Search Games on Hypergraphs

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K.C.A.M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op dinsdag 28 januari 2014 om 12:30 uur

door

Christos PELEKIS

Master of Science in Applied Mathematics
National Technical University of Athens
geboren te Agrinio, Griekenland

Dit proefschrift is goedgekeurd door de promotor:
Prof. dr. F.M. Dekkinng

Copromotor: Dr. R.J. Fokkink

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
Prof. dr. F.M. Dekking,	Technische Universiteit Delft, promotor
Dr. R.J. Fokkink,	Technische Universiteit Delft, copromotor
Prof. dr. S. Alpern,	University of Warwick, Engeland
Prof. dr. S. Gal,	University of Haifa, Israël
Prof. dr. K. Roos,	Technische Universiteit Delft
Prof. dr. G. Schäfer,	Centrum voor Wiskunde en Informatica
Dr. C. Kraaikamp,	Technische Universiteit Delft
Prof. dr. P. Van Mieghem,	Technische Universiteit Delft, reservelid

Contents

Chapter 1

Conjectures of Kikuta-Ruckle, Erdős and Samuels

1.1 Basic notions and the hypergraph incidence game

This thesis is motivated by a certain type of two-person win-lose games that are played on hypergraphs. Such games are, by definition, instances in which there are two players, say Alice and Bob, and each player has a set of possible strategies, or moves. The set of all possible strategies of each player is called her/his *strategy space* and each element of the strategy space is referred to as a *pure strategy*. Let us denote by A the strategy space of Alice, B the strategy space of Bob and suppose that both A and B are finite sets, say $A = \{a_1, \ldots, a_m\}$ and $B = \{b_1, \ldots, b_k\}$. In this case the game is called *finite*. Alice, in private, chooses an element $a \in A$. Bob, also in private, chooses an element $b \in B$. The players then announce their choices and for any such pair of choices there is a corresponding *payoff*. That is, there is a function $f : A \times B \to \mathbb{R}$ such that $f(a, b)$ indicates the amount that Alice has to pay to Bob, after the pure strategies a and b have been played. In case $f(a, b) \in \{0, 1\}$ for any $a \in A$ and any $b \in B$, the game is called *win-lose*. Both players control the game in the sense that their choices influence the outcome and both logically study the way to achieve their best possible payoff. So Bob would like the game to end with a payoff that is as large as possible and Alice would like the game to end with a payoff that is as small as possible. We may view the values $f(a, b)$ as being entries of a $|A| \times |B|$ matrix, M. It might be that the matrix M has a *saddle point*, i.e., an element $f(a_0, b_0)$ such that

$$\min_{a \in A} f(a, b_0) = f(a_0, b_0) = \max_{b \in B} f(a_0, b).$$

Hence $f(a_0, b_0)$ is maximum in its row and minimum in its column. In case the players have chosen a saddle point, we say that the game is in *equilibrium*, in the sense that no player has an intension to change her/his strategy, given that the other player plays in the same way. In such a case we say that the pure strategies a_0, b_0 solve the game. However, there are cases in which a saddle point does not exist. This fact leads to the idea of using *mixed* strategies, i.e., to allow the players to choose pure strategies according to some probability distribution over their strategy space. Suppose that the game is played again and again. Let $\alpha = (\alpha_1, \ldots, \alpha_m)$, where α_i is the proportion of times that Alice chooses the pure strategy $a_i \in A$, let $\beta = (\beta_1, \ldots, \beta_k)$, where β_i is the proportion of times that Bob chooses the pure strategy $b_i \in B$. Then the expected payoff to Bob is given by

$$\phi(\alpha, \beta) := \sum_i \sum_j f(a_i, b_j)\alpha_i \beta_j.$$

A fundamental result of von Neumann states that

$$\min_\alpha \max_\beta \phi(\alpha, \beta) = \max_\beta \min_\alpha \phi(\alpha, \beta).$$

This minimax value, η, is called the *value* of the game. An equivalent result is that there exist mixed strategies $\bar{\alpha}$ and $\bar{\beta}$ such that

$$\min_\alpha \phi(\alpha, \bar{\beta}) = \phi(\bar{\alpha}, \bar{\beta}) := \eta = \max_\beta \phi(\bar{\alpha}, \beta).$$

Thus η is a "saddle point" in mixed strategies, or in other words, if Alice chooses the mixed strategy $\bar{\alpha}$ and Bob chooses the mixed strategy $\bar{\beta}$, then both players can guarantee an expected payoff of η. A different way to state this is that if Alice chooses $\bar{\alpha}$ then she never has to pay more than η, no matter how Bob plays. Similarly, there is a mixed strategy, $\bar{\beta}$, for Bob such that his expected payoff is at least η, no matter what Alice plays. We will apply von Neumann's result to find the value a finite game in Theorem 1.1.1 below.

We will also need some definitions from the theory of finite sets. A *hypergraph*, $\mathcal{H}$, is a pair $(V, \mathcal{E})$, where V is a finite set and $\mathcal{E}$ is a family of subsets of V. The set V is called the *vertex set* of $\mathcal{H}$. The set $\mathcal{E}$ is called the *edge set* of

$\mathcal{H}$ and its elements are called *hyperedges*, or just edges. Note that in case all edges of $\mathcal{E}$ are doubletons, then we are in the case of a graph. An *edge covering* of $\mathcal{H}$ is a collection of hyperedges $E_1, \ldots, E_t$ such that $V \subseteq E_1 \cup \cdots \cup E_t$. The smallest t for which this is possible is called the *edge covering number* of $\mathcal{H}$ and is denoted by $\kappa(\mathcal{H})$. A vertex $v \in V$ is called *exposed* if it belongs to no hyperedge. Thus, if $\mathcal{H}$ has an exposed vertex, we have $\kappa(\mathcal{H}) = \infty$ and from now on we will assume that the hypergraphs under consideration have no exposed vertices. Finding the edge covering number of a hypergraph is an optimization problem. To see this, denote by $\mathbf{A}$ the *incidence* matrix of $\mathcal{H}$. That is, the matrix whose rows are represented by the vertices, $v_1, \ldots, v_n$, the columns are represented by the edges, $E_1, \ldots, E_m$, and whose elements, a_{ij}, are equal to 1 if $v_i \in E_j$ and equal to 0 otherwise. Let also $\mathbf{x}$ be an indicator vector of the sets that have been selected for the edge cover. Given such an indicator vector, $\mathbf{x}$, let $\mathcal{E}_{\mathbf{x}}$ be the set of edges from $\mathcal{E}$ that correspond to this vector. Then $\mathbf{x}$ is an indicator vector of an edge covering if and only if $\mathbf{A} \cdot \mathbf{x} \geq 1$. To see this just note that, for every $\mathbf{x}$, coordinate j of $\mathbf{A} \cdot \mathbf{x}$ equals the number of edges from $\mathcal{E}_x$ that contain v_j. Hence $\kappa(\mathcal{H})$ is the value of the optimization problem

$$\begin{aligned} &\text{minimize } \mathbf{1}^t \mathbf{x}, \\ &\text{subject to: } \mathbf{A} \cdot \mathbf{x} \geq 1, \mathbf{x} \in \{0,1\}^m \end{aligned}$$

where $\mathbf{1}$ is the vector of all ones. Problems of minimizing/maximizing a certain linear function under linear constraints and under the assumption that the variables are restricted to be integers belong to the field of Integer Programming (IP). The problem of finding the edge covering number of a hypergraph has a natural dual. A *vertex packing* in a hypergraph, $\mathcal{H} = (V, \mathcal{E})$, is a subset $X \subseteq V$ with the property that no two elements of X belong to the same element of $\mathcal{E}$. The *vertex packing number*, of $\mathcal{H}$, denoted $p(\mathcal{H})$, is defined as the largest cardinality of a vertex packing. In the case of a graph, the vertex packing number is its independence number, i.e., the maximum cardinality of a set of vertices no two of which are adjacent.

Finding the vertex packing number of a hypergraph is also an IP problem. To see this, let $\mathbf{y}$ be an indicator vector of the vertices that have been selected for the vertex packing. Then X is a vertex packing if and only if $\mathbf{A}^t \cdot \mathbf{y} \geq 1$, where $\mathbf{A}$ is as above and so $p(\mathcal{H})$ is the value of the IP

$$\begin{aligned} &\text{maximize } \mathbf{1}^t \mathbf{y}, \\ &\text{subject to: } \mathbf{A}^t \cdot \mathbf{y} \leq 1, \mathbf{y} \in \{0,1\}^n \end{aligned}$$

Integer Programming problems are in general difficult. In contrast to this, the field of Linear Programming (LP) in which a linear function has to be minimized/maximized, under linear constraints and under the assumption that the variables are *real* numbers, is easier. One approach to integer programming problems is by its linear relaxation to a Linear Programming problem. This approach to hypergraph problems is called fractional graph theory (see [47]). Thus, if we allow the coordinates of the vectors $\mathbf{x}$ and $\mathbf{y}$ above to take values in the set $[0, 1]$, then we have fractional analogues of the edge covering number and and vertex packing number. Thus define $\kappa_f(\mathcal{H})$ to be the value of the LP

$$\begin{aligned}&\text{minimize } \mathbf{1}^t\mathbf{x},\\ &\text{subject to: } \mathbf{A}\cdot\mathbf{x} \geq 1, \mathbf{x} \in [0,1]^m\end{aligned}$$

and $p_f(\mathcal{H})$ to be the value of the LP

$$\begin{aligned}&\text{maximize } \mathbf{1}^t\mathbf{y},\\ &\text{subject to: } \mathbf{A}^t\cdot\mathbf{y} \leq 1, \mathbf{y} \in [0,1]^n\end{aligned}$$

The Duality Theorem from the theory of Linear Programming implies that

$$p(\mathcal{H}) \leq p_f(\mathcal{H}) = \kappa_f(\mathcal{H}) \leq \kappa(\mathcal{H}).$$

As an example consider the following win-lose game that is quite close to the type of games that are studied in this thesis and can be found in [47]. It is called the *hypergraph incidence game* and is played on a hypergraph $\mathcal{H} = (V, \mathcal{E})$. Alice chooses a vertex $v \in V$. Bob chooses an edge $E \in \mathcal{E}$. The payoff to Bob is 1, if $v \in E$. Otherwise the payoff is 0. This is a finite game and so, by von Neumann's theorem, it has a well defined value.

Theorem 1.1.1. *For any hypergraph, $\mathcal{H}$, the value of the hypergraph incidence game played on $\mathcal{H}$ is equal to $1/\kappa_f(\mathcal{H}) = 1/p_f(\mathcal{H})$.*

Proof. Recall that we assume that $\mathcal{H}$ has no exposed vertices. Let η be the value of this game and $\mathbf{A}$ the incidence matrix of $\mathcal{H}$. Let α be an optimal mixed strategy for Alice, β an optimal mixed strategy for Bob. Thus if Bob plays the mixed strategy β then, no matter what Alice does, he can guarantee an expected payoff of at least η. Thus $\mathbf{A}\cdot\beta \geq \eta$. Similarly, if Alice plays the mixed strategy α, we have $\mathbf{A}^t\cdot\alpha \leq \eta$. Set $\alpha' = \frac{1}{\eta}\alpha$ and

$\beta' = \frac{1}{\eta}\beta$. Then $\mathbf{A} \cdot \alpha' \leq 1$, $\mathbf{A}^t \cdot \beta' \geq 1$ and so α' and β' are feasible solutions of the above LPs. Hence

$$\kappa_f(\mathcal{H}) \leq \mathbf{1}^t \cdot \beta' = \frac{1}{\eta}\mathbf{1}^t \cdot \beta = \frac{1}{\eta},$$

since β is a probability vector, and

$$p_f(\mathcal{H}) \geq \mathbf{1}^t \cdot \alpha' = \frac{1}{\eta}\mathbf{1}^t \cdot \alpha = \frac{1}{\eta},$$

since α is a probability vector. The fact that $\kappa_f(\mathcal{H}) = p_f(\mathcal{H})$ finishes the proof. □

Part of this thesis is concerned with weighted versions of the hypergraph incidence game which we call *poisoning games*, for reasons that will become clear in the next section. A poisoning game is an instance of a win-lose game in which there are two players, say Alice and Bob, and a fixed ground space, X. To fix our ideas, let us assume that X is a hypergraph $\mathcal{H} = (V, \mathcal{E})$, where V is a finite set. The strategy space of Alice is the set of all functions $w : V \to [0, 1]$ such that $\sum_v w(v) \leq h$. Such a function will be called a *weighting* over $\mathcal{H}$. Bob chooses an edge $E \in \mathcal{E}$. Once an weighting, w, and an edge, E, have been chosen the players announce their choices and the payoff, $\langle w, E \rangle$, to Bob is 1, if $\sum_{v\in E} w(v) \geq 1$. Otherwise his payoff is zero.[1]
The value of the game is the probability that Bob wins under optimal play on both sides. Note that the strategy space of Alice is not a finite set which means that this is *not* a finite game. Since Bob's strategy space is finite the poisoning game is *semi-finite* and it is not immediately clear that it has a well defined value. However, as the following result shows, the game is equivalent to a finite game and so, by von Neumann's theorem, its value is well defined.

Lemma 1.1.2. *Any semi-finite win-lose game is equivalent to a finite win-lose game.*

Proof. Let A, B be the strategy spaces of the players and suppose that A is an infinite set while B is a finite set. Define an equivalence relation on A by setting $a_1 \sim a_2$ if and only if $f(a_1, b) = f(a_2, b)$, for all $b \in B$. Let A^* be the set of equivalence classes of A under this relation. We claim that A^* is a finite set. For the sake of contradiction, suppose that A^* is

[1] We remark that the roles of Alice and Bob are going to be interchanged in the poisoning games that we will consider in subsequent sections.

an infinite set and let $\{a_i\}_{i=1}^{\infty}$ be a set of representatives from a countable set of different equivalence classes. Fix some representative, say, a_1. As $a_1 \nsim a_i, i = 2, 3, \ldots$, it follows that there exist $b_i \in B, i = 2, 3, \ldots$, for which $f(a_1, b_i) \neq f(a_i, b_i)$. As $f(\cdot, \cdot)$ takes only the values 0 or 1, it follows $f(a_i, b_i) = f(a_j, b_j)$, for $i, j \geq 2$ and thus, since B is finite, there is an infinite set of indices, I_1, such that for $i, j \in I_1$, we have $f(a_i, b_i) = f(a_j, b_j)$ and $b_i = b_j := b_{\ell_1}$. Now fix some $i_1 \in I_1$ Since $a_{i_1} \nsim a_j$, for $j \in I_1$, it follows, similarly, that there is an infinite set of indices $I_2 \subseteq I_1$ such that for $i, j \in I_2$, we have $f(a_i, b_i) = f(a_j, b_j)$ and $b_i = b_j := b_{\ell_2} \neq b_{\ell_1}$. Note that for any $i, j \in I_2$ we have $f(a_i, b_{\ell_k}) = f(a_j, b_{\ell_k})$, for $k = 1, 2$. Continuing this way and since B is a finite set, we find that there is a countable set of indices, J, such that for all $i, j \in J$ we have $f(a_i, b) = f(a_j, b)$, for all $b \in B$, a contradiction. Hence A^* is finite, and we may start removing strategies from A without changing the problem at all, until we end up with a finite set of strategies, one for every equivalence class. □

Hence the poisoning game on $\mathcal{H}$ has a value and optimal strategies exist. Note that Alice is free to choose any weighting over V. We could also consider a poisoning game in which Alice has further restrictions on her weighting. As an example, suppose that Alice is only allowed to use *unit* weights. That is, she can only choose weightings for which every vertex gets either weight 1, or zero. Thus Alice chooses a weighting that gives weight 1 in $\lfloor h \rfloor$ vertices. Then, if $h < 2$, this restricted poisoning game is the same as the hypergraph incidence game. In case $h \geq 2$ the restricted game suggests the following generalization of the hypergraph incidence game that seems to be new.

Generalized hypergraph incidence game: Let $\mathcal{H} = (V, \mathcal{E})$ be a fixed hypergraph and $i \in \mathbb{Z}_{>0}$. Alice chooses i vertices $v_1, \ldots, v_i$. Bob chooses an edge $E \in \mathcal{E}$. The payoff to Bob is 1, if there exists $j \in \{1, \ldots, i\}$ such that $v_j \in E$. Otherwise his payoff is zero.

Note that the solution of the hypergraph incidence game is based on the Duality theorem of Linear Programming. Similarly, in order to find the value of the generalized hypergraph incidence game, we will formulate an appropriate Integer Program. Before doing so, we need some definitions.

Given a hypergraph, $\mathcal{H} = (V, \mathcal{E})$, we denote by $\binom{V}{i}$ the family containing all subsets of V of cardinality i. An *edge i-covering* of $\mathcal{H}$ is a collection of edges $\mathcal{E}_0 \subseteq \mathcal{E}$ such that for every $T \in \binom{V}{i}$, there exists $E \in \mathcal{E}_0$ for which

$T \cap E \neq \emptyset$. The smallest cardinality of an edge i-covering is called the *edge i-covering number* of $\mathcal{H}$ and is denoted by $\kappa(i, \mathcal{H})$. Note that $\kappa(1, \mathcal{H}) = \kappa(\mathcal{H})$, if $\mathcal{H}$ has no isolated vertices. Again, the problem of finding the edge i-covering number of a hypergraph is an IP problem. To see this denote by $\mathbf{A}_i$ the *i-incidence* matrix of $\mathcal{H}$, that is the $0,1$-matrix whose rows are represented by the sets $T_1, \ldots, T_l$ in $\binom{V}{i}$, the columns are represented by the edges $E_1, \ldots, E_m$ in $\mathcal{E}$ and whose elements, $a_{k,j}$, are equal to 1 if $T_k \cap E_j \neq \emptyset$, and equal to 0 otherwise. Let **x** be an indicator vector of the sets that have been selected for the edge i-covering and set $\mathcal{E}_x$ be the set of edges from $\mathcal{E}$ that correspond to this vector. Then **x** is the indicator vector of an edge i-covering if $\mathbf{A}_i \cdot \mathbf{x} \geq 1$ and so $\kappa(i, \mathcal{H})$ is the value of the IP

$$\begin{aligned} &\text{minimize } \mathbf{1}^t \cdot \mathbf{x}, \\ &\text{subject to: } \mathbf{A}_i \cdot \mathbf{x} \geq 1, \mathbf{x} \in \{0,1\}^m \end{aligned}$$

Again, this problem has a natural dual. A *vertex i-packing* of $\mathcal{H}$ is a subset $\mathcal{T} \subseteq \binom{V}{i}$ with the property that no two element in $\mathcal{T}$ intersect the same member of $\mathcal{E}$ or, in other words, every member of $\mathcal{E}$ intersects at most one element from $\mathcal{T}$. The *vertex i-packing number* of $\mathcal{H}$, denoted $p(i, \mathcal{H})$, is defined as the largest cardinality of a vertex i-packing. If **y** is an indicator vector of the sets in $\binom{V}{i}$ that are contained in the i-packing, then the IP formulation of the vertex i-packing number is

$$\begin{aligned} &\text{maximize } \mathbf{1}^t \cdot \mathbf{y}, \\ &\text{subject to: } \mathbf{A}_i^t \cdot \mathbf{y} \leq 1, \mathbf{y} \in \{0,1\}^l \end{aligned}$$

Hence $p(i, \mathcal{H}) \leq \kappa(i, \mathcal{H})$ and the linear relaxation of these two IPs gives rise to the fractional analogues of the edge i-covering number and vertex i-packing number, denoted $\kappa_f(i, \mathcal{H})$ and $p_f(i, \mathcal{H})$, respectively. The Duality theorem of Linear Programming then implies that

$$p_f(i, \mathcal{H}) = \kappa_f(i, \mathcal{H}).$$

In exactly the same way as in theorem 1.1.1 one can prove the following.

Theorem 1.1.3. *For any hypergraph, $\mathcal{H}$, the value of the generalized hypergraph incidence game on $\mathcal{H}$ is equal to $1/\kappa_f(i, \mathcal{H}) = 1/p_f(i, \mathcal{H})$.*

This thesis is motivated by poisoning games on hypergraphs. In particular, we will be concerned with poisoning games on the complete uniform hypergraph and the complete cyclic hypergraph. We begin the next section by considering the first case.

1.2 A conjecture of Kikuta and Ruckle

Suppose you want to poison your mother-in-law. She comes over for tea and takes s biscuits from a tray containing n in total. She has no preference and picks her biscuits randomly. You posses a bottle of arsenic containing h grams of it, where h is a real number, and the lethal dose is, say, 1 gram. You can distribute the poison any way you want over the biscuits. Unfortunately, you cannot put the poison in her tea, you have to put it in the biscuits. Which distribution has the highest probability of killing the old lady?

Conjecture 1.2.1 (Kikuta-Ruckle, 2000)**.** *It is optimal to use j equal positive dosages of $\frac{h}{j}$ grams and $n-j$ zero gram dosages, for some $j \leq n$ that depends on h, n, s.*

This problem is due to Kikuta and Ruckle (see [33]) who, driven by less devious motives, formulated it in terms of "accumulation games" between two players. It will be referred to as the *poisoning game* or poisoning problem.

The parameter h will remain fixed throughout this chapter and will always represent the amount of poison. Similarly, n is fixed and represents the total number of biscuits. Finally, s is fixed and represents the number of biscuits taken away. Notice that in case $h \geq \frac{n}{s}$ the problem is trivial. A dose of $\frac{1}{s}$ in each biscuit kills the mother-in-law for sure. Also, if $h < 1$ then the mother-in-law can never be poisoned. So, from now on, suppose that $1 \leq h < \frac{n}{s}$.

In [33] Kikuta and Ruckle consider the following win-lose game between a Hider, which from now on will be called Bob or the poisoner, and a

Seeker, henceforth called Alice or mother-in-law. Suppose that there is a fixed set of locations, $[n] := \{1, 2, \ldots, n\}$, and a given initial amount of poison $h, 1 \leq h < \frac{n}{s}$. Bob's strategy space, Σ, is a distribution of h over the locations. That is, Bob chooses a function $w : [n] \to [0, 1]$ such that $\sum w_i \leq h$ which gives rise to the vector $(w_1, \ldots, w_n)$. Such a function will be referred to as a *weighting* (or poisoning) on $[n]$ and w_i as the weight, or amount of poison, at location $i, i = 1, \ldots, n$. Thus the strategy space of the poisoner is the set of all weightings $\Sigma = \{(w_1, \ldots, w_n) : \sum w_i \leq h, 0 \leq w_i \leq h\} \subseteq \mathbb{R}^n$ and so Σ is a convex set.
The strategy space of Alice is the family of all subsets of $[n]$ of cardinality s, denoted $\binom{[n]}{s}$. Once a weighting w and a set $I \in \binom{[n]}{s}$ have been chosen, then Alice is poisoned (and Bob's payoff is 1) if $w(I) \geq 1$, where $w(I) = \sum_{i \in I} w_i$. If $w(I) < 1$ then Bob's payoff is 0, i.e. he loses the game and Alice wins. Any s-set, J, for which $w(J) \geq 1$ will be called a *heavy* (or *lethal*) set. Otherwise the set is *light*. The value of the game is the probability that Alice is poisoned under optimal play on both sides and is denoted by $V(h, n, s)$. Note that the strategy space of Bob is not a finite set which means that this is *not* a finite game. Since Alice's strategy space is finite the poisoning game is semi-finite and so, by Lemma 1.1.2, the game is equivalent to a finite game which has a well defined value.

Recall some definitions from the theory of finite sets. An s-*uniform* hypergraph is a pair $\mathcal{H} = (V, \mathcal{E})$, where V is a finite set of vertices and $\mathcal{E}$ is a family of subsets of V with s elements, called *edges*. Notice that for $s = 2$ a 2-uniform hypergraph is just a graph. Also note that the strategy space of Alice is the complete s-uniform hypergraph, i.e. the hypergraph consisting of all subsets of $[n]$ of cardinality s. Any weighting, w, on $[n]$ gives rise to a hypergraph, $\mathcal{H}_w$, whose vertex set is $[n]$ and whose edge set is $\mathcal{E}_w := \{I \subseteq [n] : |I| = s \;\&\; w(I) \geq 1\}$, the heavy s-sets under the weighting. Thus $|\mathcal{E}_w|$ is the number of lethal s-sets under w and the condition $h < \frac{n}{s}$ guarantees that there is always *at least one* light edge, i.e. $|\mathcal{E}_w| < \binom{n}{s}$, for any weighting w.

Suppose that A, B are strategy spaces of a finite win-lose game between Alice and Bob. The game is called *invariant* under a bijective map $g : A \to A$ if for every $b \in B$ there is a unique $b' \in B$ such that

$$f(a, b) = f(g(a), b'), \text{ for all } a \in A,$$

where $f(\cdot, \cdot)$ is the payoff function. It is known (see [17]) that the set of all bijections under which a game is invariant forms a group under the composition operation and that if a game is invariant under a group, G, of bijections on A then there exists an optimal mixed strategy for Alice that assign the same probability to elements $a_1, a_2 \in A$ for which $g(a_1) = a_2$, for some $g \in G$.

It will be shown that the poisoning game is invariant under the group of automorphisms of the underlying hypergraph. An *automorphism* of a hypergraph, $\mathcal{H} = (V, \mathcal{E})$ is a pair (π, σ) where π is a permutation of V, σ is a permutation of $\mathcal{E}$ such that for all $v \in V$ and all $E \in \mathcal{E}$ it holds $v \in E$ if and only if $\pi(v) \in \sigma(E)$. Recall that the set of automorphisms of a hypergraph forms a group under the operation of composition.

The following result says that even if your mother-in-law was informed about your intensions (though, still eager to eat your biscuits) the optimal way to play would be to pick an s-set uniformly at random.

Lemma 1.2.2. *It is optimal for Alice to choose an s-set at random.*

Proof. This is an invariance argument (see [17], Theorem 3.4). The game is invariant under the group of automorphisms, $\mathcal{A}$, of the hypergraph, $\mathcal{H} = (V, \binom{[n]}{s})$ of all subsets of $[n]$ of size s. To see this let $p(w, S)$ be Bob's payoff (so either 0 or 1) provided that he chooses the weighting w and Alice chooses the s-set S. Then for any automorphism $(\pi, \sigma) \in \mathcal{A}$ we have that $p(w, S) = p(\pi(w), \sigma(S))$, where $\pi(w) := (w_{\pi(1)}, \ldots, w_{\pi(n)})$. Since the game is invariant under the group $\mathcal{A}$, there exist invariant optimal strategies for the players. Since for any pair $(v_1, E_1), (v_2, E_2) \in (V, \binom{[n]}{s})$ there exists (π, σ) that maps (v_1, E_1) to (v_2, E_2), a mixed strategy for the mother-in-law is invariant if it assigns the same probability to all elements of $\binom{[n]}{s}$. $\square$

So we know the optimal strategy of Alice. Given any poisoning, w, on $[n]$, the probability that she is poisoned equals

$$\mathbb{P}_w := \frac{|\mathcal{E}_w|}{\binom{n}{s}}.$$

Hence the solution of the game comes down to the following optimization problem: Find a weighting w over $[n]$ such that the number of lethal s-sets is maximal.

Since we are interested in maximizing the number of lethal s-sets, we may assume that the weights are arranged in decreasing order, i.e. $w_1 \geq \cdots \geq w_n$. This means that the strategy space of Bob reduces to

$$\Sigma' = \left\{ (w_1, \dots, w_n) : \sum w_i \leq h, w_1 \geq \cdots \geq w_n \geq 0 \right\}.$$

The next result shows that the Kikuta-Ruckle conjecture says that an optimal strategy occurs at an extreme point of Σ'.

Lemma 1.2.3. *Any point of Σ' is a convex combination of the following vectors:*

$$(h, 0, \dots, 0), \left(\frac{h}{2}, \frac{h}{2}, 0, \dots, 0\right), \left(\frac{h}{3}, \frac{h}{3}, \frac{h}{3}, 0, \dots, 0\right) \dots, \left(\frac{h}{n}, \frac{h}{n}, \dots, \frac{h}{n}\right).$$

Proof. Denote by e_i the vector $(\frac{h}{i}, \frac{h}{i}, \dots, \frac{h}{i}, 0, \dots, 0), i = 1, \dots, n$. It is enough to show that every vector in the boundary of Σ' is a convex combination of the vectors $e_i, i = 1, \dots, n$. Suppose $w := (w_1, \dots, w_n) \in \partial\Sigma'$, so that $\sum_i w_i = h$ and $w_1 \geq w_2 \geq \cdots \geq w_n$. Let m be the maximum index for which $w_m > 0$ and set $x_i := w_i - w_{i+1} > 0$, for $i = 1, \dots, m-1$. Now the fact that $h = mw_m + (m-1)x_{m-1} + \cdots + 2x_2 + x_1$ and $w_i = w_m + x_{m-1} + \cdots + x_i, i = 1, \dots, m-1$ implies that

$$w = \frac{x_1}{h} e_1 + \frac{2x_2}{h} e_2 + \cdots + \frac{(m-1)x_{m-1}}{h} e_{m-1} + \frac{mw_m}{h} e_m,$$

which means that w is a convex combination of $e_i, i = 1, \dots, n$. □

A strategy, w, of Bob *dominates* some other strategy, w', if and only if $w'(I) \geq 1$ implies that $w(I) \geq 1$, for any s-set I. In other words, if a set is lethal under w' then it is also lethal under w. Some of the extreme points of Σ' are dominated. To see this note that every vector e_j for which $\frac{h}{j} \geq 1$ is dominated by the vector $v = (1, \dots, 1, 0 \dots, 0)$ consisting of j unit doses since any s-set that that is lethal under e_j is also lethal under v. Furthermore, the vectors e_j for which $\frac{h}{j} < 1$ are also dominated. To see this note that $\frac{j}{h} > 1$ and so there is a positive integer, k, such that $k-1 < \frac{j}{h} \leq k$, or equivalently $\frac{1}{k-1} > \frac{h}{j} \geq \frac{1}{k}$. This means that an s-set has to have at least k doses of $\frac{h}{j}$ in order to be lethal. But then Bob can just replace the weights $\frac{h}{j}$ in e_j by $\frac{1}{k}$ and achieve the same probability of winning. Hence if the Kikuta-Ruckle conjecture is true then the following statement is also correct.
The optimal distribution of poison over the biscuits uses dosages of $\frac{1}{j}$ in as many biscuits as possible, for a positive integer j that depends on h, n, s.

The conjecture of Kikuta and Ruckle arose from a series of papers on search games and optimal allocation over a number of years (see [31],[32] and [33]). Some properties of the value of the game along with a proof that the conjecture holds true in some particular cases can be found in [3].

We close this section by mentioning that in the poisoning game just defined the mother-in-law chooses an edge from the complete s-uniform hypergraph on n vertices. The choice of the mother-in-law depends on her eating habits. So one may consider poisoning games in which the mother-in-law picks an s-set from a fixed s-uniform hypergraph that is different from the complete. For example, the mother-in-law might arrange the n biscuits cyclically and choose s consecutive elements from that circle. This case is part of an entire chapter in this thesis.

1.3 A conjecture of Erdős

A family of subsets (or hypergraph), $\mathcal{H}$, of a finite vertex set is called *intersecting* if any two sets from $\mathcal{H}$ have non-empty intersection. The following result of Erdős, Ko and Rado is well known.

Theorem 1.3.1 (EKR). *Let $\mathcal{H}$ be an intersecting family of s-subsets of some vertex set containing n elements. If $n \geq 2s$ then $\mathcal{H}$ cannot have more than $\binom{n-1}{s-1}$ elements.*

See Appendix A for a proof. Hence an example of a maximal intersecting family is given by all s-sets containing some fixed element of the vertex set. In other words, EKR says that one cannot do better than the obvious solution.

Notice that this theorem settles the Kikuta-Ruckle conjecture for a certain range of parameters. If $h < 2$ and $n \geq 2s$, then any two lethal s-sets must have non-empty intersection and the EKR theorem implies that putting a unit weight is optimal.

There are several ways to generalize the EKR theorem. One of them is by putting constraints on the number of disjoint edges of the hypergraph. A set of pairwise disjoint edges in a hypergraph, $\mathcal{H}$, is called a *matching*. We denote by $\mu(\mathcal{H})$ the cardinality of the largest matching in $\mathcal{H}$ or, in short, the *matching number*. Notice that in an intersecting family of s-sets there are no disjoint edges and so its matching equals 1. Hence the EKR theorem says that the maximum cardinality of a uniform hypergraph with matching number 1 is at most $\binom{n-1}{s-1}$, i.e. the cardinality of a family of sets that contain 1 fixed element. So what if we take matching number a?

The problem of finding an s-uniform hypergraph, $\mathcal{H} = ([n], \mathcal{E})$, with the maximum number of edges under the constraint $\mu(\mathcal{H}) < a$ is well studied, though only partially solved, and goes back to a question of Paul Erdős (see [15]) that he raised in 1965. Erdős conjectured that the maximum is attained by two extremal hypergraphs. The first one is the hypergraph, $\mathcal{H}_{1/s}$, consisting of all the s-subsets on $sa-1$ vertices, whose matching number is clearly $a-1$. The second one is an s-uniform hypergraph, $\mathcal{H}_1$, on n vertices that includes all s-sets that contain at least one element from a fixed set of $a-1$ vertices and whose matching number is $a-1$ as well.

Conjecture 1.3.2 (Erdős, 1965)**.** *The number of edges in an s-uniform hypergraph, $\mathcal{H}$, on n vertices with matching number $\mu(\mathcal{H}) < a \leq \frac{n}{s}$ is at most*

$$\max\left\{\binom{sa-1}{s}, \binom{n}{s} - \binom{n-a+1}{s}\right\}.$$

In case $n \geq a(s+1)$, then $\binom{n}{s} - \binom{n-a+1}{s} > \binom{sa-1}{s}$. To see this use the inequality $\binom{n}{s} - \binom{n-a+1}{s} \geq (a-1)\binom{n-a+1}{s-1}$ and do some elementary calculation. Erdős proved the following result.

Theorem 1.3.3 (Erdős)**.** *There exists some constant c_s, that depends on s, such that the hypergraph $\mathcal{H}_1$ is maximal for all $n > c_s \cdot a$.*

The proof can be found in [15]. There has been considerable work on finding the constant c_s (see [23]). The best known lower bound on c_s is $2s - \frac{s}{a} - 1$ and was established very recently (see [22]).

Erdős' conjecture is related to the Kikuta-Ruckle conjecture. The family of lethal s-sets under a weighting, w, on $[n]$ forms a hypergraph $\mathcal{H}_w = ([n], \mathcal{E}_w)$ and we are interested in maximizing the number of its edges. The matching number of $\mathcal{H}_w$ is $< a := \lceil h \rceil$. If $n > c_s a$ and Bob puts dosage $\frac{1}{s}$ in as many biscuits as possible then we get the hypergraph $\mathcal{H}_{1/s}$. If Bob puts a dosage of 1 in as many biscuits as possible then we get the hypergraph $\mathcal{H}_1$. Kikuta and Ruckle include more fractional doses in their conjecture. It could be optimal to put dosages $1, 1/2, 1/3, \ldots, 1/s$. Notice that not all of these dosages are included in the conjecture of Erdős. The reason is that the conjecture of Erdős concerns an integer, a, while the conjecture of Kikuta and Ruckle concerns a fractional, h. If both conjectures are correct, then the optimal dosage is either 1 or $1/s$ when the amount of poison is an integer.

The conjecture of Erdős is an optimization problem in which the number of edges of a hypergraph, $\mathcal{H}$, needs to be maximized under a constraint on the matching number. This is a linear constraint. To see this, denote by $\mathbf{A}$ the *incidence* matrix of $\mathcal{H}$. That is, the matrix whose rows are represented by the vertices, $v_1, \ldots, v_n$, the columns are represented by the edges, $E_1, \ldots, E_m$, and whose elements, a_{ij}, are equal to 1 if $v_i \in E_j$ and equal to 0 otherwise. A matching in $\mathcal{H}$ is then a binary vector $\mathbf{x} = (x_1, \ldots, x_m)$ such that $\mathbf{A} \cdot \mathbf{x} \leq \mathbf{1}$, where $\mathbf{1}$ is the vector consisting of 1's only. The linear relaxation of the matching number problem gives rise to the following definition.

Let $\mathcal{H} = (V, \mathcal{E})$ be an s-uniform hypergraph. A *fractional matching* in $\mathcal{H}$ is a function $w : \mathcal{E} \to [0, 1]$ for which $\sum_{E \ni v} w(E) \leq 1$, for each $v \in V$. In other words, w is a weighting on the *edges*, instead of the vertices. The *fractional matching number* of $\mathcal{H}$, denoted $\mu^*(\mathcal{H})$, is defined as $\max_w \sum_{E \in \mathcal{E}} w(E)$. In other words, the fractional matching number is the maximum total weight on the edges.

The fractional version of Erdős' matching conjecture reads as follows.

Conjecture 1.3.4. *Fix positive integers n, s, a. Then the maximum number of edges in an s-uniform hypergraph, $\mathcal{H}$, on n vertices whose fractional matching number is an* integer *and satisfies $\mu^*(\mathcal{H}) < a \leq \frac{n}{s}$ is at most*

$$\max\left\{\binom{sa-1}{s}, \binom{n}{s} - \binom{n-a+1}{s}\right\}.$$

This conjecture was introduced only recently by Alon et al. (see [2]). Note that both the Kikuta-Ruckle and the fractional Erdős' matching conjecture address the following general problem.

Problem 1.3.5. *Fix positive integers n, s and a real number $h \geq 1$. Find a s-uniform hypergraph, $\mathcal{H} = (V, \mathcal{E})$, on n vertices whose fractional matching number is $\leq h$ and for which $|\mathcal{E}|$ is maximum.*

Using definitions analogous to those of the first section, one can generalize this problem. Fix a hypergraph $\mathcal{H} = (V, \mathcal{E})$ on n vertices. Denote by $\partial_i(\mathcal{E})$ the family containing all subsets of cardinality i, of some sets in $\mathcal{E}$. Formally,

$$\partial_i(\mathcal{E}) = \{T \subseteq V : |T| = i, T \subseteq E \text{ for some } E \in \mathcal{E}\}.$$

$\partial_i(\mathcal{E})$ is referred to as the i-th *shadow* of $\mathcal{E}$ in the literature. An *i-matching* of $\mathcal{H}$ is a is collection of edges $\mathcal{E}_0 \subseteq \mathcal{E}$ such that every $T \in \partial_i(\mathcal{E})$ is contained in *at most one* $E \in \mathcal{E}_0$. The largest cardinality of an i-matching of $\mathcal{H}$ is called the i-matching number of $\mathcal{H}$ and is denoted by $\mu_i(\mathcal{H})$. It is not difficult to formulate an Integer Program corresponding to the problem of finding the i-matching number of a hypergraph. We leave the details to the reader. The linear relaxation of this IP gives rise to the fractional analogue of the i-matching number. We can generalize the last problem by introducing the following one.

Problem 1.3.6. *Fix positive integers n, s and a real number $h \geq 1$. Find a s-uniform hypergraph, $\mathcal{H} = (V, \mathcal{E})$, on n vertices whose fractional i-matching number is $\leq h$ and for which $|\mathcal{E}|$ is maximum.*

1.4 A conjecture of Samuels

The conjecture of Kikuta and Ruckle entails two problems. The first one is the conjecture itself which, as we already saw, reduces to an optimization problem. The second is the problem of determining the optimal dose, given the validity of the conjecture. This is a probability problem. In order to illustrate this suppose that the amount of poison, h, is an *integer*. Recall that in this case Erdős' conjecture, if true, implies that the optimal dose is either equal to 1 or equal to $1/s$. Let us assume further that unit doses are better than doses of $1/s$. So, if the Kikuta-Ruckle conjecture is true, the optimal distribution of poison over the biscuits uses $h \cdot j$ poisonous biscuits that contain $\frac{1}{j}$ grams of arsenic and $n - j$ zero gram biscuits, for some $j < s$. This means that the amount of poison taken by the mother-in-law is equal to $\frac{1}{j} \cdot H_j$, for a hypergeometric random variable, H_j, that counts the number of poisonous biscuits taken away when we sample s biscuits from a set containing n in total and the number of poisonous biscuits is $h \cdot j$. This means that the probability that Alice is poisoned equals $\mathbb{P}[H_j \geq j]$ and so Bob faces the following (tail) probability problem.

$$\textbf{maximize } \mathbb{P}[H_j \geq j] \textbf{ where } 1 \leq j \leq s - 1.$$

Recall that we require the parameters h, n, s to satisfy $\frac{h \cdot s}{n} < 1$, which gives that $\mathbb{E}[H_j] = \frac{hjs}{n} < j$. Finding the maximum of the hypergeometric tails turns out to be difficult. As a first case, we find the maximum tail under a stronger restriction on $\mathbb{E}[H_j]$. Namely, we assume that $\frac{h \cdot s}{n} \leq \frac{1}{s-1}$. Then, for every $j = 1, \ldots, s-1$, we have that $\mathbb{E}[H_j] = \frac{h \cdot j \cdot s}{n} \leq \frac{j}{s-1} \leq 1$ We now prove that in this case the optimal j equals 1. This will require the following elementary result.

Lemma 1.4.1. *Let Z be any random variable taking non-negative integer values for which $\mathbb{E}[Z] \leq 1$. Then*

$$\mathbb{P}[Z = 0] \geq \mathbb{P}[Z \geq 2] + \mathbb{P}[Z \geq 3] + \cdots \geq \mathbb{P}[Z \geq 2].$$

Proof. Notice that

$$\mathbb{P}[Z = 0] + \mathbb{P}[Z \geq 1] = 1 \geq \mathbb{E}[Z] \geq \mathbb{P}[Z \geq 1] + \mathbb{P}[Z \geq 2] + \cdots .$$

□

Some evidence that the maximum tail equals $\mathbb{P}[H_1 \geq 1]$ is given by the next result.

Lemma 1.4.2. *Let H be a hypergeometrically distributed random variable with parameters n, s, a, i.e. $H \sim$ Hyp(n, s, a). Suppose that H is such that $\mathbb{E}[H] \leq 1$. Then*

$$\mathbb{P}[H = 1] \geq \mathbb{P}[H > 1].$$

Proof. An equivalent form of the inequality is

$$\mathbb{P}[H = 0] + 2\mathbb{P}[H = 1] \geq 1.$$

For fixed n, s, a write

$$\begin{aligned} \mathbb{P}[H = 0] + 2\mathbb{P}[H = 1] &= \frac{\binom{n-a}{s} + 2a\binom{n-a}{s-1}}{\binom{n}{s}} \\ &= \frac{\binom{n-a}{s}}{\binom{n}{s}} \cdot \left(1 + \frac{2sa}{n - a - s + 1}\right). \end{aligned}$$

Straightforward calculation shows that the last quantity is increasing for $n \leq n_0 := 2na - a - n$ and decreasing for $n \geq n_0$. Since in the limit as $n \to \infty$ we get that $\mathbb{P}[H = 0] = 1$, it is enough to show that the inequality holds true in case $n = sa$, or, equivalently $\mathbb{E}[H] = 1$. For this case we compute

$$\frac{\mathbb{P}[H = 1]}{\mathbb{P}[H = 0]} = \frac{sa}{(s-1)(a-1)} > 1.$$

Thus, when $n = sa$, the inequality $\mathbb{P}[H = 1] \geq \mathbb{P}[H \geq 2]$ follows by Lemma 1.4.1. □

Theorem 1.4.3. *Let $m := \frac{h}{n}$. If m is less than $\frac{1}{s(s-1)}$ then the solution of the probability problem is $j = 1$. That is,*

$$\mathbb{P}[H_j \geq 1] \leq \mathbb{P}[H_1 \geq 1], \textit{ for } j = 1, \ldots, s - 1.$$

Proof. Since $\mathbb{E}[H_j] = \frac{h \cdot j \cdot s}{n} \leq 1$ for all $j = 1, \ldots, s-1$ it follows that $\mathbb{P}[H_j = 0] \geq \mathbb{P}[H_j \geq j]$, by Lemma 1.4.1. For $j = 1, \ldots, s-1$ it is immediate that $\mathbb{P}[H_j = 0] \geq \mathbb{P}[H_{j+1} = 0]$. Thus $\mathbb{P}[H_1 = 0] \geq \mathbb{P}[H_j = 0] \geq \mathbb{P}[H_j \geq j]$. Since H_1 is hypergeometrically distributed of mean less than one, Lemma 1.4.2 gives that

$$\mathbb{P}[H_1 \geq 1] \geq \mathbb{P}[H_1 = 1] \geq \mathbb{P}[H_1 = 0] \geq \mathbb{P}[H_j \geq j],$$

as required. □

In words, if the Kikuta-Ruckle conjecture is true and the average number of poisonous biscuits in the sample of Alice is less than 1, then it is optimal to use *unit* weights. A fact that is intuitively obvious.

Denote by $X_i, i = 1, \ldots, s$, the amount of poison in the i-th biscuit of Alice. Then $\mathbb{E}[X_i] = \frac{h}{n}$ and the total amount of poison that Alice eats equals $\Sigma_s = X_1 + \cdots + X_s$. This is a *dependent* sum of random variables and so the poisoning game addresses the following problem.

Fix s real numbers $m_1, \ldots, m_s$ such that $0 \leq m_1 \leq \cdots \leq m_s$ and $\sum m_i < 1$ and find

$$\Xi(m_1, \ldots, m_2) := \sup \mathbb{P}[X_1 + \cdots + X_s \geq 1],$$

where the supremum is over all s-tuples of (dependent) random variables $X_1, \ldots, X_s$ with means $m_1, \ldots, m_s$, respectively.

This problem has been studied since Hoeffding for sums of *independent* random variables. In the 60's Samuels published a number of papers in which he considered the following question. Let $X_i, i = 1, \ldots, s$ be independent random variables of mean m_i such that $\sum m_i < 1$. What is the maximum value of $\mathbb{P}[X_1 + \cdots + X_s \geq 1]$?

Samuels only partially solved this problem, which by now has remained open for more than fifty years, but he did conjecture a full solution: order the means in increasing order, $m_1 \leq \cdots \leq m_s$. Then the random variables, $X_1, \ldots, X_s$, that maximize the tail have the following property. There exists a $t \in \{0, 1, \ldots, s-1\}$ such that X_i is constant and equal to m_i for $i \leq t$ and for $i > t$ each X_i is a $0/1$-valued random variable of mean m_i. More formally, Samuels' problem reads as follows.
Fix s real numbers, $m_1, \ldots, m_s$ such that $0 \leq m_1 \leq \cdots \leq m_s$ as well as $\sum_{i=1}^{s} m_1 < 1$, and denote

$$\Psi(m_1, \ldots, m_s) := \inf \mathbb{P}[X_1 + \cdots + X_s < 1] = \sup \mathbb{P}[X_1 + \cdots + X_s \geq 1],$$

where the infimum and supremum is over all s-tuples of non-negative independent random variables $X_1, \ldots, X_s$ with means $m_1, \ldots, m_s$, respectively. Now for each $t = 0, 1, \ldots, s-1$ set

$$Q_t(m_1, \ldots, m_s) := \prod_{i=t+1}^{s} \left(1 - \frac{m_i}{1 - \sum_{j=1}^{t} m_j}\right).$$

Now suppose that $X_i = m_i$ with probability 1, for $i \le t$ and for $i > t$ each X_i takes the values $1 - \sum_{j=1}^{t} m_j$ with probability $\frac{m_i}{1-\sum_{j=1}^{t} m_j}$ and 0 otherwise, i.e., its mean equals m_i. Then $X_1 + \cdots + X_s < 1$ if and only if $X_{t+1} + \cdots + X_s < 1 - \sum_{j=1}^{t} m_i$ and the last inequality is satisfied in case all X_i are equal to 0, for $i \ge t+1$. Hence for this choice of random variables we have

$$\mathbb{P}[X_1 + \cdots + X_s < 1] = Q_t(m_1, \ldots, m_s).$$

Conjecture 1.4.4 (Samuels, 1966)**.** *Suppose that X_i are non-negative, independent random variables of mean $m_i, i = 1, \ldots, s$. Then the tail probability $\mathbb{P}[X_1 + \cdots + X_s \ge 1]$ is maximized by random variables X_i such that each of them is either constant and equal to m_i, or takes only the values 0 and 1. Formally, for all real numbers, $m_1, \ldots, m_s$ satisfying $0 \le m_1 \le \cdots \le m_s$, and $\sum_{i=1}^{s} m_1 < 1$,*

$$\Psi(m_1, \ldots, m_s) = \min_{t=0,\ldots,s-1} Q_t(m_1, \ldots, m_s).$$

In [46] Samuels obtained, as a corollary of a more general theorem, the following result.

Theorem 1.4.5. *Let $X_1, \ldots, X_s$ be i.i.d. with common mean, m. If m is less than $\frac{1}{\max\{4s, s(s-1)\}}$, then the tail probability $\mathbb{P}[X_1 + \cdots + X_s \ge 1]$ is maximized by i.i.d. Bernoulli random variables X_i with common mean m. That is,*

$$\mathbb{P}[X_1 + \cdots + X_s \ge 1] \le 1 - (1-m)^s.$$

Compare this result with Theorem 1.4.3. Notice also that this result confirms the conjecture in the case of i.i.d. random variables X_i with common mean. This case has an interpretation in terms of a poisoning game. Consider a poisoning game in which Bob has made s identical trays each containing n biscuits. In each tray he has distributed h grams of arsenic in exactly the same way. Alice will take one biscuit from each tray, so s in total. Let's refer to this game as the poisoning game *with replacement*. If X_i is the amount of poison taken by Alice from the i-th tray, then each

$X_i, i = 1, \ldots, s$, is a random variable of mean $m := \frac{h}{n} < \frac{1}{s}$. If Samuels conjecture holds true, then finding the optimal distribution of poison requires to determine the minimum value of $Q_t(m)$, for $t = 0, \ldots, s-1$. Note that the expression for $Q_t(m)$ simplifies a lot. For each $t = 0, \ldots, s-1$ we have

$$Q_t(m) = \left(1 - \frac{m}{1 - tm}\right)^{s-t}.$$

For $t = 0, \ldots, s-1$ define Z_t to be a binomially distributed random variable of parameters $s - t$ and $\frac{m}{1-tm}$. In short, $Z_t \sim \text{Bin}(s - t, \frac{m}{1-tm})$. Then finding the minimum $Q_t(m)$ is a (tail) probability problem.

$$\textbf{minimize } \mathbb{P}[Z_t = 0], \ \textbf{where } t = 0, 1, \ldots, s-1.$$

The following result is proven in [2], using the arithmetic-geometric means inequality.

Lemma 1.4.6. *If $m \leq \frac{1}{s+1}$ then, for all $t = 0, \ldots, s-1$, we have*

$$\mathbb{P}[Z_0 = 0] \leq \mathbb{P}[Z_t = 0].$$

Proof. We want to prove that

$$(1 - m)^s \leq \left(1 - \frac{m}{1 - tm}\right)^{s-t} = \left(\frac{1 - tm - m}{1 - tm}\right)^{s-t},$$

for $t = 1, \ldots, s-1$. Note that the inequality is true when m equals 0 or $\frac{1}{s+1}$. For the intermediate values, $0 < m < \frac{1}{s+1}$, we prove instead that

$$f(m) := s \cdot \log(1 - m) - (s - t)\log(1 - tm - m) + (s - t)\log(1 - tm) \leq 0.$$

Now notice that

$$f'(m) = \frac{t(1 - tm)(sm - 1 + m)}{(1 - m)(1 - tm - m)(1 - tm)} < 0,$$

for $0 < m < \frac{1}{s+1}$. This means that $f(\cdot)$ is decreasing for $0 < m < \frac{1}{s+1}$. □

In other words, if Samuels' conjecture is true and the parameters h, n, s satisfy $\frac{h}{n} \leq \frac{1}{s+1}$, then the optimal weighting in the poisoning game with replacement uses *unit* weights.

1.5 Monotonicity and the binomial game

In this section we discuss asymptotic approaches of the poisoning game. Throughout this section it is *assumed that the Kikuta-Ruckle conjecture holds true*. The probability problem that is associated to the conjecture asks for the optimal dose $1/j$. Which j should the poisoner choose? There are three parameters, h, n, s. Suppose we keep two of them fixed and vary the third. If h increases then this is to the advantage of Bob. His resources improve and, intuitively, it may be better to spread the poison. If s increases then, again, this is to the advantage of the Bob since it is more likely that Alice will get a lethal dose and, intuitively, it might be as well better to spread the poison. However, if n increases then this is to the advantage of Alice whose probability of getting a lethal dose decreases. This suggests that the following monotonicity might be true.

$$j \text{ increases with } h,\ j \text{ increases with } s,\ j \text{ decreases with } n.$$

This is a statement about hypergeometric random variables. Suppose that H_j is a random variable of sampling s times without replacement from a tray containing n biscuits in total and with $\lfloor h \rfloor j$ of the biscuits being poisoned. Then the integer j that maximizes $\mathbb{P}[H_j \geq j]$ increases with h and s, but decreases with n. This is a technical statement that is not easy to handle. In order to simplify the problem a bit suppose we vary two parameters and keep the third fixed. We now discuss two versions of this approach.

Suppose first that we let h and n go to infinity while keeping $\frac{h}{n}$ fixed and equal to μ and keeping s fixed as well. As n is getting larger, and in order to simplify matters, we may assume that the dependence between different samplings of Alice vanishes. Let $X_i, i = 1, \ldots, s$ be the amount of poison in the i-th choice of Alice. So, in the limit, Bob is facing the following problem. Find i.i.d. random variables $X_i, i = 1, \ldots, s$ of mean μ such that the tail probability $\mathbb{P}[X_1 + \cdots + X_s \geq 1]$ is maximal. Note that this is a special case of Samuels' problem.

We may also consider another asymptotic approach by letting n and s go to infinity while keeping $\frac{s}{n}$ fixed and equal to μ and keeping h fixed. Since the n and s go to infinity, and in order to simplify matters, we might again suppose that there is no dependence between different samples of Alice. That is, we might suppose that Alice chooses her biscuits with replacement and so, in the limit, the players can be thought of as participating in the following win-lose game on the interval $[0, 1]$. Alice chooses a subset, S, of $[0, 1]$ of Lebesgue measure μ. Bob puts poison on the interval and wins if the amount of poison in S is ≥ 1. That is, Bob chooses a measure γ on $[0, 1]$ such that $\gamma([0, 1]) = h$, where h is a real number greater than or equal to 1. Bob wins if S is lethal under γ, i.e. if $\gamma(S) \geq 1$.

If $\mu h \geq 1$ then Bob wins for sure by choosing a uniform measure. So we can assume from now on that $\mu h < 1$. This is not a finite game. The strategy space of both players is infinite and it is not at all obvious that the game has a well defined value. The next results imply that the value is well defined for a certain range of parameters.

Lemma 1.5.1. *If $\mu \leq \frac{1}{2}$ and $h < 2$ then the optimal strategy for Bob is to choose a point uniformly randomly on $[0, 1]$ and put a dosage of 1 gram in this point. The optimal strategy of Alice is to identify the points 0 and 1, thus turning the interval into a circle, and choose an interval from the circle of length μ uniformly at random.*

Proof. Clearly, the suggested strategy of Bob guarantees that he wins with probability μ, against any pure strategy of Alice.
Now fix any pure strategy, γ, of Bob. Alice picks a point, x, uniformly at random from the circle and chooses the set $S_x := [x, x + \mu), \mathrm{mod}1$. Alice is poisoned if $\gamma(S_x) \geq 1$. We show that the probability that Alice is poisoned is $\leq \mu$. That is, we need to prove that $\lambda(\{x : \gamma(S_x) \geq 1\}) \leq \mu$, where λ denotes Lebesgue measure. Since $h < 2$, any two lethal intervals S_x must have non-empty intersection. Take any interval S_x that is lethal and note that any interval that intersects S_x is one of the intervals S_t, $t \in (x - \mu, x + \mu)\ \mathrm{mod}1$. Now if S_t is lethal, then $S_{t+\mu}$ cannot be lethal since it has non-empty intersection with S_t. Thus at most half of the interval that intersect S_x can be lethal which means that $\lambda(\{t : \gamma(S_t) \geq 1\} \cap (x-\mu, x+\mu)) \leq \mu$. □

Observe that the previous lemma is related to EKR theorem. Its proof is a modified version of Katona's proof of the EKR theorem (see [29]).

Lemma 1.5.2. *If $\mu = \frac{m-1}{m}$ then the optimal strategy of Bob is to put a unit dosage at a randomly chosen point of $[0, 1]$. The optimal strategy of Alice is to divide $[0, 1]$ into m equal subintervals, $I_1, \ldots, I_m$, where $I_j = \left[\frac{j-1}{m}, \frac{j}{m}\right]$ and choose $[0, 1] \setminus I_j$ uniformly randomly.*

Proof. The probability that Bob wins is $1 - \frac{1}{m}$, for any pure strategy of Alice.

Now fix a pure strategy of Bob, γ. Alice chooses one of the sets $A_i := [0,1] \setminus I_i, i = 1, \ldots, m$ with equal probability. Now note that

$$\sum_{i=1}^{m} \gamma(A_i) = (m-1)h = m\mu h < m,$$

where the first equality follows from the fact that each subinterval, I_i, has been counted $m-1$ times. Hence there is an index i_0, for which $\gamma(A_{i_0}) < 1$. In case Alice chooses this set, then she survives and this happens with probability $\frac{1}{m}$. $\square$

Notice that for $m = 2$ the last two Lemmata show that there are two different optimal strategies for Alice.

The Kikuta-Ruckle conjecture predicts that in the asymptotic game Bob divides the poison into doses $\frac{1}{j}$ and so the probability that a "random" subset S of $[0,1]$ contains at least j of these doses is $\mathbb{P}[B_j \geq j]$, for a binomially distributed random variable $B_j \sim \text{Bin}(\lfloor jh \rfloor, \mu)$. Hence the problem of finding the optimal j reduces to the problem of maximizing the tails $\mathbb{P}[B_j \geq j]$ for $j \geq 1$. Such problems have been around for a long time. To illustrate this, suppose that h is a positive *integer*. Then the optimal dose is determined through the following probability problem.

maximize $\mathbb{P}[B_j \geq j]$ **where** $j \in \mathbb{Z}_{>0}$, **and** $B_j \sim \text{Bin}(jh, \mu)$.

Suppose further that we were allowed to choose $\mu = \frac{1}{h}$. Recall that in fact $\mu < \frac{1}{h}$. Then we would have considered a well known problem.
In 1693 Samuel Pepys wrote a letter to Isaac Newton (see chapter 12 of [13]) asking which of the following events is more likely to happen:

- throw a dice six times and gamble on at least one 6,
- throw a dice twelve times and gamble on at least two 6's, or
- throw a dice eighteen times an gamble on at least three 6's.

The answer, which is to take six dies and gamble on at least one 6, is contained in the next result.

Theorem 1.5.3. *Fix a positive integer h and let $Z_j \sim Bin(jh, \frac{1}{h})$. Then, for any $j \in \mathbb{N}$,*

$$\mathbb{P}[Z_j \geq j] \geq \mathbb{P}[Z_{j+1} \geq j+1].$$

Proof. First note that $\mathbb{E}[Z_j] = j \in \mathbb{N}$, for all j, in which case it is easy to see that the mean of the binomial distribution equals the mode.
I.e. $\mathbb{P}[Z_j = i] \leq \mathbb{P}[Z_j = j]$ for all i. If we regard Z_{j+1} as the independent sum of Z_j and Z_1 and compute $\mathbb{P}[Z_j + Z_1 \geq j+1]$ conditional on Z_j, we get

$$\begin{aligned} \mathbb{P}[Z_{j+1} \geq j+1] &= \sum_{i=0}^{\infty} \mathbb{P}[Z_1 \geq j+1-i] \cdot \mathbb{P}[Z_j = i] \\ &= \mathbb{P}[Z_1 \geq j+1] \cdot \mathbb{P}[Z_j = 0] + \mathbb{P}[Z_1 \geq j] \cdot \mathbb{P}[Z_j = 1] + \cdots \\ &+ \mathbb{P}[Z_1 \geq 1] \cdot \mathbb{P}[Z_j = j] + \mathbb{P}[Z_j \geq j+1]. \end{aligned}$$

Hence it is enough to show that

$$\mathbb{P}[Z_1 \geq j+1] \cdot \mathbb{P}[Z_j = 0] + \cdots + \mathbb{P}[Z_1 \geq 1] \cdot \mathbb{P}[Z_j = j] \leq \mathbb{P}[Z_j = j],$$

or, equivalently, that

$$\mathbb{P}[Z_1 \geq j+1] \cdot \mathbb{P}[Z_j = 0] + \cdots + \mathbb{P}[Z_1 \geq 2] \cdot \mathbb{P}[Z_j = j-1] \leq \mathbb{P}[Z_1 = 0] \cdot \mathbb{P}[Z_j = j].$$

Since $\mathbb{P}[Z_j = i] \leq \mathbb{P}[Z_j = j]$ for all i we get that

$$\mathbb{P}[Z_1 \geq j+1] \cdot \mathbb{P}[Z_j = 0] + \cdots + \mathbb{P}[Z_1 \geq 2] \cdot \mathbb{P}[Z_j = j-1]$$

is less than or equal to

$$\mathbb{P}[Z_j = j] \cdot (\mathbb{P}[Z_1 \geq j+1] + \cdots + \mathbb{P}[Z_1 \geq 2]).$$

Now that fact that $\mathbb{E}[Z_1] \leq 1$ and Lemma 1.4.1 gives that the last quantity is

$$\leq \mathbb{P}[Z_j = j] \cdot \mathbb{P}[Z_1 = 0]$$

and finishes the proof. □

Hence it is optimal to use unit weights. Note that, since h is an integer, this coincides with the optimal dose that is suggested by Erdős' conjecture.
The same dose is optimal in case $\mu < \frac{1}{h}$, as we now show. For $x \in (0,1)$ and $j \in \mathbb{N}$ define the function

$$f_j(x) := \mathbb{P}[Z_{j+1} \leq j] - \mathbb{P}[Z_j \leq j-1],$$

where $Z_j \sim \text{Bin}(hj, x)$ It is clear that $f_j(0) = f_j(1) = 0$ and from Theorem 1.5.3 we have that $f_j(\frac{1}{h}) \geq 0$. Also it is easy to verify that $f_j'(x)$ equals

$$x^{j-1}(1-x)^{hj-j} \left[-\binom{hj+h}{j} x(hj+h-j)(1-x)^{h-1} + \binom{hj}{j-1}(hj-j+1) \right]$$

Theorem 1.5.4. *For any $h \in \mathbb{N}$ and any $0 \le x \le 1/h$, we have*

$$\mathbb{P}[Z_j \ge j] \ge \mathbb{P}[Z_{j+1} \ge j+1], \text{ for all } j = 1, 2, \dots,$$

where $Z_j \sim Bin(jh, x)$.

Proof. Fix some $j \in \{1, 2, \dots\}$. Define $g_j(x) := -\binom{hj+h}{j} x(hj+h-j)(1-x)^{h-1}$. Then

$$g_j'(x) = -\binom{hj+h}{j}(hj+h-j)(1-x)^{h-2}(-1+hx).$$

which gives that $g_j(\cdot)$ is decreasing for $x < \frac{1}{h}$ and increasing for $x > \frac{1}{h}$. Also $g_j(0) = g_j(1) = 0$. If $G_j(x) := g_j(x) + \binom{hj}{j-1}(hj-j+1) \ge 0$ for all x, then $f_j(\cdot)$ is an increasing function which contradicts the fact that $f_j(0) = f_j(1) = 0$. Thus, by monotonicity of $g_j(\cdot)$, there exist $x_0 < 1/h$ and $x_1 > 1/h$ for which $G_j(x_0) = G_j(x_1) = 0$. Thus $G_j(\cdot)$ is positive for $x < x_0$, it is negative for $x_0 \le x \le x_1$ and positive again for $x > x_1$. This gives that for $x < x_0$ and $x > x_1$ the function $f_j(\cdot)$ is increasing. But as $f_j(0) = f_j(1) = 0$, it follows that $f_j(x_0) > 0$ and $f_j(x_1) < 0$. From Theorem 1.5.3 we know that $f_j(1/h) \ge 0$. Since $f_j'(\cdot)$ changes sign at the points x_0 and x_1 and $f_j(\cdot)$ is decreasing for $x \in [x_0, x_1]$, we conclude that $f_j(x) \ge 0$, for $x \in [0, 1/h]$. □

So in case h is an integer and the Kikuta-Ruckle conjecture holds true for the asymptotic game, then it is optimal to use unit weights. The monotonicity of tail probabilities of the binomial distribution is well studied. For a general result see [27].

1.6 A geometric poisoning-problem

In this section we consider a poisoning "game" between Alice and Bob in which Bob has made a pie that has the shape of a circular disk. Alice eats a circular piece of the pie of radius, say, 1 that she will choose at random. Bob has $h \geq 1$ grams of arsenic and the lethal dose is 1 gram.
Denote by Π^* the pie, which is a closed disk of radius $R > 1$. Suppose that this disk is centered at the origin, O, of a plane. Denote by $D(p, r)$ the disk with center p and radius $r > 0$. Alice eats a circular piece of the pie of radius 1 and this piece is chosen uniformly at random. More precisely, she chooses a point, p, uniformly at random from the disk $\Pi := D(O, R - 1)$ and eats the closed disk $D(p, 1)$. Which distribution of poison over the pie has the highest probability of doing the job?

Note that this is not a game. At least not the way it is defined. Alice is given no choice. She is restricted in choosing her piece uniformly at random and so this is an optimization problem.

Every distribution of poison over the pie gives rise to a measure μ defined on Π such that $\mu(\Pi) = h$. This means that Bob's strategy space is the set of all measures on Π of total mass h, denoted $\mathcal{M}_h$. For every choice of Bob, $\mu \in \mathcal{M}_h$, define

$$\alpha(\mu) := \lambda_2(p \in \Pi : \mu(D(p, 1)) \geq 1),$$

where λ_2 denotes the two-dimensional Lebesgue measure. Hence the problem of finding the optimal distribution of poison reduces to the problem of finding a $\mu_* \in \mathcal{M}_h$ for which $\alpha(\mu_*) \geq \alpha(\mu)$, for every $\mu \in \mathcal{M}_h$.
This is a difficult problem that addresses non-trivial geometric questions. In this section we focus on these geometric questions. In order to illustrate this, suppose first that $1 \leq h < 2$. Then, for every $\mu \in \mathcal{M}_h$, any two disks

D_1, D_2, that are lethal under μ must have non-empty intersection. Define the *size*, $\sigma(\mathcal{F})$, of a family $\mathcal{F}$ of unit disk in the plane to be the Lebesgue measure of the set consisting of the centers of the disks $D \in \mathcal{F}$. Formally,

$$\sigma(\mathcal{F}) = \lambda_2(\{p : D(p,1) \in \mathcal{F}\}).$$

What is the maximum size of a family of pairwise intersecting unit disks in the plane?
Notice that two unit disks have non-empty intersection if and only if their centers are at distance ≤ 2. So one can rephrase the last question as follows. What is the maximum measure of a set of points in the plane for which any two are at distance ≤ 2? Now, any two points of a set are at distance ≤ 2 if and only if the diameter of the set is ≤ 2. So the question is equivalent to the following well known problem whose solution implies that disks maximize area under constrains on the diameter.

Isodiametric problem: Among all plane sets of diameter $\leq \Delta$, find one that has maximal Lebesgue measure.

The answer, which is that a disk of radius $\frac{\Delta}{2}$ has maximal area, can be obtained via the Brunn-Minkowski inequality.

Theorem 1.6.1 (Brunn-Minkowski)**.** *Let A and B be non-empty compact sets in $\mathbb{R}^2$. Then*

$$\lambda_2(A+B)^{1/2} \geq \lambda_2(A)^{1/2} + \lambda_2(B)^{1/2},$$

where λ_2 denotes 2-dimensional Lebesgue measure and $A + B = \{a + b : a \in A, b \in B\}$.

Proof. See [38]. □

Using this theorem one can prove the so-called isodiametric (or Bieberbach) inequality.

Theorem 1.6.2 (Isodiametric inequality)**.** *If A is a subset of the plane of diameter $\leq \Delta$, then*

$$\lambda_2(A) \leq \pi \left(\frac{\Delta}{2}\right)^2.$$

Proof. The trick is to look at the set $A - A := \{a_1 - a_2 : a_1, a_2 \in A\}$. From Brunn-Minkowski inequality we have

$$4\lambda_2(A) \leq \lambda_2(A - A)$$

and so any upper bound on $\lambda_2(A - A)$ gives an upper bound on $\lambda_2(A)$. Now the fact that the diameter of A is $\leq \Delta$ implies that $A - A \subseteq D(0, \Delta)$ and so

$$4\lambda_2(A) \leq \lambda_2(A - A) \leq \pi\Delta^2,$$

as required. $\square$

Let $\mathcal{F}$ be an intersecting family of unit disks in the plane. Let $\sigma(\mathcal{F})$ be its size. The isodiametric inequality implies that $\sigma(\mathcal{F}) \leq \pi$. That is, the maximum size is achieved by the family of all disks that contain a specific point.

Theorem 1.6.3. *Suppose that $1 \leq h < 2$. Then it is optimal to put 1 gram of poison at the center, O, of Π. That is, the optimal measure is a discrete one that concentrates mass 1 at a point of Π.*

Proof. For any distribution of poison over the pie we have that the family of lethal disks is intersecting. Suppose first that $R \geq 2$. Then the Lebesgue measure of Π is $\geq \pi$. This means that a unit disk is contained in Π. The theorem follows since the area of such a disk is the maximum possible size of an intersecting family of unit disks.
If $R < 2$, then every disk of radius 1 in Π contains O and the proof is complete. $\square$

Notice the similarity with the case of the poisoning game that uses biscuits. If $1 \leq h < 2$ and $n \geq 2s$, then the family of lethal s-sets is an intersecting family and the Erdős-Ko-Rado theorem gives that it is optimal to use a unit gram dose.

More generally, if $a - 1 \leq h < a$ then, given any $\mu \in \mathcal{M}_h$, one cannot find a disjoint unit disks that are lethal under μ and we want to maximize the size of such a family of unit disks. This problem is now much more complicated. In order to illustrate this, suppose now that $2 \leq h < 3$. Call a family of unit disks in the plane $\frac{3}{2}$-*intersecting* if for any 3 disks from the family, at least two have non-empty intersection. Since $2 \leq h < 3$ we have that for any distribution of poison over the pie, μ, the family of lethal disks under μ is $\frac{3}{2}$-intersecting. Hence, if we consider the set consisting of the centers of the lethal disks under μ, the case $2 \leq h < 3$ addresses the following geometry problem.

A generalized isodiametric problem: Fix $\Delta > 0$ and suppose that A is a plane set of diameter Δ for which among any three points, at least two are at distance ≤ 2. What is the maximum measure of A?

Notice that in case $\Delta \leq 2$ the problem reduces to the isodiametric one and so the maximal measure is achieved by a disk of radius Δ. The same is true in case $2 < \Delta \leq \frac{4}{\sqrt{3}}$. To see this, first note that the smallest disk that contains an equilateral triangle of side length 2 is one of radius $\frac{2}{\sqrt{3}}$. Furthermore, the inclusion $A - A \subseteq D(0, \Delta)$ is always true and the Brunn-Minkowski inequality implies

$$4\lambda_2(A) \leq \lambda_2(A) \leq \pi\Delta^2.$$

Thus $\lambda_2(A) \leq \frac{\Delta^2}{4}$ and this bound it achieved by a disk of radius $\Delta/2$.

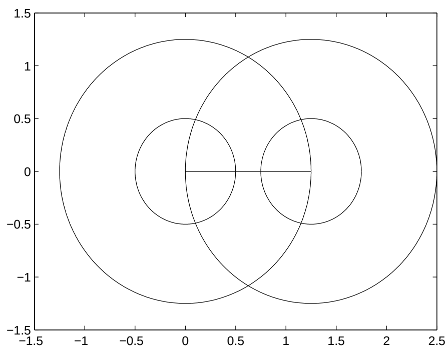

Lemma 1.6.4. *Suppose that $R > 2$. Then is is optimal to put unit gram doses in two points of Π that are at distance 4 apart.*

Proof. We prove that the maximum size of a $\frac{3}{2}$-intersecting family, $\mathcal{F}$ of disks in the plane for which the maximum distance of their centers is > 4 is at most 2π. Notice that this value is attained by the size of the family of disks whose centers form two disjoint unit disks. Let A be the set consisting of the centers of the disks in $\mathcal{F}$. Choose two points $p_1, p_2 \in C$, which are at maximal distance, say $r > 4$. Now all other points of C must belong to $D := D(p_1, r) \cap D(p_2, r)$, by maximality of the distance between p_1 and p_2. Since $r > 4$, it follows that the two disks $D(p_1, 2)$ and $D(p_2, 2)$ are disjoint. Set $A_1 := D(p_1, r) \cap D(p_2, 2)$ and $A_2 := D(p_2, r) \cap D(p_1, 2)$ and note that $A_1 \cap A_2 = \emptyset$ (see the figure above). Also notice that no point of C belongs to $D \setminus (A_1 \cup A_2)$, since any point of this set forms with the points p_1 and p_2 a triangle whose sides have length > 2. Hence all points of C belong to either A_1 or A_2. Now notice that no two points of A_1 (resp. of A_2) can be at distance > 2 because in that case they would form with p_2 (resp. with p_1)

a triangle with side lengths > 2. Thus any two points in A_1 and any two points in A_2 are at distance ≤ 2. This means that $\lambda_2(A_1), \lambda_2(A_2) \leq \pi$, from the isodiametric inequality, and so $\lambda_2(A) = \lambda_2(A_1) + \lambda_2(A_2) \leq 2\pi$. □

We close this section with a geometric problem that arises from the above and is interesting on its own.

Problem 1.6.5. *A set $A \subseteq \mathbb{R}^d$ of (fixed) diameter Δ is called $(n/k, \delta)$*-intersecting *if every n-tuple of points in A contains a k-tuple of diameter δ. What is the maximum Lebesgue measure of A?*

1.7 List of publications and manuscripts

The research for this thesis has led to the following publications and manuscripts.

1. S. Alpern, R. J. Fokkink and C. Pelekis, *A proof of the Kikuta-Ruckle conjecture on cyclic caching of resources*, J. Optim. Theory Appl., 153, p. 650–661, (2012).

2. R. J. Fokkink, J. op den Kelder and C. Pelekis, *How to poison your mother-in-law and other caching problems*, Chapter 10 in Search Theory: A Game-Theoretic Prespective, Alpern et al. (ed.), Springer, (2013).

3. C. Pelekis and M. Schauer, *Network coloring and colored coin games*, Chapter 4 in Search Theory: A Game-Theoretic Prespective, Alpern et al. (ed.), Springer, (2013).

4. C. Pelekis, *Hypergraphs with the König-Kikuta-Ruckle property*, (in preparation).

5. C. Pelekis, *Poison trials of fixed parity, random and randomly oriented graphs*, 23 pages. (submitted to Combinatorics, Probability and Computing).

Chapter 2

The Kikuta-Ruckle conjecture

In this Chapter we give further evidence for the validity of the Kikuta-Ruckle conjecture. Some instances of the conjecture have been settled in [33] and [3]. In this chapter we will settle a few more cases. In [33] it is proven that the conjecture is true when $s = 1$ or $s = n - 1$. In [3] the conjecture has been verified for the case $s = 2$ or $s = n - 2$ and for the case $n \leq 7$. In this chapter we will settle the case $n = 2s - 1$ as well a few more instances.

2.1 The conjecture for the odd graph

In this section we prove that the Kikuta-Ruckle conjecture holds true in case $n = 2s - 1$. Throughout this chapter n will always represent the total number of biscuits in the tray. Similarly, s will represent the number of biscuits taken by the mother-in-law. We begin with an example.

Example 2.1.1. *Suppose that* $\frac{3}{2} \leq h < \frac{5}{3}, n = 5$ *and the mother-in-law takes* 3 *biscuits at random. Number the biscuits from* 1 *to* 5*. Three of these biscuits are taken away. Equivalently, two of the biscuits won't be taken. Consider all* $\binom{5}{2}$ *such pairs of biscuits as vertices of a graph and identify each pair with it's complement. Put an edge whenever two pairs are disjoint. The result is the Petersen graph.*

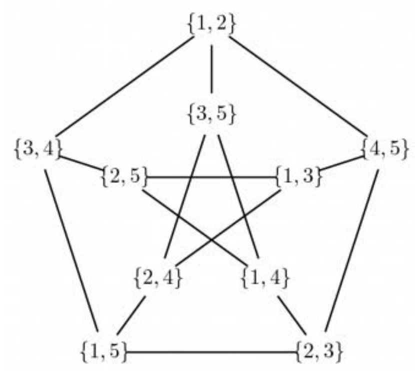

Now, any cycle of length 5 *contains each biscuit exactly twice. If we look at the complementary sets we get that each cycle of length* 5 *corresponds to five triples for which each biscuit has been counted three times and thus the amount of poison in those five triples equals* $3h < 5$*. So any cycle of length* 5 *contains at least one vertex that corresponds to a doubleton whose complement is not lethal. Notice that for any pair of vertices, there is a* 5*-cycle that avoids them. This implies that there are at least* 3 *vertices whose complement is non-lethal. Now it is not hard to verify that the distribution* $\{\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, 0, 0\}$ *creates exactly* 3 *vertices with non-lethal complement and so is optimal.*

The odd graph $\mathcal{O}_s$ has one vertex for each of the s-element subsets of a $(2s-1)$-element set. Two vertices are connected by an edge if and only if the corresponding subsets have one common element.[1] The Petersen graph is equal to the odd graph for $s = 3$. Norman Biggs [8] already remarked that if one wants to understand a graph theory problem, the odd graph is a good place to start. So we consider the Kikuta-Ruckle conjecture for the values of s and n that correspond to odd graphs: in this section, it is our standing assumption that $\mathbf{n} = \mathbf{2s} - \mathbf{1}$.

Lemma 2.1.2. *Suppose that the Kikuta-Ruckle conjecture is correct when* $n = 2s - 1$*. Then it is optimal to put a dose of* $1/j$ *if and only if* $h \in [2 - \frac{1}{j}, 2 - \frac{1}{j+1})$.

Proof. We have to determine the optimal dose $\frac{1}{j}$, depending on h. Let N_j be the number of lethal s-subsets if we put a dose of $\frac{1}{j}$ and $h \geq 2 - \frac{1}{j}$. We claim that $N_1 < \cdots < N_\ell$. Assume that we put a dose $\frac{1}{j}$ in the first $2j - 1$ biscuits and denote this strategy by S_j. To compare N_j to N_{j+1} we

[1]The original definition of the odd graph takes $(k-1)$-element subsets as its vertices. They are connected by an edge if and only if they are disjoint. So for each edge there is one element that is not contained in the two vertices: the odd one out. This is where the graph gets its name from. Our definition is equivalent and more convenient for the poisoning problem. An edge represents the odd one in.

need to consider the effect of reducing the amount of poison in the first $2j-1$ biscuits from $\frac{1}{j}$ to $\frac{1}{j+1}$, while putting a dose of $\frac{1}{j+1}$ in the next two biscuits that previously did not have any poison. A lethal subset under S_j becomes non-lethal under S_{j+1} if it contains j elements from $\{1,\ldots,2j-1\}$ and none from $\{2j, 2j+1\}$. There are exactly

$$\binom{2j-1}{j}\binom{2s-2j-2}{s-j}$$

such subsets. Conversely, a non-lethal subset under S_j becomes lethal under S_{j+1} if it contains $j-1$ elements from $\{1,\ldots,2j-1\}$ and both $2j$ and $2j+1$. There are exactly

$$\binom{2j-1}{j-1}\binom{2s-2j-2}{s-j-1}$$

such subsets. Dividing the first binomial product by the second gives $\frac{s-j-1}{s-j} < 1$, so the number of s-subsets that become lethal exceeds those that become non-lethal. Which proves that $N_j < N_{j+1}$.
If we put a dose of $1/j$ while $h < 2 - \frac{1}{j}$, then there are at most $2j-2$ poisonous biscuits. Let M_j be the number of lethal k-subsets in this case. We claim that $M_1 < \cdots < M_s$ is again an increasing sequence. To compare M_j to M_{j+1} we need to consider the effect of reducing the amount of poison in the first $2j-2$ biscuits, while putting a dose of $1/(j+1)$ in biscuit $2j-1$ and $2j$. The number of lethal subsets that become non-lethal now is

$$\binom{2j-2}{j}\binom{2s-2j-1}{s-j}$$

while the number of subsets that become lethal is

$$\binom{2j-2}{j-1}\binom{2s-2j-1}{s-j-1}.$$

and, once again, the quotient of these two binomial products is $\frac{j-1}{j} < 1$, so the number of subsets that become lethal upon redistribution again exceeds the number of those that become non-lethal. Now we claim that $M_s < N_1$, so it is better to put a single unit dose. Indeed $M_s = \binom{2s-2}{s}$ while $N_1 = \binom{2s-2}{s-1}$. So putting a dose $1/j$ for $j > 1$ is only optimal once $h \geq 2 - \frac{1}{j}$. $\square$

We have an amount of poison h that we distribute over the biscuits, putting a dose w_i in the i-th biscuit. An s-subset V is lethal if and only if $w(V) =$

$\sum_{i\in V} w_i \geq 1$. We number the biscuits in decreasing order of their doses, putting the most poisonous biscuit first, i.e., $w_1 \geq \cdots \geq w_{2s-1}$. Let $\mathcal{P}$ be the family of poisonous s-subsets. We want to distribute the poison in such a way that $\mathcal{P}$ has maximum cardinality. We adopt hypergraph terminology. We say that $V \in \mathcal{P}$ is an edge, and $\deg_{\mathcal{P}}(i)$ is equal to the number of edges that contains i.

Lemma 2.1.3. *If $h < 2 - \frac{1}{j+1}$ then $\deg_{\mathcal{P}}(2j+1) \leq \frac{1}{2}\binom{2s-2}{s-1}$.*

Proof. By the decreasing dosage of poison

$$(2j+1)w_{2j+1} \leq w_1 + \cdots + w_{2j+1} \leq h < \frac{2j+1}{j+1},$$

and so $h + w_{2j+1} < 2$. If V is any s-subset that contains $2j+1$ then let $\bar{V} = V^c \cup \{2j+1\}$. In other words, $\bar{V}$ is the neighbor of V in the odd graph $\mathcal{O}_s$ that is connected by the edge that has $2j+1$ as the odd one in. Then $w(V) + w(\bar{V}) = h + w_{2j+1} < 2$. So if V is poisonous then $\bar{V}$ is not, and we conclude that $\deg_{\mathcal{P}}(2j+1)$ is at most half of the degree of $2j+1$ in the complete hypergraph on all k subsets. The degree of the complete hypergraph is $\binom{2s-2}{s-1}$. □

Lemma 2.1.4. *If $h < 2 - \frac{1}{j+1}$ then the number of lethal edges is at most*

$$\frac{1}{2}\binom{2j}{j}\binom{2s-2j-1}{s-j} + \sum_{i=j+1}^{s}\binom{2j}{i}\binom{2s-2j-1}{s-i}.$$

Proof. We maximize the number of edges V under the constraint that the hypergraph has maximal $\sum_{i\geq 2j+1}\deg(i)$, which by the previous lemma is bounded by

$$\frac{n-2j}{2}\binom{2s-2}{s-1}.$$

The greedy solution is to first take all s-subsets that have no elements in $\{2j+1,\ldots,2s-1\}$, then to take all s-subsets that have one element in $\{2j+1,\ldots,2s-1\}$, etc, until the sum of the degrees exceed the given bound. We need to show that this happens exactly when we have taken all s-subsets that contain $> j$ elements from $\{1,\ldots,2j\}$ and half of the s-subsets that contain exactly j elements from this set. In other words, we need to show that

$$\frac{1}{2}\binom{2j}{j}\binom{2s-2j-1}{s-j}(s-j) + \sum_{i=j+1}^{\min\{2j,s-1\}}\binom{2j}{i}\binom{2s-2j-1}{s-i}(s-i)$$

is equal to $\frac{n-2j}{2}\binom{2s-2}{s-1}$. This can be rewritten to

$$\frac{\binom{2j}{j}\binom{2s-2j-2}{s-j-1}}{2\binom{2s-2}{s-1}} + \sum_{i=j+1}^{\min\{2j,s-1\}} \frac{\binom{2j}{i}\binom{2s-2j-2}{s-i-1}}{\binom{2s-2}{s-1}} = \frac{1}{2}.$$

Let X be a hypergeometric random variable that describes the number of successes in $s-1$ draws from a population of $N = 2s-2$ with $2j$ successes. Then this equation is equal to

$$\frac{1}{2}\mathbb{P}[X = j] + \mathbb{P}[X > j] = \frac{1}{2}.$$

To see why this last equation is true, notice that drawing $s-1$ from $2s-2$ is equivalent to leaving $s-1$ from $2s-2$. Since the number of successes is $2j$ this means that for every drawing for which $X > j$ there is a unique drawing for which $X < j$, which implies that $\mathbb{P}[X > j] = \mathbb{P}[X < j]$ and finishes the proof. □

Theorem 2.1.5. *The Kikuta-Ruckle conjecture is true for odd graphs, i.e., if* $n = 2s-1$.

Proof. If we put $2j-1$ doses of $1/j$ then an edge is lethal if and only if it contains at least j out of the first $2j-1$ biscuits. So the number of lethal edges is equal to

$$\sum_{i=j}^{s}\binom{2j-1}{i}\binom{2s-2j}{s-i}.$$

By the previous lemma, it suffices to show that this is equal to

$$\frac{1}{2}\binom{2j}{j}\binom{2s-2j-1}{s-j} + \sum_{i=j+1}^{s-1}\binom{2j}{i}\binom{2s-2j-1}{s-i}.$$

If we divide both sums by $\binom{2s-1}{s}$ then the first quotient is $\mathbf{P}[X_1 \geq j]$ for a hypergeometric random variable that counts the number of successes if we draw s times with $2j-1$ successes. The second quotient is $\frac{1}{2}\mathbf{P}[X_2 = j] + \mathbf{P}[X_2 \geq j+1]$ if the number of successes is $2j$. To see why these probabilities are the same, start with the population that has $2j$ successes and call one of them a failure, which transforms X_2 into X_1. Let U be the event that the draw does not contain the success which turns into a failure. Then $X_1 \geq j$ is equal to

$$(X_2 \geq j+1) \,\cup\, \{\mathbf{U} \cap X_2 = j\}.$$

Now observe that

$$\mathbb{P}[\mathbf{U} \cap X_2 = j] = \mathbb{P}[\mathbf{U} \mid X_2 = j]\mathbb{P}[X_2 = j] = \frac{1}{2}\mathbb{P}[X_2 = j].$$

□

Note that Lemma 2.1.2 works whenever $n = 2s-k, k = 1, 2, \ldots, s$. We state the result and leave the details to the reader.

Partition the interval $I := [1, \frac{n}{s})$ by writing it as a union of intervals $[1, \frac{3}{2}) \cup [\frac{3}{2}, \frac{5}{3}) \cup \cdots \cup [\frac{2\ell-3}{\ell-1}, \frac{2\ell-1}{\ell}) \cup I_1$, for the maximum possible value of ℓ. I_1 is interval that remains. Note that $\ell \leq m$, for the integer m for which $n = m \cdot k + \upsilon,\ 1 \leq \upsilon \leq k-1$.

Lemma 2.1.6. *Suppose that the Kikuta-Ruckle conjecture is true for $n = 2s - k$. Then it is optimal to put a dose of $\frac{1}{j}$ if and only if $h \in [2 - \frac{1}{j}, 2 - \frac{1}{j+1})$.*

2.2 The conjecture for a few more cases

In this section we give further evidence for the validity of the Kikuta-Ruckle conjecture by proving that it holds true in some additional cases. Denote by $V(n, s, h)$ the value of the poisoning game on the complete uniform hypergraph.

Lemma 2.2.1. *If $n = 0$ mod s, then $V(n, s, h) \leq \frac{s\lfloor h \rfloor}{n}$.*

Proof. Suppose that $n = d \cdot s$. Let $\mathcal{P}$ be the set of all partitions of $[n] := \{1, \ldots, n\}$ into d subsets of cardinality s. We first need to know the proportion of partitions $P \in \mathcal{P}$ that contain a fixed s-set. By symmetry, each of the $\binom{n}{s}$ choices for s-set, H, belongs to the same number, say a, of partitions in $\mathcal{P}$. Also, each member of $\mathcal{P}$ contains d s-sets. Thus $\binom{n}{s}a = |\mathcal{P}|d$ and so $a = |\mathcal{P}|/\binom{n-1}{s-1}$.
Let $\mathcal{H}_w$ be the number of lethal s-sets under a weighting, w, over $[n]$. Count pairs (H, P), where $H \in \mathcal{H}_w$ and $H \in P \in \mathcal{P}$. Each of the $|\mathcal{H}_w|$ choices for H belongs to a choices for P. Since the amount of poison that is available to Bob is h, each partition $P \in \mathcal{P}$ has at most $\lfloor h \rfloor$ lethal s-sets. Thus

$$a \cdot |\mathcal{H}_w| \leq \lfloor h \rfloor \cdot |\mathcal{P}|$$

and so $|\mathcal{H}_w| \leq \lfloor h \rfloor \cdot \binom{n-1}{s-1}$, for any weighting w. Hence $V(n, s, h) \leq \lfloor h \rfloor \frac{s}{n}$, as required. □

Corollary 2.2.2. *If $n = 0 \bmod s$ and $h < 2$, then it is optimal for the poisoner to put a unit weight*

Proof. The result follows from the previous lemma and the fact that $\lfloor h \rfloor = 1$. □

Corollary 2.2.3. *Suppose that $n = ks, k = 1, 2, \ldots$ and $k - \frac{1}{s} \leq h < k$. Then it is optimal for the poisoner to put weight $\frac{1}{s}$ in $ks - 1$ vertices.*

Proof. If the poisoner uses this strategy then the probability that he wins equals $1 - \binom{ks-1}{s-1}/\binom{ks}{s} = 1 - \frac{1}{k}$. By the previous lemma the value of the game is $\leq (k-1)\frac{s}{sk} = 1 - \frac{1}{k}$. Hence the suggested strategy is optimal. □

We can settle some more cases using game-theoretic arguments.

Lemma 2.2.4. *Suppose that $n = \frac{k}{k-1}s$, where $k \geq 2$. Then it is optimal for the poisoner to use a single unit weight.*

Proof. Note that $h < \frac{k}{k-1}$. If the poisoner plays this strategy, then the probability that he wins is $1 - \binom{n-1}{s}/\binom{n}{s} = \frac{k-1}{k}$, no matter how the mother-in-law plays. Now we look the game from the mother-in-law's point of view. Since $s = \frac{k-1}{k}n = n - \frac{n}{k}$, it follows that the set $\{1, \ldots, n\}$ can be partitioned into k sets of cardinality n/k. Let $\mathcal{P}$ be the set of all such partitions of $\{1, \ldots, n\}$. She selects $k-1$ of these sets at random. Now for every member of $\mathcal{P}$ there are at least $k-1$ parts that are non-lethal. This follows from the fact that $h < 2$. Hence the probability that the mother-in-law is not poisoned is at least $\frac{1}{k}$, no matter how the poisoner plays. □

Lemma 2.2.5. *The Kikuta-Ruckle conjecture holds true when $n \geq 2s + 1$ and $h < 2 + \frac{1}{s}$.*

Proof. If $h < 2$ then the Erdős-Ko-Rado theorem implies that it is optimal to put a unit gram dosage. So suppose that $2 \leq h < 2 + \frac{1}{s}$. We may assume that the doses of an optimal distribution, w, of poison are ordered $w_1 \geq \cdots \geq w_n$, so that $w_1 \geq 1/s$. If the poisoner puts two unit dosages then he creates

$$\binom{n}{s} - \binom{n-2}{s}$$

lethal s-sets. Let $\mathcal{F}_w$ be the family of lethal s-sets under w, and set

$$\mathcal{F}_1 := \{F \in \mathcal{F}_w : 1 \notin F\}.$$

Since $w_1 \geq 1/s$, it follows that the family $\mathcal{F}_1$ is intersecting and so, by the Erdős-Ko-Rado theorem, we have $|\mathcal{F}_1| \leq \binom{n-2}{s-1}$. Now, the fact that $|\mathcal{F} \setminus \mathcal{F}_1| \leq \binom{n-1}{s-1}$ implies

$$|\mathcal{F}| = |\mathcal{F} \setminus \mathcal{F}_1| + |\mathcal{F}_1| \leq \binom{n-1}{s-1} + \binom{n-2}{s-1} = \binom{n}{s} - \binom{n-2}{s},$$

as required. □

We can also settle a case using the result from the Odd graph.

Lemma 2.2.6. *Suppose that both n and s are divided by k and $n = 2s - k$. Assume further that $h \in [1, 3/2)$. Then it is optimal for the poisoner to use a unit weight.*

Proof. Let $n = k \cdot n_1$ and $s = k \cdot s_1$, so that $n_1 = 2s_1 - 1$. Denote by $\Gamma(a, b)$ the game on a vertices in which b of them are taken away by Alice. Suppose that Alice adopts the following strategy. She partitions the set $\{1, \ldots, n\}$ into n_1 sets of cardinality k and chooses s_1 of them uniformly at random. This strategy of Alice forces the poisoner to play in the game $\Gamma(n_1, s_1)$ and so, under optimal play, the probability, $P_{n,s}$, that Alice is poisoned in $\Gamma(n, s)$ is at most the probability that she is poisoned in the game $\Gamma(n_1, s_1)$. Thus $P_{n,s} \leq \mathbb{P}[X = 1]$, where $X \sim \text{Hyp}(n_1, s_1, 1)$. Now look the game $\Gamma(n, s)$ from the poisoner's point of view. Suppose that he puts a unit dosage at a vertex that is chosen uniformly at random. Then the probability that Alice is poisoned equals $\mathbb{P}[Y = 1]$, where $Y \sim \text{Bin}(1, s/n)$, and so $P_{n,s} \geq \mathbb{P}[Y = 1]$. The result follows form the fact that $\mathbb{P}[X = 1] = \mathbb{P}[Y = 1]$. □

The following result is taken from [3] and its proof is included for the sake of completeness.

Lemma 2.2.7. *The Kikuta-Ruckle conjecture holds true when $h < 2 + \frac{1}{s-1}$ and $n \geq 2s$.*

Proof. If $h < 2$, then the collection of lethal s-sets forms an intersecting family and so, by the Erdős-Ko-Rado theorem it is optimal for Bob to put one unit weight. So suppose that $2 \leq h < 2 + \frac{1}{s-1}$ and arrange the weights of an optimal distribution of poison in decreasing order, $w_1 \geq \cdots \geq w_n$. Suppose first that $w_1 < \frac{1}{s-1}$. Then all weights are smaller than $\frac{1}{s-1}$ and so every lethal edge must contain s weights that are $\geq \frac{1}{s}$. Hence the poisoner may as well use weight $\frac{1}{s}$, and the conjecture holds true.
Assume now that $w_1 \geq \frac{1}{s-1}$. Thus $h - w_1 < 2$ which implies that the family

of lethal s-sets that do not contain the vertex, v, with the biggest weight is an intersecting family, which we denote by $\mathcal{F}$. By the Erdős-Ko-Rado theorem, $|\mathcal{F}| \leq \binom{n-2}{s-1}$. There are $\binom{n-1}{s-1}$ s-sets that contain vertex v, and so the number of lethal s-sets is

$$\leq \binom{n-2}{s-1} + \binom{n-1}{s-1} = \binom{n}{s} - \binom{n-2}{s},$$

which is exactly the number of lethal edges if the poisoner puts two unit weights. □

In [3] one can find a proof that the Kikuta-Ruckle conjecture holds true when $n \leq 7$. We end this section by verifying that the conjecture holds also true in case $n = 9$. Hence $s \in \{1, \dots, 9\}$ and $h < \frac{9}{s}$. It is proven in [33] that the conjecture is true when $s = 1$, or $s = n - 1$, for any n. In [3] it is proven that the conjecture is also true when $s = 2$, or $s = n - 2$, for any n. We also know that the conjecture holds true for the odd graph. Hence, if $n = 9$, we already know that the conjecture holds true for $s \in \{1, 2, 5, 7, 8, 9\}$ and so we only have to consider the case $s \in \{3, 4, 6\}$.
The case $s = 6$ follows from lemma 2.2.4 and lemma 2.2.7 settles the case $s = 4$. We are only left with the case $s = 3$.
If $s = 3$ and $h < 2$ then, by the Erdős-Ko-Rado theorem, the conjecture holds true and one unit weight is optimal. The case $8/3 \leq h < 3$ follows from corollary 2.2.3 and Bob uses eight $\frac{1}{3}$'s. The case $2 \leq h < 5/2$ follows form lemma 2.2.7 and Bob uses two unit weights. Thus we only have to check the conjecture for the following case.

Suppose that $n = 9, s = 3$ and $5/2 \leq h < 8/3$. We prove that it is optimal for the poisoner to use five $\frac{1}{2}$'s. In this case the probability that the poisoner wins equals $25/42$ and there are 34 non-lethal 3-sets. We prove that, in this case, there is no weighting over $\{1, \dots, 9\}$ creating less than 34 non-lethal 3-sets. So suppose that this is not true and so there exists a weighting w, of total weight $h \in [5/2, 8/3)$, having ≤ 33 non-lethal 3-sets. We may assume that $w_1 \geq \dots \geq w_9$. Now in this case, the 3-sets $\{1, 6, 7\}, \{2, 4, 8\}, \{3, 4, 5\}$ are all lethal. To see this, suppose that it is not true and count the number of 3-sets that are smaller than these sets in the lexicographic order. For example, if $\{1, 6, 7\}$ is non-lethal then any 3-set that contains exactly one element from $\{1, 2, 3, 4, 5\}$ and any two elements from the set $\{6, 7, 8, 9\}$ is also non-lethal as well as any 3-set that is contained in set $\{6, 7, 8, 9\}$ is also non-lethal. But there are $\binom{5}{1}\binom{4}{2} + 4 = 34$ such 3-sets, which contradicts our assumption that the number of non-lethal 3-sets is ≤ 33. Similarly the other cases.

Now now suppose that $\{3,5,6\}$ is lethal. To finish the proof, we show that the value of the following optimization problem is $8/3$.

$$\begin{aligned} \text{minimize } & w_1 + \cdots + w_9, \\ \text{subject to } & w_1 + w_6 + w_7 \geq 1 \\ & w_2 + w_4 + w_8 \geq 1 \\ & w_3 + w_5 + w_6 \geq 1 \\ & w_1 \geq \cdots \geq w_9 \geq 0 \end{aligned}$$

Let w, i.e. $w_1 \geq \cdots \geq w_9$, be a feasible solution of the above optimization problem. Set $\Sigma(w) := \sum_{i=1}^{9} w_i$. By adding the three inequality constraints we get that $\Sigma(w) \geq 3 - w_6 + w_9$. This shows that, for an optimal w, we have $w_9 = 0$ since otherwise we could replace it by 0 and get a feasible solution having less weight. Now we claim that an optimal weighting w is such that $w_6 \geq 1/3$. To see this, suppose that $w_6 < 1/3$. Then $\Sigma(w) \geq 3 - w_6 > 8/3$. But the weighting $x_1 = \cdots x_6 = 1/3,\ x_7 = x_8 = x_9 = 0$ is a feasible solution with total sum $8/3$, a contradiction to the optimality of w. Thus we might also suppose that $w_6 \geq 1/3$ holds true for an optimal weighting w. We now claim that $w_6 = 1/3$. Suppose not, so that there exists an optimal weighting w for which $w_6 = 1/3 + \epsilon$ for a positive ϵ. Then, by monotonicity of the weights, we get $w_3, w_5 \geq 1/3 + \epsilon$ and thus $\Sigma(w) \geq 3 + 3\epsilon - w_6 = 8/3 + 2\epsilon > 8/3$. Thus ϵ should be zero and the value of the above optimization problem is $8/3$.

A similar analysis works if $\{3,5,6\}$ is non-lethal. Note that in this case the are 30 sets that are lexicographically smaller than or equal to $\{3,5,6\}$, thus also non-lethal. We claim that in this case $\{2,6,8\}$ is lethal. If not, then all 3-sets that contain 2, one element from $\{6,7\}$ and one element from $\{8,9\}$ is also non-lethal. This adds at least four more lethal 3-sets to the existing set of 30 non-lethal 3-sets.

$$\begin{aligned} \text{minimize } & w_1 + \cdots + w_9, \\ \text{subject to } & w_1 + w_6 + w_7 \geq 1 \\ & w_2 + w_6 + w_8 \geq 1 \\ & w_3 + w_4 + w_5 \geq 1 \\ & w_1 \geq \cdots \geq w_9 \end{aligned}$$

The value of this LP is again $8/3$.

We also mention another instance in which the conjecture holds true and can be obtained using results from the next chapter. It uses Theorem 3.5.7 and so can be skipped on a first reading.

Corollary 2.2.8. *Consider the poisoning game on the complete s-uniform hypergraph on n vertices with amount of poison h. Suppose that the parameters h, n, s are such that the fraction of smallest denominator in $\left[\frac{s}{n}, \frac{1}{h}\right)$ equals $\frac{j-1}{j}$, for some $j \in \{2, 3, \ldots\}$. Then it is optimal for the poisoner to use a single unit weight.*

Proof. Note that restriction on the parameters implies that $h < 1 + \frac{1}{j-1}$. Suppose that the poisoner picks an element from $\{1, \ldots, n\}$ uniformly at random and puts a unit weight. Then his winning probability is at least $\frac{s}{n}$, no matter how the mother-in-law plays. Now the mother-in-law has to come up with a strategy to compensate on this probability. Suppose that she plays as follows. She arranges the vertices cyclically and picks one path of length s from this cycle, uniformly at random. This strategy of the mother-in-law forces the poisoner to play the poisoning game on the cyclic graph $G(n, s)$. Now, by Theorem 3.5.7, the value of the cyclic game $G(n, s)$ equals $1 + q\frac{s}{n} - p$, where $\frac{p}{q}$ is the Farey successor in $\mathcal{F}_j$ of the fraction, $\frac{j-1}{j}$, of smallest denominator in $\left[\frac{s}{n}, \frac{1}{h}\right)$. The Farey successor of $\frac{j-1}{j}$ is $\frac{p}{q} = \frac{1}{1}$. It follows that, under optimal play, the mother-in-law is poisoned with probability $\leq \frac{s}{n}$, as required. $\square$

Chapter 3

The Kikuta-Ruckle conjecture on circular hypergraphs

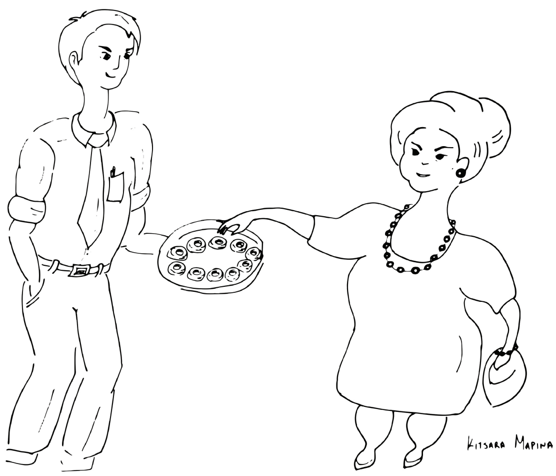

3.1 Introduction

So far we have considered the poisoning game on the complete s-uniform hypergraph. In this chapter we change the ground space and play the poisoning game on the circular hypergraph. We prove that the Kikuta-Ruckle conjecture holds true in this case. The games that we consider in this chapter are instances of the following.

Poisoning game on the hypergraph $\mathcal{H}$: Fix a hypergraph $\mathcal{H} = (V, \mathcal{E})$ and let $h > 1$ be a fixed real number. Bob (or the poisoner) chooses a measure μ on V such that $\mu(V) = h$. Alice chooses a set $E \in \mathcal{E}$. Bob wins if and only if $\mu(E) \geq 1$.

This is a win-lose game in which the payoff to Bob is 1, in case he wins, and 0 in case he loses. The measure μ can be thought of as representing the way Bob distributes the poison over V. For fixed μ and $E \in \mathcal{E}$, we will say that E is *heavy* (or lethal) under μ if $\mu(E) \geq 1$. Otherwise the set is *light*. We will also refer to $\mu(E)$ as the weight (or amount of poison) of P under μ. Notice that this a semi-finite game and so has a well defined value.

Let $\mathcal{H} = (V, \mathcal{E})$ be a fixed hypergraph. For $v \in V$, define $N(v)$ to be the set of edges from $\mathcal{E}$ that contain v. A *fractional covering* (or fractional transversal) of $\mathcal{H}$ is a real valued function $w : V \to [0, 1]$ such that

$$\sum_{v \in E} w(v) \geq 1, \text{ for all } E \in \mathcal{E}.$$

We will refer to $w(v)$ as the *weight* of v under $w(\cdot)$ and call an edge $E \in \mathcal{E}$ *heavy* if $w(E) := \sum_{v \in E} w(v) \geq 1$, otherwise E is *light*. Thus a fractional covering on $(V, \mathcal{E})$ is a function on V that makes all $E \in \mathcal{E}$ heavy. The *fractional covering number* of $\mathcal{H}$ is defined as

$$\tau^*(\mathcal{H}) := \min_{w} \sum_{v \in V} w(v),$$

where the minimum is over all fractional coverings, w, of $\mathcal{H}$.

A fractional covering on $\mathcal{H}$ can be thought of as a way to put poison over its vertices and so, provided we know that Alice chooses an edge $E \in \mathcal{E}$ uniformly at random, the poisoning game on $\mathcal{H}$ addresses the following problem. Among all subsets $\mathcal{E}_0$ of $\mathcal{E}$ whose fractional covering number is $\leq h$, find one that has the maximum cardinality.

We will also be interested in the fractional covering number of hypergraphs that are based on cyclic graphs. More precisely, fix two positive integers $a < b$ and let C_b denote the cyclic graph on b vertices, which we identify with $\mathbb{Z}_b = \mathbb{Z}/b\mathbb{Z}$. Thus the edges of C_b connect vertex i to $i + 1$. Let $\mathcal{C}_{a,b}$ be the hypergraph whose vertex set, V, is the vertex set of C_b and whose edge set, $\mathcal{E}_a$, is the set of all paths in C_b of length a. Hence $\mathcal{C}_{a,b} = (V, \mathcal{E}_a)$ and $|\mathcal{E}_a| = b$. Thus $\mathcal{C}_{a,b}$ is the *complete cyclic* hypergraph on b vertices and path length equal to a. In this chapter we study the poisoning game on $\mathcal{C}_{a,b}$. We prove that the Kikuta-Ruckle conjecture holds true for this game, i.e., that the optimal distribution of poison uses weights $1/j$, for a positive integer

j that depends on n, s, h. In order to solve this game we first need to settle the Kikuta-Ruckle conjecture on tree-like hypergraphs. This is the content of the following section.

3.2 The Kikuta-Ruckle conjecture on tree-like hypergraphs

In this section we solve the poisoning game on *tree-like* hypergraphs, which are defined as follows. Let T be a tree with vertex set V. Let $\mathcal{E} = \{E_1, \ldots, E_m\}$ be a non-empty family of subsets of V each of which induces induces a subtree of T. Then the hypergraph $\mathcal{H} = (V, \mathcal{E})$ is called tree-like.

Let us recall some more definitions from the theory of finite sets. Fix a hypergraph $\mathcal{H} = (V, \mathcal{E})$. A set $V' \subseteq V$ is a *covering* (or transversal) of $\mathcal{H}$ if it meets all edges. That is, $V' \cap E \neq \emptyset$, for all $E \in \mathcal{E}$. The *covering number* of $\mathcal{H}$, denoted $\tau(\mathcal{H})$, is the smallest cardinality of a covering. A *matching* in $\mathcal{H}$ is a family of pairwise disjoint edges. The *matching number* of $\mathcal{H}$, denoted $\nu(\mathcal{H})$, is the maximum cardinality of a matching. Note that, for any hypergraph, $\nu(\mathcal{H}) \leq \tau(\mathcal{H})$. A hypergraph is said to have the *König property* if $\nu(\mathcal{H}) = \tau(\mathcal{H})$. A *leaf* of a tree is a vertex of degree 1. The following result is taken from [7, p. 67] and its proof is included for the sake of completeness.

Lemma 3.2.1. *If $\mathcal{H} = (V, \mathcal{E})$ is a tree-like hypergraph then it has the König property.*

Proof. The proof is by induction on $\tau(\mathcal{H}) = t$. If $t = 1$, then clearly $\nu(\mathcal{H}) = \tau(\mathcal{H})$. So we may suppose that $\mathcal{H}$ has a covering, $C = \{v_1 \ldots, v_t\}$, of minimal cardinality for which $t \geq 2$. We may assume that C is such that the sub-tree, T_C, spanned by C is minimal, i.e., no other covering C' of cardinality t spans a tree that is properly contained in T_C. Furthermore, choose the covering C so that the number of vertices in T_C is minimum; this implies that any leaf of T_C belongs to C. Fix a leaf v_1 of T_C and set

$$\mathcal{E}_1 := \{E \in \mathcal{E} : E \cap C = \{v_1\}\} \text{ and } \mathcal{E}_2 := \{E \in \mathcal{E} : v_1 \in E\}.$$

The assumption that C is a covering of minimal cardinality implies that $\mathcal{E}_1$ is non-empty (indeed, if it would be empty we could remove v_1 from C and get a covering having less vertices) and the assumption that $|S|$ is minimum implies that there exists an edge $E_1 \in \mathcal{E}_1$ such that $E_1 \cap (S \setminus \{v_1\}) = \emptyset$ (otherwise we could replace T_C with $T_C \setminus \{v_1\}$ and make $|S|$ smaller). The hypergraph $\mathcal{H}' = (V, \mathcal{E} \setminus \mathcal{E}_2)$ has a covering of cardinality

$\leq t-1$; in fact a covering of cardinality equal to $t-1$ since otherwise we would have a contradiction with the assumption that the covering number of $\mathcal{H}$ is t. Thus $\nu(\mathcal{H}') = t-1$, by induction hypothesis. A maximal matching of $\mathcal{H}'$ together with the edge E_1 forms a matching of $\mathcal{H}$ of cardinality t, which implies that $\nu(\mathcal{H}) \geq t = \tau(\mathcal{H})$ and finishes the proof. □

As an example of a tree-like hypergraph, suppose that $V = \{1, \ldots, n\}$ and the elements are arranged on a line so that V is the line graph on n vertices. Let $\mathcal{A}_k$ be the set of all paths of length k in V and all paths of length $\leq k$ that have either 1 as a starting point or n as an endpoint. The poisoning game on $(V, \mathcal{A}_k)$ is defined and fully analyzed in [33, Theorem 26]. The next result generalizes its solution to tree-like hypergraphs.

Theorem 3.2.2. *Let $\mathcal{H} = (V, \mathcal{E})$ be e tree-like hypergraph and consider the poisoning game on $\mathcal{H}$. Then, an optimal strategy for Alice is to choose an element from a maximal matching uniformly at random. An optimal strategy for Bob is to put unit weights in the elements of a random subset of a minimal covering.*

Proof. Let h be the amount of poison that Bob possesses. Set $q = \lfloor h \rfloor$. Suppose that Bob chooses a minimum covering, C, of $\mathcal{H}$ and puts unit weights at the elements of a random subset of C with q vertices. Then the probability that Alice get a lethal dose is $\geq \frac{q}{|C|}$, no matter how Alice plays. On the other hand, suppose that Alice decides to play an edge of a maximum matching, M, of $\mathcal{H}$. Then the probability that she is poisoned is $\leq \frac{q}{|M|}$, since at most q edges of M are lethal. The result follows from the fact that $\mathcal{H}$ has the König property and so $|C| = |M|$. □

In this Chapter we will be concerned with poisoning games that are played on the unit circle and the cyclic graph. The analysis of these games requires some results from elementary number theory that are collected together in the next section.

3.3 Farey sequences

In this section we recall some basic facts about Farey sequences that can be found in [25, p. 28]. For any positive integer k the Farey sequence of order k, denoted $\mathcal{F}_k$, is the increasing sequence of reduced fractions between 0 and 1 whose denominators are $\leq k$. The Farey sequence of order k is defined recursively as follows. Set $\mathcal{F}_1 := \{\frac{0}{1}, \frac{1}{1}\}$ and for $i = 2, \ldots, k$ define $\mathcal{F}_i$ by doing the following: start from $\mathcal{F}_{i-1}$ and insert the fraction $\frac{a+c}{b+d}$ between consecutive fractions $\frac{a}{b}$ and $\frac{c}{d}$ if $b + d \leq i$. For example, the Farey sequence of order 5 is

$$\mathcal{F}_5 = \left\{ \frac{0}{1}, \frac{1}{5}, \frac{1}{4}, \frac{1}{3}, \frac{2}{5}, \frac{1}{2}, \frac{3}{5}, \frac{2}{3}, \frac{3}{4}, \frac{4}{5}, \frac{1}{1} \right\}.$$

The following is a well known theorem on Farey sequences. We provide a proof for the sake of completeness.

Theorem 3.3.1. *If $\frac{a}{b}$ and $\frac{a'}{b'}$ are consecutive terms of $\mathcal{F}_k$ then*

$$|ab' - a'b| = 1.$$

Proof. We induct on k. For $k = 1$ the statement is clearly correct. Suppose that it is true for all consecutive fractions in $\mathcal{F}_{k-1}$ and consider two consecutive fractions, $\frac{a}{b}$ and $\frac{a'}{b'}$, in $\mathcal{F}_k$. Suppose, without loss of generality, that $\frac{a}{b} < \frac{a'}{b'}$. If $\frac{a}{b}$ and $\frac{a'}{b'}$ are consecutive fractions in $\mathcal{F}_{k-1}$ then we are done so suppose that $\frac{a}{b}$ and $\frac{a'}{b'}$ are not consecutive in $\mathcal{F}_{k-1}$. This can only happen if one of the fractions, say $\frac{a}{b}$, belongs to $\mathcal{F}_{k-1}$ and the other fraction does not belong to $\mathcal{F}_{k-1}$. Let $\frac{c}{d}$ be the successor of $\frac{a}{b}$ in $\mathcal{F}_{k-1}$, so that $ad - cb = -1$, by the inductive hypothesis. The definition of $\mathcal{F}_k$ implies that $\frac{a'}{b'} = \frac{a+c}{b+d}$ and so $ab' - a'b = a(b + d) - (a + c)b = -1$, as required. □

In particular, this implies that successive elements of a Farey sequence have co-prime denominators.

Lemma 3.3.2. *Let $0 < x < 1$ and consider a half-open interval, I, of the positive real line, say $I = [x, y)$, that is contained in $[0, 1]$. Let $\frac{p}{q}$ be the unique rational of minimal denominator in I and $\frac{r}{s}$ be its successor in the Farey sequence of order q. If $p = 1$ then $\frac{r}{s} = \frac{1}{q-1} > y$. If $p \neq 1$, then $\frac{r}{s+1} < x < y < \frac{r}{s}$.*

Proof. The claim that $\frac{p}{q}$ is unique follows form the fact that consecutive elements in $\mathcal{F}_q$ have co-prime denominator. As $\frac{r}{s}$ is the term that succeeds $\frac{p}{q}$ in $\mathcal{F}_q$ it follows that $s < q$ and so we always have $\frac{r}{s} > y$, since otherwise $\frac{r}{s}$ would be a rational in $[x, y)$ of denominator that is smaller than the denominator of $\frac{p}{q}$. Let $\frac{r'}{s'}$ be the term before $\frac{p}{q}$ in $\mathcal{F}_q$. Similarly, the minimality of q implies that $\frac{r'}{s'} < x$.

If $p = 1$, then $\frac{r'}{s'} = \frac{0}{1}$ and the term after $\frac{1}{q}$ in $\mathcal{F}_q$ is clearly $\frac{1}{q-1}$. Now suppose that $p \neq 1$ and notice that $\frac{r}{s+1} \in \mathcal{F}_q$. If $s + 1 < q$ then $\frac{r}{s+1} < x$ and the lemma follows. If $s + 1 = q$ then $\frac{r}{s+1} < x$ follows from the uniqueness of the fraction $\frac{p}{q}$. □

We will need the above results in our analysis of the poisoning game on the cyclic graph.

3.4 A poisoning game on the unit circle

3.4.1 Definition of the game

In this section we consider the following instance of the general poisoning game. Suppose that $\mathbb{T} := \mathbb{R}/\mathbb{Z}$ is the circle of circumference 1. Bob has a fixed amount, $h \geq 1$ of poison that he can distribute over $\mathbb{T}$. So the strategy space of Bob is the set $\mathcal{M}_h$ of all Borel measures, μ, on $\mathbb{T}$ such that $\mu(\mathbb{T}) = h$. Notice that $\mathcal{M}_h$ is a convex set of measures, where μ is the convex combination of μ_1 and μ_2 if $\mu(A) = \alpha\mu_1(A) + (1-\alpha)\mu_2(A)$, for all Borel subsets, A, of $\mathbb{T}$. The strategy space of Alice is the set of all half-open intervals, $I_t := [t, t+x), t \in \mathbb{T}$, of fixed length x. Throughout this chapter we will assume that x is a *rational* number. Notice that the strategy space of Alice is the set $\mathbb{T}$, since we can identify the interval I_t with t. Let's assume for now that this game has a well defined value and optimal strategies exist. In fact, this will be a consequence of the result that this game is equivalent to a poisoning game on the cyclic graph (see lemma 3.5.3 below).
Notice that a mixed strategy of Alice is a probability measure on $\mathbb{T}$ and that a mixed strategy for Bob is a probability measure on $\mathcal{M}_h$.
If Bob plays the pure strategy $\mu \in \mathcal{M}_h$ and Alice plays the pure strategy $I_t = [t, t+x)$, then the payoff, $< \mu, t >$, to Bob is 1 if $\mu(I_t) \geq 1$, otherwise his payoff is zero. The value of the game equals the probability that Bob wins under optimal play on both sides. Given a measure $\mu \in \mathcal{M}_h$ and a set $A \in \mathcal{A}$, we will say that A is *heavy* (or lethal) under μ if $\mu(A) \geq 1$. We will also refer to $\mu(A)$ as the *weight* of A under μ.
If $h \cdot x \geq 1$ and Bob plays the uniform measure of total mass h, then any interval of length x has weight ≥ 1 and so Bob always wins. Hence, from now on, we will always assume that $h \cdot x < 1$.
If Alice plays the mixed strategy γ and Bob plays the pure strategy μ, then the probability that Bob wins equals

$$< \mu, \gamma >:= \gamma(t \in \mathbb{T} : \mu(I_t) \geq 1).$$

For fixed $t \in \mathbb{T}$ and any Borel measure β on $\mathbb{T}$ define the measure β_t by $\beta_t(\cdot) = \beta(\cdot - t)$. Suppose that Bob plays the pure strategy μ and Alice plays the mixed strategy γ. Then for any $t \in \mathbb{T}$ we have

$$\begin{aligned} < \mu_t, \gamma > &= \gamma(s \in \mathbb{T} : \mu_t(I_s) \geq 1) \\ &= \gamma(s \in \mathbb{T} : \mu(I_{s-t}) \geq 1) \\ &= \gamma(s + t \in \mathbb{T} : \mu(I_s) \geq 1) \\ &= < \mu, \gamma_{-t} >, \end{aligned}$$

and similarly

$$< \mu, \gamma_t > = < \mu_{-t}, \gamma > .$$

Notice that both μ and μ_t are pure strategies for Bob and both γ and γ_t are mixed strategies for Alice.

Lemma 3.4.1. *It is optimal for Alice to choose an interval according to the uniform probability measure.*

Proof. We prove that the optimal mixed strategy of Alice is invariant under rotations. The result will then follow from the fact that the Haar measure is the unique measure on $\mathbb{T}$ that is invariant under rotations. Let γ be a mixed strategy of Alice and let μ be an optimal pure strategy of Bob. Suppose that Alice averages γ over all rotations. That is, suppose that Alice adopts the mixed strategy γ^* defined by

$$\gamma^*(\cdot) = \int_{\mathbb{T}} \gamma_t(\cdot)dt.$$

Then the payoff to Bob against his optimal strategy is

$$\begin{aligned} < \mu, \gamma^* > &= \int_{\mathbb{T}} < \mu, \gamma_t > dt = \int_{\mathbb{T}} < \mu_{-t}, \gamma > dt \\ &\leq \int_{\mathbb{T}} < \mu, \gamma > dt = < \mu, \gamma >, \end{aligned}$$

where the inequality follows from the fact that μ is an optimal pure strategy for Bob. The fact that γ is an optimal mixed strategy for Alice implies that $< \mu, \gamma > \leq < \mu, \gamma^* >$. Hence

$$< \mu, \gamma > = < \mu, \gamma^* >,$$

and so γ^* is also an optimal strategy for Alice. This means that the optimal strategy of Alice is invariant under rotations and the result follows. □

Hence we know the optimal strategy of Alice and the solution of the game reduces to the following optimization problem:

$$\text{FIND } \mu \in \mathcal{M}_h \text{ THAT MAXIMIZES } \lambda(s \in \mathbb{T} : \mu(I_s) \geq 1),$$

where λ is Lebesgue measure.

3.4.2 Discrete measures with equidistant equal weights

In this sub-section we assume that Bob plays a discrete measure whose atoms have the same weight and determine the optimal weight of the atoms of that measure. If Bob plays such a discrete measure, μ, then an interval I satisfies $\mu(I) \geq 1$ if and only if it contains enough poisonous points. For $n = 1, 2, \ldots$ define $r_n = \frac{m_n}{n}$ to be the largest rational of denominator n that is $\leq h$ and μ_n to be the discrete measure on $\mathbb{T}$ with atoms having weight equal to $\frac{1}{n}$ at each point of the set $A_n = \{0, \frac{x}{n}, \frac{2x}{n}, \ldots, \frac{(m_n-1)x}{n}\} \subseteq \mathbb{T}$. That is, μ_n gives weight $\frac{1}{n}$ in as many equidistant points as possible. Note that the gap between these points is equal to $\delta_n := \frac{x}{n}$ except the final gap that is equal to $\Delta_n := 1 - \frac{(m_n-1)x}{n}$. Notice also that any interval $[t, t+x), t \in \mathbb{T}$, contains at most n of the points in A_n and that $\Delta_n > \delta_n$.

Alice is going to choose an interval $I_t = [t, t+x)$ uniformly at random. This means that, for $n = 1, 2, \ldots$, the probability that I_t is heavy under μ_n equals

$$P_n := \lambda(t \in \mathbb{T} : \mu_n(I_t) \geq 1),$$

where λ is Lebesgue measure. Thus, in order to find the optimal measure μ_n, we have to maximize P_n over all $n = 1, 2, \ldots$.

Let $L_n = \{t \in \mathbb{T} : \mu_n(I_t) < 1\}$ be the set of light intervals under μ_n. Then $L_n = (1 - x - \Delta_n + \delta_n, 1 - \delta_n]$. To see this take an interval $I_t = [t, t+x)$ and let t increase from 0 to 1. Any interval contains at most n of the points of A_n and it is heavy whenever it contains exactly n points of A_n. Now note that for any $t \in L_n$ the interval I_t contains strictly less than n points from A_n. Hence $\lambda(L_n) = x + \Delta_n - 2\delta_n$ and so

$$P_n = 1 - \lambda(L_n) = 1 - x - (\Delta_n - 2\delta_n).$$

From the last equation we can conclude that the integer n that maximizes P_n has to be such that $\Delta_n < 2\delta_n$.

Lemma 3.4.2. *If n is such that $\Delta_n \geq 2\delta_n$, then $P_n \leq P_1$.*

Proof. We know that $\lambda(L_n) = x + \Delta_n - 2\delta_n = 1 + x - \frac{(m_n+1)x}{n}$ and so

$$P_n = x\left(\frac{m_n + 1}{n} - 1\right).$$

Hence, in order to maximize P_n we have to maximize the fraction $\frac{m_n+1}{n}$ under the constraint $\frac{m_n+1}{n} \leq h$. If $n = 1$, then $\frac{m_n+1}{n} = \lfloor h \rfloor + 1$. If $n \neq 1$, then notice that $\frac{m_n}{n} \leq h < \lfloor h \rfloor + 1 = \frac{(\lfloor h \rfloor + 1)n}{n}$. As $\lfloor h \rfloor + 1$ is a multiple of $\frac{1}{n}$ that is $> \frac{m_n}{n}$, it follows that $\frac{m_n}{n} \leq \lfloor h \rfloor + 1 - \frac{1}{n}$ and so $\frac{m_n+1}{n} \leq \lfloor h \rfloor + 1$. Thus for all n for which $\Delta_n \geq 2\delta_n$ we have $\frac{m_n+1}{n} \leq \lfloor h \rfloor + 1$ and the maximum value of $\frac{m_n+1}{n}$ is achieved for $n = 1$. The lemma follows. □

Lemma 3.4.3. *If n is such that $\Delta_n < 2\delta_n$ then $P_n = 1 + m_n x - n$.*

Proof. We claim that in case $\Delta_n < 2\delta_n$ then every interval I_t contains either n or $n-1$ poisonous points. To see this, first notice that a heavy interval contains exactly n poisonous points and that any light interval must intersect the final gap. A light half-open interval has a minimum number of poisonous points in case it strictly contains the final gap. Since the final gap has length $\Delta_n < 2\delta_n$ it follows that a light interval has exactly $n-1$ poisonous points.
The expected number of poisonous points in a random interval of length x equals $m_n x$. Hence

$$m_n x = nP_n + (n-1)(1-P_n)$$

which implies that $P_n = 1 + m_n x - n$. □

Hence in order to maximize P_n we need to maximize $1+m_n x - n$ under the constraint $\frac{n}{m_n} \geq \frac{1}{h}$. We solve this optimization problem in the next result but first we need to fix some notation. Denote by $\frac{p}{q}$ the unique rational of smallest denominator in $[x, \frac{1}{h})$ and by $\frac{r}{s}$ the successor of $\frac{p}{q}$ in the Farey sequence of order q, $\mathcal{F}_q$.

Lemma 3.4.4. *The fraction $\frac{n}{m_n}$ that maximizes $1 + m_n x - n$, for $n = 1, 2, \ldots,$, is the successor fraction $\frac{r}{s}$ in $\mathcal{F}_q$ of the unique rational, $\frac{p}{q}$, of smallest denominator in $[x, \frac{1}{h})$.*

Proof. Note that we assume $x \leq \frac{p}{q} < \frac{1}{h}$. For any rational $\frac{a}{b} > \frac{p}{q}$ we have $0 > bp - aq \in \mathbb{Z}$ and so $bp - aq \leq -1$. This implies that for $\frac{a}{b} > \frac{p}{q}$ we have

$$1 + b\frac{p}{q} - a = 1 + \frac{bp - aq}{q} \leq 1 - \frac{1}{q}.$$

Since $\frac{r}{s}$ is the successor of $\frac{p}{q}$ in $\mathcal{F}_q$, Theorem 3.3.1 implies that $ps - qr = -1$ and so $\frac{r}{s}$ is the rational $\frac{a}{b}$ that maximizes $1 + b\frac{p}{q} - a$. That is, for any fraction $\frac{a}{b} > \frac{p}{q}$ we have $1 + b\frac{p}{q} - a \leq 1 + s\frac{p}{q} - r$. To finish the proof we show that for all rationals $\frac{a}{b} \geq \frac{1}{h}$ we have $1 + bx - a \leq 1 + sx - r$. Suppose that this is not true and so there exists a rational $\frac{a}{b} \geq \frac{1}{h}$ such that $1 + bx - a > 1 + sx - r$. We already know that $1 + s\frac{p}{q} - r \geq 1 + b\frac{p}{q} - a$. If we add the last two inequalities we get

$$b\left(x - \frac{p}{q}\right) > s\left(x - \frac{p}{q}\right)$$

and so $b < s$. Hence $\frac{a}{b} \in \mathcal{F}_q$ and since $\frac{r}{s}$ is the successor of $\frac{p}{q}$ in $\mathcal{F}_q$ it follows that $\frac{a}{b} > \frac{r}{s}$. Now note that $1 + bx - a > 1 + sx - r$ is equivalent to

$x < \frac{r-a}{s-b}$ and that $1 + s\frac{p}{q} - r \geq 1 + b\frac{p}{q} - a$ is equivalent to $\frac{r-a}{s-a} \leq \frac{p}{q}$. Since $s - a < s < q$ it follows that $\frac{r-a}{s-a}$ is a rational in $[x, \frac{1}{h})$ of denominator $< q$, which contradicts the choice of $\frac{p}{q}$. $\square$

3.5 A poisoning game on the cyclic graph

3.5.1 Definition of the game

The game on the unit circle, $\mathbb{T}$, has the following discrete analogue. Let $\mathbb{Z}_b := \mathbb{Z}/b\mathbb{Z}$ be the set of integers modulo b and $\mathcal{E}_a$ be the set of all subsets of $\mathbb{Z}_b$ that contain a consecutive integers modulo b. Equivalently, the ground space is the cyclic graph on b vertices and $\mathcal{E}_a$ consists of all paths of length a. Note that the path of length a can be thought of as a discrete interval of length a and consequently the game on the cyclic graph can be thought of as the discrete analogue of the game on the unit circle. Bob possesses $h \geq 1$ grams of poison that he may distribute over the graph. That is, his strategy space is the set of measures μ on $\mathbb{Z}_b$ such that $\mu(\mathbb{Z}_b) = h$. Note that such a measure is a vector $(\mu_0, \ldots, \mu_{b-1})$ whose coordinates add up to h and so the strategy space of Bob is a convex set. It will be referred to as *weighting* over $\mathbb{Z}_b$ of total mass h.

Alice chooses an interval, I, from $\mathcal{E}_a$. As we can identify any interval in $\mathbb{Z}_b$ with its starting point, the strategy space of Alice is $\mathbb{Z}_b$. That is, Alice chooses an element $j \in \mathbb{Z}_b$ and plays the interval $I_j := \{j, j+1, \ldots, j+a-1\}$. If Bob plays μ and Alice plays j, then the payoff, $\langle \mu, j \rangle$, to Bob is 1, if $\mu(I_j) \geq 1$, otherwise his payoff is zero. Notice that this is a semi-infinite game and so by Lemma 1.1.2 it is equivalent to a finite game. In particular, this implies that the game has a well defined value in mixed strategies. The value, $V_{a,b}$, of the game equals the probability that Bob wins, assuming optimal play on both sides. We will denote this game by $G(a, b)$.

A mixed strategy for Alice is a probability distribution on $\mathbb{Z}_b$. That is, Alice chooses a vector $\bar{p} = (p_0, \ldots, p_{b-1})$ whose coordinates add up to 1 and plays the interval I_j with probability $p_j, j = 0, 1 \ldots, b-1$. If Alice plays the *mixed* strategy $\bar{p}$ and Bob plays the *pure* strategy μ, then the expected payoff to Bob equals

$$\langle \mu, \bar{p} \rangle := \sum_{j=0}^{b-1} p_j \langle \mu, j \rangle \leq V_{a,b}.$$

Given a pure strategy, μ, of Bob and $j \in \mathbb{Z}_b$, define the measure μ_j by $\mu_j(\cdot) = \mu(\cdot - j)$. Note that μ_j is also a pure strategy for Bob. Similarly, given a mixed strategy, $\bar{p}$, of Alice and $j \in \mathbb{Z}_b$, define $\bar{p}_j$ to be the vector $(p_{-j}, \ldots, p_{b-1-j})$, which is a mixed strategy of Alice. Notice that for all $i \in \mathbb{Z}_b$ and all $j \in \mathbb{Z}_b$ we have $\langle \mu_i, j \rangle = \langle \mu, j - i \rangle$, since both are equal to 1 if $\mu(I_{j-i}) \geq 1$, and zero otherwise. Hence

$$\langle \mu_i, \bar{p} \rangle = \sum_{j=0}^{b-1} p_j \langle \mu_i, j \rangle = \sum_{j=0}^{b-1} p_j \langle \mu, j - i \rangle = \langle \mu, \bar{p}_{-i} \rangle$$

as well as

$$\langle \mu, \bar{p}_i \rangle = \langle \mu_{-i}, \bar{p} \rangle.$$

Lemma 3.5.1. *It is optimal for Alice to choose an element from $\mathbb{Z}_b$ uniformly at random.*

Proof. Let $\bar{p}$ be an optimal mixed strategy for Alice and suppose that she averages $\bar{p}$ over all rotations of $\mathbb{Z}_b$. That is, she plays the mixed strategy

$$\bar{p}_* := \frac{1}{b} \sum_{j=0}^{b-1} \bar{p}_j.$$

If μ is an optimal pure strategy for Bob, we have

$$\begin{aligned} \langle \mu, \bar{p}_* \rangle &= \frac{1}{b} \sum_{j=0}^{b-1} \langle \mu, \bar{p}_j \rangle = \frac{1}{b} \sum_{j=0}^{b-1} \langle \mu_{-j}, \bar{p} \rangle \\ &\leq \frac{1}{b} \sum_{j=0}^{b-1} \langle \mu, \bar{p} \rangle = \langle \mu, \bar{p} \rangle \leq \langle \mu, \bar{p}_* \rangle. \end{aligned}$$

Hence $\bar{p}_*$ is also an optimal mixed strategy for Alice and thus there exists an optimal mixed strategy of Alice that is invariant under rotations. The lemma follows. □

So we know the optimal strategy of Alice and thus Bob has to find a weighting over $\mathbb{Z}_b$ that creates the maximum possible number of heavy intervals. Notice that the game addresses the following problem. For given positive integers a, b and a real number $h \geq 1$, find

$$\tau_h := \max_{\mathcal{P} \subseteq \mathcal{E}_a} \{ |\mathcal{P}| : \tau^*(\mathcal{P}) \leq h \}.$$

The following result basically says that we can always assume that a and b are co-prime.

Lemma 3.5.2. *Suppose that a and b have a common divisor, say d. Then then games $G(a,b)$ and $G(a/d,b/d)$ are equivalent.*

Proof. Suppose that $a = k_1 d$ and $b = k_2 d$, where $k_1 < k_2$ are positive integers. Consider the game $G(a,b)$ and assume that Alice plays as follows. She chooses the partition $V_1,\dots,V_{k_2}$, where $V_i = \{(i-1)d+1,\dots,id\}, i = 1,\dots,k_2$, and picks k_1 consecutive sets (modulo k_2) from this partition uniformly at random. This strategy of Alice forces Bob to play on the game $G(k_1,k_2)$ and so the probability that Alice chooses a light path in $G(a,b)$, under optimal play on both sides, is at least the probability that she chooses a light path in $G(k_1,k_2)$.
On the other hand, suppose that Bob chooses an optimal weighting, γ, in $G(k_1,k_2)$ and plays the measure, μ, on $\mathbb{Z}_b$ by setting $\mu(id) = \gamma(i)$, for $i = 0,1,\dots,k_2-1$ and zero otherwise. Notice that any path of length a in $\mathbb{Z}_b$ contains k_1 vertices with non-zero weight under μ. Thus, the probability that Bob wins in $G(a,b)$ is at least the probability that he wins in $G(k_1,k_2)$. It follows that the two games are equivalent. □

3.5.2 The game on the unit circle with a rational interval

There is a close relation between the poisoning game on the unit circle and the poisoning game on the cyclic graph. This connection is clarified in the next result.

Lemma 3.5.3. *The game on the unit circle for rational $x = \frac{a}{b}$ is equivalent to the game on the cyclic graph on b vertices and discrete intervals of length a.*

Proof. Couple the two games by placing the b vertices of the cyclic graph on the unit circle at the positions $\frac{j}{b}, 0 \le j \le b-1$. Set $G_b = \{\frac{j}{b} : 0 \le j \le b-1\}$. Note that *any* interval $I_t = [t, t+x)$ contains exactly a consecutive points from G_b. Any distribution of poison in the game on the cyclic graph can also be used in the game on the unit circle by putting the poison only to the set of points G_b, of the unit circle. Hence the strategy space of the poisoner in the game on the unit circle contains the strategy space of the poisoner in the game on the cyclic graph. Thus the probability that Bob wins, assuming optimal play by both sides, in the game on the unit circle is at least the probability that he wins in the game on the cyclic graph.
On the other hand, suppose that Alice plays only intervals of the form $[\frac{j}{b}, \frac{j}{b}+x)$ in the game on the unit circle. Let $\mu \in \mathcal{M}_h$ be an optimal pure strategy of Bob. Recall that μ is a Borel measure on $\mathbb{T}$ of total mass h. Con-

sider the measure $\bar{\mu}$ that has support on G_b and is defined by

$$\bar{\mu}\left(\frac{j}{b}\right) = \mu\left(\left[\frac{j}{b}, \frac{j+1}{b}\right)\right), j = 0, 1, \ldots, b-1.$$

The weight of an interval $[\frac{j}{b}, \frac{j}{b} + x)$ under μ is equal to the sum of $\bar{\mu}\left(\frac{j}{b}\right)$ for all points $\frac{j}{b}$ that are contained in the interval. Hence if Alice plays such intervals, then Bob may as well put poison to the points $\frac{j}{b}$ only. In other words, by selecting these intervals Alice forces Bob to play the game on the cyclic graph. This implies that the probability that Alice loses (hence Bob wins), under optimal play on both sides, in the game on the unit circle is at most the probability that she loses the game on the cyclic graph. We conclude that the two games are equivalent. $\square$

Notice that the fact that the two games are equivalent, for rational x, implies that the game on the unit circle has a well defined value and that an optimal pure strategy for Bob is a discrete measure on $\mathbb{T}$. Furthermore, an optimal strategy of the game on the cyclic graph is also an optimal strategy of game on the unit circle for rational x. Notice also that we can chose any a and b such that $x = \frac{a}{b}$. In particular, we may assume that a and b are co-prime.

3.5.3 Solution of both cyclic games

Let $V = \mathbb{Z}_b$ be the vertex set of the cyclic graph, C_b, on b vertices and $\mathcal{E}_a$ be set consisting of all paths in C_b of length a. For a fixed subset $\mathcal{P}$ of $\mathcal{E}_a$, define the hypergraph $\mathcal{C}_{a,b}(\mathcal{P}) := (V, \mathcal{P})$. Note that there are $2^b - 1$ hypergraphs whose edge set consists of paths of length a in the cyclic graph. Let $\tau^*(\mathcal{P})$ be the fractional covering number of $\mathcal{C}_{a,b}(\mathcal{P})$. Given a real number $h \geq 1$, let

$$\tau_h := \max_{\mathcal{P} \subseteq \mathcal{E}_a} \{|\mathcal{P}| : \tau^*(\mathcal{P}) \leq h\}$$

be the maximum number of edges over all hypergraphs $\mathcal{C}_{a,b}(\mathcal{P})$ whose fractional covering number is at most h. Note that, since $\mathcal{E}_a$ is a finite family, τ_h is well defined for any $h \geq 1$ and the maximum is attained by a unique hypergraph.

Lemma 3.5.4. *Let $\mathcal{P} \subseteq \mathcal{E}_a$, $h \geq 1$ be such that $\tau^*(\mathcal{P}) \leq h$ and $|\mathcal{P}| = \tau_h$ and let w be a fractional covering of $\mathcal{C}_{a,b}(\mathcal{P})$ such that $\sum_v w(v) = \tau^*(\mathcal{P})$. Assume further that there is a vertex j such that $w(j) = 1$. Then all non-zero weights of w are equal to 1.*

Proof. As $w(j) = 1$ it follows that all intervals that contain j are heavy. Since $|\mathcal{P}|$ is maximum we have that w is a fractional covering of $\mathcal{C}_{a,b}(\mathcal{P})$ for which the number of heavy intervals that do not contain j is maximal. Let $V_j = V \setminus \{j\}$ and $\mathcal{P}_j = \{P \in \mathcal{P} : j \notin P\}$. Thus $\{w(i)\}_{i \neq j}$ is a fractional covering of $(V_j, \mathcal{P}_j)$ such that $|\mathcal{P}_j|$ is maximal and $\tau^*(\mathcal{P}_j) \leq h - 1$. Now from Lemma 3.2.1 we have that $(V_j, \mathcal{P}_j)$ has the König propetry and so w is a fractional covering whose non-zero weights are equal to 1. □

Lemma 3.5.5. *Let $\mathcal{P} \subseteq \mathcal{E}_a$ and suppose that w is a fractional covering of $\mathcal{C}_{a,b}(\mathcal{P})$ for which $\sum_{v \in V} w(v)$ is minimal. Suppose further that the number of vertices with non-zero weight under w is minimal. Then the endpoints of all light paths under w have zero weight.*

Proof. We argue by contradiction. Suppose that $P_j \notin \mathcal{P}$ is a light path for which $w(j) > 0$. Choose a minimal k such that $w(j-k) > 0$ and transfer the weight from j to $j-k$. This transfer reduces the number of non-zero weights by one and defines a function $\bar{w} : V \to [0,1]$. We claim that no interval that is heavy under w becomes light under $\bar{w}$. This will imply that $\bar{w}$ is a minimal fractional covering on $\mathcal{C}_{a,b}(\mathcal{P})$ having less vertices with non-zero weight, a contradiction. To prove the claim notice that $\bar{w}(P_i) < w(P_i)$ if and only if $j \in P_i$ and $j - k \notin P_i$, which holds if and only if $i \in \{j-k+1, \ldots, j\}$. But only one of these vertices has non-zero weight under w, namely vertex j. This shows that $w(P_i) \leq w(P_j)$, for $i \in \{j-k+1, \ldots, j\}$ and so all paths whose weight is reduced after the transfer were light under w. The claim follows and, hence, so does the lemma. □

For a given edge $P_j \in \mathcal{C}_{a,b}(\mathcal{P})$ let $\zeta_w(j)$ be the function that counts the number of vertices with non-zero weight in P_j under a fractional covering w.

Lemma 3.5.6. *Suppose that a and b are co-prime. Fix $h \geq 1$ and let $\mathcal{P} \subseteq \mathcal{E}_a$ be the family of paths for which $|\mathcal{P}|$ is maximum and $\tau^*(\mathcal{P}) \leq h$. Let $w : V \to [0,1]$ be a fractional covering on $\mathcal{C}_{a,b}(\mathcal{P})$, for which $\sum_{v \in V} w(v) = \tau^*(\mathcal{P})$ and assume further that $w(\cdot)$ has a minimal number of vertices with non-zero weight. Then*
(i) there exists a positive integer, n, such that $\zeta_w(j) = n - 1$ if P_j is light and $\zeta_w(j) = n$ if P_j is heavy;
(ii) If $v_1 \ldots, v_n$ are n consecutive vertices having non-zero weight, then

$$\sum_{i=1}^{n} w(v_i) \geq 1.$$

Proof. By the previous lemma we know that the endpoints of the paths that do not belong to $\mathcal{P}$ have zero weight. This implies that $\zeta_w(j) = \zeta_w(j+1)$ if

both P_j and P_{j+1} are light and $\zeta_w(j)+1=\zeta_w(j+1)$ if P_j is light and P_{j+1} is heavy. If there is a path whose vertices have zero weight then, by Lemma 3.5.4, all non-zero weights equal to 1 and the lemma holds true. So we may assume that there are no $a-1$ consecutive vertices with zero weight.
Suppose that the paths $P_1,\ldots,P_k$ are light and P_{k+1} is heavy. The previous lemma implies that the vertices $a,\ldots,a+k-1$ have zero weight and so $w(a+k)>0$, since otherwise P_{k+1} would be light. Transfer the weight from $a+k$ to a. Then all paths $P_j, j=1,\ldots,k$, become heavy since they all contain vertex a. On the other hand, the paths $P_j, j=a+1,\ldots,a+k$, reduce in weight. Since k paths became heavy after the transfer, the assumption that $|\mathcal{P}|$ is maximum implies that the paths $P_j, j=a+1,\ldots,a+k$, were heavy and become light after the transfer. All other paths preserve the same weight after the transfer. Set $J=\{1,\ldots,k\}$ and let $L=\{j : w(P_j)<1\}$ be the set of light paths under w. We saw that by a single weight transfer a maximal subinterval $J\subseteq L$ is replaced with the subinterval $a+J$ and all other vertices remain unaffected. Since a and b are co-prime we can keep on rotating maximal subintervals of L until we end up with a weighting for which all light paths of length a are consecutive. Since $\zeta(\cdot)$ is constant on consecutive light paths, it follows that all light paths contain the same number of vertices with non-zero weight. Furthermore, we can rotate L over any multiple of a. Each time we rotate L, we apply a single weight transfer and heavy paths become light by losing just one non-zero weight. This means that all heavy paths can be made light by removing a single non-zero weight. So $\zeta(\cdot)$ is constant on the heavy paths as well and the first statement follows.
To prove the second statement we may assume that the starting points of all light paths are consecutive, say $L:=\{1,\ldots,k\}$. We know that $\zeta(j)=n-1$ for $j\in L$ and $\zeta(j)=n$ for $j\notin L$. By the previous lemma we may assume that the vertices in L have zero weight and so all paths P_j with $j\in L$ contain the same $n-1$ consecutive non-zero weights. Since P_0 and P_{k+1} are heavy, it follows that the vertices 0 and $a+k$ have non-zero weight and there is no other non-zero weight between the $n-1$ non-zero weights contained in the light paths and the vertices 0 and $a+k$. It follows that any n consecutive non-zero weights are contained in a heavy path of length a. □

Theorem 3.5.7. *Consider the poisoning game on the unit circle, $\mathbb{T}$, with $\mathcal{A}$ the family of half open intervals $[t,t+x)$ of fixed* rational *length $x=\frac{a}{b}$. Let $\frac{r}{s}$ be the successor fraction in $\mathcal{F}_q$ of the unique rational, $\frac{p}{q}$, of smallest denominator in $[x,\frac{1}{h})$. Then the optimal strategy of Bob is a discrete measure with s equidistant atoms of weight $\frac{1}{r}$ and the value of the game is $1+sx-r$.*

Proof. Notice that, by Lemma 3.3.2, we have $\frac{s}{r} < h$ and so there is enough poison for Bob to play the suggested strategy. We already know that this game is equivalent to the poisoning game $G(a, b)$ on the cyclic graph. Let $\mathcal{P}$ be a subfamily of $\mathcal{E}_a$ such that $|\mathcal{P}| = \tau_h$ and $\tau^*(\mathcal{P}) \leq h$. The previous lemma implies that an optimal fractional covering on $\mathcal{C}_{a,b}(\mathcal{P})$ is such that every heavy path contains n non-zero weights and every light path contains $n-1$ non-zero weights. Let m be the total number of non-zero weights. The expected number of non-zero weights in a random path of length a equals $\frac{am}{b}$. There are τ_h paths that contain n non-zero weights and the remaining $b - \tau_h$ paths contain $n - 1$ non-zero weights. Hence

$$m\frac{a}{b} = n\frac{\tau_h}{b} + (n-1)\frac{b-\tau_h}{b} \to 1 + m\frac{a}{b} - n = \frac{\tau_h}{b}.$$

Note that $\frac{m}{n} \leq h$. As τ_h is the maximum number of edges in a sub-hypergraph of the complete cyclic hypergraph whose matching number is $\leq h$, it follows that $1 + mx - n$ is the maximum value of $1 + k\frac{a}{b} - l$, over all $\frac{k}{l} \leq h$. By Lemma 3.4.4, this maximum value is attained by the successor fraction and we know from the analysis of discrete measures with equidistant equal weights that this value is attainable. $\square$

3.6 Fractional coverings of circular hypergraphs

3.6.1 Introduction and basic result

This section deals with the following problem. Suppose that $\mathcal{H} = (V, \mathcal{E})$ is a fixed hypergraph and that you are interested in putting poison over V in such a way that you make all edges $E \in \mathcal{E}$ lethal. What is the minimum amount of poison you need? Off course, the answer depends on the underlying hypergraph. In this section we study this problem in the case of fixed sub-systems, $\mathcal{P} \subseteq \mathcal{C}_{a,b}$, of the complete cyclic hypergraph.

We have already seen that a hypergraph $\mathcal{H}$ has the König property if its covering number is equal to its matching number. An alternative way to define this is to say that $\mathcal{H}$ has a minimal fractional covering in which all non-zero weights are unit weights; just put unit weights on a covering of $\mathcal{H}$ to get a fractional covering and put unit weight on a matching of $\mathcal{H}$ to get a fractional matching.

Extend this idea, we say that $\mathcal{H}$ is a *König-Kikuta-Ruckle hypergraph*, or simply that it is KKR, is there exists a minimal fractional covering such that all vertices with positive weight have the same weight.

We consider hypergraphs that are based on cyclic hypergraphs. More precisely, consider the cyclic graph, C_b, on b vertices, identified with $\mathbb{Z}_b$. We say that a proper subset $I \subset \mathbb{Z}_b$ is a *path* if it forms a connected sub-graph of the C_b. We call $\mathcal{H}$ a *circular hypergraph* if (up to isomorphism) its set of vertices is equal to $\mathbb{Z}_b$ and its edges are paths of C_b.

The basic result of this section is the following.

Theorem 3.6.1. *A circular hypergraph is a KKR hypergraph.*

We prove this theorem in a series of lemmata. We first need to fix some definitions. A hypergraph is *redundant* if it contains edges E_1, E_2 such that E_1 is a proper subset of E_2. If we remove E_2 from the edge set of $\mathcal{H}$ then, clearly, the resulting hypergraph $\mathcal{H}'$ has the same fractional coverings as the original hypergraph. So we may restrict our attention to non-redundant hypergraphs.
We orient the circular hypergraph so that all paths are directed and have an initial point and an end-point.

Lemma 3.6.2. *A non-redundant circular hypergraph on b vertices has at most b edges*

Proof. If two paths have the same initial point then one of them is properly contained in the other. Thus a non-redundant hypergraph has at most one edge for every vertex, i.e., at most b edges. □

Lemma 3.6.3. *Suppose $\mathcal{H}$ is circular. If a vertex, v, is not an initial point of any edge in $\mathcal{H}$ then there exists a minimal fractional covering such that v has zero weight.*

Proof. Every edge that contains v also contains $v-1$ so in every fractional covering we can transfer the weight from v to $v-1$ without affecting the heavy edges. □

If v is *not* an initial point, then we can remove v from the vertex set and all the edges without changing the fractional covering number of $\mathcal{H}$. Let $\mathcal{H}_v$ be the resulting hypergraph.

Lemma 3.6.4. *Suppose that v is not an initial point. If $\mathcal{H}_v$ is KR then so is $\mathcal{H}$.*

Proof. Any minimal fractional covering of $\mathcal{H}_v$ gives rise to a minimal fractional covering of $\mathcal{H}$. □

Hence to prove our theorem, we may assume that $\mathcal{H}$ is non-redundant and that all its vertices are initial points of an edge. In particular, $\mathcal{H}$ has equally many vertices and edges.

Lemma 3.6.5. *If $\mathcal{H}$ is non-redundant and has m edges then it is KKR.*

Proof. A complete cyclic hypergraph, $\mathcal{C}_{a,b}$ is KKR. Now observe that a circular non-redundant hypergraph with m edges must be complete. □

3.6.2 One edge out

We saw in the previous subsection that every circular hypergraph is KKR and so has a minimal fractional covering such that all vertices with positive weight have the same weight. Hence a natural problem is to compute the fractional covering number of circular hypergraphs. In this section we compute the fractional covering number of the hypergraph obtained from the complete cyclic hypergraph by removing exactly one edge. That is, fix an element in $V := \mathbb{Z}_b$, say j, and define $\mathcal{C}_{a,b}(j)$ to be the hypergraph obtained from $\mathcal{C}_{a,b}$ by removing exactly one edge, namely $P_j := \{j, j+1, \ldots, j+a-1\} (\text{mod } b)$.

We begin by recalling the notions of the dual of a hypergraph and of hypergraph isomorphism. The *dual* of $\mathcal{H} = (V; E_1, \ldots, E_m)$, denoted $\mathcal{H}^*$, is a hypergraph whose vertex set $e_1, \ldots, e_m$, corresponds to the edges of $\mathcal{H}$, and with edges $X_i := \{e_j : v_i \in E_j \text{ in } \mathcal{H}\}$. That is, an edge of $\mathcal{H}^*$ consists of all edges in $\mathcal{H}$ containing a fixed element of V or, in other words, the edge set of $\mathcal{H}^*$ is the collection of sets $\{N(v) : v \in V\}$. A hypergraph $\mathcal{H}_1 = (V_1; E_1, \ldots, E_m)$ is *isomorphic* to the hypergraph $\mathcal{H}_2 = (V_2, F_1, \ldots, F_m)$, written $\mathcal{H}_1 \simeq \mathcal{H}_2$, if they have the same number m of edges, and if there exists a bijection $\phi : V_1 \to V_2$ and a permutation σ on $\{1, \ldots, m\}$ such that

$$\phi(E_i) = F_{\sigma(i)}, \ i = 1, \ldots, m.$$

A hypergraph, $\mathcal{H}$, is called *self-dual* if $\mathcal{H} \simeq \mathcal{H}^*$.

Lemma 3.6.6. *Suppose that $\mathcal{H} = (V, \mathcal{E})$ is self-dual and we can find a fractional covering, w, on $\mathcal{H}$ for which $w(E) := \sum_{v \in E} w(v) = 1$, for all $E \in \mathcal{E}$, Then the fractional covering number of $\mathcal{H}$ equals $\sum_v w(v)$.*

Proof. Notice that in this case, w gives rise to a fractional matching of $\mathcal{H}^*$ and thus, by self-duality, to a fractional matching on $\mathcal{H}$ that is equal to $\sum_v w(v)$. □

For the rest of this section we will assume that a and b are *co-prime*.

By symmetry, it is enough to find the fractional covering number of $\mathcal{C}_{a,b}(0)$. Notice that $\mathcal{C}_{a,b}(0)$ consists of b vertices and $b-1$ edges. We are interested in finding an optimal weighting $w : \mathbb{Z}_b \to [0,1]$ that makes all paths $P_j, j \neq 0$ heavy and is such that $\sum_v w(v)$ is minimal. We claim that we may suppose that such an optimal weighting satisfies $w(0) = 0$. To see this, choose an optimal weighting, w, with a minimal number of non-zero weights and suppose that $w(0) > 0$. Choose a minimal k for which $w(0-k) > 0$ and transfer the weight from 0 to $-k$ to get a new weighting. Now, as in Lemma 3.5.5 above, we can conclude that this transfer of weight does not decrease the number of heavy edges.
Hence an optimal fractional covering of $\mathcal{C}_{a,b}(0)$ gives zero weight on 0. This means that, for the purpose of finding the covering number, we might interpret $\mathcal{C}_{a,b}(0)$ as being a hypergraph on $b-1$ vertices and $b-1$ edges with vertex set $V' := \{1, 2, \ldots, b-1\}$ and edge set consisting of $P'_i := P_i \setminus \{0\}$, for $i \in \{b-a+1, \ldots, b\}$ and $P'_i := P_i$, for $i \in \{1, \ldots, b-a\}$.

Lemma 3.6.7. *The hypergraph $(V', \{P'_i\}_{i=1}^{b-1})$ is self-dual.*

Proof. Both $(V', \{P'_i\}_{i=1}^{b-1})$ and its dual have a vertex set consisting of $b-1$ vertices and an edge set consisting of $b-1$ edges. Let $E_1, \ldots, E_{b-1}$ be the edge set of the dual hypergraph. Recall that, for $\ell \in \{1, \ldots, b-1\}$, E_ℓ consists of all paths P'_j that contain vertex ℓ. Define the map $\phi : \{1, \ldots, b-1\} \to \{P'_i\}_{i=1}^{b-1}$ by $\phi(j) = P'_{-j}$. The map is clearly a bijection and it is easy to see that under this map we have

$$\phi(P'_i) = E_{b-i-a}, i = 1, \ldots, b-1.$$

The result follows. □

Hence it remains to find a fractional covering, w, on $(V', \{P'_i\}_{i=1}^{b-1})$ for which $w(P'_i) = 1$, for all $i = 1, \ldots, b-1$.

Define the map $R_a(\cdot) : \mathbb{Z}_b \to \mathbb{Z}_b$ by $R_a(j) = j + a$ and consider the set of iterates of this map,

$$\{R_a^i(a)\}_{i=0}^{t} = \{a, 2a+1, \ldots, ta\},$$

where t is the smallest integer for which $ta = -1, \mod b$. Note that such a t exists since a and b are co-prime. Notice also that $ta + 1 = 0, \mod b$ is equivalent to $ta + 1 = nb$, for some n and so $nb - ta = 1$. This implies that $\frac{a}{b} < \frac{n}{t}$ and that $\frac{n}{t}$ is the Farey successor of $\frac{a}{b}$ in the Farey sequence of order

b, $\mathcal{F}_b$.
Let $w : \mathbb{Z}_b \to [0,1]$ be the map defined as

$$w(j) = \frac{1}{n}, \text{ when } j \in \{R_a^i(a+1)\}_{i=0}^t,$$

and $w(j) = 0$, otherwise.

Lemma 3.6.8. *Let w be the map defined above. Then $w(P_j) = 1$, if $j \neq 0$ and $w(P_1) = \frac{n-1}{n}$. Moreover, all paths $P_j, j \neq 0$, contain n vertices with non-zero weight under w and P_0 contains $n-1$ vertices with non-zero weight.*

Proof. We first prove that all heavy paths of length a contain the same number of non-zero weights. Let P_i and P_{i+1} be two consecutive heavy paths. If $w(i) > 0$, then $w(i+a) > 0$ and $i + a \notin P_i$. Hence P_i and P_{i+1} contain the same number of non-zero weights. If $w(i) = 0$ then $w(i+a) = 0$ and again P_i and P_{i+1} contain the same number of non-zero weights. Hence all heavy paths contain the same number of non-zero weights. This means that, in order to determine the exact number of non-zero weights in a heavy path, it is enough to determine the number of non-zero weights in path P_a. Now the number of non-zero weights in P_a equals the number of indices i for which $R_a^i(a) \in P_a$. Since we start iterating the map $R_a(\cdot)$ from $a \in P_a$, it follows that the number of indices, i, for which $R_a^i(a) \in P_a$ equals the number of times the iterate crosses over the vertex 0, until it hits vertex -1. Thus, the number of indices, i, for which $R_a^i(a) \in P_a$ is a positive integer, d, such that $ta = db - 1$. The uniqueness of the Farey successor of $\frac{a}{b}$ in $\mathcal{F}_b$ implies that $d = n$. Since w gives weight $\frac{1}{n}$ to the elements of the set $\{R_a^i(a)\}_{i=0}^t$, it follows that all paths $P_j, j \neq 0$ satisfy $w(P_j) = 1$.
To finish the proof, we have to show that there are $n-1$ non-zero weights in the path P_0. First note that if $i \in P_0$ has non-zero weight, then $i + a \in P_a$ has also non-zero weight. The lemma now follows from the fact that the last iterate of $R_a(a)$, namely $R_a^t(a)$, is equal to -1. $\square$

Hence we have found a fractional covering, w, on $\mathcal{C}_{a,b}(0)$ such that $w(P_j) = 1$, for $j \neq 0$. Since w gives weight zero to vertex 0, it follows that w induces a fractional covering on $(V', \{P_i'\}_{i=1}^{b-1})$ for which $w(P_j') = 1$, for $j \neq 0$. The fact that $(V', \{P_i'\}_{i=1}^{b-1})$ is self-dual implies that $\sum_v w(v)$ is the fractional covering number of $(V', \{P_i'\}_{i=1}^{b-1})$. Since an optimal fractional covering of $\mathcal{C}_{a,b}(0)$ gives weight zero to vertex 0, it follows that $\sum_v w(v)$ is also the fractional covering number of $\mathcal{C}_{a,b}(0)$. Summarizing, we have proven the following result.

Theorem 3.6.9. *Suppose that a and b are co-prime. Then, for any $j \in \mathbb{Z}_b$, we have*
$$\tau^*(\mathcal{C}_{a,b}(j)) = \frac{t}{n},$$
where $\frac{n}{t}$ is the successor of $\frac{a}{b}$ in the Farey sequence $\mathcal{F}_b$, of order b.

Chapter 4

Network coloring and randomly oriented graphs

4.1 Network coloring game

4.1.1 Introduction and related work

Suppose you can color n *fair* coins with n colors. It is forbidden to color both sides of a coin with the same color, but all other colors are allowed. Let X be the number of different colors after a toss of the coins. In what way should you color the coins such that you maximize the median of X? Can you find a non-trivial upper bound on the median of X?

This is a problem that arises in the analysis of the *network coloring game*. The network coloring game is motivated by conflict resolution situations and was first defined and studied empirically in [30]. The first theoretical study of the game appears in Chaudhuri et al. (see [10]), which is the main reference of this section. The network coloring game is played on a graph, $G = (V, E)$, on n vertices and maximum degree Δ. Each vertex of the graph is thought of as a player that has k available colors. Each player has the same set of colors. As in [10] we assume that $k \geq \Delta + 2$. Fix any starting assignment of colors. The game is then played in rounds and in each round all players simultaneously and individually choose a color. They can only observe the colors chosen by their neighbors. We say that a player is *happy* if she chooses a color that is different from the colors of her neighbors, otherwise she is *unhappy*. We assume that once a player is happy, she chooses the same color in the next round. Knowing this, players will never choose a color that has been used by a neighbor in the previous round. Therefore, once a player is happy, she continues to be happy in all consecutive rounds by sticking to her color, i.e., happiness can only increase. Note that happy players are essentially removed from the game.

Suppose each player adopts the following strategy: if the player is happy she sticks to her color, if she is unhappy she changes her color and chooses equiprobably between the remaining colors that are not used by her neighbors. We call this the *simple* strategy. In [10] it is shown that under this strategy the expected number of unhappy players decays exponentially in each round. Note that the condition $k \geq \Delta + 2$ guarantees that for every unhappy player, there are always at least two colors that are not chosen by the neighbors.

Note that the condition $k \geq \Delta + 2$ is a crucial assumption. If, say, $k = \Delta + 1$ then the previous strategy might lead to a game that never ends, as can been seen by playing the game on a triangle.
For an individual player, $v \in V$, denote by τ_v the first round in which she is happy. The first round in which all players are happy, τ, is the maximum over all τ_v. In particular, the main result of Chaudhuri et al. says that

$$\mathbb{P}\left[\tau \leq O\left(\log\left(\frac{n}{\delta}\right)\right)\right] \geq 1 - \delta,$$

for arbitrary small δ. In other words, the graph is properly colored within τ steps and $\tau < c\log\left(\frac{n}{\delta}\right)$ with high probability for some constant c. It is remarkable that this estimate does *not* depend on the maximum degree of the network. The proof of this theorem depends on the following key lemma [10, p. 526]

Lemma 4.1.1 (Key Lemma)**.** *There exists a constant c such that if player v is unhappy after round t, then the probability that she is happy after* two *rounds is $\geq c$. Formally,*

$$\mathbb{P}[\tau_v \leq t + 2 | \tau_v > t] \geq c, \tag{4.1}$$

for every v.

It turns out that the constant c according to the estimates of Chaudhuri et al. is equal to $\frac{1}{1050e^9}$. Notice that this estimate does not depend on Δ. Also notice that the estimate is over *two rounds* instead of one, which is because of a two-step approach to obtain the constant c. The analysis over two steps is crucial. Over a single round, it is possible that an unhappy player gets happy in the next round with probability $\frac{1}{2^\Delta}$, an estimate which depends on maximum degree.

The game reaches a proper coloring of the graph when all players are happy. Note that when all players are happy, none of them has a motive to change her strategy and so the game reaches a Nash equilibrium.

We are going to improve the bound on the rate of convergence to Nash equilibrium by first improving the constant c in the Key Lemma. The probability that an unhappy player v gets happy after two rounds depends on two factors: the number of colors that v can choose from and the number of unhappy neighbors. Roughly, the proof of the Key Lemma is in two steps and goes as follows.

The first step concerns the event that v, who is unhappy after round t, gets many available colors in round $t + 1$. Fix a player v who is unhappy after round t. Call a color *active* if it is available to v in the next round with *positive* probability. Let A be the set of active colors for player v. Let also Y be the number of colors available to v after round $t + 1$ has been played. In [10] it is proven that $Y \geq \frac{|A|}{6}$ with probability at least $\frac{1}{25}$, by using (reverse) Markov's inequality.

The second step concerns the event that v gets happy in round $t + 2$, given that v has at least $\frac{|A|}{6}$ available colors. The probability of this event is estimated by using Markov's inequality and the Key lemma is proven by combining the estimates of the two steps. Now, using the Key Lemma, the main result in [10] is proven by applying the so-called Bayes sequential formula and an union bound. One of the intentions of this work is to avoid the use of Markov's inequality (mean estimate) and instead use ideas from search games. Below we define a simple search game that turns out to be of use to estimate the constant c in the Key Lemma. We find that the optimal strategy of the searchers involves tossing colored coins. This leads to a combinatorial probability problem whose solution allows to prove that $Y \geq \frac{|A|}{4}$ with probability at least $\frac{1}{2}$, using median estimates. The problem on colored coins, that was stated at the beginning of this chapter, rises in the analysis of the first round using this search game. Then we apply the arithmetic-geometric mean inequality to obtain a better estimate of the second step in the proof of the Key Lemma. This allow us to replace the constant c in the Key Lemma by $\frac{1}{2^9}$. Finally, we apply results on maximally dependent random variables to show that the global time to equilibrium, τ, is stochastically dominated by the maximum of n exponential random variables. In particular, we show that τ is stochastically smaller than a random variable T, such that $\mathbb{E}[T] \leq c' \log n$ (see Theorem 4.1.9 below).

Another line of research that is related to network coloring is the literature

on dispersion games. Dispersion games model situations that are similar to those modeled by the network coloring game. To give an example, suppose that there are five toilettes of comparable stylishness and ladies, attending a party every Saturday night, want to wear the one that has been chosen by the minority. In this case the ladies are happy if their dress is relatively unique.
The dispersion game was defined in [6]. It is played in rounds by a group of n players that can choose between k locations, where k divides n. In each round players simultaneously and individually choose a location and they announce their choice to all other players. A player receives payoff 1 if she has chosen a location that contains $\leq n/k$ players, otherwise she receives 0. Denote this game by $\Gamma(n,k)$. In case $n = k$ and if the locations are interpreted as colors, the dispersion game $\Gamma(n,n)$ is equivalent to the network coloring game on a complete graph on n vertices, where the number of colors available to each player is n, it's degree plus 1. In this case one can work out a one-step analysis. The following strategy is considered in [6] and is rephrased here in terms of colors:

Basic strategy. Let us say that a color is *unique* if it is chosen by one and only one player. If a player has chosen a unique color, she sticks to her choice for all consecutive rounds. If not, then she changes color and in the next round chooses equiprobably from the set of non-unique colors.

Note that this strategy uses the fact that every player knows the status of all other players. Also note that, if U_t is the total number of unhappy players after round t has been played, the probability that an unhappy player becomes happy in the next round is equal to

$$\left(1 - \frac{1}{U_t}\right)^{U_t - 1} > \frac{1}{e}.$$

By employing results on maximally dependent random variables one can conclude that the expected time to equilibrium in $\Gamma(n,n)$ under the basic strategy is less than $e + e\log n$.

There exists a vast literature on graph coloring algorithms. Some related work is in [37], where a graph coloring is provided in $O(\log n)$ rounds via a distributed algorithm which uses $\Delta + 1$ colors, or more, but requires that the neighbors have information on the status of a vertex. Attempts to properly color a graph via strategic games can be found in [39] and [16]. For an informal discussion on the network game see [11].

4.1.2 A very simple search game

In order to estimate the first time player v is happy, τ_v, we define the following search game. A player, H, the Hider, chooses an element (color) from $\Omega = \{1, 2, \ldots, d\}$ with $d \geq 2$. So the strategy space of H is the set Ω. The opponent of H consists of a team of $m \geq d$ searchers (agents) that each choose a subset Ω_j containing *at least two* colors from Ω. We denote the searchers by $S_j, 1 \leq j \leq m$. Subsequently, each searcher draws a color ω_j uniformly randomly from his own Ω_j. The searchers may communicate their choice of Ω_j. If H has chosen a color that is different from all ω_j he wins, otherwise he looses. This is a finite, one round zero-sum game that has a value, which is the probability that H wins under optimal play on both sides.

Lemma 4.1.2. *The optimal strategy for H is to choose his color uniformly at random.*

Proof. This is an invariance argument (see [17], page 24). We claim that the game is invariant under the group, $\mathcal{S}_d$, of permutations. To see this, denote by $\pi(\ell, \Omega_1, \ldots, \Omega_m)$ the payoff to H (i.e. his winning probability) provided that H has chosen ℓ and S_j has chosen $\Omega_j, j = 1, 2, \ldots, m$. Then, for any $\sigma \in \mathcal{S}_n$ we have that

$$\pi(\ell, \Omega_1, \Omega_2, \ldots, \Omega_m) = \pi(\sigma(\ell), \sigma(\Omega_1), \sigma(\Omega_2), \ldots, \sigma(\Omega_m)).$$

As the game is invariant under the group $\mathcal{S}_d$, there exist invariant optimal strategies for the players. Since for any two $\ell_1, \ell_2 \in \{1, 2, \ldots, d\}$ there exists a permutation σ that maps ℓ_1 to ℓ_2, a mixed strategy for H is invariant if it assigns the same probability to all elements of Ω. $\square$

The value of the game equals the expected proportion of the number of colors chosen by the searchers.

Lemma 4.1.3. *There exists an optimal pure strategy in which all searchers use a doubleton.*

Proof. Any searcher, S_j, picks a color uniformly at random from his own Ω_j, i.e. with probability $\frac{1}{|\Omega_j|}$. This is equivalent to first pick a doubleton from Ω_j uniformly at random and then equiprobably choose one of the two colors from that doubleton. This means that every pure strategy of S_j is equivalent to a mixed strategy on doubletons. Now we prove that it is optimal for each searcher to choose one doubleton. Since the game is finite, there exists an optimal mixed strategy for the searchers which can be described by a probability distribution on doubletons (pure strategies). Fix some searcher, say S_1, and suppose that he chooses a collection of doubletons, $D_1, D_2, \ldots, D_k$ with probabilities $p_1, p_2, \ldots, p_k$ that add up to 1. Let, P, denote the winning probability of the searchers. Then $P = \sum p_i P_i$, where P_i denotes the probability that the searchers win, given that S_1 chooses D_i and the other searchers do not change their strategy. Choose an i_0 for which $P_{i_0} = \max_i P_i$. Then $P_{i_0} \geq \sum p_i P_i$. This means that there is a doubleton such that if it is chosen by S_1, the expected payoff does not decrease, provided that the rest of the searchers do not change their strategy. $\square$

Notice that choosing a color from a doubleton uniformly at random is the same as tossing a fair colored coin. Although we will not need it, we include the computation of the value of the game, for the sake of completeness.

Theorem 4.1.4. *If $2m = ad + b$, for integers a and $0 < b < d$, then the value of the game equals $\frac{2d-b}{2^{a+1}d}$.*

Proof. Clearly, it is optimal for the searchers to use coins (doubletons) that contain every color at least once. Let Z be the set of colors chosen by the searchers after flipping their coins, let $X_{d,m} = |Z|$. That is, $X_{d,m}$ is the number of different colors after a toss. The value of the game is equal to the expected proportion of the complement of Z, $\frac{\mathbb{E}[|Z^c|]}{d} = 1 - \frac{\mathbb{E}[X_{d,m}]}{d}$. Fix some strategy, s, of the searchers, let G_s be the set of colors corresponding to this strategy and let C_i be the event that color i is chosen by the searchers after they toss their coins. Note that $|G_s| = d$. Then $\mathbb{E}[X_{d,m}] = \sum_{i \in G_s} \mathbb{P}[C_i] = \sum_{i \in G_s}(1 - (\frac{1}{2})^{c(i)})$, where $c(i)$ is the number of times that color i appears on a coin. The searchers seek to minimize the sum $\sum_{i \in G_s} \left(\frac{1}{2}\right)^{c(i)}$ under the constraint $\sum_i c(i) = 2m$. Note that whenever $l - j \geq 2$ then $\left(\frac{1}{2}\right)^l + \left(\frac{1}{2}\right)^j \geq \left(\frac{1}{2}\right)^{l-1} + \left(\frac{1}{2}\right)^{j+1}$. Iteration of this inequality shows that the minimum is achieved by choosing G_s such that all $c(i), i \in G_s$, are as equal as possible, i.e. b of them equal to $a + 1$ and the remaining $d - b$ equal to a. Then we get $\sum_{i \in G_s} \left(\frac{1}{2}\right)^{c(i)} = \frac{b}{2^{a+1}} + \frac{d-b}{2^a} = \frac{2d-b}{2^{a+1}}$. $\square$

4.1.3 Maximizing the median

Picking an element from a doubleton is just flipping a coin and so the searchers are using d colors to create m coins that do not use the same color on both sides. Note that for each array of coins used by the searchers, one can draw a *graph* whose vertices correspond to the colors and whose edges correspond to the coins. More explicitly, for each color put a vertex in the graph and join two vertices if and only if they are sides of the same coin. Note that the graph is loop-less[1] and that it might have parallel edges, because the same colored coin may occur more than one time. In addition, note that the graph may not be connected and that there is a one-to-one correspondence between array of coins and graphs and so one can choose not to distinguish between vertices and colors as well as between coins and edges. We call this graph the *dependency* graph of the set of coins. Notice that in case $m = d$, the searchers strategy $\{1,2\}\{2,3\}\ldots\{d-1,d\}\{d,1\}$ corresponds to the cycle-graph on d vertices and Theorem 4.1.4 implies that if the searchers want to maximize the mean of X_d, the number of different colors after a toss, then they have to choose coins in such a way that the corresponding graph is a cycle or a union of cycles. But what if the searchers want to maximize the *median* of X_d? What is the maximum value of the median of X_d? By *median* of a random variable, X, we mean any number μ satisfying $P[X \geq \mu] \geq 1/2$ and $P[X \leq \mu] \geq 1/2$. Notice that this μ might not be unique. It turns out that the following theorem is true.

Theorem 4.1.5. *A median of X_d is $\leq \frac{3d+2}{4}$.*

The proof of this Theorem is involved and builds on ideas from combinatorial probability. We prove this theorem in the next section. Having this result, we are then able to improve on the constant of the Key Lemma. This is the content of the following sub-section.

4.1.4 Probability of individual happiness

We now return to the network coloring game. The lemma below improves on Lemma 4, from [10].

Lemma 4.1.6. *Consider a single player, i.e., a vertex v in the network game at a given round, t, and suppose that v is unhappy. Let Y be the set of available colors to v in the next round, $t+1$, and let f be the number of happy neighbors of v in the next round, $t+1$. Then*

$$\mathbb{P}\left[|Y| \geq \frac{k-f-2}{4}\right] \geq \frac{1}{2}.$$

Proof. Let h be the number of happy neighbors of v at the start of round $t+1$. Let ξ be the degree of v. Then only $\xi - h$ unhappy neighbors are active in the game. Let I be the set of colors that are not used by the happy neighbors. For these colors there is a positive probability of being available

[1] Recall that a *loop* in a graph is an edge joining a vertex to itself.

to v in the next round. In particular I contains at least $k-h$ elements. In the worst case there are $\Delta - h \leq k-h$ unhappy neighbors all choosing a color from I. That is, the neighbors can be thought of as being searchers in the search game that was defined in the previous section. We may even add more neighbors (searchers) and suppose that the number of unhappy neighbors is $|I| \geq k-h$. If Z is the set of colors chosen by the Searchers, then we have that with probability $\geq \frac{1}{2}$, the cardinality of Z is less than $\frac{3|I|+2}{4}$, by Theorem 4.1.5. That is, with probability more than $\frac{1}{2}$ we have that $|Y| \geq \frac{|I|-2}{4} \geq \frac{k-h-2}{4} \geq \frac{k-f-2}{4}$, since the number of happy players can only increase, i.e., $f \geq h$. □

Recall that τ_v is the number of rounds needed for player v to become happy in the Network Coloring Game.

Lemma 4.1.7. *For every player, v, in the Network Coloring Game we have that*

$$\mathbb{P}[\tau_v \leq t+2 | \tau_v > t] \geq \frac{1}{2^9}.$$

Proof. Suppose that v is unhappy after round t has been played. Let Y be the set of available colors to v after round $t+1$ and f the number of happy neighbors after this round. So v is choosing a color with probability $\frac{1}{|Y|}$. Suppose that U is the set of unhappy neighbors of v after round $t+1$. Thus $|U| \leq k-f-2$. For each $u \in U$, let $p_u(i)$ be the probability with which player u chooses color i. Define also Y_u to be the set of available colors to each $u \in U$. From the previous lemma we know that with probability more than $\frac{1}{2}$ the cardinality of Y is more than $\frac{k-f-2}{4}$. The probability that a fixed color $i \in Y$ is not chosen by the neighbors is

$$\prod_{\{u \in U : i \in Y_u\}} (1 - p_u(i)).$$

Thus the probability P_v that v is happy in the next round equals

$$P_v = \frac{1}{|Y|} \sum_{i \in Y} \prod_{u \in U : i \in Y_u} (1 - p_u(i)) \geq \left(\prod_{i \in Y} \prod_{u \in U : i \in Y_u} (1 - p_u(i)) \right)^{\frac{1}{|Y|}},$$

by the arithmetic-geometric mean inequality. For each player in $u \in U$ that has i as a choice we have that $1 - p_u(i)$ equals $1 - \frac{1}{\ell}$, for some $\ell \geq 2$. If i is not a choice of $u \in U$, then $p_u(i) = 0$. Thus $1 - p_u(i) = 1 - \frac{1}{|Y_u|} \geq \frac{1}{2}$ for every

i and so

$$\begin{aligned}
\left(\prod_{i\in Y}\prod_{u\in U: i\in Y_u}(1-p_u(i))\right)^{\frac{1}{|Y|}} &\geq \left(\prod_{u\in U}\prod_{i\in Y_u}(1-p_u(i))\right)^{\frac{1}{|Y|}} \\
&\geq \left(\prod_{u\in U}\left(1-\frac{1}{|Y_u|}\right)^{|Y_u|}\right)^{\frac{1}{|Y|}} \\
&\geq \frac{1}{4^{|U|/|Y|}},
\end{aligned}$$

since $|Y_u| \geq 2$. Now on the event $|Y| \geq \frac{k-f-2}{4}$, and since $|U| \leq k-f-2$, we find $\frac{1}{4^{|U|/|Y|}} \geq \frac{1}{4^4} = \frac{1}{2^8}$. The result follows by noticing that $\mathbb{P}[\tau_v \leq t+2|\tau_v > t]$ is at least $\mathbb{P}\left[\tau_v \leq t+2|\tau_v > t, |Y| \geq \frac{k-f-2}{4}\right] \cdot \mathbb{P}\left[|Y| \geq \frac{k-f-2}{4}|\tau_v > t\right]$. $\square$

Our lower bound of $\frac{1}{2^9} \approx 0.0019531$ improves on the lower bound of $\frac{1}{1050e^9} \approx 0.0000013$ that is derived in [10]. In the next subsection we use this lower bound to estimate the expected time to global happiness.

4.1.5 Time to Nash equilibrium

So far, we have obtained a bound on the time τ_v of an individual player. Now we want to obtain a bound on the global time to happiness $\tau = \max_v \tau_v$. Unfortunately, we know nothing about the dependence structure between the τ_v, so the estimate on $\max_v \tau_v$ has to be a worst case estimate. It turns out that this worst case estimate is covered by the case of maximally dependent random variables. This is a notion that comes up in the study of stochastic order relations.

Recall that a random variable, X, is said to be stochastically smaller than another random variable, Y, if $\mathbb{P}[X > t] \leq \mathbb{P}[Y > t]$, for all t. Denote this as $X \leq_{st} Y$. It is known (see [48], Theorem $1.A.1$) that $X \leq_{st} Y$ if and only if there exist two random variables $\hat{X}, \hat{Y}$ such that $\hat{X} \sim X$, $\hat{Y} \sim Y$ and $\hat{X} \leq \hat{Y}$ with probability 1. This will apply in our case because we will show that τ_v is stochastically smaller than S_v, where $S_v \sim 2 \cdot Exp(\lambda)$ and $\lambda := -\log(1-\frac{1}{2^9})$. In that case $\max_v \hat{\tau}_v \leq \max_v \hat{S}_v$ with probability 1 and $\tau \sim \max_v \hat{\tau}_v$.

To see that $\tau_v \leq_{st} S_v$, note that the estimate of the previous subsection shows that $\mathbb{P}[\tau_v > t+2|\tau_v > t] \leq 1-\frac{1}{2^9}$. Notice also that, for every player v,

$$\mathbb{P}[\tau_v > 1|\tau_v > 0] = 1-(1-\frac{1}{k})^{\deg(v)} \leq 1-(1-\frac{1}{k})^{k-1} \leq 1-\frac{1}{e} \leq 1-\frac{1}{2^9}.$$

Hence, if t is odd,

$$\begin{aligned}\mathbb{P}[\tau_v > t] &= \mathbb{P}[\tau_v > 1|\tau_v > 0] \cdot \mathbb{P}[\tau_v > 3|\tau_v > 1] \cdots \mathbb{P}[\tau_v > t|\tau_v > t-2] \\ &\leq \left(1 - \frac{1}{2^9}\right)^{t/2} \\ &= \mathbb{P}[Exp(\lambda) > \frac{t}{2}] \\ &= \mathbb{P}[2 \cdot Exp(\lambda) > t],\end{aligned}$$

and similarly if t is even.
Thus $\tau_v \leq_{st} S_v$ which gives that $\max_v \hat{\tau}_v \leq \max_v \hat{S}_v$ with probability 1. Define $M_n := \max_v \hat{S}_v = 2\max_v X_v$, where $X_v \sim Exp(\lambda)$. Since $\tau \sim \max_v \hat{\tau}_v$ and $\max_v \hat{\tau}_v \leq M_n$ with probability 1, we conclude that $\mathbb{E}[\tau] \leq \mathbb{E}[M_n]$.
This means that, in order to estimate $\mathbb{E}[\tau]$, it is enough to estimate the maximum possible value of $\mathbb{E}[M_n] = 2\mathbb{E}[\mu_n]$, where μ_n is the maximum of n (dependent) $Exp(\lambda)$ random variables. Such ensemble maxima occur often in practical problems and have been well studied both in the independent and the dependent case (see [14], [34] and [35]).
We estimate $\mathbb{E}[\mu_n]$ using ideas from [34]. Let F be the distribution function of $X_v, v \in V$. For any real number t, we have that

$$\mu_n \leq t + \sum_v (X_v - t)^+,$$

which gives that $\mathbb{E}[\mu_n] \leq h(t) := t + n\int_t^\infty [1 - F(x)]\ dx$, for any $t \in R$. Differentiating $h(\cdot)$ one finds that its minimum is at $t_n := F^{-1}(1 - \frac{1}{n})$ and so $\mathbb{E}[\mu_n] \leq t_n + n\int_{t_n}^\infty [1 - F(x)]\ dx$. Since $1 - F(x) = e^{-\lambda x}$ it follows that $\mathbb{E}[\mu_n] \leq \frac{1}{\lambda}(1 + \log n)$. Hence

$$\mathbb{E}[\tau] \leq 2 \cdot \mathbb{E}[\mu_n] \leq \frac{2}{\lambda}(1 + \log n).$$

Similarly, in order to estimate the variance of τ, it is enough to estimate the variance of $\mu_n = \max_v X_v$. Let $X_{(1)}, \ldots, X_{(n)}$ be the order statistics of the random variables $X_v, v \in G$. The following holds true.

Lemma 4.1.8.

$$\sum_{i=1}^n Var(X_{(i)}) \leq \sum_v Var(X_v).$$

We prove this lemma in Appendix B. Having this result, we can then estimate the variance of the maximum since

$$\mathrm{Var}(\mu_n) = \mathrm{Var}(X_{(n)}) \leq \sum_v \mathrm{Var}(X_v).$$

This gives that $\mathrm{Var}(M_n) = \mathrm{Var}(2\mu_n) = 2^2\mathrm{Var}(\mu_n) \leq 2^2 \sum_v \mathrm{Var}(X_v) = \frac{4}{\lambda^2}n$. We summarize the preceding results into a theorem which is an improvement of the main result from [10].

Theorem 4.1.9. *Let G be a graph on n vertices and maximum degree Δ. If the number of available colors is at least $\Delta + 2$ and if all players adopt the simple strategy, then for any starting assignment of colors, the network coloring game reaches a proper coloring at time τ that is stochastically smaller than a random variable T, with $\mathbb{E}[T] \leq \frac{2}{\lambda}(1 + \log n)$ and Var$(T) \leq \frac{4}{\lambda^2}n$, where $\lambda \approx 0.001955$.*

4.2 Colored coin tosses

4.2.1 Fair coins

This section is devoted to the proof of Theorem 4.1.5. We want to show that median of X_d is $\leq \frac{3d+2}{4}$, where X_d is the number of different colors after a toss of d coins that are colored using d colors. Before proving this theorem we need some notation and remarks.

Suppose that we have d coins that are colored with d colors. Let G be the dependency graph corresponding to this set of coins. We are going to orient G as follows. Toss all the coins and orient each edge towards the vertex (color) that came up in the toss. Thus a toss of the coins gives rise to an orientation on the edges of G. As a consequence, $X_d = j$ corresponds to the fact that j vertices have positive in-degree, which means that $d-j$ vertices must have in-degree 0. Also note that none of the vertices of zero in-degree can be adjacent.

We denote the in-degree of a vertex v by $\deg^-(v)$ and by Z_d the number of vertices of zero in-degree. Thus $X_d = d - Z_d$.

It turns out that the median of X_d can be estimated through the median of E_d, the number of even in-degree vertices, whose distribution is easier to determine. We will need the following two graph-theoretic results.

Lemma 4.2.1. *Suppose that G is a (possibly disconnected) graph on d vertices and m edges. Fix some orientation on the edges and let $O_{d,m}$, $E_{d,m}$ be the number of odd and even in-degree vertices respectively. Then the parity of $E_{d,m}$ equals the parity of $m - d$.*

Proof. The in-degree sum formula states that

$$\sum_{v \in G} \deg^-(v) = m.$$

From this we have that the parity of $O_{d,m}$ equals the parity of m. Note that $d - E_{d,m} = O_{d,m}$. Hence the parity of m equals the parity of $d - E_{d,m}$ and the lemma follows. $\square$

For any real number r, we denote $r^+ = \max\{r, 0\}$.

Lemma 4.2.2. *For every oriented graph on d vertices, m edges and Z_d vertices of zero in-degree,*

$$-Z_d + \sum_v (deg^-(v) - 1)^+ = m - d.$$

Proof. We use again the in-degree sum formula, $\sum_v \deg^-(v) = m$. Thus

$$\sum_v (\deg^-(v) - 1) = m - d$$

and so $-Z_d + \sum_v (\deg^-(v) - 1)^+ = m - d$, since the sum contributes a "-1" for every vertex of in-degree zero. $\square$

We denote by $\mathrm{Med}(Y)$ the median of the random variable Y.

Lemma 4.2.3. *If $Med(E_d) \geq \frac{d-2}{2}$, for* any *graph on d vertices and d edges, then Theorem 4.1.5 holds true.*

Proof. Let $Y_d := E_d - Z_d$, then Lemma 4.2.2 gives that $Z_d = \sum_v (\deg^-(v) - 1)^+$, since $m = d$. Note that

$$\sum_v (\deg^-(v) - 1)^+ \geq \sum_{\{v : \deg^-(v) \geq 2\}} (\deg^-(v) - 1)^+ \geq \sum_{\{v : \deg^-(v) \geq 2\}} 1 \geq Y_d.$$

Since $Y_d + Z_d = E_d$, it follows that $Z_d \geq \frac{1}{2} E_d$. Now $X_d + Z_d = d$ so that $X_d = d - Z_d \leq d - \frac{1}{2} E_d$ and $\mathrm{Med}(X_d) \leq d - \frac{d-2}{4} = \frac{3d+2}{4}$. $\square$

So it remains to prove that $\mathrm{Med}(E_d) \geq \frac{d-2}{2}$. To prove this, we first compute the distribution of the number of even in-degree vertices in the case of a *connected* graph on d vertices and $m \geq d - 1$ edges. We then extend this computation to the general case by considering the connected components of the graph.
We denote by $\mathrm{Bin}(s, p)$ a Binomially distributed random variable of parameters s and p. In case $p = \frac{1}{2}$ we just write $\mathrm{Bin}(s)$. The parity of the in-degree of each particular vertex is related to the parity of the Binomial distribution for which the following is well known.

Lemma 4.2.4. *Suppose that $X_s := Bin(s)$ mod 2. Then X_s is a $Bin(1)$ random variable regardless of s.*

Proof. The proof is by induction on s. When $s = 1$ the conclusion is true. Suppose that it is true for all integers up to $s-1$ and consider X_s. Observe that $X_s \sim X_{s-1} + \text{Bin}(1), \mod 2$. The induction hypothesis gives that $X_{s-1} + \text{Bin}(1)$ equals $\text{Bin}(1) + \text{Bin}(1) \mod 2$, for two independent $\text{Bin}(1)$ random variables which finishes the proof of the lemma. □

The next lemma is also well known and follows immediately from the symmetry of the $\text{Bin}(s)$ distribution.

Lemma 4.2.5. *A median of a Bin(s) random variable is its mean.*

Lemma 4.2.6. *Fix some vertex v of the graph. Let C be any set of edges (coins) that does not contain some edge incident to v. Then the parity of deg$^-(v)$ is independent of the orientation of the edges in C.*

Proof. Suppose the coins corresponding to C have been flipped. Let C^- be the number of edges in C which are oriented towards v after the toss. By the previous lemma, C^- is even or odd with probability $\frac{1}{2}$. Since there is at least one edge incident to v that does not belong to C, we have that

$$\mathbb{P}[\deg^-(v) \text{ even}|C^-] = \frac{1}{2} \cdot \mathbf{1}_{\{C^- \text{ odd}\}} + \frac{1}{2} \cdot \mathbf{1}_{\{C^- \text{ even}\}} = \frac{1}{2},$$

where $\mathbf{1}_{\{\cdot\}}$ denotes indicator function. So this conditional probability does not depend on C^-. Similarly for the odd outcomes. □

We will also need a special enumeration on the vertices and edges of a tree which, combined with the previous lemma, allows us to compute the distribution of the number of even in-degree vertices.

Lemma 4.2.7. *For any tree, T, on d vertices, there exists an enumeration, $v_1, \ldots, v_d$, of the vertices and an enumeration, $e_1, \ldots, e_{d-1}$, of the edges such that, for $i = 1, \ldots, d-1$, the only edge incident to vertex v_i among the set of edges $\{e_i, e_{i+1}, \ldots, e_{d-1}\}$ is e_i.*

Proof. Fix a tree, T, on $d > 1$ vertices and choose any of its vertices. Call this vertex v_d. If v_d is a leaf, then consider the vertex set L of leaves in T except v_d and enumerate them $v_1, v_2, \ldots, v_\ell$. If v_d is not a leaf, then consider all leaves of T and enumerate them in the same manner. Note that L is not empty even if v_d is a leaf since any tree with at least two vertices has at least two leaves. Enumerate each edge incident to v_j by e_j, $j = 1, 2, \ldots, \ell$. Now consider the tree $T' := T \setminus \{v_1, v_2, \ldots, v_\ell\}$ and repeat this process on the leaves of T' again sparing v_d if it is a leaf of T'. We continue enumerating the leaves and edges of the subtrees until we end up with the graph consisting of vertex v_d only. It is evident that the enumeration satisfies the required condition. □

We are now ready to compute the distribution of the number of even in-degree vertices for any connected graph.

Theorem 4.2.8. *Suppose that G is a connected graph on d vertices and $m \geq d-1$ edges. Let $E_{d,m}$ be the number of even in-degree vertices after a random orientation on the edges. Then $E_{d,m}$ has the probability distribution of a Bin(d) random variable conditional on the event that the outcome of Bin(d) has the parity of $m-d$. Formally,*

$$\mathbb{P}[E_{d,m} = k] = \binom{d}{k}\frac{1}{2^{d-1}},$$

where k runs over the odd integers up to d if $m-d$ is odd, and over the even integers if $m-d$ is even.

Proof. Fix some spanning tree, T, of G and toss the coins corresponding to the edges that do not belong to T. Enumerate the vertices of the tree $v_1, v_2, \ldots, v_d$ and the edges $e_1, e_2, \ldots, e_{d-1}$ as in Lemma 4.2.7. Now toss the coins $e_1, e_2, \ldots, e_{d-1}$ in that order. The enumeration on the vertices and edges gives that once the coin e_j is flipped, then the parity of vertex v_j is determined. Lemma 4.2.6 gives that once the parity of some vertex v_j is determined, the parity of the next vertex v_{j+1} is independent of the parity of $v_1, v_2, \ldots, v_{j-1}$. Only the parity of v_d is deterministic given the parities of the previous vertices. Thus, if we set $\delta_i := \deg^-(v_i) \mod 2$, for $i = 1, 2, \ldots, d$, we have that each δ_i is distributed as a Bin(1) random variable which, by independence, means that $\sum_{i=1}^{d-1} \delta_i \sim$ Bin$(d-1)$. Let $O_{d,m}$ be the number of odd in-degree vertices. Then $O_{d,m} = \delta_1 + \cdots + \delta_{d-1} + \delta_d \sim X + \delta_d$, where $X \sim$ Bin$(d-1)$ and δ_d depends on the outcome of X. From the relation $O_{d,m} + E_{d,m} = d$ and the fact that X is symmetric, i.e. $X \sim d-1-X$, we get that $E_{d,m} = d - X - \delta_d \sim X + 1 - \delta_d$. Suppose that $m-d$ is even. In case $m-d$ is odd, the argument is similar. Then $E_{d,m}$ is also even, by Lemma 4.2.1, and thus $1-\delta_d$ equals 0, if X is even and equals 1, if X is odd. Hence, we have that $E_{d,m} = k$, for some even k, if and only if either $X = k$ or $X = k-1$. This means that

$$\mathbb{P}[E_{d,m} = k] = \mathbb{P}[\text{Bin}(d-1) = k] + \mathbb{P}[\text{Bin}(d-1) = k-1] = \binom{d}{k}\frac{1}{2^{d-1}}.$$

□

For any positive integer, s, we write $W \sim$ Bin(s, even) (resp. Bin(s, odd)) whenever the random variable W is distributed as a Bin(s) random variable conditioned to be even (resp. odd). We will also write Bin$(s, \odot)$ whenever we don't want to specify the exact parity and refer to it as a *half-binomial*.

Note that the proof of the last theorem says that if we are interested in an outcome of, say, Bin(s, even) (resp. Bin(s, odd)), we can toss $s-1$ fair $0/1$ coins and if the result is even, add a 0 (resp. a 1), if it is odd add 1 (resp. a 0). Call such a toss an *even-sum* (resp. *odd-sum*) toss of s coins.
We now consider the general case of a disconnected graph, G. Suppose that it consists of connected components, $G_1, G_2, \dots, G_t$ each having d_i vertices and m_i edges such that $\sum d_i = d$ and $\sum m_i = m$. Recall that we assume $d = m$. Let $E_i, 1 \leq i \leq t$ be the number of vertices of even in-degree in each graph after a toss. The E_i's are independent random variables and the total number of even in-degree vertices is given by $E = E_1 + \cdots + E_t$. Now, the distribution of each E_i is given by the previous theorem and thus E is the sum of independent Bin($d_i, \odot$) random variables. Note that if these were pure binomials instead of half-Binomials, then we would be done. In that case E would also be binomial whose median is known. The problem is that we have a sum of independent half-binomials and it is not immediately clear how to analyze a sum like Bin(7, odd)+Bin(6, even). We analyze such sums by breaking down each term of the sum, Bin($s, \odot$), into a sum of Bin($2, \odot$) and Bin($3, \odot$). More specifically, Bin($s, \odot$) will be a convex combination (mixture) of such sums. Recall that a mixture of random variables Z_i is defined as a random selection of one of the Z_i according to a probability distribution on the index set of i's. It is clear that if all these Z_i have a median that is $\geq \mu$, then also the mixture has a median $\geq \mu$.

Lemma 4.2.9. *For any $s \geq 2$, let $s = s_1 + s_2 + \cdots + s_l$ be a partition of s into $s_i = 2$ or $s_i = 3$, with at most one part equal to 3 in case s is odd. Then Bin($s, \odot$) is a mixture of sums Bin($s_1, \odot$) $+ \cdots +$ Bin($s_l, \odot$), where the parities of all these half-binomials, Bin($s_i, \odot$), add up to the given parity of Bin($s, \odot$).*

Proof. Suppose we want to decompose a Bin(s, even) random variable. The other case is similar. We get an outcome of such a half-binomial by tossing $s-1$ independent coins and add a deterministic one to fix the parity, i.e., by tossing s even-sum $0/1$ fair coins. This is equivalent to partition s into $s_1, \dots, s_l$, where all s_i are equal to 2, except possibly one that is equal to 3, and then toss l even-sum $0/1$ fair coins, assign the parity of the j-th coin, $j = 1, \dots, l$, to s_j and then this parity to Bin($s_j, \odot$). To be more precise, suppose that $Y_j \in \{\text{even}, \text{odd}\}$ is the parity of the j-th coin. Then for each $j = 1, \dots, l$, toss s_j Y_j-sum coins to get an outcome from Bin(s_j, Y_j). Then the parity of $\sum_{j=1}^{l} Y_j$ is even and thus the independent sum $\sum_{j=1}^{l}$ Bin(s_j, Y_j) has as even number of terms of the form Bin(s_j, odd) which means that it is an outcome from Bin(s, even).
To see that this is equivalent, notice that the probability of each particular outcome equals $\frac{1}{2^{l-1}} \cdot \frac{1}{2^{s_1-1}} \cdot \frac{1}{2^{s_2-1}} \cdots \frac{1}{2^{s_l-1}} = \frac{1}{2^{s-1}}$, which is exactly the

probability of each particular outcome from $\text{Bin}(s, \text{even})$. So it remains to prove that the number of outcomes for which $\text{Bin}(s, \text{even}) = k$, for some even k, equals the number of outcomes for which $\sum_{i=1}^{l} \text{Bin}(s_i, Y_i) = k$, given the parities $Y = (Y_1, \ldots, Y_l)$. But this is immediate. Every outcome of $\text{Bin}(s, \text{even})$, that is, every toss of s even-sum $0/1$ fair coins with k 1's gives rise to a vector of parities $Y = (Y_1, \ldots, Y_l)$ such that the parity of $\sum_{j=1}^{l} Y_j$ is even, $\sum_{i=1}^{l} \text{Bin}(s_i, Y_i) = k$ and vice versa. □

If we apply the last Lemma to each $E_i \sim \text{Bin}(d_i, \odot), i = 1, \ldots, t$ we get the following.

Corollary 4.2.10. *E is a mixture of sums of independent half-binomials Bin$(2, \odot)$ and Bin$(3, \odot)$.*

The reason to partition each d_i into sums of 2's and at most one 3 is the following.

Lemma 4.2.11. *Bin$(2, \odot)$ and Bin$(3, \odot)$ can be interpreted as binomials of biased coins. More precisely, they are distributed like the sum of a binomial and a scalar.*

Proof. It is easy to check that $\text{Bin}(3, \text{odd}) \sim 1+2\cdot\text{Bin}(1, \frac{1}{4})$ and $\text{Bin}(2, \text{odd}) \sim \text{Bin}(1, 1)$, as well as $\text{Bin}(3, \text{even}) \sim 2 \cdot \text{Bin}(1, \frac{3}{4})$ and $\text{Bin}(2, \text{even}) \sim 2 \cdot \text{Bin}(1, \frac{1}{2})$. □

Corollary 4.2.12. *E has the distribution of a mixture of a sum of a scalar and a sum of independent binomials.*

Having this corollary, we can then apply a well known result of Hoeffding (see [26]).

Theorem 4.2.13 (Hoeffding). *If $X_{p_1}, X_{p_2}, \ldots, X_{p_\ell}$ are independent Bernoulli trials with parameters $p_1, p_2, \ldots, p_\ell$ respectively, then*

$$\mathbb{P}[b \leq \sum_{i=1}^{\ell} X_{p_i} \leq c] \geq \mathbb{P}[b \leq \mathit{Bin}(\ell, \bar{p}) \leq c], \ \mathit{when}\ 0 \leq b \leq \ell\bar{p} \leq c \leq \ell,$$

where $\bar{p} = \frac{1}{\ell} \sum_{i=1}^{\ell} p_i$.

Recall that we are interested in a lower bound on the median of the independent sum $E \sim \sum_{i=1}^{t} E_i \sim \sum_{i=1}^{t} \text{Bin}(d_i, \odot)$. We know that E is a mixture of independent sums of $\text{Bin}(2, \odot)$ and $\text{Bin}(3, \odot)$, which are (rescaled) biased coins. We finish the proof of Theorem 4.1.5 by proving that every particular independent sum of this mixture has a median that is $\geq \frac{d-2}{2}$. Suppose that we have an independent sum, Ξ, consisting of $r, z, a, w \in \{0, 1, 2, \ldots\}$ terms from $\text{Bin}(3, \text{odd})$,$\text{Bin}(3, \text{even})$, $\text{Bin}(2, \text{even})$ and $\text{Bin}(2, \text{odd})$ respectively. Notice that $3r + 3z + 2a + 2w = d$.

Lemma 4.2.14. *A median of* Ξ *is* $\geq \frac{d-2}{2}$.

Proof. Suppose first that $z \geq r$. In that case we show that $\mathrm{Med}(\Xi) \geq \frac{d-1}{2}$. Denote by Ψ the independent sum $\mathrm{Bin}(r, \frac{1}{4})+\mathrm{Bin}(a, \frac{1}{2})+\mathrm{Bin}(z, \frac{3}{4})$. Then $\Psi = j$ if and only if $\Xi = r+2j+w$. Thus a median of Ξ can be estimated through a median of Ψ and so a median of Ξ is $\geq \frac{d-1}{2}$ if and only if a median of Ψ is $\geq \frac{r+2a+3z-1}{4}$. We apply Hoeffding's result with $\bar{p} = \frac{1}{r+a+z}\left(\frac{r+2a+3z}{4}\right), \ell = r+a+z$ and $c = r+a+z, b = \frac{r+2a+3z-1}{4}$. This gives that

$$\begin{aligned} \mathbb{P}\left[\Psi \geq \frac{r+2a+3z-1}{4}\right] &\geq \mathbb{P}\left[\mathrm{Bin}(r+a+z,\bar{p}) \geq \frac{r+2a+3z-1}{4}\right] \\ &\geq \mathbb{P}\left[\mathrm{Bin}(r+a+z,\bar{p}) \geq \frac{r+2a+3z}{4}\right]. \end{aligned}$$

Hence the lemma will follow once we prove that $\mathbb{P}[\mathrm{Bin}(r+a+z,\bar{p}) \geq \frac{r+2a+3z}{4}] \geq \frac{1}{2}$. Note that the mean of $\mathrm{Bin}(r+a+z,\bar{p})$ equals $\frac{r+2a+3z}{4}$. Now, if $z \geq r$ then $\bar{p} \geq \frac{1}{2}$ and thus $\mathrm{Bin}(r+a+z,\bar{p})$ is stochastically larger than $\mathrm{Bin}(r+a+z, \frac{1}{2})$. This means that a median of $\mathrm{Bin}(r+a+z,\bar{p})$ is bigger than or equal to a median of $\mathrm{Bin}(r+a+z, \frac{1}{2})$. But a median of $\mathrm{Bin}(r+a+z, \frac{1}{2})$ is $\frac{r+a+z}{2}$, it's mean. Since $\frac{r+a+z}{2} \leq \frac{r+2a+3z}{4}$ when $z \geq r$, the result follows.
Suppose now that $z < r$. We consider two case.
(a) Assume that $r-z$ is even. In that case we prove again that $\mathrm{Med}(\Xi) \geq \frac{d-1}{2}$. Define $\Phi_1 := \mathrm{Bin}(r, \frac{1}{4})+\mathrm{Bin}(a, \frac{1}{2})+\mathrm{Bin}(z, \frac{3}{4})+\frac{r-z}{2}$. Then $\mathrm{Med}(\Xi) \geq \frac{d-1}{2}$ if and only if $\mathrm{Med}(\Phi_1) \geq \frac{3r+2a+z}{4}$. By the result of Hoeffding we have that $\mathrm{Med}(\Phi_1) \geq \mathrm{Med}(\mathrm{Bin}(\hat{n},\hat{p}))$, where $\hat{p} = \frac{1}{r+a+z+\frac{r-z}{2}}(\frac{r}{4}+\frac{a}{2}+\frac{3z}{4}+\frac{r-z}{2}) = \frac{1}{2}$ and $\hat{n} = r+a+z+\frac{r-z}{2}$. Since $\hat{p} = 1/2$, we get that a median of $\mathrm{Bin}(\hat{n},\hat{p})$ is its mean which in turn equals $\hat{n} \cdot \hat{p} = \frac{3r+2a+z}{4}$.
(b) Assume that $r-z$ is odd. In a similar way as above we show that $\mathrm{Med}(\Xi) \geq \frac{d-2}{2}$. Define $\Phi_2 := \mathrm{Bin}(r, \frac{1}{4})+\mathrm{Bin}(a, \frac{1}{2})+\mathrm{Bin}(z, \frac{3}{4})+\frac{r-z-1}{2}$. Then $\mathrm{Med}(\Xi) \geq \frac{d-2}{2}$ if and only if $\mathrm{Med}(\Phi_2) \geq \frac{3r+2a+z-2}{4}-\frac{2}{4}$. Again, by Hoeffding, we conclude that $\mathrm{Med}(\Phi_2) \geq \mathrm{Med}(\mathrm{Bin}(\check{n},\check{p}))$, where $\check{p} = \frac{1}{r+a+z+\frac{r-z-1}{2}}(\frac{r}{4}+\frac{a}{2}+\frac{3z}{4}+\frac{r-z-1}{2})$ and $\check{n} = r+a+z+\frac{r-z-1}{2}$. Now the mean of $\mathrm{Bin}(\check{n},\check{p})$ equals $\check{n} \cdot \check{p} = \frac{3r+2a+z-2}{4}$. It is known (see [24]) that the smallest uniform (with respect to both parameters) distance of the mean and a median of a Binomial distribution is $\leq \ln 2 \approx 0.69 < \frac{3}{4}$. This means that if $\check{n} \cdot \check{p}$ equals $\mu+\frac{1}{4}$, for some integer μ, then μ is a median of $\mathrm{Bin}(\check{n},\check{p})$. If $\check{n}\cdot\check{p}$ equals $\mu+\frac{3}{4}$, for some integer μ, then $\mu+1$ is a median of $\mathrm{Bin}(\check{n},\check{p})$ and if $\check{n} \cdot \check{p} = \mu+\frac{1}{2}$, then a median of $\mathrm{Bin}(\check{n},\check{p})$ is $\geq \mu$. If the mean, $\check{n} \cdot \check{p}$ is an integer, then it is well known (see [28]) that mean and median coincide. In all cases a median is $\geq \check{n} \cdot \check{p} - \frac{1}{2}$ and the result follows. □

The colored coin tossing problem that is considered above is interesting in its own right. A natural extension of the problem that deserves further study is the following: suppose you can color n *biased* coins with n colors, all coins having the same bias. It is forbidden to color both sides of a coin with the same color, but all other colors are allowed. Let X be the number of different colors after a toss of the coins. In what way should you color the coins such that you maximize the median of X? What about a non-trivial upper bound on the median of X?
We also mention that there are possible extensions of this problem. One might try to use dies with k faces, instead of coins, and ask for the coloring over the dies for which the number of different colors after a toss is stochastically larger. We will discuss some aspects of these questions in the next sub-section.

4.2.2 Bernoulli trials of fixed parity

Denote by $B(n, p)$ a binomially distributed random variable of parameters n and p. That is, $B(n, p)$ is the number of successes in n independent and identical Bernoulli trials, Ber(p). A random variable that generalizes the binomial is defined in the following way. Fix a set of n parameters, $I = \{p_1, \ldots, p_n\}$, from $(0, 1)$ we denote by $\mathcal{H}(I)$ the random variable that counts the number of successes in n independent, non-identical Bernoulli trials, Ber(p_i), $i = 1 \ldots, n$. In other words, $\mathcal{H}(I)$ counts the number of 1's after an independent toss of n $0/1$-coins, $c_i, i = 1, \ldots, n$, having the property that coin c_i shows 1, or is a *success*, with probability p_i. The distribution of $\mathcal{H}(I)$ is well studied and is referred to as *Poisson binomial distribution*, or as *Poisson trials* (see [26],[50]). Our first result, concerning the parity of such a random variable, will be used repeatedly.

Lemma 4.2.15. *Let* $I = \{p_1, \ldots, p_n\}$ *and* $h_n := \mathcal{H}(I)$ *mod* 2. *Then* h_n *is a biased* $0/1$ *coin that lands* 1 *with probability* $\frac{1}{2}(1 - \prod_{i=1}^{n}(1 - 2p_i))$. *That is, the probability that a* $\mathcal{H}(I)$ *random variable is even equals* $\frac{1}{2}(1 + \prod_{i=1}^{n}(1 - 2p_i))$ *and the probability that it is odd equals* $\frac{1}{2}(1 - \prod_{i=1}^{n}(1 - 2p_i))$.

Proof. The proof is by induction on n. When $n = 1$ the conclusion is true. Suppose that it is true for every set of parameters having $n - 1$ numbers and consider a $\mathcal{H}(I)$ random variable such that $|I| = n$. Since $\mathcal{H}(I)$ is as an independent toss of n coins, by conditioning on the outcome of the n-th coin, we get

$$\begin{aligned}
\mathbb{P}[\mathcal{H}(I) \text{ even}] &= p_n \cdot \mathbb{P}[\mathcal{H}(I \setminus \{p_n\}) \text{ odd}] \\
&+ (1-p_n) \cdot \mathbb{P}[\mathcal{H}(I \setminus \{p_n\}) \text{ even}] \\
&= p_n \cdot \frac{1}{2}(1 - \prod_{i=1}^{n-1}(1-2p_i)) + (1-p_n) \cdot \frac{1}{2}(1 + \prod_{i=1}^{n-1}(1-2p_i)) \\
&= \frac{1}{2}(1 + \prod_{i=1}^{n}(1-2p_i))
\end{aligned}$$

Since $\mathbb{P}[\mathcal{H}(I) \text{ odd}] = 1 - \mathbb{P}[\mathcal{H}(I) \text{ even}]$ the lemma follows. □

For a fixed set of parameters $I = \{p_1, \ldots, p_n\}$, we set $\alpha(I) := \mathbb{P}[\mathcal{H}(I) \text{ even}]$ and $\beta(I) = 1 - \alpha(I)$. Note that if there is a parameter, $p_i \in I$, that is equal to $\frac{1}{2}$ then $\alpha(I) = \beta(I) = \frac{1}{2}$.

Now fix a set of parameters $I = \{p_1, \ldots, p_n\}$ and define a random variable whose outcomes have fixed parity, in the following way. First consider the case of even outcomes. Place the $0/1$ coins $c_1, \ldots, c_n$ on a line. Roll a *biased* die with n faces that shows $i \in \{1, \ldots, n\}$ with probability π_i. That is, let $\pi = (\pi_1, \ldots, \pi_n)$ be such that $\sum \pi_i = 1$ and choose i with probability π_i. If the result of the die is $i \in \{1, \ldots, n\}$, then *toss all coins except* c_i. If the outcome after the toss has an even number of 1's, then fix the parity by letting c_i to be 0. If the outcome has an odd number of 1's, then fix the parity by letting c_i to be 1. The number of 1's that we see after this (slightly dependent) toss is random. Denote it by $\mathcal{E}(I, \pi)$ and call this dependent toss an *even-sum* toss of n coins. Similarly we define the *odd-sum* toss of n coins and denote by $\mathcal{O}(I, \pi)$ the number of 1's that we see after an odd-sum toss of n coins. Formally, for an even k, the probability distribution $\mathcal{E}(I, \pi)$ is defined by

$$\mathbb{P}[\mathcal{E}(I, \pi) = k] = \sum_{i=1}^{n} \pi_i \cdot \{\mathbb{P}[\mathcal{H}(I \setminus \{p_i\}) = k] + \mathbb{P}[\mathcal{H}(I \setminus \{p_i\}) = k-1]\}$$

and similarly for an odd ℓ, the distribution of $\mathcal{O}(I, \pi)$ is defined by

$$\mathbb{P}[\mathcal{O}(I, \pi) = \ell] = \sum_{i=1}^{n} \pi_i \cdot \{\mathbb{P}[\mathcal{H}(I \setminus \{p_i\}) = \ell] + \mathbb{P}[\mathcal{H}(I \setminus \{p_i\}) = \ell-1]\}$$

Note that in case all parameters $p_i \in I, i = 1, \ldots, n$, are equal to p, then the probability distribution of an even-sum toss equals

$$\mathbb{P}[\mathcal{E}(I, \pi) = k] = \mathbb{P}[B(n-1, p) = k] + \mathbb{P}[B(n-1, p) = k-1],$$

and so does not dependent on the vector $\pi = (\pi_1, \ldots, \pi_n)$. Similarly for the odd-sum toss. In case all parameters p_i are equal to p we will denote the random variables that count the number of successes in an even-sum (resp. odd-sum) toss of n coins by $A(n,p)$ (resp. $P(n,p)$).

Notice also that in case $p_i = \frac{1}{2}$, for all $i \in \{1, \ldots, n\}$, the above formulas reduce to

$$\begin{aligned}
\mathbb{P}[A(n,1/2)=k] &= \mathbb{P}\left[B(n-1,1/2)=k\right]+\mathbb{P}\left[B(n-1,1/2)=k-1\right] \\
&= \binom{n-1}{k}\frac{1}{2^{n-1}}+\binom{n-1}{k-1}\frac{1}{2^{n-1}} \\
&= \binom{n}{k}\frac{1}{2^{n-1}}
\end{aligned}$$

and similarly for $P(n,1/2)$.
The random variables just defined are related to the random variable $\mathcal{H}(I)$, conditional on the event that its outcomes have fixed parity. More precisely, denote by $\mathcal{H}(I,0)$ (resp.$\mathcal{H}(I,1)$) the random variable that has the same distribution as $\mathcal{H}(I)$ conditional on the event that it's outcome is even (resp. odd). That is, for even k

$$\mathbb{P}[\mathcal{H}(I,0)=k]=\frac{1}{\alpha(I)}\mathbb{P}[\mathcal{H}(I)=k]$$

and, for an odd ℓ,

$$\mathbb{P}[\mathcal{H}(I,1)=\ell]=\frac{1}{\beta(I)}\mathbb{P}[\mathcal{H}(I)=\ell].$$

Hence we can obtain an outcome of a, say, $\mathcal{H}(I,0)$ random variable by tossing the coins again and again until we see an even outcome. In case I consists of n parameters all equal to p, we will write $B(n,p,0)$ for $\mathcal{H}(I,0)$ and $B(n,p,1)$ for $\mathcal{H}(I,1)$. Thus $B(n,p,0)$ is the random variable whose distribution function is binomial, conditional on the event that the outcomes are even. Similarly for $B(n,p,1)$.
The following results shows the relation between conditional Poison trials and the random variables that are under consideration.

Lemma 4.2.16. *If $I = \{p_1, \ldots, p_n\}$ and $\pi = (\pi_1, \ldots, \pi_n)$ is a probability vector then the distribution of $\mathcal{E}(I,\pi)$ is the same as the distribution of the random variable that takes even outcomes according to the following procedure. Roll a biased die with n faces. If the result of the die is $i \in \{1, \ldots, n\}$ with probability π_i, then toss a $0/1$ coin having probability of showing 1 equal to $1-\alpha(I \setminus \{p_i\}) =$*

$\beta(I \setminus \{p_i\})$. If the outcome of this coin is 0, then draw from a $\mathcal{H}(I \setminus \{p_i\}, 0)$ random variable and add 0. If the outcome is 1, then draw from a $\mathcal{H}(I \setminus \{p_i\}, 1)$ random variable and add 1.

Proof. For an even k, write

$$\begin{aligned}\mathbb{P}[\mathcal{E}(I,\pi) = k] &= \sum_{i=1}^{n} \pi_i \cdot \left\{ \alpha(I \setminus \{p_i\}) \cdot \frac{\mathbb{P}[\mathcal{H}(I \setminus \{p_i\}) = k]}{\alpha(I \setminus \{p_i\})} \right\} \\ &+ \sum_{i=1}^{n} \pi_i \cdot \left\{ \beta(I \setminus \{p_i\}) \cdot \frac{\mathbb{P}[\mathcal{H}(I \setminus \{p_i\}) = k-1]}{\beta(I \setminus \{p_i\})} \right\},\end{aligned}$$

which can be rewritten as

$$\begin{aligned}\mathbb{P}[\mathcal{E}(I,\pi) = k] &= \sum_{i=1}^{n} \pi_i \cdot \alpha(I \setminus \{p_i\}) \cdot \mathbb{P}[\mathcal{H}(I \setminus \{p_i\}, 0) = k] \\ &+ \sum_{i=1}^{n} \pi_i \cdot \beta(I \setminus \{p_i\}) \cdot \mathbb{P}[\mathcal{H}(I \setminus \{p_i\}, 1) = k-1]\end{aligned}$$

and finishes the proof of the lemma. □

For random variables Y, W that take values on the same sets, we will write $Y \sim W$ whenever Y and W have the same distribution. Note that, in case all parameters p_i are equal to p, the previous lemma says that $A(n,p)$ has the same distribution as the random variable that takes even outcomes according to the following procedure. Toss a $0/1$ coin whose probability of showing 1 equals $\beta(\{p\}_{n-1})$. If the outcome is a 1, then toss n independent $0/1$ coins that show up 1 with probability 1 until you see an odd outcome, and add a 1. If the outcome is 0, then toss n independent $0/1$ coins that show 1 with probability p until you see an even outcome, and add a 0 to this outcome. We can formally express this as

$$A(n,p) \sim B(1, \beta(\{p\}_{n-1})) + B(n-1, p, B(1, \beta(\{p\}_{n-1}))).$$

Similarly one can prove the following result for $\mathcal{O}(I,\pi)$.

Lemma 4.2.17. *If $I = \{p_1, \ldots, p_n\}$ and $\pi = (\pi_1, \ldots, \pi_n)$ is a probability vector, then the distribution of $\mathcal{O}(I,\pi)$ is the same as the distribution of the random variable that takes odd outcomes according to the following procedure. Roll a biased die with n faces. If the result of the die is $i \in \{1, \ldots, n\}$ with probability π_i, then toss a $0/1$ coin having probability of showing 1 equal to $\alpha(I \setminus \{p_i\}) = 1 - \beta(I \setminus \{p_i\})$. If the outcome of the coin is a 0, then draw from a $\mathcal{H}(I \setminus \{p_i\}, 1)$ random variable while and add a 0. If the outcome is a 1 then draw from a $\mathcal{H}(I \setminus \{p_i\}, 0)$ random variable and add a 1.*

Again, in case all parameters p_i are equal to p, the previous lemma can be formally expressed as

$$P(n,p) \sim B(1,\alpha(\{p\}_{n-1})) + B(n-1,p,1-B(1,\alpha(\{p\}_{n-1}))).$$

Lemma 4.2.16 and 4.2.17 imply that the distributions of $\mathcal{E}(I,\pi), \mathcal{O}(I,\pi)$ can be analyzed via the distributions $\mathcal{H}(I\setminus\{p_i\},0)$ and $\mathcal{H}(I\setminus\{p_i\},1)$. The next result can be used in case one is interested in adding independent copies of $\mathcal{E}(\cdot,\cdot)$ and $\mathcal{O}(\cdot,\cdot)$.

Lemma 4.2.18. *Let $I=\{p_1,\dots,p_n\}$ and consider a partition of I into disjoint, non-empty sets I_1, I_2. Then the distribution of $\mathcal{H}(I,0)$ is a mixture of the independent sums $\mathcal{H}(I_1,0)+\mathcal{H}(I_2,0)$ and $\mathcal{H}(I_1,1)+\mathcal{H}(I_2,1)$. More precisely, for an even k, we have*

$$\begin{aligned}\mathbb{P}[\mathcal{H}(I,0)=k] &= \frac{\alpha(I_1)\cdot\alpha(I_2)}{\alpha(I)}\mathbb{P}[\mathcal{H}(I_1,0)+\mathcal{H}(I_2,0)=k] \\ &+ \frac{\beta(I_1)\cdot\beta(I_2)}{\alpha(I)}\mathbb{P}[\mathcal{H}(I_1,1)+\mathcal{H}(I_2,1)=k].\end{aligned}$$

Proof. Write $\mathbb{P}[\mathcal{H}(I,0)=k]=\frac{\mathbb{P}[\mathcal{H}(I)=k]}{\alpha(I)}$ and note that if we regard $\mathcal{H}(I)$ as an independent sum of $\mathcal{H}(I_1)$ and $\mathcal{H}(I_2)$, then $\mathbb{P}[\mathcal{H}(I)=k]$ equals

$$\sum_{i:i \text{ even}} \mathbb{P}[\mathcal{H}(I_1)=i]\cdot\mathbb{P}[\mathcal{H}(I_2)=k-i] + \sum_{i:i \text{ odd}} \mathbb{P}[\mathcal{H}(I_1)=i]\cdot\mathbb{P}[\mathcal{H}(I_2)=k-i].$$

Multiply and divide the sum that runs over even indices by $\alpha(I_1)\cdot\alpha(I_2)$ and the sum that runs over odd indices by $\beta(I_1)\cdot\beta(I_2)$ to get the result. $\square$

Similarly, one can prove the following.

Lemma 4.2.19. *Let $I=\{p_1,\dots,p_n\}$ and consider a partition of I into disjoint, non-empty sets I_1, I_2. Then the distribution of $\mathcal{H}(I,1)$ is a mixture of the independent sums $\mathcal{H}(I_1,1)+\mathcal{H}(I_2,0)$ and $\mathcal{H}(I_1,0)+\mathcal{H}(I_2,1)$. More precisely, for an odd k, we have*

$$\begin{aligned}\mathbb{P}[\mathcal{H}(I,1)=k] &= \frac{\alpha(I_1)\cdot\beta(I_2)}{\beta(I)}\mathbb{P}[\mathcal{H}(I_1,0)+\mathcal{H}(I_2,1)=k] \\ &+ \frac{\beta(I_1)\cdot\alpha(I_2)}{\beta(I)}\mathbb{P}[\mathcal{H}(I_1,1)+\mathcal{H}(I_2,0)=k].\end{aligned}$$

The last two lemmata can be iterated. By doing so one gets that every $\mathcal{H}(I,0)$ or $\mathcal{H}(I,1)$ random variable is a mixture of independent sums consisting only of summands of the form $\mathcal{H}(\{a,b\},0)$, $\mathcal{H}(\{c,d\},1)$, $\mathcal{H}(\{e,f,g\},0)$

and $\mathcal{H}(\{k,l,m\},1)$, where $a,b,c,d,e,f,g,k,l,m \in (0,1)$. That is, one may apply the last two lemmata by partitioning I into $I_1 \cup D_1$, where D_1 is a doubleton. Then apply the lemma again by partitioning I_1 into $I_2 \cup D_2$, for some doubleton D_2 and so on.
The reason to partition I this way is the next result that says that all terms of the previous mixture are rescaled biased coins. Its proof is immediate.

Lemma 4.2.20. *Let* $I = \{p_1, p_2\}$ *and* $J = \{q_1, q_2, q_3\}$. *Then* $\mathcal{H}(I,0) \sim 2 \cdot B(1, \frac{p_1 \cdot p_2}{\alpha(I)})$, $\mathcal{H}(I,1) \sim B(1,1)$, $\mathcal{H}(J,0) \sim 2 \cdot B(1, 1 - \frac{(1-q_1)\cdot(1-q_2)\cdot(1-q_3)}{\alpha(J)})$ *and* $\mathcal{H}(J,1) \sim 1 + 2 \cdot B(1, \frac{q_1 \cdot q_2 \cdot q_3}{\beta(J)})$

The next result is an inequality on conditional binomial random variables. Set $\alpha_n = \mathbb{P}[B(n,p) \text{ even}]$ and $\beta_n = 1 - \alpha_n$.

Lemma 4.2.21. *Fix a positive integer* n *and a real number* $p \in (0,1)$. *Then*

$$\mathbb{P}[B(n,p,1) \geq 2k-1] \geq \mathbb{P}[B(n,p,0) \geq 2k]$$

and

$$\mathbb{P}[B(n,p,0) \geq 2k] \geq \mathbb{P}[B(n,p,1) \geq 2k+1].$$

Proof. We induct on n. For $n = 2$ it is easy to check that both inequalities hold true, so suppose that both inequalities hold true for all positive integers that are $\leq n-1$. Let $q = 1-p$. The fact that $1 - 2q = -1 + 2p$ and the symmetry of the binomial distribution imply that it is enough to check the inequalities for $p \in (0, 1/2]$. In order to simplify notation, set $X_n = B(n,p,0)$ and $Y_n = B(n,p,1)$. From Lemma 4.2.18 and Lemma 4.2.19 we know that

$$\mathbb{P}[Y_n \geq 2i-1] = \frac{p\alpha_{n-1}}{\beta_n}\mathbb{P}[X_{n-1} \geq 2i-2] + \frac{(1-p)\beta_{n-1}}{\beta_n}\mathbb{P}[Y_{n-1} \geq 2i-1],$$

$$\mathbb{P}[X_n \geq 2i] = \frac{(1-p)\alpha_{n-1}}{\alpha_n}\mathbb{P}[X_{n-1} \geq 2i] + \frac{p\beta_{n-1}}{\alpha_n}\mathbb{P}[Y_{n-1} \geq 2i-1].$$

and that

$$\mathbb{P}[Y_n \geq 2i+1] = \frac{p\alpha_{n-1}}{\beta_n}\mathbb{P}[X_{n-1} \geq 2i] + \frac{(1-p)\beta_{n-1}}{\beta_n}\mathbb{P}[Y_{n-1} \geq 2i+1].$$

Since $p \leq 1/2$ it is easy to check that

$$\frac{p}{\beta_n} \leq \frac{1-p}{\alpha_n} \quad \text{and} \quad \frac{p}{\alpha_n} \leq \frac{1-p}{\beta_n}.$$

Hence

$$\mathbb{P}[Y_n \geq 2i-1] \geq \mathbb{P}[X_n \geq 2i]$$

if and only if

$$\mathbb{P}[Y_{n-1} \geq 2i-1] \cdot \beta_{n-1} \cdot \left(\frac{1-p}{\beta_n} - \frac{p}{\alpha_n}\right) \geq$$
$$\mathbb{P}[X_{n-1} \geq 2i] \cdot \alpha_{n-1} \cdot \left(\frac{1-p}{\alpha_n} - \frac{p}{\beta_n}\right) - \frac{p\alpha_{n-1}}{\beta_n} \cdot \mathbb{P}X_{n-1} = 2i-2]$$

Elementary calculations, and the fact that $\alpha_n = p + (1-2p)\alpha_{n-1}$, imply

$$\beta_{n-1} \cdot \left(\frac{1-p}{\beta_n} - \frac{p}{\alpha_n}\right) = \alpha_{n-1} \cdot \left(\frac{1-p}{\alpha_n} - \frac{p}{\beta_n}\right)$$

and the result follows from the inductional hypothesis. Similarly,

$$\mathbb{P}[X_n \geq 2i] \geq \mathbb{P}[Y_n \geq 2i+1]$$

if and only if

$$\mathbb{P}[X_{n-1} \geq 2i] \cdot \alpha_{n-1} \cdot \left(\frac{1-p}{\alpha_n} - \frac{p}{\beta_n}\right) \geq$$
$$\mathbb{P}[Y_{n-1} \geq 2i+1] \cdot \beta_{n-1} \cdot \left(\frac{1-p}{\beta_n} - \frac{p}{\alpha_n}\right) - \frac{p\beta_{n-1}}{\alpha_n} \cdot \mathbb{P}[Y_{n-1} = 2i-1].$$

Elementary calculations, and the fact that $\alpha_n = p + (1-2p)\alpha_{n-1}$, imply

$$\alpha_{n-1} \cdot \left(\frac{1-p}{\alpha_n} - \frac{p}{\beta_n}\right) = \beta_{n-1} \cdot \left(\frac{1-p}{\beta_n} - \frac{p}{\alpha_n}\right)$$

and, once again, the inductional hypothesis finishes the proof. □

As a corollary we obtain the following result that will be used in our analysis of colored coin tosses. Recall that a random variable X is said to be stochastically larger than another random variable Y, denoted by $X \geq_{st} Y$, if $\mathbb{P}[X \geq t] \geq \mathbb{P}[Y \geq t]$, for all t.

Corollary 4.2.22. *Let $p_1 \geq p_2 \geq p$ be three real number from $(0,1)$ and fix a positive integer n. Then*

$$B(1,p_1) + B(n,p,B(1,p_1)) \geq_{st} B(1,p_2) + B(n,p,B(1,p_2))$$

and

$$B(1,p_1) + B(n,p,1-B(1,p_1)) \geq_{st} B(1,p_2) + B(n,p,1-B(1,p_2)).$$

Proof. We only prove the first inequality, the other can be proved similarly. Set $X_1 = B(1, p_1) + B(m, p, B(1, p_1))$ and $X_2 = B(1, p_2) + B(m, p, B(1, p_2))$. We want to prove that, for every even integer, say $2k$, in $\{0, 1, \ldots, n\}$, we have $\mathbb{P}[X_1 \geq 2k] \geq \mathbb{P}[X_2 \geq 2k]$. This inequality is equivalent to

$$p_1 \cdot \mathbb{P}[B(n, p, 1) \geq 2k - 1] + (1 - p_1) \cdot \mathbb{P}[B(n, p, 0) \geq 2k] \geq \\ p_2 \cdot \mathbb{P}[B(n, p, 1) \geq 2k - 1] + (1 - p_2) \cdot \mathbb{P}[B(n, p, 0) \geq 2k]$$

and the later holds true if and only if

$$\mathbb{P}[B(n, p, 1) \geq 2k - 1] \geq \mathbb{P}[B(n, p, 0) \geq 2k].$$

Lemma 4.2.21 finishes the proof. □

The following result gives a lower on a median of the random variables $A(n, p)$ and $P(n, p)$. Recall that a *median* of a random variable, Y, is any number μ satisfying $P[Y \geq \mu] \geq 1/2$ and $P[Y \leq \mu] \geq 1/2$. Notice that this μ might not be unique. By abuse of notation, we will denote any median of Y by $\mathrm{Med}(Y)$.

Lemma 4.2.23. *Fix a $p \in (0, 1)$ and a positive integer n. Then a median of a $A(n, p)$ random variable is $\geq (n - 1)p - 1$. Similarly, a median of a $P(n, p)$ random variable is $\geq (n - 1)p - 1$.*

Proof. We prove the result for $A(n, p)$. A similar argument works for $P(n, p)$. For any even k, we have

$$\mathbb{P}[A(n, p) \geq k] = \mathbb{P}[B(n - 1, p) \geq k - 1].$$

Now it is well known (see [28]) that a median of a $B(n - 1, p)$ random variable is $\geq \lfloor (n - 1)p \rfloor$. If $\lfloor (n - 1)p \rfloor$ is odd, then a median of $A(n, p)$ is $\geq \lfloor (n - 1)p \rfloor + 1 \geq (n - 1)p$. If $\lfloor (n - 1)p \rfloor$ is even, then $a := \lfloor (n - 1)p \rfloor - 1$ is odd and is such that $\mathbb{P}[B(n - 1, p) \geq a] \geq 1/2$. Thus a median of $A(n, p)$ is $\geq \lfloor (n - 1)p \rfloor \geq (n - 1)p - 1$. □

We also mention an important theorem, obtained by Hoeffding (see [26]), that will be used in the next section.

Theorem 4.2.24. *If $I = \{p_1, \ldots, p_n\}$ is a set of parameters in $(0, 1)$, then*

$$\mathbb{P}[b \leq \mathcal{H}(I) \leq c] \geq \mathbb{P}[b \leq B(n, \bar{p}) \leq c], \text{ when } 0 \leq b \leq n\bar{p} \leq c \leq n,$$

where $\bar{p} = \frac{1}{n} \sum_{i=1}^{n} p_i$.

The following result is well known (see [48]). We include a proof for the sake of completeness.

Lemma 4.2.25. *Let $I = \{p_1, \ldots, p_n\}$ and $J = \{q_1, \ldots, q_n\}$ be sets of parameters in $(0, 1)$ such that $p_i \geq q_i$, for all $i = 1, \ldots, n$. Then $\mathcal{H}(I) \geq_{st} \mathcal{H}(J)$*

Proof. Since $p_i \geq q_i$ for all i, it follows that $q_i = \frac{q_i}{p_i} p_i$ and $r_i := \frac{q_i}{p_i} \leq 1$. Begin by tossing n coins, $c_i, i = 1, \ldots, n$, such that coin c_i has probability of success p_i. Let H be the number of successes. Then $H \sim \mathcal{H}(I)$. Now replace each coin c_i that was a success, by the outcome of an independent coin c_i^* that has probability of success r_i. The final number of successes, H', is distributed like an $\mathcal{H}(J)$ random variable, by independence. If $H' \geq t$, then $H \geq t$ and thus $\mathbb{P}[H' \geq t] \leq \mathbb{P}[H \geq t]$, as required. □

For a given set of parameters, $I = \{p_1, \ldots, p_n\}$, let $p_0 = \min_{1 \leq i \leq n} p_i$ and denote by $I_0 = \{p_0, \ldots, p_0\}$ the set consisting of n copies of p_0. The following holds true.

Lemma 4.2.26. *For the sets I and I_0 defined above and any probability vector π, we have*

$$\mathcal{E}(I, \pi) \geq_{st} \mathcal{E}(n, p_0) \textit{ and } \mathcal{O}(I, \pi) \geq_{st} \mathcal{O}(n, p_0).$$

Proof. We only consider the case $\mathcal{E}(I, \pi) \geq_{st} \mathcal{E}(I_0)$, the other being similar. Notice that, for an even k, we have

$$\mathbb{P}[\mathcal{E}(I) \geq k] = \sum_{i=1}^{n} \pi_i \cdot \mathbb{P}[\mathcal{H}(I \setminus \{p_i\}) \geq k - 1].$$

By the previous lemma we have $\mathcal{H}(I \setminus \{p_i\}) \geq_{st} \mathcal{H}(I_0 \setminus \{p_0\})$, for each $i = 1, \ldots, n$. Hence $\mathbb{P}[\mathcal{H}(I \setminus \{p_i\}) \geq k - 1] \geq \mathbb{P}[\mathcal{H}(I_0 \setminus \{p_0\}) \geq k - 1]$. This gives that

$$\mathbb{P}[\mathcal{E}(I, \pi) \geq k] \geq \mathbb{P}[\mathcal{H}(I_0 \setminus \{p_0\}) \geq k - 1] = \mathbb{P}[\mathcal{E}(n, p_0) \geq k],$$

as required. □

We end with two results on Bernoulli random variables that will be used in the next subsection.

Lemma 4.2.27. *Let $p \in (0, 1)$ and suppose that $X_i, i = 1, \ldots, s$ are $\{0, 1\}$-valued random variables such that $\mathbb{P}[X_1 = 1] \geq p$ and*

$$\mathbb{P}[X_i = 1 | X_1, \ldots, X_{i-1}] \geq p, \textit{ for all } i = 2, \ldots, s. \tag{4.2}$$

Then $\Sigma_s := X_1 + \cdots + X_s$ is stochastically larger than a $B(s,p)$ random variable. Furthermore, it is possible to define random vectors $\mathbf{U} = (U_1, \ldots, U_s)$ and $\mathbf{V} = (V_1, \ldots, V_s)$ on a common probability space so that the law of $(U_1, \ldots, U_s)$ is the same as the law of $(X_1, \ldots, X_s)$, each coordinate of $\mathbf{V}$ is an independent Ber(p) random variable and

$$V_i \leq U_i, \text{ for all } i = 1, \ldots, s, \text{ with probability } 1.$$

Proof. We want to prove that

$$\mathbb{P}[\Sigma_s \geq t] \geq \mathbb{P}[B(s,p) \geq t], \text{ for all } t \in \{0, 1, \ldots, s\}.$$

Note that every outcome of the random variables $X_i, i = 1, \ldots, s$ is an s-tuple $(x_1, \ldots, x_s) \in \{0,1\}^s$. We associate a binary vector $\mathbf{b} = (b_1, \ldots, b_s)$ to every outcome of $X_i, i = 1, \ldots, s$ in such a way that the number of 1's in $\mathbf{b}$ has the same distribution as a $B(s,p)$ random variable.
To do so, begin by drawing from X_1. Let $q_1 = \mathbb{P}[X_1 = 1]$. If $X_1 = 0$, then set $b_1 = 0$. If $X_1 = 1$, then let b_1 be the outcome of a $0/1$ coin that shows up 1 with probability $\frac{p}{q_1}$. Note that $b_1 = 1$ with probability p. Now, for $i = 2, \ldots, s$ do the following: Suppose that we have sampled from $X_1, ..., X_{i-1}$ and thus have formed an $(i-1)$-tuple $(x_1, ..., x_{i-1})$. Let $q_i = \mathbb{P}[X_i = 1 | X_1 = x_1, ..., X_{i-1} = x_{i-1}] \geq p$ and now sample from X_i. If $X_i = 0$, then set $b_i = 0$. If $X_i = 1$, then let b_i be the outcome of a $0/1$ coin that shows up 1 with probability $\frac{p}{q_i}$. Notice again that $b_i = 1$ with probability p and this does not depend on the previous values $b_1, ..., b_{i-1}$, by (4.3). Thus the number of 1's in the vector $b = (b_1, \ldots, b_s)$ is binomially distributed. If the vector b has more than t 1's, then also the vector $(X_1, ..., X_n)$ has more than t 1's and first statement of the lemma follows. As $x_i \geq b_i$, for all $i = 1, \ldots, s$, the second statement is immediate. □

The next result can be proved in a similar way.

Lemma 4.2.28. *Let $p \in (0,1)$ and suppose that $X_i, i = 1, \ldots, s$ are $\{0,1\}$-valued random variables such that $\mathbb{P}[X_1 = 1] \leq p$ and*

$$\mathbb{P}[X_i = 1 | X_1, \ldots, X_{i-1}] \leq p, \text{ for all } i = 2, \ldots, s. \tag{4.3}$$

Then $\Sigma_s := X_1 + \cdots + X_s$ is stochastically smaller than a $B(s,p)$ random variable. Furthermore, it is possible to define random vectors $\mathbf{U} = (U_1, \ldots, U_s)$ and $\mathbf{V} = (V_1, \ldots, V_s)$ on a common probability space so that the law of $(U_1, \ldots, U_s)$ is the same as the law of $(X_1, \ldots, X_s)$, each coordinate of $\mathbf{V}$ is an independent Ber(p) random variable and

$$V_i \geq U_i, \text{ for all } i = 1, \ldots, s, \text{ with probability } 1.$$

4.2.3 Biased Colored Coin Tosses

Take n *biased* coins that are colored with n colors. Let $p \in (0,1)$ be the bias of the coins. Recall that for each array of coins, one can draw its dependency graph, $G = (V, E)$, whose vertex set, V, correspond to the colors and whose edge set, E, correspond to the coins, so that $|V| = |E| = n$. Recall also that G might not be connected and that a toss of the coins gives rise to an orientation on the edges of G. As a consequence, if X_G is the number of different colors after the toss, then $X_G = j$ corresponds to the fact that j vertices have positive in-degree, which in turn means that $n - j$ vertices must have in-degree 0. Note that none of the vertices of zero in-degree can be adjacent. Hence if Z_G is the number of vertices of zero in-degree after a toss then $X_G = n - Z_G$.

Some more notation is needed. For every vertex v of the graph, let P_v be the set of edges incident to v that are oriented towards v with probability p. Denote also by Q_v the set of edges incident to v that are oriented towards v with probability $q := 1 - p$. Set $x_v = |P_v|$ and $y_v = |Q_v|$. Thus $x_v + y_v = \deg(v)$. After each toss, denote by x_v^- the number of edges in P_v that are oriented towards v, and by y_v^- the number of edges in Q_v that are oriented towards v. Without loss of generality we may assume that $0 < p < \frac{1}{2}$. The case $p = \frac{1}{2}$ has already been considered.

In this section we discuss, though not solve entirely, the following.

Problem 4.2.29. *For a fixed positive integer, n, find a graph, G_0, on n vertices and n edges such that X_{G_0} is stochastically larger than X_G, for any other graph G on n vertices and n edges. If such a graph does not exist, find an upper bound on the t-th quantile of X_G, for any G.*

By *t-th quantile* of X_G it is understood any number μ_t such that $\mathbb{P}[X_G \geq \mu_t] \geq \frac{1}{t}$ and $\mathbb{P}[X_G \leq \mu_t] \geq 1 - \frac{1}{t}$. Note that this μ_t might not be unique and that the 2-nd quantile is just the median of X_G.
We fix the integer n for the rest of this section. It will always refer to the number of vertices (= number of edges) of the graph that is under consideration.
Note that if X_G is stochastically larger, then X_G has the largest mean. So in order to make an educated guess on the graph, one might first try to maximize $\mathbb{E}[X_G]$.

Lemma 4.2.30. *The maximum value of $\mathbb{E}[X_G]$ is $n(1 - p + p^2)$. This value is achieved by a set of coins that uses every color twice and every color in this set* *appear exactly once in a p-side of a coin and exactly once in a q-side of some other coin.*

Proof. Fix a graph G and for every $v \in G$ denote by C_v the event that vertex v gets positive in-degree after a toss. Then

$$\mathbb{E}[X_G] = \sum_{v \in G} \mathbb{P}[C_v] = \sum_{v \in G} (1 - (1-p)^{x_v} p^{y_v}).$$

The arithmetic-geometric mean inequality implies that

$$\sum_{v \in G} (1-p)^{x_v} p^{y_v} \geq n \cdot (\prod_{v \in G} (1-p)^{x_v} p^{y_v})^{1/n} = np(1-p),$$

since $\sum_v x_v = \sum_v y_v = n$. We conclude that $\mathbb{E}[X_G] \leq n - np(1-p) = n(1-p+p^2)$. The second statement is immediate. □

Notice that the graph G for which the mean of X_G is maximum is a union of cycles. Note also that the function $f(p) = 1 - p + p^2, p \in (0,1)$ is convex and attains its minimum at $p = \frac{1}{2}$. This means that the maximum mean is minimized when $p = \frac{1}{2}$.

Lemma 4.2.31. *Suppose that G is a (possibly disconnected) graph on n vertices and m edges. Fix some orientation on the edges and let O_G, E_G be the number of odd and even in-degree vertices respectively. Then the parity of E_G equals the parity of $m - n$.*

Proof. See lemma 4.2.1. □

A similar result holds for the vertices v for which y_v is even.

Lemma 4.2.32. *Suppose that G is the (possibly disconnected) graph on n vertices and m edges that corresponds to a set of coins. Let A_G be the number of vertices for which y_v is even and Π_G the number vertices for which y_v is odd. Then the parity of A_G is the same as the parity of $m - n$.*

Proof. Since

$$\sum_v y_v = m$$

it follows that Π_G has the same parity as m. Now the fact that $A_G = n - \Pi_G$ gives the result. □

Lemma 4.2.3 holds true for any oriented graph. We rewrite it here, giving a different proof. Denote by $Q_t(Y)$ a t-th quantile of the random variable Y.

Lemma 4.2.33. *For any oriented graph, G, on n vertices and n edges, we have*

$$Z_G \geq \frac{1}{2} E_G.$$

A lower bound on $Q_t(E_G)$ gives an upper bound on $Q_t(X_G)$. More precisely,

$$Q_t(X_G) \leq n - \frac{1}{2} Q_t(E_G).$$

Proof. Let $Y_G = E_G - Z_G$. For $i = 1, 2, \ldots$, set $I_i := \{v \in G : \deg^-(v) = i\}$. From the in-degree sum formula we have that

$$n = \sum_{v \in G} \deg^-(v) = \sum_{i \geq 1} i|I_i|.$$

In addition, $n = Z_G + \sum_{i \geq 1} |I_i|$. Hence

$$\begin{aligned} n - n &= \sum_{i \geq 1} i|I_i| - \sum_{i \geq 1} |I_i| - Z_G \\ &= \sum_{i \geq 1} (i-1)|I_i| - Z_G \\ &\geq Y_G - Z_G \\ &= E_G - 2Z_G, \end{aligned}$$

which implies that $2Z_G \geq E_G$, thus proving the first statement. From this we can conclude that

$$X_G = n - Z_G \leq n - \frac{1}{2} E_G,$$

and so $Q_t(X_G) \leq n - \frac{1}{2} Q_t(E_G)$, as required. □

The idea behind looking at the number of vertices of even in-degree is the following. Recall that we are interested in determining the graph, G, on n vertices and n edges for which

$$\mathbb{P}[X_G \geq t] \geq \mathbb{P}[X_{G'} \geq t],$$

for all t and all other graphs G' on n vertices and n edges. Since $X_G = n - Z_G$, the problem is equivalent to determining the graph G for which

$$\mathbb{P}[Z_G \geq t] \leq \mathbb{P}[Z_{G'} \geq t],$$

for all t and all other graphs, G'. Thus, equivalently, we may find the graph, G, for which Z_G is stochastically smaller. From the previous lemma we know that $Z_G \geq \frac{1}{2}E_G$, for all graphs G. This means that if we can determine the graph G for which E_G is stochastically smaller and if for this graph $Z_G = \frac{1}{2}E_G$ holds true, then G will also be the graph for which X_G is stochastically larger. Notice that a graph G for which $Z_G = \frac{1}{2}E_G$ is a union of cycles. Moreover, even if the stochastically smaller graph, G, satisfies $Z_G > \frac{1}{2}E_G$, lower bounds on the quantiles of $\frac{1}{2}E_G$ will give upper bounds on the quantiles of $X_{G'}$, for any graph, G', with n vertices and n edges. Furthermore, in case G is connected, one can "estimate" the distribution of E_G, in a way that we make precise in Theorem 4.2.8 below. To do so we will use again the following version of 4.2.7. Recall that a *leaf* in a tree is a vertex of degree 1.

Lemma 4.2.34. *Let T be a tree on n vertices and fix any edge $f \in T$. Then there exists a labeling, $v_1, \ldots, v_n$, of the vertices and a labeling, $e_1, \ldots, e_{n-1}$, of the edges of T such that*
(i) edge f has label e_{n-1};
(ii) the only edge incident to vertex $v_i, i = 1, \ldots, n-1$, among the edges with labels $\{e_i, e_{i+1}, \ldots, e_{n-1}\}$ is the edge with label e_i.

Proof. The statement is clearly true if $n = 2$, so suppose that $n > 2$. Fix a tree, T, on $n > 2$ vertices and choose any of its edges. Label this edge e_{n-1} and label its endpoints v_n and v_{n-1} arbitrarily. Notice that not both v_n and v_{n-1} can be leaves. If v_n or v_{n-1} is a leaf, say v_n, then consider the vertex set L of leaves in T except v_n and label them $v_1, v_2, \ldots, v_\ell$. If v_n is not a leaf, then consider all leaves of T and label them in the same manner. Note that L is not empty even if v_n is a leaf since any tree with at least two vertices has at least two leaves. Now label each edge incident to v_j with e_j, for $j = 1, 2, \ldots, \ell$. Now consider the tree $T' := T \setminus \{v_1, v_2, \ldots, v_\ell\}$ and repeat this process on the leaves of T' again sparing v_n or v_{n-1} if it is a leaf of T'. We keep on labeling the leaves and edges of the subtrees until we end up with the graph consisting of the edge e_{n-1} only. It is evident that the labeling satisfies the required condition. □

Note that we can label any edge of T with e_{n-1} and any endpoint of e_{n-1} with v_n. We will call a labeling on the vertices and edges of a tree, a *good labeling* if it satisfies the conditions of Lemma 4.2.34. Notice also that if we are given a good labeling of a tree and we interchange the labels v_n and v_{n-1} then we get another good labeling of the same tree. We collect this observation in the following.

Lemma 4.2.35. *Let T be a tree on n vertices and fix two adjacent vertices u_1, u_2 of T. Suppose that T has a good labeling such that u_1 has label v_{n-1} and u_2 has* *label v_n. Then the labeling that interchanges the labels of u_1 and u_2 and keep all other labels the same is also a good labeling.*

Note that the previous lemma says that for any edge $f = (u, w)$ of T there is a one-to-one correspondence between good labelings for which u gets the label v_n and w gets label v_{n-1} and good labelings for which u gets the label v_{n-1} and w gets label v_n. We will also need the following observation on the spanning trees of connected graphs.

Lemma 4.2.36. *Suppose that $G = (V, E)$ is a connected graph and fix any edge $e \in E$. Then there exists a spanning tree, T, of G such that e is an edge of T, i.e. $e \in T$.*

Proof. Let $T = (V, E')$ be a spanning tree of G. If $e \in E'$ then we are done, so suppose that $e \notin E'$. This means that if we add e to E' then we create a cycle. Now note that if we delete any edge, $e' \neq e$, from this cycle we get a spanning tree T' of G for which e belongs to T'. □

In the following result we compute the probability that a certain vertex has even in-degree.

Lemma 4.2.37. *If $v \in V$ is such that y_v is even, then*

$$\mathbb{P}[\mathit{deg}^-(v) \text{ } \mathit{even}] = \mathbb{P}[B(\mathit{deg}(v), p) \text{ } \mathit{even}].$$

If $v \in V$ is such that y_v is odd, then

$$\mathbb{P}[\mathit{deg}^-(v) \text{ } \mathit{even}] = \mathbb{P}[B(\mathit{deg}(v), p) \text{ } \mathit{odd}]$$

Proof. We only prove the first equality. The second can be proved similarly. Note that $\deg^-(v)$ is even if and only if both x_v^- and y_v^- are even, or both are odd. Thus $\mathbb{P}[\deg^-(v) \text{ even}]$ equals

$$\mathbb{P}[x_v^- \text{ even}] \cdot \mathbb{P}[y_v^- \text{ even}] + \mathbb{P}[x_v^- \text{ odd}] \cdot \mathbb{P}[y_v^- \text{ odd}]$$

or, equivalently, to

$$\mathbb{P}[B(x_v, p) \text{ even}] \cdot \mathbb{P}[B(y_v, 1-p) \text{ even}] + \mathbb{P}[B(x_v, p) \text{ odd}] \cdot \mathbb{P}[B(y_v, 1-p) \text{ odd}]$$

and thus equal to

$$\frac{1}{2}(1 + (1-2p)^{x_v}) \cdot \frac{1}{2}(1 + (1-2q)^{y_v}) + \frac{1}{2}(1 - (1-2p)^{x_v}) \cdot \frac{1}{2}(1 - (1-2q)^{y_v}).$$

Now from the fact that $q = 1 - p$ and y_v is even we can conclude that the last sum is the same as

$$\frac{1}{2}(1 + (1-2p)^{x_v}) \cdot \frac{1}{2}(1 + (1-2p)^{y_v}) + \frac{1}{2}(1 - (1-2p)^{x_v}) \cdot \frac{1}{2}(1 - (1-2p)^{y_v})$$

which in turn is equal to

$$\frac{1}{2} + \frac{1}{2}(1-2p)^{\deg(v)}$$

and proves the lemma. □

We will also need the following version of lemma 4.2.6. Recall that we assume $p \leq 1/2$.

Lemma 4.2.38. *Fix some vertex v of the graph, fix an edge, e, that is incident to v and let C be the set consisting of all edges edges incident to v except e. Let C^- denote the number of edges from C that are oriented towards v after a toss. Then*

$$\mathbb{P}[deg^-(v) \text{ } even | C^-] \geq p.$$

Proof. Suppose the coins corresponding to C have been flipped. Let C^- be the number of edges in C which are oriented towards v after the toss. Suppose that the edge e corresponds to a coin that is oriented towards v with probability p. The other case is similar. Then

$$\begin{aligned}
\mathbb{P}[\deg^-(v) \text{ even} | C^-] &= (1-p) \cdot \mathbf{1}_{\{C^- \text{ even}\}} + p \cdot \mathbf{1}_{\{C^- \text{ odd}\}} \\
&= p + (1-2p) \cdot \mathbf{1}_{\{C^- \text{ even}\}} \\
&\geq p,
\end{aligned}$$

as required. Note that in case $p = \frac{1}{2}$ the last inequality is in fact equality. □

For every vertex $v \in G$, denote by θ_v the probability that the in-degree of v is even. Note that, by Lemma 4.2.37, θ_v is either $\mathbb{P}[B(\deg(v), p) \text{ even}]$ or $\mathbb{P}[B(\deg(v), p) \text{ odd}]$. Thus $\theta_v \geq p$, for all $v \in V$.

We now have all the necessary tools to "estimate" the distribution of E_G in the case of a *connected* graph G on n vertices and m edges.

Theorem 4.2.39. *Suppose that $G = (V, E)$ is a connected graph on n vertices and $m \geq n - 1$ edges. Let d_v be the degree of vertex v, $\pi_v := \frac{d_v}{2m}$ and π be the probability vector with coordinates $\pi_v, v \in V$. Assume $p < 1 - p$ and let $\{p\}_n$ be the set consisting of n copies of p. Then, if $m - n$ is even, E_G is stochastically larger than a $\mathcal{E}(\{p\}_n, \pi)$ random variable. If $m-n$ is odd, then E_G is stochastically larger than a $\mathcal{O}(\{p\}_n, \pi)$ random variable.*

Proof. Recall that for every edge we toss a coin to decide on its orientation. All these m coins, $c_i, i = 1\dots, m$, are independent. Since the order with which we toss the coins doesn't matter we may, equivalently, suppose that we toss the coins in the following way: we choose a coin, say coin c_i, with probability $\frac{1}{m}$, flip the remaining $m-1$ coins in any way we want and then toss the coin c_i. Tossing this way does not affect the distribution of E_G but allows us to use Lemma 4.2.34. More precisely, we may suppose that once the coin c_i is chosen, then we toss the remaining $m-1$ coins according to a good labeling, $v_1, \dots, v_n; e_1, \dots, e_{n-1}$, of a spanning tree T of G that contains the edge corresponding to c_i, say this edge is $f_i = [u, w]$, and with the good labeling of T chosen in such a way that the edge f_i gets label e_{n-1}; we can use this specific good labeling of T and first toss the coins corresponding to edges that do *not* belong to T in any way we like and then toss the coins that correspond to edges $e_1, \dots, e_{n-1}$ in that specific order. This way the coin c_i is flipped last and we do not affect the distribution of E_G. Note that, by Lemma 4.2.36, there exists a spanning tree, T, of G containing edge f_i and we can always construct a good labeling of T for which f_i gets label e_{n-1}, by Lemma 4.2.34. Furthermore, the edge f_i has two endpoints, u, w, and the probability that vertex u has label v_n equals $1/2$, by Lemma 4.2.35. Since we fix coin c_i with probability $1/m$ it follows that, for every vertex $v \in V$, the probability that we toss the coins according to a good labeling of a spanning tree T of G for which vertex v gets label v_n equals $\frac{d_v}{2m}$.
So let T be a spanning tree of G with a good labeling and recall that we are going to do the following: first we randomly orient the edges that do *not* belong to T and then randomly orient the edges $e_1, e_2, \dots, e_{n-1}$ in that order. Note that the probability that the vertex with label v_1 has even in-degree equals $\theta_{v_1} \geq p$. The fact that T has a good labeling implies that, for $j = 1, \dots, n-1$, once the edge e_j is given an orientation, then the parity of vertex v_j is determined. Lemma 4.2.38 gives that once the parity of vertex v_j is determined, the probability that vertex v_{j+1} has even in-degree is $\geq p$. Only the parity of the vertex with label v_n is deterministic given the parities of the previous vertices. Let δ_i be the indicator of the event $\{\deg^-(v_i)$ is even$\}$, for $i = 1, 2, \dots, n$. Thus $E_G = \delta_1 + \cdots + \delta_n$ and each $\delta_i, i = 1, \dots, n-1$ is stochastically larger than a $B(1, p)$ random variable. From Lemma 4.2.27 we know that there exist random binary vectors $\mathbf{U} = (U_1, \dots, U_{n-1})$ and $\mathbf{V} = (V_1, \dots, V_{n-1})$ defined on a common probability space such that the law of $\mathbf{U}$ is the same as the law of $(\delta_1, \dots, \delta_{n-1})$, each V_i is an independent Bernoulli Ber(p) random variable and

$$\sum_{i=1}^{n-1} U_i \geq \sum_{i=1}^{n-1} V_i \quad \text{with probability } 1.$$

In addition we know that $\sum_{i=1}^{n-1} V_i \sim B(n-1,p)$. To end the proof, suppose that $m-d$ is even. The other case is similar. Thus E_G is even as well and $E_G \sim U_1+\cdots+U_{n-1}+\delta_n$, where $\delta_n=1$ if $U_1+\cdots+U_{n-1}$ is odd and $\delta_n=0$ if $U_1+\cdots+U_{n-1}$ is even. Now let $\gamma_n=1$ if $V_1+\cdots+V_{n-1}$ is odd and $\gamma_n=0$ if $V_1+\cdots+V_{n-1}$ is even, in order to guarantee that $V_1+\cdots+V_{n-1}+\gamma_n$ is always even. Since $U_1+\cdots+U_{n-1} \geq V_1+\cdots+V_{n-1}$ with probability 1, we also have that $U_1+\cdots+U_{n-1}+\delta_n \geq V_1+\cdots+V_{n-1}+\gamma_n$ with probability 1 and the result follows. □

Note that in case $p=\frac{1}{2}$ Lemma 4.2.38 gives that once the parity of vertex v_j is determined, the probability that vertex v_{j+1} has even in-degree is equal to $\frac{1}{2}$, and so the parity of v_{j+1} is *independent* of the parity of $v_1, v_2, \ldots, v_{j-1}$. Only the parity of v_n is deterministic given the parities of the previous vertices. This implies that the random variables $\delta_i, i=1,\ldots,n-1$ in the proof of Theorem 4.2.39 satisfy $\delta_1+\cdots+\delta_{n-1} =_{st} B(n-1,1/2)$ and the following result (which is Theorem 4 in [43]) follows.

Corollary 4.2.40. *Suppose that $p=\frac{1}{2}$. If $m-n$ is even, then E_G has the same distribution as a $A(n,1/2)$ random variable. If $m-n$ is odd, then E_G has the same distribution as a $P(n,1/2)$ random variable.*

Using Lemma 4.2.23 and Lemma 4.2.33 we have the following result on X_G, in case G is connected.

Corollary 4.2.41. *Let G be a connected loop-less multi-graph on n vertices and n edges. Then a median of X_G is $\leq n-\frac{1}{2}(n-1)p+\frac{1}{2}$.*

We end this subsection by proving the following result on a median of X_G.

Theorem 4.2.42. *For any graph G on n vertices and n edges, a median of X_G is $\leq n-\frac{p^2}{1+(1-2p)^2}n+\frac{3}{4}$.*

Recall that the dependency graph $G=(V,E)$ of the colored coins might not be connected. Suppose it consists of t connected components, $G_1,\ldots,G_t$, each having n_i vertices and m_i edges such that $\sum n_i = n$ and $\sum m_i = n$. Let also E_{G_i} be the number of vertices of even in-degree in each component, after a toss. Hence the total number of vertices of even in-degree after a toss, E_G is equal to the independent sum $E_{G_1}+\cdots+E_{G_t}$. As $|V|=|E|=n$, it follow from Lemma 4.2.1 that E_G is even. By Theorem 4.2.39, the distribution of each E_{G_i} is stochastically larger than a $A(\cdot,p)$ or $P(\cdot,p)$ random variable. More precisely, suppose that the first t_1 components of G correspond to a $A(\cdot,p)$ random variable and the remaining t_2 components correspond to a $P(\cdot,p)$ random variable, so that $t_1+t_2=t$ and t_2 is even. Let $\{p\}_k$ denote the set consisting of k parameters that are all equal to p. From Theorem 4.2.39 we know that

$$E_{G_i} \geq_{st} A(n_i, p), \text{ for } i = 1, \dots, t_1$$

and

$$E_{G_i} \geq_{st} P(n_i, p), \text{for } i = t_1 + 1, \dots, t.$$

Hence, the total number of even in-degree vertices, E_G, is stochastically larger than the independent sum

$$\sum_{i=1}^{t_1} A(n_i, p) + \sum_{i=t_1+1}^{t} P(n_i, p).$$

Since $p \in (0, 1/2]$ we have $\beta(\{p\}_{n_i-1}) \geq p$ and $\alpha(\{p\}_{n_i-1}) \geq p$ and thus Corollary 4.2.22 implies that

$$\begin{aligned} A(n_i, p) &\sim B(1, \beta(\{p\}_{n_i-1})) + B(n_i - 1, p, B(1, \beta(\{p\}_{n_i-1}))) \\ &\geq_{st} B(1, p) + B(n_i - 1, p, B(1, p)) \end{aligned}$$

and

$$\begin{aligned} P(n_i, p) &\sim B(1, \alpha(\{p\}_{n_i-1})) + B(n_i - 1, p, 1 - B(1, \alpha(\{p\}_{n_i-1}))) \\ &\geq_{st} B(1, p) + B(n_i - 1, p, 1 - B(1, p)) \end{aligned}$$

and so E_G is stochastically larger than

$$\sum_{i=1}^{t_1} B(1, p) + B(n_i - 1, p, B(1, p)) + \sum_{i=t_1+1}^{t} B(1, p) + B(n_i - 1, p, 1 - B(1, p))$$

This independent sum takes even values (recall t_2 is even) and, by Lemma 4.2.16 and Lemma 4.2.17, is equivalently described as follows. Toss t independent $0/1$ coins, $c_i, i = 1, \dots, t$, each having probability p of landing on 1. Let $\Gamma = (\gamma_1, \dots, \gamma_t) \in \{0, 1\}^t$ be a particular outcome of the toss. This is a binary vector of length t. If B_Γ is the number of 1's in this vector, then add B_Γ to the outcome of the independent sum

$$\mathcal{H}|\Gamma := \sum_{i=1}^{t_1} B(n_i - 1, p, \gamma_i) + \sum_{i=t_1+1}^{t} B(n_i - 1, p, 1 - \gamma_i),$$

thus forming the sum $B_\Gamma + \mathcal{H}|\Gamma$. Note that $B_\Gamma \sim B(t, p)$. Now each particular vector Γ can be equivalently obtained in the following way. Fist toss a coin with probability of success $\frac{1}{2}(1+(1-2p)^t)$. If the outcome is a success, then arrange t independent $0/1$ coins (whose probability of landing on 1 equals p) on a line and toss them until you see an *even* number of 1's. If Γ_e is the resulting binary vector and B_e is the number of 1's in Γ_e, then $B_e \sim B(t,p,0)$ and $B_\Gamma + \mathcal{H}|\Gamma$ equals $B_e + \mathcal{H}|\Gamma_e$ with probability $\frac{1}{2}(1+(1-2p)^t)$. If the outcome is a failure, then toss t independent $0/1$ coins until you see an *odd* number of 1's. If Γ_o is the resulting binary vector and B_o is the number of 1's in Γ_o, then $B_o \sim B(t,p,1)$ and $B_\Gamma + \mathcal{H}|\Gamma$ equals $B_o + \mathcal{H}|\Gamma_o$ with probability $\frac{1}{2}(1-(1-2p)^t)$. Hence $B_\Gamma + \mathcal{H}|\Gamma$ is a mixture of the sums $B_e + \mathcal{H}|\Gamma_e$ and $B_o + \mathcal{H}|\Gamma_o$.

Lemma 4.2.43. *A median of* $B_\Gamma + \mathcal{H}|\Gamma$ *is* $\geq n\bar{p} - \frac{3}{2}$*, where* $\bar{p} := \frac{2p^2}{1+(1-2p)^2}$.

Proof. First toss a coin to decide whether you take a vector, Γ_e, with an even number of 1's or a vector, Γ_o, with an odd number of 1's. Suppose that we end up with a vector Γ_e. The other case is similar. This vector gives rise to the sum $B_e + \mathcal{H}|\Gamma_e$. Then $B_e \sim B(t,p,0)$ and each term in $\mathcal{H}|\Gamma_e$ is of the form $B(n_i-1,p,0)$ or $B(n_i-1,p,1)$. Apply lemmata 4.2.18 and 4.2.19 repeatedly to write each term of the sum $B_e + \mathcal{H}|\Gamma_e$ as a mixture of independent sums consisting only of terms $\mathcal{H}(J,0)$ and $\mathcal{H}(J,1)$ for which $|J|$ equals 2 or 3. Thus the initial sum, $B_e + \mathcal{H}|\Gamma_e$, is a mixture of independent sums of terms $\mathcal{H}(J,0)$ and $\mathcal{H}(J,1)$ for which $|J|$ equals 2 or 3. Suppose that Ξ is a particular independent sum consisting of a terms of the form $B(2,p,0)$, b terms of the form $B(2,p,1)$, c terms of the form $B(3,p,0)$ and d terms of the form $B(3,p,1)$. Thus $2a+2b+3c+3d=n$. Lemma 4.2.20 implies that

$$B(2,p,0) \sim 2 \cdot B\left(1,\bar{p}\right),\ B(2,p,1) \sim B(1,1),$$

where $\bar{p} = \frac{2p^2}{1+(1-2p)^2}$, and that

$$B(3,p,0) \sim 2 \cdot B\left(1,\hat{p}\right),\ B(3,p,1) \sim 1 + 2 \cdot B\left(1,\tilde{p}\right),$$

where $\hat{p} = \frac{6p^2(1-p)}{1+(1-2p)^3}$ and $\tilde{p} = \frac{2p^3}{1-(1-2p)^3}$. Denote

$$\Psi := B(a,\bar{p}) + B(c,\hat{p}) + B(d,\tilde{p}).$$

Then $\Xi = 2\Psi + b + d$ and so a median of Ξ can be estimated via a median of Ψ. Hence a median of Ξ is $\geq n\bar{p} - \frac{3}{2}$ if and only if a median of Ψ is $\geq \frac{n\bar{p}-b-d}{2} - \frac{3}{4}$. Using the fact the $2a+2b+3c+3d=n$ we can write

$$\frac{n\bar{p}-b-d}{2} = a\bar{p} + b\left(\bar{p} - \frac{1}{2}\right) + c\frac{3\bar{p}}{2} + d\left(\frac{3\bar{p}}{2} - \frac{1}{2}\right) := \pi_*.$$

Note that $\bar{p} - \frac{1}{2} \le 0$. As $0 \le p \le 1/2$, elementary calculations show that $\hat{p} \ge \frac{3\bar{p}}{2}$ and $\tilde{p} \ge \frac{3\bar{p}}{2} - \frac{1}{2}$. This implies that

$$\mathbb{E}[\Psi] = a\bar{p} + c\hat{p} + d\tilde{p} \ge \pi_*.$$

From Hoeffding's result (Theorem 4.2.24) we know that

$$\mathbb{P}\left[\Psi \ge \pi_* - \frac{3}{4}\right] \ge \mathbb{P}\left[B(a+c+d, p_0) \ge \pi_* - \frac{3}{4}\right],$$

where $p_0 = \frac{1}{a+c+d}(a\bar{p} + c\hat{p} + d\tilde{p})$ and so it is enough to show that a median of a $B(a+c+d, p_0)$ random variable is $\ge \pi_* - \frac{3}{4}$. Now, it is well known (see [24]) that the smallest uniform (with respect to both parameters) distance between the mean and a median of a binomial distribution is $\le \ln 2 < \frac{3}{4}$. This means that a median of $B(a+c+d, p_0)$ is $\ge a\bar{p} + c\hat{p} + d\tilde{p} - \frac{3}{4} \ge \pi_* - \frac{3}{4}$ and the lemma follows. □

The proof of Theorem 4.2.42 is almost complete.

Proof of Theorem 4.2.42. Since E_G is stochastically larger than $B_\Gamma + \mathcal{H}|\Gamma$ and a median of $B_\Gamma + \mathcal{H}|\Gamma$ is $\ge n\bar{p} - \frac{3}{2}$, we conclude that the median of E_G is $\ge n\bar{p} - \frac{3}{2}$. Theorem 4.2.42 follows since, from Lemma 4.2.33, we have

$$\text{Med}(X_G) \le n - \frac{1}{2}\text{Med}(E_G) \le n - \frac{n}{2}\bar{p} + \frac{3}{4}.$$

□

We end this section by mentioning that our method works also in case one is interested in estimating X_G from below. Since $X_G \ge O_G = n - E_G$, for all graphs G it is enough to estimate the probability distribution of E_G from above, i.e., to find a random variable that is stochastically larger than E_G. Now we know that $\theta_v \le 1 - p$, for all $v \in V$ and a modification of the proof of Theorem 4.2.39 along with Lemma 4.2.28 shows that the following is true.

Theorem 4.2.44. *Suppose that $G = (V, E)$ is a connected multi-graph on n vertices and $m \ge n - 1$ edges. Let d_v be the degree of vertex v, set $\pi_v := \frac{d_v}{2m}$ and let π be the probability vector with coordinates $\pi_v, v \in V$. Assume $p < 1 - p$ and let $\{1-p\}_n$ be the set consisting of n copies of $1 - p$. Then, if $m - n$ is even, E_G is stochastically smaller than a $\mathcal{E}(\{1-p\}_n, \pi)$ random variable. If $m - n$ is odd, then E_G is stochastically smaller than a $\mathcal{O}(\{1-p\}_n, \pi)$ random variable.*

4.2.4 Random graphs

In this section we apply our method to the distribution of the number of vertices with odd degree in random sub-graphs of fixed graphs. More precisely, let G be any *connected graph* on n vertices and for each edge of G toss a coin that shows up tails with probability p, independently for all edges. If the result of the coin is tails, then keep the edge. If the result is heads, delete the edge. The distribution of the vertex degree in such models has been well studied (see [9] for a whole chapter on this topic). The resulting sub-graph of G that remains after the toss of the coins is random. Let $q = 1 - p$ and denote by $O_{n,p}(G)$ the number of vertices of *odd degree* in the resulting graph. The following folds true.

Theorem 4.2.45. *If $0 \leq p \leq \frac{1}{2}$ then the random variable $O_{n,p}(G)$ is stochastically larger than a $A(n,p)$ random variable. If $\frac{1}{2} \leq p \leq 1$, then $O_{n,p}(G)$ is stochastically larger than a $A(n,q)$ random variable.*

Proof. The proof is similar to the proof of Theorem 4.2.39, so we only sketch it. Suppose that $0 \leq p \leq \frac{1}{2}$. The other case is similar. Let T be a spanning subgraph of G with a good labeling, $v_1, \ldots, v_n; e_1, \ldots, e_{n-1}$ on its vertices and edges given by Lemma 4.2.34. By Lemma 4.2.15 we know that the probability that $d_i := \deg(v_i)$ is odd is equal to $\frac{1}{2}(1 - (1-2p)^{d_i})$, for $i = 1, \ldots, n$. Toss all coins to decide which edges are included in the sub-graph, except the coins corresponding to the edges $e_i, i = 1, \ldots, n-1$. Now begin from vertex v_1 and toss a coin to decide whether edge e_1 is included or not. Then proceed to vertex v_2 and toss a coin to decide on the edge e_2, and in general, at step $j, j = 1, \ldots, n-1$ move from vertex $j-1$ to vertex j and toss a coin to decide if edge e_j is included or not. Let C_j be the set of edges that are included in the graph and are incident to v_j at step $j-1$. As in Lemma 4.2.38, by conditioning on whether $|C_j|$ is even or odd we conclude that

$$\mathbb{P}[\deg(v_j) \text{ odd}|C_j] \geq p,$$

Hence the random variable $\mathbf{1}_{\{\deg(v_j) \text{ odd}\}}, j = 1, \ldots, n-1$, is stochastically larger than a $B(1,p)$ random variable. Only the parity of vertex v_n is deterministic, given the parities of the previous vertices. The result follows from the fact that the degree-sum formula implies that $O_{n,p}$ is even. $\square$

Notice that in case $p = \frac{1}{2}$ we obtain the following result.

Corollary 4.2.46. *If $p = \frac{1}{2}$ then, for any connected graph G, $O_{n,1/2}(G)$ has the same distribution as a $A(n, 1/2)$ random variable.*

4.2.5 Some applications

Let $G = (V, E)$ be a connected undirected graph and fix $T \subseteq V$. An orientation of G, is an assignment of direction to each edge of G. An orientation of G is called T-odd if the vertices in T are the only ones having odd in-degree. We allow T to be the empty set in which case $\emptyset$-odd orientation simply means that all vertices of G have even in-degree. The following result is obtained in [19], using induction.

Lemma 4.2.47. *A connected graph, $G = (V, E)$, on n vertices and m edges has a T-odd orientation if and only if $|T| + |E|$ is even.*

Proof. Suppose first that G has a T-odd orientation. Let E_G be the number of even in-degree vertices, O_G the number of odd in-degree vertices. From Lemma 4.2.1 we know that $E_G \equiv m - n \bmod 2$ and $O_G \equiv m \bmod 2$. This implies that $O_G = |T| \equiv m = |E| \bmod 2$ and so $|T| \equiv |E| \bmod 2$, which is equivalent to $|T| + |E|$ is even.
On the other hand, fix some set of vertices T such that $|T| \equiv |E| \bmod 2$ and consider a random orientation on G obtained by directing each edge in G independently of the others and with probability $\frac{1}{2}$ in each direction. Let E_G, O_G be as above. We prove that there is a positive probability that the vertices of T are the only ones having odd degree. Since $E_G \equiv m - n \equiv |T| - n \bmod 2$ it follows that $n - |T|$ belongs to the range of E_G. The result follows from Corollary 4.2.40, since $\mathbb{P}[E_G = n - |T|] = \frac{1}{2^{n-1}}\binom{n}{n-|T|} > 0$, and from the fact that any set, T, of $|T| \equiv |E| \bmod 2$ vertices can be such that all vertices in T have odd in-degree. □

We can also deduce a result on enumeration of oriented graphs.

Lemma 4.2.48. *Let $G = (V, E)$ be a graph on n vertices and m edges. Then the number of orientations on the edges of G for which there are exactly t vertices of even in-degree equals $2^{m-n+1}\binom{n}{t}$.*

Proof. Let ν_t be the number of orientations of G having exactly t vertices of even in-degree. Note that t has to be such that $t \equiv m - n \bmod 2$. From the set of all possible orientations of G, choose one uniformly at random and let A_t be the event that the orientation has t vertices of even in-degree. Then

$$\mathbb{P}[A] = \frac{\nu_t}{2^m}.$$

Now consider a random orientation on the edges of G by directing each edge in G independently of the others and with probability $\frac{1}{2}$ in each direction. The result follows since, by Corollary 4.2.40, the probability that there are t vertices of even in-degree equals $\frac{1}{2^{n-1}}\binom{n}{t}$. □

For similar results see [49]. In a similar way, using Corollary 4.2.46, one can obtain a result on enumeration of labeled graphs. We leave the details to the reader.

Lemma 4.2.49. *The number of labeled graphs on n vertices for which there are exactly t (where t is even) vertices of odd degree equals $2^{m-n+1}\binom{n}{t}$.*

Note that the case $t = 0$ of the previous result appears as problem 16 in $\S 5$ of [36].

4.2.6 A question on unimodality and a related conjecture

Suppose that we are given a set of m *fair* coins that are colored with n different colors, where $m \geq n-1$ and let let G be it's dependency graph. Assume that G is connected and denote by X_G the number of different colors that we see after a toss. We saw in the previous sections that every toss of the coins gives rise to an orientation on the edges of G by directing each edge $e = (uv)$ towards the color that occurred after the toss. Thus after each toss, the edges of G get (random) directions. This means that if $X_G = j$, then $n - j$ vertices of the graph must have in-degree 0. This observation gives a way to describe the distribution of X_G via independent sets of vertices of G. By an *independent set* of a graph it is meant a set of vertices no two of which are adjacent. Suppose $\alpha := \alpha(G)$ is the number of vertices in a maximum independent set of G, i.e., the *independence number* of the graph. If all vertices of a maximum independent have zero in-degree after a toss, then all other vertices (colors) have positive in-degree. This means that the smallest value that X_G can achieve is $n - \alpha$ and so $X_G \in \{n-\alpha, \ldots, n\}$. Now we ask the following.

Question: For which graphs, G, is the distribution of X_G unimodal?

The distribution of X_G is related to the collection of independent sets in G. If $X_G = j$, then $n-j$ vertices have in-degree zero and these $n-j$ vertices form an independent set. That is, $X_G = j$ gives rise to an independent set of vertices in G of cardinality $n-j$ and, in general, there will be many (different) independent sets of cardinality $n-j$. So we might also ask the following.

Question: For $j = 0, 1, \ldots, n$, denote by $\alpha_j(G)$ the number of independent set of vertices of G of cardinality j. Is the sequence $\{\alpha_i(G)\}_{j=0}^{n}$ unimodal?

This problem is considered in [1] where it is proven that the answer to the last question is no, for general graphs. However, it remains an open question to determine whether the question is true in the case of trees. In [1] one can find the following.

Conjecture 4.2.50 (Alavi, Erdős, Malde, Schwenk, 1987)**.** *If G is a tree, then the independent set sequence $\{\alpha_i(G)\}_{j=0}^{n}$ is unimodal.*

Our conjecture is that the same is true for the distribution of X_G.

Conjecture 4.2.51. *If G is a tree, then the distribution of X_G is unimodal.*

Appendix A

Erdős-Ko-Rado Theorem

Erdős-Ko-Rado theorem is well known. It has been proved and generalized in many different ways. See [21] for a collection of eight different proofs. The latest proof appears in [20]. In this appendix we provide a proof of EKR that uses the idea of Katona (see [29]) on cyclic permutations. Our proof only differs from Katona's proof in the counting argument that provides the bound on the probability that a randomly chosen s-subset belongs to a given intersecting family.

Let $[n] := \{1, \ldots, n\}$ and denote by $\binom{[n]}{s}$ the set of all s-subsets of $[n]$. A family $\mathcal{F} \subseteq \binom{[n]}{s}$ is called *intersecting* if $F \cap F' \neq \emptyset$ holds for all $F, F' \in \mathcal{F}$.

Theorem A.1 Suppose that $n \geq 2s$ and $\mathcal{F} \subseteq \binom{[n]}{s}$ is an intersecting family. Then

$$|\mathcal{F}| \leq \binom{n-1}{s-1}.$$

Proof. Write $n = ks + i$, where $k \geq 2$ and $i \in \{0, 1, \ldots, s-1\}$. We prove that the probability, π, that a randomly chosen s-set belongs to $\mathcal{F}$ is at most $\frac{s}{n}$. The result then follows since

$$\pi = \frac{|\mathcal{F}|}{\binom{n}{s}} \leq \frac{s}{n}.$$

Now choosing an s-set uniformly at random is equivalent to cyclically arranging the n vertices, uniformly at random, and then choosing an interval of length s from that cyclic arrangement, uniformly at random. We claim that, for any such cyclic arrangement, there are i intervals of length s on the circle that do not belong to $\mathcal{F}$. To see this, suppose that there are at most $i-1$ intervals of length s that do not belong to $\mathcal{F}$. This means that

there are at least $ks+1$ remaining intervals all belonging to $\mathcal{F}$ and so form an intersecting family. For all pairs I, J of the remaining intervals consider the distance between the anti-clockwise endpoint of I and the clockwise endpoint of J. The fact that this family is intersecting implies that the maximum of these distances is $\leq 2s-1$. Hence these intervals consist of at most $2s-1$ vertices and so $ks+1 \leq 2s-1$, a contradiction.
We finish the proof by showing that, for any cyclic arrangement of the vertices, there are at most s intervals of length s that belong to $\mathcal{F}$. So fix a cyclic arrangement. We may identify the vertices in this arrangement with $\mathbb{Z}_n$ and identify every interval of length s on this circle with its anti-clockwise endpoint. Let $I_{m_1}, \ldots, I_{m_i}$ be i intervals that do not belong to $\mathcal{F}$. Choose an element $a_1 \in \mathbb{Z}_n \setminus \{m_1, \ldots, m_i\}$ and consider the set $V := \{a_1, \ldots, a_{ks}\}$, where, for $j = 2, \ldots, ks$, a_j is the first element of $\mathbb{Z}_n \setminus \{m_1, \ldots, m_i\}$ that is located after a_{j-1} in the clockwise motion. Then

$$I_{a_1} \cup I_{a_{s+1}} \cup \cdots I_{a_{(k-1)s+1}}$$

is a disjoint union and so at most one of these intervals belongs to $\mathcal{F}$. The same argument applies to

$$I_{a_j} \cup I_{a_{s+j}} \cup \cdots I_{a_{(k-1)s+j}} \text{ for } 1 \leq j \leq s$$

and so at most s of the intervals $I_{a_{j+\ell s}}$ belong to $\mathcal{F}$, where $1 \leq j \leq s$ and $0 \leq \ell \leq k-1$. For the given restrictions on j and ℓ we find all intervals of $\mathbb{Z}_n$ except $I_{m_1}, \ldots, I_{m_i}$. But these intervals do not belong to $\mathcal{F}$. $\square$

Appendix B

Sum of variances of order statistics

This appendix is devoted to the proof of Lemma 4.1.8, i.e. to the inequality

$$\sum_{i=1}^{n} \text{Var}(X_{(i)}) \leq \sum_{i=1}^{n} \text{Var}(X_i),$$

where $X_1, \ldots, X_n$ are (possibly dependent) non-negative random variables, of *finite* mean, and $X_{(1)}, \ldots, X_{(n)}$ is their order statistics. We will prove the inequality by induction on n. The base case is $n = 2$. Fix two random variables, X_1, X_2. Let $m = \min\{X_1, X_2\}$ and $M = \max\{X_1, X_2\}$ and note that $m + M = X_1 + X_2$ and $m \cdot M = X_1 \cdot X_2$. Thus $\text{Var}(m + M) = \text{Var}(X_1 + X_2)$ and so

$$\text{Var}(m) + \text{Var}(M) + \text{Cov}(m, M) = \text{Var}(X_1) + \text{Var}(X_2) + \text{Cov}(X_1, X_2).$$

Hence, in order to prove that $\text{Var}(m) + \text{Var}(M) \leq \text{Var}(X_1) + \text{Var}(X_2)$, it suffices to show that $\text{Cov}(m, M) \geq \text{Cov}(X_1, X_2)$, or, equivalently, that

$$\mathbb{E}[m \cdot M] - \mathbb{E}[m] \cdot \mathbb{E}[M] \geq \mathbb{E}[X_1 \cdot X_2] - \mathbb{E}[X_1] \cdot \mathbb{E}[X_2].$$

Now the fact that $m \cdot M = X_1 \cdot X_2$ implies $\mathbb{E}[m \cdot M] = \mathbb{E}[X_1 \cdot X_2]$ and so it is enough to prove that

$$\mathbb{E}[X_1] \cdot \mathbb{E}[X_2] \geq \mathbb{E}[m] \cdot \mathbb{E}[M].$$

To end the proof of the base case, we compute

$$\begin{aligned}
\mathbb{E}[X_1]\cdot\mathbb{E}[X_2]-\mathbb{E}[m]\cdot\mathbb{E}[M] &= \mathbb{E}[X_1]\cdot\mathbb{E}[X_2]-\mathbb{E}[m]\cdot\mathbb{E}[X_2]\\
&+ \mathbb{E}[m]\cdot\mathbb{E}[X_2]-\mathbb{E}[m]\cdot\mathbb{E}[M]\\
&= \mathbb{E}[X_2]\cdot(\mathbb{E}[X_1]-\mathbb{E}[m])\\
&+ \mathbb{E}[m]\cdot(\mathbb{E}[X_2]-\mathbb{E}[M])\\
&\geq \mathbb{E}[m]\cdot(\mathbb{E}[X_1]+\mathbb{E}[X_2]-\mathbb{E}[m]-\mathbb{E}[M])\\
&= 0,
\end{aligned}$$

where the inequality follows from the fact $X_2 \geq m$ and the last line from the fact that $m+M = X_1+X_2$.
We now proceed to the general case. Suppose that $X_1,\dots,X_n$ are n random variables. We begin with the sum $\sum_{i=1}^n \text{Var}(X_i)$. From the case $n=2$ we know that if we replace two terms of this sum, say $\text{Var}(X_i)+\text{Var}(X_j)$, by $\text{Var}(m_{i,j})+\text{Var}(M_{i,j})$, where $m_{i,j} := \min\{X_i, X_j\}$ and $M_{i,j} := \max\{X_i, X_j\}$, then the sum does not increase. We iterate this procedure by applying the base case to the resulting sum. After a finite number of steps we will end up with a sum in which the variances are ordered. The result follows.

Note that in case all random variables, $X_i, i = 1,\dots,n$, have the same distribution then the inequality reduces to

$$\sum_{i=1}^n \text{Var}(X_{(i)}) \leq n\text{Var}(X)$$

which implies that

$$\text{Var}(X_{(i)}) \leq n\cdot\text{Var}(X), \text{ for all } i\in\{1,\dots,n\}.$$

The following question arises naturally. Fix $n\in\mathbb{Z}_{>0}$, $i\in\{1,\dots,n\}$ and a random variable X. What is the smallest constant $c_{n,i}$ such that

$$\text{Var}(X_{(i)}) \leq c_{n,i}\cdot\text{Var}(X)\ ?$$

This problem is solved in [40], under the assumption of independence. The general question was answered, quite recently, in [44].

Appendix C

A generalized notion of hypergraph matchings

In this appendix we generalize the notions of matchings and covers in uniform hypergraphs. We define their fractional analogues and prove a theorem that relates all these notions.

Let $\mathcal{H} = (V, \mathcal{E})$ be a hypergraph where $|V| = n$ and suppose further that $\mathcal{H}$ is s-uniform, that is, each $E \in \mathcal{E}$ satisfies $|\mathcal{E}| = s$. Recall the notion of the i-th shadow of $\mathcal{E}$ defined as

$$\partial_i(\mathcal{E}) := \{T \subseteq V : |T| = i,\ T \subseteq E \text{ for some } E \in \mathcal{E}\}.$$

For $i = 1, \dots, s$ and $k \in \mathbb{Z}_{>0}$, a (k,i)-*matching* of $\mathcal{H}$ is a collection of edges $\mathcal{E}_0 \subseteq \mathcal{E}$ (the same edge may occur more than once) such that each $T \in \partial_i(\mathcal{E})$ belongs to at most k edges in $\mathcal{E}_0$. The maximum cardinality of a (k,i)-matching is denoted by $\nu_k(i, \mathcal{H})$. Thus $\nu_1(1, \mathcal{H}) = \nu(\mathcal{H})$ is the maximum number of disjoint edges in $\mathcal{H}$ or, in short, the matching number of $\mathcal{H}$. A (k,i)-matching is called *simple* if no edge occurs in it more than once. We denote by $\tilde{\nu}_k(i, \mathcal{H})$ the maximum number of edges in simple (k,i)-matchings of $\mathcal{H}$.

A (k,i)-*cover* of $\mathcal{H}$ is a collection $\mathcal{T} \subseteq \partial_i(\mathcal{E})$ (the same set may occur more than once) such that any edge $E \in \mathcal{E}$ contains at least k sets from $\mathcal{T}$. The minimum number of sets in a (k,i)-cover is denoted by $\tau_k(i, \mathcal{H})$. Hence $\tau_1(1, \mathcal{H}) = \tau(\mathcal{H})$ is the covering number of $\mathcal{H}$.

Given $T \subseteq V$, denote by $N[T]$ the neighbor of T, i.e., $N[T] := \{E \in \mathcal{E} : T \subseteq E\}$. A *fractional i-matching* of $\mathcal{H}$ is a function $\omega : \mathcal{E} \to [0,1]$ such that

$$\sum_{E \in N[T]} \omega(E) \leq 1, \text{ for all } T \in \partial_i(\mathcal{E}).$$

Set $\nu^*(i,\mathcal{H}) := \max_\omega \sum_E \omega(E)$, where the maximum is over all fractional i-matchings of $\mathcal{H}$.
A *fractional i-cover* is a function $w : \partial_i(\mathcal{E}) \to [0,1]$ such that

$$\sum_{T\in\partial_i(\mathcal{E})\cap N[E]} w(T) \geq 1, \text{ for all } E \in \mathcal{E}.$$

Let $\tau^*(i,\mathcal{H}) := \min_w \sum_{T\in\partial_i(\mathcal{E})} w(T)$, where the minimum is over all fractional i-covers of $\mathcal{H}$.

Theorem B.1 Every s-uniform hypergraph $\mathcal{H} = (V,\mathcal{E})$ satisfies

$$\begin{aligned}\nu_1(i,\mathcal{H}) &= \min_k \frac{\nu_k(i,\mathcal{H})}{k} \leq \max_k \frac{\nu_k(i,\mathcal{H})}{k} \leq \nu^*(i,\mathcal{H}) \\ &= \tau^*(i,\mathcal{H}) \leq \min_k \frac{\tau_k(i,\mathcal{H})}{k} \leq \max_k \frac{\tau_k(i,\mathcal{H})}{k} = \tau_1(i,\mathcal{H}).\end{aligned}$$

Proof. Suppose that $\mathcal{E}_0$ is a $(1,i)$-matching of $\mathcal{H}$. Then, for any k, the collection $\mathcal{E}_0'$ obtained from $\mathcal{E}_0$ by taking every edge k times is an (k,i)-matching. Thus $\nu_k(i,\mathcal{H}) \geq k\nu_1(i,\mathcal{H})$. If $k=1$, then the last inequality is in fact equality and so

$$\nu_1(i,\mathcal{H}) = \min_k \frac{\nu_k(i,\mathcal{H})}{k}.$$

Now let $\mathcal{E}_1$ be a maximum (k,i)-matching of $\mathcal{H}$. For every $E \in \mathcal{E}$ let $m(E)$ denote the number of times that E appears in $\mathcal{E}_1$ and define the function $\omega : \mathcal{E} \to [0,1]$ by setting $\omega(E) := \frac{m(E)}{k}$. The fact that $\mathcal{E}_1$ is a (k,i)-matching implies that $\omega(\cdot)$ is a fractional i-matching of $\mathcal{H}$ and thus

$$\max_k \frac{\nu_k(i,\mathcal{H})}{k} \leq \nu^*(i,\mathcal{H}).$$

The equality $\nu^*(i,\mathcal{H}) = \tau^*(i,\mathcal{H})$ comes from the duality theorem of Linear Programing. Now let $\mathcal{T}$ be a (k,i)-cover of $\mathcal{H}$ and define $w : \partial_i(\mathcal{E}) \to [0,1]$ by $w(T) = \frac{1}{k}$ if $T \in \mathcal{T}$, and $w(T) = 0$ otherwise. The fact that $\mathcal{T}$ is a (k,i)-cover implies that $w(\cdot)$ is a fractional i-cover of $\mathcal{H}$ and so

$$\tau^*(i,\mathcal{H}) \leq \min_k \frac{\tau_k(i,\mathcal{H})}{k}.$$

Finally, let $\mathcal{T}_0$ be a $(1,i)$-cover of $\mathcal{H}$ and denote by $\mathcal{T}_0'$ the family obtained from $\mathcal{T}_0$ by taking every of its sets k times. Then $\mathcal{T}_0'$ is a (k,i)-cover and so $k\tau_1(i,\mathcal{H}) \geq \tau_k(i,\mathcal{H})$. This implies that

$$\max_k \frac{\tau_k(i,\mathcal{H})}{k} \leq \tau_1(i,\mathcal{H})$$

and the theorem follows. $\square$

Bibliography

[1] Y. Alavi, P. Erdős, P. Malde and A. Schwenk, The vertex independence sequence of a graph is not constrained, *Congr. Numer.*, 58, p 15–23, (1987).

[2] N. Alon, P. Frankl, H. Huang, V. Rödl, A. Ruciński, B. Sudakov, *Large matchings in uniform hypergraphs and the conjectures of Erdös and Samuels*, J. Combin. Theory, Series A, 119, p. 1200–1215, (2012).

[3] S. Alpern, R. Fokkink, K. Kikuta, *On Ruckle's conjecture on accumulation games*, SIAM J. Control Opt., (48)8, p. 5073–5083, (2010).

[4] S. Alpern, R. J. Fokkink and C. Pelekis, *A proof of the Kikuta-Ruckle conjecture on cyclic caching of resources*, J. Optim. Theory Appl., 153, p. 650–661, (2012).

[5] S. Alpern and S. Gal, *The theory of Search Games and Rendezvous*, Springer, Berlin (2003).

[6] S. Alpern, D.J. Reyniers, *Spatial Dispersion as a Dynamic Coordination Problem*, Theory and Decision 53, p. 29–59, (2002).

[7] C. Berge, *Hypergraphs: combinatorics of finite sets*, North Holland, Amsterdam, (1973).

[8] N. Biggs, *Some odd graph theory*, Annals New York Academy of Sciences, 319 no 1, 71–81 (1979).

[9] Béla Bollobás, *Random Graphs*, Cambridge Studies in Advanced Mathematics, **73**, Cambridge University Press, 2001.

[10] K. Chaudhuri, F. Chung Graham , M. Shoaib Jamall, *A Network Coloring Game*, WINE '08, p. 522–530, (2008).

[11] F. Chung Graham, *Graph Theory in the Information Age*, Notices of AMS, 57, no. 6, p. 726–732, (July 2010).

[12] K.P. Choi *On the medians of gamma distributions and an equation of Ramanujan*, Proc. Amer. Math. Soc., 121, no. 1, p. 245–251, (1994).

[13] F. N. David, *Games, Gods and Gambling*, Griffin publishing company, London, (1962).

[14] B. Eisenberg, *On the Expectation of the Maximum of IID Geometric Random Variables*, Statistics & Probability Letters, **78**, p. 135–143, (2008).

[15] P. Erdős, *A problem on independent r-tuples*, Ann. Univ. Sci. Budapest Eőtvős Sect. Math. 8, p. 93–95, (1965)

[16] B. Escoffier, L. Gourvés, J. Monnot, *Strategic Coloring of a Graph*, Algorithms and Complexity, Lecture Notes in Computer Science, 6087, Springer, Berlin, p.155-166 (2010)

[17] T.S. Ferguson: *Game Theory*, part II, www.math.ucla.edu/tom/

[18] R. J. Fokkink, J. op den Kelder and C. Pelekis, *How to poison your mother-in-law and other caching problems*, Chapter 10 in Search Theory: A Game-Theoretic Prespective, Alpern et al. (ed.), Springer, (2013).

[19] A. Frank, T. Jordán, Z. Szigeti, *An orientation theorem with parity conditions*, Discrete Appl. Math. 115, p. 37–47, (2001).

[20] P. Frankl, Z. Füredi, *A new short proof of the EKR theorem*, J. Combin. Theory Ser A, 119, no. 6, p. 1388–1390, (2012).

[21] P. Frankl, R. L. Graham, *Old and new proof of the Erdős-Ko-Rado theorem*, Sichuan Daxue Xuebao, 26, Special Issue, p. 112–122, (1982).

[22] P. Frankl, *Improved bounds for Erdős' Matching Conjecture*, Journal of Combinatorial Theory, Series A, 120, p. 1068–1072, (2013).

[23] P. Frankl, T. Luczak, K. Mieczkowska, *On matchings in hypergraphs*, The Electronic Journal of Combinatorics, 19(2), # P42, (2012).

[24] K. Hamza, *The smallest uniform upper bound on the distance between the mean and the median of the Binomial and Poisson distributions*, Statistics & Probability Letters, 23, p. 21–25, (1995).

[25] G. H. Hardy and E. M. Wright, *An introduction to the theory of numbers*, Oxford Univ. Press, 6th edition, (2008).

[26] W. Hoeffding, *On the Distribution of the Number of Successes in Independent Trials*, An. Math. Statistics, **27**, p. 713–721, (1956).

[27] K. Jogdeo, S. M. Samuels, *Monotone convergence of Binomial probabilities and a generalization of Ramanujan's equation*, Ann. Math. Statistics, Vol. 39, No. 4, p. 1191–1195, (1968).

[28] R. Kaas and J.M. Buhrman, *Mean, Median and Mode in Binomial Distributions*, Statistica Neerlandica **34**(1), p. 13–18, (1980).

[29] G. O. H. Katona, *A simple proof of the Erdős-Ko-Rado theorem*, J. Combinatorial Theory, Ser. B, 13, p. 183–184, (1972).

[30] M. Kearns, S. Suri, N. Montfort, *An Experimental Study of the Coloring Problem on Human Subject Networks*, Science 313, (5788), p. 824–827, (2006)

[31] K. Kikuta, W. Ruckle, *Accumulation games, Part 1: Noisy search*, J. Optim. Theory Appl., 94, No. 2, p. 395–408, (1997).

[32] K. Kikuta, W. Ruckle, *Continuous accumulation games in continuous regions*, J. Optim. Theory Appl., 106, No. 3, p. 581–601, (2000).

[33] K. Kikuta, W. Ruckle, *Continuous accumulation games on discrete locations*, Naval Res. Logistics, 49, No. 1, 60–77, (2002).

[34] T. L. Lai and H. Robbins, *Maximally Dependent Random Variables*, Proc.Nat.Acad.Sci.USA, Vol. 73, No. 2, p. 286–288, February 1976, Statistics.

[35] T. L. Lai and H. Robbins, *A Class of Dependent Random Variables and their Maxima*, Z. Wahrscheinlichkeitstheorie verw. Gebiete, **42**, p. 89-111, (1978).

[36] L. Lovász, *Combinatorial Problems and Excercises*, North-Holland, Amsterdam, 1979.

[37] M. Luby, *Removing Randomness in Parallel Computation without a Processor Penalty*, FOCS, p. 162-173, (1988).

[38] J. Matousek, *Lectures on Discrete Geometry*, Springer Verlag, Heidelberg, (2002).

[39] P.N. Panagopoulou and P. G. Spirakis, *A Game Theoretic Approach for Efficient Graph Coloring*, ISAAC 2008, LNCS 5369, p. 183-195, (2008).

[40] N. Papadatos, *Maximum variance of order statistics*, Annals of the Institute of Statistical Mathematics 47, p. 185–193, (1995).

[41] C. Pelekis, *Bernoulli trials of fixed parity, random and randomly oriented graphs*, (submited).

[42] C. Pelekis, *The König-Kikuta-Ruckle property on hypergraphs*, (in preparation).

[43] C. Pelekis and M. Schauer, *Network coloring and colored coin games*, Chapter 4 in Search Theory: A Game-Theoretic Prespective, Alpern et al. (ed.), Springer, (2013).

[44] T. Rychlik, *Extreme variances of order statistics in dependent samples*, Statistics and Probability Letters, 78, p. 1577–1582, (2008).

[45] S. M. Samuels, *On a Chebyshev-type inequality for sums of independent random variables*, Ann. Math. Stat., 37, p. 248–259, (1966).

[46] S. M. Samuels, *The Markov inequality for sums of independent random variables*, Ann. Math. Statistics, Vol. 40, No. 6, p. 1980–1984, (1969).

[47] E. R. Scheinerman and D. H. Ullman, *Fractional graph theory: A rational approach to the theory of graphs*, Wiley Interscience Series, (1997).

[48] M. Shaked and J.G. Shanthikumar, *Stochastic Orders and their Applications*. Springer, New York (2007).

[49] Shinsei Tazama, Teruhiro Shirakura and Saburo Tamura, *Enumeration of digraphs with given number of vertices of odd out-degree and vertices of odd in-degree*, Discrete Mathematics **90**, p. 63–74, (1991).

[50] Y. H. Wang: *On the number of successes in independent trials*, Statistics Sinica, 3(2),p. 295–312, (1993).

Summary

The main motivation behind this thesis is a certain type of win-lose games that are played on hypergraphs and can be translated into the following puzzle. Suppose there are two persons, say Alice and Bob. There are n biscuits, where n is a positive integer, and Alice chooses s of them uniformly at random. Bob possesses h grams of poison, where $h \geq 1$, and the lethal dose is 1 gram. How should Bob distribute the poison over the biscuits in order to maximize the probability of poisoning Alice?
This problem is due to Ken Kikuta and William Ruckle who, driven by less devious motives, formulated it in terms of accumulation games between two players. They conjectured that the optimal distribution of poison uses dosages of $1/j$ grams in as many biscuits as possible, where j is a positive integer that depends on h, n, s.

In Chapter 1 we introduce the poisoning problem and discuss its relation to known results from the literature. The conjecture of Kikuta and Ruckle is related to two other conjectures, one from extremal combinatorics and one from the theory of probability. The combinatorial flavor of the Kikuta-Ruckle conjecture is its relation to the matching conjecture of Paul Erdős and its fractional analogue. Its probabilistic flavor is its relation to a conjecture of Stephen Samuels on a tail probability problem. We also consider a poisoning problem on more general ground spaces. This leads to a geometric problem that generalizes the isodiametric one.

In Chapter 2 we settle the Kikuta-Ruckle conjecture in case $n = 2s - 1$. This case corresponds to the, so called, Odd graph. We also settle the conjecture for a few more cases using elementary combinatorial and game-theoretic arguments.

In Chapter 3 we consider the poisoning game on the cyclic graph. In this game the n biscuits are arranged cyclically and Alice chooses s consecutive of them uniformly at random. We find the value of this game along with

the optimal strategies of both players. In addition, we give a characterization of the fractional covering number of uniform hypergraphs obtained from the cyclic graph.

Chapter 4 deals with the analysis of a network coloring game. This is a game that is motivated by conflict resolution situations and is played on a graph. The vertices of the graph are thought of as players having a fixed set of available colors. The game is played in rounds and in each round all players simultaneously and individually choose a color with the perspective of ending up with a color that is different from the colors chosen by their neighbors. We analyze the network game by introducing a very simple search game. The optimal strategy of the searchers in this game involves tosses of fair colored coins and leads to the following combinatorial probability problem that is interesting on its own. Suppose that you can color n fair coins with n colors. It is not allowed to colors both sides of a coin with the same color, but all other combinations are allowed. Let X be the number of different colors after a toss of the coins. In what way should you color the coins such that you maximize the median of X? We solve this problem and consider its natural generalization to the case of biased coins.

Acknowledgements

I would like to thank everyone who has contributed to making my stay in Delft so much pleasurable. Especially, I would like to thank the following persons:

My mother, Argyro, my father, Konstantinos, and my brother, Nikolaos, for their love and for being always there for me.

Cindy Bosman, Evelyn Sharabi, Léonie Van der Zwan and Carl Schneider for taking care of us.

Dr. Hendrik Don, Dimitrios Kanellopoulos, Konstantinos Karousis, Dr. Derong Kong, Richard Kraaij, Moritz Schauer, Georgios Stamoulis and Dr. Julio Velazquez for the many entertaining moments and discussions.

Catharina Don, Marina Kitsara and Kind Van De Tekening for drawing for me so very nice pictures that vitalize the text. Vasilis Antonopoulos for his help in designing a cover for the thesis.

Prof. Michel Dekking for being my promotor and for many useful comments on early drafts of the thesis.

Prof. Steve Alpern, Prof. Shmuel Gal, Dr. Cornelis Kraaikamp, Prof. Kees Roos, Prof. Guido Schäfer and Prof. Piet Van Mieghem for accepting to be members of the jury of this thesis and for reading the manuscript.

Prof. Steve Alpern and Prof. Shmuel Gal for a very enjoyable visit to Israel.

Dr. Tobias Müller for his help, his advice, for saving me from troubles and for being such a nice guy.

Prof. Michalis Kolountzakis and Dr. Themistokles Mitsis for answering my questions.

All the past and present members of the Applied Probability and Statistics group at TU Delft for creating such a friendly and enjoyable atmosphere.

I am mostly indebted to my copromotor, Dr. Robbert Fokkink. For teaching me mathematics with so much patience and generosity. For introducing me to many interesting problems. For his kindness, his availability and his sense of humor. Most importantly, for hiring me and so giving me the opportunity to make a living out of my hobby. This is a rare privilege which is enjoyed by a small minority among the citizens of this world.

Curriculum Vitae

The author of this thesis was born on January 14, 1982 in Agrinio, Greece. In 2000 he got his diploma from the Unified (Polykladiko) Lyceum of Agrinio. In 2005 he got his degree (ptychion) in Mathematics from the University of Crete. From 2005 to 2006 he worked as a mathematics teacher at a high school in Crete. From 2006 to 2008 he joined the master programme in Mathematics at the National Technical University of Athens. In July 2009 he commenced his PhD research in the Probability group of the Delft Institute of Applied Mathematics. The relevant research has materialized in this thesis.