Mechanisms with Unique Learnable Equilibria

PAUL DÜTTING, Stanford University
THOMAS KESSELHEIM, Cornell University
ÉVA TARDOS, Cornell University

The existence of a unique equilibrium is the classic tool for ensuring predictiveness of game theory. Typical uniqueness results, however, are for Nash and Bayes-Nash equilibria and do not guarantee that natural game playing dynamic converges to this equilibrium. In fact, there are well known examples in which the equilibrium is unique, yet natural learning behavior does not converge to it. Motivated by this, we strive for stronger uniqueness results. We do not only require that there is a unique equilibrium, but also that this equilibrium must be learnable. We adopt correlated equilibrium as our solution concept, as simple and natural learning algorithms guarantee that the empirical distribution of play converges to the space of correlated equilibria. Our main result is to show uniqueness of correlated equilibria in a large class of single-parameter mechanisms with matroid structure. We also show that our uniqueness result extends to problems with polymatroid structure under some conditions. Our model includes a number of special cases interesting on their own right, such as procurement auctions and Bertrand competitions. An interesting feature of our model is that we do not need to assume that the players have quasi-linear utilities, and hence can incorporate models with risk averse players and certain forms of externalities.

Categories and Subject Descriptors: F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

Additional Key Words and Phrases: Unique Equilibrium, Correlated Equilibrium, Regret Minimization

1. INTRODUCTION

A desirable situation in game theory as well as in mechanism design is when the game or mechanism has a unique equilibrium, as the unique equilibrium is typically viewed as the outcome that is likely to emerge. Previous work in showing uniqueness has focused on uniqueness of *Nash* and *Bayes-Nash equilibria*. These equilibria, however, may not be learnable. In some cases natural learning dynamic converges to them in polynomial time; in other cases it takes exponentially long to converge or does not converge at all. For example, Kleinberg et al. [2011] give an example where the unique Nash equilibrium of a game has low social welfare, while natural learning behavior in repeated play ensures high average value for all players. Further, no polynomial-time algorithms for computing Nash and Bayes-Nash equilibria are known for these games, and more more generally, computing Nash equilibria is known to be PPAD complete [Daskalakis et al. 2009]. An alternate equilibrium concept that does not share these drawbacks is the concept of a *correlated equilibrium*. Correlated equilibria form a convex set, and correlated equilibria can be found in polynomial time

Authors' addresses: P. Dütting, Department of Computer Science, Stanford University, 353 Serra Mall, Stanford, CA 94305, USA; email: duetting@cs.stanford.edu; T. Kesselheim and É. Tardos, Department of Computer Science, Cornell University, Gates Hall, Ithaca, NY 14850, USA; email: {kesselheim,eva}@cs.cornell.edu.

EC'14, June 8–12, 2014, Stanford, CA, USA. ACM 978-1-4503-2565-3/14/06 ...$15.00.

http://dx.doi.org/10.1145/2600057.2602838

even in compactly described games [Papadimitriou and Roughgarden 2008; Jiang and Leyton-Brown 2011]. Maybe more importantly, correlated equilibria emerge quite naturally from sequences of repeated play in which the players minimize a form of regret, and simple and natural polynomial-time algorithms are known for minimizing regret (see [Blum and Mansour 2007] for a survey). Motivated by this, we ask: Under which conditions does a mechanism have a unique correlated equilibrium? We thus strive for stronger uniqueness results: We do not only want the equilibrium to be *unique*, we also want it to be the *outcome of natural learning dynamic*.

We consider single-parameter mechanism design settings. In these settings the private information held by the players—commonly referred to as their type—can be described by a single number. We analyze mechanisms that collect reports by the players, and compute an outcome and payments. In the first part of the paper, we focus on a large class of mechanisms that have binary outcomes for each player; they either are selected or not selected. For such mechanisms, the outcome describes which players win and which players lose, and the payments describe how much the players are charged. In the second part of the paper we extend our results to mechanisms whose outcome for each player is an amount they win. These mechanisms compute the amounts won by the players, and how much they have to pay for it.

We use a very general player utility function. We assume that a player's utility is zero if he loses, and it is given by some continuous function of the bid vector if he wins, that is bounded away from zero if he reports less than its true type. Clearly, first- and second-price mechanisms with quasi-linear utilities satisfy these requirements; but they are also met by other mechanisms and utility functions.

— Our model incorporates risk averse players, as we do not assume quasi-linear utilities. Rather the utility can be any continuous function of the bids (and hence the price).
— Our model incorporates certain types of externalities, as the utility of the player can depend on the bid vector, and not only on his own price. For example, the utility of one player can be dependent on the prices paid by others.

An interesting special case of our framework is Bertrand competition, where our model includes risk averse firms and price-dependent demand. Another interesting special case of our model are multi-unit auctions with unit-demand bidders where our model—among other things—allows to model fairness by taking the highest-to-lowest price ratio into account.

1.1. Our Results

Our main result in Section 3 is that mechanisms selecting an independent set of players maximizing the sum of reported types in a matroid have a unique correlated equilibrium if the player utilities satisfy the mild assumptions above. The unique equilibrium we exhibit achieves maximum social welfare. Our main result is applicable to a a vast number of practical problems, including

— procurement auctions (such as first or second price, or any combination) in which a single buyer buys from multiple sellers,
— problems in communication networks in which communication links have to be bought so that the bought communication links form a spanning tree, or
— instances of Bertrand competition in which several firms compete to sell their good to a single buyer.

When applied to these problems our result shows that a large class of mechanisms has a unique, socially optimal, and learnable equilibrium. In each case, our result applies with very many different pricing rules, and we allow different players to be risk averse to different extents and have externalities on the prices paid by others.

We also show that our uniqueness result for matroids is tight by demonstrating that a unilateral relaxation of any of its requirements—mechanism is optimizing, feasibility structure is a matroid, solution concept is that of a correlated equilibrium—can lead to more than one equilibrium.

In Section 4, we extend our considerations to polymatroids. Polymatroids are an important generalization of matroids, and arise, for example, in the context of

— sponsored search auctions [Goel et al. 2012] in which advertisers seek to be assigned clicks on so-called sponsored search results that are shown along with the organic search results produced by a search engine, or
— Bertrand networks [Babaioff et al. 2013] in which firms compete for markets consisting of a single buyer and markets are either captive so that only one firm has access to it or shared between two firms.

While there can be multiple equilibria in general polymatroid settings, we show conditions under which uniqueness carries over. These apply for example in Bertrand networks without captive markets. Applied to this problem, our result shows that there is a unique, learnable equilibrium with zero payoffs.

Together our uniqueness results thus imply stronger uniqueness results for models of Bertrand competition than the ones that were previously known, extending the results to correlated equilibria, to networked markets, and to the general player utilities studied here.

1.2. Proof Technique

The technique underlying our uniqueness results has similarities to iterated elimination of strictly dominated strategies, but in our case no strategies are strictly dominated. Therefore, we do not argue on a per-player basis considering unilateral deviations of a single player only. Instead, we carefully combine effects of unilateral deviations by groups of players. As none of the involved players may benefit from these deviations, we argue that certain ranges of the combined strategy space will never be entered when playing a correlated equilibrium. This proof technique may be applicable to settings beyond the settings studied here.

1.3. Related Work

Most obviously related is work that establishes *uniqueness results*. Work that falls into this category includes work by Amann and Leininger [1996], Bajari [2001], Maskin and Riley [2003], Lebrun [2006], and Chawla and Hartline [2013]. These works differ from our work in that they typically have a specific mechanism in mind. They also differ from the approach taken here in that they generally focus on Nash and Bayes-Nash equilibria, and do not consider correlated equilibria.

There has been work on auctions on domains with *matroid and polymatroid structure*. Bikhchandani et al. [2011] present an ascending price auction for matroids and polymatroids in which truthtelling is an ex-post equilibrium. Milgrom and Segal [2014] and Dütting et al. [2014a,b] argue that for matroids this auction can be implemented as a deferred-acceptance algorithm, which implies that it is weakly group-strategyproof. In a sequence of papers, Hajiaghayi et al. [2007], Chawla et al. [2010], and Alaei [2011] show how algorithms for proving so-called prophet inequalities can be turned into truthful approximation mechanisms. Prophet inequalities for matroids and polymatroids are given in [Kleinberg and Weinberg 2012] and [Dütting and Kleinberg 2013]. The main difference between these works and ours is that they only consider truthful mechanisms, while we also consider non-truthful ones.

Another related work that regards uniqueness of equilibria as a desirable design criterion is work on *implementation theory* (see, e.g., [Jackson 2001], [Maskin and Sjöström 2002], and [Palfrey 2002]). This theory studies games with multiple equilibria and seeks to alter the rules of these games so that they have a unique equilibrium. A fundamental difference between this and our approach is that this theory aims at changing a given game, while we regard the mechanism as given. Another important difference is that this literature focuses on Nash and Bayes-Nash equilibria, while our results are for the broader class of correlated equilibria.

Closest to our work in spirit (but not in techniques) are the work of Blum et al. [2006] and Even-Dar et al. [2009], who show that *no-regret learning dynamic* converges to the unique Nash equilibrium in routing games, and in concave games respectively. We extend this literature to a broad class of mechanisms, and show our uniqueness of correlated equilibria result under a very general model of user utilities that includes both risk averse players and certain forms of externalities.

There has also been work on showing unique outcomes in *Bertrand competition*. The standard model assumes that there is a single market that is accessed by all firms and utilities are quasi-linear. For this model Baye and Morgan [1999] and Kaplan and Wettstein [2000] show that there is a unique mixed Nash equilibrium, and Wu [2008] claims that there is a unique correlated equilibrium. The player utilities considered in these papers are less general than the ones we consider here, and do not allow to model risk averseness or externalities included in our model. Nadav and Piliouras [2010] consider learning outcomes in the standard model. In contrast to our results, they find that learning behavior can sustain a non-zero pricing. Our results show that this conclusion is the result of the weak model of learning used. Nadav and Piliouras [2010] consider coarse correlated equilibrium as their model of learning outcome. We show that with the slightly more sophisticated learning required to reach correlated equilibrium, the non-zero price equilibria are no longer sustained. Babaioff et al. [2013] consider a generalization of the classic setting with multiple markets and restricted access to these markets. They show that in the absence of captive markets, there is a unique mixed Nash equilibrium for quasi-linear utilities. We extend this result to correlated equilibria and the more general player utilities studied here.

Our approach is also related to the literature on *supermodular games* [e.g., Topkis 1968; Vives 1990; Milgrom and Roberts 1990]. This approach is similar to our approach in that the structure of these games allows for iterated elimination of strictly dominated strategies. This, in turn, allows to isolate smallest and largest equilibrium reports, which enables qualitative statements about correlated equilibria. Unlike our approach, however, it generally does not lead to uniqueness results. In contrast, games belonging to our class typically do not contain strictly dominated strategies.

A final related direction is work on *smooth games* [Roughgarden 2009, 2012] and *smooth mechanisms* [Syrgkanis and Tardos 2013]. This literature is related in that it seeks to analyze the quality of correlated equilibria, but unlike our approach it typically only allows to prove approximation guarantees. Our work is different in that its primary goal is to prove the uniqueness of the equilibria achieving socially optimal outcome.

2. PRELIMINARIES

Our main uniqueness results apply to a large class of problems and mechanisms. The first set of problems that we consider are binary single-parameter problems, in which players are characterized by a single privately held number and can either “win” or “lose”. Which players can simultaneously win is governed by feasibility constraints, where we focus on problems in which the feasible sets form a matroid. This feasibility structure covers many interesting applications, including various forms of forward and procurement auctions. Then in Section 4 we derive conditions under which our results extend to problems in which players can win different quantities and the feasible quantities form a polymatroid, extending the class of applications.

The class of mechanisms that we consider include a broad range of pricing policies. Rather than defining the policies considered explicitly, we define them implicitly via the utilities that they induce for the winning players. Our assumptions regarding the utilities are very general. We only require that the utility for winning is strictly positive if a player under-reports, that the utility is continuous in the private information held by the players, and that over-reporting is weakly dominated by reporting truthfully. These requirements are satisfied by a large class of mechanisms and utilities. While the generalized-first price mechanism with quasi-linear utilities falls into this category, prices can generally depend on the bids in more complex ways and utilities need not be quasi-linear. This allows, amongst others, to model risk averse players or players sensitive to fairness constraints.

2.1. Definitions

Binary mechanism design problem. A binary single-parameter mechanism design problem is defined by the triple $(N, \mathcal{F}, \Theta)$. The set N is the set of players. The family of sets of players $\mathcal{F} \subseteq 2^N$ describes the sets of players that can be feasibly accepted. We refer to the players that are picked as winners, and to the remaining ones as losers. The fact that a player can either win or lose makes it a binary problem. Finally, $\Theta = \prod_{i=1}^{n} \Theta_i$ is the type space, where $\theta_i \in \Theta_i = [\theta_{\min}, \theta_{\max}] \subseteq \mathbb{R}$ is the private information held by player i. Note that we do not require types to be non-negative. Positive types correspond to valuations, and negative types to costs. The welfare achieved by a feasible set of players $F \in \mathcal{F}$ is $\sum_{i \in F} \theta_i$.

Direct mechanism. We consider direct mechanisms. This means that the player's bid to the mechanism is a reported type. Note that, as the type is a player's private information, this report need not be truthful. We use $b \in \Theta$ to denote the reported types of the players, distinguishing them from their true types, which we denote by θ. A direct mechanism $M = (f, p)$ consists of an outcome rule $f : \Theta \to \mathcal{F}$ and a payment rule $p : \Theta \to \mathbb{R}^n$, both use the reported types b as input and then either compute a set of winners or payments.

Player utility. For the outcome of a mechanism, the player's utility depends on the fact whether he wins or not, on his type, and on the bids of all players (which determine payments). We consider full information games, so the types θ are a fixed property of the players. Therefore, we express the utility only as a function of b. We assume that $u_i(b) = w_i(b)$ if $i \in f(b)$ and 0 otherwise. Expressing the utility through the functions w_i allows us to impose conditions on w_i. We impose the following requirements on the function w_i:

(1) If underreporting the type and winning, the utility is strictly positive, i.e., $w_i(b) > 0$ if $b_i < \theta_i$.
(2) The function w_i is continuous in b.
(3) Reporting a type $b_i > \theta_i$ is weakly dominated by reporting θ_i, i.e., $u_i(b) \leq u_i(\theta_i, b_{-i})$ if $b_i \geq \theta_i$.

Note that we do not assume that $w_i(b)$ depends only on the price i pays. The utility of a player can depend on the price paid by others. This allows us to model certain externalities, such as fairness. The formulation also can model risk-averse or risk-seeking players, as the utility is not required to be quasi-linear. In general, we do not even assume the functions w_i to be monotonic. However, if w_i is strictly decreasing in b_i for all i, we will get stronger uniqueness results.

Matroid optimization. Feasibility is defined by matroid constraints. A pair consisting of N and $\mathcal{I} \subseteq 2^N$ is a matroid if

(1) $\emptyset \in \mathcal{I}$,
(2) $S \subseteq T \in \mathcal{I}$ implies $S \in \mathcal{I}$, and
(3) $S, T \in \mathcal{I}$ and $|T| > |S|$ implies the existence of $t \in T$ such that $S \cup \{t\} \in \mathcal{I}$.

The sets $S \in \mathcal{I}$ are referred to as *independent sets* and a maximal independent set is a *basis*. A set $S \subseteq N$ is called a *spanning set* if it is the superset of a basis.

In a value-maximization problem types are positive and the set of feasible solutions are independent sets. In a cost-minimization problem, types are negative and feasible solutions are spanning sets. In either case optimal solutions are bases. Hence we can unify the treatment of the two cases by considering the problem of finding a welfare-maximizing basis. We refer to a mechanism that computes such a basis as optimizing.

To illustrate this definition and to develop some intuition for what types of problems have matroid structure consider the following examples:

— In a uniform matroid of rank k the independent sets are the subsets of size at most k, where k is some non-negative integer. The bases are the sets of cardinality exactly k and the spanning sets are the set of at least k elements.
— A transversal matroid is defined by means of an undirected bipartite graph (V_1, V_2, E); its ground set is V_1 and a subset $S \subseteq V_1$ is independent if the vertices in S can be simultaneously matched to vertices in V_2.
— A graphic matroid is defined by an undirected graph $G = (V, E)$; the ground set is E and the independent sets are the acyclic subsets of E.

Solution concept. We assume that the players act strategically and that they seek to maximize their utility. The game-theoretic solution concept that we will impose is that of a correlated equilibrium. A correlated equilibrium is a distribution $\mathcal{D}$ over reports $b \in \Theta$ such that for every player i and possible deviation $\phi \colon \Theta_i \to \Theta_i$,

$$\mathbb{E}_{b \sim \mathcal{D}}[u_i(b_i, b_{-i})] \geq \mathbb{E}_{b \sim \mathcal{D}}[u_i(\phi(b_i), b_{-i})].$$

This generalizes mixed Nash equilibria by allowing the individual players' strategies to be correlated. An example of a possible deviation is $\phi(b_i) = \min\{0.5, b_i\}$, which replaces high bids by 0.5.

Correlated equilibria were introduced by Aumann [1974], using a form of correlation device—a coordinator outside the game—that gives players a recommendation on what to play. This leads naturally to the equilibrium condition that requires that no player can benefit from unilateral deviations of the form "each time the correlation device asks me to report x, I always report y instead."

Correlated equilibria are attractive because they can also be defined as limit points of sequences of repeated play in which the players use learning strategies, strategies that guarantee that they have no swap regret in the limit. See survey of Blum and Mansour [2007] and the recent book of Hart and Mas-Colell [2013] for more on the notion of correlated equilibria, and for simple and natural learning algorithms that guarantee the required form of vanishing regret.

We follow the convention to only consider undominated equilibria [e.g., Caragiannis et al. 2012]. Together with our assumptions regarding the utility functions this translates into a "no overbidding" assumption.

2.2. Examples

Single-item auction with risk-averse players. Probably the simplest example that our analysis applies to is a single-item auction in which a single seller sells to exactly one of several buyers. The feasibility structure is a uniform matroid of rank one. For first-price payments with quasi-linear utilities the utilities for winning are $w_i(b) = \theta_i - b_i$ for all $i \in N$. Second-price payments are captured by $w_i(b) = \theta_i - \min\{b_i, \max_{j \neq i} b_j\}$.[1] In both cases we can model risk averseness by applying a non-decreasing concave function to the quasi-linear term.

Bertrand competition with price-dependent demand. A similarly simple example is a Bertrand competition in which a single buyer buys from exactly one of several firms. The feasibility structure is again a uniform, rank one matroid. If first-price payments are used and utilities are quasi-linear, then the utility for winning is $w_i(b) = \theta_i - b_i$ for all $i \in N$. Here both θ_i and b_i should be thought of as negative. This can be generalized by setting $w_i(b) = d_i(b_i)(\theta_i - b_i)$, where d_i is a non-increasing continuous demand function.

Multi-unit auction with fairness-sensitive players. All these examples can be generalized to more sophisticated matroids. As a concrete example consider a multi-unit auction in which a single seller sells to exactly k unit-demand buyers. This is a uniform matroid of rank k. With first-price payments and quasi-linear utilities the utilities for winning are again $w_i(b) = \theta_i - b_i$ for all $i \in N$. Denote the 1-highest bid by $b_{(1)}$ and the k-highest bid by $b_{(k)}$. Then a desire for prices to be as equal as possible can be expressed by multiplying the $w_i(b)$ function above with $1 - (b_{(1)} - b_{(k)})/b_{(1)}$, which is maximized if the 1-highest and k-highest bids are identical.

3. MATROID FEASIBILITY STRUCTURE

In this section we show our main result: Given a matroid feasibility structure, an optimizing mechanism, and utility functions as defined above, there is a unique correlated equilibrium, which is socially optimal. We first provide some intuition for the proof of our main result by considering the special case of a single-item, first-price auction with two buyers. Then we prove the main result for the general case. We conclude by showing that a unilateral relaxation of any of our assumptions—correlated equilibrium, matroid feasibility structure, exact optimization—can lead to the existence of multiple correlated equilibria, which shows that our assumptions are not only sufficient but also necessary.

3.1. Intuition for Proof of Main Result

We begin by illustrating the idea behind the proof of our main result. To this end consider a single-item, first-price auction with two buyers. Assume that the buyers have quasi-linear utilities with a valuation of one for winning. We claim that in the unique correlated equilibrium both buyers bid one with certainty.

We will show by induction that if $\mathcal{D}$ is a correlated equilibrium, then for all $x < 1$ the probability that both buyers use a bid $b_i \in [0, x]$ for $i \in \{1, 2\}$ is zero. Given this claim for all $x < 1$, taking the limit as $x \to 1$ we get that the probability that both buyers use a bid $b_i < 1$ is also 0. To see that both buyers bid $b_i = 1$ with probability 1, note that with a bid of $b_i = 1$ a buyer has 0 utility in a first-price auction. If one of the buyers is bidding $b_i < 1$ with positive probability, the other buyer has an incentive to bid less than 1 and win with positive probability, showing the claim that both buyers must bid one with certainty.

See Figure 1 for an illustration of the inductive argument that the probability that both buyers use a bid $b_i \in [0, x]$ for $i \in \{1, 2\}$ must be zero. We will use A through F

[1]Note that the minimum operator is only necessary because we assume the functions w_i to be defined on the entire type space Θ, even for report vectors where i is a loser. This will help us to avoid cumbersome notation.

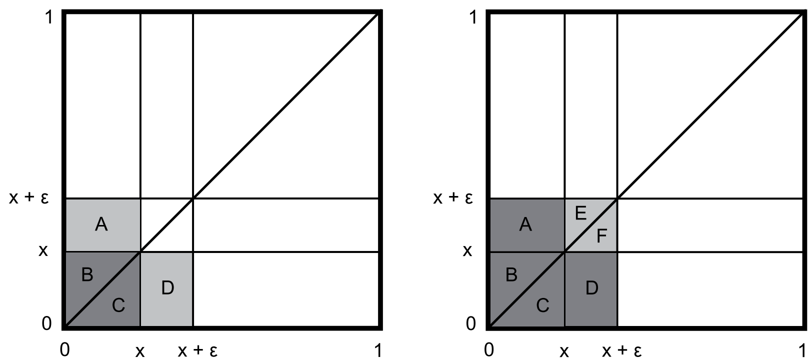

Fig. 1. Single-Item, First-Price Auction with Two Buyers

to denote the probability that the pair of bids falls into the corresponding region on the figure. We will show that this is the case for a sequence of x values with $x \to 1$. The induction hypothesis is that the probability that both buyers bid $b_i \in [0, x]$ for $i \in \{1, 2\}$ is zero. That is $B = C = 0$ in the figure. For the inductive step we choose $\epsilon > 0$ such that $x + 2\epsilon < 1$. The inductive step is to show that the probability that both buyers bid $b_i \in [0, x + \epsilon]$ for $i \in \{1, 2\}$ is zero. The proof of the inductive step consists of two parts. In the first part we argue that the probability that one of the buyers bids $b_i \in [x, x + \epsilon]$ while the other buyer bids $b_{-i} \in [0, x]$ for $i \in \{1, 2\}$ is zero. That is $A = D = 0$ in the figure. In the second part we show that the probability that both players bid $b_i \in [x, x + \epsilon]$ for $i \in \{1, 2\}$ is zero. That is $E = F = 0$ in the figure.

For the first part we consider the deviation of buyer i for $i \in \{1, 2\}$ that maps bids in $[0, x]$ to $x + \epsilon$ and leaves all other bids unchanged. Since we assume that the bids constitute a correlated equilibrium neither deviation should be beneficial. The player who used to lose in areas A and B will now win, a gain in utility, but will pay a higher price in area C, while the deviation is not effecting the outcome in other areas. Similar argument applies to the other player. We obtain that

$$0 \geq (A + B)(1 - (x + \epsilon)) - C\epsilon, \text{ and}$$
$$0 \geq (D + C)(1 - (x + \epsilon)) - B\epsilon.$$

By the induction hypothesis $B = C = 0$. So the right-hand sides of the above inequalities simplify to $A(1 - (x + \epsilon))$ and $D(1 - (x + \epsilon))$. Since $x + \epsilon < 1$ both $A > 0$ and $D > 0$ would lead to a contradiction. We conclude that $A = D = 0$.

For the second part we consider the deviation of buyer $i \in \{1, 2\}$ that maps bids in $[x, x + \epsilon]$ to $x + \epsilon$. The equilibrium condition again dictates that neither of these deviations should be beneficial. We obtain that

$$0 \geq E(1 - (x + \epsilon)) - (F + D)\epsilon, \text{ and}$$
$$0 \geq F(1 - (x + \epsilon)) - (E + A)\epsilon.$$

As we have shown in the first step $A = D = 0$. Hence the right-hand sides of the above inequalities simplify to $E(1 - (x + \epsilon)) - F\epsilon$ and $F(1 - (x + \epsilon)) - E\epsilon$. So by summing up the two inequalities we obtain

$$0 \geq (E + F)(1 - (x + 2\epsilon)).$$

Since we have chosen $\epsilon > 0$ such that $x + 2\epsilon < 1$ we get a contradiction if $E + F > 0$. We conclude that $E = F = 0$.

3.2. Uniqueness of Winners

The proof of our main result follows a similar pattern. We inductively show that certain regions of the strategy space must be empty by carefully combining individual deviations of the players, which—by the equilibrium property—need not be beneficial. The main difficulties in generalizing the above proof sketch are: First, we need the argument to work for very general utility functions and mechanisms. Second, we need to ensure that the combination of the individual deviations works for more general feasibility structures.

Our main technical tool is Lemma 3.1. To be able to work with general utility functions, we use the Heine-Cantor Theorem, which shows that continuous functions on a compact set are uniformly continuous, hence utility functions are uniformly continuous. To work with more general feasibility structures we use properties of matroids. We show that a player i can only win with a bid b_i if this player i is contained in all bases of the matroid that arises if we restrict the original matroid to players of type at least b_i. In any other case, i.e., if there is another player j whose type is also at least b_i and who could substitute player i, then player i will never win when reporting b_i.

We then use Lemma 3.1 to show that in any correlated equilibrium the set of winners is always a social-welfare maximizing basis OPT. This is due to the fact that, by Lemma 3.1, for players in OPT it makes no sense to choose a bid that does not ensure outbidding the other competing players.

LEMMA 3.1. *Let $i \in N$ be some arbitrary player and let $b_i < \theta_i$. Then, in any correlated equilibrium, i never wins with bid b_i unless i is contained in all bases of the matroid restricted to players of type at least b_i.*

PROOF. Without loss of generality assume that the players are ordered such that $\theta_1 \leq \theta_2 \leq \ldots \leq \theta_n$ with ties broken arbitrarily.[2] For each $m \in [n]$, let $A_m \subseteq \{m, \ldots, n\}$ be the set of players who are not contained in every basis of the matroid restricted to $m, \ldots, n$. Observe that $A_1 \supseteq A_2 \supseteq \ldots \supseteq A_n$. We want to show that no player in A_m will ever win while reporting a type below θ_m. That is,

$$\Pr[b_i < \theta_m, i \in f(b)] = 0 \text{ for every } i \in A_m.$$

We will prove this by induction on m.

For the induction basis, set $\theta_0 = \theta_{\min}$ and $A_0 = N$. Given these definitions, the case $m = 0$ is trivially true because $b_i < \theta_0$ is impossible by definition of possible reports.

Let us turn to the induction step. By induction hypothesis, we know that the respective statements are true for some m. As $A_m \supseteq A_{m+1}$, this implies that $\Pr[b_i < \theta_m, i \in f(b)] = 0$ for all $i \in A_{m+1}$. Considering now a fixed m, let $B(x) \subseteq \Theta$ be the set of bid vectors b such that there is some $i \in A_{m+1}$ with $b_i < x$ and $i \in f(b)$. The above statement now translates to $\Pr[b \in B(\theta_m)] = 0$ and we would like to show that $\Pr[b \in B(\theta_{m+1})] = 0$. For this purpose, we will define an increasing sequence of bids $(\theta_{m,t})_{t \in \mathbb{N}}$ such that $\theta_{m,0} = \theta_m$ and $\lim_{t \to \infty} \theta_{m,t} = \theta_{m+1}$. By using another, nested induction, we will show that $\Pr[b \in B(\theta_{m,t})] = 0$ for all $t \in \mathbb{N}$. As $B(\theta_{m+1}) = \cup_{t=0}^{\infty} B(\theta_{m,t})$, this shows the claim.

The idea behind the sequence of bids and the inductive step in the inner induction is as follows: By induction hypothesis, we have $\Pr[b \in B(\theta_{m,t})] = 0$. We then define $\theta_{m,t+1} = \theta_{m,t} + \delta_{t+1}$ for a sufficiently small δ_{t+1}. If there were now a player $i \in A_{m+1}$ winning with a bid below $\theta_{m,t+1}$, there would also be a corresponding loser $j \in A_{m+1}$, who would win by reporting $\theta_{m,t+1}$ instead. We therefore consider for each player the unilateral deviation in which he exchanges reports smaller than $\theta_{m,t+1}$ with

[2]Note that the following argument holds for any tie-breaking rule applied here. Therefore the statement is shown for any player i, even if he ties with another player.

report $\theta_{m,t+1}$. For a previously winning bid, this can only mean a minor reduction in utility because this bid must have been at least $\theta_{m,t}$ by induction hypothesis. On the other hand, if this deviation turns i into a winner, he faces a major gain. When summing up the utility changes of all players, these gains dominate the losses. As the sum of utility changes cannot be positive by equilibrium property, this then shows $\Pr[b \in B(\theta_{m,t+1})] = 0$

For a formal proof, consider the sequence $(\theta_{m,t})_{t \in \mathbb{N}}$ as defined by $\theta_{m,t+1} = \theta_{m,t} + \delta_{t+1}$, where δ_{t+1} is defined as follows. As all w_i are continuous functions on a compact set, they attain their minimum value. Let therefore be $U_{t+1} = \min_{i \in A_{m+1}} \min_{\theta_{-i} \in \Theta_{-i}} w_i(\theta_{m,t}, b_{-i})$. By definition $U_{t+1} > 0$. All w_i are continuous functions on a compact set. By Heine-Cantor theorem, this implies that they are uniformly continuous. This means, there is a δ_{t+1} such that for all $i \in N$, all b_{-i}, and all $b_i \in [\theta_{m,t}, \theta_{m,t} + \delta_{t+1}]$, we have $w_i(\theta_{m,t}, b_{-i}) - \epsilon_{t+1} \leq w_i(\theta_i, b_{-i}) \leq u_i(\theta_{m,t}, b_{-i}) + \epsilon_{t+1}$, where $\epsilon_{t+1} = \frac{U_{t+1}}{2n+1}$. In words: Each player $i \in N$ will, when increasing his bid unilaterally from $\theta_{m,t}$ by at most δ_{t+1}, change his w_i function by at most ϵ_{t+1}.

For each player $i \in A_{m+1}$, we consider the deviation $\phi\colon \Theta_i \to \Theta_i$ such that $\phi(b_i) = \theta_{m,t+1}$ for $\theta_i \in [\theta_{m,0}, \theta_{m,t+1}]$ and $\phi(b_i) = b_i$ otherwise. Let $G_i = u_i(\phi(b_i), b_{-i}) - u_i(b)$ be player i's gain in utility by this deviation. By equilibrium property, we know that $\mathbb{E}[G_i] \leq 0$ for all $i \in A_{m+1}$. In the following, we will lower-bound $\mathbb{E}[\sum_{i \in A_{m+1}} G_i]$ in terms of $\Pr[b \in B(\theta_{m+1})]$. In combination with the equilibrium property, this will then show our claim.

The expectation in $\mathbb{E}[G_i]$ is over the bid vector b. In the following, we will first leave out this expectation and fix some bid vector $b \in \Theta$. Depending on b, we derive certain bounds on G_i. To be more precise, we distinguish the three cases $b \in B(\theta_{m,t})$, $b \in B(\theta_{m,t+1}) \setminus B(\theta_{m,t})$, and $b \notin B(\theta_{m,t+1})$. Note that, in general, G_i can only be negative if $i \in f(b)$. Due to monotonicity, a winning bid remains a winning bid under the deviation, i.e., $i \in f(b)$ implies $i \in f(\phi(b_i), b_{-i})$.

In case $b \in B(\theta_{m,t})$, we bound the utility decrease trivially by $D_i = \max_{\theta \in \Theta} w_i(\theta) - \min_{\theta \in \Theta} w_i(\theta)$. This is well-defined as w_i is a continuous function on a compact set. Therefore, we have $\sum_{i \in A_{m+1}} G_i \geq -nD_i$.

Let us now turn to the case that $b \in B(\theta_{m,t+1}) \setminus B(\theta_{m,t})$. As $b \in B(\theta_{m,t+1}) \setminus B(\theta_{m,t})$, there is an $i \in A_{m+1}$, winning with a bid $b_i \in [\theta_{m,t}, \theta_{m,t+1})$. So let $S_1 = f(b) \cap \{m+1, \ldots, n\}$ be the corresponding basis of the matroid restricted to players $m+1, \ldots, n$. By definition of A_{m+1}, there has to be another basis $S_2 \subseteq \{m+1, \ldots, n\}$ such that $i \notin S_2$. Considering the two independent sets $S_1 \setminus \{i\}$ and S_2, we observe that there is some $j \in S_2 \setminus S_1$ such that $(S_1 \setminus \{i\}) \cup \{j\}$ is independent by augmentation property. This player j is obviously also contained in A_{m+1} and would become a winner by outbidding i. That is, we have $u_j(b) = 0$ and $u_j(\phi(b_j), b_{-j}) \geq U_{t+1} - \epsilon_{t+1}$, which means $G_j \geq U_{t+1} - \epsilon_{t+1}$. For all other players, we observe that each winning bid remains a winning bid under the deviation. In this case, the utility may decrease but at most by $2\epsilon_{t+1}$. A losing bid might become a winning bid after the deviation. As $\theta_{i'} \geq \phi(b_{i'})$, we get $G_{i'} \geq 0$ in these cases. Thus, summing over all players in A_{m+1}, we get $\sum_{i \in A_{m+1}} G_i \geq U_{t+1} - 2n\epsilon_{t+1}$.

Finally, we observe that if $b \notin B(\theta_{m,t+1})$, then $G_i = 0$ for all $i \in A_{m+1}$.

Taking now the expectation of $\sum_{i \in A_{m+1}} G_i$, we get

$$\mathbb{E}[\sum_{i \in A_{m+1}} G_i] \geq -nD_i \Pr[b \in B(\theta_{m,t})] + (U_{t+1} - 2n\epsilon_{t+1}) \Pr[b \in B(\theta_{m,t+1}) \setminus B(\theta_{m,t})] \enspace .$$

By induction hypothesis $\Pr[b \in B(\theta_{m,t})] = 0$. As $U_{t+1} - 2n\epsilon_{t+1} > 0$ by definition of ϵ_{t+1} and $\mathbb{E}[G_i] \leq 0$ for all $i \in A_{m+1}$ by equilibrium property, this means that $\Pr[b \in B(\theta_{m,t+1}) \setminus B(\theta_{m,t})] = 0$. ☐

THEOREM 3.2. *In any correlated equilibrium, the set $f(b)$ is a social-welfare maximizing basis with probability* 1.

PROOF. For the sake of readability, for this proof, we assume that there are no two players having the exact same type. This way, the social-welfare maximizing basis OPT is unique. The proof with multiple optimal bases works the same way but is notationally cumbersome.

Again, without loss of generality, let the players be ordered such that $\theta_1 < \theta_2 < \ldots < \theta_n$. We now show by complete downward induction that $\Pr[i \in f(b)] = 1$ if $i \in OPT$ and $\Pr[i \in f(b)] = 0$ otherwise. So, let us fix some i and assume that we have already shown this claim for $i+1, \ldots, n$. If $i \notin OPT$, clearly $\Pr[i \in f(b)] = 0$ because with probability 1 among $i+1, \ldots, n$ exactly the players in OPT win. As OPT can be considered the result of a greedy selection, we know that i cannot be added. So, it only remains to show that $\Pr[i \in f(b)] = 1$ if $i \in OPT$. Let now be $j < i$ the largest index such that i is not included in every basis of the matroid restricted to $j, \ldots, n$. Observe that for every $b_i > \theta_j$, player i will be included in the output due to no-overbidding. Furthermore, by the above lemma, we know that $\Pr[i \in f(b), b_i < \theta_j] = 0$. Note that, as we have not fixed a tie-breaking rule, it may occur that i is sometimes included in the output and sometimes is not if $b_i = \theta_j$.

Let us define $p = \Pr[i \notin f(b), b_i \leq \theta_j]$ and $q = \Pr[i \in f(b), b_i = \theta_j]$. We would like to show that $p = 0$. So, let us assume that $p > 0$. In this case, let $U = \min_{\theta_{-i} \in \Theta_{-i}} w_i(\theta_j, \theta_{-i})$ and let $\delta > 0$ such that $w_i(\theta_j + \delta, b_{-i}) \geq w_i(\theta_j, b_{-i}) - \frac{Up}{2}$ for all $b_{-i} \in \Theta_{-i}$. Such a δ exists due to uniform continuity.

Now consider the deviation $\phi\colon \Theta_i \to \Theta_i$ such that $\phi(b_i) = \theta_j + \delta$ for $b_i \leq \theta_j$ and $\phi(b_i) = b_i$ otherwise. Under this deviation, every bid becomes a winning bid because i is included in every basis of the matroid restricted to $j+1, \ldots, n$ and we have no-overbidding. There is gain in utility of at least $(U - \frac{Up}{2})p$ from winning, and a small loss of utility of at most $\frac{Up}{2}q$ due to the increased price when i was winning. The overall expected gain in utility is at least $(U - \frac{Up}{2})p - \frac{Up}{2}q \geq \frac{Up}{2}$. This is a contradiction if $p > 0$. □

3.3. Uniqueness of Winning Bids

So far, we have shown that under very mild conditions the set of winners is unique. Let us now additionally assume that the functions w_i are strictly increasing in b_i, like, for example, in first-price auctions with quasi-linear utilities, or a smooth function modeling risk aversion. In this case, uniqueness is even stronger because not only the set of winners is unique but also their respective bids.

THEOREM 3.3. *Let $i \in OPT$ be a player such that w_i is strictly increasing in b_i. Furthermore, let $j \in N$, $j \neq i$ be the player maximizing θ_j such that $(OPT \setminus \{i\}) \cup \{j\}$ is a basis. Then in any correlated equilibrium $\Pr[b_i \in [\theta_j, \theta_j + \delta]] = 1$ for all $\delta > 0$.*

PROOF. We have already shown that i wins with probability 1 but never wins with a bid below θ_j, thus $\Pr[b_i < \theta_j] = 0$. To show $\Pr[b_i > \theta_j + \delta] = 0$ for $\delta > 0$, consider the deviation $\phi\colon \Theta_i \to \Theta_i$ such that $\phi(b_i) = \theta_j + \frac{\delta}{2}$ for $b_i \geq \theta_j$ and $\phi(b_i) = b_i$ otherwise. In addition, define $m\colon \Theta_{-i} \to \mathbb{R}$ by $m(b_{-i}) = w_i(\theta_j + \frac{\delta}{2}, b_{-i}) - w_i(\theta_j + \delta, b_{-i})$. As w_i is a continuous function, m also is. Furthermore, $m(b_{-i}) > 0$ for all b_{-i} and Θ_{-i} is a compact set. Therefore $M := \min_{b_{-i} \in \Theta_{-i}} m(b_{-i})$ exists and is positive.

Observe that any bid $b_i > \theta_j$ ensures that player i wins. This means that under the deviation, every bid remains a winning bid. Due to monotonicity of w_i, the player's utility increases by at least M in this case. So, we can lower-bound the gain by the deviation by $M \Pr[b_i \geq \theta_i + \delta]$. This implies $\Pr[b_i \geq \theta_i + \delta] = 0$. □

3.4. Necessity of Conditions

We conclude this section with three examples that show that each of the assumptions on which our uniqueness result is based—correlated equilibria, matroid feasibility structure, and exact optimization—is necessary. More specifically, we demonstrate that a unilateral relaxation of each of these assumptions can lead to the existence of more than one equilibrium.

Example 3.4 (*Correlated Equilibria*). Nadav and Piliouras [2010] consider the following model of Bertrand competition: The utility of player i is $u_i(p) = D_i(p)(p_i - c_i)$, where $D_i(p) = 0$ if $p_i > p_j$ for some j and $D_i(p) = (a - p_i)/(b(m+1))$ otherwise, where m is the number of players j that are tied with player i. Let $\pi(p)$ denote the utility for price p if there is a unique winner. Then there are several coarse correlated equilibria parameterized by $c < \alpha < \beta < \gamma \leq (a+c)/2$ in which $\pi(\gamma) \in (0, \pi((a+c)/2)]$, $\pi(\beta) = (1/(2n-1))^{n-1} \cdot \pi(\gamma)$, and $\pi(\alpha) = (n-1)/(2n-1) \cdot \pi(\beta)$. The payoff of each player in the corresponding coarse correlated equilibrium is $1/2 \cdot \pi(\alpha)/n + 1/2 \cdot [(1-(n-1)/n \cdot \rho) \cdot \pi(\beta)]$, where $\rho = (\pi(\beta)/\pi(\gamma))^{1/(n-1)}$.

Example 3.5 (*Matroid Feasibility Constraints*). Consider a setting with three players $N = \{1, 2, 3\}$ and types $\theta_1 = \theta_2 = 2$, $\theta_3 = 3$. Feasible solutions are $\{1, 2\}$ and $\{3\}$ and subsets. This is clearly not a matroid because the maximal feasible sets have different sizes. Furthermore, suppose that Vickrey-Clarke-Groves (VCG) mechanism is used. Then there are multiple pure Nash equilibria: E.g. $b = (0, 0, 3)$ and $b' = (2, 2, 0)$. The equilibrium b' is socially optimal, b is not.

Example 3.6 (*Exact Optimization*). Consider a setting with five players $N = \{a, b, c, d, e\}$ and feasible set of players 2^N. Suppose the goal is to maximize social welfare, and that the players have a value of 3, 3, 2, 1 and 1 for winning. An algorithm that does not optimize social welfare considers the sets $\{a, b\}$, $\{c\}$, or $\{d, e\}$ and chooses the set of bidders among these sets of bidders with the highest sum of bids. There are two pure Nash equilibria, namely $(2, 0, 2, 0, 0)$ and $(0, 2, 2, 0, 0)$. In these equilibria bidders $\{a, b\}$ win and pay 2 and 0 resp. 0 and 2. Also note that there exists a correlated equilibrium that is not a mixture over pure Nash equilibria, namely bids $(2, 0, 2, 0, 0)$, $(0, 2, 2, 0, 0)$, and $(0, 0, 1, 1, 1)$ played with probabilities $2/5$, $2/5$, and $1/5$.

4. POLYMATROID FEASIBILITY STRUCTURE

Being able to prove uniqueness for binary single-parameter problems with matroid structure, the natural next setting to consider are integer single-parameter problems with polymatroid structure. In these problems players can win multiple units, and the feasibility constraint is a submodular set function. We focus on settings in which the utilities are additively separable across units, with the contribution from each unit satisfying the same conditions as in the binary case.

We start by formally defining the problems and mechanisms that we study in this section. Afterwards we present two counterexamples that show that the uniqueness results that we obtained for matroids generally do not carry over to polymatroids. We conclude by introducing a restricted class of polymatroids, and showing that for this restricted class our results do carry over.

4.1. Definitions

Integer mechanism design problem. An integer single-parameter mechanism design problem is defined by the triple (N, X, Θ). As before N denotes the set of players. A feasible outcome $x \in X \subseteq \mathbb{N}_+^n$ now corresponds to an n-dimensional vector, where $x_i \in \mathbb{N}_+$ specifies how many units player i wins. The type space is again $\Theta = \prod_{i=1}^n \Theta_i$, where $\theta_i \in \Theta_i = [\theta_{\min}, \theta_{\max}] \subseteq \mathbb{R}$ represents the private information held by agent i. The welfare of a feasible outcome $x \in X$ is given by $\sum_{i \in N} x_i \cdot \theta_i$.

Direct mechanism. As in the binary case we consider direct mechanisms in which the players' bid to the mechanism is a reported type, and we use b for the bids and θ for the types. More formally, a direct mechanism $M = (f, p)$ now consists of an outcome rule $f : \Theta \to X$ and a payment rule $p : \Theta \to \mathbb{R}^n$.

Player utilities. We assume that player i's utility for winning x_i units is $x_i \cdot w_i(b)$, where w_i satisfies the same conditions as in the binary case. We assume that player i's utility is zero otherwise. This is consistent with our assumptions in case of matroid structures as player i has zero utility if $x_i = 0$ and utility $w_i(b)$ if $x_i = 1$.

Polymatroid optimization. We assume that the feasible vectors form a polymatroid. That is, a vector $x \in X$ is feasible if $\sum_{i \in S} x_i \leq g(S)$ for all $S \subseteq N$, where g is an integer-valued submodular function. A function g is submodular if $g(S \cap T) + g(S \cup T) \leq g(S) + g(T)$ for all $S, T \subseteq N$.

Our goal is to compute a maximal feasible vector $x \in X$ that has maximal welfare. We refer to a mechanism that computes such a vector as optimizing.

Many practical problems have polymatroid structure (see, e.g., Section 5 of [Bikhchandani et al. 2011]). We use the following two problems as running examples, where our uniqueness proof extends to the second example but not to the first one:

— In a sponsored search auction [Goel et al. 2012] n advertisers must be assigned clicks coming from k slots. Slot j receives α_j clicks, and slots are sorted such that $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_k$. An allocation of clicks x is feasible if and only if $\sum_{i \in S} x_i \leq g(S) = \sum_{j=1}^{|S|} \alpha_j$ for all subsets of advertisers S. It is easy to verify that the function g is submodular.

— In a Bertrand network [Babaioff et al. 2013] n firms have access to different markets having different sizes. The market structure is given by an undirected network $G = (V, E)$, where each node u represents a firm, self edges (u, u) represent captive markets, and all other edges (u, v) represent shared markets. An allocation of buyers x is feasible if and only if $\sum_{i \in S} x_i \leq g(S) = \sum_{e=(u,v): u \in S \vee v \in S} \text{size}(e)$ for all subsets of firms S. It is again easy to verify that g is submodular.

Solution concept. As before we assume that players act strategically, and analyze behavior in correlated equilibria.

4.2. Counterexamples

It turns out that our uniqueness results for matroids do not carry over to polymatroids; even in the simple case in which the mechanism uses first-price payments and the utilities are quasi linear.

A first—and relatively simple—observation is that correlated equilibria no longer coincide with a socially optimal solution even though the mechanism is optimizing. To see this consider a sponsored search auction with two advertisers and two slots. The first advertiser has a value of 1 and the second advertiser has a value of $1/2$. Suppose that ties are broken in favor of the first advertiser. The first slot receives 2 clicks and the second slot receives 1 click. The socially optimal solution would be to assign advertiser 1 to slot 1 and advertiser 2 to slot 2. An inefficient mixed Nash (and hence correlated) equilibrium has the advertisers bid $b_1, b_2 \in [0, 1/4]$ such that the cumulative distribution of bids is

$$F_1(b_1) = \frac{b_1}{1/2 - b_1} \quad \text{and} \quad F_2(b_2) = \frac{1/2 + b_2}{1 - b_2}.$$

In fact, a slightly more complicated example shows that correlated equilibria need not be unique. For example, in Section 6.5 of their paper Babaioff et al. [2013] describe a Bertrand network. As they outline in Appendix A.4 of their paper, there are two distinct mixed Nash (and hence correlated) equilibria.

It is vital to these counter examples that the players are confronted with a trade-off: They can choose to take a safe strategy, securing them some fraction at a "cheap price", or they can decide for a more offensive strategy, competing for a larger share. Indeed we will show that it is this trade-off that prevents the uniqueness results that we obtained for the matroid setting.

4.3. Uniqueness Results

We will show that the uniqueness results for matroids carry over to polymatroids if each player's competition is homogeneous:[3] Fix some player $i \in N$ and for $a < \theta_i$ consider the polymatroid restricted to players of type at least a. Then one of the following conditions is fulfilled:

— There is a basis x with $x_i = 0$.
— For every basis x, the value of x_i is identical.

The key step in the proof of the more general uniqueness results is the following lemma which generalizes Lemma 3.1 from matroids to polymatroids with homogeneous competition. The lemma naturally leads to generalizations of Theorem 3.2 and Theorem 3.3, which we prove in the full version.

LEMMA 4.1. *Given a polymatroid with homogeneous competition, let $i \in N$ be some arbitrary player and let $b_i < \theta_i$. If there is a basis x of the polymatroid restricted to players of type at least b_i with $x_i = 0$, then i never wins anything with bid b_i.*

PROOF. To show the claim, we reduce the polymatroid to a matroid using the transformation described by Schrijver [2002, Section 44.6b]. This way, from a polymatroid in n dimensions, we get a matroid of nk elements, where each agent controls k elements. That is, each agent i still only reports a single type $b_i \in \Theta_i$ but this weight is applied to all elements (i, ℓ), $\ell \in [k]$, in this matroid.

In general, it is not possible to apply the proof of Lemma 3.1 to this matroid because an agent controls multiple elements. However, as we will see, in this special case the proof is still applicable.

So, again let the players be ordered such that $\theta_1 \leq \theta_2 \leq \ldots \leq \theta_n$. For each $m \in [n]$, let $A_m \subseteq \{m, \ldots, n\}$ be the set of players i such that there is a basis not containing any element (i, ℓ) of the matroid restricted to $m, \ldots, n$. We claim that $\Pr[b_i < \theta_m, f_i(b) > 0] = 0$ for every $i \in A_m$. Again, we proceed by induction on m. So let us assume the statement has already been shown for some fixed m. Let $B(x) \subseteq \Theta$ be the set of bid vectors b such that there is some $i \in A_{m+1}$ with $b_i < x$ and $f_i(b)$. We will show that $\Pr[b \in B(\theta_{m+1})] = 0$.

We define the sequence $(\theta_{m,t})_{t \in \mathbb{N}}$ exactly as stated in the proof of Lemma 3.1 with the only difference that $\epsilon_{t+1} = \frac{U_{t+1}}{2nk+1}$.

For a fixed bid vector $b \in \Theta$, we consider for each agent $i \in A_{m+1}$ the deviation $\phi \colon \Theta_i \to \Theta_i$ such that $\phi(\theta_i) = \theta_{m,t+1}$ for $\theta_i \in [\theta_{m,0}, \theta_{m,t+1}]$ and $\phi(\theta_i) = \theta_i$ otherwise. Let $G_i = u_i(\phi(b_i), b_{-i}) - u_i(b)$ be agent i's gain in utility by this deviation.

[3]Note that matroids naturally fulfill this condition because as soon as there is no basis x with $x_i = 0$, we know that always $x_i = 1$.

If $b \in B(\theta_{m,t})$, we can again use the trivial bound $\sum_{i \in A_{m+1}} G_i \geq -nkD_i$, where $D_i = \max_{\theta \in \Theta} w_i(\theta) - \min_{\theta \in \Theta} w_i(\theta)$. Furthermore, if $b \notin B(\theta_{m,t+1})$, then $G_i = 0$ for all $i \in A_{m+1}$

So, let us turn to the case that $b \in B(\theta_{m,t+1}) \setminus B(\theta_{m,t})$. By definition, there is an $i \in A_{m+1}$, winning with a bid $b_i \in [\theta_{m,t}, \theta_{m,t+1})$. Let $S_1 \subseteq \{m+1, \ldots, n\} \times [k]$ be the corresponding basis of the matroid restricted to agents $m+1, \ldots, n$. There has to be $(i, \ell) \in S_1$ for some ℓ but we know that there is another basis of the restricted matroid such that $(i, \ell) \notin S_2$ for all ℓ. Let us define $S_1' = S_1 \setminus \{(i, \ell) \mid \ell \in [k]\}$. By augmentation property, there is an $(j, \ell') \in S_2 \setminus S_1'$ such that $S_1' \cup \{(j, \ell')\}$ is independent. We now have $G_j \geq U_{t+1} - \epsilon_{t+1} - (k-1)2\epsilon_{t+1}$ and $G_{i'} \geq -k2\epsilon_{t+1}$ for all $i' \neq j$. That is, $\sum_{i \in A_{m+1}} G_i \geq U_{t+1} - 2nk\epsilon_{t+1}$. So, we get

$$\mathbb{E}\left[\sum_{i \in A_{m+1}} G_i\right] \geq -nkD_i \Pr[b \in B(\theta_{m,t})] + (U_{t+1} - 2nk\epsilon_{t+1}) \Pr[b \in B(\theta_{m,t+1}) \setminus B(\theta_{m,t})] \ .$$

By induction hypothesis $\Pr[b \in B(\theta_{m,t})] = 0$. Furthermore, $U_{t+1} - 2nk\epsilon_{t+1} > 0$ by definition of ϵ_{t+1} and $\mathbb{E}[G_i] \leq 0$ for all $i \in A_{m+1}$ by equilibrium property. This means that $\Pr[b \in B(\theta_{m,t+1}) \setminus B(\theta_{m,t})] = 0$. □

THEOREM 4.2. *Given a polymatroid with homogeneous competition and social-welfare maximizing basis x^{OPT}. Then $\Pr[f_i(b) = x_i^{OPT}] = 1$ for all $i \in N$.*

PROOF. Let the players be ordered such that $\theta_1 \leq \theta_2 \leq \ldots \leq \theta_n$. We show by complete downward induction on i that $f_i(b) = x_i^{OPT}$. In Lemma 4.1, we have already shown that there is some $j < i$ such that $\Pr[f_i(b) = 0 \mid b_i < \theta_j] = 1$ and for any basis x on the matroid restricted to $j+1, \ldots, n$ the value x_i is identical—in other words, it is x_i^{OPT}. By the same arguments as used in the proof of Theorem 3.2, we now have that $\Pr[f_i(b) = x_i^{OPT}] = 1$. □

THEOREM 4.3. *Given a polymatroid with homogeneous competition and social-welfare maximizing basis x^{OPT}. Let $i \in N$ be a player such that $x_i^{OPT} > 0$ and w_i is strictly increasing in b_i. Furthermore, let $j \in N$, $\theta_j < \theta_i$, be the player maximizing θ_j with the following property: There is a basis x of the polymatroid restricted to players of type at most θ_j with $x_i = 0$. Then $\Pr[b_i \in [\theta_j, \theta_j + \epsilon]] = 1$ for all $\epsilon > 0$.*

PROOF. By homogeneous-competition property, for all bases x of the polymatroid restricted to players of type at most θ_j, we know that $x_i = x_i^{OPT}$. Furthermore, Theorem 4.2 shows that $\Pr[f_i(b) = x_i^{OPT}] = 1$. By the same considerations as applied when proving Theorem 3.3, player i's reports never exceed $\theta_j + \epsilon$ for any ϵ because $\theta_j + \frac{\epsilon}{2}$ ensures winning the same amount at a higher utility. □

4.4. Applications

An interesting case in which our uniqueness results for polymatroids applies are Bertrand networks in which no firm has a captive market.

THEOREM 4.4. *Bertrand networks without captive markets satisfy the homogeneity requirement of our uniqueness result. The results above imply that in the unique correlated equilibrium all firms bid zero.*

This theorem can be seen as a generalization and strengthening of the classic Bertrand paradox for duopolies and pure Nash equilibria to a Bertrand paradox for networked markets and correlated equilibria.

ACKNOWLEDGMENTS

We greatly benefitted from discussions with Pooya Jalaly. Part of the work of Paul Dütting was done while he was a Postdoctoral Fellow at Cornell University. He is supported by an SNF Postdoctoral Fellowship. Thomas Kesselheim is supported by a DAAD Postdoctoral Fellowship and NSF grant CCF-0910940. Éva Tardos is supported in part by NSF grants CCF-0910940 and CCF-1215994, ONR grant N00014-08-1-0031, a Yahoo! Research Alliance Grant, and a Google Research Grant.

REFERENCES

S. Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *Proceedings of the 52nd Symposium on Foundations of Computer Science*, pages 512–521, 2011.

E. Amann and W. Leininger. Asymmetric all-pay auctions with incomplete information: The two- player case. *Games and Economic Behavior*, 14:1–18, 1996.

R. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.

M. Babaioff, B. Lucier, and N. Nisan. Bertrand networks. In *Proceedings of the 14th Conference on Electronic Commerce*, pages 33–34, 2013.

P. Bajari. Comparing competition and collusion: a numerical approach. *Economic Theory*, 18:187–205, 2001.

M. R. Baye and J. Morgan. A folk theorem for one-shot Bertrand games. *Economics Letters*, 65:59–65, 1999.

S. Bikhchandani, S. de Vries, J. Schummer, and R. V. Vohra. An ascending vickrey auction for selling bases of a matroid. *Operations Research*, 59(2):400–413, 2011.

A. Blum and Y. Mansour. Learning, regret minimization, and equilibria. In N. N. et al, editor, *Algorithmic Game Theory*. Cambridge University Press, 2007.

A. Blum, E. Even-Dar, and K. Ligett. Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games. In *Proceedings of the 5th Symposium on Principles of Distributed Computing*, pages 45–52, 2006.

I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. P. Leme, and É. Tardos. On the efficiency of equilibria in generalized second price auctions. *Working paper available from: http://arxiv.org/abs/1201.6429*, pages 1–41, 2012.

S. Chawla and J. D. Hartline. Auctions with unique equilibria. In *Proceedings of the 14th Conference on Electronic Commerce*, pages 181–196, 2013.

S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42th Symposium on Theory of Computing*, pages 311–320, 2010.

C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39:195–259, 2009.

P. Dütting and R. Kleinberg. Polymatroid prophet inequalities. *Working paper available from: http://arxiv.org/abs/1307.5299*, pages 1–13, 2013.

P. Dütting, V. Gkatzelis, and T. Roughgarden. The performance of deferred-acceptance auctions. In *Proceedings of the 15th Conference on Economics and Computation*, 2014a. To appear.

P. Dütting, T. Roughgarden, and I. Talgam-Cohen. Incentives and greed in double auctions. In *Proceedings of the 15th Conference on Economics and Computation*, 2014b. To appear.

E. Even-Dar, Y. Mansour, and U. Nadav. On the convergence of regret minimization dynamics in concave games. In *Proceedings of the 41st Symposium on Theory of Computing*, pages 523–532, 2009.

G. Goel, V. S. Mirrokni, and R. Paes Leme. Polyhedral clinching auctions and the adwords polytope. In *Proceedings of the 44th Symposium on Theory of Computing*,

pages 107–122, 2012.
M. Hajiaghayi, R. Kleinberg, and T. W. Sandholm. Automated mechanism design and prophet inequalities. In *Proceedings of the 22nd Conference on Artificial Intelligence*, pages 58–65, 2007.
S. Hart and A. Mas-Colell. *Simple Adaptive Strategies: From Regret-Matching to Uncoupled Dynamics*. World Scientific Publishing, 2013.
M. O. Jackson. A crash course in implementation theory. *Social Choice and Welfare*, 18:655–708, 2001.
A. X. Jiang and K. Leyton-Brown. Polynomial-time computation of exact correlated equilibrium in compact games. In *Proceedings of the 12th Conference on Electronic Commerce*, pages 119–126, 2011.
T. R. Kaplan and D. Wettstein. The possibility of mixed-strategy equilibria with constant-returns-to-scale technology under Bertrand competition. *Spanish Economic Review*, 2:65–71, 2000.
R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities. In *Proceedings of the 44th Symposium on Theory of Computing*, pages 123–136, 2012.
R. Kleinberg, K. Ligett, G. Piliouras, and Éva Tardos. Beyond the Nash equilibrium barrier. In *Proceedings of the 2nd Innovations in Theoretical Computer Science*, pages 125–140, 2011.
B. Lebrun. Uniqueness of the equilibrium in first-price auctions. *Games and Economic Behavior*, pages 131–151, 2006.
E. Maskin and J. Riley. Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior*, 45:395–409, 2003.
E. Maskin and T. Sjöström. Chapter 5: Implementation theory. In A. S. K.J Arrow and K. Suzumura, editors, *Handbook of Social Choice and Welfare, Volume 1*, pages 238–288, 2002.
P. Milgrom and J. Roberts. Rationalizability, learning and equilibrium in games with strategic complementarities. *Econometrica*, 58:1255–1278, 1990.
P. Milgrom and I. Segal. Deferred-acceptance auctions and spectrum reallocation. In *Proceedings of the 15th Conference on Economics and Computation*, 2014. To appear.
U. Nadav and G. Piliouras. No regret learning in oligopolies: Cournot vs. Bertrand. In *Proceedings of the 3rd Symposium on Algorithmic Game Theory*, pages 300–311, 2010.
T. R. Palfrey. Chapter 61: Implementation theory. In R. Aumann and S. Hart, editors, *Handbook of Game Theory, Volume 3*, pages 2273–2326, 2002.
C. H. Papadimitriou and T. Roughgarden. Computing correlated equilibria in multiplayer games. *Journal of the ACM*, 2008.
T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st Symposium on Theory of Computing*, pages 513–522, 2009.
T. Roughgarden. The price of anarchy in games of incomplete information. In *Proceedings of the 13th Conference on Electronic Commerce*, pages 862–879, 2012.
A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer Verlag, 2002.
V. Syrgkanis and É. Tardos. Composable and efficient mechanisms. In *Proceedings of the 45th Symposium on Theory of Computing*, pages 211–220, 2013.
D. Topkis. *Ordered Optimal Decisions*. Stanford University, 1968.
X. Vives. Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics*, 19:305321, 1990.
J. Wu. Correlated equilibrium of Bertrand competition. In *Proceedings of the 4th International Workshop on Internet and Network Economics*, pages 166–177, 2008.