

GRAPHICAL POTENTIAL GAMES

YAKOV BABICHENKO AND OMER TAMUZ

ABSTRACT. We study the class of potential games that are also graphical games with respect to a given graph G of connections between the players. We show that, up to strategic equivalence, this class of games can be identified with the set of Markov random fields on G .

From this characterization, and from the Hammersley-Clifford theorem, it follows that the potentials of such games can be decomposed to local potentials. We use this decomposition to strongly bound the number of strategy changes of a single player along a better response path. This result extends to generalized graphical potential games, which are played on infinite graphs.

CONTENTS

1. Introduction.	2
1.1. Related Literature.	3
Acknowledgments.	4
2. Example.	4
3. Definitions.	5
3.1. Games.	5
3.2. Graphical games.	5
3.3. Potential games.	6
3.4. Markov random fields.	6
4. Characterization of graphical potential games.	6
4.1. Decomposing graphical potential games.	6
4.2. Potentials of graphical potential games.	8
4.3. Payoffs of graphical potential games.	10
4.4. Equivalence of graphical potential games and Markov random fields.	11
5. Better-response paths in graphical potential games.	12
5.1. Proof of Theorem 5.3.	16
6. Generalized graphical potential games.	19
References	21
Appendix A. Markov properties of positive random fields.	22

Date: March 23, 2016.

1. INTRODUCTION.

Potential games form an important class of strategic interactions. They includes fundamental interactions such as Cournot oligopolies (see, e.g., [13]), congestion games (see, e.g., Monderer and Shapley [13] or Rosenthal [18]), routing games (see, e.g., Rosenthal [18]) and many others. The above mentioned interactions are frequently *local* in nature. Namely, there exists an underlying graph such that the payoff of a player depends on her own strategy and on the strategies of her neighbors, but does not depend on the strategies of the opponents who are not neighbors. For instance, the locality of the interaction could be geographical: In routing games, the outcome of a driver depends only on her own route and the routes that were chosen by drivers who are geographically close to her. In a Cournot oligopoly where transportation costs are high, a firm is competing only with firms which are geographically close (for instance, this is the case with natural gas market; see Victor, Jaffe and Hayes [22]). The idea of the locality of an interaction is captured by the notion of a *graphical game*, introduced by Kearns et al. [12]. These games and similar ones are also sometimes called *network games*; see Jackson and Zenou [10] for an extensive survey.

The goal of this paper is to understand the class of graphical potential games. First, we address the following questions: What characterizes the potential function of a graphical potential game? What characterizes the payoffs of the players in a potential graphical game? In Theorems 4.1, 4.4 and 4.5 we provide a complete answer to these questions. We show that

- (1) The potential function of a graphical potential game can be expressed as an additive function of *local potentials*, where each local potential corresponds to a maximal clique in the underlying graph, and the value of the local potential is determined by the strategies of the players in the maximal clique only. This condition is necessary for a potential game to be graphical, and every such potential is the potential of some graphical game.
- (2) Up to strategically equivalent transformations (see definition 4.3), the payoff of a player is the sum of the local potentials of the cliques to which she belongs.

The proof of these results is achieved by showing that, for a fixed graph G , the set of potentials of graphical games on G can be identified with the set of *Markov random fields* on G (see Section 3.4). The latter are a well studied class of probability distributions with certain graphical Markov properties. Having established this correspondence, the

Hammersley-Clifford theorem [9], a classical result on Markov random fields, yields the above characterization of graphical potential games.

Next, we use this characterization to study dynamics. A central class of dynamics that has been studied in the context of potential games is (strict) better-response dynamics, where at each period of time a single player updates her strategy to a (strictly) better one, with respect to the current strategies of the opponents; the basic observation regarding all potential games is that better-response dynamics always converge.

Such a sequence of unilateral improvements is called a *better response path*. We address the question: Are there properties of better response paths that are unique for *graphical* potential games? We focus on the *number of updates of a single player* along a better response path, and we prove in Theorem 5.3 and Corollary 5.6 that under mild uniformity assumptions on the game, and for graphs with slow enough growth (see Definition 5.5), the number of updates of each player is bounded by a constant, which in particular is independent of the number of players. This result is general and holds for all the above mentioned examples. Example 5.7 demonstrates that the slow-enough-growth condition is tight, in some sense.

Finally, in Section 6, we introduce *generalized graphical potential games* on infinite graphs. These are graphical games which are not necessarily potential games, but still have a local potential structure. We prove here the same bounds on strict better-response paths. We note that these results, as well as the results on finite graphs, translate to novel results on Markov random fields, which may be of interest outside of game theory.

1.1. Related Literature. Potential games and graphical games are both fundamental classes of games (see, e.g., Nisan [15]). There has recently been a growing interest in the intersection of these two classes, because many interesting potential games have a graphical structure (see e.g., Bilò et al. [3] or Bimpikis, Ilklic and Shayan [4]), and many interesting graphical games have a potential function (see e.g., Auletta et al. [1], Bramoullé, Kranton and D’amours [5]). However, to the best of our knowledge, this paper is the first to fully characterize the intersection of these two classes¹.

Best-response and better-response paths in potential games were studied in Fabrikant, Papadimitriou and Talwar [7], Skopalik and Vöcking [20] and Awerbuch et al. [2], where the focus is on the *length* of the paths. Our result (Theorem 5.3 and Corollary 5.6) focuses on the *number* of

¹In an independent, later work Ortiz [16] addresses a similar characterization problem.

changes of a single player. This aspect of better- and best-response dynamics also plays an interesting role in graphical games of strategic complements; see Jackson and Zenou [10].

The connection between graphical potential games and Markov random fields was previously observed for specific cases; see Auletta et al. [1], who establish a connection between the logit dynamic in a particular class of graphical potential coordination games and Glauber dynamics in the Markov random field of the Ising model. We show that this connection is much more general and extends to all graphical potential games.

Another intriguing connection between graphical games and Markov random fields was established by Kakade et al. [11], who show that every correlated equilibrium of every graphical game is a Markov random field. Daskalakis and Papadimitriou [6] use Markov random fields to compute pure Nash equilibria in graphical games.

There is a vast literature on *majority dynamics*, which can be interpreted as best-response dynamics of the majority game on a graph, which is a graphical potential game. Tamuz and Tessler [21] upper bound the number of changes in majority dynamics on slowly growing graphs. We adapt their technique in the proof of Theorem 5.3.

Acknowledgments. The authors would like to thank Elchanan Mossel for some enlightening comments.

2. EXAMPLE.

Before diving into the definitions and results, we provide the reader with a simple, informal example of a graphical potential game, and of a generalized graphical potential game. These will be games of pure *network externalities*.

Let V be a finite set of players located at the nodes of a graph $G = (V, E)$. Let some of the edges in E be blue and the rest be red. Pairs of players connected by a blue edge will benefit from choosing the same strategy, and pairs connected by a red edge will benefit from choosing different strategies. In particular, there are two strategies, and each player gains a unit of utility for each blue neighbor that chooses the same strategy, and loses a unit of utility for each red neighbor that chooses the same strategy.

This is a graphical game: the utility of each player is independent of the choices of those who are not her neighbors.

Given a strategy profile, consider the number of blue edges along which the two players choose the same strategy, minus the number of blue edges along which the two players choose different strategies. This

function of strategy profiles is readily seen to be a potential for this game, and thus this game is a graphical potential game.

We can define the same game for the case in which the set of players is countably infinite and the graph G is locally finite (i.e., each player has a finite number of neighbors). In this case the game clearly still is a graphical game. However, it is no longer a potential game. Still, in some ways this game resembles a potential game, and in particular it has a “local potential structure”. Accordingly, it falls into the class of generalized potential games (Definition 6.1).

An additional sense in which this game resembles a potential game is that it has pure equilibria (Theorem 6.2). However, unlike potential games, this game has infinite better response paths. A natural question is: given a particular player, does there exist a better response path in which this player makes infinitely many strategy changes? It turns out that the answer to this question depends on the graph G . We show that for some infinite graphs (for example, graphs with subexponential growth rates) each player changes strategy only a finite number of times, as in a finite potential game (Theorems 6.3 and 6.4). This bound on the number of changes also applies to large finite graphs (Theorem 5.3), so that, for some sequences of growing graphs, the maximum number of strategy changes does not depend on the size of the graph.

3. DEFINITIONS.

3.1. Games. Let \mathcal{G} be a game played by a finite set of players V , and denote $n = |V|$. Let \mathcal{S}_i be the finite set of strategies available to $i \in V$, and denote the set of strategy profiles by $\mathcal{S} = \prod_i \mathcal{S}_i$. Let $u_i: \mathcal{S} \rightarrow \mathbb{R}$ be player i 's utility function, which maps strategy profiles to payoffs. Note that we do not allow payoffs in $\{-\infty, \infty\}$.

3.2. Graphical games. We identify V with the set of nodes of a simple, undirected graph $G = (V, E)$. The game \mathcal{G} is a *graphical game* on G if each player's utility is a function of the strategies of her neighbors in G . This can be formally expressed by the following condition on the utility functions. Choose from each \mathcal{S}_i an arbitrary distinguished strategy o_i . Given $a \in \mathcal{S}$, denote $a^i = (a_1, \dots, a_{i-1}, o_i, a_{i+1}, \dots, a_n)$.

Definition 3.1. \mathcal{G} is a graphical game on G if

$$(3.1) \quad i \neq j \text{ and } (i, j) \notin E \quad \Rightarrow \quad \forall a \in \mathcal{S}, u_i(a) = u_i(a^j).$$

It is easy to verify that this definition does not depend on the choice of $\{o_i\}_{i \in V}$.

3.3. Potential games.

Definition 3.2. \mathcal{G} is a *potential game* if there exists a (potential) function $\Phi: \mathcal{S} \rightarrow \mathbb{R}$ such that for all $i \in V$ and $a \in \mathcal{S}$

$$(3.2) \quad u_i(a) - u_i(a^i) = \Phi(a) - \Phi(a^i).$$

Note that if Φ is a potential for \mathcal{G} then $\Phi + C$ is also a potential for \mathcal{G} , for any $C \in \mathbb{R}$. We make a canonical choice and assume always that

$$(3.3) \quad \sum_{a \in \mathcal{S}} e^{\Phi(a)} = 1.$$

The rationale behind this choice will become apparent below. Here we use the fact that utilities (and hence potentials) cannot be infinite. It is easy to see that given a potential game, all its potentials differ by a constant.

3.4. Markov random fields. Let $G = (V, E)$ be a finite, simple, undirected graph. Associate with each $i \in V$ a random variable X_i , and denote $X = (X_1, \dots, X_n)$. X is called a *random field* on G .

Let U and W be subsets of V . A subset $A \subset V$ is a (U, W) -cut in the graph if every path from U to W must pass through A . For any subset $Z \subset V$ define the random variable $X_Z = \{X_i : i \in Z\}$ to be the restriction of X to Z .

Definition 3.3. X is a *Markov random field* (MRF) if, for every (U, W) -cut A it holds that conditioned on X_A , X_U is independent of X_W .

The random field X is said to be *positive* if its distribution is equivalent (in the sense of mutual absolute continuity) to the product of its marginal distributions. For discrete distributions, this is equivalent to requiring the event $(X_1, \dots, X_n) = (a_1, \dots, a_n)$ to have positive probability whenever all of the events $X_i = a_i$ have positive probability.

For positive random fields, Markov random fields can be characterized by a weaker condition, namely that for every pair $i \neq j \in V$ such that $(i, j) \notin E$ it holds that, conditioned on $X_{V \setminus \{i, j\}}$, X_i is independent of X_j . While this seems to be a well-known fact, we were not able to find a reference for its proof, and therefore provide it (for finite distributions) in Appendix A.

4. CHARACTERIZATION OF GRAPHICAL POTENTIAL GAMES.

4.1. Decomposing graphical potential games. Our first result shows that the potential of a graphical potential game can be decomposed into local potentials.

For $W \subseteq V$ and $a \in \mathcal{S}$, let $a^W \in \prod_{i \in W} \mathcal{S}_i$ be the restriction of a to W , given by $(a^W)_i = a_i$, for $i \in W$.

Denote by $\mathcal{C}(G)$ the set of *maximal cliques* in G ; these are cliques which are not subsets of strictly larger cliques.

Theorem 4.1. *Let a game \mathcal{G} be both a graphical game on G and a potential game with potential Φ . Then Φ can be written as*

$$(4.1) \quad \Phi(a) = \sum_{C \in \mathcal{C}(G)} \Phi^C(a^C),$$

for some functions $\Phi^C: \prod_{i \in C} \mathcal{S}_i \rightarrow \mathbb{R}$.

The functions Φ^C are called *local potentials*.

Before proving this proposition we introduce the *Hessian* Φ_{ij} and prove a lemma. For $a \in \mathcal{S}$ and $i, j \in V$, denote

$$a^{ij} = (a_1, \dots, a_{i-1}, o_i, a_{i+1}, \dots, a_{j-1}, o_j, a_{j+1}, \dots, a_n).$$

We recall that o_i is an arbitrary distinguished strategy of player i . Note that $(a^i)^j = (a^j)^i = a^{ij} = a^{ji}$.

Let $\Phi_i: \mathcal{S} \rightarrow \mathbb{R}$ be given by

$$\Phi_i(a) = \Phi(a) - \Phi(a^i).$$

Alternatively, by (3.2), $\Phi_i(a) = u_i(a) - u_i(a^i)$.

The Hessian $\Phi_{ij}: \mathcal{S} \rightarrow \mathbb{R}$ is given by

$$\Phi_{ij}(a) = \Phi_i(a) - \Phi_i(a^j) = \Phi(a) - \Phi(a^i) - \Phi(a^j) + \Phi(a^{ij}).$$

Lemma 4.2. *Let a game \mathcal{G} be both a graphical game on G and a potential game with potential Φ . Then $\Phi_{ij} = 0$ for all $(i, j) \notin E$.*

Proof. Choose $(i, j) \notin E$. Then by (3.1) we have that $u_i(a) = u_i(a^j)$ and that $u_i(a^i) = u_i(a^{ij})$. Hence

$$\begin{aligned} 0 &= [u_i(a) - u_i(a^j)] - [u_i(a^i) - u_i(a^{ij})] \\ &= [u_i(a) - u_i(a^i)] - [u_i(a^j) - u_i(a^{ij})] \\ &= \Phi_i(a) - \Phi_i(a^j) \\ &= \Phi_{ij}(a). \end{aligned}$$

□

Proof of Theorem 4.1. Let \mathbb{P} be a probability distribution over the set of strategy profiles \mathcal{S} given by

$$(4.2) \quad \mathbb{P}[a] = e^{\Phi(a)}.$$

By (3.3) this is indeed a probability distribution. Let $X = (X_1, \dots, X_n)$ be a random field on G with law \mathbb{P} . Then X_i is a (random according

to \mathbb{P}) strategy played by player i . Note that X is a positive random field, although a priori it may not be Markov.

We will prove the theorem by showing that for every $(i, j) \notin E$, the random variables X_i and X_j are independent, conditioned on $\{X_k\}_{k \notin \{i, j\}}$. This will prove that X is a Markov random field over the graph G , and so the claim will follow immediately from the Hammersley-Clifford Theorem (see Hammersley and Clifford [9], and also Grimmett [8], Preston [17] and Sherman [19]) for positive MRFs, which states that \mathbb{P} can be decomposed as

$$\log \mathbb{P}[X = a] = \sum_{C \in \mathcal{C}(G)} \Phi^C(a^C),$$

for some functions $\Phi^C: \prod_{i \in C} \mathcal{S}_i \rightarrow \mathbb{R}$.

Choose $(i, j) \notin E$. Then by Lemma 4.2 we have that, for every $a \in \mathcal{S}$,

$$\Phi(a) + \Phi(a^{ij}) = \Phi(a^i) + \Phi(a^j).$$

By (4.2), this can be written as

$$(4.3) \quad \mathbb{P}[X = a] \cdot \mathbb{P}[X = a^{ij}] = \mathbb{P}[X = a^i] \cdot \mathbb{P}[X = a^j].$$

Denote by $\mathcal{X}^{ij}(a) = \{X_k = a_k\}_{k \notin \{i, j\}}$ the event that $X_k = a_k$ for all $k \notin \{i, j\}$. Then we can write (4.3) as

$$\begin{aligned} & \mathbb{P}[X_i = a_i, X_j = a_j, \mathcal{X}^{ij}(a)] \cdot \mathbb{P}[X_i = o_i, X_j = o_j, \mathcal{X}^{ij}(a)] \\ &= \mathbb{P}[X_i = o_i, X_j = a_j, \mathcal{X}^{ij}(a)] \cdot \mathbb{P}[X_i = a_i, X_j = o_j, \mathcal{X}^{ij}(a)]. \end{aligned}$$

Hence

$$\begin{aligned} & \mathbb{P}[X_i = a_i, X_j = a_j | \mathcal{X}^{ij}(a)] \cdot \mathbb{P}[X_i = o_i, X_j = o_j | \mathcal{X}^{ij}(a)] \\ &= \mathbb{P}[X_i = o_i, X_j = a_j | \mathcal{X}^{ij}(a)] \cdot \mathbb{P}[X_i = a_i, X_j = o_j | \mathcal{X}^{ij}(a)]. \end{aligned}$$

Summing over all possible values of a_i and a_j we arrive at

$$\begin{aligned} & \mathbb{P}[X_i = o_i, X_j = o_j | \mathcal{X}^{ij}(a)] \\ &= \mathbb{P}[X_i = o_i | \mathcal{X}^{ij}(a)] \cdot \mathbb{P}[X_j = o_j | \mathcal{X}^{ij}(a)]. \end{aligned}$$

Finally, we recall that the choice of $\{o_i\}_{i \in V}$ was arbitrary, and so X_i and X_j are independent, conditioned on $\{X_k\}_{k \notin \{i, j\}}$. \square

4.2. Potentials of graphical potential games. Theorem 4.1 states that if a potential game is graphical, then the potential has the form (4.1). We cannot hope to have the opposite direction (i.e., if the potential function has the form (4.1) then the game is graphical) because of the following observation: If we add a constant payoff of c to player i whenever player $j \neq i$ plays a certain strategy a_j , then the potential

function does not change; on the other hand, if in the original game player j 's strategy does not influence player i 's payoff (i.e., there is no edge (i, j) in G), then in the new game (with the additional payoff of c) this is no longer the case. Such a change preserves the potential function but not the graphical structure. Therefore, we cannot deduce a result on the graphical structure of the game from the potential function.

However, the above mentioned change is a special type of change of the game in the following sense: it does not cause any strategic change in the game, because whether or not player i receives an additional payoff of $c \neq 0$ does not depend on her own behavior. This motivates the notion of *strategically equivalent games* introduced by Monderer and Shapley [13]. Denote by \mathcal{S}^{-i} the set of strategy profiles of players in $V \setminus \{i\}$.

Definition 4.3. Two games \mathcal{G} and \mathcal{G}' with payoff functions u and w over the same strategy profile set \mathcal{S} are called *strategically equivalent* if there exist functions f_1, \dots, f_n where $f_i : \mathcal{S}^{-i} \rightarrow \mathbb{R}$ such that $u_i(a) = w_i(a) + f_i(a_{-i})$. We say that \mathcal{G} is a *strategically equivalent transformation* of \mathcal{G}' .

Note again, that since the additional payoff of $f_i(a_{-i})$ does not depend on player i 's behavior, the transformation causes no strategic change. In particular, any (mixed) Nash equilibrium of \mathcal{G} is a (mixed) Nash equilibrium of \mathcal{G}' .

We next show that up to a strategically equivalent transformation of the game, the decomposition (4.1) is a necessary and sufficient condition for a potential game to be graphical.

Theorem 4.4. *A potential game \mathcal{G} with potential Φ is strategically equivalent to a graphical game on G if and only if the potential Φ has the form*

$$(4.4) \quad \Phi(a) = \sum_{C \in \mathcal{C}(G)} \Phi^C(a^C).$$

Proof. If \mathcal{G} is strategically equivalent to a graphical game \mathcal{G}' on G , then \mathcal{G}' is also a potential game with the same potential function Φ , because a strategic equivalent transformation preserves the potential function (see Monderer and Shapley [13]). By Theorem 4.1 Φ has the form (4.4).

For the opposite direction, let Φ be of the form (4.4), and consider the game \mathcal{G}' where player i 's payoff is defined by

$$(4.5) \quad u_i(a) = \sum_{C \in \mathcal{C}(G) : i \in C} \Phi^C(a^C).$$

We show that \mathcal{G}' is a potential game on the graph G with the potential Φ . This will complete the proof, because \mathcal{G} and \mathcal{G}' have the same potential, and therefore are strategically equivalent (see Monderer and Shapley [13] again).

Since $(a^i)^C = a^C$ whenever $i \notin C$, we have that $\Phi^C(a^C) - \Phi^C((a^i)^C) = 0$ whenever $i \notin C$. Hence

$$\begin{aligned} u_i(a) - u_i(a^i) &= \sum_{C \in \mathcal{C}(G) : i \in C} \Phi^C(a^C) - \Phi^C((a^i)^C) \\ &= \sum_{C \in \mathcal{C}(G)} \Phi^C(a^C) - \Phi^C((a^i)^C) \\ &= \Phi(a) - \Phi(a^i), \end{aligned}$$

where the first equality follows from (4.5), the second equality follows from the observation above, and the third from the fact that Φ has the form (4.4).

Finally, to see that \mathcal{G} is a graphical game over G , note that for $(i, j) \notin E$ we have that $(a^j)^C = a^C$ for all $C \in \mathcal{C}$ such that $i \in C$. Hence

$$\begin{aligned} u_i(a) - u_i(a^j) &= \sum_{C \in \mathcal{C}(G) : i \in C} \Phi^C(a^C) - \Phi^C((a^j)^C) \\ &= \sum_{C \in \mathcal{C}(G) : i \in C} \Phi^C(a^C) - \Phi^C(a^C) \\ &= 0. \end{aligned}$$

□

4.3. Payoffs of graphical potential games. Theorem 4.1 characterizes the *potential function* of a graphical potential game, but it does not characterize the *payoff functions* of the players in such a game. We next present an exact characterization of the payoff functions in a graphical potential game. For every player $i \in V$ let $N(i)$ be the set of neighbors of player i , excluding i . Denote by $\mathcal{S}^{N(i)}$ the set of strategy profiles of players in $N(i)$.

Theorem 4.5. *Let G be a graph, and let Φ be a potential of the form (4.1). For every choice of functions f_1, \dots, f_n where $f_i : \mathcal{S}^{N(i)} \rightarrow$*

\mathbb{R} , the game with the payoffs

$$(4.6) \quad u_i(a) = \sum_{C \in \mathcal{C}(G) : i \in C} \Phi^C(a^C) + f_i(a^{N(i)})$$

is a graphical potential game over G with the potential Φ .

Conversely, for every graphical potential game over G with potential Φ there exist functions f_1, \dots, f_n such that the payoffs are given by (4.6).

Proof. Let \mathcal{G} be a game with payoffs u_i , and let \mathcal{G}' be the game with payoffs

$$(4.7) \quad w_i(a) = \sum_{C \in \mathcal{C}(G) : i \in C} \Phi^C(a^C).$$

\mathcal{G}' is a graphical potential game over G with the potential Φ ; this was proved in Theorem 4.4. The strategically equivalent transformation $u_i = w_i + f_i(a^{N(i)})$ preserves the potential function Φ . Moreover it preserve the graphical structure as well, because f_i depends only on the strategy of the neighbors of i . Therefore \mathcal{G} is a graphical potential game over G with the potential Φ .

For the opposite direction, given a graphical potential game \mathcal{G} over G with the potential Φ , we know that the payoffs can be written as:

$$(4.8) \quad u_i(a) = \sum_{C \in \mathcal{C}(G) : i \in C} \Phi^C(a^C) + g_i(a_{-i})$$

for some functions $g_i : \mathcal{S}^{-i} \rightarrow \mathbb{R}$. This follows from the fact that the the games \mathcal{G} and \mathcal{G}' have the same potential function Φ , and therefore are strategically equivalent.

The game \mathcal{G} is graphical, therefore for every $j \notin N(i) \cup \{i\}$ it holds that $g(a_{-i}) = g((a_{-i})^j)$. Otherwise, j will have influence on i 's payoff. Therefore g does not depend on a_j , and it can be written as $g(a_{-i}) = f(a^{N(i)})$, which completes the proof. \square

4.4. Equivalence of graphical potential games and Markov random fields. In the proof of Theorem 4.1 we showed how every graphical potential game on G can be mapped to a Markov random field on G . We next show that this map is a bijection, so that every Markov random field on G can be mapped back to the potential of a graphical game on G .

Let $\text{MRF}(G, \mathcal{S}) \subset \Delta(\mathcal{S})$ denote the set of probability measures on \mathcal{S} which describe positive Markov random fields with underlying graph G . Let $\text{PGG}(G, \mathcal{S}) \subset \mathbb{R}^{\mathcal{S}}$ be the set of normalized (as in (3.3)) potentials

of graphical potential games over G with strategy profiles in \mathcal{S} . Define $\psi: \text{PGG}(G, \mathcal{S}) \rightarrow \text{MRF}(G, \mathcal{S})$ by

$$(4.9) \quad [\psi(\Phi)](a) = e^{\Phi(a)}.$$

This is the same mapping that we use in the proof of Theorem 4.1.

Proposition 4.6. *The map $\psi: \text{PGG}(G, \mathcal{S}) \rightarrow \text{MRF}(G, \mathcal{S})$ is a bijection.*

Proof. It was shown in the proof of Theorem 4.1 that the image of ψ is indeed in $\text{MRF}(G, \mathcal{S})$. Since ψ is clearly one-to-one, it remains to be shown that it is onto.

Let $\mathbb{P} \in \text{MRF}(G, \mathcal{S})$ be a distribution over \mathcal{S} that is a positive MRF over G . Then, by the Hammersley-Clifford Theorem it can be written as

$$\mathbb{P}[a] = \prod_{C \in \mathcal{C}(G)} e^{\Phi^C(a^C)},$$

for some functions $\Phi^C: \prod_{i \in C} \mathcal{S}_i \rightarrow \mathbb{R}$.

Define a game \mathcal{G} on V with strategies in \mathcal{S} by

$$u_i(a) = \sum_{C \in \mathcal{C}(G): i \in C} \Phi^C(a^C).$$

In the proof of Theorem 4.4 we showed that \mathcal{G} is a graphical potential game with the potential

$$(4.10) \quad \Phi(a) = \sum_{C \in \mathcal{C}(G)} \Phi^C(a^C),$$

and therefore $\psi(\mathcal{G}) = \mathbb{P}$. □

5. BETTER-RESPONSE PATHS IN GRAPHICAL POTENTIAL GAMES.

In this section we make the simplifying assumption that utilities take integer values, and in fact demand that each local potential is integer². Under these restrictions, we show how, in some families of games, one can bound the total number of changes of strategy *per player*, independently of the number of players in the game.

²The extension to real values is straightforward if instead of considering better-response paths (see definition below) one considers ε -better-response paths, which are paths where each updating player improves her payoff by at least ε .

Definition 5.1. A game \mathcal{G} with potential Φ over a graph G has *M-bounded local potentials* if in the decomposition of the potential (4.1) all the local potentials satisfy $|\Phi^C(a^C)| \leq M$ for all a^C . Here we do not require Φ to be normalized as in (3.3). \mathcal{G} is called an *integral game* if all the local potentials are always integer.

A *better-response path* of length L is defined to be a sequence of strategy profiles $(a(t))_{t=0}^L$ such that $a(t+1)$ differs from $a(t)$ in exactly one coordinate $i = i(t)$ and $a_i(t+1)$ is a *strict* better-response to $a_{-i}(t)$. Player $i = i(t)$ is called *the updating player at time t* . Note that strict best-response dynamics are a special case, and therefore our results will apply to them, too.

An immediate observation regarding better-response dynamics is that in an integral game, the length of a better-response path is bounded by the range of the potential. This follows from the fact that in each update the potential increases by at least one.

We analyze general better-response paths, without committing to a particular order in which the player sequence $\{i(t)\}$ is chosen. This analysis, in particular, applies to continuous time better-response dynamics (or best-response dynamics) in which the path is drawn by Poisson arrivals.

We define the *clique-degree* of a graph as the largest number of maximal cliques that any vertex in the graph participates in.³

Claim 5.2. *For every n -player potential game on the graph G with M -bounded local potentials, where the clique-degree of G is D , the potential of the game is bounded by $|\Phi(a)| < nDM$ for all $a \in \mathcal{S}$.*

The claim follows from the fact that each vertex may participate in at most D cliques.

An immediate consequence is that the length of any better-response path is at most nDM . Hence, the average number of updates performed by a player is at most DM , regardless of the size of the graph. The next claim, which is the main theorem of this section, shows that for graphs of slow enough growth (in a particular sense) we can bound by a constant the total number of updates performed by any single player, without having to average over all players.

Given a graph G , denote by $\Delta(i, j)$ the graph distance between two nodes i and j in G . This is the length of a shortest path between i

³Note that in a graph of constant degree d , the clique-degree is trivially bounded by the constant 2^d , because the number of cliques that a vertex i participates in is bounded by the number of subsets of neighbors of i . More interestingly, the number of maximal cliques is sharply bounded by $4 \cdot 3^{d/3}$, as was shown by Moon and Moser [14].

and j . Denote by $S_r(G, i)$ the number of vertices at exactly distance r from i :

$$S_r(G, i) = \#\{j : \Delta(i, j) = r\}.$$

Theorem 5.3. *Let \mathcal{G} be an integral potential game on the graph G with M -bounded local potentials, where G has clique-degree D . Then the number of times that a player k updates her strategy in any better-response path is at most*

$$2DM \sum_{r=0}^{\infty} \left(1 - \frac{1}{2DM}\right)^r S_r(G, k).$$

The proof appears in Section 5.1. Before that, we give an example of an application, introduce a consequence of this theorem on graphs that satisfy a “slow enough growth” condition, and an example that demonstrates the tightness of this condition.

The next example is a particular case of the game introduced in Section 2.

Example 5.4. Let \mathbb{Z}_n^2 be the $n \times n$ two dimensional grid, and consider the following game \mathcal{G}_n in which each player has two strategies. Label some subset of the edges blue and the rest red. The utility of each player is equal to the number of neighbors along blue edges that match her strategy, minus the number of neighbors along red edges that match her strategy. It is easy to see that \mathcal{G} is a graphical potential game on \mathbb{Z}_n^2 , that local potentials are 1-bounded, that the clique-degree of \mathbb{Z}_n^2 is four (since all maximal cliques are of size two and correspond to the edges), and that $S_r(\mathbb{Z}^2, i) \leq 4r$.

Hence by Theorem 5.3, in any better-response path, no player changes her strategy more than 1792 times, independently of n .

We generalize this example to a larger class of games and graphs.

Definition 5.5. Given a function $f : \mathbb{N} \rightarrow \mathbb{R}$ we say that the *growth of a graph G* is bounded by f if for every vertex i and every $r \in \mathbb{N}$ it holds that $S_r(G, i) \leq f(r)$.

In order to gain some intuition about the notion of the growth of the graph let us mention several graphs and their growths:

- For a line or a cycle with n vertices, the growth is bounded by $f(r) = 2$.
- The growth of the m -dimensional grid is bounded by $f(r) = (2mr)^{m-1}$, which is *polynomial in r* for a constant m .
- The growth of a binary tree is at least 2^r , and therefore is *exponential in r* .

A straightforward corollary from Theorem 5.3 is the following:

Corollary 5.6. *Fix M , and let G be a graph with click-degree D and with growth that is bounded from above by the following exponentially increasing function*

$$f(r) = c \left(1 + \frac{1}{4DM}\right)^r$$

for some constant $c > 0$. Then for every potential game on the graph G with M -bounded local potentials, the number of updates of every player along every better-response path is bounded by the constant $8cD^2M^2$. In particular, this bound does not depend on the size of the graph.

Note that Corollary 5.6 holds, in particular, for the case where the graph has any subexponential growth.

Proof of Corollary 5.6. By Theorem 5.3 we can bound the number of updates by

$$\begin{aligned} & 2DM \sum_{r=0}^{\infty} \left(1 - \frac{1}{2DM}\right)^r c \left(1 + \frac{1}{4DM}\right)^r \\ & \geq 2cDM \sum_{r=0}^{\infty} \left(1 - \frac{1}{4DM}\right)^r = 8cD^2M^2. \end{aligned}$$

□

The bound on the growth of the graph is crucial for the result of Theorem 5.3 (and Corollary 5.6), as demonstrated by the following example of a sequence of graphs of exponential growth, in which there is no uniform bound on the number of strategy changes a player makes.

Example 5.7. We denote by $BT_k = (V, E)$ the binary tree graph of depth k and $n = 2^{k+1} - 1$ vertices. For $l = 0, \dots, k$ we denote by $V_l = \{v_1^l, \dots, v_{2^l}^l\}$ the set of vertices at depth-level l .

We consider the majority game \mathcal{G} on BT_k in which each player has two strategies, and her utility is equal to the number of neighbors who played the same strategy. As in Example 5.4, it is easy to see that \mathcal{G} is a graphical potential game on BT_k , that local potentials are 1-bounded, and that the clique-degree of G is three, since again all maximal cliques correspond to edges. However, in this case $S_r(BT_k, i) \geq 2^r$ for $r < k$, and so we cannot hope to use Theorem 5.3 to bound the number of changes without dependence on n . Indeed, we show that no such bound exists, so that, as k grows, the number of strategy changes performed by the level 0 player diverges.

Consider the following better-response path.⁴ In the initial configuration $a(0)$ we set the strategy of all players to be 1 at even depth-levels and (-1) at odd depth-levels (i.e., $a(0)_v = (-1)^l$ for $v \in V_l$). The players update their strategies according to the following order, comprising k “waves” of updates. The first wave consists of an update of the player at depth 0 only. The second wave starts at all players at depth 1, and then continues with the player at depth 0. The j 'th wave starts with all players at depth $j - 1$, then continues with all players at depth $j - 2$, then $j - 3$ and so on, and finally ends with the player at depth 0. Formally, the order of updates is as follows:

$$\begin{aligned}
 & \left\{ \begin{array}{l} \text{Step 1.1: All the players in } V_0. \end{array} \right. \\
 & \left\{ \begin{array}{l} \text{Step 2.1: All the players in } V_1. \\ \text{Step 2.2: All the players in } V_0. \end{array} \right. \\
 & \left\{ \begin{array}{l} \text{Step 3.1: All the players in } V_2. \\ \text{Step 3.2: All the players in } V_1. \\ \text{Step 3.3: All the players in } V_0. \end{array} \right. \\
 & \vdots \\
 & \text{Step } l.j \text{ for } 1 \leq j \leq l \leq k: \text{ All the players in } V_{l-j}. \\
 & \vdots \\
 & \left\{ \begin{array}{l} \text{Step } k.k: \text{ All the players in } V_0. \end{array} \right.
 \end{aligned}$$

where at each step $l.j$ the players in V_{l-j} update their strategies in an arbitrary order.

Note that at each step $l.j$, every player $v \in V_{l-j}$ indeed updates her strategy, because the order is designed in such a way that both v 's neighbors at depth-level $l - j + 1$ play the opposite strategy, and player v has at most three neighbors. Note also that the root player v_1^0 updates her strategy $k = \log_2(\frac{n+1}{2})$ times. Namely the number of her updates is *not* bounded by a constant. This is possible on a sequence of growing binary trees because their growth rate is too high. As Example 5.4 above shows, no such construction is possible on a sequence of (polynomially) growing grids.

5.1. Proof of Theorem 5.3. Let G be a graph of clique-degree D with n nodes, and let i be a vertex in G . Let C be a clique in G . We

⁴In fact, the presented path is also a best-response path.

define $\Delta(i, C)$, the distance between i and C , as the minimal graph distance between i and a vertex in C .

Given a potential Φ with M -bounded local potentials on G , let

$$\lambda = 1 - \frac{1}{2DM}.$$

Fix a vertex k in G , and define the potential Θ by

$$\Theta = \sum_{C \in \mathcal{C}(G)} \lambda^{\Delta(k, C)} \Phi^C.$$

By the definition of better-response dynamics, Φ increases along any better-response path. We claim that the same (weakly) holds for Θ . This follows from the next claim.

Lemma 5.8. *For every $a \in \mathcal{S}$ and $b = (b_i, a_{-i}) \in \mathcal{S}$ it holds that if $\Phi(b) > \Phi(a)$ then $\Theta(b) \geq \Theta(a)$.*

In other words, the lemma claims that Θ is an *ordinal potential* (see Monderer and Shapley [13]) of the game \mathcal{G} .

Proof. Fix a and b such that $\Phi(b) > \Phi(a)$. Since they differ only in the strategy of i then

$$(5.1) \quad \Phi(b) - \Phi(a) = \sum_{C \in \mathcal{C}(G) : i \in C} [\Phi^C(b) - \Phi^C(a)],$$

and likewise

$$(5.2) \quad \Theta(b) - \Theta(a) = \sum_{C \in \mathcal{C}(G) : i \in C} \lambda^{\Delta(k, C)} [\Phi^C(b) - \Phi^C(a)].$$

Now, note that for any clique C that includes i it holds that either $\Delta(k, C) = \Delta(k, i)$ or else $\Delta(k, C) = \Delta(k, i) - 1$. Hence, if we denote

$$\mathcal{C}_1 = \{C \in \mathcal{C}(G) : i \in C, \Delta(k, C) = \Delta(k, i)\}$$

and

$$\mathcal{C}_2 = \{C \in \mathcal{C}(G) : i \in C, \Delta(k, C) = \Delta(k, i) - 1\}$$

then we can write (5.2) as

$$\begin{aligned} \frac{\Theta(b) - \Theta(a)}{\lambda^{\Delta(k, C) - 1}} &= \lambda \sum_{C \in \mathcal{C}_1} [\Phi^C(b) - \Phi^C(a)] + \sum_{C \in \mathcal{C}_2} [\Phi^C(b) - \Phi^C(a)] \\ &= \sum_{C \in \mathcal{C}(G) : i \in C} [\Phi^C(b) - \Phi^C(a)] - (1 - \lambda) \sum_{C \in \mathcal{C}_1} [\Phi^C(b) - \Phi^C(a)]. \end{aligned}$$

By (5.1) this implies that

$$\frac{\Theta(b) - \Theta(a)}{\lambda^{\Delta(k,C)-1}} \geq \Phi(b) - \Phi(a) - (1 - \lambda) \sum_{c_1} [\Phi^C(b) - \Phi^C(a)],$$

and by the definition of λ

$$\geq \Phi(b) - \Phi(a) - \frac{1}{2DM} \sum_{c_1} [\Phi^C(b) - \Phi^C(a)].$$

Now, $\Phi(b) > \Phi(a)$ implies $\Phi(b) - \Phi(a) \geq 1$, since the potential is integer. Hence

$$\geq 1 - \frac{1}{2DM} \sum_{c_1} [\Phi^C(b) - \Phi^C(a)].$$

But the summand has at most D terms, each one of which is strictly less than $2M$ (because the local potentials are bounded by M). Hence

$$\frac{\Theta(b) - \Theta(a)}{\lambda^{\Delta(k,C)-1}} \geq 0,$$

which completes the proof of the claim. \square

By the definition of Θ we have that

$$|\Theta(a)| \leq \sum_{C \in \mathcal{C}(G)} \lambda^{\Delta(k,C)} |\Phi^C(a)|.$$

Since $|\Phi^C(a)| \leq M$ then

$$\begin{aligned} |\Theta(a)| &\leq \sum_{C \in \mathcal{C}(G)} \lambda^{\Delta(k,C)} M \\ &= M \cdot \sum_{r=0}^{\infty} \lambda^r \cdot \#\{C \in \mathcal{C}(G) : \Delta(k,C) = r\}. \end{aligned}$$

Now, the number of cliques at distance r from k is at most D times the number of vertices at distance r from k . Hence

$$|\Theta(a)| \leq DM \sum_{r=0}^{\infty} \lambda^r S_r(G, k).$$

Whenever player k updates her strategy to a better one, Φ increases by at least 1, because the potential is integral. Note that a change of strategy by player k changes only the value of the local potentials Φ^C such that $k \in C$ (i.e., $\Delta(k,C) = 0$). Therefore, an update of player

k causes an increment by at least 1 (again because the potential is integral) of the potential

$$(5.3) \quad \sum_{C \in \mathcal{C}(G) : \Delta(k,C)=0} \Phi^C = \sum_{C \in \mathcal{C}(G) : \Delta(k,C)=0} \lambda^{\Delta(k,C)} \Phi^C$$

and causes no change in the potential

$$(5.4) \quad \sum_{C \in \mathcal{C}(G) : \Delta(k,C) \geq 1} \lambda^{\Delta(k,C)} \Phi^C.$$

Namely, every better response of player k causes an increment by at least 1 of Θ , which is the sum the potentials in (5.3) and (5.4).

By Lemma 5.8 Θ is non-decreasing along a better response path. Hence the total number of times that player k updates her strategy is at most

$$2DM \sum_{r=0}^{\infty} \lambda^r S_r(G, k).$$

Finally, since $\lambda = 1 - 1/(2DM)$, it follows that the total number of times that player k updates her strategy is at most

$$2DM \sum_{r=0}^{\infty} \left(1 - \frac{1}{2DM}\right)^r S_r(G, k).$$

This completes the proof of Theorem 5.3.

6. GENERALIZED GRAPHICAL POTENTIAL GAMES.

Let $G = (V, E)$ be a countably infinite graph, where the degree of each node is finite.

Definition 6.1. A graphical game \mathcal{G} on G is a *generalized* graphical potential game if it is strategically equivalent to a game whose utilities are given by

$$u_i(a) = \sum_{C \in \mathcal{C}(G) : i \in C} \Phi^C(a^C)$$

for some local potentials $\{\Phi^C\}_{C \in \mathcal{C}}$.

Note that \mathcal{G} is not necessarily a potential game, as the sum of the local potentials may diverge. However, given the decomposition of graphical potential games into local potentials (Theorem 4.1), we find that this is a natural generalization. Indeed, we show that like potential games, generalized potential games always have pure Nash equilibria. This follows from a compactness argument.

Theorem 6.2. *Every generalized potential game has a pure Nash equilibrium.*

Proof. Equip $\mathcal{S} = \prod_i \mathcal{S}_i$ with the product topology. Fix a player i , and let $B_r = \cup_{s \leq r} \mathcal{S}_s$ be the set of players at distance at most r from i .

For each $r > 0$, choose a strategy profile $a_r \in \mathcal{S}$ such that every player in $B_r(i)$ is already best responding. This is possible, since if we arbitrarily fix the strategies of the players outside $B_r(i)$, the induced game on the players in $B_r(i)$ is a potential game, and therefore has a pure Nash equilibrium. This follows from the definition of generalized potential games by local potentials.

Since \mathcal{S} is compact, the sequence a_r will have a converging subsequence, with some limit $a \in \mathcal{S}$. It is easy to see that a is a Nash equilibrium. \square

Note that unlike finite potential games, in a generalized potential game a better-response path does not necessarily reach an equilibrium. However, under the appropriate conditions on the game and growth rate of the graph, we now show that every better-response path converges to an equilibrium.

To this end, we define the clique-degree of an infinite graph as for finite graphs, although in this case it may diverge. Likewise, the definition of M -boundedness can be used, since it only depends on the local potentials.

The statement of the next theorem is identical to that of Theorem 5.3, with the exception that it refers more widely to generalized graphical potential games. An inspection of the proof of Theorem 5.3 will reveal that it applies here too.

Theorem 6.3. *Let \mathcal{G} be an integral generalized potential game on the graph G with M -bounded local potentials, where G has clique-degree D . Then the number of times that a player k updates her strategy in any better-response path is at most*

$$2DM \sum_{r=0}^{\infty} \left(1 - \frac{1}{2DM}\right)^r S_r(G, k).$$

Note that the bound in Theorem 6.3 may be infinite (for example, in graphs of sufficiently fast exponential growth) or finite (for example, in graph of subexponential growth), in which case it is easy to show that it is finite for all $k \in V$. When the bound is infinite, the statement of the theorem is vacuous. However, when the bound is finite, it follows that in any better response path the strategies of all players converge. We end the paper by formally stating this observation.

Theorem 6.4. *Let \mathcal{G} be an integral generalized potential game on the graph G with M -bounded local potentials, where G has clique-degree D , and where*

$$2DM \sum_{r=0}^{\infty} \left(1 - \frac{1}{2DM}\right)^r S_r(G, k)$$

is finite for some (equivalently, all) $k \in V$. In any better-response path (which may be infinite), each player makes only a finite number of strategy changes.

REFERENCES

- [1] Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, Paolo Penna, and Giuseppe Persiano, *Convergence to equilibrium of logit dynamics for strategic games*, Proceedings of the 23rd acm symposium on parallelism in algorithms and architectures, 2011, pp. 197–206.
- [2] Baruch Awerbuch, Yossi Azar, Amir Epstein, Vahab Seyed Mirrokni, and Alexander Skopalik, *Fast convergence to nearly optimal solutions in potential games*, Proceedings of the 9th acm conference on electronic commerce, 2008, pp. 264–273.
- [3] Vittorio Bilò, Angelo Fanelli, Michele Flammini, and Luca Moscardelli, *Graphical congestion games*, Algorithmica **61** (2011), no. 2, 274–297.
- [4] Kostas Bimpikis, Rahmi Ilkilic, and Ehsani Shayan, *Cournot competition in networked markets*, Proceedings of the 15th acm conference of economics and coputation, 2014.
- [5] Yann Bramoullé, Rachel Kranton, and Martin D’amours, *Strategic interaction and networks*, The American Economic Review **104** (2014), no. 3, 898–930.
- [6] Constantinos Daskalakis and Christos H Papadimitriou, *Computing pure Nash equilibria in graphical games via markov random fields*, Proceedings of the 7th acm conference on electronic commerce, 2006, pp. 91–99.
- [7] Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar, *The complexity of pure Nash equilibria*, Proceedings of the thirty-sixth annual acm symposium on theory of computing, 2004, pp. 604–612.
- [8] Geoffrey R Grimmett, *A theorem about random fields*, Bulletin of the London Mathematical Society **5** (1973), no. 1, 81–84.
- [9] John M Hammersley and Peter Clifford, *Markov fields on finite graphs and lattices*, 1968. Unpublished.
- [10] Matthew O. Jackson and Yves Zenou, *Games on networks*, Technical Report 9127, C.E.P.R. Discussion Papers, 2012.
- [11] Sham Kakade, Michael Kearns, John Langford, and Luis Ortiz, *Correlated equilibria in graphical games*, Proceedings of the 4th acm conference on electronic commerce, 2003, pp. 42–47.
- [12] Michael Kearns, Michael L Littman, and Satinder Singh, *Graphical models for game theory*, Proceedings of the seventeenth conference on uncertainty in artificial intelligence, 2001, pp. 253–260.
- [13] Dov Monderer and Lloyd S Shapley, *Potential games*, Games and economic behavior **14** (1996), no. 1, 124–143.

- [14] J.W. Moon and L. Moser, *On cliques in graphs*, Israel Journal of Mathematics **3** (1965), no. 1, 23–28 (English).
- [15] Noam Nisan, *Algorithmic game theory*, Cambridge University Press, 2007.
- [16] Luis E Ortiz, *Graphical potential games*, arXiv preprint arXiv:1505.01539 (2015).
- [17] CJ Preston, *Generalized Gibbs states and Markov random fields*, Advances in Applied probability (1973), 242–261.
- [18] Robert W Rosenthal, *A class of games possessing pure-strategy nash equilibria*, International Journal of Game Theory **2** (1973), no. 1, 65–67.
- [19] S Sherman, *Markov random fields and Gibbs random fields*, Israel Journal of Mathematics **14** (1973), no. 1, 92–103.
- [20] Alexander Skopalik and Berthold Vöcking, *Inapproximability of pure Nash equilibria*, Proceedings of the 40th annual acm symposium on theory of computing, 2008, pp. 355–364.
- [21] Omer Tamuz and Ran J Tessler, *Majority dynamics and the retention of information*, Israel Journal of Mathematics, forthcoming (2014).
- [22] David G Victor, Amy M Jaffe, and Mark H Hayes, *Natural gas and geopolitics: From 1970 to 2040*, Cambridge University Press, 2006.

APPENDIX A. MARKOV PROPERTIES OF POSITIVE RANDOM FIELDS.

Let $G = (V, E)$ be a finite simple graph with $|V| = n$, and let $X = (X_1, \dots, X_n)$ be random field on G . Recall that we say that X is positive if $\mathbb{P}[X_1 = a_1, \dots, X_n = a_n] > 0$ whenever $\mathbb{P}[X_i = a_i] > 0$ for $i = 1, \dots, n$. We say that X has the *pairwise Markov property* if for all $(i, j) \notin E$ it holds that X_i is independent of X_j , conditioned on $X_{V \setminus \{i, j\}}$. In this section we prove the following proposition.

Proposition A.1. *If X is positive and has the pairwise Markov property then it is a Markov random field.*

To prove this proposition we will need the following lemma. Let $G^{1,2}$ be the graph derived from G by amalgamating the nodes 1 and 2 into a node $(1, 2)$; the vertices of $G^{1,2}$ are $V^{1,2} = \{(1, 2), 3, \dots, n\}$, and the edges of $G^{1,2}$ are

$$E^{1,2} = \{(i, j) \in E : i, j > 2\} \cup \{((1, 2), j) : j > 2, (1, j) \in E \text{ or } (2, j) \in E\}.$$

Accordingly, let $X^{1,2} = ((X_1, X_2), X_3, \dots, X_n)$ be the associated random field, where we unite the two random variables X_1 and X_2 into the random variable (X_1, X_2) , and associate it to the node $(1, 2)$. Generally, define $G^{i,j}$ and $X^{i,j}$ likewise.

Lemma A.2. *When X is positive and has the pairwise Markov property with respect to G , then $X^{i,j}$ is positive and has the pairwise Markov property, with respect to $G^{i,j}$.*

Proof. Without loss of generality, we assume that $i = 1$ and $j = 2$. It is immediate that $X^{1,2}$ is positive.

Every vertex k that is connected to either 1 or 2 in the graph G , is connected to $(1, 2)$ in $G^{1,2}$. Therefore, for proving the pairwise property on the graph $G^{1,2}$ we will consider a vertex k such that neither $(1, k)$ nor $(2, k)$ is an edge in G , so that $((1, 2), k)$ is not an edge in $G^{1,2}$. Without loss of generality, we assume that $k = 3$. We will show that, conditioned on $Y = X_{V \setminus \{1,2,3\}}$, (X_1, X_2) is independent of X_3 . This will prove the claim.

Denote by \mathcal{S}_i the (finite) support of X_i . We would like to show that for all $a_1 \in \mathcal{S}_1, a_2 \in \mathcal{S}_2$ and $a_3 \in \mathcal{S}_3$ it holds that

$$\begin{aligned} & \mathbb{P}[(X_1, X_2, X_3) = (a_1, a_2, a_3)|Y] \\ &= \mathbb{P}[(X_1, X_2) = (a_1, a_2)|Y] \cdot \mathbb{P}[X_3 = a_3|Y]. \end{aligned}$$

Since X is positive $\mathbb{P}[X_3 = a_3|Y, (X_1, X_2) = (a_1, a_2)]$ is well defined, and we can apply the pairwise Markov property of X once to X_1, X_3 and once to X_2, X_3 to arrive at

$$\begin{aligned} & \mathbb{P}[X_3 = a_3|Y, (X_1, X_2) = (a_1, a_2)] \\ \text{(A.1)} \quad &= \mathbb{P}[X_3 = a_3|Y, X_1 = a_1] \\ \text{(A.2)} \quad &= \mathbb{P}[X_3 = a_3|Y, X_2 = a_2]. \end{aligned}$$

Since X is positive we can apply the same argument with any $b_2 \in \mathcal{S}_2$ substituted for a_2 , and conclude, as in (A.2), that

$$\mathbb{P}[X_3 = a_3|Y, X_1 = a_1] = \mathbb{P}[X_3 = a_3|Y, X_2 = b_2]$$

for all $b_2 \in \mathcal{S}_2$, and that therefore

$$\mathbb{P}[X_3 = a_3|Y, X_1 = a_1] = \mathbb{P}[X_3 = a_3|Y].$$

Applying (A.1) now yields

$$\mathbb{P}[X_3 = a_3|Y, (X_1, X_2) = (a_1, a_2)] = \mathbb{P}[X_3 = a_3|Y],$$

which proves the claim. \square

Proof of Proposition A.1. Let A be a (U, W) -cut. We would like to show that X_U is independent of X_W , conditioned on X_A .

Amalgamate all the vertices in U to a single vertex, and likewise amalgamate A and W . The resulting graph is (a subgraph of)

$$G' = (V', E') = (\{U, A, W\}, \{(U, A), (A, W)\}),$$

and the associated random field is precisely (X_U, X_A, X_W) . By repeated applications of Lemma A.2 this graph has the pairwise Markov property, and hence X_U is independent of X_W , conditioned on X_A . \square

