

On the Existence of Low-Rank Explanations for Mixed Strategy Behavior

Siddharth Barman* Umang Bhaskar † Federico Echenique ‡ Adam Wierman §

Abstract

Nash equilibrium is used as a model to explain the observed behavior of players in strategic settings. For example, in many empirical applications we observe player behavior, and the problem is to determine if there exist payoffs for the players for which the equilibrium corresponds to observed player behavior. Computational complexity of Nash equilibria is an important consideration in this framework. If the instance of the model that explains observed player behavior requires players to have solved a computationally hard problem, then the explanation provided is questionable. In this paper we provide conditions under which Nash equilibrium is a reasonable explanation for strategic behavior, i.e., conditions under which observed behavior of players can be explained by games in which Nash equilibria are easy to compute. We identify three structural conditions and show that if the data set of observed behavior satisfies any of these conditions, then it is consistent with payoff matrices for which the observed Nash equilibria could have been computed efficiently. Our conditions admit large and structurally complex data sets of observed behavior, showing that even with complexity considerations, Nash equilibrium is often a reasonable model.

1 Introduction

The computational complexity of equilibria in economic models is at the core of recent research in algorithmic game theory. In general, the basic message of this research is negative: Computing Nash equilibria is PPAD-complete even for 2-player games [10], and computing Walrasian equilibria is PPAD-hard [9, 27]. The fact that the standard notions of equilibrium used by economists are hard to compute raises a concern that they are flawed models for economic behavior, because they lack a plausible rationale for how one would arrive at equilibrium.

These hardness results have motivated the study of instances in which equilibrium can be computed efficiently, e.g., [1, 13, 14]. These results — both positive and negative — assume a fixed and literal interpretation of economic models. For example, in two-player games the assumption is that payoff matrices are explicitly specified, and agents seek to maximize their payoffs given the strategy of the other player.

*Center for the Mathematics of Information, California Institute of Technology. barman@caltech.edu

†Center for the Mathematics of Information, California Institute of Technology. umang@caltech.edu

‡Humanities and Social Sciences, California Institute of Technology. fede@hss.caltech.edu

§Computing and Mathematical Sciences, California Institute of Technology. adamw@caltech.edu

However, models are frequently used as explanations of behavior, rather than literal descriptions. In empirical applications of game theory, for example, one rarely observes payoffs of the players. Rather, what we observe is the behavior of the players, and a first-order task is to determine if the model explains observed behavior or not (see, e.g., [7, 8] for classical examples of this exercise). In particular, a good model instance (i.e., payoffs specification) is one where the observed behavior of the players is as if they were playing equilibrium. This is certainly how economists use game theory: as a modeling tool to explain observed behavior.

This perspective however cannot ignore computational complexity. Even if a model explains observed behavior, it is unreasonable if it requires players to solve computationally hard problems. Thus, an important question that arises from the perspective of models as explanations of observed behavior and the computational complexity of equilibrium is: *when is an economic model a reasonable explanation for observed behavior?* We address this question in the context of Nash equilibrium as a model for two-player games, and computational complexity as a measure of reasonableness.

This framework of starting with observed behavior is consistent with much of economics, and is formalized in *revealed preference theory*, which was pioneered by Samuelson in 1938 [22] and has a long tradition in economics (see, e.g., [3, 23–26]). Classical revealed preference theory asks, “Does there exist an instance of the model that is consistent with the observed behavior?” If there is such a model instance, then the observed behavior is said to be consistent with the theory, and thus *rationalizable*. Our work augments the fundamental question of rationalizability in revealed preference theory with complexity considerations. The question we address is then, “Does there exist an instance of the model that is consistent with the observed behavior *and* for which the observations could have been computed efficiently?”

We focus on addressing this question in the context of bimatrix games. We consider a setting where mixed strategy behavior (i.e., probability distributions over the actions of the players) is observed. In particular, with multiple instances of mixed strategy behavior as input, we want to know if there is some tractable instance of the game such that the given observations correspond to Nash equilibria. We focus on low *player rank* — the minimum of the ranks of the payoff matrices of the players — as our notion of tractability given recent algorithmic results, e.g., [13]. We defer a discussion of our modeling assumptions and their relaxations to Section 6.

It is important to be clear about what we do not do. We do not address a problem of inference; we do not claim to estimate or back out any instance of the game. Our focus is on the problem of *testing*: When can a dataset be explained by equilibrium behavior for which there is a story (an efficient algorithm) for how players could have arrived at the equilibrium? The existence of such an explanation does not imply that the explanation is unique; indeed there may be other explanations with different properties, a common situation in revealed preference theory. Our results say that one can quite often find explanations for economic behavior which do not require agents to solve computationally hard problems. Hence in such settings Nash equilibria, even with complexity considerations, remains an applicable model of behavior.

Summary of results. Broadly, the results in this paper show that large and structurally complex data sets of observations — including data sets with overlapping observed strategies, each of which

may be an arbitrary distribution — can be rationalized by games that admit efficient computation of the observed Nash equilibria. Specifically, we identify three measures of structural complexity in data sets, and show that if a data set has a low value on any one of these measures for either player, or can be partitioned in a manner such that each partition has a low value on any one of these measures, then it has a rationalization with low player rank. The three measures we study are (i) the *dimensionality* of the observed strategies (Theorem 2), (ii) the *support size* of the observed strategies (Theorem 3), and (iii) the *chromatic number* of the data set (Theorem 4). We believe these are natural and complementary measures to evaluate the structural complexity of a data set, and our contribution is to show that each of these measures individually translates to the existence of rationalizing games with low player rank.

Finally, we also show that the bounds we obtain on the player rank for these measures are nearly tight by giving an example of a data set that necessitates high player rank for any rationalization (Theorem 5).

Related work. This paper seeks to incorporate computational complexity as a measure of reasonableness of models (which are used to explain observed behavior) in the context of bimatrix games. This general direction was initiated in the context of consumer choice theory [11]. In this context it was shown that a data set of n observations of a consumer choosing among d indivisible goods can always be explained by a consumer utility function that can be optimized in $O(nd)$ time. Thus, despite the fact that the consumer choice problem is NP-hard, consumer choice data is not rich enough to expose computationally hard utility functions without an exponentially large data set.

The question was also asked in the context of bimatrix games in [5]; however the focus in that work was on *pure strategy equilibria* and so differs from the current paper, which focuses on mixed strategy equilibria. In [5], the core problem is not computationally hard, and so the focus is on the structural complexity of the rationalizing game, as formalized via game rank.¹ In contrast, the purpose of the current paper is to develop a connection between the structural richness of the data and the *computational* complexity of the game.

This problem at hand is significantly more difficult and technically involved than the ones solved in [11] and [5]. To highlight this, note that, there are no results in the revealed preference literature characterizing which data sets can be explained as mixed Nash equilibria. So, our results represent a contribution to pure economic theory, as well as to algorithmic game theory.

2 Preliminaries

Bimatrix Games. Bimatrix games are two player games in normal form. Such games are specified by a pair of matrices (A, B) of size $n \times n$, which are termed the payoff matrices for the players. The first player, also called the row player, has payoff matrix A , and the second player, or the column player, has payoff matrix B . The strategy set for each player is $[n] = \{1, 2, \dots, n\}$, and, if the row

¹Game rank is defined to be the rank of the sum of the payoff matrices of the players. On the other hand, player rank is the minimum of the ranks of the payoff matrices of the players.

player plays strategy i and column player plays strategy j , then the payoffs of the two players are A_{ij} and B_{ij} respectively. The player rank of game (A, B) is defined to be $\min\{\text{rank}(A), \text{rank}(B)\}$. Our focus on player rank stems from a number of important properties. In particular, if a game has player rank $k \in \{1, 2, \dots, n\}$, then an equilibrium can be computed in time $O(n^{O(k)})$ [13, 15].

Let Δ^n be the set of probability distributions over the set of pure strategies $[n]$. For $x \in \Delta^n$, we define $\text{Supp}(x) := \{i : x_i > 0\}$. Further, $e_i \in \mathbb{R}^n$ is the vector with 1 in the i th coordinate and 0's elsewhere, and $u_k \in \Delta^n$ is the uniform distribution over the set $\{1, 2, \dots, k\}$. The players can randomize over their strategies by selecting any probability distribution in Δ^n , called a mixed strategy. When the row and column players play mixed strategies x and y respectively, the expected payoff of the row player is $x^T A y$ and the expected payoff of the column player is $x^T B y$.

Given a mixed strategy $y \in \Delta^n$ for the column player, the best-response set of the row player, β_r , is defined as $\beta_r(y) := \{i \in [n] \mid e_i^T A y \geq e_k^T A y \ \forall k \in [n]\}$. Similarly, the best-response set, β_c , of the column player (against mixed strategy $x \in \Delta^n$ of the row player) is defined as $\beta_c(x) := \{j \in [n] \mid x^T B e_j \geq x^T B e_k \ \forall k \in [n]\}$. The best response sets β_r and β_c are defined with respect to the payoff matrices A and B . When we want to emphasize this fact we use superscripts: β_r^A and β_c^B .

Definition 1 (Nash Equilibrium). *A pair of mixed strategies (x, y) , $x, y \in \Delta^n$, is a Nash equilibrium if and only if:*

$$\begin{aligned} x^T A y &\geq e_i^T A y && \forall i \in [n] \text{ and} \\ x^T B y &\geq x^T B e_j && \forall j \in [n]. \end{aligned}$$

The Nash equilibrium is strict if additionally the support and the best-response sets are equal, i.e., $\text{Supp}(x) = \beta_r(y)$ and $\text{Supp}(y) = \beta_c(x)$. Thus, for strict Nash equilibrium $x^T A y > e_i^T A y$ for all $i \notin \text{Supp}(x)$ and $x^T B y > x^T B e_j$ for all $j \notin \text{Supp}(y)$.

Observed behavior. We consider settings in which the payoff matrices A and B are not explicitly specified. Instead, we are given a collection of observed mixed strategy pairs. Our framework of observing mixed strategies is entirely analogous to the theory of individual stochastic choice: see for example [16–19]. The idea is that repeated observation of pure play allows one to infer a probability distribution over pure strategies.

A data set \mathcal{D} (of size m) is a collection of mixed-strategy pairs, $\mathcal{D} = \{(x_k, y_k) \in \Delta^n \times \Delta^n \mid k \in [m]\}$. We refer to mixed-strategy pairs (x_k, y_k) as observations.

We denote the set of observed mixed strategies of the row and column player in a data set \mathcal{D} by $\mathcal{O}_r(\mathcal{D})$ and $\mathcal{O}_c(\mathcal{D})$ respectively: $\mathcal{O}_r(\mathcal{D}) := \{x \in \Delta^n \mid \exists y \in \Delta^n \text{ such that } (x, y) \in \mathcal{D}\}$ and $\mathcal{O}_c(\mathcal{D}) := \{y \in \Delta^n \mid \exists x \in \Delta^n \text{ such that } (x, y) \in \mathcal{D}\}$. When there is a single data set under consideration, for ease of notation, we simply refer to these sets as \mathcal{O}_r and \mathcal{O}_c .

3 Warmup result: Rationalization

Given observed behavior (data) as described in the previous section, the first-order goal of classical revealed preference theory is to understand whether the data is rationalizable.

Definition 2 (Rationalizable Data). *A data set $\mathcal{D} = \{(x_k, y_k)\}_k$ is rationalizable if there exist payoff matrices A and B such that for all k , (x_k, y_k) is a strict Nash equilibrium in the game (A, B) .*

Strictness is required in the above definition in order to avoid rationalization by trivial games.

The first result of this paper comes from the observation that rationalizability of a given data set can be expressed as a linear program, and can thus be determined efficiently. In fact, the use of a linear program is a robust tool, and allows us to determine rationalizability in many more settings, such as:

- Rationalizability by payoff matrices that satisfy given linear constraints, e.g., enforcing values for certain payoffs, or inequalities between payoffs.
- Games with multiple players. Each player now has a m -dimensional payoff tensor, where m is the number of players.
- Data sets that consist of observed player behavior in subgames. Each observation thus consists of a tuple (I_k, J_k, x_k, y_k) , where $I_k, J_k \subseteq [n]$ are subsets of the available pure strategies and $\text{Supp}(x_k) \subseteq I_k, \text{Supp}(y_k) \subseteq J_k$.
- Rationalization of the data set as ϵ -approximate Nash equilibria.

For simplicity of notation, we give the linear program for the basic case of two players; this can easily be extended to any of the above cases. The variables in the linear program are the entries of the payoff matrices A and B , and the payoffs π_k, π'_k obtained by the players for each observation (x_k, y_k) in the data set.

Proposition 1. *A data set $\mathcal{D} = \{(x_k, y_k) \in \Delta^n \times \Delta^n \mid 1 \leq k \leq m\}$ is rationalizable iff the optimal value of (LP) is strictly greater than zero.*

$$\begin{aligned}
& \text{maximize} && \delta \\
& \text{subject to} && (Ay_k)_i = \pi_k && \forall k, \forall i \in \text{Supp}(x_k) \\
& && (Ay_k)_j \leq \pi_k - \delta && \forall k, \forall j \notin \text{Supp}(x_k) \\
& && (x_k^T B)_i = \pi'_k && \forall k, \forall i \in \text{Supp}(y_k) \\
& && (x_k^T B)_j \leq \pi'_k - \delta && \forall k, \forall j \notin \text{Supp}(y_k) \\
& && 0 \leq A_{i,j}, B_{i,j} \leq 1 && \forall i, j \in [n] \\
& && \delta \geq 0.
\end{aligned} \tag{LP}$$

Proof. If the optimal value of the linear program is strictly greater than zero then we have matrices A and B (corresponding to an optimal solution) under which the observations (x_k, y_k) are strict Nash equilibria. Hence \mathcal{D} is rationalizable.

On the other hand, if \mathcal{D} is rationalizable, say via game (\hat{A}, \hat{B}) , then we can scale the entries of \hat{A} and \hat{B} (by a large enough positive constant) and add a fixed number to all of them to obtain a (normalized) game (A, B) with entries between 0 and 1. Since such an affine transformation preserves the set of strict Nash equilibria, \mathcal{D} is also rationalized by (A, B) . In other words, (A, B) gives us a feasible solution to the linear program with objective function value strictly greater than zero. This establishes the claim. \square

4 Main Results: Rationalizations with Low Player Rank

This section includes the main results of the paper, which identify structural properties of data sets that guarantee the existence of rationalizations for which the observed Nash equilibria can be computed efficiently. To ensure that the observed equilibria can be computed efficiently in the rationalizations we construct rationalizations with low player rank and make use of the following result from [13].

Theorem 1 ([13]). *If the player rank of a bimatrix game is k then all extreme Nash equilibria can be computed in time $O(n^{O(k)})$.*

The algorithm implicit in this theorem computes all extreme equilibria, and so it can be used to compute the observations in the data set rather than simply some arbitrary equilibria. This fact is crucial to the exercise, and provides strong motivation for a focus on player rank, since the goal is to explain the specific observations in the data set.

4.1 Observations from a Low Dimensional Subspace

The first property we connect to player rank is the dimensionality of the observed strategies in the data set. Given a finite set $S \subset \mathbb{R}^n$ of m vectors s_1, \dots, s_m , write $\dim(S)$ to denote the maximum number of linearly independent vectors in S . Observed strategies that form a low dimensional subspace are natural candidates for low player rank rationalizations and, the following theorem shows that — independent of the size of the data set — if the observations form a low dimensional subspace then they can be rationalized by a game of low player rank.

Theorem 2. *If a data set \mathcal{D} is rationalizable then it can be rationalized by a game of player rank at most $\min\{\dim(\mathcal{O}_r), \dim(\mathcal{O}_c)\}$.*

An immediate consequence of the above theorem is that, if a data set \mathcal{D} is rationalizable, then it can be rationalized by a game of player rank at most $|\mathcal{D}|$. Additionally, later, in Theorem 5, we prove a lower bound that highlights that this result is tight.

Importantly, Theorem 2 has rationalizability as one of its hypotheses, and thus implies that, for data with low-dimensional strategies, the computational constraints have no added empirical or observational content. In other words, any low dimensional data set that is rationalizable without computational constraints is also rationalizable with them.

To prove Theorem 2, a key technical piece is the following lemma about reconstructing the product of an arbitrary matrix and a low-rank matrix.

Lemma 1. *Suppose $Y \in \mathbb{R}^{n \times m}$ is a matrix of rank t . Then for every matrix $\hat{A} \in \mathbb{R}^{n \times n}$, there exists a matrix $A \in \mathbb{R}^{n \times n}$ of rank at most t that satisfies $AY = \hat{A}Y$.*

Proof. Let $\{y_1, y_2, \dots, y_t\}$ be a set of linearly independent columns in Y , and define \hat{Y} as the $n \times t$ matrix $[y_1 \ y_2 \ \dots \ y_t]$. Since each of the columns of \hat{Y} are linearly independent, we can obtain a matrix Γ of size $n \times t$ that satisfies $\Gamma^T \hat{Y} = I_t$. Thus if we denote the i th column of Γ by γ_i , then for all $j \in [t]$,

$$\begin{aligned}\gamma_j^T y_j &= 1 & \text{and} \\ \gamma_j^T y_k &= 0 & \forall k \in [t] \setminus \{j\}.\end{aligned}$$

We define matrix A as the following sum of t rank-1 outer products:

$$A = \sum_{j=1}^t \hat{A} y_j \gamma_j^T.$$

By construction, the rank of A is at most t . Further, for all y_j for $j \in [t]$:

$$A y_j = \sum_{k=1}^t \hat{A} y_k \gamma_k^T y_j = \hat{A} y_j. \quad (1)$$

Since any column y of matrix Y can be expressed as a linear combination of the columns of \hat{Y} , we can write $y = \sum_{j=1}^t \lambda_j y_j$. Then $\hat{A} y$ is given by

$$\begin{aligned}\hat{A} y &= \hat{A} \left(\sum_{j=1}^t \lambda_j y_j \right) = \sum_{j=1}^t \lambda_j \hat{A} y_j = \sum_{j=1}^t \lambda_j A y_j \\ &= A \left(\sum_{j=1}^t \lambda_j y_j \right) = A y.\end{aligned}$$

where the third equality is obtained from (1). Overall we have the desired claim, $A Y = A \hat{Y}$. \square

Using this lemma, Theorem 2 can be established as follows.

Proof of Theorem 2. Consider a game (\hat{A}, \hat{B}) that rationalizes the data set $\mathcal{D} = \{(x_k, y_k)\}_{k=1}^m$. Write X (Y) to denote the matrix whose k th row (column) is equal to x_k (y_k), for all $k \in [m]$. Note that $\text{rank}(X) = \dim(\mathcal{O}_r)$ and $\text{rank}(Y) = \dim(\mathcal{O}_c)$. By Lemma 1, there exist matrices A, B so that $\text{rank}(A) \leq \text{rank}(Y)$, $\text{rank}(B) \leq \text{rank}(X)$, and $A Y = \hat{A} Y$, $X B = X \hat{B}$. Then (A, B) is the rationalization required by the theorem. We have already shown that A and B are of the required rank. To see that (A, B) rationalize \mathcal{D} , note that since $A Y = \hat{A} Y$, for all $(x, y) \in \mathcal{D}$ we have $\beta_r^A(y) = \beta_r^{\hat{A}}(y)$. Hence $\text{Supp}(x) = \beta_r^A(y)$. Similarly, since $X B = X \hat{B}$, $\text{Supp}(y) = \beta_c^B(x)$. Hence (x, y) is a strict Nash equilibrium in (A, B) . \square

4.2 Observations with Small Support Size

The second structural property of the data set we consider is the support size of the observations. In spirit, the following theorem complements the result of Lipton et al. [15] wherein they establish that if the rank of both the payoff matrices is low then the game contains a small-support equilibrium. The following result highlights that there is a connection in the other direction as well.

Theorem 3. *Let $\mathcal{D} = \{(x_k, y_k) \in \Delta^n \times \Delta^n \mid 1 \leq k \leq m\}$ be a data set in which $|\text{Supp}(x_k)| \leq s$ for all $k \in [m]$ or $|\text{Supp}(y_k)| \leq s$ for all $k \in [m]$. If the observed strategies $\mathcal{O}_r(\mathcal{D})$ and $\mathcal{O}_c(\mathcal{D})$ are generic then \mathcal{D} can be rationalized by a game with player rank $\leq 2s + 1$.*

Note that, later, Theorem 5 highlights that the bound in Theorem 3 is tight to within a factor of 2.

Our proof of Theorem 3 uses a construction based on polynomials, and the following lemma is a key technical piece in the argument. The lemma states that if for all $j \in [m]$, the j th column of an $n \times m$ matrix M is obtained by evaluating a degree d polynomial p_j at $1, 2, \dots, n$ (i.e., the j th column of M is equal to $(p_j(1), p_j(2), \dots, p_j(n))^T$), then the rank of M is at most $d + 1$.

Lemma 2. *Let p_1, p_2, \dots, p_m be m univariate polynomials over \mathbb{R} , and suppose that the degree of each of them is at most d . If the (i, j) th entry of an $n \times m$ matrix M is equal to $p_j(i)$, for all $i \in [n]$ and $j \in [m]$, then the rank of M is at most $d + 1$.*

Proof. Write $p_j(x) = a_d^{(j)}x^d + a_{d-1}^{(j)}x^{d-1} + \dots + a_1^{(j)}x + a_0$ for all $j \in [m]$. M can be expressed as a sum of $d + 1$ outer products:

$$M = \begin{pmatrix} 1^d \\ 2^d \\ \vdots \\ n^d \end{pmatrix} \begin{pmatrix} a_d^{(1)} & a_d^{(2)} & \dots & a_d^{(m)} \end{pmatrix} + \begin{pmatrix} 1^{d-1} \\ 2^{d-1} \\ \vdots \\ n^{d-1} \end{pmatrix} \begin{pmatrix} a_{d-1}^{(1)} & a_{d-1}^{(2)} & \dots & a_{d-1}^{(m)} \end{pmatrix} + \dots \\ \dots + \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} \begin{pmatrix} a_0^{(1)} & a_0^{(2)} & \dots & a_0^{(m)} \end{pmatrix}.$$

Note that the rank of an outer product is one and the rank of the sum of two matrices satisfies $\text{rank}(X + Y) \leq \text{rank}(X) + \text{rank}(Y)$. Hence the rank of M is no more than $d + 1$. \square

Using this lemma, Theorem 3 can be established as follows.

Proof of Theorem 3. We prove the claim for the case in which the mixed strategies of the row player in the data set $\mathcal{D} = \{(x_k, y_k) \in \Delta^n \times \Delta^n \mid 1 \leq k \leq m\}$ are of support size at most s . A construction similar to the one presented below takes care of the alternate case wherein $|\text{Supp}(y_k)| \leq s$ for each $k \in [m]$.

We consider a polynomial p_k that satisfies $\arg \max_x p_k(x) = \text{Supp}(x_k)$ and has degree $2|\text{Supp}(x_k)|$. In particular,

$$p_k(x) := - \prod_{i \in \text{Supp}(x_k)} (x - i)^2.$$

Say $\text{Supp}(x_k) = \{i_1, i_2, \dots, i_s\} \subset [n]$, then the polynomial p_k vanishes exactly at i_1, i_2, \dots, i_s and is negative elsewhere. Hence, $\text{Supp}(x_k)$ is the set of points at which p_k attains its maximum value. In addition, the degree of p_k is $2|\text{Supp}(x_k)|$.

Consider the $n \times m$ matrix P in which the k th column is equal to $(p_k(1), p_k(2), \dots, p_k(n))^T$. By construction, for all $k \in [m]$, degree of the polynomial p_k is no more than $2s$. Therefore, Lemma 2 implies that the rank of P is at most $2s + 1$. Moreover, the set of the largest components of the k th column of P (i.e., $\arg \max_i P_{i,k}$) is exactly equal to $\text{Supp}(x_k)$. Since, $\arg \max_{i \in [n]} P_{i,k} = \arg \max_{i \in [n]} p_k(i) = \text{Supp}(x_k)$.

Recall that the mixed strategies in $\mathcal{O}_c(\mathcal{D})$ are generic. Therefore, we can find an $m \times n$ matrix V that satisfies the following equality for all $y_k \in \mathcal{O}_c(\mathcal{D})$: $Vy_k = e_k^{(m)}$. Here $e_k^{(m)}$ is the m -dimensional vector with a 1 in the k th coordinate and 0s elsewhere.

Set the payoff matrix of the row player $A = PV$. Rank of the product of two matrices satisfies: $\text{rank}(XY) \leq \min\{\text{rank}(X), \text{rank}(Y)\}$. Hence $\text{rank}(A) \leq 2s + 1$.

For all $(x_k, y_k) \in \mathcal{D}$, we have $Ay_k = Pe_k^{(m)} = (p_k(1), p_k(2), \dots, p_k(n))^T$. Hence, the set of the largest components of the vector Ay_k is equal to $\text{Supp}(x_k)$. Overall, under the payoff matrix A , we have $\beta_r(y_k) = \text{Supp}(x_k)$, for all $(x_k, y_k) \in \mathcal{D}$. That is, A rationalizes the mixed strategies of the row player. \square

4.3 Observations with Low Chromatic Number

The third, and final, structural property of data sets that we consider is the chromatic number. Intuitively, the chromatic number quantifies the degree of intersection between the observed mixed strategies, and hence it is a relevant measure of the structural complexity of data.

For a data set \mathcal{D} , we define the row chromatic number $\kappa_r(\mathcal{D})$ and the column chromatic number $\kappa_c(\mathcal{D})$ as the chromatic numbers of graphs G_r and G_c , defined as follows. For the row chromatic number, $\kappa_r(\mathcal{D})$, construct graph G_r with a vertex corresponding to each observation in \mathcal{O}_r . For distinct observations (x, y) and (x', y') in \mathcal{D} , if $\text{Supp}(x) \cap \text{Supp}(x') \neq \emptyset$ then the graph G_r has an edge between the corresponding vertices. Then set $\kappa_r(\mathcal{D}) = \chi(G_r)$, i.e., the chromatic number of graph G_r . The column chromatic number is defined similarly using intersections $\text{Supp}(y) \cap \text{Supp}(y')$. The chromatic number of the data set, $\kappa(\mathcal{D})$, is defined to be the minimum of $\kappa_r(\mathcal{D})$ and $\kappa_c(\mathcal{D})$.

Theorem 4. *Let \mathcal{D} be a data set with chromatic number equal to $\kappa(\mathcal{D})$. If the observed mixed-strategy sets $\mathcal{O}_r(\mathcal{D})$ and $\mathcal{O}_c(\mathcal{D})$ are generic then \mathcal{D} can be rationalized by a game of player rank at most $2\kappa(\mathcal{D})$.*

Note that, later, Theorem 5 highlights that the bound in Theorem 3 is tight to within a factor of 2.

Importantly, like the case of support size, the bound on the player rank in Theorem 4 is not exactly dependent on the size of the data set, but rather only on the “richness” of the observations in terms of the structure of the underlying graph.

Of course, in general the chromatic number of a graph is hard to compute. However, an easy upper bound is the maximum degree of any vertex in G_r and G_c plus one, which then can be interpreted as follows: $\kappa_r(\mathcal{D}) \leq \max_{(x,y) \in \mathcal{D}} |\{(x',y') \in \mathcal{D} : \text{Supp}(x) \cap \text{Supp}(x') \neq \phi\}|$ and $\kappa_c(\mathcal{D}) \leq \max_{(x,y) \in \mathcal{D}} |\{(x',y') \in \mathcal{D} : \text{Supp}(y) \cap \text{Supp}(y') \neq \phi\}|$. Though these bounds provide intuition, it is obvious the chromatic numbers can be much less than these upper bounds, e.g., if the graph G_r is a star.

The proof of Theorem 4 starts by focusing on data sets where the row chromatic number is one, i.e., the supports of all the observations for the row player are pairwise disjoint. In this case we show, in the following Lemma, that the data set can be rationalized by a game with row player rank 2. We start with a proof of the initial lemma.

Lemma 3. *Let \mathcal{D}' be a data set with row chromatic number $\kappa_r(\mathcal{D}') = 1$. If the set of observations $\mathcal{O}_c(\mathcal{D}')$ is generic then there exists a rank 2 matrix A' that rationalizes the row player's strategies in \mathcal{D}' .*

Proof. Say \mathcal{D}' consists of m observations $(x_1, y_1), \dots, (x_m, y_m)$. Since the support sets of the strategies of the row player are disjoint, any (pure) strategy $i \in [n]$ of the row player is in at most one such support set. For pure strategy $i \in [n]$, let $\sigma(i) := k$ such that $i \in \text{Supp}(x_k)$. If row i is not in the support of any strategy, let $\sigma(i) = m + 1$.

We construct three vectors u , w and f in \mathbb{R}^n . Values are assigned to the components of the vectors u and w directly: for $i \in [n]$, entry $u_i = -\sigma(i)^2$, and $w_i = \sigma(i)$. We define vector f so that for all $y_k \in \mathcal{O}_c(\mathcal{D}')$, $f^T y_k = k$. Since the set of observations $\mathcal{O}_c(\mathcal{D}')$ are generic, such a vector f exists and can be obtained. Matrix A' is then defined as follows: $A' := u\mathbf{1}_n^T + 2wf^T$, where $\mathbf{1}_n$ is the n -column vector consisting of all 1's.

Since A' is the sum of two outer products, it has rank 2. We show that with payoffs from A' the strict Nash requirement for the row player, $\text{Supp}(x_k) = \beta_r(y_k)$, is satisfied for all $(x_k, y_k) \in \mathcal{D}'$. Hence, we get the desired lemma.

For any $y_k \in \mathcal{O}_c(\mathcal{D}')$, by our construction, $A'y_k = u + 2wf^T y_k = u + 2kw$. For a fixed k , consider the j th component of the vector $A'y_k$, which by the construction is $-\sigma(j)^2 + 2k\sigma(j)$. Note that this expression is maximized when $\sigma(j) = k$ and is strictly less for other values of $\sigma(j)$. Further, by construction, if $\sigma(j) = k$ then $j \in \text{Supp}(x_k)$. Thus, the maximum components of $A'y_k$ are exactly those that correspond to the support of x_k , and hence under payoff matrix A' , $\beta_r(y_k) = \text{Supp}(x_k)$. \square

We show now how this lemma can then be used to construct games with low player rank for data sets with larger chromatic numbers.

Proof of Theorem 4. We constructively show that there exists a matrix A that rationalizes the mixed strategies of the row player in \mathcal{D} and has rank no more than $2\kappa_r(\mathcal{D})$. Similarly, we can

construct a payoff matrix B of rank at most $2\kappa_c(\mathcal{D})$ for the column player. This establishes the existence of the game (A, B) that rationalizes \mathcal{D} with the required player rank.

Let t be the row chromatic number of the data set, i.e., $t = \kappa_r(\mathcal{D})$, and let $\chi(G_r)$ be the graph coloring that defines the row chromatic number. We partition the observations in data set \mathcal{D} into t sets $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_t$ according to the color assigned to the vertex corresponding to each observation, so that observations with the same color are in the same partition. Note that by this technique, for all observations within the same partition, the supports of the observations for the row player are disjoint.

By Lemma 3, for each data set \mathcal{D}_j , we can obtain a rank-2 matrix A_j so that $\beta_r(y_i) = \text{Supp}(x_i)$ with A_j as the payoff matrix. In order to combine these matrices, for $j \in [t]$, we define matrix V_j as satisfying the following property:

$$\begin{aligned} V_j y &= y \text{ for } y \in \mathcal{O}_c(\mathcal{D}_j), \text{ and} \\ V_j y &= \mathbf{0} \text{ otherwise.} \end{aligned}$$

Since the set of strategies of the players are generic, we can obtain such matrices. We then define the payoff matrix A as $A = \sum_{j=1}^t A_j V_j$. Since each A_j is of rank 2, matrix A is of rank $2t = 2\kappa_r(\mathcal{D}')$. To see that A rationalizes the data set \mathcal{D} , note that for $j \in [t]$ and $(x, y) \in \mathcal{D}_k$, $Ay = \sum_{j=1}^t A_j V_j y = A_k y$, and by construction of A_k , $\beta_r(y) = \text{Supp}(x)$. \square

4.4 A Unifying Result

The previous sections have identified three structural properties for data sets that ensure the existence of low player rank rationalizations. In this section, we present a unifying result that extends the previous three theorems to provide a more robust low-rank construction and, in particular, shows that addition of a small number of observations to a data set does not have a big impact on the player rank necessary to rationalize the data. Specifically, we establish low-rank rationalizations for data sets that can be partitioned into three sets which are structurally simple in terms of dimensionality, support size, and chromatic number, respectively. For example, say we have a data set \mathcal{D} in which all but t observed mixed strategies are of support size s , then it can be partitioned into a set that has support bounded by s and a set that has dimensionality bounded by t (and an empty set that has chromatic number zero). The following corollary then shows how to construct a rationalization for \mathcal{D} of player rank at most $(2s + 1) + t$.²

To obtain such a generalization we introduce the notion of the *composite number* of a data set, which considers a 3-partition of the data set and combines the dimensionality of the first partition, the support size of the second partition, and the chromatic number of the third partition.

Definition 3 (Composite number). *The row composite number $\sigma_r(\mathcal{D})$ of a data set \mathcal{D} is defined to be the smallest number for which there exists a 3-partition of \mathcal{D} , $\{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$, that satisfies*

²The precise bound from the corollary is $2(s + t) + 1$, but this can be strengthened to $(2s + 1) + t$. For ease of presentation, we present the corollary with slightly loose factors.

$\dim(\mathcal{O}_c(\mathcal{D}_1)) + \max_{(x,y) \in \mathcal{D}_2} |\text{Supp}(x)| + \kappa_r(\mathcal{D}_3) = \sigma_r(\mathcal{D})$. The column composite number $\sigma_c(\mathcal{D})$ is defined similarly.

The composite number of a data set, $\sigma(\mathcal{D})$, is the minimum of the row and column composite number: $\sigma(\mathcal{D}) := \min\{\sigma_r(\mathcal{D}), \sigma_c(\mathcal{D})\}$.

Corollary 1. *Let \mathcal{D} be a data set with composite number $\sigma(\mathcal{D})$. If the observed mixed-strategy sets $\mathcal{O}_r(\mathcal{D})$ and $\mathcal{O}_c(\mathcal{D})$ are generic then \mathcal{D} can be rationalized by a game of player rank at most $2\sigma(\mathcal{D}) + 1$.*

Proof. Below we show that there exists a payoff matrix A of rank at most $2\sigma_r(\mathcal{D})+1$ that rationalizes the row player's strategies in \mathcal{D} . A similar argument establishes the existence of a matrix B (which rationalizes the column player's strategies) of rank no more than $2\sigma_c(\mathcal{D}) + 1$, and hence we get the desired claim.

Say $\{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$ is a 3-partition that satisfies $\dim(\mathcal{O}_c(\mathcal{D}_1)) + \max_{(x,y) \in \mathcal{D}_2} |\text{Supp}(x)| + \kappa_r(\mathcal{D}_3) = \sigma_r(\mathcal{D})$. The constructions of Theorems 2, 3, and 4, imply that there exist matrices A_1, A_2, A_3 of rank $\dim(\mathcal{O}_c(\mathcal{D}_1))$, $2 \max_{(x,y) \in \mathcal{D}_2} |\text{Supp}(x)| + 1$, and $2\kappa_r(\mathcal{D}_3)$ respectively such that A_i rationalizes the row player's observations in \mathcal{D}_i for all $i \in [3]$.

Since $\mathcal{O}_c(\mathcal{D})$ is generic, there exists matrix V_i for all $i \in [3]$ that satisfies the following equalities:

$$\begin{aligned} V_i y &= y & \forall y \in \mathcal{O}_c(\mathcal{D}_i) \\ V_i y &= \mathbf{0} & \forall y \in \mathcal{O}_c(\mathcal{D}) \setminus \mathcal{O}_c(\mathcal{D}_i) \end{aligned}$$

Note that payoff matrix $A = \sum_{i=1}^3 A_i V_i$ rationalizes the row players observations in \mathcal{D} and is of rank at most $2\sigma_r(\mathcal{D}) + 1$. \square

This result serves as another illustration of why player rank is an appealing choice for the revealed preference exercise in this paper, since it allows us to merge the constructions used in Theorems 2, 3, and 4.

5 A Lower Bound on Player Rank

The results to this point of the paper have focused on constructing rationalizing games with low player rank, thus guaranteeing the observed equilibria can be computed efficiently. It is also natural to ask if there exist data sets that *require* rationalizations to have large player rank. In the following, we show that such data sets do exist. In particular, there exists a data set that requires any rationalization to have player rank at least $n - 1$.

Theorem 5. *Any game (A, B) that rationalizes $\mathcal{D} = \{(u_k, u_k) \mid k \in \{2, 3, \dots, n\}\}$ has player rank at least $n - 1$, i.e., $\text{rank}(A) \geq n - 1$ and $\text{rank}(B) \geq n - 1$.*

Proof. For $2 \leq k \leq n$, recall that $u_k \in \Delta^n$ denotes the uniform distribution over the set $\{1, 2, \dots, k\}$. Consider the following data set with $n - 1$ observations, $\mathcal{D} = \{(u_k, u_k) \mid k \in \{2, 3, \dots, n\}\}$.

Note that the data set \mathcal{D} is rationalizable. In particular, the game obtained by setting the payoff matrices of both the players to I_n (the $n \times n$ identity matrix) rationalizes \mathcal{D} . The player

rank of the rationalization, (I_n, I_n) , is n . Below we establish that in fact the player rank of any game that rationalizes \mathcal{D} is at least $n - 1$.

Say (A, B) is a rationalization of \mathcal{D} . Let $A_{(i)}$ be the i th row of the matrix A and for $2 \leq j \leq n$, we define $n - 1$ vectors v_j s as follows: $v_j := A_{(1)} - A_{(j)}$. Note that v_j s lie in the row space of A . We will show that v_j s are linearly independent and hence get that the dimension of the row space of A is at least $n - 1$. This, in turn, proves that the rank of A is at least $n - 1$. Since the data set \mathcal{D} is symmetric, via a similar argument, we can establish that the rank of B is no less than $n - 1$. This overall establishes the stated claim that the player rank of (A, B) is at least $n - 1$.

Since mixed strategy pair (u_k, u_k) is a strict Nash equilibrium in (A, B) , we have $e_1^T A u_k = e_j^T A u_k$ for all $j \in \{2, \dots, k\}$ and $e_1^T A u_k > e_j^T A u_k$ for all $j \in \{k + 1, \dots, n\}$. That is, $A_{(1)}^T u_k = A_{(j)}^T u_k$ for all $j \in [k]$ and $A_{(1)}^T u_k > A_{(j)}^T u_k$ for all $j \notin \{k + 1, \dots, n\}$. We can rewrite these equalities and inequalities using the definition of v_j s as follows:

$$\begin{aligned} v_j^T u_k &= 0 & \forall j \in \{2, \dots, k\} \text{ and} \\ v_j^T u_k &> 0 & \forall j \in \{k + 1, \dots, n\}. \end{aligned}$$

Hence, for all $2 \leq k < n$, vector v_2, v_3, \dots, v_{k+1} are linearly independent. Say for contradiction that they are linearly dependent. Then we can write v_{k+1} as a linear combination of v_2, \dots, v_k , i.e., $v_{k+1} = \sum_{j=2}^k \lambda_j v_j$. Taking inner product of both sides of this equation with u_k leads to a contradiction.

Overall, we get that the vectors v_2, \dots, v_n are linearly independent and this completes the proof. \square

This result is important for two reasons. First, the theorem highlights that Theorems 2, 3, and 4 are (nearly) tight. Specifically, by construction, data set \mathcal{D}' satisfies: (i) the observed strategies of each player lie in a subspace of dimension $n - 1$; (ii) each observed strategy has support size at most n ; and (iii) the chromatic number of the data set is $n - 1$. It follows immediately that the bounds in Theorem 2 are exactly tight, and the bounds in Theorems 3 and 4 are tight to within a factor of 2.

Second, the lower bound strongly suggests that adding computational constraints to the theory of Nash equilibrium has testable implications, i.e., it is likely that there exist data sets for which any payoff matrices that explain the data sets require players to solve computationally hard problems when computing the observed equilibrium. This is in contrast with single-person consumer theory [11]. It is still possible that rationalizing games could be simple, but it seems unlikely. Investigating this issue further is an intriguing direction for future work.

6 Discussion and Future Work

Our work is the first to consider the implications of computational complexity for mixed strategy behavior in games with multiple players. For our results we make certain simplifying assumptions that are natural given prior work on revealed preference theory for a single agent and on the computational complexity of equilibrium. In this section we discuss these assumptions, their relaxation, and interesting open problems.

Observations of exact mixed strategy behavior. We assume that exact mixed strategy behavior is observable, and ignore the presence of noise in our observations. This assumption is justified as our results would be easier to establish under the extra flexibility afforded by noisy observations. Further, while a number of other observations can be considered as possible inputs, e.g., pure strategy samples from a mixed Nash equilibrium, the model of observations we consider for mixed strategies is the natural first step. The assumption that exact mixed strategy behavior can be observed is also made in the theory of individual stochastic choice, e.g., [16–19]. In the theory of individual stochastic choice, these models have developed into actual empirical tools that are very heavily used among economists (for example, they are used as a standard tool by most empirical economists). See [4] for an exposition of the theory as used by empiricists.

Repeated observations from the same game. The data sets considered in this paper comprise of multiple observations of mixed-strategy behavior that potentially correspond to different equilibria. Superficially, this may suggest that one could consider observed player behavior as resulting from a learning dynamic, or as outcomes from a repeated game. However, given the “anything goes” message of the folk theorem (see [12]), studies in experimental economics are often designed explicitly to avoid repeated game behavior. This is typically enforced through anonymous and random matching experimental designs. Thus, Nash equilibria is often the model of interest. Further, considering repeated games or learning dynamics would again involve relaxing some of the constraints considered in the current paper, since it would no longer be required that the payoff be maximized at each observed strategy profile. It is likely that these relaxations would make it easier to provide explanation of the data via tractable payoff matrices.

Alternate notions of tractability. Our results provide strong motivation for the use of player rank as a notion of tractability; however there are many other properties which ensure the existence of efficient algorithms for computing Nash equilibria. These include: a *game rank* of zero or one,³ where the game rank is the rank of $C := A + B$; the existence of a potential function; or the existence of a pure Nash equilibrium. None of these properties is appropriate for use in the exercise here because each is binary: either a data set possesses the particular property, or it does not; and absence of the property renders algorithms based on the property useless in computing equilibria. Further, in Appendix A, we show that for each property other than player rank, simple data sets with a small number of observations necessitate rationalizations that do not satisfy the property.

Genericity of the observations. Theorems 3 and 4 assume that the set of observations are generic. For games with n pure strategies for the players, if the number of observations are less than n , this is a mild assumption, since a tiny perturbation is sufficient to ensure genericity. The assumption does necessitate that the number of observations be at most n . Since computational complexity is interesting for large values of n , our work is still relevant for large classes of data. From a technical point of view however, it would be very interesting to see if this assumption could

³If the game rank is 3 or larger then the computation of a Nash equilibria is PPAD-hard [20].

be removed; or if our computational complexity results could be extended to other classes of data sets.

Learning versus testing. Our focus in this paper is on testing. We do not address the related problem of inference in this paper. The problem of actually learning the underlying payoffs of players by observing player behavior has previously been studied for the case of a single consumer [6, 29] and for correlated equilibrium with multiple players [28]. Our work, while primarily focused on testing, does offer some insight into how structural properties of the observed behavior affect the rank of the payoff matrices we are trying to learn. A more comprehensive study of learning the underlying utilities of players in game-theoretic settings is an important direction for future work.

Acknowledgments

This research was supported by NSF grants CNS-0846025, EPAS-1307794, and CCF-1101470, along with a Linde/SISL postdoctoral fellowship.

References

- [1] Bharat Adsul, Jugal Garg, Ruta Mehta, and Milind Sohoni. Rank-1 bimatrix games: a homeomorphism and a polynomial time algorithm. In *Proceedings of the 43rd annual ACM symposium on Theory of computing*, STOC '11, pages 195–204. ACM, 2011.
- [2] Bharat Adsul, Jugal Garg, Ruta Mehta, and Milind A. Sohoni. Rank-1 bimatrix games: a homeomorphism and a polynomial time algorithm. In *STOC*, 2011.
- [3] Sydney N. Afriat. The construction of utility functions from expenditure data. *International Economic Review*, 8(1):67–77, feb 1967.
- [4] Simon P Anderson, Andre De Palma, and Jacques François Thisse. *Discrete Choice Theory of Product Differentiation*. The MIT press, 1992.
- [5] Siddharth Barman, Umang Bhaskar, Federico Echenique, and Adam Wierman. The empirical implications of rank in bimatrix games. In *ACM Conference on Electronic Commerce*, 2013.
- [6] Eyal Beigman and Rakesh Vohra. Learning from revealed preference. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 36–42. ACM, 2006.
- [7] Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pages 841–890, 1995.
- [8] Timothy F Bresnahan. Empirical studies of industries with market power. *Handbook of industrial organization*, 2:1011–1057, 1989.
- [9] Xi Chen, Decheng Dai, Ye Du, and Shang-Hua Teng. Settling the complexity of arrow-debreu equilibria in markets with additively separable utilities. In *FOCS*, 2009.

- [10] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *J. ACM*, 56(3), 2009.
- [11] Federico Echenique, Daniel Golovin, and Adam Wierman. A revealed preference approach to computational complexity in economics. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 101–110. ACM, 2011.
- [12] Drew Fudenberg and Jean Tirole. Game theory. *Cambridge, Massachusetts*, 1991.
- [13] Jugal Garg, Albert Xin Jiang, and Ruta Mehta. Bilinear games: Polynomial time algorithms for rank based subclasses. In *WINE*, 2011.
- [14] Kamal Jain. A polynomial time algorithm for computing an arrow-debreu market equilibrium for linear utilities. *SIAM Journal on Computing*, 37(1):303–318, 2007.
- [15] Richard J Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 36–41. ACM, 2003.
- [16] R.. Duncan Luce. *Individual Choice Behavior a Theoretical Analysis*. John Wiley and sons, 1959.
- [17] D. McFadden. Conditional logit analysis of qualitative choice behavior. In P. Zarembka, editor, *Frontiers in econometrics*, page 105. Academic Press New York, 1974.
- [18] D. McFadden and M.K. Richter. Stochastic rationality and revealed stochastic preference. In *Preferences, Uncertainty, and Optimality, Essays in Honor of Leo Hurwicz*, Westview Press: Boulder, CO, pages 161–186. 1990.
- [19] D.L. McFadden. Revealed stochastic preference: a synthesis. *Economic Theory*, 26(2):245–264, 2005.
- [20] Ruta Mehta. Constant rank bimatrix games are PPAD-hard. To appear in the ACM Symposium on Theory of Computing (STOC), 2014.
- [21] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [22] Paul A. Samuelson. A note on the pure theory of consumer’s behaviour. *Economica*, 5(17):61–71, 1938.
- [23] Hal R. Varian. The nonparametric approach to demand analysis. *Econometrica*, 50(4):945–974, jul 1982.
- [24] Hal R. Varian. Non-parametric tests of consumer behaviour. *Review of Economic Studies*, 50(1):99–110, jan 1983.

- [25] Hal R. Varian. The nonparametric approach to production analysis. *Econometrica*, 52(3):579–598, may 1984.
- [26] Hal R. Varian. Revealed preference. *Samuelsonian economics and the twenty-first century*, pages 99–115, 2006.
- [27] Vijay V. Vazirani and Mihalis Yannakakis. Market equilibrium under separable, piecewise-linear, concave utilities. *J. ACM*, 58(3), 2011.
- [28] Kevin Waugh, Brian D. Ziebart, and Drew Bagnell. Computational rationalization: The inverse equilibrium problem. In *International Conference on Machine Learning (ICML)*, 2011.
- [29] Morteza Zadimoghaddam and Aaron Roth. Efficiently learning from revealed preference. In *Internet and Network Economics*, pages 114–127. Springer, 2012.

A Game Rank, Potential Games, and Pure Strategy Equilibria

The goal of this paper is to understand when it is possible to rationalize data via payoff matrices for which the mixed strategies observed are efficiently computable. Given the hardness of computing equilibria in general, this requires that the rationalizations we generate must have some special property that allows for efficient computation. We do this by focusing on rationalizations with small player rank; however there are a number of other properties that could be considered. For example, a small game rank, the existence of a potential function, or the existence of a pure Nash equilibrium. In the following, we highlight that these alternatives are not well-suited for use in this paper.

A.1 Game Rank

The connection between game rank and computational efficiency has only recently begun to be understood. To this point, polynomial-time algorithms to compute Nash equilibria are known when the game rank is either zero [21] or one [2], and it has recently been shown that when the game rank is four or more computing an equilibrium is PPA-hard [20]. Though incomplete, these results are already problematic for the use of game rank in this paper. In particular, the following result highlights that only very small data sets can be guaranteed to have game rank small enough to ensure that a computationally efficient algorithm exists, e.g., there is a data set with 9 observations that necessitates game rank of at least two.

Theorem 6. *There exists a rationalizable data set \mathcal{D} with $2n + 1$ observations such that any game (A, B) that rationalizes \mathcal{D} has game rank at least $n - 2$, i.e., $\text{rank}(A + B) \geq n - 2$.*

Proof. Let $u_n \in \Delta^n$ be the uniform distribution over $[n]$ and $e_k \in \Delta^n$ be the vector with a 1 in the k th coordinate and 0's elsewhere. Write v_k to denote the uniform distribution over $[n] \setminus \{k\}$. We consider the following data set with $2n + 1$ observations, $\mathcal{D} = \{(e_k, e_k) \mid 1 \leq k \leq n\} \cup \{(v_k, v_k) \mid 1 \leq k \leq n\} \cup \{(u_n, u_n)\}$. Note that \mathcal{D} can be rationalized by the game (I_n, I_n) , where I_n is the $n \times n$ identity matrix.

Say game (A, B) rationalizes \mathcal{D} . First we show that in every column of A all the off-diagonal entries are equal to each other. That is, for all $k \in [n]$ and for all $i, i' \in [n] \setminus \{k\}$ we have $A_{i,k} = A_{i',k}$. A similar result holds for the rows of matrix B .

Since (u_n, u_n) is a strict Nash equilibrium in (A, B) we have $\text{Supp}(u_n) = \beta_r(u_n)$. This implies that all the components of the vector Au_n are equal, i.e., the row sums of A are equal to each other. Formally,

$$\sum_j A_{i,j} = \sum_j A_{i',j} \quad \forall i, i' \in [n]. \quad (2)$$

Similarly, the fact that (v_k, v_k) is a strict Nash equilibrium implies $\text{Supp}(v_k) = \beta_r(v_k)$. In particular, for all $i, i' \in \text{Supp}(v_k)$ the i th and the i' th component of Av_k must be equal to each other. Since the i th component of the vector Av_k is equal to $\frac{1}{n-1} \sum_{j \neq k} A_{i,j}$ and $\text{Supp}(v_k) = [n] \setminus \{k\}$, we have the following equality for all $i, i' \in [n] \setminus \{k\}$:

$$\sum_{j \neq k} A_{i,j} = \sum_{j \neq k} A_{i',j}. \quad (3)$$

Subtracting (3) from (2) gives us $A_{i,k} = A_{i',k}$ for $i, i' \in [n] \setminus \{k\}$.

Finally, using the fact that $(e_k, e_k) \in \mathcal{D}$ we get that (k, k) is a pure and strict Nash equilibrium in (A, B) for all k . Therefore, $A_{k,k} > A_{i,k}$ for all $i \neq k$. Say the off-diagonal entries of the k th column of A are equal to α_k . We have $A_{k,k} > \alpha_k$ and matrix A has the following form:

$$A = \begin{pmatrix} A_{1,1} & \alpha_2 & \alpha_3 & \cdots & \alpha_n \\ \alpha_1 & A_{2,2} & \alpha_3 & \cdots & \alpha_n \\ \alpha_1 & \alpha_2 & A_{3,3} & \cdots & \alpha_n \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \alpha_1 & \alpha_2 & \alpha_3 & \cdots & A_{n,n} \end{pmatrix}$$

We can write A as the sum of a diagonal matrix and an outer product.

$$A = \begin{pmatrix} A_{1,1} - \alpha_1 & & & & \\ & A_{2,2} - \alpha_2 & & & \\ & & A_{3,3} - \alpha_3 & & \\ & & & \ddots & \\ & & & & A_{n,n} - \alpha_n \end{pmatrix} + \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} \begin{pmatrix} \alpha_1 & \alpha_2 & \cdots & \alpha_n \end{pmatrix}$$

Write D to denote the above diagonal matrix and P to denote the outer product. We have $A = D + P$. Note that all the diagonal entries of D are positive, since $A_{k,k} > \alpha_k$ for all k . Similarly, we can decompose column player's payoff matrix B into a diagonal matrix D' and an outer product P' , i.e., $B = D' + P'$. Like D , the all the diagonal entries of D' are positive.

Overall, we have $A + B = D + D' + P + P'$. The rank of the sum of two matrices satisfies $\text{rank}(X + Y) \leq \text{rank}(X) + \text{rank}(Y)$. Therefore, $\text{rank}(D + D') \leq \text{rank}(A + B) + \text{rank}(-(P + P'))$. Since P and P' are outer products, $\text{rank}(-(P + P')) \leq 2$. The diagonal entries of both D and D' are positive, hence matrix $D + D'$ has full rank. This gives us the desired bound, $\text{rank}(A + B) \geq n - 2$. \square

A.2 Potential Games

When a game has a potential function, it is termed a *potential game*, and an appealing property of such games is that a pure strategy equilibrium is guaranteed to exist (e.g., [21]). Not surprisingly, this property is limiting for the purposes of this paper. That is, if we were to use the existence of a pure strategy equilibria as a property to yield efficient computability of an equilibrium in the rationalizing game, then we would be restricted to extremely limited data sets. To see this, note that there are very simple data sets that cannot be rationalized by a game that has a pure Nash equilibrium, and consequently cannot be rationalized by a potential game.

Theorem 7. *There exists a rationalizable data set \mathcal{D} with three observations such that any game (A, B) that rationalizes \mathcal{D} does not possess a pure Nash equilibrium.*

Proof. We consider a game where each player has 3 strategies, and a data set consisting of the following three observations: $((1, 0, 0), (0, 1/2, 1/2))$; $((0, 1, 0), (1/2, 0, 1/2))$, and $((0, 0, 1), (1/2, 1/2, 0))$. Thus the row player plays a different pure strategy in each observation, while the column player randomizes uniformly over two strategies. To see that any rationalization by matrices A, B does not admit a pure Nash equilibrium, suppose for a contradiction that (i, j) is in fact a pure Nash equilibrium and consider the matrix B . The data set enforces that the maximum entry in each row is not unique, and hence no entry can be a strict pure Nash equilibrium. \square