^*, y^*), we have a disjoint Nash equilibrium for $\mathcal{G}'$ by the construction and strict domination of σ. Now suppose (x', y') is a Nash equilibrium with disjoint supports in $\mathcal{G}'$. We easily can map these distributions to a Nash equilibrium (x^*, y^*) of $\mathcal{G}$ by defining, for $i \in S$, $x_i^* = x_i'$ and $y_i^* = y_{i+n}'$.

□

Remark 5.2.6. The preceding technique can only be applied to the exact versions of our problems. We show that, under different approximation schemes that we reviewed (in Chapter 2), this technique can connect present hardness results on Nash equilibrium to our strategic constraints if applied carefully. Theorem 5.2.2 is also a fairly direct corollary of this lemma. It is easy to see that we can scale the payoffs to be in the standard range $[0, 1]$.

Lemma 5.2.7 (Penalty Lemma)**.** *Suppose that (x^*, y^*) is a Nash equilibrium in a bi-matrix game whose utility matrix entries are in the range $[\alpha, \beta]$. For each i, we define x_i^- to be the same as x_i^* except we take $\epsilon_i \leq x_i^*$ from the weight of strategy i and redistribute it to some strategies in $[n] - \{i\}$. Assuming $\sum_{i=1}^n \epsilon_i \leq \epsilon$, we have the following*[7]*:*

1. *(x^-, y^*) and (x^*, y^-) are $2\epsilon(\beta - \alpha)$-Nash equilibria.*
2. *$|U_1(x^*, y^*) - U_1(x^-, y^-)| \leq 2\epsilon(\beta - \alpha)$*
3. *$|U_2(x^*, y^*) - U_2(x^-, y^-)| \leq 2\epsilon(\beta - \alpha)$*

[7] We assume $\alpha = 0$, $\beta = 1$ unless stated.

Proof. To prove this lemma, we show that modifying the strategy profile of the players in a Nash equilibrium to another strategy profile can guarantee an approximate Nash equilibrium in which the approximation error is bounded by the statistical distance of the modified strategy and the original Nash equilibrium.

It is not hard to see $||x^- - x^*||_1 \leq 2\epsilon$. We prove that we have $|U_1(x^*, y^*) - U_1(x^-, y^*)| \leq \epsilon(\beta - \alpha)$. Taking and redistributing (regardless of how we do so) ϵ_i from x_i^* can lower a payoff at most $\epsilon_i(\beta - \alpha)$. So, the overall utility loss will be less than $\epsilon(\beta - \alpha)$.

First, we show that x^- is a $\epsilon(\beta - \alpha)$-best response to y^* since x^* is a best response to y^* which is proven by the following argument:

$$\forall x \in \Delta_n, \quad U_1(x^*, y^*) \geq U_1(x, y^*)$$

$$\forall x \in \Delta_n, \quad U_1(x^-, y^*) + \epsilon(\beta - \alpha) \geq U_1(x, y^*)$$

We now show that y^* is a $2\epsilon(\beta - \alpha)$-best response to x^- . For all $y \in S$ we have:

$$U_2(x^*, y^*) \geq U_2(x^*, y)$$

$$U_2(x^-, y^*) + \epsilon(\beta - \alpha) \geq U_2(x^*, y) \geq U_2(x^-, y) - \epsilon(\beta - \alpha) \text{ (for any pure strategy } y)$$

$$U_2(x^-, y^*) + 2\epsilon(\beta - \alpha) \geq U_2(x^-, y)$$

This means (x^-, y^*) is a $2\epsilon(\beta - \alpha)$-Nash equilibrium. We can prove this for (x^*, y^-) similarly.

Inequality (2) is established as follows:

$$|U_1(x^*, y^*) - U_1(x^-, y^-)| \leq |U_1(x^*, y^*) - U_1(x^*, y^-)| + |U_1(x^*, y^-) - U_1(x^-, y^-)|$$

$$\leq \epsilon(\beta - \alpha) + \epsilon(\beta - \alpha) = 2\epsilon(\beta - \alpha)$$

(3) is established similarly. □

Remark 5.2.8. Let $\beta = 1$ and $\alpha = 0$. Here is an example that we use throughout the research. We take δ fraction of a pure strategy spread it equally to all other pure strategies.

$$\left|U_1(x^*,y^*) - U_1(x^-,y^*)\right| = \left|\sum_{1\le j\le n} \delta y_j^* U_1(t,j) - \sum_{1\le i\ne t\le n}\sum_{1\le j\le n} \frac{\delta}{(n-1)} y_j^* U_1(i,j)\right| \quad (5.1)$$

$$= \left|\sum_{1\le i\ne t\le n}\sum_{1\le j\le n} \frac{\delta}{(n-1)} y_j^* U_1(t,j) - \sum_{1\le i\ne t\le n}\sum_{1\le j\le n} \frac{\delta}{(n-1)} y_j^* U_1(i,j)\right|$$

$$= \left|\frac{\delta}{(n-1)}\left(\sum_{1\le i\ne t\le n}\sum_{1\le j\le n} y_j^* U_1(t,j) - \sum_{1\le i\ne t\le n}\sum_{1\le j\le n} y_j^* U_1(i,j)\right)\right|$$

It is obvious that $|U_1(t,j) - U_1(i,j)| \le 1$ since we can assume games' payoff is restricted to $[0,1]$.

$$|U_1(x^*,y^*) - U_1(x^-,y^*)| \le \left|\frac{\delta}{(n-1)}\sum_{1\le i\ne t\le n}\sum_{1\le j\le n} y_j^*\right|$$

Finally, since y^* is a distribution:

$$\left|U_1(x^*,y^*) - U_1(x^-,y^*)\right| \le \delta$$

We can do the same for the column player too:

$$\left|U_2(x^*,y^*) - U_2(x^*,y^-)\right| \le \delta$$

We calculated the exact difference between the utilities, but the reason that this forms a 2δ-approximate Nash equilibrium follows from Lemma 5.2.7.

To inspect the computational complexity of generalized far equilibrium, we use a technique that penalizes the diagonals of the payoff matrices. Using this approach, we demonstrate that Nash equilibrium with certain properties can be related to generalized equilibrium with farness-related constraints (see Proposition 5.4.9). The following definition identifies a class of games that we will show are guaranteed to have a generalized far equilibrium, under some conditions.

Definition 5.2.9. *For a game $\mathcal{G}$, the* diagonally modified *version of $\mathcal{G}$, denoted $\mathcal{D}^M(\mathcal{G})$, is a game where the entries of the payoff matrices $U_i^{\mathcal{D}^M(\mathcal{G})}$ are the same as $\mathcal{G}$*

except the diagonal entries as follows:

$$\forall t \in S : U_i^{\mathcal{D}^M(\mathcal{G})}(t,t) = -M$$

Remark 5.2.10. Note that if we find a Nash equilibrium in $\mathcal{D}^{\mathrm{M}}(\mathcal{G})$, it does not necessarily form a Nash equilibrium for the original game $\mathcal{G}$ since we removed all the information that the diagonals of the payoff matrices carry. We might assume that as M becomes sufficiently large, the likelihood of disjoint equilibria increases. While this observation may be useful in most cases, it does not necessarily hold true in general.

Deriving NP-hardness

The NP-hardness proofs of the problems stated in this chapter are inspired by the generic proof provided by [52] for NP-hardness of deciding the existence of a Nash equilibrium with constraints. This generic reduction (from SAT) improved the NP-completeness results for the exact Nash equilibrium with certain constraints given in [89]. However, the hardness of approximation proof given in [52] does not apply to the form of approximation that we consider in this work (the additive approximation error introduced in [59]). Instead, we start with a reduction given by Schoenebeck and Vadhan [169] (which modifies the proof provided in [52]) showing that the approximate Nash equilibrium with a guaranteed payoff for all players is NP-complete. Most NP-hardness reductions for constrained Nash equilibrium problems prove that the existence of a solution for the given SAT instance implies the existence of a desirable equilibrium solution of the constructed game, and the absence of an acceptable solution for the first problem implies that *any* existing Nash equilibrium does not satisfy the given constraints. In the absence of a solution for the SAT problem, existing reductions usually force the players to choose a single strategy while our reduction proofs depend on the use of more than one strategy.

Our proofs begin with the hardness of disjoint Nash, obtained by applying a non-trivial combination of the reductions in [169] and [52]. To clarify the structure of the proof, we break down the proof of [169] into a series of lemmas – see Appendix B.4. We use modest variations of these lemmas to prove the hardness of disjoint Nash by combining the technical lemmas introduced above and the ideas given in both reductions given in [52] (Corollary 6) and [169] (Theorem 8.6) – See Section 5.3. The hardness of partition Nash and the existence of an approximate far Nash in certain cases will be directly derived by using the technical lemmas – see Theorem

5.2.2 and Proposition 5.3.13. Next, by an adaptation of the duplication technique in the constructed game, force any equilibrium corresponding to a SAT solution to have the farness property along with other properties making it suitable to establish the NP-hardness of far Nash equilibrium – see Section 5.3.4. Proving NP-hardness results for the restricted disjoint equilibrium requires a further extension of the games that we construct based on the SAT problem in which we also embed a sub-game that is a generalization of *rock-paper-scissors*. Also, we will require a different analysis of the players' social welfare to establish NP-hardness – see Section 5.4.6. For the restricted far equilibrium, we will need to consistently integrate all the techniques and construction elements previously employed – see Section 5.4.7.

5.3 Proof of Theorem 5.2.1 and 5.2.2 (Nash Equilibrium with Strategic Constraints)

We now proceed to formally prove Theorem 5.2.1. It is not hard to prove that all of the stated problems (in Theorem 5.2.1) are in NP, since given a strategy profile, (unconstrained) Nash equilibrium conditions and our constraints can be verified in polynomial time.

Our proof for the NP-hardness results stated in Theorem 5.2.1 begins with the easiest case proving that approximate disjoint Nash (Theorem 5.2.1 part (a)) is NP-hard (see Proposition 5.3.2). Specifically, the hardness of approximate partition Nash equilibrium (Theorem 5.2.1 part (b)) can be derived by a modest modification of Proposition 5.3.2 (see Proposition 5.3.13). To prove Theorem 5.2.1 part (c), we show that the hardness can be extended to any δ (for δ-far Nash) polynomial in the size of the strategies in the game (see Section 5.3.4). The proof requires a more complicated form of the technique we introduced in Lemma 5.2.5. For simplicity, we break the proof down into two sub-cases where the statistical distance δ can be either a small or a big number (see 5.3.17 and 5.3.18).

The hardness of close Nash (part (d)) is also derived by combining the original game introduced by [169] (Appendix B.4) and Proposition 5.3.2 (see Proposition 5.3.20). The hardness of the problem of whether a game has more than one Nash equilibrium with a fixed statistical distance (part (e)) is proved in Proposition 5.3.15. Also, the proof of the hardness of the approximating Nash equilibrium with no probabilities is relegated to Section 5.4 as it is completely related to the hardness of restricted disjoint equilibrium (see Corollary 5.4.14). Proof of Theorem 5.2.2 (existence of far Nash) is relegated to Section 5.3.6.

5.3.1 Hardness of Disjoint Nash (Theorem 5.2.1 Part (a))

Before proving the NP-hardness of disjoint Nash, we recall the satisfiability problem for 3CNF formulas and some associated terminology.

Definition 5.3.1. *A* Boolean formula ϕ in CNF *(conjunctive normal form) is specified by a set V of* variables *(with $|V| = n$), a set of L of* literals *consisting of variables and their negations, and a set C of clauses, where each clause is a set of literals. A* 3CNF formula *is a CNF formula in which each clause has exactly 3 literals. 3CNFSAT is the problem of deciding whether there is a satisfying assignment for a 3CNF formula ϕ (i.e. a setting of the variables to* true *or* false *under which ϕ evaluate to* true.*)*

To show the hardness of approximate disjoint Nash, we give a poly-time mapping of any 3CNF ϕ to a non-symmetric bi-matrix game $\mathcal{G}(\phi, \epsilon)$ and prove Proposition 5.3.2. The game $\mathcal{G}(\phi, \epsilon)$ is a modification of the game given in [169], which we refer to as $\mathcal{SV}(\phi, \epsilon)$[8]. We use the duplication technique (see Lemma 5.2.5) in which we make two copies of some of the strategies where each player will have its own *associated* strategy while the other copy (which we call *unassociated* strategy) will be a strictly dominated strategy.

Let $S \equiv S_1 = S_2 = L_1 \cup L_2 \cup V \cup C \cup \{f\}$ be the strategy set for both players, where L_1 and L_2 are copies of L. Let $\mathrm{v} : L \to V$ be the function that gives the variable corresponding to a literal, e.g. $\mathrm{v}\,(x_1) = \mathrm{v}\,(-x_1) = x_1$. For example, if x_1 is a variable, x_1 and $-x_1$ are literals that are **representatives** of variables x_1 being true or false respectively. Also, let $\mathrm{g} : L_1 \cup L_2 \to L$ be the function (which we call the *unlabelling function*) that maps copies L_1 and L_2 to L. The utility matrices are defined as follows:

1. $u_1\,(l^1, l^2) = u_2\,(l^1, l^2) = n - 1$ for all $l^1 \in L_1,\ l^2 \in L_2$ with $\mathrm{g}(l^1) \neq -\mathrm{g}(l^2)$;
2. $u_1(l^1, -l^2) = u_2(-l^1, l^2) = n - 4$ for all $l^1 \in L^1,\ ,l^2 \in L^2$ & $\mathrm{g}(l^1) = \mathrm{g}(l^2)$;
3. $u_1(l^1, x) = u_2(x, l^2) = n - 4$ for all $l^1 \in L_1,\ l^2 \in L_2,\ x \in V \cup C$;

[8]The proof of hardness of deciding whether $\mathcal{SV}(\phi, \epsilon)$ has a ϵ-approximate Nash with a guaranteed payoff $n - 1 - \epsilon$ is given in Appendix B.4. We also do require this proof for other hardness results as well.

4. $u_1(v, l^2) = u_2(l^1, v) = n$ for all $v \in V,\ l^1 \in L_1,\ l^2 \in L_2$ & $\mathrm{g}(l^1) = \mathrm{g}(l^2)$ with $\mathrm{v}(\mathrm{g}(l^1)) \neq v$;

5. $u_1(v, l^2) = u_2(l^1, v) = 0$ for all $v \in V,\ l^1 \in L_1,\ l^2 \in L_2$ & $\mathrm{g}(l^1) = \mathrm{g}(l^2)$ with $\mathrm{v}(\mathrm{g}(l^1)) = v$;

6. $u_1(v, x) = u_2(x, v) = n - 4$ for all $v \in V,\ x \in V \cup C$;

7. $u_1(c, l^2) = u_2(l^1, c) = n$ for all $c \in C,\ l^1 \in L_1,\ l^2 \in L_2$ with $\mathrm{g}(l^1) = \mathrm{g}(l^2) \notin c$;

8. $u_1(c, l^2) = u_2(l^1, c) = 0$ for all $c \in C,\ l^1 \in L_1,\ l^2 \in L_2$ with $\mathrm{g}(l^1) = \mathrm{g}(l^2) \in c$;

9. $u_1(c, x) = u_2(x, c) = n - 4$ for all $c \in C,\ x \in V \cup C$;

10. $u_1(x, f) = u_2(f, x) = 0$ for all $x \in S - \{f\}$;

11. $u_1(f, f) = u_2(f, f) = 2n$;

12. $u_1(f, x) = u_2(x, f) = n - 1$ for all $x \in S - \{f\}$.

We need to define the following cases in such a way that players do not try to play from their associated literal set as they become strictly dominated strategies[9]:

13. $u_1(l^2, *) = u_2(*, l^1) = -2n$ for all $* \in S,\ l^2 \in L_2$ and $l^1 \in L_1$;

14. $u_2(l^2, *) = u_1(*, l^1) = 0$ for all $* \in S - \{f\},\ l^2 \in L_2$ and $l^1 \in L_1$.

Proposition 5.3.2. *Given an instance of 3CNF ϕ with n variables, there is a ϵ-disjoint Nash equilibrium (where $\epsilon = \frac{1}{2n^3}$) in $\mathcal{G}(\phi, \epsilon)$ iff ϕ is satisfiable.*

Before proving Proposition 5.3.2, we need several lemmas that are adaptations of the lemmas used to prove the hardness of guaranteed Nash (for the game $\mathcal{SV}(\phi, \epsilon)$).

Essential Lemmas for Proposition 5.3.2

Lemma 5.3.3. *In any ϵ-approximate Nash equilibrium of $\mathcal{G}(\phi, \epsilon)$, the row player cannot play literals from L_2 with probability greater than $\frac{\epsilon}{2n}$. A symmetric argument also applies to the column player and all other strategies that have a payoff of $-2n$.*

[9]Note that in [52], $u_i(f, f)$ is an arbitrary positive number.

Proof. If one player plays a literal that is not from its associated literal set, they will get a payoff of $-2n$. Thus, in any ϵ-approximate Nash equilibrium, these strictly dominated strategies can be played with a probability of at most $\frac{\epsilon}{2n}$, as playing such a strategy with probability greater than $\frac{\epsilon}{2n}$ will cause the player to lose at least ϵ compared to other strategies and this violates the definition of ϵ-best response. □

Lemma 5.3.4. *In any ϵ-approximate Nash equilibrium of $\mathcal{G}(\phi, \epsilon)$ if neither player plays f with a positive probability, literals that are not associated with the players, as well as clauses and variables, can be played with a probability of at most ϵ.*

Proof. The proof is similar to Lemma B.4.3. □

Indeed, strictly dominated strategies, variables, and clauses cannot be played with a combined probability greater than ϵ. The next lemma will show the fact that if one of the players cannot meet the guaranteed payoff $n - 1 - \epsilon$, an ϵ-disjoint Nash equilibrium cannot be obtained.

Lemma 5.3.5. *In any ϵ-approximate Nash equilibrium in $\mathcal{G}(\phi, \epsilon)$, all players have a guaranteed payoff of at least $n - 1 - \epsilon$.*

Proof. Suppose that one of the players does not have a guaranteed payoff of at least $n - 1 - \epsilon$. This player can deviate from the assumed strategy and play f to get $n - 1$ no matter what the other player does. Consequently, this strategy is not a ϵ-best response to the opponent player. □

Lemma 5.3.6. *Suppose that the formula ϕ is satisfiable. The game $\mathcal{G}(\phi, \epsilon)$, has at least one disjoint Nash equilibrium.*

Proof. An argument that is similar to Theorem B.4.2 can be applied to $\mathcal{G}(\phi, \epsilon)$ to show that there exists a disjoint Nash equilibrium. The only difference here is the duplicated literal sets and the unlabelling function. In other words, the strategy where each player randomizes uniformly over n out of $2n$ (based on a satisfying assignment) of their associated literals in L_1 and L_2 respectively, is a disjoint Nash equilibrium (where the expected payoff to each player is $n - 1$). □

Lemma 5.3.7. *Suppose that ϕ is not satisfiable. If neither player plays f with a positive probability, then for any $l^1 \in L_1$, the probability that the row player plays $l^1 \in L_1$ or $-l^1 \in L_1$ is at least $\frac{1}{n} - 2\epsilon$.*

Proof. Suppose this is not the case. Assume that the probability is smaller than $\frac{1}{n} - \epsilon - \frac{2\epsilon}{n} \geq \frac{1}{n} - 2\epsilon$. Let $l^2 \in L_2$ be a literal such that $\mathrm{g}(l^2) = \mathrm{g}(l^1)$. Recall that in Lemma 5.3.4, both players play from their associated literal sets with a probability of at least $1 - \epsilon$ while all other strategies can be played with a probability of at most ϵ. Then, the expected payoff for the column player, when playing $\mathrm{v}(\mathrm{g}(l^2))$ will be at least $n - 1 + 2\epsilon$ by Lemma B.4.4, where the only modification is that we use Lemma 5.3.4 instead of Lemma B.4.3. Since the maximum social welfare is $2n - 2$, the other (row) player fails to meet the guarantee $n - 1 - \epsilon$ and will deviate to f. □

Remark 5.3.8. A similar argument can be applied to the column player as well.

Lemma 5.3.9. *Suppose that ϕ is not satisfiable. If neither player plays f with a positive probability, then for each player and literal l from their associated set, either l or $-l$ is played with probability $\geq \frac{1}{n} - 2\epsilon - \frac{1}{n^2}$ while the other is played with probability less than $\frac{1}{n^2}$.*

Proof. The proof is exactly the same as Lemma B.4.5 except each player will play from their respective literal set with a high probability. □

Lemma 5.3.10. *Suppose that ϕ is not satisfiable. If neither player plays f with a positive probability, at least one player cannot guarantee a payoff of $n - 1 - \epsilon$.*

Proof. We prove that in any ϵ-approximate Nash equilibrium in the constructed game except for the ones that have f in their support, there still exists a one-to-one correspondence between literals and truth assignments. We know that for each player, the player will play one of l or $-l$ (from their associated literals) with a probability of at least $1 - \frac{1}{n^2} - 2\epsilon$ while the other is played with a probability of at most $\frac{1}{n^2}$. We can consider the variables to be true if l^1 from L_1 and l^2 from L_2 are played with a higher probability compared to $-l^1$ and $-l^2$ for each player respectively (note that $\mathrm{g}(l^1) = \mathrm{g}(l^2) = l$). If an assignment does not satisfy the formula, by changing the strategy to a clause that is not satisfiable, one of the players will receive at least $\left(1 - \epsilon - \frac{3}{n^2}\right) n > n - 1 + 2\epsilon$ by Proposition B.4.2. So, the other player cannot attain $n - 1 - \epsilon$.

□

Proof of Proposition 5.3.2

After proving the essential lemmas, we finally proceed to the proof of Proposition 5.3.2.

Proof. We showed that if ϕ is satisfiable, a disjoint Nash equilibrium (also an approximate solution) exists. We also showed that if ϕ is not satisfiable, if neither player plays f with a positive probability, at least one player cannot guarantee a payoff of $n-1-\epsilon$. In conclusion, by Lemma 5.3.5, one player will have to play f, and we show that this forces the other player to use f with a positive probability.

Now, assume that one of the players plays f with probability α while the other player does not play f[10]. There are only two possible cases for α where the first case is that $\alpha < \frac{\epsilon}{n}$ and the second case is $\alpha \geq \frac{\epsilon}{n}$. We show that for either of the cases, an ϵ-disjoint Nash equilibrium cannot be generated in $\mathcal{G}(\phi, \epsilon)$ in the following lemma. □

Lemma 5.3.11. *Let $\epsilon = \frac{1}{2n^3}$ where n denotes the number of the variables in ϕ. If ϕ is unsatisfiable, $\mathcal{G}(\phi, \epsilon)$ has no ϵ-approximate Nash equilibrium in which exactly one of the players plays f with a positive probability.*

Proof. Without loss of generality, suppose that the row player plays f with probability α. It is easy to observe that no ϵ-approximate Nash equilibrium with a payoff less than $n-1-\epsilon$ exists for both players, similar to Lemma 5.3.5. Indeed, the column player can actually guarantee a better payoff of $[\alpha \cdot (2n) + (1-\alpha)(n-1))] - \epsilon$ by deviating to f. As argued above, if one of the players does not play f, the maximum possible social welfare is at most $2n-2$. So, if one of the players can achieve a payoff greater than $n-1+2\epsilon$, the other player will get less than $n-1-\epsilon$. If one of the players plays f with a positive probability, variables, clauses, and literals that are not associated with players, can be played with a probability of at most ϵ (we will use this bound when determining payoffs later). This follows as in Lemma 5.3.4 since the only modification in this setting is that the column player has a better guaranteed expected payoff.

Now assume that (x^*, y^*) is an approximate disjoint Nash equilibrium where the row player plays f with probability $0 < \alpha < \frac{\epsilon}{n}$ and the column player plays f with probability 0. We show (similar to Lemma B.4.4) that the probability that the row player plays l^1 or $-l^1$ is at least $\frac{1}{n} - 2\epsilon$: Suppose that the probability is less than $\frac{1}{n} - 2\epsilon$. The expected payoff for the column player playing $\mathrm{v}(\mathrm{g}(l^2))$ (we know $\mathrm{g}(l^1) = \mathrm{g}(l^2)$) with probability 1 is at least $n-1+2\epsilon$, computed as follows:

- When the row player plays l^1 or $-l^1$ the payoff is zero.

[10] We can only assume one of the players can play f due to the disjointness constraint.

- When the row player plays an associated literal in L_1 other than l^1 or $-l^1$, the payoff is

$$(1-\epsilon-(\frac{1}{n}-2\epsilon)-\alpha)n=(1+\epsilon-\frac{1}{n}-\alpha)n\geq n+\epsilon n-1-\epsilon=n-1+\epsilon(n-1)$$

- When the row player does not play a literal the payoff is at least $\epsilon\cdot 0$.

- When the row player plays f, the portion of the payoff is $\alpha\cdot 0$.

The summation will be at least $n-1+2\epsilon$ (for $n>2$) which indicates the row player cannot meet the guarantee $n-1-\epsilon$. In conclusion, the row player should now play f with a greater probability, namely 1, not α. For any $l^2\in L_2$, the probability that the column player plays either l^2 or $-l^2$ is at least $\frac{1}{n}-2\epsilon$ by Lemma 5.3.7. This is because the column player does not play f. So, for each player, the probability of playing either l or $-l$ from their associated literal set is at least $\frac{1}{n}-2\epsilon$. This implies that one of $l^1\in L_1$ or $-l^1$ (for any $l^1\in L_1$) is played with probability greater than $\frac{1}{n}-2\epsilon-\frac{1}{n^2}$ while the other is played with probability smaller than $\frac{1}{n^2}$ (see Lemma B.4.5 and 5.3.9). The preceding holds for the column player and the associated literal as well. Using an argument similar to those used in Lemma 5.3.10 and Proposition B.4.2, we conclude that the row player can change its strategy to play an unsatisfied clause (with a high probability $1-\alpha$) and receive at least $(1-\epsilon-\frac{3}{n^2}-\alpha)n>n-1+2\epsilon$. Since the maximum social welfare is $2n-2$, it is better for the column player to play f and obtain a payoff of $n-1$ instead.

Now assume the row player plays f with probability $\alpha>\frac{\epsilon}{n}$ while the column player plays f with probability zero. Recall that the column player can only play strategies other than its associated literals with a probability of at most ϵ. If the column player deviates to f with probability 1, the player will increase their payoff by at least ϵ which violates the definition of ϵ-best response:

$$[2\alpha\cdot n+(1-\alpha)(n-1)]-[(1-\epsilon-\alpha)(n-1)+\epsilon\cdot n+\alpha\cdot 0]=2\alpha\cdot n-\epsilon>\epsilon$$

□

Remark 5.3.12. The minimum and the maximum payoff of the game $\mathcal{G}(\phi,\epsilon)$ are $-2n$ and $2n$ respectively. We can scale the payoffs of this game to be in the standard range $[0,1]$.

5.3.2 Hardness of Partition Nash (Theorem 5.2.1 Part (b))

While proving EXACT partition Nash is NP-hard appears to be quite challenging, an NP-hardness result for 2-player approximate partition Nash equilibrium can easily be derived from the NP-hardness result that we have for disjoint Nash equilibrium and Lemma 5.2.7.

Proposition 5.3.13. *Given an instance of 3CNF ϕ with n variables, the game $\mathcal{G}(\phi, \epsilon)$ has a ϵ-partition Nash equilibrium (for $\epsilon = \frac{1}{2n^3}$) iff the formula is satisfiable.*

Proof. As stated, it is easy to check whether the two conditions that we need for a partition Nash are satisfied or not. So, this problem is in NP. Proposition 5.3.2 shows that even under polynomial approximation, this problem is NP-hard. We know that in $\mathcal{G}(\phi, \epsilon)$, by Proposition 5.3.2, any ϵ-approximate Nash equilibrium that does not have f in the support of both players, has the following property. For each player $i \in \{1, 2\}$ and any specific literal l_i such that $\mathrm{g}(l_2) = \mathrm{g}(l_1) = l$, one of $l_i \in L_i$ or $-l_i \in L_i$ is played with probability greater than $\frac{1}{n} - \frac{1}{n^2} - 2\epsilon$ while the other is played with probability smaller than $\frac{1}{n^2}$. All other possible strategies are played with a probability of at most ϵ since the literals have to be played with probability $1 - \epsilon$.

In $\mathcal{G}(\phi, \epsilon)$, any disjoint Nash equilibrium should have the property that for all $l_1 \in L_1$ and $l_2 \in L_2$ such that $g(l_1) = g(l_2) = l$, one of l_1 or $-l_1$ for the row player and one of l_2 or $-l_2$ should be played with probability $\frac{1}{n}$ while the other has to be played with probability zero directly followed by [52]. Note that if one of the players plays f, the other player has to play f too, which we are not concerned with in this proof. In this case, all players will play exclusively from their associated literals since all other literals are dominated by all other strategies.

If the formula ϕ is satisfiable, an (EXACT) disjoint Nash can be generated by 5.3.2. Assume that (x^*, y^*) is a disjoint Nash equilibrium of $\mathcal{G}(\phi, \epsilon)$. We modify this strategy to (x^-, y^-) and generate a ϵ-partition Nash equilibrium. Consider $\epsilon' = \frac{\epsilon}{cn}$ and recall that $\epsilon = \frac{1}{2n^3}$. The reason that we chose ϵ' to be very small is that we have some strictly dominated strategies (c is large enough to cover all these rules for a partition) and also f cannot have a probability greater than $\frac{\epsilon}{n}$ (due to the fact that both players will have to play f). We take a small fraction $\frac{\epsilon'}{2}$ of one pure strategy from the column player and distribute it equally to strategies that are not in the support of the row and the column player for the column player. It is easy to see that (x^-, y^-) is an ϵ-approximate partition Nash equilibrium by Lemma 5.2.7. Now, both properties of a partition Nash instance are satisfied. It is obvious that if we have an approximate partition Nash (x^-, y^-), we can simply generate a disjoint Nash equilibrium that can imply a satisfying assignment. This is because all literals l that are played with probability greater than $\frac{1}{n} - 2\epsilon - \frac{1}{n^2}$ can have probability $\frac{1}{n}$ can determine an assignment of the formula ϕ.

□

Remark 5.3.14. Proof of Proposition 5.1.13 (resource allocation application) is simply followed as a corollary.

5.3.3 Proof Theorem 5.2.1 Part (e)

The following proposition establishes the hardness of deciding whether a game has more than one equilibrium with a statistical distance at least $\delta = 2\epsilon$ between the equilibria of the game.

Proposition 5.3.15. *When the formula ϕ is unsatisfiable, both players will have to play f with probability greater than $1 - \epsilon$ in any equilibria of the game $\mathcal{G}(\phi, \epsilon)$.*

Proof. We can proceed as in the proof of Proposition 5.3.2. Assume that the row player plays f with probability α_1 and the column player plays f with α_2. It is easy to see that one of α_1 or α_2 must be greater than $\frac{\epsilon}{n}$. If both of these probabilities are smaller than $\frac{\epsilon}{n}$, one of the players cannot reach their guaranteed payoff, by a modification of Lemma 5.3.11 (both players play f with a small probability). In conclusion, we can assume that $min(\alpha_1, \alpha_2) \geq \frac{\epsilon}{n}$. It is easy to show that both players have a guaranteed payoff of at least $[min(\alpha_1, \alpha_2) \cdot (2n) + (1 - min(\alpha_1, \alpha_2)(n-1))] - \epsilon$ by just playing the pure strategy f. The players can only achieve a payoff greater than their approximate guaranteed payoff if they play f with probability greater than $1 - \epsilon$. If one player plays f with probability smaller than $1 - \epsilon$ (we know that must be greater than $\frac{\epsilon}{n}$), the player loses $2\min(\alpha_1, \alpha_2) \cdot n - \epsilon$ which is greater than ϵ. □

Remark 5.3.16. Note that when the formula is satisfiable, a disjoint Nash equilibrium exists in addition to the Nash equilibrium with the strategy profile (f, f).

5.3.4 Hardness of Far Nash (Theorem 5.2.1 Part (c))

In this section, we prove that for any δ and some ϵ poly-bounded in the size of the strategies in a game, the problem of deciding whether there exists a (δ, ϵ)-far Nash remains hard. The construction that we provide indeed even works for a stronger result is also presented in Corollary 5.3.19.

Given a formula ϕ with n variables, we generate the game $\mathcal{H}(\phi, \delta, \epsilon)$, which is similar to $\mathcal{G}(\phi, \epsilon)$ with the exception that not all literals are duplicated, depending on the given value δ (see Proposition 5.3.17 and 5.3.18). The strategy space for both players is $S^{H_\delta} \equiv S_1^{H_\delta} = S_2^{H_\delta} = L' \cup L'_1 \cup L'_2 \cup V \cup C \cup \{f\}$. We divide the original literals (L from the formula ϕ) into three sets. The set L' will contain the literals from L that are not duplicated, while L'_1 and L'_2 contain the duplicated literals, associated respectively with the row and the column player. The function $\mathrm{g}' : L' \cup L'_1 \cup L'_2 \rightarrow L$ can be defined similarly to g except that it maps elements of $L' \subset L$ to themselves. The rules for literals in L' are similar to rules in $\mathcal{SV}(\phi, \epsilon)$ and the ones are in $L'_1 \cup L'_2$ similar to those in $\mathcal{G}(\phi, \epsilon)$.

Proposition 5.3.17. *Given any 3CNF formula ϕ with n variables, for any $\delta > \frac{1}{n} - \frac{1}{n^2} - \frac{1}{n^3}$, there exists a game $\mathcal{H}(\phi, \delta, \epsilon)$ such that for some ϵ poly-bounded in n, there exists a (δ, ϵ)-far Nash equilibrium in $\mathcal{H}(\phi, \delta, \epsilon)$ iff ϕ is satisfiable.*

Proof. The proof follows by combining Proposition 5.3.2 and Proposition B.4.2. Recall that by Lemma 5.3.9, 5.3.11 and B.4.5, for any literal $l \in L'$ (and also the associated copies $l_1 \in L'_1$ and $l_2 \in L'_1$), either the literal or the negation of the literal $-l$ will be played with a probability of at least $\frac{1}{n} - 2\epsilon - \frac{1}{n^2}$ when both players do not play f with a positive probability greater than $\frac{\epsilon}{n}$.

To generate $\mathcal{H}(\phi, \delta, \epsilon)$ for the given δ, we determine i such that $\delta = \frac{1}{2}(2i(\frac{1}{n} - 2\epsilon - \frac{1}{n^2}))$ and select $\lfloor i \rfloor$ some arbitrary literals of the original formula and generate the sets L'_1 and L'_2. The rest of the literals will form L' (will not be duplicated). Note that for $\delta' \leq \delta$, NP-hardness of (δ', ϵ)-far Nash implies the hardness of (δ, ϵ)-far Nash. An argument similar to Proposition 5.3.15 implies that if the given formula is unsatisfiable, there will be no (δ, ϵ)-far Nash equilibrium (both players have to play f with a probability of at least $1 - \epsilon$ and we know that $\delta > \epsilon$). □

Proposition 5.3.18. *Given any 3CNF formula ϕ with n variables, for any δ poly-bounded in n, there exists a game $\mathcal{D}(\phi, \delta, \epsilon)$ such that for some ϵ poly-bounded in n, there exists a (δ, ϵ)-far Nash equilibrium in $\mathcal{D}(\phi, \delta, \epsilon)$ iff ϕ is satisfiable.*

Proof. For $\delta > \frac{1}{n} - \frac{1}{n^2} - \frac{1}{n^3}$, the hardness result follows by Proposition 5.3.17 with the game $\mathcal{H}(\phi, \delta, \epsilon)$. Otherwise, we create a new game $\mathcal{D}(\phi, \delta, \epsilon)$ similar to the game $\mathcal{H}(\phi, \delta, \epsilon)$ with the difference that we make more copies of one literal. For some literal

$l^t \in L$ (and also its negation $-l^t$) from the original formula ϕ, we make d copies and i out of these d copies will be duplicated in a way that each player will be inclined to play only from the associated copies. This can be attained by using a game structure similar to the one discussed in Proposition 5.3.2 and the technique used in Lemma 5.2.5).

We consider $L^{\mathcal{D}} = L \cup (\bigcup_{j=i}^{d} l_j^t) \cup (\bigcup_{j=i}^{d} (-l_j^t)) - \{l^t, -l^t\}$ to be all literal strategies that are shared between the players where for each j, l_j^t is just a copy of l^t. Similarly, $L_1^{\mathcal{D}} = (\bigcup_{j=1}^{i} l_{1,j}^t) \cup (\bigcup_{j=1}^{i} -l_{1,j}^t) - \{l_1^t, -l_1^t\}$ and $L_2^{\mathcal{D}} = (\bigcup_{j=1}^{i} l_{2,j}^t) \cup (\bigcup_{j=1}^{i} -l_{2,j}^t) - \{l_2^t, -l_2^t\}$ will be considered as the sets that contain the literals that are associated with the row and the column player respectively. Finally, we define the strategy set of the game $\mathcal{D}(\phi, \delta, \epsilon)$ to be $S^{\mathcal{D}} = L^{\mathcal{D}} \cup L_1^{\mathcal{D}} \cup L_2^{\mathcal{D}} \cup V \cup C \cup \{f\}$.

In this game, for the row player, the summation of the probabilities assigned to literals that are copies of l^t and $-l^t$ of the original formula, must be at least $\frac{1}{n} - 2\epsilon$ followed by Lemma 5.3.7, 5.3.11, and B.4.4 (if both players play f with a positive probability less than $\frac{\epsilon}{n}$). In addition, the summation of the probabilities that are assigned to literals that are copies of l^t or $-l^t$ must be at least $\frac{1}{n} - \frac{1}{n^2} - 2\epsilon$ while the other becomes less than $\frac{1}{n^2}$.

We set the payoff $u_1^{\mathcal{D}}(f, f) = u_2^{\mathcal{D}}(f, f)$ to be high enough so that all ϵ-approximate Nash equilibria that include the case that both players play f cannot form a δ-far Nash equilibrium (see Proposition 5.3.15). In conclusion, if the formula ϕ is unsatisfiable, there will be no (δ, ϵ)-far Nash equilibrium. If the formula ϕ is satisfiable, then we can find a $\frac{i}{d}(\frac{1}{n} - \frac{1}{n^2} - \frac{1}{n^3})$-far Nash equilibrium (by spreading the weight evenly among the copies of l^t) and set i to be small enough so that this fraction becomes greater than the given δ. □

This concludes the fact that the problem of whether there exists a (δ, ϵ)-far Nash equilibrium is NP-hard. Stronger results can also be implied by using similar techniques as follows:

Corollary 5.3.19. *The problem of whether there exists a Nash equilibrium (x^*, y^*) for a bi-matrix game $\mathcal{G}$ with the following property is NP-complete.*

- $2\alpha \leq ||x^* - y^*||_1 \leq 2\beta$ *for any* $1 \geq \beta > \alpha \geq 0$ *poly-bounded in the size of the strategies of* $\mathcal{G}$.

5.3.5 Hardness of Close Nash (Theorem 5.2.3 Part (d))

Proposition 5.3.20. *The problem of deciding whether an arbitrary bi-matrix game has a δ-close Nash equilibrium for any poly-bounded δ is NP-complete.*

Proof. Proof is followed by similar arguments used in Proposition 5.3.2 and Proposition 5.3.15 while instead of duplicating the literals in L, we duplicate the strategy f as follows:

1. $u_1(x, f_2) = u_2(f_1, x) = 0$ for all $x \in S - \{f_1, f_2\}$;

2. $u_1(f_1, f_2) = u_2(f_1, f_2) = 2n$;

3. $u_1(f_1, x) = u_2(x, f_2) = n - 1$ for all $x \in S - \{f_1, f_2\}$;

4. $u_1(f^2, *) = u_2(*, f^1) = -2n$ for all $* \in S$;

5. $u_2(f^2, *) = u_1(*, f^1) = 0$ for all $* \in S - \{f_1, f_2\}$.

All other rules will be the same as the game $\mathcal{SV}$ in Appendix B.4. By adding these five rules to the game $\mathcal{SV}$, a δ-close Nash equilibrium exists iff the given formula is satisfiable. Note that $0 \leq \delta < 2 - 2\epsilon$ and the hardness result can be improved to any δ poly-bounded in n by using the same argument used in Proposition 5.3.15 (setting higher payoffs).

□

5.3.6 Total Far Nash (Proof of Theorem 5.2.2)

We prove that a far Nash equilibrium exists for a sufficiently large approximation error. This follows by the existence of a Nash equilibrium and simple techniques. We begin with a lemma and then proceed to the proof.

An Essential Lemma

Lemma 5.3.21. *Let $\mathcal{G}$ be a game that has a symmetric Nash equilibrium (z^*, z^*) such that some pure strategy t has a probability greater than δ in z^*. Then, there is a strategy z^- such that (z^-, z^*) and (z^*, z^-) are 2δ-approximate Nash equilibria and $||z^* - z^-||_1 \geq 2\delta$.*

Proof. In this game, we have at least one pair of strategies (z^*, z^*) that forms a Nash equilibrium. For one player, we give δ share of z_t^* (arbitrary t) evenly to all other strategies to generate z^-. This modification on strategies has the property that z^- and z^* have 2δ-farness. Then the strategy profile (z^-, z^*) has the property $||z^- - z^*|| \geq 2\delta$. This also holds for (z^*, z^-) by a symmetric argument. Both (z^-, z^*) and (z^*, z^-) are 2δ-approximate Nash equilibrium by Lemma 5.2.7.

□

Proof of Theorem 5.2.2

Proof. Suppose (x^*, y^*) is a Nash equilibrium in $\mathcal{G}$. If $||x^* - y^*||_1 > 2\delta$, we are finished with the problem. Otherwise, $||x^* - y^*||_1 \leq 2\delta$, we can change x^* to y^* (or the other way around) so that there exists z^* such that (z^*, z^*) will form a $(0, 2\delta)$-far equilibrium by an argument similar to Lemma 5.2.7. For the next step, we simply can use Lemma 5.3.21 to generate z^- and $||z^- - z^*||_1 = 2\delta$.

By Lemma 5.2.7 and 5.3.21, we can conclude that (z^-, z^*) is a 4δ-approximate Nash equilibrium ($||x^* - z^-|| \leq 4\delta$) and we know $||z^- - z^*||_1 = 2\delta$ which indicates that we can generate $(\delta, 4\delta)$-far Nash equilibrium. It is obvious that there exists one strategy that is assigned with probability greater than $\frac{1}{n}$ in each distribution, and we can modify this strategy accordingly. Our proof reduces this problem to an instance of approximate Nash equilibrium, which means that this problem is in PPAD.

□

Remark 5.3.22. To get a better approximation factor compared to 4δ, we could look at the following optimization problem:

$$Min_{x \in \Delta_n} ||x - x^*||_1$$

$$Subject\ to\ ||x - y^*||_1 \geq 2\delta$$

5.4 Proof of Theorem 5.2.3 (Generalized Equilibria with Strategic Constraints)

We now proceed to the proof of Theorem 5.2.3, which addresses the computational aspects of the generalized equilibria version of the problems we study. In the presence of constraints such as disjointness and farness, the existence of a generalized equilibrium

Table 5.3: Rock-Paper-Scissors, a zero-sum game

	Player 1		
	(0,0)	(-1,1)	(1,-1)
Player 2	(1,-1)	(0,0)	(-1,1)
	(-1,1)	(1,-1)	(0,0)

cannot be guaranteed[11]. However, we show that under some simple assumptions, a trivial solution (part (a)) and a non-trivial solution (part (c)) can be guaranteed.

5.4.1 A Simple Example

A generalized disjoint equilibrium does not necessarily exist in all games. Consider the well-known *rock-paper-scissors* which is a zero-sum game (table 5.3). This game has a unique Nash equilibrium where both players uniformly select all three pure strategies and put them in their support. For each strategy, the game has one *anti-strategy* where one of the players gets 1 when the player plays the anti-strategy while the opponent gets -1. If both players play the same pure strategy, they both get 0 payoff. There are two possible cases for each player if we limit the players to respect the constraint set $\mathcal{R}^n_{\text{disjoint}}$. The first case is that both players pick one pure strategy, and the second case happens when one player plays one pure strategy while the other player plays two pure strategies. In either case, both players prefer to play the anti-strategies as this option is always available for at least one player. So, this game does not have a generalized disjoint equilibrium. But for example, if one player chooses 'rock' and the other player players 'paper' with probability $1-\epsilon$ and 'scissors' with probability ϵ. We can form an ϵ-generalized disjoint equilibrium.

5.4.2 Proof of Theorem 5.2.3 Part (a)

Proof. Choose a pure strategy l and let that be the column player's strategy. The row player will select a pure strategy t (other than l) that maximizes the row player's payoff. Let x^* be the row player's strategy. For all $i \in [n] - \{t, l\}$, put $x^-_i = \frac{\epsilon}{(n-2)}$ and $x^-_t = 1 - \epsilon$. By Lemma 5.2.7, the row player can lose at most ϵ. Therefore, x^- is a ϵ-best response under the disjointness constraint for the row player. In this situation,

[11]The following example shows that the exact version of generalized disjoint equilibrium does not necessarily admit a solution. Also, note that similar proofs could not be derived for the generalized δ-far equilibrium in general.

the column player cannot obtain an improved payoff, unless the column player plays strategies that are in the support of x^-, which violates the constraints. □

Remark 5.4.1. This result shows that the disjointness of strategies as a common (coupled) constraint is flawed as one selfish player can degrade the social welfare of all other players if all players comply with the common constraint. This indeed provides a justification for considering (θ, ϵ)-restricted disjoint equilibrium.

5.4.3 Hardness and Totality of Generalized Far Equilibrium and Theorem 5.2.3 Part (b) and (c)

Proof of Theorem 5.2.3 Part (b)

For $\delta = 0$, the proof can be trivially implied by the hardness of Nash equilibrium. We know that essentially, there exists a game $\mathcal{G}$ that is known to be hard for finding an equilibrium in [45] even for inverse polynomial approximation, as we restate it in the following:

Theorem 5.4.2 ([45])**.** *For any $c > 0$, the problem of finding an n^{-c}-approximate Nash equilibrium of a bi-matrix game is PPAD-complete.*

In the following proposition, we conclude that the hardness holds for the exact case when $\delta = 1$ (generalized disjoint equilibrium) while the approximate version is tractable (see 5.2.3 part (a)).

Proposition 5.4.3. *For any game $\mathcal{G}$, the problem of finding a Nash equilibrium in the bi-matrix game $\mathcal{G}$ is poly-time reducible to the problem of finding a generalized disjoint equilibrium in $\mathcal{D}^M(\mathcal{G})$.*

Proof. The proof follows from the fact that the problem of finding a Nash equilibrium is a PPAD-hard problem by [45]. A complete proof is given in Appendix A.4.3. □

Remark 5.4.4. The proof of the previous proposition also implies that this modification of the diagonal entries does not change the fact that finding a Nash equilibrium remains hard.

Proof of Theorem 5.2.3 Part (c)

In order to prove Theorem 5.2.3 part (c), we need to prove some properties of *diagonally modified games.* We begin with the following definition and establish then establish a relationship between this definition and generalized far equilibrium.

Definition 5.4.5. *Suppose that we are given a bi-matrix game and the goal is to find a Nash equilibrium with the following property:*

- *For all $t \in [n]$, if both players play t with a positive probability, then at least one of the players has to play t with probability smaller than $\frac{1}{M}$.*

*This equilibrium is called M-*semi-disjoint Nash *equilibrium. A* generalized M-semi-disjoint *equilibrium can be defined in a similar manner where we only investigate deviations with respect to strategies that have the property.*

Lemma 5.4.6. *Suppose that any in Nash equilibrium of $\mathcal{D}^M(\mathcal{G})$, at least one of the players does not play a fully mixed strategy. For any pure strategy t that is played with a positive probability by both two players, there exists at least one player that has to play strategy t with probability less than $\frac{1}{M}$. Furthermore, for ϵ-approximate Nash, this holds for probability less than $\frac{1+\epsilon}{M}$.*

Proof. By the assumption, for any Nash equilibrium of $\mathcal{D}^{\mathrm{M}}(\mathcal{G})$, at least one player does not play a fully mixed strategy. Assume that for one Nash equilibrium (x^*, y^*), there exists a pure strategy t such that $y^*_t > \frac{1}{M}$ and $x^*_t > \frac{1}{M}$. Without loss of generality, assume that the column player does not play $t' \in [n]$. We obtain a contradiction by showing that the row player can deviate from t to t' to obtain a better payoff. Letting (t, y^*) denote the strategy profile in which the row player plays the pure strategy t with probability 1 and the column player plays the mixed strategy y^*, we have:

$$U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(t, y^*) = \sum_{1 \le j \le n} y^*_j U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(t, j) < 1 + y^*_t U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(t, t) \le 0 \tag{5.2}$$

The last inequality comes from the fact that $y^*_t \ge \frac{1}{M}$. Assuming y^* is fixed, if we let $U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x^*_t, y^*)$ (where x^*_t does not denote a mixed strategy but rather the probability with which the pure strategy t is played in x^*) denote $x^*_t \cdot U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(t, y^*)$, we have:

$$U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x^*_t, y^*) = x^*_t \sum_{1 \le j \le n} y^*_j U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(t, j) \le x^*_t(1 + y^*_t U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(t, t)) < 0 \tag{5.3}$$

Now suppose that the row player moves all the weight x^*_t from t to t'. Call the modified strategy x'. It is obvious that $U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x'_{t'}, y^*) \ge 0$ because $y^*_{t'} = 0$.

$$U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x^*, y^*) - U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x', y^*) = x^*_t[\sum_{j=1}^{n} y^*_j U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(t, j) - \sum_{j=1}^{n} y^*_j U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(t', j)]$$

$$U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x^*, y^*) - U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x', y^*) = U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x_t^*, y^*) - U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x'_{t'}, y^*) < 0$$

In conclusion, assuming that for all j, we have $y_j^* > \frac{1}{M}$, we reach a contradiction. We can easily extend the proof to work with the approximate Nash equilibrium as well. □

Proposition 5.4.7. *For a diagonally modified game $\mathcal{D}^M(\mathcal{G})$ where at least one player does not play a fully mixed strategy, the problem of finding a M-semi-disjoint Nash equilibrium is in PPAD.*

Proof. We know that the problem of finding a Nash equilibrium of an arbitrary bi-matrix game is in PPAD. By Lemma 5.4.6, *any* Nash equilibrium of $\mathcal{D}^{\mathrm{M}}(\mathcal{G})$ satisfying the condition is a M-semi-disjoint Nash equilibrium. □

Proposition 5.4.8. *For any bi-matrix game $\mathcal{G}$, a M-semi generalized disjoint equilibrium is a $(1 - \frac{n}{M})$ generalized far equilibrium.*

Proof. The proof follows immediately by the definition of generalized far equilibrium. □

Proof of Theorem 5.2.3 Part (c)

We are now ready to prove Theorem 5.2.3 Part (c).

Proof. Suppose $\mathcal{G}$ is a game whose diagonally modified version has no fully-mixed Nash equilibrium. We first show that the problem of finding a M-semi generalized disjoint equilibrium with approximation error $\frac{6n}{M}$ is in PPAD. Next, by Proposition 5.4.7 and 5.4.8, the problem of finding a generalized $(1 - \frac{n}{M}, \frac{6n}{M})$-far equilibrium is in PPAD.

Recall that the problem of finding a M-semi disjoint Nash is in PPAD for diagonally modified games if the diagonally modified game does not have a fully mixed Nash equilibrium. Let $\mathcal{S}_{\mathrm{M}}^n(y^*)$ be the row player's constraint set for a M-semi generalized disjoint equilibrium. Now suppose that (x^*, y^*) is a M-semi disjoint Nash for $\mathcal{D}^{\mathrm{M}}(\mathcal{G})$. To prove that (x^*, y^*) is a M-semi generalized disjoint equilibrium with approximation error $\frac{6n}{M}$ in $\mathcal{G}$, we need to prove the following for all $f \in \mathcal{S}_{\mathrm{M}}^n(y^*)$ (The proof for the column player is similar):

$$U_1^{\mathcal{G}}(x^*, y^*) + \frac{6n}{M} > U_1^{\mathcal{G}}(f, y^*)$$

Write $x^* = x^b + x^s$ and $y^* = y^b + y^s$ where the superscript b includes the strategies that are used with probability greater than $\frac{1}{M}$ in x^* while the superscript s is used for the other strategies (similarly for y^b and y^s). It is obvious that $U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x^b, y^b) = U_1^{\mathcal{G}}(x^b, y^b)$ since x^b and y^b both are associated with strategies that have disjoint supports. Since (x^*, y^*) is a M-semi disjoint Nash equilibrium for $\mathcal{D}^{\mathrm{M}}(\mathcal{G})$, it will satisfy the following:

$$\forall f \in \mathcal{S}_{\mathrm{M}}^n(y^*) \subset \Delta_n, \quad U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x^*, y^*) \geq U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(f, y^*) \tag{5.4}$$

The following inequality holds by applying Lemma 5.2.7 to Equation 5.4. If we write $\mathcal{D}^{\mathrm{M}}(\mathcal{G})_1(x^b, y^*)$ to denote the share of the row player's payoff coming from the sub-distribution x^b, we have:

$$\forall f \in \mathcal{S}_{\mathrm{M}}^n(y^*) \subset \Delta_n, \quad U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x^b, y^*) + \frac{n}{M} \geq U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(f, y^*)$$

Similarly:

$$\forall f \in \mathcal{S}_{\mathrm{M}}^n(y^*) \subset \Delta_n, \quad U_1^{\mathcal{G}}(x^b, y^b) + \frac{n}{M} + \frac{n}{M} \geq U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(f, y^*)$$

and by the fact that $|U_1^{\mathcal{G}}(f, y^b) - U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(f, y^*)| \leq \frac{n}{M}$:

$$\forall f \in \mathcal{S}_{\mathrm{M}}^n(y^*) \subset \Delta_n, \quad U_1^{\mathcal{G}}(x^b, y^b) + \frac{2n}{M} + \frac{n}{M} \geq U_1^{\mathcal{G}}(f, y^b)$$

Now, we select one of the elements in the support of x^b and y^b and give them the summation of all of the probabilities in x^s and y^s, respectively. These two new distributions x' and y' form a M-semi generalized disjoint equilibrium with $\frac{6n}{M}$ as the approximation error by Lemma 5.2.7 (since $||x' - x^b||_1 \leq \frac{n}{M}$ and $||y' - y^b||_1 \leq \frac{n}{M}$):

$$\forall f \in \mathcal{S}_{\mathrm{M}}^n(y') \subset \Delta_n, \quad U_1^{\mathcal{G}}(x', y') + \frac{6n}{M} \geq U_1^{\mathcal{G}}(f, y')$$

□

5.4.4 An Interesting Connection

The following proposition shows how diagonally modified games provide a connection between generalized disjoint equilibrium and Nash equilibrium.

Proposition 5.4.9. *For $\theta \geq \frac{1}{M}$, any (θ, ϵ)-restricted disjoint equilibrium in $\mathcal{G}$ forms an ϵ-approximate disjoint Nash equilibrium in $\mathcal{D}^M(\mathcal{G})$.*

Proof. If we have a θ-restricted disjoint equilibrium (x^*, y^*) for $\mathcal{G}$, we prove that this equilibrium will be a disjoint Nash equilibrium for $\mathcal{D}^{\mathrm{M}}(\mathcal{G})$. For all pure strategies $s \in S$, we need to prove:

$$U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x^*, y^*) \geq U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(s, y^*)$$

We can easily show that this inequality is correct for all $s \notin Supp(y^*)$ since we know that $U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(x^*, y^*) = U_1^{\mathcal{G}}(x^*, y^*)$ and $U_1^{\mathcal{G}}(s, y^*) = U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(s, y^*)$. Next, consider any $s \in Supp(y^*)$. In (x^*, y^*), it is impossible to have probabilities smaller than $\frac{1}{M}$ except zero. Then, we can conclude that $U_1^{\mathcal{D}^{\mathrm{M}}(\mathcal{G})}(s, y^*) < 0$. This is followed by the fact that s is played with probability 1 and the column player has to s with probability greater than $\frac{1}{M}$.

We know that the payoff of (x^*, y^*) is positive in $\mathcal{D}^{\mathrm{M}}(\mathcal{G})$, which means it is impossible to deviate to s in this case. This is correct for the column player by a symmetric argument. This also holds for the approximation notion by a similar argument.

□

5.4.5 Restricted Far and Disjoint Equilibrium and Proof of Theorem 5.2.3 Part (d) and (e)

We begin with the following proposition that uses the hardness of disjoint Nash equilibrium with no minor probabilities to establish the hardness result for generalized disjoint equilibrium.

Essential Elements to Prove Theorem 5.2.3 Part (d)

Definition 5.4.10. *(Major Nash): A strategy profile is a θ-*major Nash equilibrium *if all strategies in the support of all players are played with probability greater than θ. For the approximation version, we consider (θ, ϵ)-*major Nash equilibrium.

Proposition 5.4.11. *Given an instance of 3CNF ϕ, $\mathcal{C}(\phi, \epsilon)$ has an EXACT disjoint Nash equilibrium in which all strategies that are in the support of either of the players are played with probability strictly greater than $\frac{1}{c}$ iff the given formula ϕ is satisfiable.*

Proof. When the formula is satisfiable, a solution similar to Proposition 5.3.2 can be attained where all players play some literals with probability $\frac{1}{n}$. We need to show that when a formula is not satisfiable, the only Nash equilibrium of the game $\mathcal{C}(\phi, \epsilon)$ can have compared to $\mathcal{G}(\phi, \epsilon)$ is when players play the following strategy:

$$x^c = y^c = (\frac{1}{c}f_1, \dots, \frac{1}{c}f_c)$$

It is easy to check that (x^c, y^c) is a Nash equilibrium since the sub-game induced by strategies in $F^{\mathcal{C}}$ is just a simple generalization of *rock-paper-scissors* (where the payoffs are also scaled). It is also not hard to see that in any (EXACT) Nash equilibrium, if none of the players play any strategy in $F^{\mathcal{C}}$, all strategies in $C \cup V$ cannot be played with a positive probability by an analysis similar to [52] and Lemma 5.3.4 since social welfare with the use of these strategies must be less than $2n - 2$ and at least one of the players will try to achieve $n - 1$ by deviating to any $f^i \in F^{\mathcal{C}}$. Using any strategy $f^i \in F^{\mathcal{C}}$ will cause the opponent player to play f^j where $j = i + 1 \pmod{c}$ and this can continue forever.

Without loss of generality, suppose that the row player plays one strategy $f^i \in F^{\mathcal{C}}$ with probability greater than $\frac{\epsilon}{n}$. A simple calculation shows that the expected payoff of the column player playing the pure strategy $f^j \in F^{\mathcal{C}}$ where $j = i + 1 \pmod{c}$ is always greater than n (f^i cannot be played with a minor probability). This indeed will prevent the column player from playing any strategy from L, C, or V. The sub-game $\mathcal{C}_F$ induced by considering the strategy set to be $F^{\mathcal{C}} \subset S^{\mathcal{C}}$ has only one unique Nash equilibrium (x^c, y^c). For this strategy profile, all strategies in the support of both players are playing with a probability of at most $\frac{1}{c}$. □

Remark 5.4.12. We can also observe that the expected guaranteed payoff of any Nash equilibrium of $\mathcal{C}(\phi, \epsilon)$ is $\frac{1}{c} \cdot \frac{n^2}{\epsilon} + \frac{1}{c} \cdot \frac{2n^2}{\epsilon} = \frac{3n^2}{c\epsilon}$.

5.4.6 Proof of Theorem 5.2.3 Part (d)

In order to prove Theorem 5.2.3 part (d), we begin with a generalization of Proposition 5.3.2 by embedding a zero-sum sub-game into the game. We will modify $\mathcal{G}(\phi, \epsilon)$ to $\mathcal{C}(\phi, \epsilon)$ ($\frac{1}{\epsilon} > c > n$) by removing rules 10-14 and adding a series of rules instead. Let $S^{\mathcal{C}} \equiv S_1^{\mathcal{C}} = S_2^{\mathcal{C}} = L_1 \cup L_2 \cup V \cup C \cup F^{\mathcal{C}}$ be the strategy set for both players where $F^{\mathcal{C}} = \{f^i \mid i \in [c]\}$.

1. $u_1^{\mathcal{C}}(x, f^i) = u_2^{\mathcal{C}}(f^j, x) = 0$ for all $i, j \in [c]$ and $x \in S^{\mathcal{C}} - F^{\mathcal{C}}$;

2. $u_1^{\mathcal{C}}(f^i, x) = u_2^{\mathcal{C}}(x, f^i) = n - 1$ for all $i, j \in [c]$ and $x \in S^{\mathcal{C}} - F^{\mathcal{C}}$;

3. $u_1^{\mathcal{C}}(f^i, f^j) = u_2^{\mathcal{C}}(f^j, f^i) = \frac{n^2}{\epsilon}$ for all $i = j \in [c]$;

4. $u_1^{\mathcal{C}}(f^i, f^j) = u_2^{\mathcal{C}}(f^j, f^i) = 2(\frac{n^2}{\epsilon})$ for all $i, j \in [c]$ and $i = j + 1 \pmod{c}$;

5. $u_1^{\mathcal{C}}(f^i, f^j) = u_2^{\mathcal{C}}(f^j, f^i) = 0$ for all $i, j \in [c]$ and $i \neq j + 1 \pmod{c}$;

6. $u_1^{\mathcal{C}}(l^2, *) = u_2^{\mathcal{C}}(*, l^1) = -2n$ for all $* \in S^{\mathcal{C}},\ l^1 \in L_1$ and $l^2 \in L_2$;

7. $u_2^{\mathcal{C}}(l^2, *) = u_1^{\mathcal{C}}(*, l^1) = 0$ for all $* \in S^{\mathcal{C}} - F^{\mathcal{C}},\ l^1 \in L_1$ and $l^2 \in L_2$;

It is not hard to see that the problem of deciding whether a bi-matrix game has a restricted (and also generalized) disjoint equilibrium is in NP (see Appendix A.4.4). The following proposition (hardness) concludes the proof of Theorem 5.2.3 part (d).

Proposition 5.4.13. *Given a formula ϕ with n variables, there exists a game such that for some θ and ϵ poly-bounded in n, there exists a (θ, ϵ)-restricted disjoint equilibrium in this game iff ϕ is satisfiable.*

Proof. We want to show that an arbitrary formula ϕ is satisfiable iff this modified version of $\mathcal{C}(\phi, \epsilon)$ has a $(\frac{1}{n} - \epsilon - \frac{1}{n^2}, \epsilon)$-restricted disjoint equilibrium. It is not hard to extend Proposition 5.4.11 to the approximate case (with $\epsilon = \frac{1}{2n^3}$) by combining it with the techniques used in the proof of Proposition 5.3.2. We duplicate all clauses and variables in addition to the strategies in $F^{\mathcal{C}}$ similarly to what we did in game $\mathcal{G}(\phi, \epsilon)$ and consider four functions that map copies of the associated strategies. We also note that we will need to assure that $\epsilon \leq \frac{1}{c} - \frac{2\epsilon^2}{n} \leq \frac{1}{n} - 2\epsilon - \frac{1}{n^2}$ by choosing a suitable c. Then, obviously $\frac{1}{c} - \frac{2\epsilon^2}{n}$ is greater than ϵ, $\frac{\epsilon}{n}$ and $\frac{2\epsilon}{n} + \frac{\epsilon^2}{2n^2}$.

Each player can play unassociated literals, clauses, and variables only with probability smaller than $\frac{\epsilon}{n}$ since a generalized disjoint equilibrium is also a disjoint Nash equilibrium in this game. This indicates that these strategies cannot be played in any $(\frac{1}{n} - 2\epsilon - \frac{1}{n^2}, \epsilon)$-restricted disjoint equilibrium. This statement also holds for the (associated copies of) variables and clauses that can be played only with probability less than ϵ by Lemma 5.3.4 and Proposition 5.3.15.

If none of the players play strategies in $F_1^{\mathcal{C}}$ or $F_2^{\mathcal{C}}$, an analysis similar to Proposition 5.3.2 can show that at least one of the players will be motivated to select some of the strategies in $F_1^{\mathcal{C}}$ or $F_2^{\mathcal{C}}$. If the players are allowed to play some strategies from $F_1^{\mathcal{C}}$ or $F_2^{\mathcal{C}}$, we will prove that the analysis of Proposition 5.4.11 with minor adjustments still holds. In this modified version of $\mathcal{C}(\phi, \epsilon)$ (with duplicated variables and clauses), in any EXACT generalized disjoint equilibrium with the condition that at least one of the players plays strategies from $F_1^{\mathcal{C}}$ or $F_2^{\mathcal{C}}$ with probability greater than $\frac{\epsilon}{2n}$, no associated literal, clause, or variable strategy can be played with a positive probability. In the approximate case, literals, clauses, or variables can be played with probability of at most $\frac{\epsilon^2}{n^2}$. This is because we know there exists one strategy from $F_1^{\mathcal{C}}$ or $F_2^{\mathcal{C}}$ with probability greater than $\frac{2\epsilon}{n}$ and playing $f^{i+1 \pmod c}$ (from $F_1^{\mathcal{C}}$ or $F_2^{\mathcal{C}}$ depending on the player) will provide ϵ difference in the payoff at least. In conclusion, these strategies cannot be played in any $(\frac{1}{n} - 2\epsilon - \frac{1}{n^2}, \epsilon)$-restricted disjoint equilibrium.

Recall that each player must guarantee a payoff of $\frac{3n^2}{c\epsilon}$ by Remark 5.4.12 (and $(\frac{3n^2}{c\epsilon}) - \epsilon$ for the approximate version). Finally, if a player prefers to play any strategy from $F_1^{\mathcal{C}}$ or $F_2^{\mathcal{C}}$ with probability smaller than $\frac{1}{c} - \frac{2\epsilon^2}{n}$, then there exists at least one strategy in $F_1^{\mathcal{C}}$ or $F_2^{\mathcal{C}}$ such that this player plays the strategy with probability greater than $\frac{1}{c} + \frac{2\epsilon^2}{n^2}$ since we showed that the first and the second player can only use strategies in $F_1^{\mathcal{C}}$ and $F_2^{\mathcal{C}}$ in the restricted equilibrium respectively. The opponent can get a payoff of at least $(\frac{3n^2}{c\epsilon}) + \frac{2\epsilon^2}{n^2}(\frac{n^2}{\epsilon})$ which is equal to $\frac{3n^2}{c\epsilon} + 2\epsilon$. This follows by the fact that the structure of the sub-game that is induced by strategies in $F_1^{\mathcal{C}}$ and $F_2^{\mathcal{C}}$ is a game similar to *rock-paper-scissors*.

The maximum social welfare of any equilibrium in this sub-game is $2\frac{3n^2}{c\epsilon}$ and this shows that one player cannot meet the specified guarantee $\frac{3n^2}{c\epsilon} - \epsilon$. In conclusion, in this case, there will be no ϵ-generalized disjoint equilibrium with a minimum probability greater than $\frac{1}{c} - \frac{2\epsilon^2}{n} < \frac{1}{n} - \epsilon - \frac{1}{n^2}$. This completes the proof that an arbitrary formula ϕ is satisfiable iff this modified version of $\mathcal{C}(\phi, \epsilon)$ has a $(\frac{1}{n} - \epsilon - \frac{1}{n^2}, \epsilon)$-restricted disjoint equilibrium.

□

Corollary 5.4.14. *For a given formula ϕ with n variables, there exists a game such that for some ϵ and θ poly-bounded in n, a (θ, ϵ)-major Nash exists for this game iff the formula ϕ is satisfiable.*

5.4.7 Proof of Theorem 5.2.3 Part (e)

In short, establishing this theorem necessitates a meticulous integration of the constructions employed to prove part (d) of Theorem 5.2.1 and part (d) of Theorem 5.2.3.

Proposition 5.4.15. *Given an arbitrary formula ϕ with n variables, there exists a* *game $\mathcal{R}(\phi, \delta, \epsilon)$ such that for some ϵ,θ and for any given δ poly-bounded in n, there exists a restricted $(\theta, \delta, \epsilon)$-far equilibrium for $\mathcal{R}(\phi, \delta, \epsilon)$ iff ϕ is satisfiable.*

Proof. We only prove the theorem for $\delta < \frac{1}{n} - \frac{1}{n^2} - \frac{1}{n^3}$ and the other case is similar (see Proposition 5.3.17). We generate the final game $\mathcal{R}(\phi, \delta, \epsilon)$ which is a combination of the games $\mathcal{D}(\phi, \delta, \epsilon)$ and $\mathcal{C}(\phi, \epsilon)$. This game has the strategy set $S^{\mathcal{R}} = L^{\mathcal{D}} \cup L_1^{\mathcal{D}} \cup L_2^{\mathcal{D}} \cup V_1 \cup V_2 \cup C_1 \cup C_2 \cup F_1^{\mathcal{R}} \cup F_2^{\mathcal{R}}$. The literal strategies of this game are the same as the game $\mathcal{D}(\phi, \delta, \epsilon)$ and we duplicated the clauses and variables similar to $\mathcal{C}(\phi, \epsilon)$. The function $\mathrm{g}_1 : L^{\mathcal{D}} \cup L_1^{\mathcal{D}} \cup L_2^{\mathcal{D}} \to L$ is defined similarly based on g' in the game $\mathcal{D}(\phi, \delta, \epsilon)$. We also have three functions $\mathrm{g}_2 : V_1 \cup V_2 \to V$ and $\mathrm{g}_3 : C_1 \cup C_2 \to C$ (similar to g_1, g and g') that map the copied strategies to their original variables and clauses in ϕ. The function $\mathrm{g}_4 : F_1^{\mathcal{R}} \cup F_2^{\mathcal{R}} \to F^{\mathcal{C}}$ can be used for strategies in $F^{\mathcal{C}}$. The payoff matrices of the game $\mathcal{R}(\phi, \delta, \epsilon)$ are defined as follows:

1. $u_1^{\mathcal{R}}(l^1, l^2) = u_2^{\mathcal{R}}(l^2, l^1) = n - 1$ for all $l^1, l^2 \in L^{\mathcal{D}}$ such that $\mathrm{g}_1(l^1) \neq -\mathrm{g}_1(l^2)$;
2. $u_1^{\mathcal{R}}(l^1, l^2) = u_2^{\mathcal{R}}(l^1, l^2) = n - 1$ for all $l^1 \in L_1^{\mathcal{D}},\ l^2 \in L_2^{\mathcal{D}}$ with $\mathrm{g}_1(l^1) \neq -\mathrm{g}_1(l^2)$;
3. $u_1^{\mathcal{R}}(l^1, l^2) = u_2^{\mathcal{R}}(l^1, l^2) = n - 1$ for all $l^1 \in L^{\mathcal{D}},\ l^2 \in L_2^{\mathcal{D}}$ with $\mathrm{g}_1(l^1) \neq -\mathrm{g}_1(l^2)$;
4. $u_1^{\mathcal{R}}(l^1, l^2) = u_2^{\mathcal{R}}(l^1, l^2) = n - 1$ for all $l^1 \in L_1^{\mathcal{D}},\ l^2 \in L^{\mathcal{D}}$ with $\mathrm{g}_1(l^1) \neq -\mathrm{g}_1(l^2)$;
5. $u_1^{\mathcal{R}}(-l, l) = u_2^{\mathcal{R}}(-l, l) = n - 4$ for all $l \in L^{\mathcal{D}}$;
6. $u_1^{\mathcal{R}}(-l^1, l^2) = u_2^{\mathcal{R}}(-l^1, l^2) = n - 4$ for all $l^1 \in L_1^{\mathcal{D}}$ $, l^2 \in L_2^{\mathcal{D}}$ and $\mathrm{g}_1(l^1) = \mathrm{g}_1(l^2)$;
7. $u_1^{\mathcal{R}}(-l^1, l^2) = u_2^{\mathcal{R}}(-l^1, l^2) = n - 4$ for all $l^1 \in L^{\mathcal{D}}$ $, l^2 \in L_2^{\mathcal{D}}$ and $\mathrm{g}_1(l^1) = \mathrm{g}_1(l^2)$;
8. $u_1^{\mathcal{R}}(-l^1, l^2) = u_2^{\mathcal{R}}(-l^1, l^2) = n - 4$ for all $l^1 \in L_1^{\mathcal{D}}$ $, l^2 \in L^{\mathcal{D}}$ and $\mathrm{g}_1(l^1) = \mathrm{g}_1(l^2)$;
9. $u_1^{\mathcal{R}}(v^1, l) = u_2^{\mathcal{R}}(l, v^2) = n$ for all $v^1 \in V,\ v^2 \in V_2,\ l \in L^{\mathcal{D}}$ with $\mathrm{g}_2(v^1) = g(v^2)$ and $\mathrm{v}(\mathrm{g}_1(l)) \neq \mathrm{g}_2(v^1)$;
10. $u_1^{\mathcal{R}}(v^1, l^2) = u_2^{\mathcal{R}}(l^1, v^2) = n$ for all $v^1 \in V_1,\ v^2 \in V_2,\ l^1 \in L_1^{\mathcal{D}},\ l^2 \in L_2^{\mathcal{D}}$ with $\mathrm{g}_1(l^1) = \mathrm{g}_1(l^2)$, $\mathrm{g}_2(v^1) = \mathrm{g}_2(v^2)$ and $\mathrm{v}(\mathrm{g}_1(l^1)) \neq \mathrm{g}_2(v^1)$;
11. $u_1^{\mathcal{R}}(l, x^2) = u_2^{\mathcal{R}}(x^1, l) = n - 4$ for all $l \in L^{\mathcal{D}},\ x^1 \in V_1 \cup C_1,\ x^2 \in V_2 \cup C_2$;
12. $u_1^{\mathcal{R}}(l^1, x^2) = u_2^{\mathcal{R}}(x^1, l^2) = n{-}4$ for all $l^1 \in L_1^{\mathcal{D}},\ l^2 \in L_2^{\mathcal{D}},\ x^1 \in V_1 \cup C_1,\ x^2 \in V_2 \cup C_2$ and $\mathrm{g}_1(l^1) = \mathrm{g}_1(l^2)$;

13. $u_1^{\mathcal{R}}(v^1, l) = u_2^{\mathcal{R}}(l, v^2) = 0$ for all $v^1 \in V_1,\ v^2 \in V_2,\ l \in L^{\mathcal{D}}$ with $\mathrm{g}_2(v^1) = \mathrm{g}_2(v^2)$ and $\mathrm{v}(\mathrm{g}_1(l)) = \mathrm{g}_2(v^1)$;

14. $u_1^{\mathcal{R}}(v^1, l^2) = u_2^{\mathcal{R}}(l^1, v^2) = 0$ for all $v^1 \in V_1,\ v^2 \in V_2,\ l^1 \in L_1^{\mathcal{D}},\ l^2 \in L_2^{\mathcal{D}}$ with $\mathrm{g}_1(l^1) = \mathrm{g}_1(l^2)$, $\mathrm{g}_2(v^1) = \mathrm{g}_2(v^2)$ and $\mathrm{v}(\mathrm{g}(l^1)) = \mathrm{g}_2(v^1)$;

15. $u_1^{\mathcal{R}}(v^1, x^2) = u_2^{\mathcal{R}}(x^1, v^2) = n - 4$ for all $v^1 \in V_1,\ v^2 \in V_2,\ x^1 \in V_1 \cup C_1,\ x^2 \in V_2 \cup C_2$ and $\mathrm{g}_2(v^1) = \mathrm{g}_2(v^2)$;

16. $u_1^{\mathcal{R}}(c^1, l) = u_2^{\mathcal{R}}(l, c^2) = n$ for all $c^1 \in C_1,\ c^2 \in C_2,\ l \in L^{\mathcal{D}}$ with $\mathrm{g}_3(c^1) = \mathrm{g}_3(c^2)$ and $\mathrm{g}_1(l) \notin \mathrm{g}_3(c^1)$;

17. $u_1^{\mathcal{R}}(c^1, l^2) = u_2^{\mathcal{R}}(l^1, c^2) = n$ for all $c^1 \in C_1,\ c^2 \in C_2,\ l^1 \in L_1^{\mathcal{D}},\ l^2 \in L_2^{\mathcal{D}}$ with $\mathrm{g}_1(l^1) = \mathrm{g}_1(l^2)$, $\mathrm{g}_3(c^1) = \mathrm{g}_3(c^2)$ and $\mathrm{g}_1(l^1) \notin \mathrm{g}_3(c^1)$;

18. $u_1^{\mathcal{R}}(c^1, l) = u_2^{\mathcal{R}}(l, c^2) = 0$ for all $c^1 \in C_1,\ c^2 \in C_2,\ l \in L^{\mathcal{D}},$ with $\mathrm{g}_1(l^1) = \mathrm{g}_1(l^2)$, $\mathrm{g}_3(c^1) = \mathrm{g}_3(c^2)$ and $\mathrm{g}_1(l) \in \mathrm{g}_3(c^1)$;

19. $u_1^{\mathcal{R}}(c^1, l^2) = u_2^{\mathcal{R}}(l^1, c^2) = 0$ for all $c^1 \in C_1,\ c^2 \in C_2,\ l^1 \in L_1^{\mathcal{D}},\ l^2 \in L_2^{\mathcal{D}}$ with $\mathrm{g}_1(l^1) = \mathrm{g}_1(l^2)$ and $\mathrm{g}_1(l^1) \in \mathrm{g}_3(c^1)$;

20. $u_1^{\mathcal{R}}(c^1, x^2) = u_2^{\mathcal{R}}(x^1, c^2) = n-4$ for all $v^1 \in V_1,\ v^2 \in V_2,\ x^1 \in V_1 \cup C_1,\ x^2 \in V_2 \cup C_2$ and $\mathrm{g}_3(c^1) = \mathrm{g}_3(c^2)$;

21. $u_1^{\mathcal{R}}(x, f^i) = u_2^{\mathcal{R}}(f^j, x) = 0$ for any $f^i \in F_2^{\mathcal{R}}$ and $f^j \in F_1^{\mathcal{R}}$ where $x \in S^{\mathcal{R}} - F_1^{\mathcal{R}} - F_2^{\mathcal{R}}$;

22. $u_1^{\mathcal{R}}(f^i, x) = u_2^{\mathcal{R}}(x, f^j) = n - 1$, for any $f^i \in F_1^{\mathcal{R}}$ and $f^j \in F_2^{\mathcal{R}}$ where $x \in S^{\mathcal{R}} - F_1^{\mathcal{R}} - F_2^{\mathcal{R}}$;

23. $u_1^{\mathcal{R}}(f^i, f^j) = u_2^{\mathcal{R}}(f^i, f^j) = \frac{dn^2}{\delta\epsilon}$ for any $f^i \in F_1^{\mathcal{R}}$ and $f^j \in F_2^{\mathcal{R}}$ such that $\mathrm{g}_4(f^i) = \mathrm{g}_4(f^j) = f$;

24. $u_1^{\mathcal{R}}(f_1^i, f_2^j) = u_2^{\mathcal{R}}(f_1^j, f_2^i) = 2(\frac{dn^2}{\delta\epsilon})$ for $f_1^i, f_1^j \in F_1^{\mathcal{R}}$ and $f_2^i, f_2^j \in F_2^{\mathcal{R}}$ where $\mathrm{g}_4(f_1^i) = \mathrm{g}_4(f_2^i) = f^i$ and $\mathrm{g}_4(f_1^j) = \mathrm{g}_4(f_2^j) = f^j$ (the superscripts i and j denote the i-th and j-th member of $F^{\mathcal{C}}$) while $i = j + 1 \pmod{c}$;

25. $u_1^{\mathcal{R}}(f_1^i, f_2^j) = u_2^{\mathcal{R}}(f_1^j, f_2^i) = 0$ for $f_1^i, f_1^j \in F_1^{\mathcal{R}}$ and $f_2^i, f_2^j \in F_2^{\mathcal{R}}$ where $\mathrm{g}_4(f_1^i) = \mathrm{g}_4(f_2^i) = f^i$ and $\mathrm{g}_4(f_1^j) = \mathrm{g}_4(f_2^j) = f^j$ while $i \neq j + 1 \pmod{c}$;

26. $u_1^{\mathcal{R}}(l^2, *) = u_2^{\mathcal{R}}(*, l^1) = -2n$ for all $* \in S^{\mathcal{R}},\ l^1 \in L_1$ and $l^2 \in L_2$;

27. $u_2^{\mathcal{R}}(l^2, *) = u_1^{\mathcal{R}}(*, l^1) = 0$ for all $* \in S^{\mathcal{R}} - F_1^{\mathcal{R}} - F_1^{\mathcal{R}},\ l^1 \in L_1$ and $l^2 \in L_2$;

28. $u_1^{\mathcal{R}}(v^2, *) = u_2^{\mathcal{R}}(*, v^1) = -2n$ for all $* \in S^{\mathcal{R}},\ v^1 \in V_1$ and $v^2 \in V_2$;

29. $u_2^{\mathcal{R}}(v^2, *) = u_1^{\mathcal{R}}(*, v^1) = 0$ for all $* \in S^{\mathcal{R}} - F_1^{\mathcal{R}} - F_2^{\mathcal{R}},\ v^1 \in V_1$ and $v^2 \in V_2$;

30. $u_1^{\mathcal{R}}(c^2, *) = u_2^{\mathcal{R}}(*, c^1) = -2n$ for all $* \in S^{\mathcal{R}},\ c^1 \in C_1$ and $c^2 \in C_2$;

31. $u_2^{\mathcal{R}}(c^2, *) = u_1^{\mathcal{R}}(*, c^1) = 0$ for all $* \in S^{\mathcal{R}} - F_1^{\mathcal{R}} - F_2^{\mathcal{R}},\ c^1 \in C_1$ and $c^2 \in C_2$;

32. $u_1^{\mathcal{R}}(f^2, *) = u_2^{\mathcal{R}}(*, f^1) = -2n$ for all $* \in S^{\mathcal{R}},\ f^1 \in F_1^{\mathcal{R}}$ and $f^2 \in F_2^{\mathcal{R}}$;

33. $u_2^{\mathcal{R}}(f^2, *) = u_1^{\mathcal{R}}(*, f^1) = 0$ for all $* \in S^{\mathcal{R}} - F_1^{\mathcal{R}} - F_2^{\mathcal{R}},\ f^1 \in F_1^{\mathcal{R}}$ and $f^2 \in F_2^{\mathcal{R}}$.

If the formula ϕ is satisfiable, then there exists a $(\frac{\delta}{d}, \delta, 0)$-restricted far equilibrium where d is driven by δ (see Proposition 5.3.18). We show that if ϕ is not satisfiable, there exists no $(\frac{\delta}{d} - \frac{2\delta\epsilon^2}{dn}, \delta, \epsilon)$-restricted far equilibrium. If no (associated) strategy from $F_1^{\mathcal{R}}$ or $F_2^{\mathcal{R}}$ is played, one of the players cannot get the guaranteed payoff $n-1-\epsilon$ and will choose one strategy from these sets, similar to all previous results. If there exists one (associated) strategy in $F_1^{\mathcal{R}}$ or $F_2^{\mathcal{R}}$ such that the row or the column player plays this strategy with probability greater than $\frac{\delta\epsilon}{2dn}$, all strategies that are related to literals, variables, and clauses cannot be played with probability greater than $\frac{\delta\epsilon^2}{dn^2}$ by following an analysis similar to Theorem 5.4.13.

Finally, if a player prefers to play any strategy from $F_1^{\mathcal{R}}$ or $F_2^{\mathcal{R}}$ with probability smaller than $\frac{1}{c} - \frac{2\delta\epsilon^2}{dn}$, then there exists at least one strategy such that the opponent can get a payoff of at least $(\frac{3dn^2}{c\delta\epsilon}) + \frac{2\delta\epsilon^2}{dn^2}(\frac{dn^2}{\delta\epsilon})$ which is equal to $\frac{3dn^2}{c\delta\epsilon} + 2\epsilon$. The maximum social welfare of any equilibrium of this game is sub-game $2\frac{3dn^2}{c\delta\epsilon}$ and this shows that one player cannot meet the guarantee $\frac{3dn^2}{c\delta\epsilon} - \epsilon$. If we set $c = \lceil \frac{d}{\delta} \rceil + 1$, then $\frac{1}{c} < \frac{\delta}{d}$ and $\frac{1}{c} - \frac{2\delta\epsilon^2}{dn^2}$ is always positive (generating $\mathcal{R}(\phi, \delta, \epsilon)$ with $|F_1^{\mathcal{R}}| = |F_2^{\mathcal{R}}| > \frac{\delta}{d}$). This completes the proof that an arbitrary formula ϕ is satisfiable iff the $\mathcal{R}(\phi, \delta, \epsilon)$ has a $(\frac{\delta}{d}, \delta, \epsilon)$-restricted far equilibrium.

□

Chapter 6

A Meta-Heuristic Approach for Strategic Fair Division Problems

In this chapter, we propose a framework for modeling item allocation with both non-strategic and strategic agents, with suitable fairness constraints. As mentioned in the introduction, key constraints ensure no player can use an item chosen by another, while further conditions guarantee minimum payoffs.

Starting from the fair allocation of indivisible items (the *simple arrangement problem*), we show NP-completeness. Extending to strategic settings, we define the *strategic arrangement problem*, a constrained Nash equilibrium framework for fair allocation, and prove its NP-completeness. The strategic arrangement problem is a more generalized version of a problem introduced in the previous chapter (see 5.1.5). Prior algorithms for approximate equilibria, such as [130] have limitations (see Section 6.4). Also, to the best of our knowledge, limited progress has been made on local search- or metaheuristic-based techniques for the computation of unconstrained Nash equilibria with mixed strategies [42, 69]. This gap motivates our use of metaheuristic algorithms in the constrained setting. We employ a human-inspired meta-heuristic, in particular, a modified *bus transportation algorithm* [34], which enables faster search for constrained equilibria, providing solutions for the strategic arrangement problem.

6.1 The Arrangement Problem

First, we provide the definition of a fair allocation, which is a general definition in fair division literature.

Definition 6.1.1. *A simple* item allocation *problem can be specified by a set I that includes indivisible items and a group of k agents. Each agent i has a specific value $V_i(t)$ (preference) for each item t. An* allocation *is a partition of I into k disjoint subsets. An allocation is* fair *if it satisfies all specified fairness constraints for all players.*

In the non-strategic setting, we consider the fairness constraint to be a minimum specified guaranteed summation of the values of all of the items in their allocation. In other words, a fair allocation will divide the items among the agents in a way that all agents can achieve a certain aggregated (additive) value. Below, we provide the definition of the simple arrangement problem.

6.1.1 Simple Arrangement Problem

Definition 6.1.2. *A* simple arrangement problem *$\mathcal{A}$ is specified by $\langle K, I, J, U, r\rangle$, where $K = [c]$ is the set of* agents (players)*, $I = [n]$ is the set of* items*, $J = [m]$ is the set of* places*, $U = [U]_{k \times n \times m}$ is the* utility matrix *with* utility threshold *r. A* multi-strategy *for each player k is a set $T_k \subset I_k \times J_k$ where $I_k \subseteq I$ and $J_k \subseteq J$.*

A multi-strategy profile $\boldsymbol{T} = (T_1, \ldots, T_k)$ that has all agents' multi-strategies is legal *if the specified items are placed in the specified places subject to the following conditions:*

- *For each place $j \in J$, there exists exactly one item $i \in I$ and agent k such that $(i, j) \in T_k$ (Each place is filled with one item).*
- *For all k and k', $I_k \cap I_{k'} = \emptyset$ (Each item can be picked by at most one agent).*
- *For each item $i \in I_k$, there can exist at most one place $j \in J_k$ such that $(i, j) \in T_k$ (Each item can be used once).*

Given a multi-strategy of an agent, we define the payoff *$P_k = \sum_{j \in J_k} U(k, T_k(j), j)$ for agent k, where $T_k(j)$ will give the item that is placed in place j by player k. The final goal of this problem is to find a multi-strategy profile such that each player gets at least the threshold payoff r.*

It is easy to observe that this problem is more complex compared to a standard resource allocation problem since there are multiple choices and constraints for the agents. In general, a fair item allocation algorithm may allocate a single item to one agent for each round. We can simply consider places to be the rounds in which one item is allocated. Then, if we have $|I|$ items, we will have $|I|$ rounds. Thus, a solution (a legal strategy profile) to the simple arrangement problem $\mathcal{A}\langle K, I, I, U, r\rangle$ can form a solution to a simple item allocation problem with the item's set I where $U(k, i, j) = V_k(i)$. In other words, in a setting without strategic behavior, the agents are not concerned with the round in which an item is allocated to them. Instead, they are only concerned with the value of the items given by a fair algorithm.

Next, we will show that even this simple version of the arrangement problem is NP-complete.

Theorem 6.1.3. *The problem of whether a given simple arrangement problem $\mathcal{A} = \langle K, I, J, U, r\rangle$, has a legal multi-strategy profile for which each agent obtains a payoff of at least r is NP-complete.*

To prove NP-hardness of the simple arrangement problem, we use the following well-known NP-complete problem.

Definition 6.1.4. Partition *is the following problem: given $I \subset [n]$, are there disjoint subsets $P_1, P_2 \subseteq I$ such that $\sum_{s\in P_1} s = \sum_{s\in P_2} s$ and $P_1 \cup P_2 = S$.*

We are now ready to proceed to the proof of NP-completeness.

Proof. It is obvious that the problem is in NP because we can check whether a given multi-strategy profile can give the required payoff threshold for all agents. To prove NP-hardness, we begin with an arbitrary instance of partition where $I \subset [n]$ is the input of the partition problem, and generate a simple arrangement problem in the following. $\mathcal{A}\langle\{1, 2\}, I, I, U, \frac{1}{2}\sum_{t\in I} t\rangle$ is a simple arrangement problem with the specified payoff threshold $\frac{1}{2}\sum_{t\in I} t$ where U is defined as follows:

$$U(1, i, j) = U(2, i, j) = \begin{cases} i & \text{if } i = j \\ -2n^2 & \text{if } i \neq j \end{cases}$$

If we have a partition $\{P_1, P_2\}$ (such that $P_1, P_2 \subset I$), we know that $P_1 \cup P_2 = I$ and $\sum_{t\in P_1} t = \sum_{t\in P_2} t = \frac{1}{2}\sum_{t\in I} t$. We define the multi-strategy of each agent k to be $T_k = \{(t, t) \mid t \in P_k\}$. Each agent will get the payoff $\frac{1}{2}\sum_{t\in I} t$ and all conditions (legal multi-strategy profile) are satisfied since $\{P_1, P_2\}$ is a partition.

Suppose that we have a solution to this simple arrangement problem and T_1 and T_2 are multi-strategies of both agents respectively. The multi-strategy of each agent

k must have the format $T_k \subset \{(t,t) \mid t \in I\}$ otherwise, we cannot guarantee the threshold $\frac{1}{2}\sum_{t\in I} t$ if we pick any strategy that give the utility $-2n^2$. For each agent k, we define $B_k = \{t \mid (t,t) \in T_k\}$ and show that B_1 and B_2 form a partition. The proof is followed by the fact that the multi-strategy profile that is guaranteed the threshold $\frac{1}{2}\sum_{t\in I} t$ is a legal strategy profile which means that each $i \in I$, we can conclude either $i \in P_1$ or $i \in P_2$. The guaranteed payoff $\frac{1}{2}\sum_{t\in I} t$ proves the second condition that we need for a partition:

$$\frac{1}{2}\sum_{t\in I} t = \sum_{t_1\in P_1} t_1 = \sum_{t_2\in P_2} t_2$$

□

6.1.2 Strategic Arrangement Problem

It is natural to consider a situation in which the agents can act strategically. In the strategic fair division problem that we define, an agent's payoff depends on the strategies chosen by the other agents. Specifically, each agent will express the value of each item and will also express the regret of losing items for each round (see Section 5.1.5).

By solving the strategic arrangement problem, and considering **the supports** of the players in a constrained Nash equilibrium with a minimum expected payoff for each agent, we generate a potential fair allocation of the items. For convenience, to reflect more commonly used game-theoretic notation, we write $u_k(i,j)$ synonymously for $U(k,i,j)$. In this setting, we also may use the term *players* for agents.

Definition 6.1.5. *A* strategic arrangement problem $\mathcal{A}_\epsilon$ *is given by a tuple of the form* $\langle K, I, J, U, \theta, \tau, r\rangle$ *where* $K = [k]$ *is the set of* players*,* $I = [n]$ *is the set of* items *and* $J = [m]$ *is the set of* places*. A pure strategy is a pair* (i,j) *such that* $i \in I$ *and* $j \in J$*. A* mixed strategy *is a distribution over pure strategies, and a* (mixed) strategy profile *is a vector of (mixed) strategies for all players. For each player* k*, we are also given the utility matrix* u_k *such that* $u_k(i,j,\boldsymbol{x}_{-k})$ *gives the payoff of putting the item* $i \in I$ *in place* $j \in J$ *assuming all other players play* $\boldsymbol{x}_{-k}$*. A desirable output of this problem is a strategy profile* $\boldsymbol{x}^*$ *such that all the following conditions are satisfied:*

- *For each player* k*, we have* $u_k(\boldsymbol{x}^*) + \epsilon \geq u_k(s, \boldsymbol{x}^*_{-k})$ *for all* $s = (i,j)$ *where* $i \in I$ *and* $j \in J$ *(A Nash equilibrium).*

- *If $(i, j) \in (I, J)$ is assigned with a positive probability in one player's strategy, then for $\forall i' \in I$ and $\forall j' \in J$, (i, j') and (i', j) cannot be assigned with a positive probability in any other player's mixed strategy. In addition, for each item $i \in I_k$, there can exist at most one place $j \in J_k$ such that $(i, j) \in T_k$ (Each item and place can be used once by only one player).*

- *For all players, each pure strategy (i, j) can be assigned with a probability greater than θ or zero.*

- *All players achieve an expected payoff of at least $r - \epsilon$.*

- *Each player must pick at least τ items.*

A strategy profile is called legal *if it satisfies all of the specified conditions.*

Remark 6.1.6. The simple arrangement problem aims to find a legal multi-strategy profile in which all players get at least the threshold payoff that was defined. In contrast to the simple arrangement problem, this problem aims to find an equilibrium (a legal mixed strategy profile) that achieves an expected guaranteed payoff for all players. We also need to have the (general) assumption that $|K|$ is a constant in the strategic arrangement problem.

Remark 6.1.7. We may consider the utility matrices being defined as follows:

$$\forall k \in K \text{ and } \forall t \in I,\ \forall j \in J \quad u_k(t, j, \mathbf{x}_{-i}) = \alpha V_k(t) + (1 - \alpha) R_k(\mathbf{x}_{-i})$$

$V_i(t)$ indicates the value of item t for player i and $R_i(\mathbf{x}_{-i})$ is a function that is computable in constant time that shows how a player can interpret the allocation of items $(t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_k)$ in different rounds $(r_1, \ldots, r_{i-1}, r_{i+1}, \ldots, r_k)$ to the other players. Note that $x_{-i} = ((t_1, r_1), \ldots, (t_{i-1}, r_{i-1}), (t_{i+1}, r_{i+1}), (t_k, r_k))$ and α is a constant.

Remark 6.1.8. Indeed, our model could reflect the strategic behavior of the agents since the agents include how they value the items and their regret of losing the items in different rounds of the allocation.

Next, we will investigate the computational complexity of this problem. We prove the hardness of the strategic arrangement problem by extending the results of the previous chapter. There, the hardness of the problem of finding an approximate

Nash equilibrium with the constraints that we had in Definition 6.1.5 for two-player games is demonstrated separately. We need to show that hardness holds while all of the provided constraints are in place, in addition to more complex strategies that have the format of $|I| \times |J|$ (see the definition and the proof of NP-completeness of Partition Nash equilibrium in the previous chapter).

Theorem 6.1.9. *The problem of deciding whether there exists a legal strategy profile for $\mathcal{A}_\epsilon(K, I, J, U, \frac{1}{\theta}, \tau, r)$ is NP-complete.*

Proof. An analysis similar to Theorem 5.2.1 Part (b) shows that this game has an approximate partition Nash equilibrium with the following properties iff the given formula ϕ is satisfiable:

- A guaranteed payoff of $n - 1 - \epsilon$.
- For each variable v and its associated literals $-l_v$ and l_v, one of the literals is played with a probability greater than $\frac{1}{n} - 2\epsilon - \frac{1}{n^2}$ while the other is played with a probability smaller than $\frac{1}{n^2}$[1].
- All other strategies will be played with a small probability $\frac{\epsilon}{cn}$ by exactly one player. Note that c is just a constant used to generate a game in Proposition 5.4.11.

We define the strategy set of the game $\mathcal{A}(\phi, \epsilon)$ to be $S^{\mathcal{A}} = S_1^{\mathcal{A}} = S_2^{\mathcal{A}} = S \times S$. The utility matrices of this game can be generated by using a similar idea (copying the literals and associating the first copy with the first player, and the second copies will be associated with the second player by penalizing the players) compared to Proposition 6.1.3 as follows:

$$u_1^{\mathcal{A}}((i,j),(s,t)) = \begin{cases} u_1(i,s) & j \neq t \\ -2cn^2 & \text{o.w} \end{cases}$$

This definition will prevent the possibility of assigning probabilities smaller than $\frac{\epsilon}{2cn^2}$ (by an argument similar to Lemma 5.3.4 and Lemma B.4.3.). Assuming $(x_1, \ldots, x_n)$ satisfies the formula ϕ with literals $l_1, \ldots, l_n$. For the row player, assign the probability $\frac{1}{n}$ to each of the strategies $((l_1^1, 1), \ldots (l_i^1, 2i+1), \ldots, (l_n^1, 2n-1)$. The column player will assign the probabilities of $((l_2^2, 2), \ldots (l_2^i, 4) \ldots (l_2^n, 2n))$ to be $\frac{1}{n}$, except that $(l_1^2, 2)$ is assigned to have a probability of $\frac{1}{n} - \frac{\epsilon}{2cn}$.

[1]See Lemma 5.3.9 which is similar to Theorem 8.6 in [169].

The column player will distribute the weight $\frac{\epsilon}{cn}$ equally to the remaining strategies to generate, and we can generate a legal strategy profile as it satisfies the conditions of a partition Nash equilibrium. More specifically, for any remaining place (round) r ($|S| \geq r > 2n$), the column player will pick the remaining strategies other than pure strategies in $L_1 \cup L_2$ with a probability of $\frac{\epsilon}{2cn(|S|-2n)}$. This shows that this game does have a legal strategy profile with a guaranteed payoff of $r = n - 1 - \epsilon$, $\frac{1}{\theta} = \frac{\epsilon}{2cn(|S|-2n)}$ (greater than $\frac{\epsilon}{2cn^2}$) where for each place (round) one item in $|S|$ is selected.

If the formula is unsatisfiable, at least one of the players cannot make a guaranteed payoff $n - 1 - \epsilon$ without violating the disjointness condition, by a proof similar to 5.3.2 in the previous chapter. □

Remark 6.1.10. It is easy to see that all strategies other than the ones that are related to the literals can be distributed arbitrarily between the players. It is also easy to see that the support of players will have n elements if the formula is satisfiable.

6.2 Simple Analysis for Random Matrices

To find a constrained Nash equilibrium, we can use well-known algorithms that generate Nash equilibria that search through the space of potential strategy profiles. We provide a simple analysis, indicating that, at least for random utility matrices, approximation algorithms that are designed based on considering random potential strategies may not guarantee a good payoff (see also [56, 162]). In particular, in this case, such algorithms are likely to find a Nash equilibrium, but are unlikely to find an equilibrium with disjoint supports that has any expected guaranteed payoff greater than $\frac{1}{2}$ (conditions 2 and 4 in Definition 6.1.5).

Proposition 6.2.1. *In a game with random utility matrices, all mixed strategy profiles will form a Nash equilibrium regardless of what the other player does. In other words, both players are indifferent to all mixed strategies.*

Proof. We show that the expected payoff for any strategy (x, y) pair, where $x = (x_1, \cdots, x_n)$ and $y = (y_1, \cdots, y_n)$ are fixed mixed strategies, is always $\frac{1}{2}$. The expected payoff for the row player is calculated as $\sum_{i=1}^{n} \sum_{j=1}^{n} x_i.y_j u_1(i, j)$. Since the expected expectation is a just a number ($E[E[x]] = E[x]$), we can conclude the following:

$$E[u_1(x, y)] = E[u_1((x_1, \cdots, x_n), (x_1, \cdots, y_n))]$$

$$= E[\sum_{i=1}^{n}\sum_{j=1}^{n} x_i.y_j u_1(i,j)]$$

$$= \sum_{i=1}^{n}\sum_{j=1}^{n} x_i.y_j E[u_1(i,j)] = \sum_{i=1}^{n} x_i \sum_{j=1}^{n} y_j E[u_1(i,j)]$$

We know that for each i and j, $E[u_1(i,j)] = \frac{1}{2}$.

$$E[u_1(x,y)] = \sum_{i=1}^{n} x_i \sum_{j=1}^{n} y_j E[u_1(i,j)] = \frac{1}{2}\sum_{i=1}^{n} x_i.(\sum_{j=1}^{n} y_j) = \frac{1}{2}$$

The argument holds for the column player symmetrically. □

The following proposition implies that it is unlikely to find a Nash equilibrium with our second constraint (Definition 6.1.5) by using randomization over possible mixed strategies, similar to the algorithms described in [56, 76].

Proposition 6.2.2. *In a two-player game that has random matrices, the probability that a player does not use any arbitrary pure strategy in the support is zero.*

Proof. There is a mechanism for generating points uniformly on simplex $(x_1, \cdots, x_n)$ where $\sum_{i=1}^{n} x_i = 1$ by following this procedure:

Suppose $X_1, \ldots, X_n$ are i.i.d exponential random variables. Next, we can conclude $(X_1/X, \ldots, X_n/X)$ is uniformly distributed on the simplex where $X = X_1 + \cdots + X_n$.

We need to show that the outcome $\frac{X_i}{X} = 0$ has a low probability to happen. We do so by using the expected value. Suppose x_i is the i-th order statistic from i.i.d exponential sample.

By using Rényi representation, we can construct x_i based on a recursive equation: We initialize it by using an additional variable $x_0 = 0$, and let $\left(\widetilde{X}_i\right)_{i\geq 1}$ be a sequence of i.i.d exponentials. Then, define:

$$x_i = x_{i-1} + \frac{\widetilde{X}_i}{n-i+1}$$

$$\Pr(x_i > 0) = Pr(x_{i-1} + \frac{\widetilde{X}_i}{n-i+1} > 0)$$

$$\geq \Pr(\frac{\widetilde{X}_1}{n-1+1} > 0) = \Pr(\widetilde{X}_1 > 0) = 1 - \Pr(\widetilde{X}_1 = 0) = 1$$

□

6.3 A Quasi-Polynomial Time Algorithm

We use the following quasi-polynomial time algorithm [130] that is close to optimal running time for finding a Nash equilibrium in normal-form games with a constant number of players. Note that we assume the exponential time hypothesis (ETH) for the (PPAD-complete) end-of-the-line problem (also called PETH) [166].

Definition 6.3.1. *Suppose that all players choose their mixed strategies from S. Assume that S' is a* ***multi-set*** *whose elements are from S (it has a number of copies of some elements of S). A mixed strategy is called* t-uniform *if it is the uniform distribution on a multi-set S' of pure strategies, with $|S'| = t$.*

Proposition 6.3.2 ([130])**.** *For any two-player game $\mathcal{G}$, for any real number ϵ between 0 and 1 , for every $t \geq \frac{12 \ln n}{\epsilon^2}$, there exists a pair of t-uniform strategies x', y', such that:*

1. *x', y' is an ϵ-Nash equilibrium for $\mathcal{G}$.*
2. *$|u_1^{\mathcal{G}}(x', y') - u_1^{\mathcal{G}}(x^*, y^*)| \leq \epsilon'$, (row player gets almost the same payoff as in the Nash equilibrium (x^*, y^*))*
3. *$|u_2^{\mathcal{G}}(x', y') - u_1^{\mathcal{G}}(x^*, y^*)| \leq \epsilon'$, (column player gets almost the same payoff as in the Nash equilibrium)*

The proof is available in [130], which is based on the probabilistic method.

Proposition 6.3.3 ([130])**.** *For a k-player game where k is a constant, for a fixed t, there exists a quasi-polynomial algorithm for computing all t-uniform ϵ-Nash equilibria.*

Proof. For a two-player game, it is easy to verify ϵ-Nash equilibrium conditions as we need to check only for deviations to pure strategies. Given an $\epsilon > 0$, fix $t = \frac{12 \ln n}{\epsilon^2}$. Using brute-force search, we can compute all t-uniform ϵ-Nash equilibria (at least one such equilibrium exists). The running time of the algorithm is quasi-polynomial since there are $C(n+t-1, t)^2$ possible pairs of multi-sets to look at. The proof is similar for any constant k with a minor difference that $t = \frac{3k^2 \ln k^2 n}{\epsilon^2}$ and checking $C(n+t-1, t)^k$ strategies (see [130]).

∎

Following the introduction of the strategic arrangement problem and the examination of several computational barriers and techniques related to general game theory, we initially begin with the following algorithm (Algorithm 2) to solve a two-player arrangement problem. This algorithm can be simply expanded to any k-player arrangement problem as long as $\frac{1}{\theta} \geq \frac{3k^2 \ln(k^2|I|\cdot|J|)}{\epsilon^2}$[2].

ALGORITHM 2: A quasi-polynomial algorithm for approximating a legal strategy of an arrangement problem

Input: $\mathcal{A}_\epsilon\langle\{1,2\}, I, J, U, \tau, r\rangle$
Calculate $\frac{1}{\theta} = \frac{12 \ln |I|\cdot|J|}{\epsilon^2}$ based on ϵ;
For any θ-uniform strategy profile that can be generated from $|I| \times |J|$: **if** *Check if the strategy is a ϵ-Nash equilibrium* **then**
  if *This strategy profile is legal and provides the minimum expected payoff guarantee r to all players* **then**
    return the strategy profile;
  end
  else
    continue the search;
  end
end

6.4 Using the Bus Transportation Algorithm

Algorithm 2 has three obvious drawbacks for finding a legal strategy profile for the arrangement problem. First, the worst-case running time is not ideal since it is a quasi-polynomial algorithm. Second, it may not find a legal strategy profile since it only searches θ-uniform strategies. Third, ϵ and θ depend on one another.

We use the bus transportation algorithm (BTA), a meta-heuristic which attempts to find a legal strategy profile that maximizes the social welfare of all players. First, we need to review some essential background of the bus transportation algorithm. Next, we show that we can transform the search process of finding a legal strategy in a strategic arrangement to a binary quadratically constrained quadratic program (QCQP) and apply an algorithm similar to the bus transportation algorithm for integer programming to the transformed problem. The formulation became complex

[2] Indeed, θ depends on ϵ. This is due to the fact that we need to find a Nash equilibrium that also satisfies the third constraint of the strategic arrangement problem and the algorithm of [130] only considers t-uniform strategies.

due to the fact that the strategies must have certain properties in contrast to simpler methods such as Lemke-Howson [127] and support enumeration [174, 184].

6.4.1 Overview of General Bus Transportation algorithm

The bus transportation algorithm (BTA) is a meta-heuristic inspired by the collective civic behavior of humans. The algorithm uses *humans*, *buses*, and *stations* as components to reflect a situation in which the purpose of using buses is to help people reach their destinations. The buses will take some passengers, who have limited knowledge of an ideal station, and will help them reach their destination based on their preferences and their prior knowledge. Figure 6.1 illustrates an overview of this meta-heuristic algorithm.

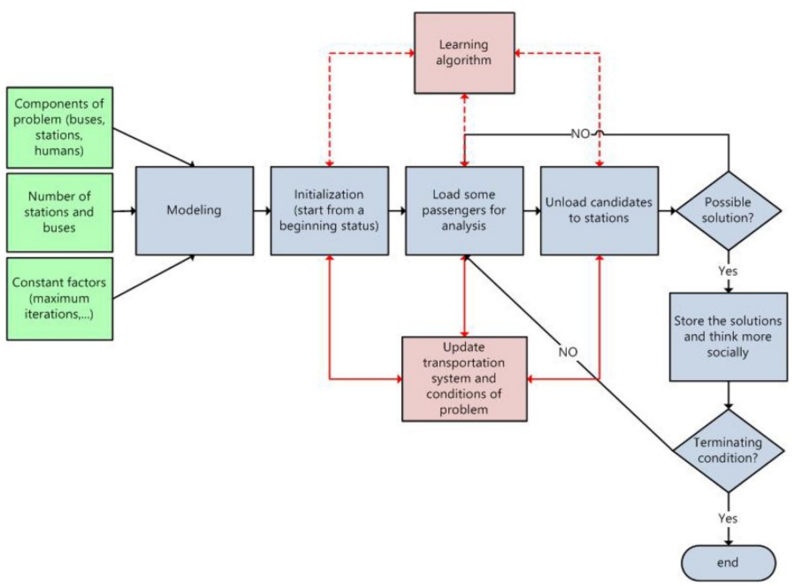

Figure 6.1: General Bus Transportation Algorithm

6.4.2 Overview of BTA for Integer Programming

We now describe a brief review of BTA for integer programming, which is presented in [34], where the learning algorithm component that was used is simple human learning optimization (SHLO) [182]. The key concept in SHLO is that humans use their individual and social experiences in order to reach their goals, while sometimes they exhibit random behavior. BTA for integer programming also adds more elements, which are supervision (the bus drivers), along with some other features of other meta-heuristic algorithms, such as tabu search and IMPRO [14, 90, 91]. Accordingly, the bus driver will try to find a simple pattern to improve the answer and try to avoid solutions that have already been examined. In other words, the bus driver will not take the passengers to stations that are not ideal.

For this specific problem (binary integer programming), the variables will be the passengers while the bus driver controls the passengers going to the right station (taking the right value, either zero or one) [34]. In the paper [34], a variety of stations were considered where the passengers can be placed either in short-term (STS), mid-term (MTS), or long-term stations (LTS). Each passenger begins with a random (zero or one) short-term station, and one of the buses takes them to other stations. The passengers follows a merit function F (which is a smarter version of IKD in [182]), a social experience matrix (SKD in [182]), and announces their preferred station (zero or one station) to the bus driver. The bus driver takes this passenger to the preferred station (zero or one) with the correct type of station (short, mid, or long term). The type of station that the passengers are sent to depends on their previous history. Figure 6.2 shows the possible transition between the stations for the passengers. The final goal is that all passengers eventually reach their ideal long-term station while the optimization goals of the problem (constraints) are satisfied. If most of the passengers reach a long-term station, it means that this configuration is likely to form a global optimum. To avoid local optimum answers, a chance is given to the passengers to go back to mid-term stations with a lower probability. We may consider constant stations (CS) (only used in [34]) for passengers whose stations can be determined based on the structure of the optimization problem easily. According to [34], BTA outperforms well-known meta-heuristic algorithms in binary integer programming problems.

6.4.3 Overview of BTA for The Arrangement Problem

Given a finite set of players K, action index sets I and J, utility functions $\{U_k\}_{k \in K}$, and parameters τ, r, and θ, the goal is to find a legal strategy profile $x = (x_1, \ldots, x_{|K|})$ satisfying the arrangement constraints (Definition 6.1.5) that maximizes the total social welfare $\sum_{k \in K} U_k(x)$. Each mixed strategy x_k is represented by $\{x_k^{(i,j)}\}_{i \in I,\, j \in J}$, with

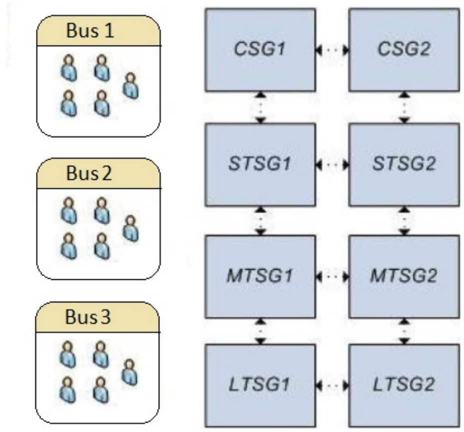

Figure 6.2: A transportation system with three buses and four types of stations

a discretized binary encoding. We encode this optimization problem as a quadratically constrained quadratic program (QCQP) as follows:

$$
\begin{aligned}
&\max z = \sum_{k=1}^{|K|} U_k(x_1, \ldots, x_k, \ldots, x_{|K|}) \\
&\quad U_k(x_1, \ldots, x_k, \ldots, x_{|K|}) > U_k(x_1, \ldots, s_k, \ldots, x_{|K|}), \quad k \leq |K|,\ s_k \in |I||J| \\
&\quad U_k(x_1, \ldots, x_k, \ldots, x_{|K|}) > r, \quad k \leq |K| \\
&\quad x_k = \sum_{i \leq |I|} \sum_{j \leq |J|} x_k^{(i,j)} = 1, \quad k \leq |K| \\
&\quad x_k^{(i,j)} = \sum_{u \leq \frac{1}{\theta}} y_k^{(i,j,u)} \cdot \theta, \quad k \leq |K|,\ i \leq |I|,\ j \leq |J| \\
&\quad y_k^{(i,j,u)} \in \{0,1\}, \quad k \leq |K|,\ j \leq |J|,\ u \leq \frac{1}{\theta} \\
&\quad x_{k_1}^{(i,j_1)} x_{k_2}^{(i,j_2)} = 0, \quad k_1, k_2 \leq |K|,\ i \leq |I|,\ j_1, j_2 \leq |J| \\
&\quad \sum_{k \leq |K|} \sum_{i \leq |I|} x_k^{(i,j)} > 0, \quad j \leq |J| \\
&\quad x_{k_1}^{(i_1,j)} x_{k_2}^{(i_2,j)} = 0, \quad k_1, k_2 \leq |K|,\ j \leq |J|,\ i_1, i_2 \leq |I| \\
&\quad \mathrm{Supp}(x_k) > \tau, \quad k \leq |K|
\end{aligned}
\tag{6.1}
$$

Remark 6.4.1. Here, the constraint $\mathrm{Supp}(x_k) > \tau$ is used as shorthand for an exact algebraic encoding of a minimum support-size requirement and can be implemented via auxiliary binary variables and quadratic constraints linking support indicators to mixed-strategy probabilities.

Note that it is easy to see that the support of each player can be extracted based on the mixed strategies of the players. We again use SHLO to help us determine how we can assign the variables similarly. Note that here we use IKD similarly to the original paper [182], which means that we do not use a merit function compared to [34]. Some basic ideas of Tabu search and IMPRO that were used in [34] are also considered here. We consider the passengers to be the items, and stations will be considered to be the places in the arrangement problem. Unlike the algorithm described in [34], here multiple buses will be involved to solve an arrangement problem (the transformed QCQP version). The midterm and long-term stations will only have a capacity of 1 while a short-term station can have all of the items (passengers). In each iteration, each driver moves some passengers (items) to some stations (places) according to the passengers' preferences (based on SHLO). We can assume that the stations (rounds) are divided among the players equally ($\tau = \frac{|I|}{c}$ since here $|I| = |J|$) and the drivers will only pick the passengers (items) if they desire to go to one of their associated stations. We now proceed to Algorithm 3, which describes the whole procedure.

ALGORITHM 3: BTA for the arrangement problem

Input: $\mathcal{A}_\epsilon\langle K, I, J, U, \theta, \tau, r\rangle$ and $maxBTA$
Initialize ϵ, and the stations;
For each driver b, pick some passengers from STS, MTS and LTS stations;
Ask for a desirable destination from each passenger;
Move the passenger to either STS, MTS or LTS based on the previous history (See 6.2);
If the destination station (mid or long-term only) is occupied, move the occupant passenger to their desirable STS station;
Form the support of each player based on the transitions that were done by each driver;
Try to Generate a θ-strategy profile that satisfies Equation 6.1 from the support;
if *A solution is found within* $maxBTA$ *rounds* **then**
    Return the strategy profile;
end
else
    Terminate the program;
end

6.5 Numerical Results

As stated, a general assumption in game theory considers games with a constant number of players since the payoff matrices will have exponential space usage. For simplicity, we only consider 2-player games similar to all present works. The performance is assessed based on a different number of items, and two different distributions[3]. The first distribution (D1) is similar to a uniform distribution. The second distribution (D2) is a biased distribution in which even the existence of a unconstrained Nash equilibrium is unlikely. Table 6.1 indicates the best, worst, and average running times of a strategic item allocation problem transformed to a strategic arrangement problem. We assume $\epsilon = 0.01$, maxBTA $= 50$, and with a guaranteed expected payoff of 0.4 (assuming all payoffs are scaled in $[0, 1]$) to all players. The practical results confirm the fact that if we take any support strategy profile that satisfies the conditions of the strategic arrangement problem, finding a solution is easy for utility matrices that have a distribution similar to a uniform distribution. It is not surprising to see that for utility matrices in which some items have strictly better values, the success rate of each iteration of BTA is low. Furthermore, in this case, the overall running time will be close to exponential as we need to iterate through all possible strategy profiles.

Table 6.1: 30 different independent runs of BTA with different sizes of I

Specifications	Best	Worst	Average	(Success,Failure)	Average Iterations
$\lvert I \rvert$=6, D1	0.0013	0.024	0.007	50,0	3
$\lvert I \rvert$=6, D2	0.0016	0.031	0.008	50,0	3.600
$\lvert I \rvert$=12, D1	0.014	0.057	0.017	50,0	1.667
$\lvert I \rvert$=12, D2	0.027	2.748	1.027	47,3	37.933
$\lvert I \rvert$=30, D1	0.5688	1.223	0.6180	50,0	1
$\lvert I \rvert$=30, D2	14.180	62.739	56.251	2,48	48.100
$\lvert I \rvert$=60, D1	10.750	11.203	10.900	50,0	1

The first distribution is generated by simply considering the valuation of each item i for player k ($V_k(i)$) to be a random number in $linspace(0.1, 0.9, 9)$. The second distribution is generated by considering the valuation of each item i for player k ($V_k(i)$) to be a random number in the following set:

$$S = Good \cup Mid \cup Bad$$

$$Bad = \{0.01, 0.04, 0.005, 0.06, 0.09, 0.01, 0.012, 0.020.026, 0.03, 0.05\}$$

[3]The version of MATLAB was R2022b. The system that the test was done on had Core i7 10870H with no turbo boost with a single thread.

$$Mid = \{0.06, 0.12, 0.13, 0.15, 0.2, 0.21\}$$

$$Good = \{0.22, 0.230.25, 0.27, 0.3, 0.32, 0.44, 0.8\}$$

The final utility is generated as follows:

$$U_k(i, j, i', j') = V_k(i) - V_k(i')/4$$

Chapter 7

Conclusion and Open Problems

In this thesis, we initiated a systematic study of the computational aspects of approximating several general equilibrium problems in game theory. In particular, we established PPAD-completeness of variational inequalities as a unifying framework for modeling equilibrium problems that require certain assumptions, such as convexity of the feasible domain, for which we perform optimization. We also extended our framework to equilibrium problems under uncertainty, showing that finding an approximation solution to the robust counterparts of classical equilibrium models is PPAD-complete. Furthermore, we studied several general equilibrium problems that do not have convexity assumptions in two general settings. We also consider a problem that is a combination of several equilibrium problems (strategic arrangement problem), which can introduce an alternative fairness criterion, and attempt to solve it by employing human-inspired metaheuristics. In Section 7.1, we summarize the known results on the computational complexity of equilibrium problems under convex and non-convex constraints. Below, we suggest some possible research directions:

We have shown that our definitions of variational inequalities are general enough for various problems in game theory. One interesting direction could be investigating other game theoretic problems, as several other papers study properties of the Nash equilibrium in network games by using variational inequalities [136, 142, 154]. The relevance of our work in variational inequalities is further reflected in subsequent research that cites it, including studies that extend its ideas beyond the traditional boundaries of game theory [5, 92, 187, 188]. For example, [187] introduces a new relaxation of the classical variational inequality (VI) framework—called Expected Variational Inequalities (EVIs). These have a close connection to coarse correlated equilibrium problem [11, 140] and maximizing a (single) function that satisfies quasarconcavity. Instead of finding a single point satisfying the VI constraints, EVIs look for a probability distribution over feasible solutions that satisfies the VI condition in expectation. Extension of their work to include variations such as EQVI and EGQVI is still an open problem. Also, we have not considered the exact version of variational inequality problems, but we conjecture that it should be possible to prove FIXP-completeness for appropriate formulations (see [78]).

Our findings and insights for games with uncertainty could be of use in deriving computational complexity results for various learning problems (see also [55, 189]). Furthermore, an interesting extension of our work could be investigating games with smart contracts and their applications to permissionless blockchains such as Ethereum [94]. Also, deriving computational complexity results for problems we studied, such as resilient Nash equilibrium under different and less restrictive assumptions, could be of potential interest.

Given the fact that constrained equilibria are provably harder than their unconstrained versions, an open question is whether natural variants of equilibrium lie in TFNP but outside PPAD[1]. For the strategic arrangement problem, we may consider less restrictive concepts such as generalized Nash equilibrium (social Nash) along with different game formats such as graphical games [124], where an equilibrium is easier to find.

[1] This question was answered in [87]. However, the question of whether there exists another variation that goes beyond CLS is still interesting

7.1 Summarization of Results

Tables 7.1 and 7.2 summarize the computational complexity of various equilibrium problems under convex and non-convex constraints, respectively. For convex constraints (Table 7.1), the results are organized into exact computation, approximate computation, and computation under uncertainty (approximate robust equilibrium). For non-convex constraints (Table 7.2), the categories are constant-factor approximation (see also [64]) and polynomial-time approximation. Several of the approximate results, including PPAD-hardness, follow directly from the known hardness of approximating Nash equilibria [45, 59]. For the exact computation, the problems are all $FIXP$-hard for more than two players, and we speculate that most problems are either FIXP-complete [70, 78] or $\exists\mathbb{R}$-complete [82] (except non-remedial L/F-equilibrium). The complexity of these problems remains open in general. In the tables, "H" denotes hardness and "C" denotes completeness.

Table 7.1: Complexity of Equilibrium Problems under Convex Constraints

Problem (more than 2 players)	**Exact**	**Approx**	**+Uncertainty**
Normal Form Games (mixed Nash)	FIXP-C [70]	PPAD-C [45, 59]	PPAD-C
General Convex Games (Generalized Nash)	FIXP-C [78]	PPAD-C	PPAD-C
Generalized Close Equilibrium	FIXP-C	PPAD-C	PPAD-C
Total Resilient Nash	FIXP-H	PPAD-C	PPAD-C
Remedial L/F-Equilibrium	FIXP-H	PPAD-C	PPAD-C
(Non-Remedial) L/F-Equilibrium	$\Sigma_2 P$-H	PPAD-H	PPAD-H

Table 7.2: Complexity of Equilibrium Problems under Non-Convex Constraints

Problem (more than 2 players)	**Const-Approx**	**Poly-Approx**
Disjoint Nash	PPAD-H	NP-C
Generalized Disjoint Equilibrium	P	P
Partition Nash Equilibrium	PPAD-H	NP-C
Far Nash Equilibrium	PPAD-H	NP-C
Close Nash Equilibrium	PPAD-H	NP-C
Generalized Far Equilibrium	PPAD-H	PPAD-H
Restricted Disjoint Equilibrium	PPAD-H	NP-C
Restricted Far Equilibrium	PPAD-H	NP-H
Resilient Nash Equilibrium	PPAD-H	NP-C

Appendix A

Auxiliary Proofs

This appendix collects shorter proofs and technical arguments that, while not central to the main exposition, support various results in the thesis. These proofs are included here for completeness and to avoid interrupting the flow of the main text. In particular, they cover auxiliary lemmas, straightforward derivations.

A.1 Some Properties of L2 Distance

We use some fundamental lemmas pertaining to the properties of the L2 distance.

Lemma A.1.1. *Suppose that f and g are γ_f and γ_g strongly concave functions where $f, g : X \to Y$. Then, $f + g$ is a $\gamma_f + \gamma_g$-strongly concave function.*

Proof. Suppose that $x, y \in X$. We have the following:

$$(f+g)(\lambda x + (1-\lambda)y) - ((\lambda(f+g)(x) - (1-\lambda)(f+g)(y))$$
$$\geq \frac{\lambda(\lambda-1)}{2}\gamma_f\|x-y\|_2^2 + \frac{\lambda(\lambda-1)}{2}\gamma_g\|x-y\|_2^2 = \frac{\lambda(\lambda-1)}{2}(\gamma_f+\gamma_g)\|x-y\|_2^2$$

□

Lemma A.1.2. *The function $f(x) = \gamma\|x\|_2^2$ is 2γ-strongly convex.*

Proof. The key lies in leveraging the fact that the $L2$ is arising from an inner product.

$$
\begin{aligned}
& f(\lambda x + (1-\lambda)y) - (\lambda f(x) - (1-\lambda) f(y)) \\
& = \gamma\lambda^2 \|x\|^2 + 2\gamma\lambda(1-\lambda)\langle x, y\rangle + \gamma(1-\lambda)^2\|y\|^2 \\
& - \gamma\lambda\|x\|^2 - \gamma(1-\lambda)\|y\|^2 \\
& = \gamma\lambda(\lambda-1)\|x\|^2 + \gamma(1-\lambda)(1-\lambda-1)\|y\|^2 s + 2\gamma\lambda(1-\lambda)\langle x, y\rangle \\
& = -\gamma\lambda(1-\lambda)\|x\|^2 + 2\gamma\lambda(1-\lambda)\langle x, y\rangle - \gamma\lambda(1-\lambda)\|y\|^2 \\
& = -\gamma\lambda(1-\lambda)\left(\|x\|^2 - 2\langle x, y\rangle + \|y\|^2\right) \\
& = -\gamma\lambda(1-\lambda)\|x-y\|^2 \\
& \leq -\frac{1}{2}\gamma\lambda(1-\lambda)\|x-y\|^2 \\
& \text{since } \gamma\lambda(1-\lambda)\|x-y\|^2 \geq 0.
\end{aligned}
$$

□

Corollary A.1.3. *The function $f(x) = -\gamma\|x\|_2^2$ is γ-strongly concave.*

A.2 On Lipschitz Continuity

The paper [152] proves that the relaxed algorithmic Lipschitzness parameters are reasonable for an $(\eta, \sqrt{m}, L)$ well-conditioned correspondence.

Lemma A.2.1. *Let $\mathcal{F} : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$ be an $(\eta, \sqrt{m}, L_{\mathcal{F}})$ well-conditioned correspondence, and two vectors $\mathbf{p}, \mathbf{q} \in \mathbb{R}^{m*}$, $w = \widetilde{\Pi}^{\epsilon}_{\mathcal{F}(q)}(q)$ and $z = \widetilde{\Pi}^{\epsilon}_{\mathcal{F}(p)}(w)$, then $\|z - w\| < L_{\mathcal{F}}\|p - q\| + 3(1 + \hat{c}_{d,\eta})\epsilon$ holds where $\hat{c}_{d,\eta}$ is the constant of Theorem 3.15 of [152].*

Remark A.2.2. For strong separation oracles, $\hat{c}_{d,\eta} = 0$ as the margin $\eta = 0$.

Remark A.2.3. Theorem 3.15 of [152] establishes the relationship between the exact projection problem and the approximate projection problem by bounding the error caused by using the approximate version instead of the exact projection.

For establishing the reduction from Kakutani's fixed point to a PPAD version of Sperner's lemma, the paper [152], ensures that Lipschitz's continuity of the given correspondence of Kakutani's fixed point is sufficient. Lemma C.6 in [152] also ensures why Lemma A.2.1 is sufficient for the violation of Lipschitz continuity by testing the requirement of Lemma A.2.1 on the given solution by Sperner's problem. For GQVI, similar to concave games in [152], it is sufficient to ensure we adapt robust Berge's maximum theorem. Any infeasible solution given by Sperner's problem (indirectly given to Kakutani's fixed point) can only be caused by the violation of emptiness or the violation of Lipschitzness (see page 31 in [152]). This is possible to detect by the given tools in [152], given the discussion here and throughout the research.

Lemma A.2.4. *Let $\mathcal{F}$ be an L-Hausdorff Lipschitz continuous and well-bounded correspondence. Additionally, let any two vectors $p, q \in \mathbb{R}^{m*}$ with distance at most $\|p - q\| \leq \xi$ and any positive constants ϵ, ϵ°, such that $\left\|\widehat{\Pi}^{\epsilon}_{F(\boldsymbol{p})}\left(\widehat{\Pi}^{\epsilon}_{F(\boldsymbol{q})}(\boldsymbol{q})\right) - \widehat{\Pi}^{\epsilon}_{F(\boldsymbol{q})}(\boldsymbol{q})\right\| \leq L\xi + \hat{\mathcal{L}}_{d,\eta} \cdot \epsilon + \hat{\mathcal{L}}^{\circ}_{d,\eta}\epsilon^\circ$ for some constants $\hat{\mathcal{L}}_{d,\eta}, \hat{\mathcal{L}}^{\circ}_{d,\eta}$. Then, it holds that :*

$$\left\|\widehat{\Pi}^{\epsilon}_{F(p)}(p) - \widehat{\Pi}^{\epsilon}_{F(q)}(q)\right\| \leq 2\sqrt[4]{d}\sqrt{2\,\mathrm{error}_{d,\eta,L}\left(\xi, \epsilon, \epsilon^\circ\right)} + \mathrm{error}_{d,\eta,L}\left(\xi, \epsilon, \epsilon^\circ\right)$$

where error $_{d,\eta,L}\left(\xi, \epsilon, \epsilon^\circ\right) = (L+1)\xi + \hat{\mathcal{L}}'_{d,\eta} \cdot \epsilon + \hat{\mathcal{L}}^{\circ}_{d,\eta}\epsilon^\circ$ *and* $\hat{\mathcal{L}}'_{d,\eta} = \hat{\mathcal{L}}_{d,\eta} + 2\hat{c}_{d,\eta}$

We restate the following lemma from (Lemma A.1 in [76]) states that any linear arithmetic circuit f mapping $\mathbb{R}^n$ to $\mathbb{R}^m$ is Lipschitz-continuous with respect to the ℓ_∞-norm. The Lipschitz constant for the circuit is bounded by $2^{\mathrm{size}(f)^2}$, where size(f) is the number of nodes in the circuit. This means that the output of the circuit changes at a controlled rate, depending on its size.

Lemma A.2.5 ([76])**.** *Any linear arithmetic circuit $F : \mathbb{R}^m \to \mathbb{R}^m$ is $2^{size(F)^2}$-Lipschitz-continuous (with respect to the ℓ_∞-norm) over $\mathbb{R}^m$.*

We cannot extend the result to linear arithmetic circuits that implement separation oracles for convex correspondences. This is due to the fact that for correspondence $\mathcal{F}$, a strong separation oracle only provides information about whether a point is inside or outside $\mathcal{F}(x)$, not about how $\mathcal{F}(x)$ changes continuously with respect to x.

Example A.2.6. Consider the correspondence $\mathcal{C} : \mathbb{R} \rightrightarrows \mathbb{R}$ defined in the following. The separation oracle for $\mathcal{C}(x)$ is piecewise linear, as the separating hyperplanes defining the sets are given by linear functions. Note that $[0, 1]$ and $[1, 2]$ are both bounded, compact, and convex. However, the correspondence $\mathcal{C}$ is not Lipschitz continuous with respect to the Hausdorff distance. Specifically, there is a discontinuity at

$x = 0$, where the Hausdorff distance between $\mathcal{C}(0) = [0, 1]$ and $\mathcal{C}(\epsilon) = [1, 2]$ is always 1 for any $\epsilon > 0$, regardless of how small ϵ is.

$$\mathcal{C}(x) = \begin{cases} [0, 1], & \text{if } x \leq 0, \\ [1, 2], & \text{if } x > 0. \end{cases}$$

A reasonable assumption that would prevent counterexamples could be *Inner Stability Condition* which means that there exists a function $g : X \to \mathbb{R}^m$ such that for all $x \in X$, the set $\mathcal{F}(x)$ can be written as $\mathcal{F}(x) = \{y \mid G(x, y) \leq 0\}$ for some Lipschitz-continuous function $G(x, y)$. This assumption ensures that the boundary of $\mathcal{F}(x)$ moves in a controlled way, ruling out sudden jumps. If $G(x, y)$ is Lipschitz continuous, then small changes in x lead to small changes in $\mathcal{F}(x)$, making $\mathcal{F}(x)$ Hausdorff-continuous with a controlled Lipschitz bound. In this situation, if we consider G to be a well-behaved function, we can effectively approximate it by linear arithmetic circuits, and Lipschitzness inherently holds using the restated lemma and the theorem from [76].

A.3 Other Variations of GQVI

We may impose more restricted convexity assumptions on our correspondences, such as $\mathcal{F}$ being strongly convex-valued. Note that for the correspondences in GQVI, we only have access to their separation oracles and tools, such as approximate projection problems. Strongly convex sets could be of interest in optimization and numerical stability because they ensure unique solutions, robust algorithms, and theoretical guarantees for stability. For the sake of completeness, we consider a variant considering strong convexity of $\mathcal{F}$. We define the Violation of (almost) μ-strong convexity as follows:

(Violation of μ-strong convexity): There vectors $x, p, q \in \mathbb{R}^{m*}$ and two constants $\epsilon > 0$ and $\lambda \in (0, 1)$ such that:

$$\left| \tilde{\Pi}^{\epsilon}_{\mathcal{F}(x)} \left(\lambda p + (1 - \lambda) q \right) - \left(\lambda p + (1 - \lambda) q - \frac{1}{2} \mu \lambda (1 - \lambda) \|p - q\|^2 \mathbb{1} \right) \right| \succ \epsilon \mathbb{1}$$

To justify the definition, we start with an equivalent convexity definition based on projection.

Lemma A.3.1. *Let C be a subset of $\mathbb{R}^m$. If C is convex, then, for any p, q and*

$\lambda \in (0,1)$, $\Pi_C(\lambda p + (1-\lambda)q) = \lambda p + (1-\lambda)q$.

Proof. The proof follows by the fact that, by definition, the set C is convex if for any $p, q \in C$ and $\lambda \in [0,1]$, the point $\lambda p + (1-\lambda)q$ lies in C and the (exact) projection maps a point inside of a set to itself. □

Next, we establish a relationship between strong convexity of sets and (almost) strong convexity in the following:

Proposition A.3.2. *Let $C \subset \mathbb{R}^m$ be closed, and let $\tilde{\Pi}^\epsilon_C$ be an approximate projection operator such that for all $z_0 \in \mathbb{R}^m$,*

$$|\tilde{\Pi}^\epsilon_C(z_0) - \Pi_C(z_0)| \preceq \epsilon \mathbb{1},$$

where $\Pi_C(z_0)$ is the exact Euclidean projection of z_0 onto C. Assume that for all $p, q \in C$ and $\lambda \in (0,1)$,

$$\left| \tilde{\Pi}^\epsilon_C(\lambda p + (1-\lambda)q) - \left(\lambda p + (1-\lambda)q - \frac{1}{2}\mu\lambda(1-\lambda)\|p-q\|^2 \mathbb{1} \right) \right| \preceq \epsilon \mathbb{1}. \qquad \text{(A.1)}$$

*Then C is μ-strongly convex, provided the following **limit/uniformity assumption** holds:*

$$\lim_{\epsilon \to 0} \tilde{\Pi}^\epsilon_C(z_0) = \Pi_C(z_0) \quad \textit{uniformly in } z_0.$$

Proof. Let $p, q \in C$ and $\lambda \in (0,1)$, and define

$$z := \lambda p + (1-\lambda)q - \frac{1}{2}\mu\lambda(1-\lambda)\|p-q\|^2 \mathbb{1}.$$

By assumption (A.1) we have

$$|\tilde{\Pi}^\epsilon_C(\lambda p + (1-\lambda)q) - z| \preceq \epsilon \mathbb{1}.$$

Now, take the limit $\epsilon \to 0$ and use the uniform convergence assumption:

$$\Pi_C(\lambda p + (1-\lambda)q) = \lim_{\epsilon \to 0} \tilde{\Pi}^\epsilon_C(\lambda p + (1-\lambda)q).$$

Therefore, we obtain

$$\Pi_C(\lambda p + (1-\lambda)q) = z.$$

Since the exact projection of any point onto a closed set lies in the set, we have

$$z = \Pi_C(\lambda p + (1-\lambda)q) \in C.$$

This holds for all $p, q \in C$ and $\lambda \in (0,1)$, which is precisely the definition of μ-strong convexity. □

Remark A.3.3. Note that, similar to [152], the violation of strong convexity is meaningful as output whenever the form of input is explicitly given; otherwise, strong convexity holds as a promise. The other direction of the above-mentioned proposition only holds under specific conditions, such as $\Pi_C(z_0) = z$.

A.4 General Game Theory Results

A.4.1 Rational Solutions for Nash Equilibrium in Bi-Matrix Games

In a bi-matrix game, the goal is to find a strategy profile (pair of mixed strategies) such that the Nash conditions are satisfied. If we already know the support of these strategies, determining the relative weight of the problem can be solved in polynomial time using linear programming. This immediately gives the following:

Proposition A.4.1. *Any* 2*-player game with rational payoffs has a rational Nash equilibrium where the probabilities are of a bit-length polynomial with respect to the number of strategies and bit-lengths of the payoffs.*

Remark A.4.2. The Proposition does not apply to more than two players [70].

A.4.2 Support Enumeration and Fully Mixed Strategies

We can also provide an equivalent definition for a Nash equilibrium in a bi-matrix game $\mathcal{G}(R, C)$ (R and C denote the payoff matrices for the row (first) and the column (second) player respectively) by doing the following[1]:

$$x^{*T} R y^* \geq e_i^T R y^*, \forall i \in \{1, \ldots, m\}$$
$$x^{*T} C y^* \geq x^{*T} C e_i, \forall i \in \{1, \ldots, n\}$$

[1]We use these notations to enhance clarity and align with the original research.

Remark A.4.3. For a bi-matrix game $\mathcal{G}$, we can represent U_1 and U_2 by matrices $R = (r_{ij})$ and $C = (c_{ij})$ respectively where $U_1(i,j) = r_{ij}$ and $U_2(i,j) = c_{ij}$. The matrix representation of such a game can be denoted $\mathcal{G}(R, C)$.

We can also define an ϵ-approximate Nash equilibrium by doing the following modification:

$$x^{*T}Ry^* + \epsilon \geq e_i^T Ry^*, \forall i \in \{1, \ldots, m\}$$
$$x^{*T}Cy^* + \epsilon \geq x^{*T}Ce_i, \forall i \in \{1, \ldots, n\}$$

Proposition A.4.4. *The mixed strategy* (x^*, y^*) *is a Nash equilibrium of* $\mathcal{G}(R, C)$ *iff these two conditions are satisfied:*

$$\forall i \in Supp(x^*), (Ry)_i = u = \max_{q \in S} \{(Ry)_q\}$$

$$\forall j \in Supp(y^*), \left(x^T C\right)_j = v = \max_{r \in S} \left\{\left(x^T C\right)_r\right\}$$

The first condition of this proposition states that a mixed strategy x^* of the row player is a best response to mixed strategy y^* of the column player if and only if all pure strategies i in the support of x^* are best responses to mixed strategy y^*. The second condition represents the best response condition corresponding to the column player.

Proposition A.4.5. *The problem of whether a game has a Nash equilibrium in which both players play a fully mixed strategy is in P.*

Proof. The following linear equations can check the case whether both players can have a Nash equilibrium with support size n. A candidate mixed strategy pair is determined by solving the equations:

$$\forall j \in [n], \sum_{i \in [n]} x_i c_{ij} = v,$$
$$\sum_{i \in [n]} x_i = 1$$
$$\forall i, \; x_i > 0$$

and

$$\forall j \in [n],\ \sum_{j \in [n]} y_j r_{ij} = u$$
$$\sum_{j \in [n]} y_j = 1$$
$$\forall i,\ y_i > 0$$

For the first player, the first set of equations tries to find a strategy x with support size n that makes the column player indifferent among playing the pure strategies in $[n]$. That means the column player obtains the same payoff, v, by playing the pure strategies in $[n]$ if the first player plays x. The second set of equations plays the same role for the column player. If the linear equations do not have a solution, then there is no Nash equilibrium where both players play a fully mixed strategy. The systems of equations can be solved using back-substitution. □

A.4.3 Diagonally Modified Games and Generalized Far Equilibrium (Proposition 5.4.3)

Proposition A.4.6 (Restatement of Proposition 5.4.3)**.** *For any game $\mathcal{G}$, the problem of finding a (EXACT) Nash equilibrium in the bi-matrix game $\mathcal{G}$ is poly-time reducible to the problem of finding a (EXACT) generalized disjoint equilibrium in $\mathcal{D}^M(\mathcal{G})$.*

Proof. This game can be modified to a diagonally modified game by using the construction that we provide in the following. Suppose this normal-form game $\mathcal{G}$ has payoff matrices U_1 and U_2 where both players' strategies come from $S = [n]$. In the following game $\mathcal{D}^{\mathrm{M}}(\mathcal{G})'$ that we generate (which is also a diagonally modified game), any Nash equilibrium of the original game $\mathcal{G}$, forms a disjoint Nash equilibrium (which is also a generalized disjoint equilibrium) in the generated game. Note that for any Nash equilibrium in this game, the row player can only assign positive probabilities on S_1, and this holds for the column player with S_2, respectively. The proof follows from Lemma 5.2.5.

$$U_1'(x,y) = \begin{cases} U_1^{\mathcal{G}}(i,j) & i \in S_1 \textit{ and } j \in S_2 \\ -M & \textit{otherwise} \end{cases}$$

$$U_2'(i,j) = \begin{cases} U_2^{\mathcal{G}}(i,j) & i \in S_1 \textit{ and } j \in S_2 \\ -M & \textit{otherwise} \end{cases}$$

□

A.4.4 Generalized Disjoint Equilibrium is in NP

Proposition A.4.7. *The problem of finding a generalized disjoint equilibrium in a bi-matrix is in NP.*

Proof. We use the equivalent matrix format to calculate the expected payoffs. Suppose that we have a generalized disjoint equilibrium $(x^*, y^*) \in \triangle_n \times \triangle_n$. We can check the following for the row (similarly for the column player):

$$\forall i \in \{1, \ldots, m\} \cap S - Supp(y^*), \quad x^{*T} R y^* \geq e_i^T R y^*$$

Suppose $x = (\alpha_{i_1}, \ldots, \alpha_{i_m})$ is an arbitrarily mixed strategy in $\mathcal{R}^n_{disjoint}(y^*)$ where we removed all strategies that are assigned with a zero probability. We show that any strategy in $\mathcal{R}^n_{disjoint}(x^*)$ could be generated as follows. We multiply α_i for each i to the previous inequality and calculate the summation as follows:

$$\sum_{i=1}^{m} \alpha_i x^{*T} R y^* \geq \sum_{i=1}^{m} \alpha_i e_i^T R y^*$$

Since x is a distribution, we have $\sum_{i=1}^{m} \alpha_i = 1$. And the right hand side will be $\sum_{i=1}^{m} \alpha_i e_i^T R y^* = x^T R y^*$. Finally, the result will be:

$$\forall x \in \mathcal{R}^n_{disjoint}(y^*),\ x^{*T} R y^* \geq x^T R y^*$$

And for the second player we can take the same approach:

$$\forall y \in \mathcal{R}^n_{disjoint}(x^*),\ x^{*T} C y^* \geq x^{*T} C y$$

□

Corollary A.4.8. *The problem of finding an approximate restricted disjoint equilibrium in a bi-matrix game is in NP.*

Proof. The restricted version of generalized disjoint has one additional condition, where checking the condition can be done in polynomial time. □

Appendix B

Additional Proofs

This appendix presents a selection of proofs that are either too long or too detailed for the main text. Some proofs may involve repetitive steps or technical arguments, and are included here to provide full rigor while keeping the main chapters concise.

B.1 A More General Form of the Robust Berge Maximum Theorem

We assume $f(x, y) = (f_1(x, y), \ldots, f_t(x, y))$. We can easily extend the robust form of Berg's maximum theorem to $f : A \times B \to \mathbb{R}^t$ for any t by some minor changes, as we now discuss[1]. We define the operator maxall to return all maximum values of all coordinates in vector-valued functions in a vector in the given format. Recall that argmaxall also returns a point in the domain that maximizes all coordinates of the vector-valued function.

Theorem B.1.1. *(Generalized Robust Berge Maximum Theorem). Let $A \subseteq \mathbb{R}^{n_1} \times \mathbb{R}^{n_2}$ and $B \subseteq \mathbb{R}^m$. Consider a continuous function $f : A \times B \to \mathbb{R}^t$ that is μ-strongly concave $\forall a \in A$, L-Lipschitz in $A \times B$ and also a L'-Hausdorff Lipschitz, non-empty, convex-set, compact-valued correspondence $g : A \rightrightarrows B$. Define $f^*(a) = \mathrm{maxall}_{b \in g(a)} f(a, b)$ and $g^*(a) = \mathrm{argmaxall}_{b \in g(a)} f(a, b)$. Then, we observe f^* is continuous and g^* is upper semi-continuous and single-valued, i.e., continuous. Furthermore, g^* is Lipschitz and $\left(L' + 2\sqrt{\frac{4}{\mu}}\sqrt{(L + L \cdot L')}\right) - (1/2, p)$-Holder continuous respectively (for sufficiently small differences)[2].*

[1] For MGQVI, we only need to prove the Lipschitzness of g^*.

[2] We only need the cases where $n_1 = m$, $n_2 = r \cdot m$ and $t = r$.

Proof. We will follow the proof of Theorem 3.20 in [152], leveraging the generalized Apollonius theorem. Note that if $f^*(a) = (f_1^*(a), \ldots, f_t^*(a)) = (f_1(a,b), \ldots, f_t(a,b))$, we can conclude that $(f_1(a,b), \ldots, f_t(a,b)) \geq (f_1(a,c), \ldots, f_t(a,c))$ for all $c \in g(a)$. In other words for all $1 \leq i \leq t$, $f_i(a,b) \geq f_i(a,c)$.

We first prove that g^* is a single-value function. For the sake of contradiction, assume that we have $b_1, b_2 \in g^*(a)$. By the definition of $f^*(a)$, we have the following:

$$f^*(a) = \text{maxall}\{f(a,b) : b \in g(a)\} = f\left(a, b_1\right) = f\left(a, b_2\right)$$

Let $b_k = b_1 \cdot k + (1-k) \cdot b_2$ for some $k \in (0,1)$. By the fact that g is a convex-valued correspondence, it also holds that $b_k \in g(a)$. Then, by concavity of f we have that:

$$\begin{aligned} f^*(a) \succeq f\left(a, b_k\right) = f\left(a, b_1 \cdot t + (1-t) - b_2\right) &\succeq f\left(a, b_1\right) \cdot k + (1-k) \cdot f\left(a, b_2\right) \\ &= f^*(a) \cdot k + (1-k) \cdot f^*(a) = f^*(a) \end{aligned}$$

So, by definition $b_k \in g^*(a)$. However, since f is strongly concave, it has a unique maximizer. This means that $g^*(a)$ is a single-valued correspondence. By Berge's original theorem, we know that g^* would be upper semi-continuous, which, for the single-valued case, corresponds to the classical notion of continuity.

Next, we will prove now that $f^*(a), g^*(a)$ will have some form of Lipschitz continuity. More precisely, for two arbitrary inputs $a_1, a_2 \in A$ it holds that:

$$\begin{cases} f^*\left(a_1\right) = \left\{\text{maxall}\, f\left(a_1, b\right) : b \in g\left(a_1\right)\right\} = f\left(g^*\left(a_1\right), a_1\right) \\ f^*\left(a_2\right) = \left\{\text{maxall}\, f\left(a_2, b\right) : b \in g\left(a_2\right)\right\} = f\left(g^*\left(a_2\right), a_2\right) \end{cases}$$

Furthermore, we have the following:

$$\begin{aligned} f^*\left(a_1\right) = \left\{\text{maxall}\, f\left(a_1, b\right) : b \in g\left(a_1\right)\right\} \succeq f\left(a_1, b\right)\ \forall b \in g\left(a_1\right) \\ f^*\left(a_1\right) \succeq f\left(a_1, \Pi_{g(a_1)}\left(g^*\left(a_2\right)\right)\right) \\ \succeq -L\left\|\left(\Pi_{g(a_1)}\left(g^*\left(a_2\right)\right), a_1\right) - \left(g^*\left(a_2\right), a_2\right)\right\| \mathbb{1} \\ + f^*\left(a_2\right) \end{aligned}$$

Note that $\Pi_{g(a_1)}\left(g^*\left(a_2\right)\right)$ denotes the exact projection of $g^*(a_2)$ to $g(a_1)$. The second inequality is implied by the fact that $f^*(a_1) = f\left(a_1, g^*\left(a_1\right)\right) \geq f(a_1, b)$ for all $b \in g(a_1)$. Finally, the last inequality follows from the Lipschitzness of f in all coordinates.

Next:

$$f^*(a_2) - f^*(a_1) \preceq L\|a_1 - a_2\|\mathbb{1} + L\left\|\Pi_{g(a_1)}(g^*(a_2)) - g^*(a_2)\right\|\mathbb{1}$$
$$f^*(a_2) - f^*(a_1) \preceq L\|a_1 - a_2\|\mathbb{1} + L\,\mathrm{d}_H(g(a_1), g(a_2))\mathbb{1}$$
$$f^*(a_2) - f^*(a_1) \preceq (L + L \cdot L')\|a_1 - a_2\|\mathbb{1}$$

Note that the second inequality again follows by the definition of Hausdorff distance and projection. Applying symmetrically the same argument for $f^*(a_2)$:

$$|f^*(a_1) - f^*(a_2)| \preceq (L + L \cdot L')\|a_1 - a_2\|\mathbb{1}$$

Or equivalently:

$$\|f^*(a_1) - f^*(a_2)\| \leq \sqrt{t} \cdot (L + L \cdot L')\|a_1 - a_2\|$$

For $g^*(a)$, in the first step, we will leverage the generalization of the Apollonius theorem[3]. More precisely, for a μ-strongly concave function f, we have the following relationship:

$$-f\left(a, \frac{x+y}{2}\right) \preceq -\frac{f(a,x) + f(a,y)}{2} - \frac{\mu}{8}\|x-y\|_2^2\mathbb{1} \quad \forall a \in A$$

or equivalently,

$$\frac{f(a,x) + f(a,y)}{2} \preceq f\left(a, \frac{x+y}{2}\right) - \frac{\mu}{8}\|x-y\|_2^2\mathbb{1} \quad \forall a \in A$$

Since $f_i\left(a, \frac{x+y}{2}\right) \leq \max\{f_i(a,x), f_i(a,y)\}$, we get that:

$$\|x - y\| \leq \sqrt{\frac{8}{\mu} \cdot \frac{\max\{f_i(a,x), f_i(a,y)\} - \min\{f_i(a,x), f_i(a,y)\}}{2}} \quad \forall a \in A \tag{B.1}$$

Or equivalently:

$$\|x - y\|\sqrt{t} \leq \sqrt{\frac{4}{\mu}}\|\sqrt{\mathrm{maxall}\{f(a,x), f(a,y)\} - \mathrm{minall}\{f(a,x), f(a,y)\}}\| \quad \forall a \in A \tag{B.2}$$

In the interest of clarity, we introduce extra variables similar to [152]. $K_1 = g^*(a_1)$,

[3] We assume that $f = (f_1, \ldots, f_m)$ and derive the bounds for each f_i (for $i \in [m]$) separately, similar to [152].

$K_2 = g^*(a_2)$ and $K_3 = \Pi_{g(a_1)}(K_2)$, we get the following:

$$\begin{aligned}\|g^*(a_1) - g^*(a_2)\| = \|K_1 - K_2\| &\leq \|K_1 - K_3\| + \|K_2 - K_3\| \\ &\leq \mathrm{d_H}(g(a_1), g(a_2)) + \|(K_1 - K_3)\|\end{aligned}$$

The second inequality again follows by the definition of Hausdorff distance and projection. Now we are ready to apply Equation B.1 knowing $f(a_1, K_1) > f(a_1, K_3)$ (see the definition of f^*).

$$\begin{aligned}&\|g^*(a_1) - g^*(a_2)\| \leq L'\|a_1 - a_2\| + \sqrt{\frac{4}{t\mu}}\|\sqrt{(f(a_1, K_1) - f(a_1, K_3))}\| \\ &\leq L'\|a_1 - a_2\| + \sqrt{\frac{4}{t\mu}}\|\sqrt{|f(a_1, K_1) - f(a_1, K_3) + f(a_2, K_2) - f(a_2, K_2)|}\| \\ &\leq L'\|a_1 - a_2\| + \sqrt{\frac{4}{t\mu}}\|\sqrt{\left\{\begin{array}{c} |f(a_1, K_1) - f(a_2, K_2)| \\ + \\ |f(a_2, K_2) - f(a_1, K_3)| \end{array}\right\}}\| \\ &\leq L'\|a_1 - a_2\| + \sqrt{\frac{4}{t\mu}}\|\left\{\begin{array}{c} \sqrt{|f^*(a_1) - f^*(a_2)|} \\ + \\ \sqrt{|f(a_2, K_2) - f(a_1, K_3)|} \end{array}\right\}\|\end{aligned}$$

Next, we apply the bound that we have for f^*:

$$
\begin{aligned}
\|g^*(a_1) - g^*(a_2)\| &\le L' \|a_1 - a_2\| + \sqrt{\frac{4}{t\mu}} \left\{ \begin{array}{c} \sqrt{\sqrt{t}\,(L + L \cdot L')\,\|a_1 - a_2\|} \\ + \\ \sqrt{L\,\|K_2 - K_3\| + L\|a_2 - a_1\|} \end{array} \right\} \\
&\le L' \|a_1 - a_2\| + \sqrt{\frac{4}{t\mu}} \left\{ \begin{array}{c} \sqrt{\sqrt{t}\,(L + L \cdot L')\,\|a_1 - a_2\|} \\ + \\ \sqrt{L\,\mathrm{d_H}\left(g(a_1), g(a_2)\right) + L\|a_2 - a_1\|} \end{array} \right\} \\
&\le L' \|a_1 - a_2\| + 2\sqrt{\frac{4}{\sqrt{t}\mu}} \sqrt{(L + L \cdot L')\,\|a_1 - a_2\|} \\
&\le \underbrace{L' + 2\sqrt{\frac{4}{\mu}}\sqrt{(L + L \cdot L')}}_{\kappa} \max\left\{ \|a_1 - a_2\|^{1/2}, \|a_1 - a_2\| \right\} \\
&\le \kappa \max\left\{ \|a_1 - a_2\|^{1/2}, \|a_1 - a_2\| \right\}
\end{aligned}
$$

□

Remark B.1.2. When we use $\sqrt{v}$ for a vector, we mean the square root of each component.

B.2 Equivalent Results for Weak Separation Oracle Variation

Weak separation oracles are useful for establishing general distinctions between problems and deriving lower bounds on computational complexity. They offer a broader perspective on the relative difficulty of problems. The choice between weak and strong separation oracles depends on the goals of the analysis and the level of detail required to draw meaningful conclusions about the computational landscape of specific problems. For some convex sets, we may not be able to provide polynomial-time strong separation oracles. Within this section, we only reproduce some of the equivalent results for weak separation oracles for completeness, maintaining a consistent methodology with Chapters 3 and 4. For example, establishing results for remedial robust Nash equilibrium in multi-leader-follower games is similar.

B.2.1 Essential Elements

In short, a weak separation oracle (WSO) for a correspondence is a circuit that takes as input a point and the accuracy of the separation oracle δ (compared to a strong separation oracle $\delta = 0$) and produces either an almost-membership or a guarantee of almost-separation.

A Weak Separation Oracle (via a circuit $C_{\mathcal{R}(x)}$)

Input: A vector $z \in \mathbb{Q}^m \cap \mathbb{R}^{m*}$ and a constant δ as inputs:

Output: $(a, b) \in \mathbb{Q}^m \times \mathbb{Q}$ such that the threshold $b \in [0, 1] \cap \mathbb{Q}$ denotes almost membership of z in $\mathcal{R}(x)$. More precisely:

- If $z \in \overline{\mathrm{B}}(\mathcal{R}(x), \delta)$ then $b > \frac{1}{2}$ and the vector a will be $\perp$. In other words, a is meaningful only when $b \leq \frac{1}{2}$
- $b \leq \frac{1}{2}$ and normalized vector a , with $\|a\|_\infty = 1$ which defines an almost separating hyperplane $\mathcal{H}(a, z) := \{y \in \mathbb{R}^{m*} : \langle a, y - z \rangle = 0\}$ between the vector z and the set $\mathcal{R}(x)$ such that $\langle a, y - z \rangle \leq -\delta$ for every $y \in \overline{\mathrm{B}}(\mathcal{R}(x), -\delta)$.

Similarly, we need the weak version of the constrained convex optimization problem.

Weak Constrained Convex Optimization Problem

Input: A zeroth and first order oracle for the convex function $F : \mathbb{R}^m \to \mathbb{R}$, two rational numbers $\delta, \epsilon > 0$ and a weak separation oracle $\mathrm{WSO}_{\mathcal{R}}$ for a non-empty closed convex set $\mathcal{R} \subseteq \mathbb{R}^{m*}$.

Output: One of the following cases:

- (Violation of non-emptiness): A failure symbol $\perp$ with a polynomial-sized certificate that certifies that $\overline{\mathrm{B}}(\mathcal{R}, -\delta) = \emptyset$.
- (Approximate maximization): A vector $z \in \mathbb{Q}^m \cap \overline{\mathrm{B}}(\mathcal{R}, \delta)$, such that $F(z) + \epsilon \geq \max_{y \in \overline{\mathrm{B}}(\mathcal{R}, -\delta)} F(y)$.

The weak approximate version of the projection problem, denoted by $\widehat{\Pi}_X^{\epsilon,\delta}(x)$, is an instance of a weak constrained optimization problem. The problem is defined as follows:

WEAK APPROXIMATE PROJECTION PROBLEM

Input: Two rational numbers $\epsilon, \delta > 0$ and a weak separation oracle WSO_X for a non-empty closed convex set $X \subseteq \mathbb{R}^{m*}$ and a vector x that belongs to $\mathbb{Q}^m \cap X$.

Output: One of the following cases:

- (Violation of non-emptiness): A failure symbol $\perp$ followed by a polynomial-sized witness that certifies that $\overline{\mathrm{B}}(X, -\delta) = \emptyset$.
- (Approximate projection): A vector $z \in \mathbb{Q}^m \cap \overline{\mathrm{B}}(X, -\delta)$, such that:

$$\|z - x\|_2^2 \leq \min_{y \in \overline{\mathrm{B}}(X,\delta)} \|x - y\|_2^2 + \epsilon.$$

For the case of weak separation oracles, a disparity issue arises (see Remark B.4 in [152]) and to alleviate this problem, the authors assumed that the L-Hausdorff correspondences to be $(\eta, \sqrt{m}, L)$-well conditioned, i.e., $\forall x \in \mathbb{R}^{m*}$, there exists $a \in \mathcal{R}(x)$ such that $\overline{\mathrm{B}}(a, \eta) \subseteq \mathcal{R}(x)$. Considering this assumption and leveraging the ellipsoid method will assure the existence of an oracle-polynomial time algorithm (Algorithm 1 in [152]) that has a polynomial non-emptiness certificate in the following modification on the weak constrained convex optimization. We also can define (r, R)-*well-boundedness* of convex set X by having the property that $\exists \boldsymbol{a}_0 \in \mathbb{R}^m : \overline{\mathrm{B}}\left(\boldsymbol{a}_0, r\right) \subseteq X \subseteq \overline{\mathrm{B}}(0, R)$.

Weak Convex (Feasibility/Projection/Optimization) Problem

Input: A zeroth and first order oracle for the convex function $F : \mathbb{R}^{m*} \to \mathbb{R}$, two rational numbers $\delta, \epsilon > 0$ and a weak separation oracle $\mathrm{WSO}_{\mathcal{R}}$ for a non-empty closed convex set $\mathcal{R} \subseteq \mathbb{R}^{m*}$.

Output: One of the following cases:

- (Violation of non-emptiness): A failure symbol $\perp$ with a polynomial-sized witness that certifies that either $\mathcal{R} = \emptyset$ or $\mathrm{vol}(\mathcal{R}) \leq \mathrm{vol}(\overline{\mathrm{B}}(0, \eta))$.
- (Approximate minimization): A vector $z \in \mathbb{Q}^m \cap \overline{\mathrm{B}}(\mathcal{R}, \delta)$, such that $F(z) \leq \min_{y \in \overline{\mathrm{B}}(\mathcal{R}, -\delta)} F(y) + \epsilon$.

The notation $\mathrm{vol}(A)$ represents the Lebesgue volume measure of the set A.

Next, we establish similar results focusing on weak separation oracles. First, we need to introduce the computational problem of approximating Kakutani's fixed point and generalized equilibrium similar to [152].

Kakutani With A Weak Separation Oracle (via $C_{\mathcal{R}(x)}$)

Input: A circuit $C_{\mathcal{R}}$ that represents weak separation oracle for an $(\eta, \sqrt{m}, L)$ well-conditioned correspondence $\mathcal{R} : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$ and an accuracy parameter α.

Output: One of the following cases:

- (Violation of η-non-emptiness): A vector $x \in \mathbb{R}^{m*}$ such that $\mathrm{vol}(\mathcal{R}(x)) \leq \mathrm{vol}(\overline{\mathrm{B}}(0, \eta))$,
- (Violation of L-Hausdorff Lipschitzness): Four vectors $p, q, z, w \in \mathbb{R}^{m*}$ and a constant $\epsilon > 0$ such that $w = \widehat{\Pi}^{\epsilon,\epsilon}_{\mathcal{R}(q)}(q)$ and $z = \widehat{\Pi}^{\epsilon,\epsilon}_{\mathcal{R}(p)}(w)$ but $\|z - w\| > L\|p - q\| + 3(1 + c_{\eta,m})\epsilon^4$,
- Vectors $x, z \in \mathbb{R}^{m*}$ such that $\|x - z\| \leq \alpha$ and $z \in \mathcal{R}(x) \Leftrightarrow d(x, \mathcal{R}(x)) \leq \alpha$.

---

[4] $c_{\eta,m}$ is a constant based on the dimension m and η. For more information, see [152].

Theorem B.2.1 ([152]). *The Kakutani fixed-point problem for a weak separation oracle is PPAD-complete.*

Definition B.2.2. *((ϵ, η)-Approximate Generalized Equilibrium). Let $(\mathcal{U}, \mathcal{S}, \mathcal{R})$ be a concave game with k players where $\mathcal{R}$ represents the common constraint set (that includes all strategies that are acceptable given this constraint). Define $\mathcal{R}_\eta = \overline{\mathrm{B}}(\mathcal{R}, \eta)$, and $\mathcal{R}_{-\eta} = \overline{\mathrm{B}}(\mathcal{R}, -\eta)$ (recall that $\mathcal{R}$ is the convex set that imposes one common convex constraint). A vector $\boldsymbol{x}^\star \in \mathcal{R}_\eta$ is an (ϵ, η)- approximate generalized equilibrium of this game if for every $i \in [k]$ and every $x_i \in S_i$ such that $\left(x_i, x^\star_{-i}\right) \in \mathcal{R}_{-\eta}$ it holds that:*

$$u_i\left(\boldsymbol{x}^\star\right) + \epsilon \geq u_i\left(x_i, \boldsymbol{x}^\star_{-i}\right) \qquad (\ddagger)$$

Remark B.2.3. When $\eta = 0$, we have ϵ-approximate generalized equilibrium.

Theorem B.2.4 ([152]). *The computational problem Strongly Concave Games for a weak separation oracle is PPAD-complete.*

STRONGLY CONCAVE GAMES PROBLEM WITH WSO

Input: We receive as input all the following:

- k circuits representing the utility functions $(u)_{i=1}^{k}$ for all k players,
- A Lipschitzness parameter L, a strong concavity parameter μ, and accuracy parameter ϵ,
- $S = \Pi_{i=1}^{k} S_i$ a convex set called *the strategy domain* where S_i represent the strategy domain for each player i,
- An arithmetic circuit representing a weak separation oracle for the well-bounded set $\mathcal{R}$ (strategies that satisfy the common constraint) that is a non-empty, convex, and compact subset of S,
- Accuracy parameters ϵ, η.

Question: We output one of the following:

- (Violation of almost non-emptiness) A certificate that $\mathrm{vol}(\mathcal{R}) \leq \mathrm{vol}(\overline{\mathrm{B}}(0, \eta))$.
- (Violation of Lipschitz Continuity) A certification that there exist at least two vectors $x, y \in S$ and an index $i \in [n]$ such that $|u_i(x) - u_i(y)| > L\|x - y\|$.
- (Violation of Strong Concavity) An index $i \in [n]$, three vectors $x_i, y_i \in S_i,\ \mathbf{x}_{-i} \in S_{-i} = \Pi_{j=1, j\neq i}^{k} S_j$ and a number $\mu \in [0, 1]$ such that:

$$u_i\left(\lambda x_i + (1-\lambda) y_i, \mathbf{x}_{-i}\right) < \lambda u_i\left(x_i, \mathbf{x}_{-i}\right) + (1-\lambda) u_i\left(y_i, \mathbf{x}_{-i}\right) + \frac{\lambda(1-\lambda)}{2}\mu \cdot \left\|(x_i, \mathbf{x}_{-i}) - (y_i, \mathbf{x}_{-i})\right\|_2^2$$

- An (ϵ, δ)-approximate generalized equilibrium in the sense of (‡).

B.2.2 Variational Inequalities with Weak Separation Oracles

Next, we investigate the variational problems for weak separation oracles.

GQVI($\mathcal{F}$, $\mathcal{R}$) With weak Separation Oracles

---

Input: We receive as input all the following:

- A circuit $C_{\mathcal{R}}$ which represents a weak separation oracle for a $(\eta, \sqrt{m}, L_{\mathcal{R}})$ well-conditioned correspondence $\mathcal{R} : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$,
- A circuit $C_{\mathcal{F}}$ which represents a $(\eta, \sqrt{m}, L_{\mathcal{F}})$ well-conditioned correspondence $\mathcal{F} : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$ which is also γ-strongly convex valued,
- Accuracy parameters β, η.

Output: One of the following cases:

- (Violation of almost non-emptiness): A vector $x \in \mathbb{R}^{m*}$ with a certificate that either $\mathrm{vol}(\mathcal{R}(x)) \leq \mathrm{vol}(\overline{\mathrm{B}}(0, \eta))$ or $\mathrm{vol}(\mathcal{F}(x)) \leq \mathrm{vol}(\overline{\mathrm{B}}(0, \eta))$.
- (Violation of $L_{\mathcal{R}}$-Hausdorff Lipschitzness of $\mathcal{R}$) Four vectors $p, q, z, w \in \mathbb{R}^{m*}$ and a constant $\epsilon > 0$ such that $w = \widehat{\Pi}^{\epsilon,\epsilon}_{\mathcal{R}(q)}(q)$ and $z = \widehat{\Pi}^{\epsilon,\epsilon}_{\mathcal{R}(p)}(w)$ but $\|z - w\| > L_{\mathcal{R}}\|p - q\| + 3(1 + c_{\eta,m})\epsilon$,
- (Violation of $L_{\mathcal{F}}$-Hausdorff Lipschitzness of $\mathcal{F}$): Four vectors $p, q, z, w \in \mathbb{R}^{m*}$ and a constant $\epsilon > 0$ such that $w = \widehat{\Pi}^{\epsilon,\epsilon}_{\mathcal{F}(q)}(q)$ and $z = \widehat{\Pi}^{\epsilon,\epsilon}_{\mathcal{F}(p)}(w)$ but $\|z - w\| > L_{\mathcal{F}}\|p - q\| + 3(1 + c_{\eta,m})\epsilon$,
- Two tuples (x, w) and (x^*, w^*) with $\|(x, w) - (x^*, w^*)\| \leq \beta$ such that $x^* \in \overline{\mathrm{B}}(\mathcal{R}(x), \eta)$ and $w^* \in \overline{\mathrm{B}}(\mathcal{F}(x), \eta)$ such that $(y - x)^T w^* + \beta \geq 0, \quad \forall y \in \overline{\mathrm{B}}(\mathcal{R}(x), -\eta)$

Remark B.2.5. We can similarly define a version with strong convexity of correspondences which is more restricted. The constant $c_{\eta,m}$ comes from Lemma A.2.1.

B.2.3 Special Cases: GVI and VI

Next, similar to the main part of the work, we investigate QVI and VI for weak separation oracles.

$QVI(F, \mathcal{R})$ WITH A WEAK SEPARATION ORACLE

Input: We receive as input all the following:

- A circuit $C_{\mathcal{R}}$ which represents a weak separation oracle for a $(\eta, \sqrt{m}, L_{\mathcal{R}})$ well-conditioned correspondence $\mathcal{R} : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$.
- A circuit C_F which represents a L_F-Lipschitz function $F : \mathbb{R}^{m*} \to \mathbb{R}^{m*}$
- Accuracy parameters β, η.

Output: One of the following cases:

- (Violation of almost non-emptiness): A certificate that $\text{vol}(\mathcal{R}(x)) \leq \text{vol}(\overline{\text{B}}(0, \eta))$.
- (Violation of $L_{\mathcal{R}}$-Hausdorff Lipschitzness of $\mathcal{R}$): Four vectors $p, q, z, w \in \mathbb{R}^{m*}$ and a constant $\epsilon > 0$ such that $w = \widehat{\Pi}^{\epsilon,\epsilon}_{\mathcal{R}(q)}(q)$ and $z = \widehat{\Pi}^{\epsilon,\epsilon}_{\mathcal{R}(p)}(w)$ but $\|z - w\| > L_{\mathcal{R}}\|p - q\| + 3(1 + c_{\eta,m})\epsilon$.
- (Violation of L_F-Hausdorff Lipschitzness of F): Two vectors $p, q \in \mathbb{R}^{m*}$ such that $\|F(p) - F(q)\| > L_F\|p - q\|$.
- Two vectors x and x^* with the condition $\|x - x^*\| \leq \beta$ such that $x^* \in \overline{\text{B}}(\mathcal{R}(x), \eta)$ such that $(y - x)^T F(x) + \beta \geq 0, \quad \forall y \in \overline{\text{B}}(\mathcal{R}(x), -\eta)$

$VI(F)$ With a weak Separation Oracle

Input: We receive as input all the following:

- A circuit $C_{\mathcal{R}}$ which represents a weak separation oracle for a non-empty closed convex set $\mathcal{R}$,
- A circuit C_F which represents a L_F-Lipschitz function $F : \mathbb{R}^{m*} \rightarrow \mathbb{R}^{m*}$,
- Accuracy parameters β, η.

Output: One of the following cases:

- (Violation of almost non-emptiness): A certificate that $\text{vol}(\mathcal{R}) \leq \text{vol}(\overline{\text{B}}(0, \eta))$
- (Violation of L_F-Hausdorff Lipschitzness of F): Two vectors $p, q \in \mathbb{R}^{m*}$ such that $\|F(p) - F(q)\| > L_F\|p - q\|$,
- One vector with the condition $x \in \overline{\text{B}}(\mathcal{R}, \eta)$ such that $(y - x)^T F(x) + \beta \geq 0, \quad \forall y \in \overline{\text{B}}(\mathcal{R}, -\eta)$

Theorem B.2.6. *The generalized quasi-variational inequality problem (GQVI), quasi-variational inequality problem (QVI) variational inequality problem (VI) for strong separation oracles are PPAD-complete.*

Proof. We prove inclusion in PPAD in the following. The PPAD-hardness follows from the hardness results in strong separation oracles. □

Using the following proposition, we show that the computational version of GQVI for weak separation oracles is in PPAD.

Theorem B.2.7. *The generalized quasi-variational inequality problem $GQVI(\mathcal{F}, \mathcal{R})$ for weak separation oracles* $\text{SO}_{\mathcal{R}}$ *and* $\text{SO}_{\mathcal{F}}$ *is in PPAD.*

Proof. Without loss of generality, we can consider $[-1, 1]^m$ instead of $\mathbb{R}^{m*}$. Define $\mathcal{R}(x) = \{y \in \overline{\text{B}}(\mathcal{R}(x), \delta)\}$ and $\mathcal{R}'(x) = \{y \in \overline{\text{B}}(\mathcal{R}(x), -\delta)\}$. By definition, $\mathcal{R}(x)$ and $\mathcal{R}'(x)$ have a weak separation oracle and are convex. The proof is organized as follows. First, we define the following:

$$\Phi(y,x,w) = -(y-x)^T w - \gamma(\|y\|_2^2)$$

$$\Pi(x,w) = \{y \in \mathcal{R}(x) \mid \Phi(y,x,w) > max_{y\in\mathcal{R}'(x)}\Phi(y,x,w) - \epsilon\}$$

Next, we show that for a constant κ', $\Psi(x,w) = (\Pi(x,w), \mathcal{F}(x))$ is κ'-Lipschitz continuous in (x,w). Then, we construct a weak separation oracle for Ψ. Finally, we show that an approximate Kakutani's fixed point of this function will provide an approximate solution to the given GQVI with WSO.

To show that Ψ is κ'-Lipschitz continuous, we show that $d_H(\Psi(x_1,w_1), \Psi(x_2,w_2)) \leq \kappa'\|(x_1,w_1)-(x_2,w_2)\|_p^q + c$ for some constants q,p and c. Similar to the strong separation oracle case, $\Phi(y,x,w)$ is (2γ)-strongly concave function of y. In addition, Φ is G-Lipschitz continuous where $G = m + 2\gamma m$ for all x and w. We also have access to the sub-gradients of Φ. Now, we define the following sets:

$$H_+(x,w) = \{y \in \mathcal{R}(x) \mid \Phi(y,x,w) = max_{y\in\mathcal{R}(x)}\Phi(y,x,w)\}$$

$$H_+^\epsilon(x,w) = \{y \in \mathcal{R}(x) \mid \Phi(y,x,w) \geq max_{y\in\mathcal{R}(x)}\Phi(y,x,w) - \epsilon\}$$

$$H_-(x,w) = \{y \in \mathcal{R}'(x) \mid \Phi(y,x,w) = max_{y\in\mathcal{R}'(x)}\Phi(y,x,w)\}$$

$$H_-^\epsilon(x,w) = \{y \in \mathcal{R}'(x) \mid \Phi(y,x,w) \geq max_{y\in\mathcal{R}'(x)}\Phi(y,x,w) - \epsilon\}$$

We also need the following lemma (Lemma E2 from [152]) that helps us relate these two sets[5].

Lemma B.2.8 ([152])**.** *For a function $f : A \subseteq \mathbb{R}^d \to \mathbb{R}^d$, which is μ-strongly concave and G-lipschitz and a well-bounded convex set S, i.e. $\exists a_0 \in \mathbb{R}^d : \overline{\mathrm{B}}(a_0, r) \subseteq S \subseteq \overline{\mathrm{B}}(0,R) \subseteq A$, it holds that:*

$$\left|\max_{\boldsymbol{a}\in\overline{\mathrm{B}}(S,\eta_1)} f(\boldsymbol{a}) - \max_{\boldsymbol{a}\in\overline{\mathrm{B}}(S,\eta_2)} f(\boldsymbol{a})\right| \leq C_{G,\mu}\,\|\eta_1-\eta_2\|$$

for some constant $C_{G,\mu}$.

By optimality KKT conditions for maximization of a concave function with respect

[5]Both $\mathcal{R}(x)$ and $\mathcal{R}'(x)$ are convex sets. Recall that we could consider $\mathbb{R}^{m*}$ alternatively for well-boundedness.

to the constraint set $\mathcal{R}(x)$ for all $y \in \mathcal{R}(x)$ we have:

$$\partial\Phi(y,x,w)^T(y^* - y) \geq 0 \text{ where } y^* = argmax_{y\in\mathcal{R}(x)}\Phi(y,x,w)$$

In addition, (2γ)-strong-concavity of $\Phi(\cdot,x,w)$ results the following inequality:

$$\Phi\left(y^*,x,w\right) - \Phi(y,x,w) \geq \partial\Phi\left(y^*,x,w\right)^\top\left(y^* - y\right) + \gamma\left\|y - y^*\right\|_2^2$$

By combining the previous inequalities, we have:

$$\Phi\left(y^*,x,w\right) - \Phi(y,x,w) \geq \gamma\left\|y - y^*\right\|_2^2$$

In conclusion, for $y^* = argmax_{y\in\mathcal{R}(x)}\Phi(y,x,w)$ and $y \in H_+^\epsilon(x,w)$, we have the following inequality:

$$\gamma\left\|y - y^*\right\|_2^2 \leq \epsilon \text{ or equivalently } \left\|y - y^*\right\|_2 \leq \sqrt{\frac{\epsilon}{\gamma}} \tag{B.3}$$

Applying Lemma B.2.8, we have that for every $y \in \Pi(x,w)$ it holds that there exists a constant k such that:

$$\Phi(y,x,w) \geq \max_{y\in\mathcal{R}'(x)}\Phi(y,x,w) - \epsilon \geq \max_{y\in\mathcal{R}(x)}\Phi(y,x,w) - \epsilon - k$$

Therefore for every $y \in \Pi(x,w)$, it holds that $y \in H_+^{\epsilon+k}(x,w)$. Consequently,

$$\begin{cases} \mathrm{d_H}\left(\Pi\left(x_1,w_1\right), H_+\left(x_1,w_1\right)\right) \leq \sqrt{\frac{\epsilon+k}{\gamma}} \\ \mathrm{d_H}\left(\Pi\left(x_2,w_2\right), H_+\left(x_2,w_2\right)\right) \leq \sqrt{\frac{\epsilon+k}{\gamma}} \\ \mathrm{d_H}\left(H_+\left(x_1\right), H_+\left(x_2\right)\right) = \mathrm{d}\left(H_+\left(x_1\right), H_+\left(x_2\right)\right) \leq \kappa\left\|x_1 - x_2\right\|_2^{1/2} \end{cases}$$

The last inequality comes from the application of the Theorem 3.2.9 to $f((x,w),y) = H_+(y,x,w)$ and $g((w,x)) = \mathcal{R}(x)$. This shows that Π is approximate Hausdorff-$(\frac{1}{2},2)$-Holder continuous:

$$d_H(\Pi(x_1,w_1),\Pi(x_2,w_2)) \leq \kappa\|(x_1,w_1) - (x_2,w_2)\|_2^{\frac{1}{2}} + 2\sqrt{\frac{\epsilon+k}{\gamma}} \tag{B.4}$$

Since $\mathcal{F}$ is $L_\mathcal{F}$-Hausdorff Lipschitz, for some constant c' by the definition of Ψ we can there exists a constant κ':

$$d_H(\Psi(x_1, w_1), \Psi(x_2, w_2)) \leq \kappa' \|(x_1, w_1) - (x_2, w_2)\|_2^{\frac{1}{2}} + c' \tag{B.5}$$

This shows that Ψ is κ'-Hausdorff Lipschitz (and also $(\frac{1}{2}, 2)$-Holder) continuous.

To establish a reduction employing the computational Kakutani's Problem variant, it is essential to establish the existence of a bounded-radius ball within the correspondence $\Psi(x, w)$. Let $y^*_{(x,-\eta)} = argmax_{y \in \mathcal{R}'(x)} \Phi(y, x, w)$. By almost non-emptiness, $\text{vol}(\mathcal{R}(x)) > \text{vol}(\overline{\text{B}}(0, \eta))$ and $\text{vol}(\mathcal{F}(x)) > \text{vol}(\overline{\text{B}}(0, \eta))$.

1. $\exists \hat{v} : \|\hat{v}\| = 1$ & $\mathcal{V}_x := \overline{\text{B}}\left(y^\star_{(x,-\eta)} - \frac{\eta}{2}\hat{v}, \min\{\eta/2, \epsilon/G\}\right) \subseteq \mathcal{R}'(x) = \overline{\text{B}}(\mathcal{R}(x), -\eta)$
2. By Lipschitzness of $\phi(\cdot, x, w)$, $y \in \mathcal{V}_x$:

$$\left|\phi(x, \boldsymbol{y}) - \phi\left(x, \boldsymbol{y}^\star_{(\boldsymbol{x},-\eta)}\right)\right| \leq G \min\{\eta/2, \epsilon/G\}$$

From (2) we can $\Phi(\mathcal{V}_x, x, w) \subseteq \left[\Phi\left(y^\star_{(x,-\eta)}, x, w\right) - \epsilon, \Phi\left(y^\star_{(x,-\eta)}, x, w\right)\right] \overset{(1)}{\Rightarrow} \mathcal{V}_x \subseteq \Pi(x, w)$. Therefore, $\Pi(x, w)$ always contains a ball of radius $\min\{\eta/2, \epsilon/G\}$. Similar to the strong separation oracle case, Ψ is non-empty as well (we know $\text{vol}(\mathcal{F}(x)) > \text{vol}(\overline{\text{B}}(0, \eta))$ and $\mathcal{F}$ is Lipschitz).

Now we are ready to construct a weak separation oracle for $\Psi(x, w)$ leveraging a slightly modified version of the Weak Convex (Feasibility/Projection/Optimization) Problem stated in the following.

Modified Weak Convex (Feasibility/Projection/Optimization) Problem

---

Input: A zeroth and first-order oracle for the concave function $G : \mathbb{R}^{m*} \to \mathbb{R}$, two rational numbers $\delta, \epsilon > 0$ and a weak separation oracle $\mathrm{SO}_{\mathcal{R}}$ for a non-empty closed convex-valued correspondence $\mathcal{R} : \mathbb{R}^{m*} \rightrightarrows \mathbb{R}^{m*}$ and one input x.

Output: One of the following cases:

- (Violation of non-emptiness): A failure symbol $\perp$ with a polynomial-sized witness that certifies that either $\mathcal{R}(x) = \emptyset$ or $\mathrm{vol}(\mathcal{R}(x)) \leq \mathrm{vol}(\overline{\mathrm{B}}(0, \eta))$.
- (Approximate maximization): A vector $z \in \mathbb{Q}^m \cap \overline{\mathrm{B}}(\mathcal{R}(x), \delta)$, such that $G(z) + \epsilon \geq \max_{y \in \overline{\mathrm{B}}(\mathcal{R}(x), -\delta)} G(y)$.

Let us define $\mathcal{R}''(x) = \{y \in \overline{\mathrm{B}}(\mathcal{R}(x), min\{\delta, \epsilon\})\} \subseteq \mathcal{R}'(x)$. We can compute a solution $y^* \in \mathcal{R}''(x)$ such that $\Phi(y^*, x, w) \geq \max_{y \in \mathcal{R}''(x)} \Phi(y, x, w) - \min\{\epsilon, \eta\} \geq \max_{y \in \mathcal{R}'(x)} \phi(x, y) - \epsilon$ using sub-gradient ellipsoid central cut method (see Appendix B.3). And as discussed, such a separation oracle exists. Thus, we can substitute a WSO for Ψ by considering a separation oracle for the following set:

$$\overline{\Psi}_s(x, w) = \{(y, w) \in \big(\mathcal{R}(x), \overline{\mathrm{B}}\left(\mathcal{F}(x), \epsilon\right)\big) \mid -\Phi(y, x, w) \leq \sigma\}$$

where $\sigma = -\Phi(y^*, x, w)$. In other words, are looking for $(y, w) \in (\mathcal{R}(x), \overline{\mathrm{B}}\left(\mathcal{F}(x), \epsilon\right))$ given the strong separation oracles representing $\mathcal{F}$ and D such that:

$$(y - x)^T w + \gamma \|y\|_2^2 \geq (y^* - x)^T w + \gamma \|y^*\|_2^2$$

Next, we give Ψ as input to the computational Kakutani problem with accuracy parameter $\alpha = \frac{\epsilon'}{\kappa'}$ where $\epsilon' = \epsilon h$. The output of this Kakutani instance will be two points $(x, w) \in ([-1, 1]^m, [-1, 1]^m)$ and $z = (x^*, w^*) \in \Psi(x, w)$ where $\|(x, w) - (x^*, w^*)\| \leq \frac{\epsilon'}{\kappa'}$ and $d((x, w), \Psi(x, w)) \leq \frac{\epsilon'}{\kappa'}$. Thus, $w^* \in \mathcal{F}(x)$ and $d(w, \mathcal{F}(x)) \leq \frac{\epsilon'}{\kappa'}$. In addition, $x^* \in \Pi(x, w)$ and $d(x, \Pi(x, w)) \leq \frac{\epsilon'}{\kappa'}$. By the definition of Π, for every $y \in \mathcal{R}'(x)$, $\Phi\left(x^*, x, w\right) \geq \Phi\left(y, x, w\right) - \epsilon$. In conclusion:

$$(y - x)^T w + \gamma \|y\|_2^2 \geq (x^* - x)^T w + \gamma \|x^*\|_2^2 - \epsilon, \quad \forall y \in \mathcal{R}'(x)$$

Recall that $x, y, w \in [-1,1]^m$:

$$(y-x)^T w \geq \pm \frac{2\sqrt{m}\epsilon'}{\kappa'} \pm 4\gamma \pm \frac{2\sqrt{m}\epsilon'}{\kappa'} - \epsilon, \quad \forall y \in \mathcal{R}'(x)$$

Finally, by knowing the facts that $\|(x,w)-(x^*,w^*)\|_2^2 \leq \frac{\epsilon'}{\kappa'}$ and $w^* \in \mathcal{F}(x)$, we can simply deduce that:

$$(y-x)^T w^* \geq \pm \frac{4\sqrt{m}\epsilon'}{\kappa'} \pm u, \quad \forall y \in \mathcal{R}'(x)$$

Considering appropriate numbers for ϵ, h (recall that $\epsilon' = \epsilon h$) and γ will imply the following:

$$(y-x)^T w^* + \beta \geq 0, \quad \forall y \in \overline{\mathrm{B}}(\mathcal{R}(x), -\eta)$$

□

B.2.4 Remedial L/F Equilibrium With Weak Separation Oracles

Building on previous contributions and theorems, we aim to further explore them within the context of weak separation oracles. For the partial responses while considering the definitions of $Z^{\mathrm{I}}(x_{\mathrm{I}}, x_{\mathrm{II}}^*)$, $W_{(i,\mathrm{I})}(x_{\mathrm{II}}^*)$ and $V_{(i,\mathrm{I})}(x_{\mathrm{II}}^*)$ for leader I's optimization problem and considering x_{II}^* as an exogenous variable, we find a solution $(x_{\mathrm{I}}, y_{\mathrm{I}})$ to:

$$\begin{aligned} &\mathrm{Min}\, \phi_{\mathrm{I}}\left(x_{\mathrm{I}}, x_{\mathrm{II}}^*, y_{\mathrm{I}}\right) \\ &\mathrm{s.t}\ x_{\mathrm{I}} \in \overline{\mathrm{B}}(X^{\mathrm{I}}, \eta) \\ &\mathrm{and} \quad (x_{\mathrm{I}}, y_{\mathrm{I}}) \in \overline{\mathrm{B}}\left(\mathrm{graph}\, Z^{\mathrm{I}}\left(\cdot, x_{\mathrm{II}}^*\right), \eta\right) \end{aligned} \tag{B.6}$$

For leader II, the optimization problem with the surrogate complementarity conditions can be defined similarly by considering $Z^{\mathrm{II}}(x_{\mathrm{I}}^*, x_{\mathrm{II}})$, $W_{(i,\mathrm{II})}(x_{\mathrm{I}}^*)$ and $V_{(i,\mathrm{II})}(x_{\mathrm{I}}^*)$ where x_{I}^* will be an exogenous variable:

$$\begin{aligned} &\mathrm{Min}\, \phi_{\mathrm{II}}\left(x_{\mathrm{I}}^*, x_{\mathrm{II}}, y_{\mathrm{I}}\right) \\ &\mathrm{s.t}\ x_{\mathrm{II}} \in \overline{\mathrm{B}}(X^{\mathrm{II}}, \eta) \\ &\mathrm{and} \quad (x_{\mathrm{II}}, y_{\mathrm{II}}) \in \overline{\mathrm{B}}\left(\mathrm{graph}\, Z^{\mathrm{II}}\left(x_{\mathrm{I}}^*, \cdot\right), \eta\right) \end{aligned} \tag{B.7}$$

Finally, we are now ready to define the computational version of finding an equilibrium in this setting:

Remedial L/F Equilibrium With Weak Separation Oracles

Input: We receive as input all the following:

- Two linear arithmetic circuits representing the convex loss functions (ϕ_{I},ϕ_{II}) for two leaders,
- Two linear arithmetic circuits representing weak separation oracles for non-empty, convex, and compact sets X^{I} and X^{II} for the leaders I and II,
- Two linear arithmetic circuits representing weak separation oracles for Z^{I} and Z^{II} that represent restricted or relaxed responses of the followers for each leader respectively that are two non-empty, convex, and compact correspondences.
- Accuracy parameters β and η.

Output: One of the following cases:

- (Violation of almost non-emptiness): A certificate indicating that at least one of the following cases is almost empty:
  - X^{I} or X^{II}
  - $Z^{\mathrm{I}}(\cdot, x)$ for some $x \in X^{\mathrm{I}}$
  - $Z^{\mathrm{II}}(x, \cdot)$ for some $x \in X^{\mathrm{II}}$
- (Violation of convexity): of any of the loss functions of the inputs,
- (Approximate Minimization): Vectors $(x_{\mathrm{I}}^*, y_{\mathrm{I}}^*, x_{\mathrm{II}}^*, y_{\mathrm{II}}^*)$ having the following relationship:
  - $\phi_{\mathrm{I}}(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*, y_{\mathrm{I}}^*) \leq \beta +$ Sol(B.6)
  - $\phi_{\mathrm{II}}(x_{\mathrm{I}}^*, x_{\mathrm{II}}^*, y_{\mathrm{II}}^*) \leq \beta +$ Sol(B.7)

Proposition B.2.9. *In a multi-leader-follower game, the problem of finding a remedial (ϵ, δ)-L/F-equilibrium where the constraints of the followers are given by weak separation oracles is PPAD-complete.*

Proof. The hardness simply follows by considering $\delta = 0$. Inclusion in PPAD is similar to 4.1.6 with the difference that we consider the weak separation oracle cases discussed in this section. □

B.3 A Slightly Generalized Oracle Polynomial-Time Subgradient Ellipsoid Central Cut

For the sake of completeness, we will present a more generalized version of the subgradient-cut method to solve the modified convex-constrained optimization problem in polynomial time, given the separation oracles. Our generalized polynomial-time oracle algorithm is general enough to capture all desired properties and also future similar applications. For example, for the exception of violation of almost non-emptiness in multi-leader-follower games, we need to output specific almost-emptiness exceptions to denote whether $\text{vol}(\mathcal{R}(x))$ or $\text{vol}(\mathcal{F}(x))$ is very small, and our algorithm gives a suitable exception output.

Assume that $f(x) = (f_1(x), \ldots, f_k(x))$ (compared to f being a single coordinate valued function in [152]) and $\mathcal{X} = (X_1, \ldots, X_s)$ (compared to X being one set). We also know that for each $i \in [s]$, X_i has a weak separation oracle. Similar to [152], for each $i \in [k]$, the approximate value and subgradient oracle for the objective function f_i are available. One main difference of our algorithm is that it changes one index of $x^{(t)} \in (X_1, \ldots, X_s)$ in each step t while other coordinates stay the same.

The cutting plane methods are distinguished by their construction of sets $M^{(t)}$ and selection of query points $x^{(t)}$. These methods exhibit an exponential decrease in the volume of $M^{(t)}$ as t increases, which leads to linear convergence guarantees in the presence of gradient and value oracles[6]. Given a desirable approximate parameter ϵ and also a margin δ, we can set some thresholds $T_{emptiness}$ and $T_{ellipsoid}$ so that we achieve the desirable outputs for the minimization (and also maximization) problems, such as the modified convex-constrained optimization problem that we discussed before.

Proposition B.3.1. *There exists a cutting plane method, referred to as the "central-*

[6]Here, we only have a more complex function with constant coordinates, and this will not ruin the exponential decrease in size guarantee.

ALGORITHM 4: Subgradient Central-Cut Ellipsoid Method

Input: Gradient and value oracles $\mathrm{O}^{f_i}_{\text{grad}}, \mathrm{O}^{f_i}_{\text{val}}$ with accuracies $(\epsilon_{\text{grad}}, \epsilon_{\text{val}})$, for each $i \in [k]$.

Weak separation oracles $\mathrm{O}^{i}_{\text{sep}}$ for set X_i with margin δ, for each $i \in [s]$

```
index ← 1;
for t ∈ [T_ellipsoid] do
    if x^(t) ∈ B̄(𝒳, δ) ;                                  /* x^(t) = (x_1^(t), ..., x_k^(t)) */
    then
        for i ∈ [k] do
            Call a gradient oracle g_i^(t) ← O^{f_i}_grad (x^(t));
        end
        if ∀i ∈ [k], ‖g_i^(t)‖ ≤ G_threshold then
            Output: x^(t);
        else
            for i ∈ [k] do
                w_i^(t) ← g_i^(t) / ‖g_i^(t)‖_∞ (Output A);
            end
        end
    else
        if mod(index, s) == 0 then
            Call the separation oracle w_i^(t) ← O^f_sep (x_i^(t));
            calls_index ← calls_index + 1;
            index ← mod (index, s) + 1;
            if calls_index > T_emptiness then
                Output: ⊥_index ;   /* Emptiness of the respective index in the
                                       domain */
            end
        end
    end
    For each i ∈ [k], construct an ellipsoid M_i^(t+1) such that :
      { x ∈ M_i^(t) : w_i^(t)⊤ (x − x^(t)) ≤ δ } ⊆ M_i^(t+1);
    Let x^(t+1) have all of the centroids of M_i^(t+1) for each i ∈ [k];
end
```

Output: The iteration $\bar{x} \in \operatorname{argmin}\left\{\mathrm{O}^{f}_{\text{val}}(x) \mid x \in \left\{x^{(1)}, \cdots, x^{(T_{\text{ellipsoid}})}\right\} \cap \overline{\mathrm{B}}(\mathcal{X}, \delta)\right\}$

cut Ellipsoid method", with a decay rate of $\theta = O(1/d)$, such that for all $i \in [k]$[7]*:*

$$\frac{\mathrm{vol}\left(M_i^{(t)}\right)}{\mathrm{vol}\left(M_i^{(1)}\right)} \leq e^{-\theta t}$$

Define the following sets:

$$\mathcal{R}_{(i,\epsilon)} = \left\{ x \in \overline{\mathrm{B}}(\mathcal{X}, -\delta) : \min_{x \in \overline{\mathrm{B}}(\mathcal{X}, -\delta)} f_i(x) \leq f_i(x) \leq \min_{x \in \overline{\mathrm{B}}(\mathcal{X}, -\delta)} f_i(x) + \frac{\epsilon}{2} \right\}$$

This set is the set of all ϵ-approximate and δ-marginally inside $\mathcal{X}$ minimum solutions of f_i. First, we need to show that, $\mathcal{R}_{(i,\epsilon)}$ has non-zero volume. If we assume that $\overline{\mathrm{B}}(\mathcal{X}, -\delta) \neq \emptyset$, then by L-lipschitzness of f_i for each $i \in [k]$, we know that $\mathcal{R}_{(i,\epsilon)}$ has a ball of radius $r(\epsilon, \delta) = \min\{\delta, \epsilon/L\}$.

Next, we define $\mathcal{T}_{\text{active}} := \left\{ t \in [T_{\text{ellipsoid}}] \mid x^{(t)} \in \overline{\mathrm{B}}(\mathcal{X}, \delta) \right\}$, and $w^{(t)} = \Pi_{i=1}^{k} w_i^{(t)}$. There are three possible cases:

- Case 1: Assume that for any $t \in \mathcal{T}_{\text{active}}$, any $i \in [k]$, and for any $x_{(i,\epsilon)} \in \mathcal{R}_{(i,\epsilon)}$, we have that ${w_i^{(t)}}^\top \left(x_{(i,\epsilon)} - x^{(t)} \right) \leq \delta$. This can imply that $\forall i \in [k]$, $\forall t \in [T_{\text{ellipsoid}}], \forall x_{(i,\epsilon)} \in \mathcal{R}_{(i,\epsilon)} : {w_i^{(t)}}^\top \left(x_{(i,\epsilon)} - x^{(t)} \right) \leq \delta$. The reasoning follows by the definition of the separation oracles we have $({w_i^{(t)}}^\top \left(x - x^{(t)} \right) \leq \delta$) for all $x \in \overline{\mathrm{B}}(\mathcal{X}, -\delta)$. Thus, for all $i \in [k]$, it holds that:

  $$\forall t \in [T_{\text{ellipsoid}}] : D_{(i,\epsilon)} \subseteq M_i^{(t)} \Rightarrow \mathrm{vol}\left(\mathcal{R}_{(i,\epsilon)}\right) \leq \mathrm{vol}\left(M_i^{(t)}\right)$$

  Next, we show that the above-mentioned condition can happen only if $t \leq C_0 \cdot d^2 \log(d/r(\epsilon, \delta))$. For any $i \in [k]$ some positive constant C_0 independent of d, δ, we have the following inequalities:

  $$\begin{cases} \frac{\mathrm{vol}\left(M_i^{(t)}\right)}{\mathrm{vol}\left(M_i^{(1)}\right)} \leq e^{-\theta t}, \quad \theta = \Theta\left(\frac{1}{d}\right) \\ \frac{\pi^d}{(d/2+1)!} r(\delta, \epsilon)^d = \mathrm{Vol}(\overline{\mathrm{B}}(\mathcal{X}, r(\delta, \epsilon))) \\ \mathrm{vol}(\overline{\mathrm{B}}(\mathcal{X}, r(\delta, \epsilon))) \leq \mathrm{Vol}\left(\mathcal{R}_{(i,\epsilon)}\right) \leq \mathrm{Vol}\left(M_i^{(t)}\right) \\ \mathrm{Vol}\left(M^{(1)}\right) \leq \mathrm{Vol}(Box) \end{cases} \Longrightarrow t \leq C_0 \cdot d^2 \left(\log\left(\frac{d}{2r(\epsilon, \delta)} \right) \right)$$

  If for any $i \in [s]$ the number of used separation oracle calls are greater than

[7] vol is the usual d-dimensional volume.

$T_{\text{emptiness}}$, then $\mathrm{Vol}(\mathcal{X}) \leq \mathrm{Vol}(\overline{\mathrm{B}}(0,\delta))$, or consequently $\overline{\mathrm{B}}(\mathcal{X},-\delta) = 0$. Otherwise, if we set $T_{\text{ellipsoid}} = \max\{C_0, 10\}\, d^2 \left(\log\left(\frac{d}{2r(\epsilon,\delta)}\right)\right)$, then either Case 2 or 3 hold for $C_0 \cdot d^2 \left(\log\left(\frac{d}{2r(\epsilon,\delta)}\right)\right) < t \leq T_{\text{ellipsoid}},$.

- Case 2: If for all $i \in [k]$, $\left\| g_i^{(t)} \right\| \leq G_{\text{threshold}}$ for appropriate choice of $G_{\text{threshold}}$, we will show that $x^{(t)}$ is an ϵ-approximate minimizer. Indeed, for each $i \in [k]$, by convexity $\min_{x \in \overline{\mathbb{B}}(X,-\delta)} f_i(x) \geq f_i\left(x^{(t)}\right) + \min_{x \in \overline{\mathbb{B}}(X,-\delta)} \partial f_i\left(x^{(t)}\right)^\top \left(x - x^{(t)}\right)$. By choosing $G_{\text{threshold}} = O\left(\mathrm{poly}\left(d, \epsilon, \epsilon_{\text{grad}}\right)\right)$ such that $\epsilon \geq \left(G_{\text{threshold}} - \epsilon_{\text{grad}}\right)\sqrt{d}$, we have, for all $i \in [k]$, $f_i\left(x^{(t)}\right) \leq \min_{x \in \overline{B}(\mathcal{X},-\delta)} f_i(x) + \epsilon$ and $f\left(x^{(t)}\right) \leq \min_{x \in \overline{B}(\mathcal{X},-\delta)} f(x) + \epsilon \mathbb{1}$.

- Case 3: Assume that for any index i, the element x_ϵ at iteration $t^\star \in [T_{ellipsoid}]$ has the property that $w_i^{(t^\star)}\left(x_\epsilon - x^{(t^*)}\right) > \delta$. In this case, using the convexity of objective can imply $f_i\left(x^{(t^*)}\right) \leq f_i\left(x_\epsilon\right) - \nabla f_i\left(x^{(t^*)}\right)^\top \left(x_\epsilon - x^{(t^*)}\right)$. The last expression is equal to:

  $f_i\left(x_\epsilon\right) - \left(\nabla f_i\left(x_{t^*}\right) - g_i^{(t^*)}\right)^\top \left(x_\epsilon - x_{t^*}\right) - g_i^{(t^*)\top}\left(x_\epsilon - x_{t^*}\right) \leq f_i\left(x_\epsilon\right) + \epsilon_{\text{grad}}\sqrt{d} - \delta.$
  If we set $\epsilon_{\text{grad}} \leq \frac{\epsilon}{\sqrt{d}}$ and $\delta \leq \frac{\epsilon}{2}$, then $f_i\left(x_{t^*}\right) \leq f_i\left(x_\epsilon\right) + \frac{\epsilon}{2} \leq \min_{x \in \bar{B}(X,-\delta)} f_i(x) + \epsilon$. This means that: $f\left(x_{t^*}\right) \leq f\left(x_\epsilon\right) + \frac{\epsilon}{2}\mathbb{1} \leq \min_{x \in \bar{B}(X,-\delta)} f(x) + \epsilon\mathbb{1}$

Remark B.3.2. It is noteworthy to highlight that the recent study on the optimal fusion of subgradient descent and the ellipsoid method, conducted by Rodomanov and Nesterov [160], has exclusively concentrated on the scenario of a strong oracle. We investigate only the most complex case similar to [152]

Remark B.3.3 (Separation oracle and almost-emptiness)**.** If the number of calls to the weak separation oracle for a coordinate X_i exceeds the threshold $T_{\text{emptiness}}$, this indicates that the ellipsoid has been repeatedly cut without enclosing any substantial feasible region. By the properties of the central-cut ellipsoid method, this can only happen if the effective volume of X_i (after accounting for the margin δ) is smaller than the volume of a δ-ball, i.e., $\overline{\mathrm{B}}(X_i, -\delta)$ is effectively empty. Therefore, returning the $\perp_i$ output correctly signals that X_i is almost-empty in the sense required for the algorithm.

Remark B.3.4 (Coordinate-wise ellipsoid volume decay)**.** In our algorithm, each iteration updates only one coordinate block X_i while keeping the other coordinates fixed. For this coordinate, the standard central-cut ellipsoid method guarantees an exponential decrease of the ellipsoid volume in dimension $\dim(X_i)$. Since we cycle through all coordinates regularly, each ellipsoid $M_i^{(t)}$ experiences asymptotically the same exponential volume decay as in the classical full-dimensional setting. Consequently, the polynomial-time convergence guarantee of the ellipsoid method remains valid for the product space $\mathcal{X} = X_1 \times \cdots \times X_s$, even though only one coordinate is updated per iteration.

Remark B.3.5 (Volume containment under cuts)**.** In Case 1, each cut satisfies $w_i^{(t)\top}(x - x^{(t)}) \leq \delta$ for all $x \in \overline{\mathrm{B}}(X_i, -\delta)$ by the property of the weak separation oracle. Since $\mathcal{R}_{(i,\epsilon)} \subset \overline{\mathrm{B}}(X_i, -\delta)$, the cuts do not remove any point from $\mathcal{R}_{(i,\epsilon)}$, and therefore $\mathcal{R}_{(i,\epsilon)} \subseteq M_i^{(t)}$ is preserved throughout the iterations. This justifies the inequality $\mathrm{vol}(\mathcal{R}_{(i,\epsilon)}) \leq \mathrm{vol}(M_i^{(t)})$.

B.4 Complexity of Nash Equilibrium with a Guaranteed Payoff

Here we revisit a hardness result in [169] and break the proof down into multiple parts in a way that clarifies the connection to our techniques. The result shows that deciding whether a bi-matrix game has a ϵ-approximate Nash equilibrium with a guaranteed payoff of $n-1-\epsilon$ is NP-complete. We call this problem *guaranteed Nash*.

B.4.1 Approximate Guaranteed Nash

The following definition of the game $\mathcal{SV}(\phi, \epsilon)$ from [169] (Theorem 8.6) is simpler than our game $\mathcal{G}(\phi, \epsilon)$, as both players use the same literal set and f was omitted from their strategy set in contrast to the game provided in [52]. The symmetric game $\mathcal{SV}(\phi, \epsilon)$ is defined as follows:

1. $u_1(l^1, l^2) = n - 1$, where $l^1 \neq -l^2$ for all $l^1, l^2 \in L$. This will ensure each player gets a high payoff for playing the aforementioned strategy.

2. $u_1(l, -l) = n - 4$ for all $l \in L$. This will ensure that each player does not play a literal and its negation at the same time.

3. $u_1(v, l) = 0$, where $\mathrm{v}(l) = v$, for all $v \in V, l \in L$. This, along with rule 4, ensures that for each variable v, each agent plays either l or $-l$ with a probability of at least $\frac{1}{n}$, where $\mathrm{v}(l) = \mathrm{v}(-l) = v$.

4. $u_1(v,l) = n$, where $\mathrm{v}(l) \neq v$, for all $v \in V, l \in L$.

5. $u_1(l,x) = n-4$, where $l \in L,\ x \in V \cup C$. This, along with rules 6 and 7, ensures that if both players do not play the literals, then the payoffs cannot meet the guarantees.

6. $u_1(v,x) = n-4$ for all $v \in V, x \in V \cup C$.

7. $u_1(c,x) = n-4$ for all $c \in C, x \in V \cup C$.

8. $u_1(c,l) = 0$ where $l \in c$ for all $c \in C, l \in L$. This, along with rule 9, ensures that for each clause c, each agent plays a literal in the clause c with probability least $\frac{1}{n}$.

9. $u_1(c,l) = n$, where $l \notin c$ for all $c \in C, l \in L$.

Remark B.4.1. The game introduced in both [52, 169] is a symmetric game which means $\forall s_1, s_2 \in S, u_1(s_1, s_2) = u_2(s_2, s_1)$.

Proposition B.4.2. *Let $\epsilon = \frac{1}{2n^3}$ and let the guarantee to each player be $n-1$. Given a 3CNF ϕ is satisfiable iff there exists a ϵ-approximate Nash equilibrium in $\mathcal{SV}(\phi, \epsilon)$ where each player has a guaranteed payoff of $n-1-\epsilon$.*

Proof. We will prove this theorem in the following. □

Essential Lemmas for Proposition B.4.2

Lemma B.4.3. *In any ϵ-approximate Nash equilibrium with the guaranteed payoff $n-1-\epsilon$ in $\mathcal{SV}(\phi, \epsilon)$, clauses and variables are played with a probability of at most ϵ.*

Proof. The social welfare of $\mathcal{SV}(\phi)$ is at most $2n-2$. If neither player plays from the L, the social welfare is at most $2n-4$. When both players play from L with probability $1-\epsilon$, the expected social welfare is at most $2n-2-2\epsilon$. So, for any $\epsilon' > \epsilon$, playing variables and clauses with probability ϵ' will give an expected social welfare of at most $2n-2-2\epsilon' < 2n-2-2\epsilon$. This means at least one player will have an expected payoff less than $n-1-\epsilon$, which violates the definition ϵ-best response.

□

Lemma B.4.4. *In any ϵ-approximate Nash equilibrium with a guaranteed payoff of $n-1-\epsilon$ in $\mathcal{SV}(\phi, \epsilon)$, for any $l \in L$, the probability that the row player plays l or $-l$ is at least $\frac{1}{n} - 2\epsilon$.*

Proof. Suppose this is not correct and there exists one $l \in L$ and the probability of l being played will be less than $\frac{1}{n} - \epsilon - \frac{2\epsilon}{n} \geq \frac{1}{n} - 2\epsilon$. Then, the expected payoff for the column player playing the pure strategy v(l) is at least the summation of these cases:

- When the row player plays l or $-l$: $\left(\frac{1}{n} - \epsilon - \frac{2\epsilon}{n}\right) \cdot 0$.
- When the row player plays a literal other than l or $-l$:

$$(1 - \epsilon - (\frac{1}{n} - \epsilon - \frac{2\epsilon}{n})) \cdot n$$

- When the row player does not play a literal from L, the expected payoff portion of this possibility will be at least: $\epsilon \cdot 0$

The summation will be at least $n-1+2\epsilon$. Since the maximum social welfare is $2n-2$, the other (row) player fails to meet the guarantee $n - 1 - \epsilon$. This is correct for the other player as well. □

Lemma B.4.5. *In any ϵ-approximate Nash equilibrium of $\mathcal{SV}(\phi, \epsilon)$ with a guaranteed payoff of $n - 1 - \epsilon$, for each player and any literal $l \in L$, either l or $-l$ is played with probability $\geq \frac{1}{n} - 2\epsilon - \frac{1}{n^2}$ while the other is played with probability less than $\frac{1}{n^2}$*

Proof. If the row player plays l and the column player plays $-l$, then according to the construction, the sum of payoffs is $2n - 8$. It is not hard to see that the probability that this happens is less than $\frac{\epsilon}{3}$ (by an argument similar to Lemma B.4.3). If both players do so with probability greater than $\frac{\epsilon}{3}$, then the social welfare will be $2n - 2 - \frac{\epsilon}{3}(2n - 2) - \frac{\epsilon}{3}(2n - 8) = 2n - 2 - 2\epsilon$ and this will cause at least one player not gain a payoff of $n - 1 - \epsilon$.

Consider the literals l and $-l$ from L and assume without loss of generality, the row player plays l more than $-l$. Recall that we showed that each player plays either l or $-l$ with probability of at least $\frac{1}{n} - \epsilon - \frac{2\epsilon}{n} \geq \frac{1}{n} - 2\epsilon$. For the sake of contradiction, let us assume that the row player plays l with probability less than $\frac{1}{n} - \left(\frac{1}{n^2} + 2\epsilon\right)$ (still greater than $\frac{1}{n^2}$) which in turn, the row player has to play $-l$ with probability more than $\frac{1}{n^2}$. The column player either plays l with probability less than $\frac{1}{n^2}$ or plays $-l$ with probability greater than $\frac{1}{n} - \epsilon - \frac{2\epsilon}{n}$ .In either case, the probability that they choose both l and $-l$ is at least:

$$\left[\frac{1}{n} - \left(\frac{1}{n^2} + 2\epsilon\right)\right]\left[\frac{1}{n^2}\right] = \frac{1}{n^3} - \frac{1}{n^4} - \frac{1}{2n^6} \geq \frac{1}{2n^3} = \epsilon$$

This is impossible because we showed the probability that both outcomes happen must be less than $\frac{\epsilon}{3}$. So, if the row player must play a literal with probability greater than $\frac{1}{n} - \left(\frac{1}{n^2} + 2\epsilon\right)$, the negation of this literal can be played with probability of at most $\frac{1}{n^2}$. By a symmetric argument, the column player has to follow the same rules. □

Remark B.4.6. The analysis of the case that the column player plays l with probability greater than $\frac{1}{n^2}$ is similar and was not provided in [169]. If the column player plays l with probability greater than $\frac{1}{n} - \frac{1}{n^2} - \epsilon$, by the fact that we know the row player plays $-l$ with a probability greater than $\frac{1}{n^2}$, the probability that they both play l and $-l$ is at least $\epsilon > \frac{\epsilon}{3}$. If the column player plays l with probability less than $\frac{1}{n} - \frac{1}{n^2} - \epsilon$, the column player has to play $-l$ with probability greater than $\frac{1}{n^2}$. By the assumption that we had, the row player plays l with probability greater than $\frac{1}{2}(\frac{1}{n} - 2\epsilon)$. The probability that both players play l and $-l$ is still greater than $\frac{\epsilon}{3}$.

Proof of Proposition B.4.2

The next step is proving Proposition B.4.2.

Proof. If ϕ is satisfiable, we show that there is a uniform Nash equilibrium of $\mathcal{SV}(\phi, \epsilon)$ where each player plays $l^i \in L$ or $-l^i \in L$ (for all i) uniformly with probability $\frac{1}{n}$.

If ϕ is satisfiable, there exists an assignment where $v_1, \ldots, v_n$ are assigned with *true* or *false* that satisfies all clauses. If v_i is assigned *true*, then l^i, otherwise $-l^i$ will be selected as the strategies in the game. Then, $l^1, \ldots, l^n$ are literals that correspond to a satisfying assignment. The expected payoff for each player is $n-1$ because they will always be playing one of l or $-l$. Secondly, there are only two rules that give more than $n-1$, namely, where one of the players plays either a variable or a clause. For example, if the row player plays any clause c with a positive probability, we know that the column player plays some literal in that clause c with probability $\frac{1}{n}$ because the other player randomizes between literals in a satisfying assignment. Then in this case, the row player's payoff is at most $\frac{1}{n} \cdot 0 + \frac{(n-1)}{n} \cdot n = n-1$, so the row player gets the same payoff $n-1$ and is indifferent. This holds for variables as well. In conclusion, assuming one player plays uniform strategies on literals that satisfy the formula, the other player takes the same approach,h and this forms a Nash equilibrium with a guaranteed payoff of $n-1$.

Now suppose that ϕ is not satisfiable. We show that in any ϵ-approximate Nash equilibrium, at least one player always receives an expected payoff of less than $n -$

$1-\epsilon$. For the approximate version, we need to do the following modification to the correspondence between literals and truth assignments compared to [52]. We consider a literal is *true* if it is played more often than its negation. If an assignment does not satisfy the formula, there is at least one clause that does not have a satisfying literal. Any of the players will turn to this clause strategy to receive a payoff of n whenever the opponent plays a literal that is not in that clause. We know that the column player plays literals with probability $> 1-\epsilon$ by Lemma B.4.3, and there are only 3 literals in each clause, each of which the column player plays with probability $\leq \frac{1}{n^2}$ by Lemma B.4.5. By changing the strategy to this clause, the row player will receive at least $\left(1-\epsilon-\frac{3}{n^2}\right)n > n-1+2\epsilon$. So either the row player can do ϵ better by changing his strategy, or he is already receiving $n-1+2\epsilon$, and so the other player does not have a guaranteed payoff of at least $n-1-\epsilon$. □

B.5 Decision Approximate Resilient Nash

Here, we use the hardness results of [169, 52] and Chapter 5 to establish NP-hardness of the decision version of t-resilient Nash equilibrium problem. Next, we define the computational problem, focusing our proof on the simple case of bi-matrix games.

---

t-Resilient Nash in Bi-matrix Games

Input: We receive as input all the following:

- A bi-matrix game $\mathcal{G}$ with k players represented by k utility matrices (u_i for all $i \in [k]$),
- The strategy sets $\mathcal{S}$,
- An accuracy parameter ϵ.

Output: The following case:

- Whether there exists a vector s^* which represents the strategy profile of all players that satisfies:

$$(\forall J \in \mathcal{J} \text{ s.t } |J| \leq t), \forall s'_J \in S_J, \forall j \in J : u_j(\mathbf{s}^*) + \epsilon \geq u_j\left(\mathbf{s}'_J, \mathbf{s}^*_{-J}\right)$$

Theorem B.5.1. *The problem of t-resilient Nash for 3-player games is NP-complete.*

Proof. We need to add the following strategies to $\mathcal{SV}(\phi, \epsilon)$. We name the modified game $\mathcal{G}(\phi, \epsilon)$ where the strategy set is now $S = L \cup C \cup V \cup \{f\}$ similar to [121]. We also assume that there exists a third player that has the same payoff on any strategies in S.

- $u_1(x, f, y) = u_2(f, x, y) = 0$ for all $x \in S - \{f\}$ and for all $y \in S$;
- $u_1(f, f, y) = u_2(f, f, y) = 0$ for all $y \in S$;
- $u_1(f, x, y) = u_2(x, f, y) = n - 1$ for all $x \in S - \{f\}$ and for all $y \in S$;
- $u_3(x, x', y) = 1$ for all $x, x', y \in S$.

Other rules can be modified to three players similarly. We now show that if ϕ is satisfiable, there is a uniform EXACT Nash equilibrium of $\mathcal{SV}(\phi, \epsilon)$ where the first and second players play $l^i \in L$ or $-l^i \in L$ (for all i) uniformly with probability $\frac{1}{n}$ getting the expected payoff $n - 1$ while the third player plays an arbitrary strategy. The proof is the same as [52, 169, 121]. Furthermore, this Nash equilibrium is an EXACT strong Nash equilibrium as well [52], and similarly can be shown that it is an EXACT 2-resilient Nash equilibrium. This is followed by the fact that there exists no other EXACT Nash equilibrium other than this equilibrium unless both the first and the second players play f [52].

Now suppose that ϕ is not satisfiable. We can show that in any ϵ-approximate Nash equilibrium, at least one player always receives an expected payoff of less than $n - 1 - \epsilon$, which causes one of the first two players to deviate to f. According to [169], for the approximate version, we need to make the following modification to the correspondence between literals and truth assignments compared to [52]. We consider a literal to be *true* if it is played more often than its negation. If an assignment does not satisfy the formula, there is at least one clause that does not have a satisfying literal. Any of the first two players will turn to this clause strategy to receive a payoff of n whenever the opponent plays a literal that is not in that clause. We know that in the game $\mathcal{SV}$ the second player plays literals with probability $> 1-\epsilon$ by Lemma B.4.3, and there are only 3 literals in each clause each of which the second player plays with probability $\leq \frac{1}{n^2}$ by Lemma B.4.5. By changing the strategy to this clause, the first player will receive at least $\left(1 - \epsilon - \frac{3}{n^2}\right) n > n - 1 + 2\epsilon$. The proofs can be extended to the case where the game includes the strategy f and a third trivial player exists if both the first and the second player do not use f (see [121]). So either the first player can do ϵ better by changing his strategy, or he is already receiving $n-1+2\epsilon$, and so the second player does not have a guaranteed payoff of at least $n-1-\epsilon$. Reaching the payoff less than $n-1-\epsilon$ will cause either of the first two players to deviate to f. If either the first or the second player uses f, no approximate 2-resilient Nash could be formed as the coalition of the first, and the second player can always gain more than 0 (particularly, $\frac{n}{2} > \epsilon$ for each player) for either of the players. □

Bibliography

[1] Ittai Abraham, Danny Dolev, Rica Gonen, and Joe Halpern. Distributed Computing Meets Game Theory. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing - PODC '06*, New York, New York, USA, 2006. ACM Press.

[2] Ittai Abraham, Danny Dolev, and Joseph Y Halpern. Lower Bounds on Implementing Robust and Resilient Mediators. In *Theory of Cryptography*, pages 302–319. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[3] Michele Aghassi and Dimitris Bertsimas. Robust game theory. *Mathematical Programming*, 107(1-2):231–273, February 2006.

[4] E Altman and E Solan. Constrained Games: The Impact of the Attitude to Adversary's Constraints. *IEEE Trans. Automat. Contr.*, 54(10):2435–2440, October 2009.

[5] Ioannis Anagnostides, Gabriele Farina, Tuomas Sandholm, and Brian Hu Zhang. A Polynomial-Time Algorithm for Variational Inequalities under the Minty Condition. *arXiv preprint arXiv:2508.03432*, 2025.

[6] Nir Andelman, Michal Feldman, and Yishay Mansour. Strong price of anarchy. *Games and Economic Behavior*, 65(2):289–317, 2009.

[7] Kenneth J. Arrow. An Extension of the Basic Theorems of Classical welfare Economics. 1951.

[8] Qamar Askari, Irfan Younas, and Mehreen Saeed. Political Optimizer: A novel socio-inspired meta-heuristic for global optimization. *Knowledge-Based Systems*, 195:105709, 2020.

[9] Robert Aumann and Adam Brandenburger. Epistemic Conditions for Nash Equilibrium. *Econometrica*, 63(5):1161–1180, 1995.

[10] Robert J. Aumann. *16. Acceptable Points in General Cooperative n-Person Games*, page 287–324. Princeton University Press, December 1959.

[11] Robert J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

[12] Robert J Aumann. What is game theory trying to accomplish? In *Frontiers of Economics, edited by K. Arrow and S. Honkapohja*. Citeseer, 1985.

[13] Per Austrin, Mark Braverman, and Eden Chlamtáč. Inapproximability of NP-Complete Variants of Nash Equilibrium. In Leslie Ann Goldberg, Klaus Jansen, R. Ravi, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 13–25, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[14] Azar, Adel, Seyed mirzaee, and Seyed moslem. Providing new meta-heuristic algorithm for optimization problems inspired by humans' behavior to improve their positions. *International Journal of Artificial Intelligence & Applications*, 4(1):1–12, January 2013.

[15] Julius B. Barbanel. Super Envy-free Cake Division and Independence of Measures. *Journal of Mathematical Analysis and Applications*, 197(1):54–60, 1996.

[16] Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, and Shivika Narang. Fair Division of Indivisible Goods Among Strategic Agents, 2019.

[17] Paul I Barton, Kamil A Khan, Peter Stechlinski, and Harry A J Watson. Computationally relevant generalized derivatives: theory, evaluation and applications. *Optimization Methods and Software*, 33(4-6):1030–1072, 2018.

[18] Tamer Başar and Geert Jan Olsder. *Dynamic Noncooperative Game Theory, 2nd Edition*. Society for Industrial and Applied Mathematics, jan 1998.

[19] Heinz H. Bauschke, Regina S. Burachik, Patrick L. Combettes, Veit Elser, D. Russell Luke, and Henry Wolkowicz, editors. *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*. Springer New York, 2011.

[20] Omar Ben-Ayed. Bilevel linear programming. *Computers & Operations Research*, 20(5):485–501, 1993.

[21] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23:769–805, 1998.

[22] A. Ben-Tal and A. Nemirovski. Robust solutions of uncertain linear programs. *Operations Research Letters*, 25:1–13, 1999.

[23] A. Ben-Tal and A. Nemirovski. Selected topics in robust convex optimization. *Mathematical Programming*, 112:125–158, 2008.

[24] A. Bensoussan and J. L. Lions. Nouvelles Methodes en Controle Impulsionnel. *Applied Mathematics & Optimization*, 1(4):289–312, December 1975.

[25] A. Bensoussan and J. L. Lions. Problemes de temps D'arret Optimaux et de Perturbations Singulieres Dans Les Inequations Variationnelles. In *Lecture Notes in Economics and Mathematical Systems*, pages 567–584. Springer Berlin Heidelberg, 1975.

[26] Alain Bensoussan and Jacques Louis Lions. Nouvelle formulation de problèmes de contrôle impulsionnel et applications. 1973.

[27] Alain Bensoussan and Jacques Louis Lions. Contrôle impulsionnel et inéquations quasi variationnelles. 1982.

[28] Martino Bernasconi, Matteo Castiglioni, Andrea Celli, and Gabriele Farina. On the Role of Constraints in the Complexity of Min-Max Optimization, 2024.

[29] B.Douglas Bernheim, Bezalel Peleg, and Michael D Whinston. Coalition-Proof Nash Equilibria i. Concepts. *Journal of Economic Theory*, 42(1):1–12, 1987.

[30] B.Douglas Bernheim and Michael D Whinston. Coalition-Proof Nash Equilibria ii. Applications. *Journal of Economic Theory*, 42(1):13–29, 1987.

[31] Marie Louisa Tølbøll Berthelsen and Kristoffer Arnsfelt Hansen. On the Computational Complexity of Decision Problems About Multi-player Nash Equilibria. *Theory Comput. Syst.*, 66(3):519–545, 2022.

[32] P. Bhattacharyya and V. Vetrivel. An Existence Theorem on Generalized Quasi-Variational Inequality Problem. *Journal of Mathematical Analysis and Applications*, 188(2):610–615, 1994.

[33] John Birge and Francois Louveaux. *Introduction to Stochastic Programming*. Springer Series in Operations Research and Financial Engineering. Springer, New York, 2nd edition, 2011.

[34] Mohammad Bodaghi and Koosha Samieefar. Meta-heuristic bus transportation algorithm. *Iran Journal of Computer Science*, 2(1):23–32, October 2018.

[35] Eleonora Braggion, Nicola Gatti, Roberto Lucchetti, and Tuomas Sandholm. Strong Nash equilibria and mixed strategies, 2015.

[36] Eleonora Braggion, Nicola Gatti, Roberto Lucchetti, Tuomas Sandholm, and Bernhard von Stengel. Strong Nash equilibria and mixed strategies. *International Journal of Game Theory*, 49(3):699–710, July 2020.

[37] Simina Brânzei, Ioannis Caragiannis, David Kurokawa, and Ariel Procaccia. An Algorithmic Framework for Strategic Fair Division. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1), February 2016.

[38] Simina Brânzei, Ioannis Caragiannis, David Kurokawa, and Ariel Procaccia. An Algorithmic Framework for Strategic Fair Division. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1), February 2016.

[39] Simina Brânzei, Vasilis Gkatzelis, and Ruta Mehta. Nash Social Welfare Approximation for Strategic Agents. *Operations Research*, 70(1):402–415, 2022.

[40] Simina Brânzei and Peter Bro Miltersen. Equilibrium Analysis in Cake Cutting. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems*, AAMAS '13, page 327–334, Richland, SC, 2013. International Foundation for Autonomous Agents and Multiagent Systems.

[41] Colin F. Camerer, Teck-Hua Ho, and Juin-Kuan Chong. A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics*, 119(3):861–898, 08 2004.

[42] Sofia Ceppi, Nicola Gatti, Giorgio Patrini, and Marco Rocco. Local Search Methods for Finding a Nash Equilibrium in Two-Player Games. In *2010*

IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology, volume 2, pages 335–342, 2010.

[43] D. Chan and J. S. Pang. The Generalized Quasi-Variational Inequality Problem. *Mathematics of Operations Research*, 7(2):211–222, 1982.

[44] A. Charnes. Constrained Games and Linear Programming. *Proceedings of the National Academy of Sciences*, 39(7):639–641, 1953.

[45] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the Complexity of Computing Two-Player Nash equilibria. *J. ACM*, 56(3):1–57, May 2009.

[46] Yihsu Chen, Benjamin F. Hobbs, Sven Leyffer, and Todd S. Munson. Leader-Follower Equilibria for Electric Power and NO x Allowances markets. *Computational Management Science*, 3(4):307–330, June 2006.

[47] S. Chan Choi, Wayne S. Desarbo, and Patrick T. Harker. Product Positioning Under Price Competition. *Management Science*, 36(2):175–199, feb 1990.

[48] Juin-Kuan Chong, Teck-Hua Ho, and Colin Camerer. A generalized cognitive hierarchy model of games. *Games and Economic Behavior*, 99:257–274, 2016.

[49] Richard Cole and Vasilis Gkatzelis. Approximating the Nash Social Welfare with Indivisible Items. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '15, page 371–380, New York, NY, USA, 2015. Association for Computing Machinery.

[50] Antonio Conejo, Miguel Carrion, and Juan Morales. *Decision Making under Uncertainty in Electricity Markets*. Springer Verlag, New York, 2010.

[51] Stefano Coniglio, Nicola Gatti, and Alberto Marchesi. Computing a Pessimistic Stackelberg Equilibrium with Multiple Followers: The Mixed-Pure Case. *Algorithmica*, 82(5):1189–1238, November 2019.

[52] Vincent Conitzer and Tuomas Sandholm. New Complexity Results about Nash Equilibria. *Games Econ. Behav.*, 63(2):621–641, July 2008.

[53] H. W. Corley. Normative Utility Models for Pareto Scalar equilibria in n-person, Semi-Cooperative Games in Strategic Form. *Theoretical Economics Letters*, 07(06):1667–1686, 2017.

[54] Miguel A. Costa-Gomes and Vincent P. Crawford. Cognition and Behavior in Two-Person Guessing games: An Experimental Study. *American Economic Review*, 96(5):1737–1768, December 2006.

[55] G. P. Crespi, D. Radi, and M. Rocca. Insights on the Theory of Robust Games. *Computational Economics*, December 2023.

[56] Artur Czumaj, Argyrios Deligkas, Michail Fasoulakis, John Fearnley, Marcin Jurdziński, and Rahul Savani. Distributed Methods for Computing Approximate Equilibria. *Algorithmica*, 81(3):1205–1231, March 2019.

[57] Constantinos Daskalakis. Non-concave Games: A Challenge for Game Theory's Next 100 Years. 2022.

[58] Constantinos Daskalakis, Alex Fabrikant, and Christos H. Papadimitriou. *The Game World Is Flat: The Complexity of Nash Equilibria in Succinct Games*, page 513–524. Springer Berlin Heidelberg, 2006.

[59] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The Complexity of Computing a Nash Equilibrium. *SIAM j. comput.*, 39(1):195–259, January 2009.

[60] Constantinos Daskalakis, Aranyak Mehta, and Christos Papadimitriou. A Note on Approximate Nash Equilibria. *Theor. Comput. Sci.*, 410(17):1581–1588, April 2009.

[61] Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. The Complexity of Constrained Min-Max Optimization. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1466–1478, 2021.

[62] Gerard Debreu. A Social Equilibrium Existence Theorem. *Proceedings of the National Academy of Sciences*, 38(10):886–893, 1952.

[63] Mohammad Dehghani, Eva Trojovská, and Pavel Trojovský. A new human-based metaheuristic algorithm for solving optimization problems on the base of simulation of driving training process. *Scientific reports*, 12(1):9924, June 2022.

[64] Argyrios Deligkas, John Fearnley, and Rahul Savani. Inapproximability Results for Approximate Nash Equilibria. In Yang Cai and Adrian Vetta, editors, *Web*

and Internet Economics, pages 29–43, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.

[65] Argyrios Deligkas, John Fearnley, and Rahul Savani. Computing constrained approximate equilibria in polymatrix games. In *Algorithmic Game Theory*, pages 93–105. Springer International Publishing, 2017.

[66] V. DeMiguel and H. F. Xu. A stochastic multiple-leader stackelberg model: Analysis, computation, and application. *Operations Research*, 57:1220–1235, 2009.

[67] Andreas Ehrenmann. *Equilibrium problems with equilibrium constraints and their applications in electricity markets.* PhD thesis, Cambridge University, 2005.

[68] Samuel Eilenberg and Deane Montgomery. Fixed Point Theorems for Multi-Valued Transformations. *American Journal of Mathematics*, 68(2):214–222, 1946.

[69] Niels Elgers, Nguyen Dang, and Patrick De Causmaecker. *A Metaheuristic Approach to Compute Pure Nash Equilibria*, pages 221–233. Springer International Publishing, Cham, 2019.

[70] Kousha Etessami and Mihalis Yannakakis. On the Complexity of Nash Equilibria and other Fixed Points. *SIAM j. comput.*, 39(6):2531–2597, January 2010.

[71] Francisco Facchinei, Andreas Fischer, and Veronica Piccialli. On generalized Nash games and variational inequalities. *Oper. Res. Lett.*, 35(2):159–164, March 2007.

[72] Francisco Facchinei and Christian Kanzow. Generalized Nash equilibrium problems. *4OR*, 5(3):173–210, September 2007.

[73] Francisco Facchinei and Christian Kanzow. Generalized Nash equilibrium problems. *Ann. Oper. Res.*, 175(1):177–211, March 2010.

[74] Francisco Facchinei and Jong-Shi Pang, editors. *Finite-Dimensional Variational Inequalities and Complementarity Problems.* Springer New York, 2004.

[75] S. C. Fang and E. L. Peterson. Generalized variational inequalities. *Journal of Optimization Theory and Applications*, 38(3):363–383, November 1982.

[76] John Fearnley, Paul W. Goldberg, Alexandros Hollender, and Rahul Savani. The complexity of gradient descent: CLS = PPAD ∩ PLS. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. ACM, June 2021.

[77] M Ferris, S Dirkse, and A Meeraus. Mathematical programs with equilibrium constraints: automatic reformulation and solution via constrained optimization. Technical report, 2002.

[78] Aris Filos-Ratsikas, Kristoffer A. Hansen, Kasper Høgh, and Alexandros Hollender. FIXP-Membership via Convex Optimization: Games, Cakes, and Markets. *SIAM Journal on Computing*, 0(0):FOCS21–30–FOCS21–84, 0.

[79] Aris Filos-Ratsikas, Kristoffer Arnsfelt Hansen, Kasper Høgh, and Alexandros Hollender. PPAD-membership for problems with Exact Rational Solutions: A General Approach via Convex Optimization. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC '24. ACM, June 2024.

[80] Steven A. Gabriel, Antonio J. Conejo, J. David Fuller, Benjamin F. Hobbs, and Carlos Ruiz. Equilibrium Problems with Equilibrium Constraints. In *Complementarity Modeling in Energy Markets*, pages 263–321. Springer New York, June 2012.

[81] Dian Gadjov and Lacra Pavel. Resilient Nash equilibrium Seeking in the Partial Information Setting, 2021.

[82] Jugal Garg, Ruta Mehta, Vijay V. Vazirani, and Sadra Yazdanbod. ETR-completeness for decision versions of multi-player (symmetric) nash equilibria. In *Automata, Languages, and Programming*, pages 554–566. Springer Berlin Heidelberg, 2015.

[83] Nicola Gatti, Marco Rocco, and Tuomas Sandholm. Algorithms for Strong Nash Equilibrium with More than Two Agents. *Proceedings of the AAAI Conference on Artificial Intelligence*, 27(1):342–349, June 2013.

[84] Nicola Gatti, Marco Rocco, and Tuomas Sandholm. Strong Nash Equilibrium Is in Smoothed P. In *AAAI Conference on Artificial Intelligence*, 2013.

[85] Nicola Gatti, Marco Rocco, and Tuomas Sandholm. On the Verification and Computation of Strong nash equilibrium, 2017.

[86] Gamow George and Marvin Stern. *Puzzle-math the Big Bang to Black Holes.* Viking Press, 1958.

[87] Abheek Ghosh and Alexandros Hollender. The Complexity of Symmetric Bimatrix Games with Common Payoffs. In *Proceedings of the 20th Conference on Web and Internet Economics (WINE 2024)*, page 5. Springer, 2024. Accepted and peer-reviewed.

[88] F. Giannessi and A. Maugeri, editors. *Variational Inequalities and Network Equilibrium Problems.* Springer US, 1995.

[89] Itzhak Gilboa and Eitan Zemel. Nash and Correlated Equilibria: Some Complexity Considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.

[90] Fred Glover. Future paths for integer programming and links to artificial intelligence. *Computers & Operations Research*, 13(5):533–549, 1986. Applications of Integer Programming.

[91] Fred Glover and Manuel Laguna. Tabu Search. In Ding-Zhu Du and Panos M. Pardalos, editors, *Handbook of Combinatorial Optimization: Volume1–3*, pages 2093–2229. Springer US, Boston, MA, 1998.

[92] Denizalp Goktas and Amy Greenwald. Tractable General Equilibrium. *arXiv preprint arXiv:2502.11449*, 2025.

[93] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2. Springer Science & Business Media, 2012.

[94] Mathias Hall-Andersen and Nikolaj I. Schwartzbach. Game Theory on the Blockchain: A Model for Games with Smart Contracts. In Ioannis Caragiannis and Kristoffer Arnsfelt Hansen, editors, *Algorithmic Game Theory*, pages 156–170, Cham, 2021. Springer International Publishing.

[95] J. Y. Halpern and R. Pass. Game theory with costly computation. Unpublished manuscript, 2008.

[96] Joseph Y. Halpern. Beyond Nash equilibrium: solution concepts for the 21st century. In *Proceedings of the Twenty-Seventh ACM Symposium on Principles of Distributed Computing*, PODC '08, page 1–10, New York, NY, USA, 2008. Association for Computing Machinery.

[97] Patrick T. Harker. Generalized Nash games and quasi-variational inequalities. *European Journal of Operational Research*, 54(1):81–94, 1991.

[98] Patrick T. Harker and Jong-Shi Pang. Finite-dimensional variational inequality and nonlinear complementarity problems: A survey of theory, algorithms and applications. *Mathematical Programming*, 48(1–3):161–220, March 1990.

[99] C Harsanyi. Games with incomplete information played by "bayesian" players, part i. the basic model. *Management Science*, 14:159–182, 1967.

[100] J. C. Harsanyi. Games with incomplete information played by "bayesian" players, part ii. bayesian equilibrium points. *Management Science*, 14:320–340, 1968.

[101] J. C. Harsanyi. Games with incomplete information played by "bayesian" players, part iii. the basic probability distribution of the game. *Management Science*, 14:486–502, 1968.

[102] Jaroslav Haslinger, Markku Miettinen, and Panagiotis D. Panagiotopoulos. *Finite Element Method for Hemivariational Inequalities*. Springer US, 1999.

[103] Shunsuke Hayashi, Nobuo Yamashita, and Masao Fukushimay. Robust Nash Equilibria and Second-Order Cone Complementarity Problems. *Journal of Nonlinear and Convex Analysis*, 6, 01 2005.

[104] B.E. Hobbs. Linear complementarity models of nash-cournot competition in bilateral and poolco power markets. *IEEE Transactions on Power Systems*, 16(2):194–202, 2001.

[105] B.F. Hobbs, C.B. Metzler, and J.-S. Pang. Strategic gaming analysis for electric power systems: an MPEC approach. *IEEE Transactions on Power Systems*, 15(2):638–645, 2000.

[106] Ming Hu and Masao Fukushima. Variational Inequality Formulation of a Class of Multi-Leader-Follower Games. *Journal of Optimization Theory and Applications*, 151(3):455–473, August 2011.

[107] Ming Hu and Masao Fukushima. Existence, Uniqueness, and Computation of Robust Nash Equilibria in a Class of Multi-Leader-Follower Games. *SIAM Journal on Optimization*, 23(2):894–916, 2013.

[108] Ming Hu and Masao Fukushima. MULTI-LEADER-FOLLOWER GAMES: MODELS, METHODS AND APPLICATIONS. *Journal of the Operations Research Society of Japan*, 58(1):1–23, 2015.

[109] Xinmin Hu. *Athematical programs with complementarity constraints and game theory models in electricity markets*. PhD thesis, university of Melbourne, 2003.

[110] Xinmin Hu and Daniel Ralph. Using EPECs to Model Bilevel Games in Restructured Electricity Markets with Locational Prices. *Operations Research*, 55(5):809–827, October 2007.

[111] Tatsuro Ichiishi. A Social Coalitional Equilibrium Existence Lemma. *Econometrica*, 49(2):369, March 1981.

[112] Gabriela Jeronimo, Daniel Perrucci, and Juan Sabia. A parametric representation of totally mixed Nash equilibria. *Comput. Math. Appl.*, 58(6):1126–1141, September 2009.

[113] Houyuan Jiang, Serguei Netessine, and Sergei Savin. Robust Newsvendor Competition Under Asymmetric Information. *Operations Research*, 59(1):254–261, 2011.

[114] Alejandro Jofré, R. Terry Rockafellar, and Roger J-B. Wets. Variational Inequalities and Economic Equilibrium. *Mathematics of Operations Research*, 32(1):32–50, 2007.

[115] Shizuo Kakutani. A generalization of brouwer's fixed point theorem. *Duke Mathematical Journal*, 8(3):457 – 459, 1941.

[116] Mamoru Kaneko and Kenjiro Nakamura. The Nash Social Welfare Function. *Econometrica*, 47(2):423–435, 1979.

[117] Irving Kaplansky. A Contribution to Von Neumann's Theory of Games. *Annals of Mathematics*, 46(3):474–479, 1945.

[118] Irving Kaplansky. A contribution to Von Neumann's Theory of Games. II. *Linear Algebra Appl.*, 226-228:371–373, September 1995.

[119] Bruce M. Kapron and Koosha Samieefar. The Computational Complexity of Variational Inequalities and Applications in Game Theory, 2024.

[120] Bruce M. Kapron and Koosha Samieefar. On the Computational Complexity of Quasi-Variational Inequalities and Multi-Leader-Follower Games. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '24, page 2324–2326, Richland, SC, 2024. International Foundation for Autonomous Agents and Multiagent Systems.

[121] Bruce M. Kapron and Koosha Samieefar. The Computational Complexity of Equilibria with Strategic Constraints. In Rastislav Královič and Věra Kůrková, editors, *SOFSEM 2025: Theory and Practice of Computer Science*, Lecture Notes in Computer Science, Bratislava, Slovak Republic, January 20–23 2025. Springer.

[122] Bruce M. Kapron and Koosha Samieefar. On The Computational Complexity of Games with Uncertainty. In Irene Finocchi and Loukas Georgiadis, editors, *Algorithms and Complexity*, pages 343–360, Cham, 2025. Springer Nature Switzerland.

[123] Erim Kardeş, Fernando Ordóñez, and Randolph W. Hall. Discounted Robust Stochastic Games and an Application to Queueing Control. *Operations Research*, 59(2):365–382, 2011.

[124] Michael J. Kearns, Michael L. Littman, and Satinder P. Singh. Graphical Models for Game Theory. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, UAI '01, page 253–260, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.

[125] G. Kol and M. Naor. Cryptography and game theory: Designing protocols for exchanging information. In *Theory of Cryptography Conference*, pages 320–339. Springer, 2008.

[126] Michal Kočvara and Jiří V. Outrata. On a class of quasi-variational inequalities. *Optimization Methods and Software*, 5(4):275–295, 1995.

[127] C. E. Lemke and J. T. Howson, Jr. Equilibrium Points of Bimatrix Games. *Journal of the Society for Industrial and Applied Mathematics*, 12(2):413–423, 1964.

[128] Sven Leyffer and Todd Munson. Solving multi-leader–common-follower games. *Optimization Methods and Software*, 25(4):601–623, aug 2010.

[129] G. H. Lin and M. Fukushima. Stochastic equilibrium problems and stochastic mathematical programs with equilibrium constraints: A survey. *Pacific Journal of Optimization*, 6:455–482, 2010.

[130] Richard J Lipton, Evangelos Markakis, and Aranyak Mehta. Playing Large Games Using Simple Strategies. In *Proceedings of the 4th ACM conference on Electronic commerce - EC '03*, New York, New York, USA, 2003. ACM Press.

[131] Zhi-Quan Luo, Jong-Shi Pang, and Daniel Ralph. *Mathematical Programs with Equilibrium Constraints*. Cambridge University Press, November 1996.

[132] Alberto Marchesi. Leadership Games: Multiple Followers, Multiple Leaders, and perfection. In Angelo Geraci, editor, *Special Topics in Information Technology*, pages 107–118. Springer International Publishing, Cham, 2021.

[133] Eric Maskin. Nash Equilibrium and Welfare Optimality. *The Review of Economic Studies*, 66(1):23–38, 1999.

[134] J.C.C. McKinsey. *Introduction to the Theory of Games*. McGraw-Hill Book Co., 1952.

[135] Andrew McLennan and Rabee Tourky. Simple complexity from imitation games. *Games Econ. Behav.*, 68(2):683–688, March 2010.

[136] Emerson Melo. A Variational Approach to Network Games. *SSRN Electronic Journal*, 2018.

[137] Monica Milasi. Existence theorem for a class of generalized quasi-variational inequalities. *Journal of Global Optimization*, 60(4):679–688, October 2013.

[138] Igal Milchtaich. Computation of Completely Mixed Equilibrium Payoffs in Bimatrix Games. *Int. Game Theory Rev.*, 08(03):483–487, September 2006.

[139] D. Motreanu and P. D. Panagiotopoulos. *Minimax Theorems and Qualitative Properties of the Solutions of Hemivariational Inequalities*. Springer US, 1999.

[140] Hervé Moulin and Jean-Pierre Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3-4):201–221, 1978.

[141] Rosemarie Nagel. Unraveling in Guessing Games: An Experimental Study. *The American Economic Review*, 85(5):1313–1326, 1995.

[142] Parinaz Naghizadeh and Mingyan Liu. On the uniqueness and stability of equilibria of network games. In *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, page 280–286. IEEE Press, 2017.

[143] Ahmad Nahhas and H. W. Corley. An Alternative Interpretation of Mixed Strategies in n-person Normal Form Games via Resource Allocation. *Theoretical Economics Letters*, 08(10):1854–1868, 2018.

[144] John Nash. Non-Cooperative Games. *Annals of Mathematics*, 54(2):286–295, 1951.

[145] John F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, January 1950.

[146] Rabia Nessah and Guoqiang Tian. On the existence of strong nash equilibria. *Journal of Mathematical Analysis and Applications*, 414(2):871–885, 2014.

[147] R. Nishimura, S. Hayashi, and M. Fukushima. Robust Nash equilibria in N-person non-cooperative games: Uniqueness and reformulations. *Pacific Journal of Optimization*, 5:237–259, 2009.

[148] Jiři Outrata, Michal Kočvara, and Jochem Zowe. *Nonsmooth Approach to Optimization Problems with Equilibrium Constraints*. Springer US, 1998.

[149] Jiří V. Outrata. A note on a class of equilibrium problems with equilibrium constraints. *Kybernetika*, 40(5):[585]–594, 2004.

[150] Jong-Shi Pang and Masao Fukushima. Quasi-variational inequalities, generalized nash equilibria, and multi-leader-follower games. *Computational Management Science*, 2(1):21–56, January 2005.

[151] Jong-Shi Pang and Jen-Chih Yao. On a Generalization of a Normal Map and Equation. *SIAM Journal on Control and Optimization*, 33(1):168–184, 1995.

[152] Christos Papadimitriou, Emmanouil-Vasileios Vlatakis-Gkaragkounis, and Manolis Zampetakis. The Computational Complexity of Multi-player Concave games and Kakutani Fixed points. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC '23, page 1045, New York, NY, USA, 2023. Association for Computing Machinery.

[153] Christos H Papadimitriou. On the Complexity of the Parity Argument and Other Inefficient Proofs of Existence. *J. Comput. Syst. Sci.*, 48(3):498–532, June 1994.

[154] Francesca Parise and Asuman Ozdaglar. A variational inequality framework for network games: Existence, uniqueness, convergence and sensitivity analysis. *Games and Economic Behavior*, 114:47–82, 2019.

[155] T Parthasarathy, Vasudha Sharma, and A R Sricharan. Completely mixed bimatrix games. *Proc. Math. Sci.*, 130(1), December 2020.

[156] Bezalel Peleg. *Game Theoretic Analysis of Voting in Committees*. Cambridge University Press, July 1984.

[157] Benjamin Plaut. Optimal Nash Equilibria for Bandwidth Allocation. In Xujin Chen, Nikolai Gravin, Martin Hoefer, and Ruta Mehta, editors, *Web and Internet Economics*, pages 77–88, Cham, 2020. Springer International Publishing.

[158] T. E. S. Raghavan. Completely Mixed Strategies in Bimatrix Games. *Journal of the London Mathematical Society*, 2(Part_4):709–712, 1970.

[159] R. Tyrrell Rockafellar and Roger J. B. Wets. *Variational Analysis*. Springer Berlin Heidelberg, 1998.

[160] Anton Rodomanov and Yurii Nesterov. Subgradient ellipsoid method for non-smooth convex problems. *Mathematical Programming*, 199(1–2):305–341, June 2022.

[161] J. B. Rosen. Existence and Uniqueness of Equilibrium Points for Concave N-Person Games. *Econometrica*, 33(3):520–534, 1965.

[162] Tim Roughgarden. Intrinsic Robustness of the Price of Anarchy. *Journal of the ACM*, 62(5):1–42, November 2015.

[163] Ola Rozenfeld. Strong equilibrium in congestion games. Master's thesis, Technion–Israel Institute of Technology, 2007.

[164] Ariel Rubinstein. Comments on the Interpretation of Game Theory. *Econometrica*, 59(4):909–924, 1991.

[165] Aviad Rubinstein. Inapproximability of Nash Equilibrium. In *Proceedings of the forty-seventh annual ACM symposium on Theory of Computing*, New York, NY, USA, June 2015. ACM.

[166] Aviad Rubinstein. Settling the Complexity of Computing Approximate Two-Player Nash Equilibria. *SIGecom Exch.*, 15(2):45–49, feb 2017.

[167] Koosha Samieefar. A Meta-hearustic Approach for Strategic Fair Division. In *In proceedings of the 7th International Conference on Intelligent Systems, Metaheuristics & Swarm Intelligence*, Kuala Lumpor, Malaysia, 2023. ACM.

[168] Larry Samuelson. Game Theory in Economics and Beyond. *Journal of Economic Perspectives*, 30(4):107–130, November 2016.

[169] Grant R Schoenebeck and Salil Vadhan. The Computational Complexity of Nash Equilibria in Concisely Represented Games. *ACM trans. comput. theory*, 4(2):1–50, May 2012.

[170] Wayne Shafer and Hugo Sonnenschein. Equilibrium in abstract economies without ordered preferences. *Journal of Mathematical Economics*, 2(3):345–348, 1975.

[171] U. V. Shanbhag, G. Infanger, and P. W. Glynn. A complementarity framework for forward contracting under uncertainty. *Operations Research*, 59:810–834, 2011.

[172] Hanif D. Sherali. A multiple Leader Stackelberg Model and Analysis. *Operations Research*, 32(2):390–404, apr 1984.

[173] Dale O. Stahl and Paul W. Wilson. Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization*, 25(3):309–327, 1994.

[174] Bernhard von Stengel. Equilibrium Computation for Two-Player Games in Strategic and Extensive form. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V.Editors Vazirani, editors, *Algorithmic Game Theory*, page 53–78. Cambridge University Press, 2007.

[175] Che-Lin Su. *EQUILIBRIUM PROBLEMS WITH EQUILIBRIUM CONSTRAINTS:STATIONARITIES, ALGORITHMS, AND APPLICATIONS*. PhD thesis, STANFORD UNIVERSITY, 2005.

[176] Che-Lin Su. Analysis on the forward market equilibrium model. *Operations Research Letters*, 35(1):74–82, 2007.

[177] Koichi Tadenuma and William Thomson. Games of Fair Division. *Games and Economic Behavior*, 9(2):191–204, 1995.

[178] Eva Trojovská and Mohammad Dehghani. A new human-based metahurestic optimization method based on mimicking cooking training. *Scientific Reports*, 12(1), September 2022.

[179] Luis N. Vicente and Paul H. Calamai. Bilevel and multilevel programming: A bibliography review. *Journal of Global Optimization*, 5(3):291–306, October 1994.

[180] Heinrich von Stackelberg. *The Theory of Market Economy*. Oxford University Press, 1952.

[181] Neha S. Wadia, Yatin Dandi, and Michael I. Jordan. A Gentle Introduction to Gradient-Based Optimization and Variational Inequalities for Machine Learning, 2024.

[182] Ling Wang, Haoqi Ni, Ruixin Yang, Minrui Fei, and Wei Ye. A Simple Human Learning Optimization Algorithm. In Minrui Fei, Chen Peng, Zhou Su, Yang Song, and Qinglong Han, editors, *Computational Intelligence, Networked Systems and Their Applications*, pages 56–65, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.

[183] William A. Webb. An Algorithm For Super Envy-free Cake Division. *Journal of Mathematical Analysis and Applications*, 239(1):175–179, 1999.

[184] Jonathan Widger and Daniel Grosu. Computing Equilibria in Bimatrix Games by Parallel Support Enumeration. In *2008 International Symposium on Parallel and Distributed Computing*. IEEE, 2008.

[185] Andrzej Wieczorek. *Constrained and indefinite games and their applications*. Instytut Matematyczny Polskiej Akademi Nauk, 1983.

[186] Jen-Chih Yao. The generalized quasi-variational inequality problem with applications. *Journal of Mathematical Analysis and Applications*, 158(1):139–160, 1991.

[187] Brian Zhang, Ioannis Anagnostides, Emanuel Tewolde, Ratip Emin Berker, Gabriele Farina, Vincent Conitzer, and Tuomas Sandholm. Expected Variational Inequalities. *arXiv preprint arXiv:2502.18605*, 2025.

[188] Brian Hu Zhang. *New Solution Concepts and Algorithms for Equilibrium Computation and Learning in Extensive-Form Games and Beyond*. Phd thesis, Carnegie Mellon University, Pittsburgh, PA, August 2025.

[189] Kaiqing Zhang, TAO SUN, Yunzhe Tao, Sahika Genc, Sunil Mallya, and Tamer Basar. Robust Multi-Agent Reinforcement Learning with Model Uncertainty. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 10571–10583. Curran Associates, Inc., 2020.