

Extremal hypergraphs for the biased Erdős–Selfridge theorem

Eric Lars Sundberg

Mathematics Department

Occidental College

Los Angeles, CA

sundberg@oxy.edu

Submitted: May 30, 2012; Accepted: Jan 27, 2013; Published: Feb 5, 2013

Mathematics Subject Classification: 91A46

Abstract

A positional game is essentially a generalization of Tic-Tac-Toe played on a hypergraph (V, \mathcal{F}) . A pivotal result in the study of positional games is the Erdős–Selfridge theorem, which gives a simple criterion for the existence of a Breaker’s winning strategy on a finite hypergraph \mathcal{F} . It has been shown that the bound in the Erdős–Selfridge theorem can be tight and that numerous extremal hypergraphs exist that demonstrate the tightness of the bound. We focus on a generalization of the Erdős–Selfridge theorem proven by Beck for biased $(p : q)$ games, which we call the $(p : q)$ –Erdős–Selfridge theorem. We show that for pn -uniform hypergraphs there is a *unique* extremal hypergraph for the $(p : q)$ –Erdős–Selfridge theorem when $q \geq 2$.

1 Introduction

A positional game is a generalization of Tic-Tac-Toe played on a hypergraph (V, \mathcal{F}) where the vertices can be considered the “board” on which the game is played, and the edges can be thought of as the “winning sets.” (In this paper we will only consider finite hypergraphs.) A *positional game* on (V, \mathcal{F}) is a two-player game where at every turn each player alternately occupies a previously unoccupied vertex from V . A *biased positional game* or a $(p : q)$ *positional game* on (V, \mathcal{F}) is a two player game where at every turn the first player occupies p previously unoccupied vertices from V and then the second player occupies q previously unoccupied vertices from V . The game is over when all vertices of \mathcal{F} have been occupied, which may require one of the players to occupy less than his/her number of allotted vertices on the last turn. In a *strong positional game*, the first player to occupy all vertices of some edge $A \in \mathcal{F}$ wins. If at the end of play no edge is completely

occupied by either player, that play is declared a draw. Normal 3×3 Tic-Tac-Toe is a strong positional game where the vertices of the hypergraph are the nine positions and the edges are the eight winning lines. In a *Maker–Breaker* positional game, the first player, *Maker*, wins if she¹ occupies all vertices of some edge $A \in \mathcal{F}$, otherwise the second player, *Breaker*, wins. Therefore, by definition there are no draw plays in Maker–Breaker games. We say that a player P has a *winning strategy* if no matter how the other player plays, player P wins by following that winning strategy. Please note that in this paper, we will use \mathcal{F} both to denote the whole hypergraph (V, \mathcal{F}) and to denote just its set of edges, where the appropriate interpretation should be understood from the context.

A pivotal result in the study of positional games is the Erdős–Selfridge theorem [4]. In their paper, Erdős and Selfridge introduced the idea of transforming a probabilistic argument into a completely deterministic potential-based strategy for positional games. Their theorem gives a simple criterion for the existence of an explicit Breaker’s winning strategy on a hypergraph \mathcal{F} . It states that if

$$\sum_{A \in \mathcal{F}} 2^{-|A|} < \frac{1}{2}, \quad (1)$$

then Breaker has an explicit winning strategy for the Maker–Breaker game played on \mathcal{F} . In the case where \mathcal{F} is n -uniform, condition (1) simplifies to $|\mathcal{F}| < 2^{n-1}$. Despite its simplicity, the Erdős–Selfridge theorem can be used to determine the correct order of magnitude for the breaking points of many games. Moreover, it laid the groundwork for using potential-based strategies in positional game theory. These potential-based strategies play a key role in determining the asymptotically exact breaking points for many games where such a result is known. See Beck [2].

In addition to the remarkable results that stem from their theorem, Erdős and Selfridge provided an example of an n -uniform hypergraph with exactly 2^{n-1} edges on which Maker has a winning strategy, thus proving that the bound in their theorem is tight. Let us call a hypergraph \mathcal{F} an *extremal* hypergraph for the Erdős–Selfridge theorem if $\sum_{A \in \mathcal{F}} 2^{-|A|} = \frac{1}{2}$ and Maker has a winning strategy on \mathcal{F} . For now, we note that there are numerous extremal hypergraphs for the Erdős–Selfridge theorem. In this paper we focus on a generalization of the Erdős–Selfridge theorem proven by Beck [1] for $(p : q)$ games, which we call the $(p : q)$ -Erdős–Selfridge theorem (or sometimes the *biased Erdős–Selfridge theorem*). The $(p : q)$ -Erdős–Selfridge theorem states that if

$$\sum_{A \in \mathcal{F}} (q + 1)^{-\frac{|A|}{p}} < \frac{1}{q + 1}, \quad (2)$$

then Breaker has an explicit winning strategy for the $(p : q)$ -Maker–Breaker game played on \mathcal{F} . In the case where \mathcal{F} is pn -uniform, condition (2) simplifies to $|\mathcal{F}| < (q + 1)^{n-1}$. Along with this theorem, Beck also gave an example of a pn -uniform hypergraph \mathcal{F} with $|\mathcal{F}| = (q + 1)^{n-1}$ on which Maker has a winning strategy, i.e., an extremal hypergraph for

¹In this paper, we will refer to Maker with feminine pronouns such as “she” and “her,” and we will refer to Breaker with masculine pronouns.

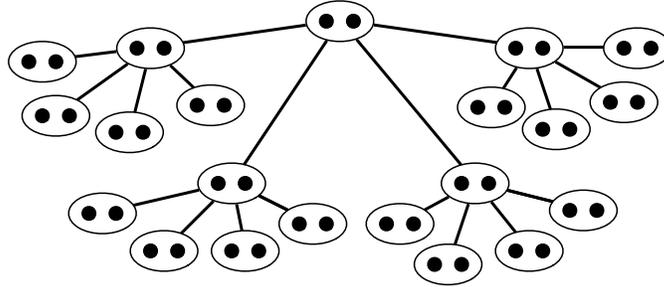


Figure 1: An example of a *complete* 3-level, 4-ary tree where each node has two vertices. This is an extremal hypergraph for the $(2 : 3)$ -Erdős–Selfridge theorem which is 6-uniform and has 4^2 edges.

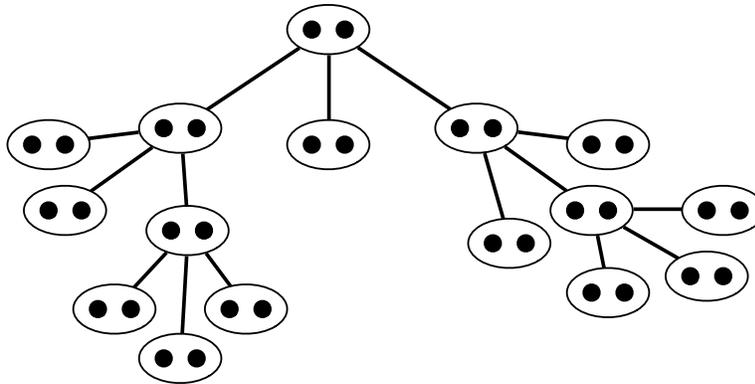


Figure 2: A 3-ary tree where every node has 2 vertices. This is an extremal hypergraph for the $(2 : 2)$ -Erdős–Selfridge theorem which is nonuniform and has 11 edges.

the $(p : q)$ -Erdős–Selfridge theorem, thus showing that the bound in the theorem is tight. In this paper we will prove that if $q \geq 2$, then the extremal hypergraph given by Beck is unique (for the pn -uniform case).

To explain the extremal hypergraph for the biased case given by Beck in [1], we first consider the following generalization of a complete binary tree. A *rooted $(q + 1)$ -ary tree* is a generalization of a rooted binary tree, where each (non-leaf) node has $q + 1$ children as a binary tree has two. The hypergraph we wish to consider can be derived from a rooted $(q + 1)$ -ary tree so that each node of the tree is identified with p distinct vertices of the hypergraph. Thus, the hypergraph has p times as many vertices as the tree has nodes. Whereas a tree-edge in the underlying tree connects two nodes of the tree, an edge in the hypergraph consists of all of the vertices from a path beginning at the root node and ending at a leaf node. See Figure 1 for a drawing of a 6-uniform hypergraph based on a complete 3-level, 4-ary tree where each node contains 2 vertices; and see Figure 2 for a drawing of a nonuniform hypergraph based on a 3-ary tree where each node contains 2 vertices.

The winning strategy that Maker has on a rooted $(q + 1)$ -ary tree can be described as follows. First Maker occupies all p vertices from the root node. Then there are (essentially) $q + 1$ disjoint $(q + 1)$ -ary trees left over. (Each tree is rooted at level 2 in the original hypergraph.) Breaker can choose his vertices from at most q of the subtrees. Thus, there is always a subtree in which Breaker has occupied no vertices. Maker occupies her next p vertices in the root of one such unoccupied subtree and continues in that manner until she reaches a leaf node. It is worth noting that Breaker can always force upon Maker the edge by which she wins if at each turn he chooses a vertex from each of the root nodes of those subtrees which do not contain the leaf node which determines the desired edge.

For a hypergraph \mathcal{F} , let us define the *potential* of an edge $A \in \mathcal{F}$ as $(q + 1)^{-\frac{|A|}{p}}$, and the potential of the whole hypergraph \mathcal{F} as $\sum_{A \in \mathcal{F}} (q + 1)^{-\frac{|A|}{p}}$. Through an inductive argument we can see that the potential of a rooted $(q + 1)$ -ary tree is $\frac{1}{q + 1}$. Indeed, a single edge with p vertices is a $(q + 1)$ -ary tree with exactly one level, and its potential is $\frac{1}{q + 1}$. A $(q + 1)$ -ary tree T with more than one level is constructed by connecting $q + 1$ disjoint $(q + 1)$ -ary trees T_1, \dots, T_{q+1} with a root node R which consists of p vertices. By induction, each subtree T_i has potential $\frac{1}{q + 1}$ as a standalone tree. However, once the root R is appended to each of its edges, as a subtree of T , each T_i has potential $\frac{1}{(q + 1)^2}$ since each edge in T_i grows by p vertices. Since T is composed of the $q + 1$ edge-disjoint, augmented subtrees, the potential of T is $\frac{1}{q + 1}$.

The main result of this paper is that when $q \geq 2$, the only extremal hypergraphs for the $(p : q)$ -Erdős-Selfridge theorem are $(q + 1)$ -ary trees. Lu [5] investigated the extremal hypergraphs for $(p : q)$ -Erdős-Selfridge theorem (no restriction on q) in the case when Maker has an *economical* winning strategy, i.e., if the hypergraph \mathcal{F} is n -uniform, then Maker has a winning strategy that wins in n turns. However, his paper contains some errors. Thus, some of the hypergraphs described in that paper are either not extremal hypergraphs or not economical extremal hypergraphs. In [6], the author addressed the case of economical extremal hypergraphs for the $(p : q)$ -Erdős-Selfridge theorem when $q \geq 2$. However, those results are superseded by the current paper. In his paper, Lu described a family of hypergraphs which became a starting point for defining numerous extremal hypergraphs for the *unbiased*, i.e., $(1 : 1)$, Erdős-Selfridge theorem. (It is fairly easy to see that any extremal hypergraph for a $(1 : 1)$ game can be transformed into an extremal hypergraph for the $(p : 1)$ game by replacing each vertex with a node containing p distinct vertices.) In Chapter 6 of [3], Beck amended Lu's work to define a family of n -uniform extremal hypergraphs for the (unbiased) Erdős-Selfridge theorem. The construction proceeds as follows. Let T be a rooted binary tree with n levels and 2^{n-1} leaves. Using standard tree terminology, we say v is an *ancestor* of w if $v \neq w$ and v is on the path between w and the root; and two nodes are called *brothers* if they are adjacent to a common ancestor (called their *father*). Let us say a labeling $L : V(T) \rightarrow \mathbb{N}$ is *good* if it satisfies the following three properties:

1. if v and w are brothers, then $L(v) \neq L(w)$;
2. if v is an ancestor of w , then $L(v) \neq L(w)$;

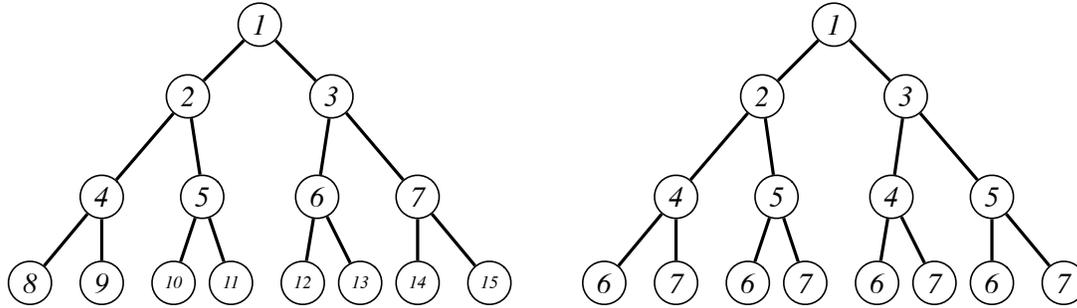


Figure 3: Two good labelings which lead to the binary tree extremal hypergraph given by Beck in [1] and the original extremal hypergraph given by Erdős and Selfridge in [4].

3. suppose v and w are brothers and v' and w' are brothers, then $L(v) = L(v')$ if and only if $L(w) = L(w')$.

Given a good labeling of T , we define a hypergraph \mathcal{H} whose vertices are the labels on the nodes of T and whose edges are all sets of the form $\{L(v_1), \dots, L(v_n)\}$, where v_1, \dots, v_n is a path from the root node of T to a leaf node of T . Properties (1) and (2) of the good labeling guarantee that \mathcal{H} is n -uniform and has 2^{n-1} edges. (Lu only used those two properties for the family of hypergraphs he described.) Property (3) of the good labeling (missing in Lu’s paper) provides Maker with the following winning strategy. Maker occupies the label of the root node for her first move. Then for each subsequent turn, Maker occupies the “brother” of the label that Breaker occupied in the previous turn. Each label will have a well-defined “brother” by Property (3). Figure 3 shows two good labelings. The labeling on the left is a trivial labeling where each node receives a unique label. In this case, the derived hypergraph is the original binary tree example given by Beck in [1]. The labeling on the right leads to the original extremal hypergraph given by Erdős and Selfridge in [4].

In [3], Beck explains how some hypergraphs derived from good labelings are economical extremal hypergraphs for the Erdős–Selfridge theorem while others are not. Currently, there are no criteria for determining if a given good labeling will yield an economical or non-economical extremal hypergraph. Additionally, Lu [5] provided an example of a non-economical extremal hypergraph for the Erdős–Selfridge theorem which is not derived from a labeling of a binary tree. A very similar example, provided by A.J. Sanders in a manuscript from 2004, is mentioned by Beck in Chapter 6 of [3]. Figure 4 shows the extremal hypergraph given by Sanders on the left and the one given by Lu on the right. Each hypergraph is 4-uniform. In both cases, the vertices of the hypergraphs are the labels on the nodes. The edges are the paths indicated by the arrows, i.e., the six downward black paths, the red (mostly) horizontal path $\{1, 2, 3, 4\}$, and the blue (mostly) horizontal path $\{1, 5, 6, 7\}$. Based on these examples and the main result of this paper, it is clear that there is a stark contrast between the case $q = 1$ and the case $q \geq 2$: when $q \geq 2$, the structure of the extremal hypergraphs for the $(p : q)$ –Erdős–Selfridge theorem is completely determined, whereas when $q = 1$ classifying all of the extremal hypergraphs

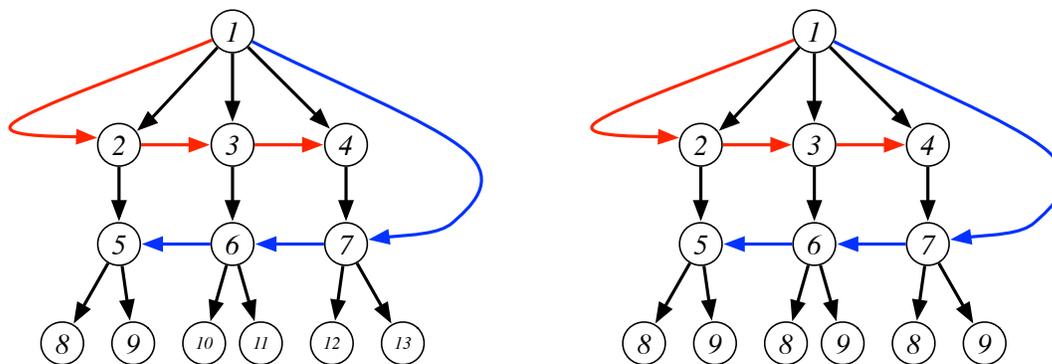


Figure 4: Two 4-uniform non-econominal extremal hypergraphs for the Erdős–Selfridge theorem which are not derived from a labeling of a binary tree.

for the $(p : q)$ –Erdős–Selfridge theorem is a wide open problem.

The remainder of the paper is organized as follows. In Section 2 we give a proof of the $(p : q)$ –Erdős–Selfridge theorem as we make substantial use of its proof in Section 3. In Section 3 we give the proof of the main theorem of the paper, namely, that in the case when $q \geq 2$, the only extremal hypergraphs for the $(p : q)$ –Erdős–Selfridge theorem are rooted $(q + 1)$ -ary trees. Section 3 is divided into three subsections. In Section 3.1 we lay the foundation of the proof by giving some fundamental results that are used throughout Sections 3.2 and 3.3. In Section 3.2, we give the proof of our main theorem, which requires Lemma 9 from Section 3.3. Section 3.3 contains the proof of Lemma 9 along with the results which build up to its proof.

2 $(p : q)$ –Erdős–Selfridge Theorem

In this section we state and prove a generalization of the Erdős–Selfridge theorem for the $(p : q)$ –Maker–Breaker game. This theorem was originally proven by Beck [1], and it reduces to the Erdős–Selfridge theorem in the case $p = q = 1$. Although Beck did not refer to his theorem by this name, we will refer to it as the $(p : q)$ –Erdős–Selfridge theorem.

Let us use the notation $\Phi(\mathcal{F}) := \sum_{A \in \mathcal{F}} \lambda^{-|A|}$, where $\lambda = (q + 1)^{1/p}$, to denote the *potential* of a finite hypergraph \mathcal{F} . Then the $(p : q)$ –Erdős–Selfridge theorem can be stated as follows.

Theorem 1 (Beck [1]) *In the $(p : q)$ –Maker–Breaker game on \mathcal{F} , if $\Phi(\mathcal{F}) < \lambda^{-p}$, where $\lambda = (q + 1)^{1/p}$, then Breaker has an explicit winning strategy.*

Proof: Let X_i be the set of vertices occupied by Maker during her i^{th} turn, and let Y_i be the set of vertices occupied by Breaker during his i^{th} turn. Let $M_i = \cup_{j=1}^i X_j$ and let $B_i = \cup_{j=1}^i Y_j$. Let $\mathcal{F}_i = \{A \setminus M_i : A \in \mathcal{F}, A \cap B_{i-1} = \emptyset\}$. Notice that \mathcal{F}_i is the hypergraph that Breaker plays on during turn i , because neither player is concerned with edges already killed by Breaker, and both players need consider only the unoccupied portions of those

edges which are still alive and partially occupied by Maker. Although Maker moves first within each turn, we will focus on the fact that Breaker can destroy more potential in turn i than Maker can create in turn $i + 1$.

Claim 1 *For each turn i , it is possible for Breaker to choose Y_i so that for any choice of X_{i+1} ,*

$$\Phi(\mathcal{F}_{i+1}) \leq \Phi(\mathcal{F}_i). \tag{3}$$

Proof of Claim 1: Consider the time in the game during turn i after Maker has occupied her p vertices, but before Breaker has begun to occupy his q vertices. Let $Y_i = \{y_1, \dots, y_q\}$ and $X_{i+1} = \{x_1, \dots, x_p\}$. We will assume that Breaker picks y_1 , then y_2 and so forth, rather than all at once. Similarly, we will assume that Maker picks the vertex x_1 , then x_2 , and so forth. We define

$$\mathcal{F}_i^{(b,0)} = \{A \in \mathcal{F}_i : A \cap \{y_1, \dots, y_b\} = \emptyset\}, \quad 1 \leq b \leq q, \tag{4}$$

to be the hypergraphs on which Breaker plays during turn i and

$$\mathcal{F}_i^{(q,m)} = \{A \setminus \{x_1, \dots, x_m\} : A \in \mathcal{F}_i^{(q,0)}\}, \quad 1 \leq m \leq p, \tag{5}$$

to be the hypergraphs on which Maker plays during turn $i + 1$. For convenience, we let $\mathcal{F}_i^{(0,0)} = \mathcal{F}_i$. The hypergraph $\mathcal{F}_i^{(b,0)}$ consists of edges from \mathcal{F}_i which have not been blocked by the first b vertices of Breaker's i^{th} move. The hypergraph $\mathcal{F}_i^{(q,m)}$ consists of "partial" edges of \mathcal{F}_i not blocked by Breaker during his i^{th} turn. By "partial" edge, we mean that each edge A has been replaced by $A \setminus \{x_1, \dots, x_m\}$.

For an arbitrary hypergraph \mathcal{H} , vertex v and set S , let $\mathcal{H}(v) = \{A \in \mathcal{H} : v \in A\}$ and $\mathcal{H}(S) = \{A \in \mathcal{H} : S \subseteq A\}$. We derive the equation

$$\Phi(\mathcal{F}_{i+1}) = \Phi(\mathcal{F}_i) - \sum_{k=1}^q \Phi(\mathcal{F}_i^{(k-1,0)}(y_k)) + \sum_{j=1}^p (\lambda - 1) \Phi(\mathcal{F}_i^{(q,j-1)}(x_j)) \tag{6}$$

as follows. We begin with $\Phi(\mathcal{F}_i)$, then subtract the potential of each edge blocked by Breaker using his vertex y_1 . We then look at the remaining hypergraph $\mathcal{F}_i^{(1,0)}$ and remove the potential of each edge containing y_2 . We continue this process until Breaker occupies q vertices. This accounts for the terms in the first summation. When Maker occupies her first vertex x_1 , she multiplies the potential of each edge in $\mathcal{F}_i^{(q,0)}$ which contains x_1 by λ . So, for each edge $A \in \mathcal{F}_i^{(q,0)}(x_1)$, we remove its old potential $\lambda^{-|A|}$ and replace it with its new potential $\lambda^{-|A|+1}$, which we accomplish by adding $(\lambda - 1)\lambda^{-|A|}$. When Maker occupies x_2 , she multiplies the potential of each edge in $\mathcal{F}_i^{(q,1)}$ containing x_2 by λ . Thus, for each edge $A \in \mathcal{F}_i^{(q,1)}(x_2)$, we add $(\lambda - 1)\lambda^{-|A|}$ to its old potential. We continue this process until Maker occupies p vertices. This accounts for the terms in the second summation and finishes our derivation of equation (6).

Note that $\mathcal{F}_i^{(k-1,0)} \setminus \mathcal{F}_i^{(k-1,0)}(y_k) = \mathcal{F}_i^{(k,0)}$, which implies that $\mathcal{F}_i^{(k-1,0)} \supseteq \mathcal{F}_i^{(k,0)}$. For each $1 \leq k \leq q$, Breaker selects y_k so that $\Phi(\mathcal{F}_i^{(k-1,0)}(y_k))$ is maximum, i.e., $\Phi(\mathcal{F}_i^{(k-1,0)}(y_k)) = \max_w \Phi(\mathcal{F}_i^{(k-1,0)}(w))$. (This is what we will refer to as the strategy from the $(p : q)$ -Erdős-Selfridge theorem.) Therefore,

$$\Phi(\mathcal{F}_i^{(k-1,0)}(y_k)) \geq \max_w \Phi(\mathcal{F}_i^{(q,0)}(w)) \quad (7)$$

since $\mathcal{F}_i^{(k-1,0)} \supseteq \mathcal{F}_i^{(q,0)}$ for $1 \leq k \leq q$. Let z maximize $\Phi(\mathcal{F}_i^{(q,0)}(w))$. Notice for all $1 \leq j \leq p$,

$$\Phi(\mathcal{F}_i^{(q,j-1)}(x_j)) = \sum_{x_j \in A \in \mathcal{F}_i^{(q,j-1)}} \lambda^{-|A|} = \sum_{x_j \in A \in \mathcal{F}_i^{(q,0)}} \lambda^{-|A \setminus \{x_1, \dots, x_{j-1}\}|} \quad (8)$$

$$\leq \sum_{x_j \in A \in \mathcal{F}_i^{(q,0)}} \lambda^{-|A|+j-1} = \lambda^{j-1} \sum_{x_j \in A \in \mathcal{F}_i^{(q,0)}} \lambda^{-|A|} \quad (9)$$

$$= \lambda^{j-1} \Phi(\mathcal{F}_i^{(q,0)}(x_j)) \leq \lambda^{j-1} \Phi(\mathcal{F}_i^{(q,0)}(z)), \quad (10)$$

i.e.,

$$\Phi(\mathcal{F}_i^{(q,j-1)}(x_j)) \leq \lambda^{j-1} \Phi(\mathcal{F}_i^{(q,0)}(z)). \quad (11)$$

Starting with equation (6) then using equations (7) and (11), we get

$$\begin{aligned} \Phi(\mathcal{F}_{i+1}) &= \Phi(\mathcal{F}_i) - \sum_{k=1}^q \Phi(\mathcal{F}_i^{(k-1,0)}(y_k)) + \sum_{j=1}^p (\lambda - 1) \Phi(\mathcal{F}_i^{(q,j-1)}(x_j)) \\ &\leq \Phi(\mathcal{F}_i) - q \Phi(\mathcal{F}_i^{(q,0)}(z)) + (\lambda - 1) \Phi(\mathcal{F}_i^{(q,0)}(z)) \sum_{j=1}^p \lambda^{j-1} \\ &= \Phi(\mathcal{F}_i) + \Phi(\mathcal{F}_i^{(q,0)}(z)) [-q + (\lambda - 1)(1 + \lambda + \lambda^2 + \dots + \lambda^{p-1})] \\ &= \Phi(\mathcal{F}_i) + \Phi(\mathcal{F}_i^{(q,0)}(z)) (-q + \lambda^p - 1). \end{aligned} \quad (12)$$

Since $\lambda = (q + 1)^{1/p}$, then $\Phi(\mathcal{F}_{i+1}) \leq \Phi(\mathcal{F}_i)$ and the proof of the claim is finished. \square

Note that if an edge of \mathcal{F}_i is completely occupied by Maker, then its potential becomes 1 during some turn for Maker, and remains 1 for the duration of the game. Therefore, if $\Phi(\mathcal{F}_i) < 1$ for all i , then Maker never occupies a complete edge and Breaker wins the game. By Claim 1, we have $\Phi(\mathcal{F}_i) \leq \Phi(\mathcal{F}_1)$ for all i . Recall that \mathcal{F}_1 is the hypergraph that the players are working on after Maker has played her first turn, i.e., occupied her first p vertices. If Maker is able to occupy p vertices which are contained in every edge, then she is able to create the most potential possible. In this case, each edge has its potential multiplied by λ^p . By our hypothesis $\Phi(\mathcal{F}) < \frac{1}{\lambda^p}$, therefore $\Phi(\mathcal{F}_1) \leq \lambda^p \Phi(\mathcal{F}) < 1$. This completes the proof our theorem since we now have, for all i ,

$$\Phi(\mathcal{F}_i) \leq \Phi(\mathcal{F}_1) < 1. \quad \blacksquare$$

3 Main Theorem

Theorem 2 *When $q \geq 2$, the only extremal hypergraphs for the $(p : q)$ -Erdős-Selfridge theorem are $(q + 1)$ -ary trees where each node contains p vertices. If the hypergraph is pn -uniform, then there is a unique extremal hypergraph, the complete $(q + 1)$ -ary tree with n levels, where each node contains p vertices.*

Let us recall some notation. As before, let X_i be the set of vertices that Maker occupies during turn i , and let Y_i be the set of vertices that Breaker occupies during turn i . Let $M_i = \cup_{j=1}^i X_j$ and let $B_i = \cup_{j=1}^i Y_j$. Let $\mathcal{F}_i = \{A \setminus M_i : A \in \mathcal{F}, A \cap B_{i-1} = \emptyset\}$, and let us introduce $\mathcal{F}_{i+\frac{1}{2}} = \{A \setminus M_i : A \in \mathcal{F}, A \cap B_i = \emptyset\}$. Notice that \mathcal{F}_i is the hypergraph that Maker leaves after her i^{th} move, and $\mathcal{F}_{i+\frac{1}{2}}$ is the hypergraph that Breaker leaves after his i^{th} move. Thus, \mathcal{F}_i is the hypergraph on which Breaker plays during turn i , and $\mathcal{F}_{i+\frac{1}{2}}$ is the hypergraph on which Maker plays during turn $i + 1$.

Let us also give some definitions. A Maker's strategy on a hypergraph \mathcal{F} is called a $(p : q)$ -Erdős-Selfridge-defeating strategy if Maker wins the $(p : q)$ -Maker-Breaker game on \mathcal{F} whenever Breaker follows the $(p : q)$ -Erdős-Selfridge strategy and Maker uses said strategy. Notice that if Maker has a winning strategy for the $(p : q)$ -Maker-Breaker game on \mathcal{F} , then that winning strategy is also a $(p : q)$ -Erdős-Selfridge-defeating strategy on \mathcal{F} because a winning strategy defeats *all* Breaker strategies. If \mathcal{F} is a hypergraph such that $\Phi(\mathcal{F}) = \frac{1}{q+1}$ and Maker has a $(p : q)$ -Erdős-Selfridge-defeating strategy on \mathcal{F} , then we say \mathcal{F} is a *conditionally extremal hypergraph* for the $(p : q)$ -Erdős-Selfridge theorem. We use the qualifier "conditionally" because Maker may lose to Breaker if Breaker does not restrict himself to using only the $(p : q)$ -Erdős-Selfridge strategy. Notice that if \mathcal{F} is an extremal hypergraph for the $(p : q)$ -Erdős-Selfridge theorem, then \mathcal{F} is also a conditionally extremal hypergraph for the $(p : q)$ -Erdős-Selfridge theorem because $\Phi(\mathcal{F}) = \frac{1}{q+1}$ and Maker not only has a $(p : q)$ -Erdős-Selfridge-defeating strategy on \mathcal{F} , but, in fact, a winning strategy on \mathcal{F} .

Our proof hinges on being able to take a conditionally extremal hypergraph \mathcal{F} for the $(p : q)$ -Erdős-Selfridge theorem and determine information about its structure by following plays where Breaker uses the strategy from the $(p : q)$ -Erdős-Selfridge theorem and Maker uses a fixed $(p : q)$ -Erdős-Selfridge-defeating strategy. While those are the underlying assumptions in Sections 3.1 and 3.3, in Section 3.2 we see what extra structure is implied when we further specify that \mathcal{F} is an extremal hypergraph for the $(p : q)$ -Erdős-Selfridge theorem and Maker uses a winning strategy. Here is a brief outline. Claim 2 from Section 3.1 states that there are exactly p vertices that are contained in every edge of \mathcal{F} . Thus, Maker's first move X_1 is precisely those p vertices. (We refer to the p vertices in X_1 as the root node of \mathcal{F} , since we eventually show, in the case where \mathcal{F} is extremal, that X_1 is the root node of a $(q + 1)$ -ary tree.) Claim 3 from Section 3.1 exploits the fact that all inequalities in the proof of Claim 1 must be equalities since \mathcal{F} is conditionally extremal, and concludes that for each i , the edges of \mathcal{F}_i can be partitioned into $\mathcal{F}_i(y_1), \mathcal{F}_i(y_2), \dots, \mathcal{F}_i(y_q), \mathcal{F}_i(X_{i+1}), S_i$, where S_i is the set of those edges of \mathcal{F}_i which do not contain any of Breaker's vertices from turn i or any of Maker's vertices from

turn $i + 1$. It also determines that in \mathcal{F}_i the potential of $\mathcal{F}_i(y_k)$ equals the potential of $\mathcal{F}_i(x_j)$ for any $1 \leq k \leq q$ and $1 \leq j \leq p$. Thus, Breaker could choose to swap out any single vertex from his move during turn i and replace it with one vertex from Maker's move during her $(i + 1)^{th}$ turn and still have a valid $(p : q)$ -Erdős-Selfridge move. Lemma 4 from Section 3.1 states that if Maker wins during turn $i + 1$, then $S_i = \emptyset$. Thus, there must be some turn i where $S_i = \emptyset$. In Section 3.2 we use induction (basically on the size of the largest edge in \mathcal{F}) to prove our main theorem. In Claim 4 from Section 3.2 we show that if $S_i = \emptyset$, then $\mathcal{F}_i(y_1), \mathcal{F}_i(y_2), \dots, \mathcal{F}_i(y_q), \mathcal{F}_i(X_{i+1})$ constitute $q + 1$ disjoint $(q + 1)$ -ary trees. We then show that it must be the case that $S_1 = \emptyset$, which allows us to conclude that \mathcal{F} is the root node X_1 connected to the $q + 1$ disjoint $(q + 1)$ -ary trees $\mathcal{F}_1(y_1), \mathcal{F}_1(y_2), \dots, \mathcal{F}_1(y_q), \mathcal{F}_1(X_2)$, i.e., \mathcal{F} is a $(q + 1)$ -ary tree. To show this, we assume that the first time $S_i = \emptyset$ is not during turn 1. We consider the first time $S_i = \emptyset$, and determine that \mathcal{F}_{i-1} must have a structure very close to $q + 1$ disjoint $(q + 1)$ -ary trees, since Claim 4 implies that \mathcal{F}_i is indeed $q + 1$ disjoint $(q + 1)$ -ary trees. In fact, in Section 3.3 we show that \mathcal{F}_{i-1} is $q + 1$ "pseudo-trees" linked by a set of common edges. In Section 3.3 we define pseudo-trees and investigate the structure of \mathcal{F}_{i-1} . This investigation culminates in Lemma 9 which shows how the structure of \mathcal{F}_{i-1} allows Breaker to deviate from the $(p : q)$ -Erdős-Selfridge strategy during turn $i - 1$ to eventually win the game. Thus, we are able to conclude that when \mathcal{F} is an extremal hypergraph for the $(p : q)$ -Erdős-Selfridge theorem and Maker is using a winning strategy, then $S_1 = \emptyset$ and \mathcal{F} is a $(q + 1)$ -ary tree.

3.1 Preliminaries

We begin by stating some lemmas, claims and observations that will be useful in the proof of our theorem. The first lemma is essentially an observation that after each Maker move, the potential of the partial edges in the remaining hypergraph must be at least 1, if Maker is following a $(p : q)$ -Erdős-Selfridge-defeating strategy.

Lemma 1 *If Maker follows a $(p : q)$ -Erdős-Selfridge-defeating strategy and Breaker follows the $(p : q)$ -Erdős-Selfridge strategy on a hypergraph \mathcal{F} , then*

$$\Phi(\mathcal{F}_i) \geq 1, \quad \text{for } i \geq 1.$$

Otherwise, since the proof of the $(p : q)$ -Erdős-Selfridge theorem shows that Breaker can force a non-increasing property on the potential, he can force the potential at the end of the game to be less than 1 which implies Maker does not win.

The next claim states that if \mathcal{F} is a conditionally extremal hypergraph for the $(p : q)$ -Erdős-Selfridge theorem, then \mathcal{F} must contain a "root node," i.e., \mathcal{F} has exactly p vertices that are contained in every edge of \mathcal{F} . (We use the terminology "root node" because when \mathcal{F} is an extremal hypergraph for the $(p : q)$ -Erdős-Selfridge theorem, we eventually show that \mathcal{F} is a $(q + 1)$ -ary tree, and those p vertices constitute its root node.)

Claim 2 *If \mathcal{F} is a conditionally extremal hypergraph for the $(p : q)$ -Erdős-Selfridge theorem, then \mathcal{F} contains a root node, that is, $|\bigcap_{A \in \mathcal{F}} A| = p$.*

Proof of Claim 2: Clearly $|\bigcap_{A \in \mathcal{F}} A| \leq p$, because if $|\bigcap_{A \in \mathcal{F}} A| > p$, then after Maker's first move, there is still at least one vertex that is contained in every edge of \mathcal{F} . Breaker would occupy one of those vertices (which would be a vertex of maximum potential) during his first move and kill all edges in \mathcal{F} .

Let X_1 be the set of p vertices that Maker occupies during turn 1 according to her $(p : q)$ -Erdős-Selfridge-defeating strategy. Let $\mathcal{F}_1 = \{A \setminus X_1 : A \in \mathcal{F}\}$ be the hypergraph played on after Maker's first turn. In calculating the potential of \mathcal{F}_1 , we see that

$$\Phi(\mathcal{F}_1) = \sum_{A \in \mathcal{F}} (q+1)^{-|A \setminus X_1|/p} \leq \sum_{A \in \mathcal{F}} (q+1)^{-(|A|-p)/p} = (q+1)\Phi(\mathcal{F}) = 1, \quad (13)$$

since at best each edge A contains p of Maker's vertices and since $\Phi(\mathcal{F}) = \frac{1}{q+1}$. By Lemma 1, $\Phi(\mathcal{F}_1) \geq 1$. Thus, we must have equality in equation (13). This implies that each edge contains exactly p of Maker's vertices, i.e., there is a root node in \mathcal{F} . \square

Our next two observations are straightforward facts about the potentials of finite hypergraphs.

Observation 1 *If \mathcal{G} and \mathcal{H} are hypergraphs so that $\mathcal{G} \subseteq \mathcal{H}$, $V(\mathcal{G}) = V(\mathcal{H})$, and*

$$\Phi(\mathcal{G}) = \Phi(\mathcal{H}),$$

then $\mathcal{G} = \mathcal{H}$. Likewise, if \mathcal{G} and \mathcal{H} are sets of edges so that $\mathcal{G} \subseteq \mathcal{H}$ and

$$\sum_{A \in \mathcal{G}} \lambda^{-|A|} = \sum_{A \in \mathcal{H}} \lambda^{-|A|},$$

then $\mathcal{G} = \mathcal{H}$.

Observation 2 *If \mathcal{G} and \mathcal{H} are hypergraphs so that the set of edges of \mathcal{G} is a proper subset of the set of edges of \mathcal{H} , i.e., $\mathcal{G} \subset \mathcal{H}$, then*

$$\Phi(\mathcal{G}) < \Phi(\mathcal{H}).$$

Now we declare a set of basic hypotheses that we will often assume.

Basic Hypotheses (for turns 1 through i):

\mathcal{F} is a conditionally extremal hypergraph for the $(p : q)$ -Erdős-Selfridge theorem, Breaker uses the strategy from the $(p : q)$ -Erdős-Selfridge theorem for turns 1 through i , and Maker uses a fixed $(p : q)$ -Erdős-Selfridge-defeating strategy for turns 1 through $i + 1$.

Additionally, let us say that a Breaker's move is *valid* if it follows (i.e., is consistent with) the strategy from the proof of the $(p : q)$ -Erdős-Selfridge theorem and that a Maker's move is *valid* if it follows a fixed $(p : q)$ -Erdős-Selfridge-defeating strategy. We will use the fact that at a given turn i there will often be more than one valid move available for Breaker; however, since Maker is following a fixed $(p : q)$ -Erdős-Selfridge-defeating strategy, there is always a unique valid response for Maker, given a move by Breaker and the history of the game up to that point.

The following lemma is used frequently throughout our proof.

Lemma 2 *If the basic hypotheses hold for turns 1 through i , then*

$$\Phi(\mathcal{F}_j) = 1, \quad \text{for } 1 \leq j \leq i + 1.$$

Lemma 2 holds because we showed, in the proof of Claim 2, that $\Phi(\mathcal{F}_1) = 1$. Thus, using Claim 1 from Section 2 and Lemma 1, we are able to conclude that $\Phi(\mathcal{F}_j) = 1$ after each valid Maker move, i.e., for $1 \leq j \leq i + 1$.

The following claim uses the notation from equations (4) and (5) from the proof of Claim 1, and is critical to our proof.

Claim 3 *Suppose the basic hypotheses hold for turns 1 through i and $\{y_1, y_2, \dots, y_q\}$ is Breaker's move in turn i and $\{x_1, \dots, x_p\}$ is Maker's move during turn $i + 1$. Then*

$$\mathcal{F}_i^{(k-1,0)}(y_k) = \mathcal{F}_i(y_k), \quad 1 \leq k \leq q, \quad (14)$$

$$\mathcal{F}_i^{(q,0)}(x_j) = \mathcal{F}_i(x_j), \quad 1 \leq j \leq p, \quad (15)$$

and

$$\Phi(\mathcal{F}_i(y_k)) = \Phi(\mathcal{F}_i(x_j)), \quad 1 \leq k \leq q, \quad 1 \leq j \leq p. \quad (16)$$

Moreover, x_1, \dots, x_p are in exactly the same edges of \mathcal{F}_i .

Proof of Claim 3: Recall that $\mathcal{F}_i = \{A \setminus M_i : A \in \mathcal{F}, A \cap B_{i-1} = \emptyset\}$ is the hypergraph composed of the set of partial edges left after Maker's i^{th} move. It is the hypergraph on which Breaker plays during turn i . Let $Y_i = \{y_1, y_2, \dots, y_q\}$ be Breaker's i^{th} move, and let $X_{i+1} = \{x_1, \dots, x_p\}$ be Maker's response to Y_i .

By Lemma 2, $\Phi(\mathcal{F}_i) = \Phi(\mathcal{F}_{i+1}) = 1$. This implies that the inequality from equation (3) in Claim 1 must, in fact, be an equality. Thus, every inequality in the proof of Claim 1 must also be an equality. Let us examine those inequalities.

Since inequality (12) must be tight, we may conclude that

$$\Phi(\mathcal{F}_i^{(k-1,0)}(y_k)) = \Phi(\mathcal{F}_i^{(q,0)}(z)), \quad \text{for } 1 \leq k \leq q, \quad (17)$$

and

$$\Phi(\mathcal{F}_i^{(q,j-1)}(x_j)) = \lambda^{j-1} \Phi(\mathcal{F}_i^{(q,0)}(z)), \quad \text{for } 1 \leq j \leq p, \quad (18)$$

where $z \in V(\mathcal{F}_i^{(q,0)})$ is such that $\Phi(\mathcal{F}_i^{(q,0)}(z)) = \max_w \Phi(\mathcal{F}_i^{(q,0)}(w))$. Equality in line (10) then implies that

$$\Phi(\mathcal{F}_i^{(q,0)}(x_j)) = \Phi(\mathcal{F}_i^{(q,0)}(z)), \quad \text{for } 1 \leq j \leq p. \quad (19)$$

Since we have equality in moving from line (8) to line (9), we may also conclude that

$$\left\{ A \in \mathcal{F}_i^{(q,0)} : x_j \in A \right\} \subseteq \left\{ A \in \mathcal{F}_i^{(q,0)} : \{x_1, \dots, x_{j-1}\} \subseteq A \right\}, \quad (20)$$

i.e., for every $1 \leq j \leq p$, every edge that contains x_j also contains x_1, \dots, x_{j-1} . But this naturally implies that for each $1 \leq j \leq p$,

$$\left\{ A \in \mathcal{F}_i^{(q,0)} : x_j \in A \right\} \subseteq \left\{ A \in \mathcal{F}_i^{(q,0)} : x_k \in A \right\} \quad \text{where } 1 \leq k \leq j - 1. \quad (21)$$

Since equation (19) implies $\Phi(\mathcal{F}_i^{(q,0)}(x_j)) = \Phi(\mathcal{F}_i^{(q,0)}(x_k))$ for all $1 \leq k, j \leq p$, the containment in line (21) and Observation 1 imply that

$$\mathcal{F}_i^{(q,0)}(x_j) = \mathcal{F}_i^{(q,0)}(x_k) \text{ for } 1 \leq k \leq j - 1. \quad (22)$$

In particular,

$$\mathcal{F}_i^{(q,0)}(x_p) = \mathcal{F}_i^{(q,0)}(x_k) \text{ for } 1 \leq k \leq p, \quad (23)$$

i.e., x_1, \dots, x_p are all in exactly the same edges of $\mathcal{F}_i^{(q,0)}$.

Clearly, $\mathcal{F}_i^{(r,0)}(v) \subseteq \mathcal{F}_i(v)$ for $0 \leq r \leq q$ and $v \in V(\mathcal{F}_i)$, and by definition, $\mathcal{F}_i^{(0,0)} = \mathcal{F}_i$. Recall that y_1 is a vertex of maximum potential in \mathcal{F}_i , i.e., $\Phi(\mathcal{F}_i(y_1)) = \max_w \Phi(\mathcal{F}_i(w))$. Since equations (17) and (19) imply that

$$\Phi(\mathcal{F}_i(y_1)) = \Phi(\mathcal{F}_i^{(k-1,0)}(y_k)) \text{ for } 1 \leq k \leq q$$

and

$$\Phi(\mathcal{F}_i(y_1)) = \Phi(\mathcal{F}_i^{(q,0)}(x_j)) \text{ for } 1 \leq j \leq p,$$

we must have that

$$\mathcal{F}_i^{(k-1,0)}(y_k) = \mathcal{F}_i(y_k) \text{ for } 1 \leq k \leq q$$

and

$$\mathcal{F}_i^{(q,0)}(x_j) = \mathcal{F}_i(x_j) \text{ for } 1 \leq j \leq p,$$

or else we could use Observation 2 to contradict the fact that y_1 is a vertex of maximum potential in \mathcal{F}_i . It is now clear that we must also have

$$\Phi(\mathcal{F}_i(y_k)) = \Phi(\mathcal{F}_i(x_j)), \quad 1 \leq k \leq q, \quad 1 \leq j \leq p.$$

□

Observation 3 *If $X_{i+1} = \{x_1, \dots, x_p\}$, then we will often use $\mathcal{F}_i(x_1)$ to denote the set of edges that contain Maker's $(i+1)^{\text{th}}$ move. This is because $\mathcal{F}_i(X_{i+1}) = \mathcal{F}_i(x_1)$, by the last line of Claim 3.*

We also have the following two corollaries of Claim 3. Corollary 1 follows from equations (14) and (15) and Observation 3; while Corollary 2 follows from Corollary 1, equation (16), and the last line of Claim 3.

Corollary 1 *Suppose the basic hypotheses hold for turns 1 through i , $Y_i = \{y_1, \dots, y_q\}$ is Breaker's valid move during turn i , $X_{i+1} = \{x_1, \dots, x_p\}$ is Maker's valid response in turn $i+1$, and for convenience, we let $y_0 = x_1$. Then we can partition the edges of \mathcal{F}_i into $\mathcal{F}_i(y_1), \mathcal{F}_i(y_2), \dots, \mathcal{F}_i(y_q), \mathcal{F}_i(y_0), S_i$, where S_i is the set of edges from \mathcal{F}_i which do not contain any of Breaker's vertices from turn i or Maker's vertices from turn $i+1$.*

Corollary 2 *Suppose the basic hypotheses hold for turns 1 through i , $Y_i = \{y_1, \dots, y_q\}$ is Breaker's valid move during turn i , and $X_{i+1} = \{x_1, \dots, x_p\}$ is Maker's valid response in turn $i+1$. Then $(Y_i \setminus \{y_k\}) \cup \{x_j\}$ would also work as a valid Breaker's move during turn i for each $1 \leq k \leq q$ and each $1 \leq j \leq p$.*

Thus, if the basic hypotheses hold for turns 1 through i , then Corollary 1 and equation (16) give us

$$\Phi(\mathcal{F}_i) = \sum_{k=0}^q \Phi(\mathcal{F}_i(y_k)) + \Phi(S_i) = (q+1)\Phi(\mathcal{F}_i(y_j)) + \Phi(S_i), \text{ for any } 0 \leq j \leq q. \quad (24)$$

Then we use Lemma 2, which says $\Phi(\mathcal{F}_i) = 1$, and equation (24), to give us the following lemma.

Lemma 3 *Suppose the basic hypotheses hold for turns 1 through i . Then $S_i \neq \emptyset$ if and only if $\Phi(\mathcal{F}_i(y_k)) < \frac{1}{q+1}$ for each $0 \leq k \leq q$.*

Claim 3, equation (24), and Lemma 3 lead to the following lemma, which tells us that at the end of the game, $S_i = \emptyset$ and Maker's last move is a single edge of cardinality p .

Lemma 4 *Suppose that the basic hypotheses hold for turns 1 through i and that Maker wins during turn $i+1$. Then $S_i = \emptyset$, and $\mathcal{F}_i(x_1) = \{A\}$ where $|A| = p$, i.e., $\mathcal{F}_i(x_1)$ contains a single edge of size p .*

Proof of Lemma 4: If Maker wins during turn $i+1$, then $\mathcal{F}_i(x_1)$ must contain at least one edge A with $|A| \leq p$. Since the potential of a single edge of size p is $\frac{1}{q+1}$, then $\Phi(\mathcal{F}_i(x_1)) \geq \frac{1}{q+1}$. Thus, by Lemma 3, $S_i = \emptyset$. By Lemma 2, $\Phi(\mathcal{F}_i) = 1$. Thus, equation (24) and $S_i = \emptyset$ imply that $\Phi(\mathcal{F}_i(y_0)) = \frac{1}{q+1}$. Since $y_0 = x_1$, then $\mathcal{F}_i(x_1) = \{A\}$ where $|A| = p$. \square

3.2 Proof of Theorem 2

Proof of Theorem 2: Our proof proceeds by induction on $\max_{A \in \mathcal{F}} [|A|/p]$ (a measure of the largest edge in \mathcal{F}). Let \mathcal{F} be an extremal hypergraph for the $(p : q)$ -Erdős-Selfridge theorem. Recall that because every extremal hypergraph is also a conditionally extremal hypergraph and because a Maker's winning strategy is also a $(p : q)$ -Erdős-Selfridge-defeating strategy, then all of the results from Section 3.1 still hold. Let $n = \max_{A \in \mathcal{F}} [|A|/p]$. Our base case will be when $n = 1$. If $n = 1$, this implies that all edges have size p or smaller. Claim 2 implies that \mathcal{F} has a root node, thus, every edge must contain at least p vertices. Since a single edge with p vertices contributes $\frac{1}{q+1}$ to the potential of \mathcal{F} and since $\Phi(\mathcal{F}) = \frac{1}{q+1}$ because it is extremal, \mathcal{F} must consist of a single edge with p vertices. This is a $(q+1)$ -ary tree with one level.

Let us now assume that our theorem holds when $\max_{A \in \mathcal{F}} [|A|/p] \leq n-1$ and show that it holds when $\max_{A \in \mathcal{F}} [|A|/p] = n$. By Claim 2, \mathcal{F} has a root node, thus, Maker must occupy all p vertices in the root for her first move or she cannot win. Moreover, after Maker occupies the root node, every partial edge A satisfies $[|A|/p] \leq n-1$, thus we are free to utilize the inductive hypothesis.

We need to declare a set of *new hypotheses* that we will assume for the duration of this proof. The only difference between the new hypotheses and the basic hypotheses is

that we now make the stronger assumptions that \mathcal{F} is extremal and Maker uses a winning strategy.

New Hypotheses (for turns 1 through i):

\mathcal{F} is an extremal hypergraph for the $(p : q)$ -Erdős-Selfridge theorem, Breaker uses the strategy from the $(p : q)$ -Erdős-Selfridge theorem for turns 1 through i , and Maker uses a fixed winning strategy for turns 1 through $i + 1$.

Notice that if the new hypotheses hold, then the basic hypotheses also hold since an extremal hypergraph is also conditionally extremal and a Maker’s winning strategy is also a $(p : q)$ -Erdős-Selfridge-defeating strategy. As before, let us say that a Breaker’s move is *valid* if it follows (i.e., is consistent with) the strategy from the proof of the $(p : q)$ -Erdős-Selfridge theorem; and for the duration of this proof, let us say that a Maker’s move is *winning* if it follows a fixed winning strategy.

We now state a claim which uses the notation from Corollary 1 and Lemma 3 and makes use of the inductive hypothesis.

Claim 4 *Suppose the new hypotheses hold for turns 1 through i . Then $S_i = \emptyset$ if and only if \mathcal{F}_i is $q + 1$ disjoint $(q + 1)$ -ary trees.*

Proof of Claim 4: Suppose \mathcal{F}_i is $q + 1$ disjoint $(q + 1)$ -ary trees T_1, T_2, \dots, T_{q+1} . Since the root of each tree T_i has potential $\frac{1}{q+1}$, and since Breaker occupies each vertex y_1, y_2, \dots, y_q so that it has maximum potential, then $\Phi(\mathcal{F}_i(y_k)) \geq \frac{1}{q+1}$, for $1 \leq k \leq q$. Thus, by Lemma 3, $S_i = \emptyset$.

Now suppose $S_i = \emptyset$. By Corollary 1, since $S_i = \emptyset$, we can partition \mathcal{F}_i into $\mathcal{F}_i(y_0), \mathcal{F}_i(y_1), \dots, \mathcal{F}_i(y_q)$. Since $\Phi(\mathcal{F}_i) = 1$ (by Lemma 2) and $S_i = \emptyset$, then by equation (24), we have $\Phi(\mathcal{F}_i(y_j)) = \frac{1}{q+1}$ for $0 \leq j \leq q$. Let $Y_i^{(k)} = \{y_0, y_1, \dots, y_q\} \setminus \{y_k\}$. If Breaker chooses $Y_i^{(k)}$ (which is a valid move by Corollary 2) as his i^{th} move, then only $\mathcal{F}_i(y_k)$ is left alive. Therefore, by the inductive hypothesis, $\mathcal{F}_i(y_k)$ must be a $(q + 1)$ -ary tree for each $0 \leq k \leq q$, because $\Phi(\mathcal{F}_i(y_k)) = \frac{1}{q+1}$ and Maker has a winning strategy on $\mathcal{F}_i(y_k)$. Thus, $\mathcal{F}_i(y_0), \mathcal{F}_i(y_1), \dots, \mathcal{F}_i(y_q)$ is a collection of $q + 1$ edge-disjoint $(q + 1)$ -ary trees. We must now show that they are vertex-disjoint.

Assume towards a contradiction that $\mathcal{F}_i(y_0), \mathcal{F}_i(y_1), \dots, \mathcal{F}_i(y_q)$ are not vertex-disjoint. W.l.o.g., $V(\mathcal{F}_i(y_0)) \cap V(\mathcal{F}_i(y_1)) \neq \emptyset$. Let $I = V(\mathcal{F}_i(y_0)) \cap V(\mathcal{F}_i(y_1))$ be the set of vertices on which those two $(q + 1)$ -ary trees intersect. For $w \in V(\mathcal{F}_i(y_0))$ we define $l_0(w)$ to be the *level* of w in $\mathcal{F}_i(y_0)$, i.e., the tree-distance from the node containing w to the root node of $\mathcal{F}_i(y_0)$ plus 1. Similarly, let $l_1(v)$ be the level of a vertex v in $\mathcal{F}_i(y_1)$. We then define the *dual-level* of $w \in I$ to be $l(w) = \min\{l_0(w), l_1(w)\}$. See Figure 5 for an example. The left tree in Figure 5 is $\mathcal{F}_i(y_0)$ and its vertices and lines are solid, while the right tree is $\mathcal{F}_i(y_1)$ and its vertices are unfilled and its lines are dashed. Also note that $\mathcal{F}_i(y_1)$ is drawn upside down in order to emphasize that there need not be any relation between the levels of the vertices in the two trees. For the two vertices of intersection u and x , we have $l_0(u) = 2$, $l_1(u) = 4$, thus $l(u) = 2$, while $l_0(x) = 4$, $l_1(x) = 3$, thus $l(x) = 3$.

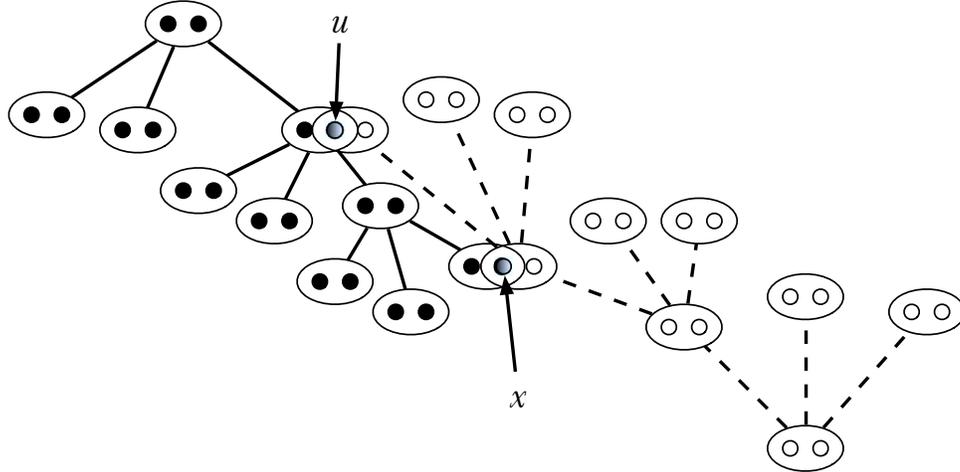


Figure 5: In this picture, $l_0(u) = 2$, $l_1(u) = 4$, thus $l(u) = 2$; while $l_0(x) = 4$, $l_1(x) = 3$, thus $l(x) = 3$. Also, $I = \{u, x\}$ and $I_{\min} = \{u\}$.

Let $\mu = \min_{v \in I} l(v)$, and let $I_{\min} = \{w \in I : l(w) = \mu\}$ be the set of vertices in I whose dual-level is minimum. Breaker chooses $\{u\} \cup \{y_2, y_3, \dots, y_q\}$ for his i^{th} move, where $u \in I_{\min}$ and satisfies $\max\{l_0(u), l_1(u)\} = \min_{v \in I_{\min}} \max\{l_0(v), l_1(v)\}$. Breaker's move kills all $(q + 1)$ -ary trees except $\mathcal{F}_i(y_0)$ and $\mathcal{F}_i(y_1)$, yet both $\mathcal{F}_i(y_0)$ and $\mathcal{F}_i(y_1)$ are damaged because $u \in V(\mathcal{F}_i(y_0)) \cap V(\mathcal{F}_i(y_1))$. Looking at Figure 5 we see that $I = \{u, x\}$ and $I_{\min} = \{u\}$, thus Breaker would select u as his vertex of intersection.

We outline an inductive argument as to how Breaker can win. W.l.o.g., let $l_0(u) = \mu$ and $l_1(u) = \nu \geq \mu$. Let $N_\mu^{(0)}$ be the node of $\mathcal{F}_i(y_0)$ that contains u , and let $N_\nu^{(1)}$ be the node of $\mathcal{F}_i(y_1)$ that contains u . Let $P^{(0)} = N_1^{(0)}, \dots, N_\mu^{(0)}$ be the path of nodes in $\mathcal{F}_i(y_0)$ from the root $N_1^{(0)}$ of $\mathcal{F}_i(y_0)$ to $N_\mu^{(0)}$, and let $P^{(1)} = N_1^{(1)}, \dots, N_\nu^{(1)}$ be the corresponding path in $\mathcal{F}_i(y_1)$. Initially, every living edge in $\mathcal{F}_i(y_0)$ contains the smallest indexed unoccupied node of $P^{(0)}$ (namely, the root $N_1^{(0)}$), and similarly for $\mathcal{F}_i(y_1)$ and $P^{(1)}$. Breaker forces this property to hold at the beginning of each turn for the duration of the game. Indeed, we may assume that at each turn Maker occupies the smallest indexed unoccupied node from either $P^{(0)}$ or $P^{(1)}$, or else, since Breaker occupies $q \geq 2$ vertices per turn, he will occupy one vertex from the smallest indexed unoccupied node from each path $P^{(0)}$ and $P^{(1)}$ and kill all remaining edges. If Maker occupies the smallest indexed unoccupied node $N_r^{(0)}$ from $P^{(0)}$, then Breaker responds by occupying one vertex from every child of $N_r^{(0)}$ except its child $N_{r+1}^{(0)}$ contained in $P^{(0)}$. Breaker does likewise if Maker occupies the smallest indexed unoccupied node $N_s^{(1)}$ from $P^{(1)}$. (This is how Breaker maintains the aforementioned property.) None of the vertices in $N_1^{(0)}, \dots, N_{\mu-1}^{(0)}$ intersect $\mathcal{F}_i(y_1)$ by the choice of u . Thus, Maker's moves in these nodes do not affect $\mathcal{F}_i(y_1)$. Additionally, Maker's moves in $N_1^{(1)}, \dots, N_{\nu-1}^{(1)}$ can only affect $\mathcal{F}_i(y_0)$ in levels $c \geq \mu + 1$ by the choice of u . Thus, Breaker is able to employ his strategy described above. Eventually, every living edge in $\mathcal{F}_i(y_0)$ contains $N_\mu^{(0)}$ and every living edge in $\mathcal{F}_i(y_1)$ contains $N_\nu^{(1)}$. Thus, since

Breaker has already occupied $u \in N_\mu^{(0)} \cap N_\nu^{(1)}$, there are no living edges, which contradicts the fact that Maker has a winning strategy. Therefore, $\mathcal{F}_i(y_0), \mathcal{F}_i(y_1), \dots, \mathcal{F}_i(y_q)$ must also be vertex-disjoint. \square

Assuming that the new hypotheses hold, since Maker has a winning strategy and \mathcal{F} is finite, Lemma 4 implies that for any alternating sequence of winning and valid moves, we must encounter a turn j where $S_j = \emptyset$. If $S_1 = \emptyset$, then Claim 4 provides the $q + 1$ disjoint $(q + 1)$ -ary trees that are combined to form a single $(q + 1)$ -ary tree with the addition of the root guaranteed by Claim 2. However, if $S_1 \neq \emptyset$, then we will show that there are $q + 1$ alternating sequences of winning and valid moves that satisfy the assumptions for Section 3.3. In which case, Lemma 9 from Section 3.3 will hold. But this will contradict the fact that Maker has a winning strategy, since Lemma 9 states that under the assumptions for Section 3.3, Breaker can deviate from the $(p : q)$ -Erdős–Selfridge strategy to win.

Let $\mathbf{Seq} = (X_1, Y_1, X_2, Y_2, \dots, X_m, Y_m, X_{m+1})$ be a maximum length alternating sequence of winning and valid moves with $S_m \neq \emptyset$, $m \geq 1$. Since $S_m \neq \emptyset$, Claim 4 implies that \mathcal{F}_m is not $q + 1$ disjoint $(q + 1)$ -ary trees. Thus, the first assumption from Section 3.3 holds. Let $Y_m = \{y_1, \dots, y_q\}$, $X_{m+1} = \{x_1, \dots, x_p\}$, and as before, let $y_0 = x_1$. As in Claim 4, let $Y_m^{(k)} = \{y_0, y_1, \dots, y_q\} \setminus \{y_k\}$, for $0 \leq k \leq q$. By Corollary 2, $Y_m^{(k)}$ is a valid Breaker’s move for $0 \leq k \leq q$, so let $X_{m+1}^{(k)} = \{x_1^{(k)}, \dots, x_p^{(k)}\}$ be Maker’s winning response to $Y_m^{(k)}$, and let $S_m^{(k)}$ be as S_i is in Corollary 1. We apply Corollary 1 and equation (16) to $Y_m^{(k)}$ and $X_{m+1}^{(k)}$ for each $0 \leq k \leq q$ to obtain equation (24) for $0 \leq k \leq q$:

$$\begin{aligned} \Phi(\mathcal{F}_m) &= \sum_{y \in Y_m^{(k)}} \Phi(\mathcal{F}_m(y)) + \Phi(\mathcal{F}_m(x_1^{(k)})) + \Phi(S_m^{(k)}) \\ &= (q + 1)\Phi(\mathcal{F}_m(v)) + \Phi(S_m^{(k)}), \end{aligned} \quad \text{for any } v \in Y_m^{(k)} \cup \{x_1^{(k)}\}.$$

Since $S_m^{(0)} = S_m \neq \emptyset$ and $Y_m^{(0)} \cap Y_m^{(k)} \neq \emptyset$ for $0 \leq k \leq q$, these equations allow us to conclude that $\Phi(S_m^{(k)}) = \Phi(S_m)$ for $0 \leq k \leq q$. Thus, $S_m^{(k)} \neq \emptyset$ for $0 \leq k \leq q$. Therefore, $\mathbf{Seq}^{(k)} = (X_1, Y_1, X_2, Y_2, \dots, X_m, Y_m^{(k)}, X_{m+1}^{(k)})$ is a maximum length alternating sequence of winning and valid moves with $S_m^{(k)} \neq \emptyset$, for $0 \leq k \leq q$. Thus, if $(X_1, Y_1, X_2, Y_2, \dots, X_m, Y_m^{(k)}, X_{m+1}^{(k)}, Y_{m+1}^{(k)}, X_{m+2}^{(k)})$ is an alternating sequence of winning and valid moves, then by the maximality of $\mathbf{Seq}^{(k)}$, we must have $S_{m+1}^{(k)} = \emptyset$, for $0 \leq k \leq q$. Therefore, by Claim 4, $\mathcal{F}_{m+1}^{(k)}$ is $q + 1$ disjoint $(q + 1)$ -ary trees for $0 \leq k \leq q$. Thus, the second assumption from Section 3.3 holds. Finally, since $S_m^{(k)} \neq \emptyset$ for $0 \leq k \leq q$, Lemma 3 implies $\Phi(\mathcal{F}_m(y)) < \frac{1}{q+1}$ for $y \in Y_m^{(k)}$. Thus, the third assumption from Section 3.3 also holds. Since $\mathbf{Seq}^{(k)}$ is a sequence of valid moves for $0 \leq k \leq q$ and all three assumptions from Section 3.3 hold, Lemma 9 tells us that Breaker can deviate from the $(p : q)$ -Erdős–Selfridge strategy during turn m to win. This contradicts the fact that Maker has a winning strategy, therefore, it must be the case that $S_1 = \emptyset$ and \mathcal{F} is a $(q + 1)$ -ary tree. \blacksquare

3.3 Pseudo-trees

As before, let us say that a Breaker's move is *valid* if it follows (i.e., is consistent with) the strategy from the proof of the $(p : q)$ -Erdős–Selfridge theorem and that a Maker's move is *valid* if it follows a fixed $(p : q)$ -Erdős–Selfridge–defeating strategy.

We begin with a set of assumptions that will hold for the duration of this section. While these assumptions on their own may seem arbitrary, we see that such a situation arises in our proof of Theorem 2 in Section 3.2.

Assumptions for Section 3.3:

The sequence $(X_1, Y_1, X_2, Y_2, \dots, X_m, Y_m, X_{m+1})$ is a sequence of valid moves with $Y_m = \{y_1, \dots, y_q\}$, $X_{m+1} = \{x_1, \dots, x_p\}$, and $y_0 = x_1$. Let $Y_m^{(k)} = \{y_0, y_1, \dots, y_q\} \setminus \{y_k\}$, for $0 \leq k \leq q$. By Corollary 2, $Y_m^{(k)}$ is a valid Breaker's move for $0 \leq k \leq q$, so let $X_{m+1}^{(k)} = \{x_1^{(k)}, \dots, x_p^{(k)}\}$ be Maker's valid response to $Y_m^{(k)}$. Notice that in the special case $k = 0$, $Y_m^{(0)} = Y_m$ and $X_{m+1}^{(0)} = X_{m+1}$. Let $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$ be the hypergraph of surviving edges after Breaker occupies $Y_m^{(k)}$ during turn m and let $\mathcal{F}_{m+1}^{(k)}$ be the residual hypergraph left after Maker occupies $X_{m+1}^{(k)}$. Moreover,

1. \mathcal{F}_m is not $q + 1$ disjoint $(q + 1)$ -ary trees;
2. $\mathcal{F}_{m+1}^{(k)}$ is $q + 1$ disjoint $(q + 1)$ -ary trees for $0 \leq k \leq q$;
3. $\Phi(\mathcal{F}_m(y)) < 1$ for $y \in Y_m^{(k)}$ and $0 \leq k \leq q$.

We will now study the structure of each $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$. Since $\mathcal{F}_{m+1}^{(k)}$ is $q + 1$ disjoint $(q + 1)$ -ary trees, it is reasonable to conjecture that the structure of $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$ is somewhat close to that of a $(q + 1)$ -ary tree. We proceed to show that this is indeed the case.

We begin by partitioning the edges of $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$ into those which contain $X_{m+1}^{(k)}$ and those which do not. Recall, by Observation 3, $\mathcal{F}_m(x_1^{(k)}) = \mathcal{F}_m(X_{m+1}^{(k)})$. Since $Y_m^{(k)}$ is a valid Breaker's move and $X_{m+1}^{(k)}$ is a valid Maker's response for $0 \leq k \leq q$, then by Corollary 1, \mathcal{F}_m can be partitioned into $\bigcup_{y \in Y_m^{(k)}} \mathcal{F}_m(y)$, $\mathcal{F}_m(x_1^{(k)})$, $S_m^{(k)}$, where $S_m^{(k)}$ are those edges which do not contain any vertices from $Y_m^{(k)}$ or $X_{m+1}^{(k)}$. Once Breaker occupies $Y_m^{(k)}$, we are left with only $\mathcal{F}_m(x_1^{(k)})$ and $S_m^{(k)}$, thus, $\mathcal{F}_{m+\frac{1}{2}}^{(k)} = \mathcal{F}_m(x_1^{(k)}) \cup S_m^{(k)}$ where the union is disjoint.

We will often refer to $S_m^{(k)}$ as the set of *short edges* of $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$ and to $\mathcal{F}_m(x_1^{(k)})$ as the set of *long edges* of $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$, because if the edges of the $q + 1$ disjoint $(q + 1)$ -ary trees in $\mathcal{F}_{m+1}^{(k)}$ are uniform, then every edge in $\mathcal{F}_m(x_1^{(k)})$ has exactly p more vertices than those in $S_m^{(k)}$, namely, the vertices in $X_{m+1}^{(k)}$. Since $\mathcal{F}_{m+1}^{(k)}$ is $q + 1$ disjoint $(q + 1)$ -ary trees, we let $T_1^{(k)}, \dots, T_{q+1}^{(k)}$ be those trees and $r_1^{(k)}, \dots, r_{q+1}^{(k)}$ be the roots of those trees. Thus, $T_j^{(k)} = \mathcal{F}_{m+1}^{(k)}(r_j^{(k)})$ is a $(q + 1)$ -ary tree, for $1 \leq j \leq q + 1$. (We will sometimes call $r_1^{(k)}, \dots, r_{q+1}^{(k)}$

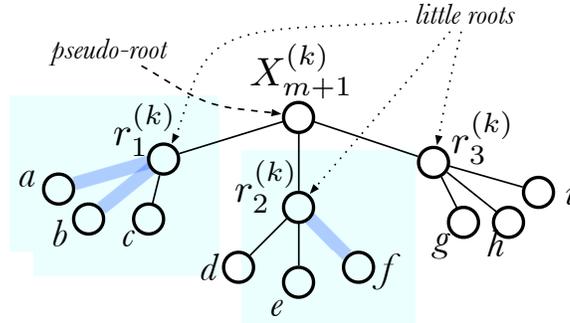


Figure 6: An example of a pseudo-3-ary tree $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$ with pseudo-root $X_{m+1}^{(k)}$ and little roots $r_1^{(k)}, r_2^{(k)}, r_3^{(k)}$.

little roots.) Stepping back, we let $t_j^{(k)} = \mathcal{F}_{m+\frac{1}{2}}^{(k)}(r_j^{(k)}) = \{A \in \mathcal{F}_{m+\frac{1}{2}}^{(k)} : r_j^{(k)} \subset A\}$. We then think of $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$ as a pseudo- $(q+1)$ -ary tree with pseudo-root $X_{m+1}^{(k)}$, and we think of $t_j^{(k)}$ as a pseudo-subtree of $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$. This is because once $X_{m+1}^{(k)}$ is occupied, $t_j^{(k)} = \mathcal{F}_{m+\frac{1}{2}}^{(k)}(r_j^{(k)})$ becomes $T_j^{(k)} = \mathcal{F}_{m+1}^{(k)}(r_j^{(k)})$, which is a true $(q+1)$ -ary tree. (We will sometimes call $t_j^{(k)}$ a *little tree*.)

Figures 6, 7, and 8 give pictorial representations of a pseudo-tree, its underlying structure, and how it evolves into $q+1$ disjoint $(q+1)$ -ary trees after Maker occupies its pseudo-root. Figure 6 shows a pseudo-3-ary tree. The pseudo-root is labeled $X_{m+1}^{(k)}$, thus all long edges contain $X_{m+1}^{(k)}$. Each short edge is a path from a leaf to a little root $r_i^{(k)}$ ($i = 1, 2, 3$) using only thick, lightly shaded edges, e.g., the path from f to $r_2^{(k)}$ is a short edge, as is the path from a to $r_1^{(k)}$, which we will denote by $(f, r_2^{(k)})$ and $(a, r_1^{(k)})$, respectively (even though an edge is technically the union of the vertices of the nodes on such a path). Each long edge is a path from a leaf to the pseudo-root $X_{m+1}^{(k)}$ using only thin, dark lines, e.g., the path from g to $X_{m+1}^{(k)}$ is a long edge, as is the path from d to $X_{m+1}^{(k)}$, denoted by $(g, r_3^{(k)}, X_{m+1}^{(k)})$ and $(d, r_2^{(k)}, X_{m+1}^{(k)})$, respectively. Figure 7 illustrates the three pseudo-subtrees $t_1^{(k)}, t_2^{(k)}, t_3^{(k)}$ that are contained in the pseudo-3-ary tree. The pseudo-subtrees are rooted at $r_1^{(k)}, r_2^{(k)}$, and $r_3^{(k)}$. Figure 8 shows how the pseudo-3-ary tree reduces to 3 disjoint 3-ary trees $T_1^{(k)}, T_2^{(k)}, T_3^{(k)}$ after Maker occupies the pseudo-root. Though uniformity is not necessary, it is sufficient for us to use an example where $\mathcal{F}_{m+1}^{(k)}$ (pictured in Figure 8) is 2-uniform in order to illustrate our definitions.

Lemma 5 *Every little tree contains at least one long edge, i.e.,*

$$t_j^{(k)} \cap \mathcal{F}_m^{(k)}(x_1^{(k)}) \neq \emptyset,$$

for $0 \leq k \leq q$ and $1 \leq j \leq q+1$,

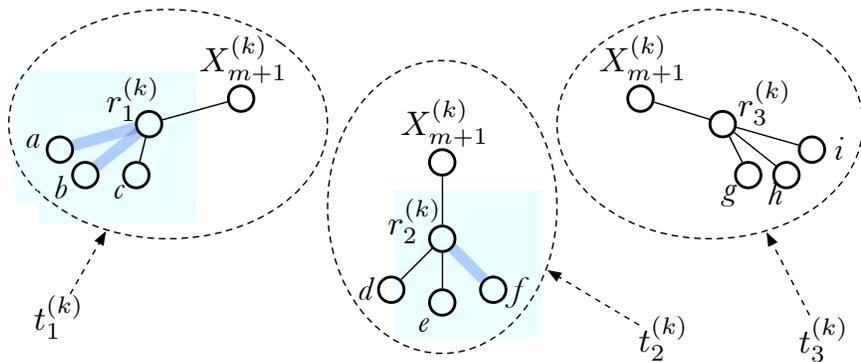


Figure 7: The three pseudo-subtrees that are contained in the pseudo-3-ary tree.

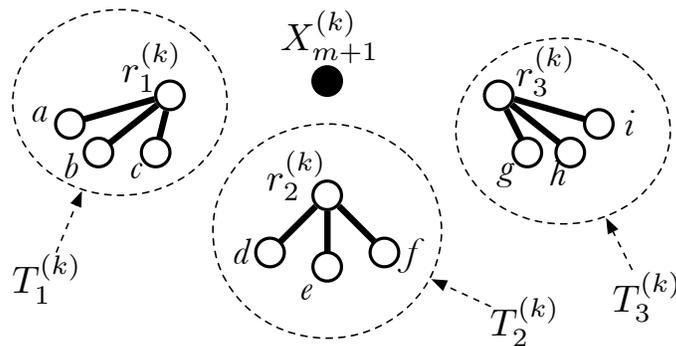


Figure 8: The three genuine 3-ary trees that are left after Maker occupies the pseudo-root.

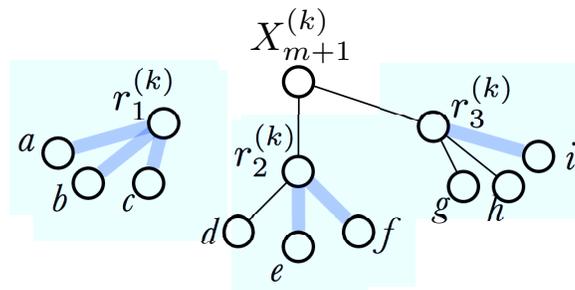


Figure 9: Example of a pseudo-3-ary tree $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$ with a forbidden little tree $t_1^{(k)}$ rooted at $r_1^{(k)}$.

Proof of Lemma 5: Recall that $\mathcal{F}_{m+\frac{1}{2}}^{(k)} = \mathcal{F}_m^{(k)}(x_1^{(k)}) \cup S_m^{(k)}$ where the union is disjoint. Assume towards a contradiction that $t_j^{(k)} \subseteq S_m^{(k)}$, i.e., $\mathcal{F}_{m+\frac{1}{2}}^{(k)}(r_j^{(k)}) \subseteq S_m^{(k)}$. (See Figure 9 where the little tree $t_1^{(k)} = \mathcal{F}_{m+\frac{1}{2}}^{(k)}(r_1^{(k)}) \subseteq S_m^{(k)}$.) The only edges of $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$ that are changed when Maker occupies $X_{m+1}^{(k)}$ are those edges in $\mathcal{F}_m^{(k)}(x_1^{(k)})$. Thus, since $\mathcal{F}_{m+\frac{1}{2}}^{(k)}(r_j^{(k)}) \subseteq S_m^{(k)}$, we have $t_j^{(k)} = \mathcal{F}_{m+\frac{1}{2}}^{(k)}(r_j^{(k)}) = \mathcal{F}_{m+1}^{(k)}(r_j^{(k)}) = T_j^{(k)}$, i.e., $t_j^{(k)}$ is already a $(q+1)$ -ary tree before Maker moves. But if this is the case, then for any $v \in r_j^{(k)}$, the potential of v in \mathcal{F}_m is too large because

$$\Phi(\mathcal{F}_m(v)) \geq \Phi(\mathcal{F}_{m+\frac{1}{2}}^{(k)}(v)) \geq \Phi(\mathcal{F}_{m+\frac{1}{2}}^{(k)}(r_j^{(k)})) = \Phi(\mathcal{F}_{m+1}^{(k)}(r_j^{(k)})) = \Phi(T_j^{(k)}) = \frac{1}{q+1}. \quad (25)$$

This contradicts Breaker's choice of $y \in Y_m^{(k)}$ as a vertex of maximum potential since $\Phi(\mathcal{F}_m(y)) < \frac{1}{q+1}$ by the third assumption for Section 3.3. \square

Let $\ell_j^{(k)} = t_j^{(k)} \cap \mathcal{F}_m(x_1^{(k)})$ be the set of long edges in $t_j^{(k)}$, e.g., in Figures 6 and 7, $\ell_2^{(k)} = \{(d, r_2^{(k)}, X_{m+1}^{(k)}), (e, r_2^{(k)}, X_{m+1}^{(k)})\}$. By Lemma 5, $\ell_j^{(k)} \neq \emptyset$ for $1 \leq j \leq q+1$, however, it is quite possible that $t_j^{(k)} \cap S_m^{(k)} = \emptyset$. Let us define $s_j^{(k)} = t_j^{(k)} \cap S_m^{(k)}$ to be the set of short edges in $t_j^{(k)}$, with the understanding that $s_j^{(k)}$ can be empty in some cases, e.g., in Figures 6 and 7, $s_1^{(k)} = \{(a, r_1^{(k)}), (b, r_1^{(k)})\}$ and $s_3^{(k)} = \emptyset$. Since $\mathcal{F}_{m+1}^{(k)}$ is $q+1$ disjoint $(q+1)$ -ary trees, we can prove the following facts about $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$.

Lemma 6 For $0 \leq k \leq q$, any two distinct pseudo-subtrees of $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$ are edge-disjoint and the intersection of their vertices is $X_{m+1}^{(k)}$. That is to say, $t_j^{(k)} \cap t_i^{(k)} = \emptyset$ and $V(t_j^{(k)}) \cap V(t_i^{(k)}) = X_{m+1}^{(k)}$, if $i \neq j$.

Proof of Lemma 6: By Lemma 5, every little tree contains at least one long edge, thus, $X_{m+1}^{(k)} \subseteq V(t_j^{(k)}) \cap V(t_i^{(k)})$. After Maker occupies $X_{m+1}^{(k)}$, $t_j^{(k)}$ and $t_i^{(k)}$ become the disjoint $(q+1)$ -ary trees $T_j^{(k)}$ and $T_i^{(k)}$. Thus, $(V(t_j^{(k)}) \setminus X_{m+1}^{(k)}) \cap (V(t_i^{(k)}) \setminus X_{m+1}^{(k)}) = \emptyset$, allowing us to conclude $V(t_j^{(k)}) \cap V(t_i^{(k)}) = X_{m+1}^{(k)}$. Since every edge in $t_j^{(k)}$ contains vertices from $V(T_j^{(k)})$ and every edge in $t_i^{(k)}$ contains vertices from $V(T_i^{(k)})$, and $T_j^{(k)}$ and $T_i^{(k)}$ are disjoint, we see that there is no edge common to both $t_j^{(k)}$ and $t_i^{(k)}$. \square

Corollary 3 For $0 \leq k \leq q$, let A and C be long edges in different pseudo-subtrees of $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$, i.e., let $A \in \ell_j^{(k)}$ and $C \in \ell_i^{(k)}$, $i \neq j$, then $A \cap C = X_{m+1}^{(k)}$.

Proof of Corollary 3: Let $A \in \ell_j^{(k)}$ and $C \in \ell_i^{(k)}$, where $i \neq j$. Since A and C are long edges, then $X_{m+1}^{(k)} \subseteq A \cap C$. However, $A \setminus X_{m+1}^{(k)} \subseteq V(T_j^{(k)})$ and $C \setminus X_{m+1}^{(k)} \subseteq V(T_i^{(k)})$, and $T_j^{(k)}$ and $T_i^{(k)}$ are disjoint, therefore $A \cap C = X_{m+1}^{(k)}$. \square

Corollary 4 For $0 \leq k \leq q$, let $A, C \in S_m^{(k)}$. Then $A \cap C \neq \emptyset$ if and only if $A, C \in s_j^{(k)}$ for some $1 \leq j \leq q + 1$.

Proof of Corollary 4: Let $A, C \in S_m^{(k)}$. For the first direction, suppose $A, C \in s_j^{(k)}$. Thus, $A, C \in T_j^{(k)}$, since the short edges of $t_j^{(k)}$ do not change when Maker occupies $X_{m+1}^{(k)}$. However, every edge in $T_j^{(k)}$ contains its root $r_j^{(k)}$, thus, $A \cap C \neq \emptyset$.

For the other direction, suppose $A \in s_j^{(k)}$ and $C \in s_i^{(k)}$, where $i \neq j$. As above, since short edges are not affected when Maker occupies $X_{m+1}^{(k)}$, we may conclude that $A \in T_j^{(k)}$ and $C \in T_i^{(k)}$. Since $T_j^{(k)}$ and $T_i^{(k)}$ are disjoint, $A \cap C = \emptyset$. \square

Recall that for $0 \leq k \leq q$, $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$ can be partitioned into its set of long edges and its set of short edges, i.e.,

$$\mathcal{F}_{m+\frac{1}{2}}^{(k)} = \mathcal{F}_m(x_1^{(k)}) \cup S_m^{(k)} \quad (\text{where } \mathcal{F}_m(x_1^{(k)}) \cap S_m^{(k)} = \emptyset). \quad (26)$$

By taking a different viewpoint, we can describe another partition of $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$. Recall that Corollary 1 says we can partition the edges of \mathcal{F}_m into $\mathcal{F}_m(y_0), \mathcal{F}_m(y_1), \dots, \mathcal{F}_m(y_q), S_m$. When Breaker occupies $Y_m^{(k)}$, he kills the edges in $\mathcal{F}_m(y)$ for each $y \in Y_m^{(k)}$, leaving only the edges in $\mathcal{F}_m(y_k)$ and S_m alive. Thus, for $0 \leq k \leq q$, we can also say

$$\mathcal{F}_{m+\frac{1}{2}}^{(k)} = \mathcal{F}_m(y_k) \cup S_m \quad (\text{where } \mathcal{F}_m(y_k) \cap S_m = \emptyset). \quad (27)$$

Luckily, we have the following lemma.

Lemma 7 For $0 \leq k \leq q$,

$$\mathcal{F}_m(x_1^{(k)}) = \mathcal{F}_m(y_k) \quad \text{and} \quad S_m^{(k)} = S_m,$$

i.e., $\mathcal{F}_m(y_k)$ is the set of long edges of $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$ and S_m is the set of short edges.

Proof of Lemma 7: Since $y_0 = x_1 = x_1^{(0)}$, equations (26) and (27) imply the lemma is true when $k = 0$. Now, consider the pseudo-tree $\mathcal{F}_{m+\frac{1}{2}}^{(k)} = \mathcal{F}_m(x_1^{(k)}) \cup S_m^{(k)}$ with pseudo-root $X_{m+1}^{(k)}$, where $k \geq 1$. If $y_k \in X_{m+1}^{(k)} = \{x_1^{(k)}, \dots, x_p^{(k)}\}$, then Claim 3 implies $\mathcal{F}_m(y_k) = \mathcal{F}_m(x_1^{(k)})$ because all of Maker's vertices are in exactly the same edges of \mathcal{F}_m ; and, by equations (26) and (27), it follows that $S_m^{(k)} = S_m$. If $y_k \notin X_{m+1}^{(k)}$, then y_k is contained in one of the little trees $t_j^{(k)}$; w.l.o.g., $y_k \in V(t_1^{(k)}) \setminus X_{m+1}^{(k)} = V(T_1^{(k)})$. By Lemma 5, $t_2^{(k)}$ and $t_3^{(k)}$ each contain at least one long edge. Consider long edges $A_2 \in t_2^{(k)}$ and $A_3 \in t_3^{(k)}$. Since they are long and in different little trees, Corollary 3 implies $A_2 \cap A_3 = X_{m+1}^{(k)}$. Since $y_k \in V(T_1^{(k)})$ and $V(T_1^{(k)})$ is disjoint from $V(t_2^{(k)})$ and $V(t_3^{(k)})$, then $y_k \notin A_2, A_3$, i.e., $A_2, A_3 \notin \mathcal{F}_m(y_k)$. Therefore, by equation (27), $A_2, A_3 \in S_m$, which allows us to conclude that $A_2, A_3 \in \mathcal{F}(y_0) \cup S_m = \mathcal{F}_{m+\frac{1}{2}}^{(0)}$ also! Since $A_2 \cap A_3 \neq \emptyset$ and $A_2, A_3 \in S_m = S_m^{(0)}$, Corollary 4 implies they are in the same little tree of $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$; w.l.o.g., $A_2, A_3 \in t_1^{(0)}$.

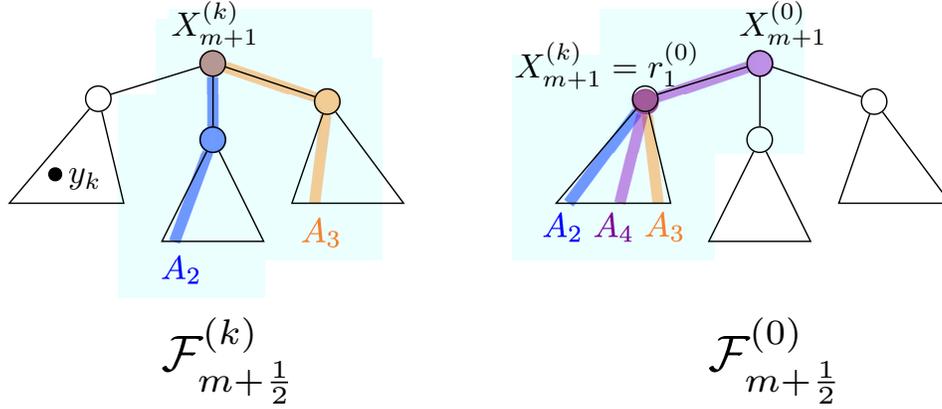


Figure 10: On the left is the pseudo-tree $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$. The long edges A_2 and A_3 are shown with thick blue and orange lines, respectively. On the right A_2 and A_3 are shown relative to $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$. The long edge A_4 is shown as a thick purple line. Notice $X_{m+1}^{(k)}, X_{m+1}^{(0)} \subseteq A_4$, which contradicts the fact that $\mathcal{F}_m(x_1^{(k)}) \cap \mathcal{F}_m(y_0) = \emptyset$.

Every edge of $t_1^{(0)}$ contains its little root $r_1^{(0)}$, thus, $r_1^{(0)} \subseteq A_2 \cap A_3 = X_{m+1}^{(k)}$. Since $|r_1^{(0)}| = p = |X_{m+1}^{(k)}|$, $r_1^{(0)} = X_{m+1}^{(k)}$. By Lemma 5, $t_1^{(0)}$ contains a long edge A_4 . Since every long edge of $t_1^{(0)}$ contains both its little root $r_1^{(0)} = X_{m+1}^{(k)}$ and the pseudo-root $X_{m+1}^{(0)}$, and since $\mathcal{F}_m(x_1^{(k)}) = \mathcal{F}_m(X_{m+1}^{(k)})$ and $\mathcal{F}_m(y_0) = \mathcal{F}_m(X_{m+1}^{(0)})$ by Observation 3, we conclude that $A_4 \in \mathcal{F}_m(x_1^{(k)}) \cap \mathcal{F}_m(y_0)$. However, $\mathcal{F}_m(x_1^{(k)}) \cap \mathcal{F}_m(y_0) = \emptyset$ because $x_1^{(k)} \in X_{m+1}^{(k)}$ and $y_0 \in Y_m^{(k)}$. (Here we used Corollary 1 applied to Breaker's valid move $Y_m^{(k)}$ and Maker's valid response $X_{m+1}^{(k)}$.) Therefore, by contradiction, it must be the case that $\mathcal{F}_m(y_k) = \mathcal{F}_m(x_1^{(k)})$ and $S_m^{(k)} = S_m$. (See Figure 10 for an illustration.) \square

Thus, Lemma 7 and equation (27) imply that each $\mathcal{F}_m(y_k) \cup S_m$ is a different pseudo-tree, and all of the pseudo-trees share the same set of short edges, S_m .

Recall that $s_i^{(0)}$ is the set of short edges from little tree $t_i^{(0)}$, i.e., $s_i^{(0)} = t_i^{(0)} \cap S_m$. W.l.o.g., let us index the little trees of $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$ so that for some $1 \leq h \leq q+1$, $s_i^{(0)} \neq \emptyset$ if $1 \leq i \leq h$, and $s_i^{(0)} = \emptyset$ if $i > h$. Using Corollary 4, we can show that if $s_i^{(0)} \cap s_{i_k}^{(k)} \neq \emptyset$, then $s_i^{(0)} = s_{i_k}^{(k)}$.

Lemma 8 *If $s_i^{(0)} \cap s_{i_k}^{(k)} \neq \emptyset$, then $s_i^{(0)} = s_{i_k}^{(k)}$.*

Proof of Lemma 8: Let $A \in s_i^{(0)} \cap s_{i_k}^{(k)}$. Since $A \in s_i^{(0)}$, then $C \in s_i^{(0)}$ if and only if $A \cap C \neq \emptyset$, by Corollary 4. Yet, by Corollary 4, $A \cap C \neq \emptyset$ if and only if $C \in s_{i_k}^{(k)}$, since $A \in s_{i_k}^{(k)}$. Therefore, $C \in s_i^{(0)}$ if and only if $C \in s_{i_k}^{(k)}$. \square

As a result of Lemma 8, we can re-index the little trees of each pseudo-tree so that $s_i^{(0)} = s_{i_k}^{(k)}$ for each $0 \leq k \leq q$ and each $1 \leq i \leq q+1$. (We are able to say $1 \leq i \leq q+1$

because for each $i > h$, $s_i^{(0)} = \emptyset = s_i^{(k)}$.) Thus, we can drop superscripts and write

$$S_m = s_1 \cup s_2 \cup \cdots \cup s_h, \text{ where } s_i = t_i^{(k)} \cap S_m \text{ and } 0 \leq k \leq q. \quad (28)$$

We now prove a claim which will help us show that Breaker has a winning strategy in our current setup if he deviates from the $(p : q)$ -Erdős-Selfridge strategy during turn m . Although the pseudo-subtrees that contain short edges will intersect because $s_i^{(k)} = s_i \neq \emptyset$ for $0 \leq k \leq q$ and $1 \leq i \leq h$, we will show for all pseudo-subtrees that we cannot have too many little roots intersecting. More specifically, we have the following claim.

Claim 5 *The intersection of any $q + 1$ distinct little roots is empty, i.e.,*

$$\bigcap_{k=0}^q r_{i_k}^{(k)} = \emptyset.$$

Proof of Claim 5: For convenience of notation, we will suppress the subscript k . If $v \in \bigcap_{k=0}^q r_i^{(k)}$, then $\mathcal{F}_m(v)$ contains every edge in $t_i^{(k)}$ for $0 \leq k \leq q$. If Maker occupies $X_{m+1}^{(k)}$, then every long edge in $t_i^{(k)}$ has its potential increased by a factor of $(q + 1)$, every short edge maintains its potential, and $t_i^{(k)}$ becomes the $(q + 1)$ -ary tree $T_i^{(k)}$. Thus,

$$\frac{1}{q + 1} = \Phi(T_i^{(k)}) = (q + 1)\Phi(\ell_i^{(k)}) + \Phi(s_i) \text{ for each } 0 \leq k \leq q. \quad (29)$$

Since equation (29) holds for $0 \leq k \leq q$, and s_i is independent of k , we see that

$$\Phi(\ell_i^{(k)}) = \Phi(\ell_i^{(j)}) \text{ for } 0 \leq k \leq j \leq q. \quad (30)$$

By equation (28), if $s_i \neq \emptyset$, then $s_i \subset t_i^{(k)}$ for $0 \leq k \leq q$. By Lemma 7, $\ell_i^{(k)} \subseteq \mathcal{F}_m(y_k)$ for $0 \leq k \leq q$, and by Corollary 1, $\mathcal{F}_m(y_k) \cap \mathcal{F}_m(y_j) = \emptyset$, for $k \neq j$, therefore, $\ell_i^{(k)} \cap \ell_i^{(j)} = \emptyset$ for $k \neq j$. Thus, the potential of the edges that contain v is at least the potential of the (shared) short edges s_i (if any) plus the potential of the long edges from each $t_i^{(k)}$, where $0 \leq k \leq q$:

$$\Phi(\mathcal{F}_m(v)) \geq \Phi(s_i) + \sum_{k=0}^q \Phi(\ell_i^{(k)}) = \Phi(s_i) + (q + 1)\Phi(\ell_i^{(0)}) = \frac{1}{q + 1}. \quad (31)$$

This contradicts Breaker's choice of $y \in Y_m^{(k)}$ as a vertex of maximum potential since $\Phi(\mathcal{F}_m(y)) < \frac{1}{q+1}$. Therefore, it must be the case that $\bigcap_{k=0}^q r_i^{(k)} = \emptyset$. \square

We will now state an important corollary of Claim 5 that will help us show that Breaker can win in our given setup. The corollary basically says that in each pseudo-subtree with short edges, there are more than p vertices that Breaker can use to kill all of the short edges in that pseudo-subtree.

Corollary 5 *If $s_i \neq \emptyset$, then there are more than p vertices that are contained in every edge in s_i , i.e.,*

$$\left| \bigcap_{A \in s_i} A \right| > p. \quad (32)$$

Proof of Corollary 5: Because $s_i = S_m \cap t_i^{(k)}$ for each $0 \leq k \leq q$ and because of the tree structure of each $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$, each edge $A \in s_i$ contains each little root $r_i^{(k)}$ for $0 \leq k \leq q$. Thus, $\bigcup_{k=0}^q r_i^{(k)} \subset \bigcap_{A \in s_i} A$. Since Claim 5 states that $\bigcap_{k=0}^q r_i^{(k)} = \emptyset$, and since each little root $r_i^{(k)}$ contains exactly p vertices, it must be the case that

$$p < \left| \bigcup_{k=0}^q r_i^{(k)} \right| \leq \left| \bigcap_{A \in s_i} A \right|. \quad (33)$$

□

We have established that \mathcal{F}_m is composed of $q + 1$ distinct pseudo-trees $\mathcal{F}_{m+\frac{1}{2}}^{(k)} = \mathcal{F}_m(y_k) \cup S_m$, where $0 \leq k \leq q$. Since $\Phi(\mathcal{F}_m(y_k)) < \frac{1}{q+1}$ for $0 \leq k \leq q$ by the third assumption for Section 3.3, Lemma 3 implies that $S_m \neq \emptyset$. Thus, any two pseudo-trees necessarily intersect via their common set of short edges S_m . However, the pseudo-root of a given pseudo-tree will be disjoint from all other pseudo-trees, as explained in the following claim.

Claim 6 For $k \neq j$, the pseudo-root $X_{m+1}^{(k)}$ of $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$ is disjoint from the pseudo-tree $\mathcal{F}_{m+\frac{1}{2}}^{(j)}$, i.e., $X_{m+1}^{(k)} \cap V(\mathcal{F}_{m+\frac{1}{2}}^{(j)}) = \emptyset$.

Proof of Claim 6: Let $v \in X_{m+1}^{(k)}$. Then, by Lemma 7 and the fact that all of Maker's vertices are in exactly the same edges, $\mathcal{F}_m(v) = \mathcal{F}_m(y_k)$. Assume $v \in V(\mathcal{F}_{m+\frac{1}{2}}^{(j)})$ also, then $v \in A$ for some $A \in \mathcal{F}_{m+\frac{1}{2}}^{(j)} = \mathcal{F}_m(y_j) \cup S_m$. Since $v \in A$, then $A \in \mathcal{F}_m(v) = \mathcal{F}_m(y_k)$. But if $k \neq j$, then this is a contradiction because $\mathcal{F}_m(y_k) \cap (\mathcal{F}_m(y_j) \cup S_m) = \emptyset$ by Corollary 1. □

The situation is somewhat similar to the one considered in the proof of Claim 4, except for the presence of the shared short edges which, by definition, cannot be killed by occupying vertices from any pseudo-root $X_{m+1}^{(k)}$. However, Corollary 5 reveals that Breaker will have a reasonable opportunity to kill the edges in S_m . See Figure 11 for a pictorial representation of how \mathcal{F}_m might appear. The hypergraph in Figure 11 corresponds to a $(2 : 2)$ game. Thus, there are three pseudo-trees, each distinguished by its own color. In this figure, $S_m = s_1$ and s_1 only contains one edge $(r_1^{(0)}, r_1^{(2)}) = (r_1^{(1)}, r_1^{(2)})$ (since $r_1^{(0)} = r_1^{(1)}$) which is represented with dotted lines.

Now we prove a lemma which states that Breaker can win in our given situation if he deviates from the $(p : q)$ -Erdős-Selfridge strategy during turn m .

Lemma 9 Under the assumptions for Section 3.3, Breaker can deviate from the $(p : q)$ -Erdős-Selfridge strategy during turn m to win.

Proof of Lemma 9: We use a proof by cases, but please note, in some of the cases below we only specify the vertices occupied by Breaker which are crucial to his winning strategy. Any unspecified vertices can be considered as chosen arbitrarily. We show that Breaker can win in all cases.

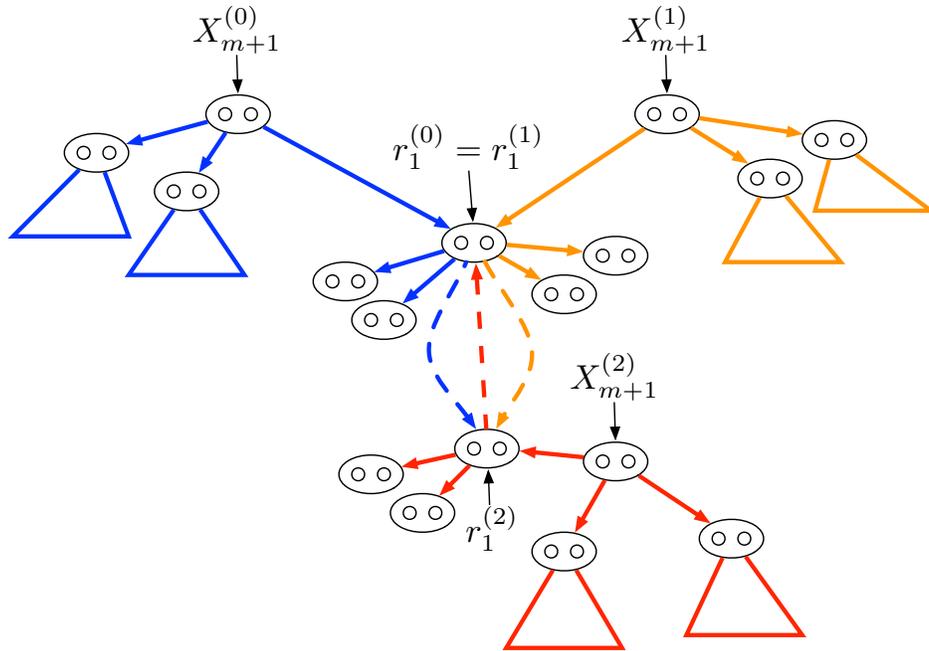


Figure 11: A possible configuration for \mathcal{F}_m in a $(2 : 2)$ game. Each pseudo-tree has its own color: $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$ is blue, $\mathcal{F}_{m+\frac{1}{2}}^{(1)}$ is orange, $\mathcal{F}_{m+\frac{1}{2}}^{(2)}$ is red. There is only one short edge (indicated by dotted lines) shared by all three pseudo-trees.

The two main cases that we need to consider are when there exist two little roots $r_{i_j}^{(j)}$ and $r_{i_k}^{(k)}$ from different pseudo-trees that intersect, and when no two little roots $r_{i_j}^{(j)}$ and $r_{i_k}^{(k)}$ from different pseudo-trees intersect.

Recall that by Corollary 1, \mathcal{F}_m is partitioned into $\mathcal{F}_m(y_0), \mathcal{F}_m(y_1), \dots, \mathcal{F}_m(y_q), S_m$, and by Lemma 7 and equation (27), $\mathcal{F}_{m+\frac{1}{2}}^{(k)} = \mathcal{F}_m(y_k) \cup S_m$ where $\mathcal{F}_m(y_k)$ is the set of long edges of $\mathcal{F}_{m+\frac{1}{2}}^{(k)}$ and S_m is the set of short edges for $0 \leq k \leq q$.

Case 1: *There exist two little roots $r_{i_j}^{(j)}$ and $r_{i_k}^{(k)}$ from different pseudo-trees that intersect, i.e., $r_{i_j}^{(j)} \cap r_{i_k}^{(k)} \neq \emptyset$.*

W.l.o.g., the two little roots which intersect are $r_{q+1}^{(0)}$ and $r_{q+1}^{(1)}$. If this is the case, for his move during turn m , Breaker occupies $\{y_2, y_3, \dots, y_q, v\}$ where $v \in r_{q+1}^{(0)} \cap r_{q+1}^{(1)}$. Breaker's move kills $\mathcal{F}_m(y_k)$ for $2 \leq k \leq q$, and it kills $t_{q+1}^{(0)}$ and $t_{q+1}^{(1)}$. All that is left of \mathcal{F}_m are the two pseudo-trees $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$ and $\mathcal{F}_{m+\frac{1}{2}}^{(1)}$ which are composed of the edges from $\mathcal{F}_m(y_0), \mathcal{F}_m(y_1)$, and S_m . In addition, each of the pseudo-trees is missing one pseudo-subtree, namely $t_{q+1}^{(0)}$ and $t_{q+1}^{(1)}$. We often use the fact that $X_{m+1}^{(0)}$ and $X_{m+1}^{(1)}$ are disjoint from all of the little roots within both surviving pseudo-trees $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$ and $\mathcal{F}_{m+\frac{1}{2}}^{(1)}$. This holds because the pseudo-root

is disjoint from its own little roots by definition, and by Claim 6, a pseudo-root from one pseudo-tree is disjoint from other pseudo-trees.

Case 1.1: *Maker responds by occupying one of the pseudo-roots $X_{m+1}^{(0)}$ or $X_{m+1}^{(1)}$.*

W.l.o.g., Maker occupies $X_{m+1}^{(0)}$. Breaker responds by occupying one vertex from each of the pseudo-subtree roots $r_1^{(0)}, \dots, r_q^{(0)}$. This finishes killing off $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$ because now all of the roots from its pseudo-subtrees have been killed by Breaker. In addition to killing $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$, Breaker also manages to kill S_m , because Breaker kills all of the edges in each $t_i^{(0)}$, including all of the short edges $s_i \subset t_i^{(0)}$. Therefore the only edges that are alive are contained in $\mathcal{F}_m(y_1)$. Thus, Maker must occupy $X_{m+1}^{(1)}$ for her next move. Then Breaker responds by occupying one vertex from each of the pseudo-subtree roots $r_1^{(1)}, \dots, r_q^{(1)}$. This finishes killing off $\mathcal{F}_{m+\frac{1}{2}}^{(1)}$, and Breaker wins.

Case 1.2: *Maker responds by occupying one of the little roots.*

It is easier to handle this case if we further divide it into two situations.

Case 1.2.1: *No little roots are identical, i.e., $r_j^{(0)} \neq r_k^{(1)}$ for all $1 \leq j, k \leq q$.*

W.l.o.g., Maker occupies $r_1^{(0)}$. Breaker responds by occupying one vertex from $X_{m+1}^{(0)}$ and one vertex from $r_1^{(1)}$. (We know there is a vertex available in $r_1^{(1)}$ because $r_1^{(0)} \neq r_1^{(1)}$.) Now each of the long edges of $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$ are dead and s_1 is dead. If Maker continues by occupying $X_{m+1}^{(1)}$, then Breaker occupies one vertex from each of the little roots $r_2^{(1)}, \dots, r_q^{(1)}$. After this, each edge in S_m is dead as is $\mathcal{F}_{m+\frac{1}{2}}^{(1)}$, and Breaker wins. If Maker does not occupy $X_{m+1}^{(1)}$, then Breaker occupies a vertex from $X_{m+1}^{(1)}$ and a vertex from $\bigcap_{A \in s_i} A$, for $2 \leq i \leq q$. This finishes killing each edge in S_m along with $\mathcal{F}_{m+\frac{1}{2}}^{(1)}$ and Breaker wins. (We use Corollary 5 to justify the availability of the $q - 1$ vertices Breaker uses to kill the short edges.)

Case 1.2.2: *There are little roots which are identical.*

When $q > 2$, we may assume that $r_1^{(0)} = r_1^{(1)}$ and Maker occupies $r_1^{(0)}$. In this case, Breaker responds by occupying a vertex from $X_{m+1}^{(0)}$, one from $X_{m+1}^{(1)}$, and a vertex $v_1 \in \bigcap_{A \in s_1} A$, whose availability is guaranteed by Corollary 5. This kills all long edges and the short edges in s_1 , thus, the only edges still alive are those in s_i for $2 \leq i \leq q$. No matter what Maker's next move is, Breaker occupies a vertex $v_i \in \bigcap_{A \in s_i} A$, for $2 \leq i \leq q$, whose availability is guaranteed by Corollary 5.

When $q = 2$, we notice that (during turn m), the pseudo-trees $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$, $\mathcal{F}_{m+\frac{1}{2}}^{(1)}$, and $\mathcal{F}_{m+\frac{1}{2}}^{(2)}$ cannot have too many "shared" little roots amongst them. For instance, if each pair of pseudo-trees $\mathcal{F}_{m+\frac{1}{2}}^{(i)}$ and $\mathcal{F}_{m+\frac{1}{2}}^{(j)}$, where $0 \leq i < j \leq 2$ share at least two little roots, then the pigeon-hole principle implies that there are in fact three little roots $r_k^{(0)} = r_k^{(1)} = r_k^{(2)}$ which are identical, contradicting Claim 5. Thus, there must be a pair of pseudo-trees $\mathcal{F}_{m+\frac{1}{2}}^{(i)}$ and $\mathcal{F}_{m+\frac{1}{2}}^{(j)}$ that share at most one little root. W.l.o.g., $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$ and $\mathcal{F}_{m+\frac{1}{2}}^{(1)}$ share at

most one little root, and the little root that they share is $r_3^{(0)} = r_3^{(1)}$. In this case, during turn m , Breaker occupies $y_2 \in X_{m+1}^{(2)}$ and $v \in r_3^{(0)}$. Thus, when Maker occupies a little root during turn $m + 1$, the situation reduces to Case 1.2.1, where *no little roots are identical* because $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$ and $\mathcal{F}_{m+\frac{1}{2}}^{(1)}$ share at most one identical little root, which Breaker already killed during turn m .

Case 1.3: *Maker does not occupy $X_{m+1}^{(0)}$, $X_{m+1}^{(1)}$, or a little root.*

If Maker occupies at least one vertex from $X_{m+1}^{(0)}$, then Breaker responds by occupying one vertex from each of $r_1^{(0)}, \dots, r_q^{(0)}$. Now Maker cannot finish occupying both $X_{m+1}^{(1)}$ and a little root before Breaker's next turn. If Maker does not fully occupy $X_{m+1}^{(1)}$, then Breaker occupies a vertex from $X_{m+1}^{(1)}$ and wins. Otherwise, every little root $r_1^{(1)}, \dots, r_q^{(1)}$ has at least one unoccupied vertex, and Breaker occupies one vertex from each of them and wins.

If Maker puts no vertices in $X_{m+1}^{(0)}$, then Breaker responds by occupying one vertex from $r_1^{(1)}, \dots, r_q^{(1)}$. Maker must use her next move to occupy $X_{m+1}^{(0)}$, or else Breaker takes a vertex from $X_{m+1}^{(0)}$ and wins. Since Maker did not occupy a little root during turn $m + 1$ or $m + 2$, each little root $r_1^{(0)}, \dots, r_q^{(0)}$ has an unoccupied vertex, and Breaker occupies one vertex from each of them and wins.

Case 2: *No two little roots $r_{i_j}^{(j)}$ and $r_{i_k}^{(k)}$ from different pseudo-trees intersect, i.e., $r_{i_j}^{(j)} \cap r_{i_k}^{(k)} = \emptyset$ for $j \neq k$.*

In this case, for his move during turn m , Breaker occupies $\{y_2, y_3, \dots, y_q, v\}$ where $v \in r_1^{(0)}$. Breaker's move kills $\mathcal{F}_m(y_k)$ for $2 \leq k \leq q$, and it kills $t_1^{(0)}$. All that is left of \mathcal{F}_m are the two pseudo-trees $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$ and $\mathcal{F}_{m+\frac{1}{2}}^{(1)}$ which are composed of the edges from $\mathcal{F}_m(y_0), \mathcal{F}_m(y_1)$, and S_m . In addition, $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$ is missing its pseudo-subtree $t_1^{(0)}$ and the edges in s_1 are dead. Notice, by Lemma 3 and the assumption that $\Phi(\mathcal{F}_m(y_k)) < \frac{1}{q+1}$ for $0 \leq k \leq q$, we have $S_m \neq \emptyset$. Thus, the indexing used in equation (28) implies that $s_1 \neq \emptyset$ also.

Case 2.1: *Maker responds by occupying $X_{m+1}^{(0)}$.*

Breaker responds by occupying one vertex from each of the pseudo-subtree roots $r_2^{(0)}, \dots, r_{q+1}^{(0)}$. This kills $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$ and S_m , so Maker must now occupy $X_{m+1}^{(1)}$. Breaker then occupies one vertex from each of $r_2^{(1)}, \dots, r_{q+1}^{(1)}$. Now $t_1^{(1)}$ is still alive and has transformed into $T_1^{(1)}$ since Maker occupied $X_{m+1}^{(1)}$. But all of the edges in $s_1 \subseteq T_1^{(1)}$ are dead. Since $s_1 \neq \emptyset$ and $s_1 \subseteq T_1^{(1)}$, then $\Phi(T_1^{(1)}) < \frac{1}{q+1}$. Thus, Breaker can win via the $(p : q)$ -Erdős-Selfridge strategy.

Case 2.2: *Maker responds by occupying $X_{m+1}^{(1)}$.*

Breaker responds by occupying one vertex from each of the pseudo-subtree roots $r_2^{(1)}, \dots, r_{q+1}^{(1)}$. This kills all edges in S_m and leaves only a damaged $T_1^{(1)}$ alive from $\mathcal{F}_{m+\frac{1}{2}}^{(1)}$. Recall

that the vertex v , which Breaker occupied in turn m , is contained in $T_1^{(1)}$ because v was used to kill the short edges $s_1 \subset T_1^{(1)}$. Thus, there is a unique path P in $T_1^{(1)}$ from its root $r_1^{(1)}$ to the node containing v . Suppose $P = N_1, N_2, \dots, N_k$ where each N_i is a node and $N_1 = r_1^{(1)}$ and N_k is the node containing v . By Corollary 6 below, since $v \in \bigcap_{A \in s_1} A$, then $N_k \subseteq \bigcap_{A \in s_1} A$. Also, by Corollary 7 below, since $N_k \subseteq \bigcap_{A \in s_1} A$, then $N_i \subseteq \bigcap_{A \in s_1} A$ for $1 \leq i \leq k$. Since $s_1 \subset T_1^{(0)}$, all vertices in $N_1 \cup \dots \cup N_k$ only intersect $\mathcal{F}_{m+\frac{1}{2}}^{(0)}$ in $t_1^{(0)}$, which is already dead. Breaker follows a strategy so that at the beginning of each turn, every living edge in $T_1^{(1)}$ contains N_i where i is the smallest index for which Maker has not occupied N_i . Initially this is true because every living edge contains the root $N_1 = r_1^{(1)}$. Breaker maintains this property by occupying one vertex from every child of N_i , except node N_{i+1} , immediately after Maker occupies N_i . The rest of Breaker's strategy is as follows. If Maker occupies $X_{m+1}^{(0)}$, then Breaker occupies one vertex from each of $r_2^{(0)}, \dots, r_{q+1}^{(0)}$. Then all that is left is $T_1^{(1)}$, in which Breaker will win by forcing Maker to node N_k . If Maker occupies N_i , then Breaker does as stated above. If Maker occupies neither $X_{m+1}^{(0)}$ nor N_i , then, since $q \geq 2$, Breaker immediately occupies one vertex from each of $X_{m+1}^{(0)}$ and N_i and wins.

Case 2.3: *Maker does not occupy $X_{m+1}^{(0)}$ or $X_{m+1}^{(1)}$.*

Breaker responds by occupying a vertex from $X_{m+1}^{(0)}$ and a vertex from $X_{m+1}^{(1)}$. Since none of the little roots intersect, $|\bigcap_{A \in s_i} A| \geq |r_i^{(0)} \cup \dots \cup r_i^{(q)}| = (q+1)p$ for each $2 \leq i \leq h$, where h is the largest index such that $s_i \neq \emptyset$. Maker occupies her next p vertices, then Breaker occupies a vertex $v_i \in \bigcap_{A \in s_i} A$, for $2 \leq i \leq h$, where $h \leq q+1$.

We have shown that Breaker wins in every situation. Thus, we are finished with the proof of Lemma 9. \square

We now provide the proofs of Corollaries 6 and 7 from Case 2.2. Let T be a finite rooted tree. We will make use of the natural partial order on the nodes of T , where $N_i \leq N_j$ if N_i is on the path between N_j and the root node. Given a set of nodes \mathcal{N} , the *meet* of \mathcal{N} is the node \hat{N} such that

1. $\hat{N} \leq N$ for every $N \in \mathcal{N}$, and
2. if $M \leq N$ for every $N \in \mathcal{N}$, then $M \leq \hat{N}$.

Given a nonempty set of nodes \mathcal{N} , notice that $\mathcal{A} = \{N' : N' \leq N \text{ for every } N \in \mathcal{N}\}$ must be a nonempty chain. Indeed, the root of T is in \mathcal{A} . Moreover, suppose $N \in \mathcal{N}$ and $N_1, N_2 \in \mathcal{A}$. Since $N_1 \leq N$ and $N_2 \leq N$, they are both on the path between N and the root, thus N_1 is related to N_2 . Since the partial order is finite and antisymmetric, the meet of a nonempty set exists and is unique. Recall that if \mathcal{T} is a $(q+1)$ -ary tree hypergraph, then an edge $A \in \mathcal{T}$ is determined by a leaf node N_A so that A is the union of all nodes on the path from N_A to the root node.

Lemma 10 *Let S be a nonempty set of hyperedges of a $(q + 1)$ -ary tree \mathcal{T} . Let \mathcal{N} be the set of leaf nodes that determine the edges in S , let \hat{N} be the meet of \mathcal{N} , and let N' be a node in \mathcal{T} . Then $N' \subseteq \bigcap_{A \in S} A$ if and only if $N' \leq \hat{N}$.*

Proof of Lemma 10: Let N' be a node, A a hyperedge, and N_A the leaf node that determines A . Then $N' \subseteq A$ if and only if $N' \leq N_A$. Thus, $N' \subseteq \bigcap_{A \in S} A$ if and only if $N' \leq N$ for every $N \in \mathcal{N}$. Therefore, $N' \subseteq \bigcap_{A \in S} A$ if and only if $N' \leq \hat{N}$, by the definition of the meet of \mathcal{N} . \square

Observation 4 *Since distinct nodes of \mathcal{T} are disjoint, every $v \in V(\mathcal{T})$ is contained in a unique node N_v .*

Lemma 11 *Let A be a hyperedge in \mathcal{T} , $v \in V(\mathcal{T})$, and N_v the unique node such that $v \in N_v$. If $v \in A$, then $N_v \subseteq A$.*

Proof of Lemma 11: Let A be the edge determined by the leaf node N_A . Then, by definition, $A = \bigcup_{N \leq N_A} N$. Thus, if $v \in A$, then $v \in N$ for some $N \leq N_A$. Moreover, $N \subseteq A$. Since N_v is the unique node that contains v , then $N = N_v$ and $N_v \subseteq A$. \square

Corollary 6 *Let S be a nonempty set of hyperedges of a $(q + 1)$ -ary tree \mathcal{T} , $v \in V(\mathcal{T})$, and N_v the unique node such that $v \in N_v$. If $v \in \bigcap_{A \in S} A$, then $N_v \subseteq \bigcap_{A \in S} A$.*

Proof of Corollary 6: Suppose $v \in \bigcap_{A \in S} A$, then $v \in A$ for each $A \in S$. Thus, by Lemma 11, $N_v \subseteq A$ for each $A \in S$, i.e., $N_v \subseteq \bigcap_{A \in S} A$. \square

Corollary 7 *Let S be a nonempty set of hyperedges of a $(q + 1)$ -ary tree \mathcal{T} , and let N be a node in \mathcal{T} . If $N \subseteq \bigcap_{A \in S} A$, then $N' \subseteq \bigcap_{A \in S} A$ for every $N' \leq N$.*

Proof of Corollary 7: Let \mathcal{N} be the set of leaf nodes that determine the edges in S , and let \hat{N} be the meet of \mathcal{N} . Suppose $N \subseteq \bigcap_{A \in S} A$, then by Lemma 10, $N \leq \hat{N}$. Let $N' \leq N$ be arbitrary. Then by transitivity, $N' \leq \hat{N}$. Thus, by Lemma 10, $N' \subseteq \bigcap_{A \in S} A$. \square

4 Conclusion

Theorem 2 settles the question of determining the structure of an extremal hypergraph \mathcal{F} for the $(p : q)$ -Erdős–Selfridge theorem for the case when $q \geq 2$: \mathcal{F} must be a $(q + 1)$ -ary tree where each node contains p vertices; and if \mathcal{F} is pn -uniform, then it must be a complete $(q + 1)$ -ary tree. As was mentioned in the Introduction, there are many extremal hypergraphs for the (unbiased) Erdős–Selfridge theorem, and the question of classifying all of the extremal hypergraphs for the Erdős–Selfridge theorem is a wide open problem. Perhaps we should begin by classifying all *economical* extremal hypergraphs for the Erdős–Selfridge theorem. Based on preliminary investigations, the author presents the following conjecture:

Conjecture 1 *Unlike non-economical extremal hypergraphs for the Erdős–Selfridge theorem, all economical extremal hypergraphs for the Erdős–Selfridge theorem are derived from special (to be determined) “good” labelings of binary trees (where a “good” labeling is as described in the Introduction).*

Acknowledgments

I would like to thank József Beck and Klay Kruczek for their valuable input and suggestions. I also wish to thank Tamás Lengyel, Greg Tollisen, and the anonymous referee for making many helpful suggestions and comments that improved the presentation of this paper.

References

- [1] J. Beck. Remarks on positional games. I. *Acta Math. Acad. Sci. Hungar.*, 40(1-2):65-71, 1982.
- [2] József Beck. *Combinatorial Games: Tic-Tac-Toe Theory*, Cambridge University Press, 2008.
- [3] József Beck. Inevitable randomness in discrete mathematics. *University Lecture Series, 49* American Mathematical Society, Providence, RI, xii+250, 2009.
- [4] P. Erdős and J. L. Selfridge. On a combinatorial game. *J. Combinatorial Theory Ser. A*, 14:298-301, 1973.
- [5] Xiaoyun Lu. A characterization on n -critical economical generalized tic-tac-toe games. *Discrete Math.*, 110(1-3):197-203, 1992.
- [6] E. L. Sundberg. Economical tight examples for the biased Erdős-Selfridge theorem. *Discrete Math.*, 308 (15): 3308-3314, 2008.