TIPSY COP AND TIPSY ROBBER: COLLISIONS OF BIASED RANDOM WALKS ON GRAPHS

PAMELA E. HARRIS, ERIK INSKO, AND FLORIAN LEHNER

Abstract. We study collisions of two interacting random walks, where one of the walks has a bias to create collisions, and the other one has a bias to avoid them. This can be seen as a variant of the well-studied cop and robber game in which both of the players get to employ their respective strategies on some (random) proportion of their moves, and move randomly otherwise.

We study different sensible strategies for both players on the infinite grid $\mathbb{Z}^2$ and on certain families of infinite trees obtained by attaching a δ-regular tree to every vertex of a Δ-regular tree, where $\Delta \geq \delta$. Our results show that the best possible cop strategy on the grid is very sensitive to change, and the best possible robber strategy on trees not only depends on the tree the game is played on, but also on the proportion of random moves by either player. We conclude with some directions for further study.

1. Introduction

A *random walk* on a graph G is a stochastic process with random variables $(X_i)_{i\in\mathbb{N}=\{1,2,3\dots\}}$ where x_{i+1} is chosen among the neighbors of X_i according to some predefined probability distribution. The study of random walks is a classical topic in discrete probability. For surveys on the study of random walks on graphs we refer the reader to [14, 20]. One question of interest is whether a given random walk is *recurrent* (almost surely $X_i = X_0$ for infinitely many i) or *transient* (almost surely $X_i = X_0$ for finitely many i). One of the first and most famous results in the study of recurrence and transience is Pólya's theorem, stating that the simple random walk on $\mathbb{Z}^d$ is recurrent for $d \leq 2$, but transient for $d \geq 3$, or as Shizuo Kakutani succinctly summarized it in a talk: "A drunk man will always find his way home, but a drunk bird may not."[1]

Pólya describes his motivation for this result as follows, see [17]. He was taking a stroll in the woods when he happened to cross paths with a young couple walking in the same woods, so often that he felt embarrassed as he suspected they thought he was snooping, of which, he assures the reader that he was not. This incident caused him to ask how likely it is that two independent random walks meet infinitely often. In $\mathbb{Z}^d$ this reduces to the problem of recurrence or transience of a single random walk, which is precisely the outcome of Pólya's theorem.

While recurrence and transience of random walks are certainly interesting properties to study, the story above suggests a different notion. Let us say that a random walk has the *finite collision property* if two independent copies of it almost surely meet only finitely many times. As noted in [13], for simple random walks on vertex transitive graphs such as $\mathbb{Z}^d$ this notion is equivalent to transience, but there are examples of graphs where this is not the case. First examples of recurrent simple random walks with the finite collision property were studied by Krishnapur and Peres [13], for further results on the finite collision property of recurrent graphs see for instance [2, 4, 6, 7, 8, 11].

There are also examples of transient random walks where two or even arbitrarily many copies meet infinitely often [5, 9]. The simplest such example is perhaps a biased random walk on $\mathbb{Z}$: the differences of two such walks yield a simple random walk on $\mathbb{Z}$ which is recurrent by Pólya's theorem and thus returns to 0 infinitely often. One key property of these examples is that two independent copies almost surely 'go to the same point at infinity' and move away from the starting vertex 'at the same speed.'

In the examples mentioned above it was always assumed that the two random walks are independent and that their distributions are identical; in the present paper we drop both of these assumptions and study

Date: October 6, 2022.

P. E. Harris was supported by a Karen Uhlenbeck EDGE Fellowship.

[1] A joke by Shizuo Kakutani at a UCLA colloquium talk as attributed in Rick Durrett's book Probability: Theory and Examples.

collisions between two biased random walks, where (roughly speaking) the first walk has a bias to reduce the distance between the positions of the two walks, whereas the second walk has a bias to increase this distance.

As the title suggests, this situation can be modeled by a pursuit-evasion game similar to the cop-and-robber game, a two-player game introduced independently by Quilliot [18] and by Nowakowski and Winkler [16] in the 1980s. In this game, two players (called cop and robber) alternate in moving their playing pieces from vertex to vertex along the edges of a given graph. The cop wins if after finitely many steps both the cop and robber are at the same vertex. The monograph [3] by Bonato and Nowakowski offers an extensive treatment of this game and numerous variants thereof.

A variant which is relevant to this paper is the cop and drunken robber game introduced by Kehagias and Prałat [12]. In this variant, the cop may follow some strategy, while the robber moves 'drunkenly' according to a random walk on the graph. More recently a variant called the *tipsy cop and drunken robber* was introduced [1, 10], where in addition to the robber moving randomly, the cop only follows a strategy on some steps and performs random moves otherwise.

In the current work, we consider the case in which both players can follow a strategy on some (randomly selected) moves and make uniform random moves otherwise. We call this game the *tipsy cop and tipsy robber* game. Roughly speaking, the tipsiness of each player governs the proportion of random moves of this player; we make this precise in Section 2. We note that if both players have tipsiness 1, then all moves are random and thus the question of who wins the game is equivalent to the question whether two independent random walks almost surely collide. The other extreme case, when both players have tipsiness 0, resembles the classical cop and robber game, but the two are not quite equivalent since in our game we choose randomly which player gets to move, and hence the moves of the two players do not necessarily alternate.

As in all two-player games, one question of interest is whether one of the two players has a winning strategy. Let us call a cop strategy winning if it almost surely leads to a win for the cop, and let us call a robber strategy winning if it leads to a win for robber with positive probability. The reason for this asymmetry in the definition of winning strategies is that (due to the randomness involved) there is always a positive probability that the cop wins the game before either of the two players gets to make a non-random move. It is worth noting that this observation also shows that on a finite graph, the cop will almost surely win the game, which is why we focus our attention on infinite graphs.

In Section 3, we study the game on the infinite Cartesian grid $\mathbb{Z}^2$. We show that the robber has a winning strategy if and only if they are less tipsy than the cop. In fact, if the robber is less tipsy, then it is not hard to show that the precise strategy does not matter so long as they increase the distance between themselves and the cop whenever they have a sober move. Perhaps surprisingly, the cop's winning strategies are more sensitive to change: there are strategies where each sober cop move decreases the distance between the cop and the robber that still allow the robber to win with positive probability despite being more tipsy than the cop.

In Section 4, we analyze the game on certain families of infinite trees. It is not hard to see that for regular trees, the question of who has a winning strategy boils down to a simple application of the gambler's ruin problem, but there are also trees exhibiting more interesting behavior. We consider the game on a family of infinite trees, $\{X(\Delta, \delta)\}$, where the tree $X(\Delta, \delta)$ is created by starting with a Δ-regular base tree and attaching a copy of a δ-regular tree to each node in this base tree. These trees provide examples where the robbers optimal strategy does not necessarily increase the distance between the players: under certain conditions the robber's optimal strategy may be to backtrack toward the cop in order to reach the base tree, where the number of possible escape routes is higher. Interestingly, these conditions not only depend on the values Δ and δ, but also on the tipsiness parameters.

We conclude the article in Section 5 by providing some possible direction for further study.

2. Background and general set up

2.1. Markov chains and random walks. A Markov chain with state space Ω is a random process $(X_\mathbf{s})_{\mathbf{s}\in\mathbb{N}}$ where each $X_\mathbf{s}$ takes values in Ω such that

$$\mathbb{P}[X_{\mathbf{s}+1} = x \mid X_1 = x_1, X_2 = x_2, \dots, X_\mathbf{s} = x_\mathbf{s}] = \mathbb{P}[X_{\mathbf{s}+1} = x \mid X_\mathbf{s} = x_\mathbf{s}]$$

whenever the event we condition on has positive probability. For a thorough introduction to the topic of denumerable Markov chains see [22]; in this section we briefly recall some well known facts. Let us say that the Markov chain *is at x at time step* $\mathbf{s}$ if $X_{\mathbf{s}} = x$, and let us denote the *transition probability* $\mathbb{P}[X_{\mathbf{s}+1} = b \mid X_{\mathbf{s}} = a]$ by $p_{ab;\mathbf{s}}$, or sometimes (for readability reasons) by $p_{a,b;\mathbf{s}}$. In order to make a Markov chain well defined, we also need to give a distribution of the first random variable X_1 in the process. Since our results do not depend on the initial distribution, we will tacitly assume that there is some $o \in \Omega$ such that $\mathbb{P}[X_1 = o] = 1$.

A Markov chain is called *time homogeneous* if $p_{ab;\mathbf{s}}$ does not depend on $\mathbf{s}$; in this case we simply denote it by p_{ab}. Most Markov chains that will appear in this paper have this property. A Markov chain is called *irreducible* if for every pair a, b of elements of Ω and every $\mathbf{s} \in \mathbb{N}$ there is some n such that the probability $p^n_{ab;\mathbf{s}}$ that $X_{\mathbf{s}+n} = b$ given that $X_{\mathbf{s}} = a$ is positive.

A random walk on a graph G is a Markov chain with state space $\Omega = V(G)$, where $p_{ab,\mathbf{s}}$ is only allowed to be non-zero if ab is an edge of G. When speaking about random walks on $\mathbb{N}$ or $\mathbb{Z}$, we always assume that edges are between elements whose difference is 1. If $p_{ab} = \frac{1}{\deg a}$ for any pair a, b of adjacent vertices, then we speak of a simple random walk; in this case the random walk is equally likely to transition from a to any of its neighbors. Another well-established way to define transition probabilities for a random walk is by assigning a weight $w(e) > 0$ to each edge e of the graph. The probability to move from a to b in this case is given by

$$p_{ab} = \frac{w(ab)}{\sum_{b' \sim a} w(ab')}, \tag{1}$$

in other words, the probability of moving from a to any of its neighbors is proportional to the weight of the edge connecting them. The simple random walk is a special case of this, where all edge weights are equal. We note that any random walk on a connected graph defined by edge weights as in (1) is time homogeneous and irreducible.

For every pair a, b of elements of Ω, let us denote by $T_{ab;\mathbf{s}}$ the random variable counting the number of steps it takes the Markov chain to reach b if it is at a at time step $\mathbf{s}$ (assuming that $\mathbb{P}[X_{\mathbf{s}} = a] > 0$), that is

$$T_{ab;\mathbf{s}} = \inf\{n > 0 : X_{\mathbf{s}+n} = b \text{ conditioned on } X_{\mathbf{s}} = a\}.$$

For a time homogeneous Markov chain this once again does not depend on $\mathbf{s}$, and we denote it by T_{ab}.

Theorem 2.1. For every time homogeneous, irreducible Markov chain the following (mutually exclusive) statements hold for every pair a, b of elements of Ω:

(1) $\mathbb{P}[T_{ab} < \infty] = 1$, in this case we call the Markov chain *recurrent*. This case further splits into:
   (a) $\mathbb{E}[T_{ab}] < \infty$ (*positive recurrent*), and
   (b) $\mathbb{E}[T_{ab}] = \infty$ (*null recurrent*).
(2) $\mathbb{P}[T_{ab} < \infty] < 1$, in this case we call the Markov chain *transient*.

In the remainder of this section we provide some well-known results about random walks and Markov chains. Most of these results can be deduced from results contained in any introductory textbook on the topics, see for instance [15, 21, 22].

The first result we mention is Pólya's theorem on recurrence and transience of simple random walks on $\mathbb{Z}^d$.

Theorem 2.2. The simple random walk on $\mathbb{Z}^d$ is null recurrent if $d \in \{1, 2\}$ and transient if $d \geq 3$.

The next result is [15, Exercise 2.1 (f)].

Theorem 2.3. A random walk on a graph defined by edge weights as in (1) is positive recurrent if and only if the sum of the edge weights is finite.

The next result is a direct consequence of [15, Theorem 2.16].

Theorem 2.4. Let w and w' be weight functions on $E(G)$ and assume that there are constants $c_1 > 0$ and c_2 such that $c_1 w(e) \leq w'(e) \leq c_2 w(e)$ for every $e \in E(G)$. The random walk defined by the weights w is recurrent if and only if the simple random walk on G is recurrent.

The last few results concern biased random walks on $\mathbb{N}_0 = \{0, 1, 2, \ldots\}$. These are random walks where the transition probability $p_{n,n+1}$ is p and the transition probability $p_{n,n-1}$ is $q = 1 - p$ for every $n > 0$. Such random walks can be interpreted as sums of random variables each taking values in $\{-1, 1\}$. If these random variables are independent, then their sum is well-understood due to the central limit theorem. We will need the following concentration result which readily follows from the central limit theorem, or from other concentration results such as the Azuma-Hoeffding inequality.

Theorem 2.5. Let X_i be a sequence of i.i.d. random variables each of which has finite mean μ and finite variance σ^2. Let $S_n = \sum_{i \leq n} X_i$. Then for every $\epsilon > 0$ there are constants k and n_0 such that

$$\mathbb{P}[|S_n - n\mu| > n\epsilon] < e^{-kn}$$

for any $n > n_0$.

The following well-known result can be found in most textbooks on probability theory, often as an introductory example to random walks.

Theorem 2.6 (Gambler's ruin)**.** Let $X_{\mathbf{s}}$ be a biased random walk on $\mathbb{N}_0$ as defined above.

- If $p > \frac{1}{2}$, then the random walk is transient.
- If $p < \frac{1}{2}$, then the random walk is positive recurrent and $\mathbb{E}[T_{n,0}] = \frac{n}{1-2p}$.
- If $p = \frac{1}{2}$, then the random walk is null recurrent.

Theorem 2.6 also holds for reasonably well-behaved inhomogeneous random walks; in particular, the following proposition can be derived by a suitable coupling of $X_{\mathbf{s}}$ with a simple random walk on $\mathbb{N}$.

Proposition 2.7. Let $X_{\mathbf{s}}$ a random walk on $\mathbb{N}$ whose transition probabilities are $p_{n,n+1;\mathbf{s}} = p_{\mathbf{s}}$ and $p_{n,n-1;\mathbf{s}} = q_{\mathbf{s}} = 1 - p_{\mathbf{s}}$ for all $n > 0$. If $p_{\mathbf{s}} \leq \frac{1}{2}$ for all $\mathbf{s} \in \mathbb{N}$, then almost surely $X_{\mathbf{s}} = 0$ for infinitely many $\mathbf{s}$. Conversely, if $p_{\mathbf{s}} > \frac{1}{2}$ for all $\mathbf{s} \in \mathbb{N}$, then there is a positive probability that there is no $\mathbf{s} > 0$ such that $X_{\mathbf{s}} = 0$.

The next result implies that the return time to the origin in the positive recurrent case of Theorem 2.6 has finite mean and variance. We provide a proof since we need a specific generalization of the biased random walk, where we allow finitely many transition probabilities to be different.

Proposition 2.8. Let $X_{\mathbf{s}}$ be a random walk on $\mathbb{N}_0$, and assume that all but finitely many transition probabilities $p_{n,n+1}$ are equal to $p < \frac{1}{2}$, and that the remaining transition probabilities are strictly less than one. Let T_i denote the random number of steps it takes to get from i to $i-1$ in this random walk. Then T_1 has finite mean and variance.

Proof. First observe that $T_1 = 1 + B \cdot (T_2 + T_1)$, where B is a random variable which takes the value 1 with probability $p_{1,2}$, and takes the value 0 with probability $1 - p_{1,2}$. Squaring both sides of this equation, we get that $T_1^2 = 1 + 2B \cdot (T_2 + T_1) + B \cdot (T_2^2 + 2T_2T_1 + T_1^2)$; note that $B^2 = B$ since B only takes values 0 and 1. Taking expected values on both sides (and noting that T_2 and T_1 are independent) yields

$$\begin{aligned}
\mathbb{E}[T_1] &= 1 + p_{1,2}(\mathbb{E}[T_1] + \mathbb{E}[T_2]) \\
\mathbb{E}[T_1^2] &= 1 + 2p_{1,2}(\mathbb{E}[T_1] + \mathbb{E}[T_2]) + p_{1,2}(\mathbb{E}[T_1^2] + 2\mathbb{E}[T_1]\mathbb{E}[T_2] + \mathbb{E}[T_2^2]).
\end{aligned}$$

We use induction on the maximal n such that $p_{n,n+1} \neq p$. If all $p_{n,n+1}$ are equal to p, then T_2 and T_1 have the same distribution, in particular $\mathbb{E}[T_2] = \mathbb{E}[T_1]$ and $\mathbb{E}[T_2^2] = \mathbb{E}[T_1^2]$. Solving the above equations for $\mathbb{E}[T_1]$ and $\mathbb{E}[T_1^2]$ gives $\mathbb{E}[T_1] = \frac{1}{1-2p}$ and $\mathbb{E}[T_1^2] = \frac{1}{1-2p}(1 + 4p\mathbb{E}[T_1] + 2p\mathbb{E}[T_1]^2)$ which is finite because $\mathbb{E}[T_1]$ is finite.

For the induction step, we may assume that $\mathbb{E}[T_2]$ and $\mathbb{E}[T_2^2]$ are both finite. Again, solving the above equations for $\mathbb{E}[T_1]$ and $\mathbb{E}[T_1^2]$ yields finite values. $\square$

Finally, we note that more precise statements about the behaviour of the biased random walk from Theorem 2.6 can be made in both the positive recurrent and the transient case. As one might expect, in the transient case the distance from 0 increases linearly in $\mathbf{s}$ (for instance by Theorem 2.5).

Theorem 2.9. Let $X_{\mathbf{s}}$ be a biased random walk on $\mathbb{N}_0$ as defined above, and assume that we are in the transient case, that is $p > \frac{1}{2}$. Then for every small enough $k > 0$ the probability $\mathbb{P}[X_{\mathbf{s}} > k\mathbf{s}]$ goes to 1 as $\mathbf{s} \to \infty$.

On the other hand, in the positive recurrent case, one can show that the random walk will spend a large proportion of steps 'close to 0'.

Theorem 2.10. Let $X_{\mathbf{s}}$ be a biased random walk on $\mathbb{N}_0$ as defined above, and assume that we are in the positive recurrent case, that is $p < \frac{1}{2}$. Let $M(n, \mathbf{s})$ be the number of $\mathbf{t} \leq \mathbf{s}$ such that $X_{\mathbf{s}} < n$. For every $\epsilon > 0$ there is an $n \in \mathbb{N}$ such that almost surely $\liminf_{\mathbf{s} \to \infty} \frac{M(n,\mathbf{s})}{\mathbf{s}} > 1 - \epsilon$.

2.2. **The tipsy cop and robber game.** In this paper, we will study pairs $\mathbf{C}_{\mathbf{s}}, \mathbf{R}_{\mathbf{s}}$ of random walks, where the probability distribution of the next step does not only depend on $\mathbf{s}$, but also on the relative positions of the random walks at the current time step. Roughly speaking, the random walk $\mathbf{C}_{\mathbf{s}}$ should be biased to maximize the probability that the two walks meet, whereas the random walk $\mathbf{R}_{\mathbf{s}}$ should aim to minimize the meeting probability.

More precisely, we have four parameters c, r, t_c, and t_r satisfying the condition $c + r + t_c + t_r = 1$, and for every $\mathbf{s}$ we have a pair of *strategy functions* $c_{\mathbf{s}} \colon V(G) \times V(G) \to V(G)$ and $r_{\mathbf{s}} \colon V(G) \times V(G) \to V(G)$, where $c_{\mathbf{s}}(u, v)$ is a neighbor of u and $r_{\mathbf{s}}(u, v)$ is a neighbor of v. Given these parameters and strategy functions, we define the transition probabilities of the pair $(\mathbf{C}_{\mathbf{s}}, \mathbf{R}_{\mathbf{s}})$ by

- with probability c: $(\mathbf{C}_{\mathbf{s}+1}, \mathbf{R}_{\mathbf{s}+1}) = (c_{\mathbf{s}}(\mathbf{C}_{\mathbf{s}}, \mathbf{R}_{\mathbf{s}}), \mathbf{R}_{\mathbf{s}})$,
- with probability r: $(\mathbf{C}_{\mathbf{s}+1}, \mathbf{R}_{\mathbf{s}+1}) = (\mathbf{C}_{\mathbf{s}}, r_{\mathbf{s}}(\mathbf{C}_{\mathbf{s}}, \mathbf{R}_{\mathbf{s}}))$,
- with probability t_c: $\mathbf{C}_{\mathbf{s}+1}$ is a uniformly chosen random neighbor of $\mathbf{C}_{\mathbf{s}}$ and $\mathbf{R}_{\mathbf{s}+1} = \mathbf{R}_{\mathbf{s}}$, and
- with probability t_r: $\mathbf{R}_{\mathbf{s}+1}$ is a uniformly chosen random neighbor of $\mathbf{R}_{\mathbf{s}}$ and $\mathbf{C}_{\mathbf{s}+1} = \mathbf{C}_{\mathbf{s}}$.

We note that this defines a Markov chain with state space $V(G) \times V(G)$ whose transition probabilities depend on c, r, t_c, t_r and the sequences $c_{\mathbf{s}}$ and $r_{\mathbf{s}}$. We will be interested in the probability of the event that there is some time step $\mathbf{s} \in \mathbb{N}$ for which $\mathbf{C}_{\mathbf{s}}$ and $\mathbf{R}_{\mathbf{s}}$ coincide.

This can be thought of as a pursuit-evasion game on a graph of which we call the the tipsy cop and robber game. This game is played between two players, the cop and the robber, each of whom controls a playing piece. A move by either player consists of taking their respective playing piece and moving it to an adjacent vertex. In every round we spin a spinner wheel with four possible outcomes which coincide with the four transition options above: (1) the cop can (but does not have to) make a move of their choosing, (2) the robber can make a move of their choosing, (3) the cop has to make a move chosen uniformly at random, (4) the robber has to make a move chosen uniformly at random. An example of one such spinner is depicted in Figure 1. The cop wins the game, if at some point both playing pieces are at the same vertex, the robber wins the game if this never happens.

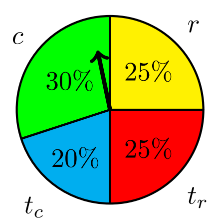

FIGURE 1. Probability of each move based on a spinner model.

We refer to t_c and t_r as the *tipsiness parameter* of the cop and robber, respectively. The family of functions $c_{\mathbf{s}} \colon V(G) \times V(G) \to V(G)$ is called a *cop strategy*; $c_{\mathbf{s}}(u, v) = w$ denotes the fact that if at time $\mathbf{s}$ the cop is at u, the robber is at v, and the cop gets to decide where to move, then the cop will move to w. Similarly, the family of functions $r_{\mathbf{s}} \colon V(G) \times V(G) \to V(G)$ is called a *robber strategy*. We refer to moves in which the cop or robber gets to employ their strategy as *sober* cop or robber moves, and to moves where a random neighbor is selected as *tipsy* moves.

We say that the cop strategy $c_{\mathbf{s}}$ is *winning* against the robber strategy $r_{\mathbf{s}}$ if and only if there almost surely is a time step $\mathbf{s}$ such that $\mathbf{C}_{\mathbf{s}} = \mathbf{R}_{\mathbf{s}}$ in the random process defined above. Note that if the Markov chain $(\mathbf{C}_{\mathbf{s}}, \mathbf{R}_{\mathbf{s}})$ is recurrent, then the cop strategy $c_{\mathbf{s}}$ is winning: in this case the probability of reaching any state in finite time is 1, and in particular we will almost surely reach any state with $\mathbf{C}_{\mathbf{s}} = \mathbf{R}_{\mathbf{s}}$ in finite time.

We note that most of the strategies we are interested in do not depend on $\mathbf{s}$ (but it will be convenient to allow the dependency on $\mathbf{s}$ in some of the proofs). Moreover, the graphs will be such that the Markov chain $(\mathbf{C_s}, \mathbf{R_s})$ defined above is irreducible. Consequently, if the robber has a positive probability of winning the game from some starting position $(\mathbf{C}_0, \mathbf{R}_0)$, then they have a positive probability of winning from any starting position, since there is a positive probability that the game state will reach $(\mathbf{C}_0, \mathbf{R}_0)$ before either player makes a sober move.

We call a cop strategy $c_\mathbf{s}$ superior to a cop strategy $c'_\mathbf{s}$ against a robber strategy $r_\mathbf{s}$ if the corresponding random processes $(\mathbf{C_s}, \mathbf{R_s})$ and $(\mathbf{C'_s}, \mathbf{R'_s})$ satisfy

$$\forall N\colon \mathbb{P}[\exists\, \mathbf{s} < N\colon \mathbf{C_s} = \mathbf{R_s}] \geq \mathbb{P}[\exists\, \mathbf{s} < N\colon \mathbf{C'_s} = \mathbf{R'_s}].$$

A cop strategy is called best possible in a class C of strategies against a given robber strategy $r_\mathbf{s}$, if it is superior to any other strategy in C against $r_\mathbf{s}$. Superior and best possible robber strategies can be defined analogously (with the converse inequality). We note that if $c_\mathbf{s}$ is superior to a strategy $c'_\mathbf{s}$ against $r_\mathbf{s}$ and $c'_\mathbf{s}$ is winning against $r_\mathbf{s}$, then so is $c_\mathbf{s}$. Thus there is no point in choosing $c'_\mathbf{s}$ over $c_\mathbf{s}$ in case the robber plays $r_\mathbf{s}$.

3. Optimal and less than optimal strategies on the infinite grid

In this section we describe the best strategy for each player to take given different values of r and c when playing on the infinite grid graph $\mathbb{Z}^2$. Following Pólya's approach, we start by simplifying the random process $(\mathbf{C_s}, \mathbf{R_s})$ defined in the last section: instead of keeping track of both random walks, it suffices to keep track of the difference $\mathbf{D_s} := \mathbf{R_s} - \mathbf{C_s}$. In other words, we always think of the cop's position as the origin $(0,0)$, and the current state is uniquely determined by the robber's position (x,y) in $\mathbb{Z}^2$.

Also the outcome of a tipsy robber move or tipsy cop move are essentially the same. So to simplify our notation, we combine t_r and t_c into one parameter which we call t. We now have the following parameters:

- c is the probability of a sober move by the cop,
- r is the probability of a sober move by the robber,
- $t = t_c + t_r = 1 - (c + r)$ is the probability of a random move by either player,
- $\mathbf{D_s} = (x,y) \in \mathbb{Z}^2$ is the difference between the cop's and robber's positions.

Intuitively, it is clear that the cop should try to decrease their distance from the robber while the robber should try to increase this distance. Note that generally (unless $x = 0$ or $y = 0$) there will be two different ways of increasing the distance and two different ways of decreasing it. Thus there are several somewhat sensible strategies for both the cop and the robber. The main result of this section shows that not all of these strategies are equally good.

It will be convenient to impose some restrictions on the strategies we consider. We assume that the strategies of both the cop and the robber depend on the current state (x,y) of the game, and the number of steps $\mathbf{s}$ made so far. It is worth noting that most of the strategies we consider do not depend on $\mathbf{s}$, but considering the possibility that they do will occasionally be convenient.

We introduce the following notation: $c_\mathbf{s}(x,y)$ and $r_\mathbf{s}(x,y)$ denote the strategies of the cop and robber, respectively. We let $(x',y') = c_\mathbf{s}(x,y)$, and $(x',y') = r_\mathbf{s}(x,y)$ denote the fact that a sober cop or robber move at step $\mathbf{s}$ in position (x,y) leads to the new position (x',y'). We assume that strategies are symmetric in both coordinates, i.e. if $c_\mathbf{s}(x,y) = (x',y')$, then $c_\mathbf{s}(-x,y) = (-x',y')$, $c_\mathbf{s}(x,-y) = (x',-y')$, and $c_\mathbf{s}(-x,-y) = (-x',-y')$ and the same holds for $r_\mathbf{s}$. Note that under the above symmetry conditions it is sufficient to define the strategies $c_\mathbf{s}$ and $r_\mathbf{s}$ for $y \geq |x|$.

Among the possible strategies satisfying the conditions above, we will be particularly interested in the following strategies which, by the above symmetry assumptions, we only need to define for $y \geq |x|$.

- Robber Strategy (RS) always increases the y-coordinate.
- Cop Strategy 1 (CS1) always decreases the y-coordinate. Note that if $x = y$ then by symmetry this is equivalent to decreasing the x-coordinate, and if $y = -x$, then it is equivalent to increasing the x-coordinate.
- Cop Strategy 2 (CS2): always decreases the y-coordinate, unless $x = y$ in which case the cop is allowed to do anything subject to the symmetry assumptions above. We point out that CS2 is a family of strategies rather than a single strategy, and that CS1 is a special instance of CS2.

As mentioned above, throughout the game the game state (x,y) can be described by a random walk on $\mathbb{Z}^2$, where the transition probabilities depend on the strategies employed by the cop and the robber, as well as on the probabilities c and r of a sober move by either player. We note that this random walk is irreducible.

It is also time homogeneous, unless the cop and robber strategies change over time. If this random walk is recurrent, then the cop almost surely wins, if it is positive recurrent then the expected time until this happens is finite. Conversely, if this random walk is transient, then the robber has a positive probability of winning.

Our main result in this section, Theorem 3.1, is as follows.

Theorem 3.1. With parameters as defined above:

(1) If $r > c$, then the robber wins with positive probability as long as they play any strategy that increases their distance from the cop in every round.
(2) If $c \geq r$, then the cop almost surely wins, provided they play the strategy to always decrease the largest coordinate. If $c > r$, then the expected time until the game ends is finite.
(3) If $c \geq r$ and the cop employs the strategy to always decrease the smallest coordinate, then the robber can win with positive probability provided that the difference $c - r$ is small enough.

The first part of this theorem follows from Lemma 3.2 below. The second and third part will follow from results in Sections 3.1 and 3.2, respectively.

Lemma 3.2. If $r > c$, then robber wins with positive probability as long as the robber strategy is such that the distance between the position of the cop and the robber increases on every sober robber move. In particular, the robber wins by playing strategy RS.

Proof. Since $r > c$ and the robber always increases the distance between themselves and the cop, we know that the probability of the distance increasing between the cop and robber in any round is $r + t/2$ if they are not on the same axis (the smaller coordinate is nonzero), and $r + 3t/4$ when they are on the same axis (the smallest coordinate is zero). The probability of the distance decreasing is at most $c + t/2$. Since $r + t/2 > c + t/2$, the probability of the distance increasing in any round is strictly larger than $1/2$. Hence, the expected distance between the cop and robber increases in each round. Implying that the robber wins with positive probability. □

In the remainder of this section, we focus on the case $c \geq r$. As one might expect, for $c \geq r$ the cop has a strategy which makes the corresponding random walk recurrent, and positive recurrent in the case $c > r$. This is shown in Subsection 3.1. On the other hand, in contrast to Lemma 3.2, this actually depends on the strategy chosen by the cop; even if $c > r$ there are cop strategies always decreasing the distance between the cop and the robber which do not almost surely lead to a cop win. One such strategy is discussed in Subsection 3.2.

3.1. **Cop playing an optimal strategy.** In this section, we show that as long as $c \geq r$ the cop has a winning strategy. The section is organized as follows: we start by showing in Lemma 3.3 that RS is the best possible symmetric robber strategy against CS1, and after that, we show that CS1 is almost surely winning against RS given that $c \geq r$. The proof of this result is divided into two lemmas: we begin by showing that CS1 is the best possible strategy against RS in the family CS2, and then we give a strategy in the family CS2 which is winning against RS.

Lemma 3.3. The robber strategy RS is superior to any symmetric robber strategy against CS1. In particular, if the cop plays CS1, then the robber has a symmetric strategy which is winning with positive probability if and only if RS is winning with positive probability.

Proof. By the symmetry assumptions we assume that every game state (x, y) satisfies $y \geq |x|$. If it does not, then we can reflect along the diagonals to obtain an 'equivalent' game state that does satisfy this assumption.

We start by showing that if the robber has a strategy which is winning with positive probability, then RS is also winning with positive probability. Let $r_{\mathbf{s}}(x, y)$ be any winning robber strategy and let $\mathbf{s}_0$ be the minimal $\mathbf{s}$ such that $r_{\mathbf{s}_0}(x, y)$ does not increase the larger coordinate for some (x, y). Assume that $r_{\mathbf{s}_0}(x, y) = (x+1, y)$ (the other cases are similar) and that at time $\mathbf{s}_0$ there is a robber move. Define a robber strategy $\hat{r}$ as follows:

$$\hat{r}_{\mathbf{s}}(x, y) = \begin{cases} r_{\mathbf{s}}(x, y) & \text{if } \mathbf{s} < \mathbf{s}_0 \text{ or } \mathbf{s} > \mathbf{s}_0 \text{ and we are in equivalent game states using } r_{\mathbf{s}} \text{ and } \hat{r}_{\mathbf{s}} \\ (x, y+1) & \text{if } \mathbf{s} = \mathbf{s}_0 \\ r_{\mathbf{s}}(x-1, y-1) & otherwise. \end{cases}$$

We show that $\hat{r}_{\mathbf{s}}$ is superior to $r_{\mathbf{s}}$ by giving a coupling between the two corresponding processes $\mathbf{D}_{\mathbf{s}}$ and $\hat{\mathbf{D}}_{\mathbf{s}}$. In this coupling, it is always true that if $\hat{\mathbf{D}}_N = (0,0)$ then $\mathbf{D}_n = (0,0)$ for some $n \leq N$. The coupling is given by playing games with the two different strategies simultaneously using the same feed of random variables. More precisely, if we have a sober cop or sober robber move in one game, then we also have a sober cop or sober robber move in the other game. If we have a random move in one game, then we have a random move in the other, and the two random moves go in the same direction (modulo our symmetry assumptions).

Note that the game states using r and $\hat{r}$ are either equivalent or differ by $(1,1)$ with the coordinates of $\hat{r}$ being greater than or equal to those of r. In particular, if the robber gets caught using $\hat{r}$, they get caught in the same step or earlier using r. Iterating this construction shows that RS is superior to any symmetric robber strategy against CS1.

We note that the other direction of the if and only if statement is a tautology – if RS is winning with positive probability, then of course, the robber has a strategy which is winning with positive probability. □

Lemma 3.4. The cop strategy CS1 is superior to any strategy of type CS2 against RS. In particular, if the cop wins almost surely by playing any strategy of type CS2 against RS, then the cop also wins almost surely by playing CS1. Moreover, if the expected time until the cop wins using said strategy of type CS2 is finite, then so is the time until they win using CS1.

Proof. By symmetry of the strategies we may assume that every game state (x,y) satisfies $y \geq |x|$. If it does not, then we can reflect along the diagonals to obtain an 'equivalent' game state satisfying this assumption.

As in the proof of Lemma 3.3, we show superiority of strategy CS1 by giving a coupling between two games, G_1 and G_2, where the robber follows RS in both games, and the cop follows strategy CS1 in G_1 and strategy CS2 in G_2. Again, the move sequences are determined by the same feed of random variables in both games, meaning that if it is a cop's turn in G_1 it is also a cop's turn in G_2 and vice versa.

We inductively ensure that if the state of G_1 is (x,y), then the state of G_2 is $(x+j, y+j)$ or $(x-j, y+j)$ for some $j \geq 0$. Note that this holds before the first step with $j = 0$.

If neither G_1 nor G_2 is on the diagonal, then the transitions are the same in G_1 and G_2 regardless of whose move it is, so the inductive claim remains valid. Also, if it is a sober robber move, then G_1 transitions from (x,y) to $(x, y+1)$ and G_2 transitions from $(x \pm j, y+j)$ to $(x \pm j, y+j+1)$ since the robber plays strategy RS in both games. Clearly the new game states satisfy the inductive hypothesis.

We now analyze what happens when it is a sober cop move or a random move when either G_1 or G_2 are on the diagonal, and before proceeding, we note that on the diagonal $x = y$, random 'right' and 'up' moves are both interpreted as 'up' (due to our symmetry assumptions), and similarly 'left' and 'down' moves are interpreted as 'left'. If both G_1 and G_2 are on the diagonal $x = y$ and it is a sober cop move, then G_1 transitions from (x,x) to $(x-1, x)$ and G_2 transitions from $(x+j, x+j)$ to $(x+j, x+j+1)$ or to $(x+j-1, x+j)$. In both cases the inductive claim is satisfied after the transition. Also if both G_1 and G_2 are on the diagonal $x = y$ and it is a random move, then both games make the same move, so the inductive claim remains valid.

Note that if only one game is on the diagonal, then it is G_1. In this case, if G_1 is at (x,x) then G_2 is at $(x-j, x+j)$ for some $j > 0$. If only G_1 is on the diagonal, and it is cop move, then G_1 transitions from (x,x) to $(x-1,x)$ and G_2 transitions from $(x-j, x+j)$ to $(x-j, x+j-1)$. The result of this round means that the original value of j has now decreased by 1, and the inductive claim is still valid.

If only G_1 is on the diagonal $x = y$, and it is a random move, then there are four random moves to consider. Case-by-case analysis shows that if the game states of G_1 and G_2 make a random move to the right or downwards, then the original value of j decreases by 1, and the inductive claim still holds. If the game states make a random move to the left or upwards, then j remains the same, and the inductive claim still holds.

By the inductive claim, if the cop wins using a strategy of type CS2, i.e., the game G_2 reaches $(0,0)$, then the game G_1 must also be at state $(0,0)$ or it reached it sooner since the respective game states satisfy $|(x,y)| \leq |(x \pm j, y+j)|$. □

Lemma 3.5. Assume that $r < c$. Take CS2 so that (on the diagonal) we follow CS1 with probability $p = \frac{t(c-r)}{2c(2r+t)}$, and move up with probability $1-p$. If cop plays this specific instance of CS2 and robber plays strategy RS, then cop wins the game in finite expected time.

Proof. Due to our symmetry assumptions, the game describes a random walk on the induced subgraph of the integer grid with vertex set $\{(x,y) \in \mathbb{Z}^2 \mid y \geq |x|\}$. The only possible transition in this graph starting at $(0,0)$ is to $(0,1)$, but this is not important since the game ends when we reach $(0,0)$. The probability of moving from $(x,y) \neq (0,0)$ to (x',y') are given by

$$p_{(x,y)\to(x',y')} = \begin{cases} r + \frac{t}{4} & \text{if } |x| \neq y \text{ and } (x',y') = (x, y+1) \\ c + \frac{t}{4} & \text{if } |x| \neq y \text{ and } (x',y') = (x, y-1) \\ \frac{t}{2} & \text{if } |x| \neq y \text{ and } (x',y') = (x \pm 1, y) \\ pc + \frac{t}{2} & \text{if } x = \pm y \text{ and } (x',y') = (x \mp 1, y) \\ (1-p)c + r + \frac{t}{2} & \text{if } |x| = y \text{ and } (x',y') = (x, y+1). \end{cases}$$

As a first step, we will show that these are precisely the transition probabilities of a random walk given by edge weights, that is, there is a function $w\colon E(G) \to \mathbb{R}_+$ such that

$$p_{a\to b} = \frac{w(ab)}{\sum_{b'\sim a} w(ab')}.$$

We claim that this function is given by

$$w((x,y),(x,y+1)) = \left(\frac{r+t/4}{c+t/4}\right)^y \text{ and}$$

$$w((x,y),(x+1,y)) = \frac{t/4}{c+t/4}\left(\frac{r+t/4}{c+t/4}\right)^{y-1}.$$

To prove this claim, first consider a vertex $a = (x,y)$ with $y \neq |x|$. Then

$$\sum_{b'\sim a} w(ab') = \left(\frac{r+t/4}{c+t/4}\right)^y + \left(\frac{r+t/4}{c+t/4}\right)^{y-1} + 2\cdot\frac{t/4}{c+t/4}\left(\frac{r+t/4}{c+t/4}\right)^{y-1} = \frac{(r+t/4)^{y-1}}{(c+t/4)^y}$$

and it is easy to verify that $p_{a\to b} = w(ab)/\sum_{b'\sim a} w(ab')$ for every neighbor b of a. Next consider a vertex $a = (x,x)$; the case $a = (x,-x)$ is analogous. We only have two incident edges and thus obtain

$$\sum_{b'\sim a} w(ab') = \left(\frac{r+t/4}{c+t/4}\right)^y + \frac{t/4}{c+t/4}\left(\frac{r+t/4}{c+t/4}\right)^{y-1} = \frac{(r+t/4)^{y-1}}{(c+t/4)^y}\cdot\left(r+\frac{t}{2}\right).$$

Letting $b = (x-1,y)$, we obtain

$$\frac{w(ab)}{\sum_{b'\sim a} w(ab')} = \frac{t/4}{r+t/2} = \frac{t(c+r+t)}{2(2r+t)} = \frac{t(c-r+2r+t)}{2(2r+t)} = \frac{t(c-r)}{2(2r+t)} + \frac{t(2r+t)}{2(2r+t)} = pc + \frac{t}{2}$$

as claimed. Since a has only two neighbors whose transition probabilities add up to 1, this concludes the proof of our claim.

By Theorem 2.3, this random walk is positive recurrent if and only if $\sum_{a,b} w(a,b) < \infty$. Hence it suffices to show that

$$\sum_{y\geq 0}\left(\sum_{-y\leq x\leq y} w((x,y),(x,y+1))\right) + \sum_{y\geq 0}\left(\sum_{-y\leq x<y} w((x,y),(x+1,y))\right) < \infty.$$

Letting $\beta = \frac{r+t/4}{c+t/4}$ and $\alpha = \frac{t/4}{c+t/4}$ and noting that $\beta < 1$, we obtain

$$\sum_{y\geq 0}\left(\sum_{-y\leq x\leq y} w((x,y),(x,y+1))\right) = \sum_{y\geq 0}\left(\sum_{-y\leq x\leq y}\beta^y\right) = \sum_{y\geq 0}(2y+1)\beta^y = \frac{\beta+1}{(\beta-1)^2} < \infty,$$

and

$$\sum_{y\geq 0}\left(\sum_{-y\leq x<y} w((x,y),(x+1,y))\right) = \sum_{y\geq 1}\left(\sum_{-y\leq x<y}\alpha\beta^{y-1}\right) = \sum_{y\geq 1} 2y\alpha\beta^{y-1} = \frac{2\alpha}{(\beta-1)^2} < \infty.$$

Thus, the random walk is positive recurrent and in particular the expected time until we reach $(0, 0)$ is finite; in other words, the cop wins the game almost surely and the expected time until this happens is finite. $\square$

Finally, let us consider the case $r = c$; note that Lemmas 3.3 and 3.4 still apply in this case.

Lemma 3.6. If $0 \leq r = c \leq 1/2$ then the game is null recurrent provided the cop plays strategy CS1 and the robber plays RS, that is, cop wins almost surely, but the expected length of the game does not exist.

Proof. First we note that when $r = c = 0$, then $\mathbf{D_s}$ describes a simple random walk on the integer lattice $\mathbb{Z}^2$ which is null-recurrent, and when $r = c = 1/2$ then $\mathbf{D_s}$ describes a simple random walk on the integers $\mathbb{Z}$ which is also null-recurrent.

For the remaining cases $0 < r = c < 1/2$, the distance of $\mathbf{D_s}$ and $(0, 0)$ is a random walk on the integers that is just slightly biased away from the origin because it increases with probability $r + t/2 = c + t/2$ when $\mathbf{D_s}$ is not on an axis and $r + 3t/4 > c + t/4$ in the rare event that it is on an axis. Since the simple random walk on the integers is null recurrent, the random walk modeling the game cannot be positive recurrent.

If the cop plays a strategy of the form CS2 where on the diagonal cop always increases the y-coordinate (this is the same strategy as in Lemma 3.5, with $p = 0$) against RS, then we get a random walk on a graph with two different edge weights

$$w((x, y), (x + 1, y)) = \left(\frac{t/4}{c + t/4}\right) \quad \text{and} \quad w((x, y), (x, y + 1)) = 1.$$

Suppose for sake of contradiction that this walk is transient, then the random walk with all edge weights equal to 1 would also be transient by Theorem 2.4, but we know that a simple random walk on $\mathbb{Z}^2$ is recurrent. So this particular instance of CS2 is winning against RS. By Lemma 3.4, if CS2 is winning against RS, then so is CS1. By Lemma 3.3, if CS1 wins against RS, then CS1 wins against any symmetric robber strategy. This shows that CS1 is a winning strategy for the cop. $\square$

3.2. **Cop playing foolish strategy.** The aim of this subsection is to show that the cop should not deviate too far from strategy CS1 if they want to maximize their chances of winning. More precisely, if the cop plays a less than optimal strategy of decreasing the smaller coordinate first, and the robber plays strategy RS (increase the larger coordinate), then we will show that the robber has positive probability of winning for some values $c > r$.

Theorem 3.7. If the cop employs the strategy of always decreasing the smaller nonzero coordinate, and the robber plays RS, then the robber has a positive probability of winning provided c, r, and t satisfy the following constraint

$$2r^2 + 2rt - \frac{3ct}{2} - 2c^2 > 0.$$

Proof. We will show that the distance function's larger coordinate is expected to increase after every pair of moves when $2r^2 + 2rt - \frac{3ct}{2} - 2c^2 > 0$. In any round of the game, the larger coordinate of the distance function between the cop and robber is expected to increase except when the smaller coordinate is zero. So if the smaller coordinate is greater than 1, the larger coordinate is expected to increase after any pair of moves.

If the smaller coordinate is zero at the start of a round, then the expected change of the larger coordinate is $r - c < 0$ after one round. To compute the expected change in the larger coordinate after two rounds when the smaller coordinate is zero, we note the following probabilities describing the possible changes within one round when the smaller coordinate is less than 2:

$$\begin{aligned} &p_{(0,y)\to(0,y+1)} = p_{(\pm1,y)\to(\pm1,y+1)} = r + t/4, && p_{(0,y)\to(\pm1,y)} = p_{(\pm1,y)\to(\pm2,y)} = t/4, \\ &p_{(0,y)\to(0,y-1)} = p_{(\pm1,y)\to(0,y)} = c + t/4, && p_{(\pm1,y),(\pm1,y-1)} = t/4. \end{aligned}$$

Thus the probabilities of moving from $(0, y)$ to any other location within two rounds are as follows:

$$\begin{aligned} &p_{(0,y)\to(0,y+2)} = (r + t/4)^2, && p_{(0,y)\to(\pm1,y+1)} = t(r + t/4), \\ &p_{(0,y)\to(\pm1,y-1)} = t/2(c + t/2), && p_{(0,y)\to(0,y-2)} = (c + t/4)^2. \end{aligned}$$

Hence the expected change in the larger coordinate after two rounds is given by

$$2(r+t/4)^2 - 2(c+t/4)^2 + (t)(r+t/4) + (t/2)(c+t/2) = 2r^2 + 2rt - \frac{3ct}{2} - 2c^2.$$

The transition probabilities after two rounds when the smaller coordinate starts at plus or minus one are calculated below:

$$\begin{aligned} p_{(\pm1,y)\to(\pm1,y+2)} &= (r+t/4)^2, & p_{(\pm1,y)\to(\pm1,y-2)} &= (t/4)^2,\\ p_{(\pm1,y)\to(0,y+1)} &= 2(c+t/4)(r+t/4), & p_{(\pm1,y)\to(\pm2,y+1)} &= (t/2)(r+t/4),\\ p_{(\pm1,y)\to(0,y-1)} &= (c+t/2)(c+t/4), & p_{(\pm1,y)\to(\pm2,y-1)} &= 2(t/4)^2. \end{aligned}$$

Hence the expected change of the larger coordinate after two rounds of the game is $2r^2+2rt+2rc-c^2-\frac{ct}{4}$ when smaller coordinate starts off as plus or minus one, and we note that

$$2r^2 + 2rt + 2rc - c^2 - \frac{ct}{4} > 2r^2 + 2rt - \frac{3ct}{2} - 2c^2 > 0.$$

By Proposition 2.7, since the distance function's larger coordinate is expected to increase after every pair of moves when $2r^2+2rt-\frac{3ct}{2}-2c^2 > 0$, there is a positive probability that the game lasts forever when c, r, and t satisfy the constraint $2r^2+2rt-\frac{3ct}{2}-2c^2 > 0$. □

4. Game on Trees

Throughout this section we use the same parameters as defined previously. However, our approach requires us to distinguish between the tipsy moves by the cop and the robber. Hence we reintroduce t_r and t_c, the probability of a tipsy robber move and the probability of a tipsy cop move, respectively. We assume that $c+t_c = r+t_r = 1/2$, or in other words, we assume that no player gets to make significantly more moves than the other; the only imbalance in the game comes from the tipsiness parameters of the two players.

We start by considering infinite δ-regular trees. On such a tree, all paths of length n look the same, so a game state can be described completely (up to isomorphism) by the distance between the two players without mentioning their exact positions. Hence the optimal cop strategy is to move in the unique direction that decreases this distance and the optimal strategy for the robber is to move in any direction that increases the distance (which are all the same up to isomorphism). If both players play their optimal strategies, then the probability that the distance between the two players increases is given by $r+(t_c+t_r)\frac{\delta-1}{\delta}$ and the probability of moving closer is $c+(t_c+t_r)\frac{1}{\delta}$.

Consequently, the cop and robber game can be reduced to the gambler's ruin problem with these transition probabilities. By substituting $c = 1/2 - t_c$ and $r = 1/2 - t_r$ we arrive at the following result.

Lemma 4.1. Assume the cop and the robber play strategies as described above on an infinite δ-regular tree. If $t_r \geq t_c(\delta-1)$, then the cop almost surely wins the game, otherwise the robber wins the game with positive probability.

We now consider a more general class of trees which can be described as δ-regular trees rooted to a $(\Delta-1)$-regular tree. To be precise, let $\delta > 1$, and let T be the tree that is δ-regular everywhere except the root, which has degree $\delta - 1$. Let $\Delta > \delta$, and let B be a $(\Delta-1)$-regular tree. Then the tree $X(\Delta,\delta)$ is constructed by connecting the root of a copy of T to each node in B, see Figure 2 for examples. We call B the *base tree* of $X(\Delta,\delta)$ and we refer to a copy of the δ-regular tree in $X(\Delta,\delta)$ as a *small tree*. Clearly, each node in the base of $X(\Delta,\delta)$ has degree Δ, and each node in a small tree has degree δ. Note that $X(\delta,\delta)$ is simply a δ-regular tree.

An intuitive cop strategy on the trees $X(\Delta,\delta)$ (and in fact on any tree) is to simply move in the unique direction toward the robber, call this strategy CS. There are two intuitively sensible robber strategies on $X(\Delta,\delta)$ which we denote RSA and RSB. In both RSA and RSB, if the robber is in the base tree B, they increase the distance between the two players while staying in B. In a smaller tree T, RSA will keep aiming to increase the distance from the cop by picking a neighbor of $\mathbf{R_s}$ which does not lie on the unique path connecting $\mathbf{R_s}$ to $\mathbf{C_s}$; among the possible choices they pick one uniformly at random. A robber playing strategy RSB will opt to backtrack toward the base tree B, that is, pick the unique neighbor of $\mathbf{R_s}$ that lies closer to the base tree B than $\mathbf{R_s}$. For different values of Δ and δ, we will see that there are certain scenarios in $X(\Delta,\delta)$ where RSA is preferable to RSB and vice versa. To simplify the analysis of the robber strategies we also introduce a cop strategy CSB which is similar to RSB: On the base tree B, a cop playing strategy CSB moves towards the copy of T in which the robber is located (and opts not to do anything, if the cop and the robber are located in the same copy). If the cop is not located in the base tree B, then similarly to a robber playing RSB they move to the unique neighbor of $\mathbf{C_s}$ that lies closer to the base tree B than $\mathbf{C_s}$. We point out that the only time CSB differs from CS is when the two players are on the same copy of the smaller tree T, and the only advantage of CSB over CS is that it is easier to analyse against RSB.

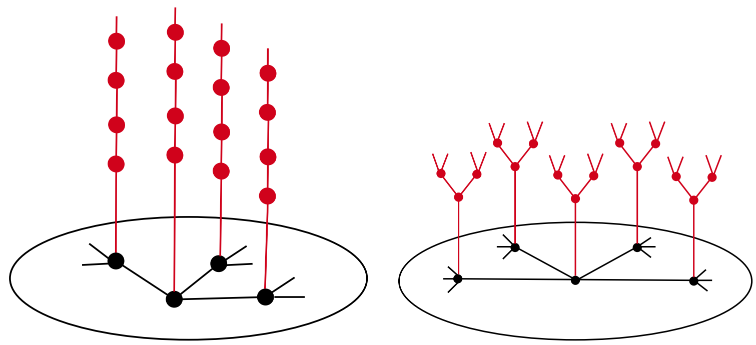

FIGURE 2. The trees $X(4,2)$ and $X(5,3)$.

The main result of this section shows that it depends on the values of Δ, δ, t_c, and t_r, which of the two strategies RSA and RSB is better for the robber. We will give precise bounds on the tipsiness parameters for which the robber is guaranteed to win playing either of the two strategies, see Theorems 4.4 and 4.6. These bounds imply the following result.

Theorem 4.2. Consider the game on $X(\Delta,\delta)$ with $\Delta \geq 4$. There is some $t_0 \in [0, \frac{\delta}{4(\delta-1)}]$ such that

(1) if $t_r \leq t_0$ and the robber wins against every cop strategy by playing RSA for some given value t_c, then they also win by playing RSB, and
(2) if $t_r > t_0$ and the robber wins against every cop strategy by playing RSB for some given value t_c, then they also win by playing RSA.

If $\delta = 2$, then $t_0 = \frac{1}{2}$. If $2\delta \geq \Delta + 1$, then $t_0 = 0$. Otherwise the value of t_0 lies strictly between 0 and $\frac{1}{2}$ and can be determined by solving a quadratic equation.

We note that the restriction $\Delta \geq 4$ only rules out the infinite binary tree $X(3,3)$ and the infinite comb $X(3,2)$. For $X(3,3)$, Lemma 4.1 tells us when each player has a winning strategy. For $X(3,2)$, the robber strategy RSA is winning if $t_c > t_r$ by the same argument as in the proof of Lemma 3.2. If $t_c < t_r$, then CS is easily seen to be winning by a slight modification of the argument in the proof of the second part of Theorem 4.4: the cop will almost surely get to the same small tree as the robber, and once this happens there is a positive probability that the cop will not leave this tree before catching the robber. In particular the small cases for which Theorem 4.2 does not provide an answer are not especially difficult.

4.1. **Analysis of RSA and RSB.** The goal of this section is to provide conditions for when RSA and RSB are winning on $X(\Delta,\delta)$. The main results of this section are Theorem 4.4 which tells us exactly when RSA is winning and Theorem 4.6 which tells us in almost all cases whether RSB is winning.

We start by setting up the notation we use to analyse the game on $X(\Delta,\delta)$. Throughout our analysis, we keep track of the projections of the cop's and robber's position to the base tree, in other words we consider the auxiliary random process $(\mathbf{C}'_\mathbf{s}, \mathbf{R}'_\mathbf{s})$ where $\mathbf{C}'_\mathbf{s}$ is the vertex of B which is closest to $\mathbf{C_s}$, and $\mathbf{R}'_\mathbf{s}$ is the vertex of B which is closest to $\mathbf{R_s}$. Further denote by $T_r(n)$ the (random) time step at which the robber makes their n-th move on the base tree B. Finally, we introduce random variables $F_r(n)$ which take the value 1 if the move at time $T_r(n)$ increases the distance between the cop and the robber and -1 if it decreases the distance. Analogously, define $T_c(n)$ and $F_c(n)$ for the cop. For a cop playing strategy CSB, we also count a sober cop move during which they opt to stay at the same vertex as a move on the base tree.

Our analysis of the game is based on the total change in distance between $\mathbf{C}'_{\mathbf{s}}$ and $\mathbf{R}'_{\mathbf{s}}$ up to time step $\mathbf{s}$, which is equal to

$$Y_{\mathbf{s}} := \sum_{n: T_r(n) < \mathbf{s}} F_r(n) + \sum_{n: T_c(n) < \mathbf{s}} F_c(n). \tag{2}$$

The basic idea behind our analysis is that if this value tends to infinity, then there is a positive probability that the distance between $\mathbf{C}'_{\mathbf{s}}$ and $\mathbf{R}'_{\mathbf{s}}$ and thus also the distance between $\mathbf{C}_{\mathbf{s}}$ and $\mathbf{R}_{\mathbf{s}}$ never reaches zero. Conversely, if the sum does not tend to infinity, then the distance between $\mathbf{C}'_{\mathbf{s}}$ and $\mathbf{R}'_{\mathbf{s}}$ is zero on a positive proportion of the steps. If after infinitely many of these steps the distance between of two players to the base tree B is bounded, then the cop almost surely wins the game because there is a positive probability that the cop wins the game within some bounded number of moves.

In order to bound $Y_{\mathbf{s}}$ it will be useful to bound the individual summands in Equation (2) by two auxiliary sequences $\tilde{F}_r(n)$ and $\tilde{F}_c(n)$ of i.i.d. random variables. These random variables only depend on the outcome of the spinner (that is, what kind of move happens), but not on the specific strategies employed by the cop and the robber. For the definition of $\tilde{F}_r(n)$, imagine that the decision of what kind of move occurs in the game at time $\mathbf{s} = T_r(n)$ is governed by a uniform random variable $S_{\mathbf{s}}$ on $[0, r + \frac{t_r(\Delta-1)}{\Delta}]$ as follows:

- If $0 \le S_{\mathbf{s}} \le r$, we have a sober robber move which will be a move along an edge of B because at time $\mathbf{s} = T_r(n)$ the robber must be located in the base tree.
- If $r < S_{\mathbf{s}} \le r + \frac{t_r(\Delta-2)}{\Delta}$, we have a tipsy robber move on B which increases the distance between the cop and the robber.
- If $S_{\mathbf{s}} > r + \frac{t_r(\Delta-2)}{\Delta}$, we have a tipsy robber move on B which decreases the distance between the cop and the robber if $\mathbf{C}'_{\mathbf{s}} \neq \mathbf{R}'_{\mathbf{s}}$, but increases the distance if $\mathbf{C}'_{\mathbf{s}} = \mathbf{R}'_{\mathbf{s}}$.

All other potential moves (including all cop moves) cannot occur since the move at time $T_r(n)$ must be a robber move along an edge of B. The random variable $\tilde{F}_r(n)$ takes the value 1, if $S_{\mathbf{s}} \le r + \frac{t_r(\Delta-2)}{\Delta}$, and -1 otherwise. We note that

$$\mathbb{P}[\tilde{F}_r(n) = 1] = \frac{r + \frac{t_r(\Delta-2)}{\Delta}}{\frac{1}{2} - \frac{t_r}{\Delta}} = \frac{1 - \frac{4}{\Delta} t_r}{1 - \frac{2}{\Delta} t_r},$$

and consequently

$$\mathbb{E}[\tilde{F}_r(n)] = \frac{1 - \frac{6}{\Delta} t_r}{1 - \frac{2}{\Delta} t_r}.$$

Clearly, the variables $\tilde{F}_r(n)$ are i.i.d., they only depend on the outcome of the spinner. If the robber plays RSA or RSB, then $F_r(n) \ge \tilde{F}_r(n)$ and equality holds if and only if $\mathbf{C}'_{\mathbf{s}} \neq \mathbf{R}'_{\mathbf{s}}$ at time step $\mathbf{s} = T_r(n)$. Moreover, note that $\mathbb{E}[\tilde{F}_r(n)]$ is larger than 0 unless $\Delta = 3$ and $t_r = \frac{1}{2}$.

Analogously, we can define a sequence of i.i.d. random variables $\tilde{F}_c(n)$ taking values in $\{1, -1\}$ with $F_c(n) \ge \tilde{F}_c(n)$ for any cop strategy: $\tilde{F}_c(n)$ takes the value -1 if a sober cop move occurs at time step $\mathbf{s} = T_c(n)$, and is otherwise defined analogously to $\tilde{F}_r(n)$. We note that

$$\mathbb{P}[\tilde{F}_c(n) = 1] = \frac{\frac{t_c(\Delta-2)}{\Delta}}{\frac{1}{2} - \frac{t_c}{\Delta}}$$

and thus

$$\mathbb{E}[\tilde{F}_c(n)] = -\frac{1 - \frac{4\Delta-6}{\Delta} t_c}{1 - \frac{2}{\Delta} t_c}.$$

Moreover, $F_c(n) = \tilde{F}_c(n)$ if the cop plays CS or CSB and $\mathbf{C}'_{\mathbf{s}} \neq \mathbf{R}'_{\mathbf{s}}$ at time step $\mathbf{s} = T_c(n)$.

It follows from the above discussion that if the robber plays RSA or RSB, then

$$Y_{\mathbf{s}} \ge \sum_{n: T_r(n) < \mathbf{s}} \tilde{F}_r(n) + \sum_{n: T_c(n) < \mathbf{s}} \tilde{F}_c(n). \tag{3}$$

This bound has the advantage that the summands are i.i.d.; if their expected values are positive, then it immediately follows from Theorem 2.6 that there is a positive probability that neither of the two sums will ever reach zero which implies that the robber has a positive probability of winning the game. In particular, if the cop is too tipsy, then the robber will win playing either of the two strategies.

Proposition 4.3. Consider the game on $X(\Delta, \delta)$. If $t_c > \frac{\delta}{4\delta-4}$ or $t_c > \frac{\Delta}{4\Delta-6}$ then the robber wins the game playing either RSA or RSB against any cop strategy.

Proof. Note that $\mathbb{E}[\tilde{F}_r(n)] > 0$, so it suffices to show that there is a non-zero probability that contribution of the sum of the $\tilde{F}_c(n)$ (in Equation (3)) is non-negative.

If $t_c > \frac{\delta}{4\delta-4}$, then there is a positive probability that this sum is empty. Indeed, when the cop is not on the base tree, the probability that the next cop move increases the distance between $\mathbf{C_s}$ and $\mathbf{C'_s}$ is at least $2t_c\frac{\delta-1}{\delta} > \frac{1}{2}$; the factor 2 comes from the fact that we have to condition on the move being a cop move. Thus by Theorem 2.6 the probability that the cop never gets to the base tree B is positive. If $t_c > \frac{\Delta}{4\Delta-6}$, then a straightforward calculation shows that $\mathbb{E}[\tilde{F}_c(n)] > 0$, and thus there is a positive probability that the sum of the $\tilde{F}_c(n)$ is always larger than zero. □

We note that neither of the two bounds in Proposition 4.3 generally implies the other; which one is stronger depends on the values of Δ and δ.

Our next goal is to describe precisely when the robber can win by playing RSA. It is not surprising that in this case the game essentially boils down to a game on the δ-regular tree (assuming that the cop is able to get to the same copy of the δ-regular tree as the robber).

Theorem 4.4. If $t_c > \min\{\frac{t_r}{\delta-1}, \frac{\Delta}{4\Delta-6}\}$, then a robber playing RSA has a positive probability of winning, regardless of the cop's strategy. If $t_c \leq \min\{\frac{t_r}{\delta-1}, \frac{\Delta}{4\Delta-6}\}$, then a cop playing CS almost surely wins against a robber playing RSA, provided that $t_r < \frac{1}{2}$, and against a robber playing RSB provided that $t_r < \frac{\delta}{4\delta-4}$.

Proof. If $t_c > \frac{\Delta}{4\Delta-6}$, then by Proposition 4.3 there is a positive probability that the robber wins the game. If $t_c > \frac{t_r}{\delta-1}$, then (assuming $\Delta > \delta$) the probability of increasing the distance between the cop and the robber in any move is at least

$$r + \frac{t_r(\delta-1)}{\delta} + \frac{t_c(\delta-1)}{\delta} \geq \frac{1}{2} - t_r + \frac{t_r(\delta-1)}{\delta} + \frac{t_r}{\delta} + \epsilon = \frac{1}{2} + \epsilon,$$

so the distance can be bounded from below by a biased random walk on $\mathbb{N}_0$ which has a positive probability of never reaching zero by Theorem 2.6.

Now assume that $t_c \leq \min\{\frac{t_r}{\delta-1}, \frac{\Delta}{4\Delta-6}\}$ and either $t_c < \frac{1}{2}$ and the robber plays RSA, or $t_c > \frac{\delta}{4\delta-4}$ and the robber plays RSB. Note that in both cases the probability that the robber moves away from the base tree is larger than the probability that the robber moves towards the base tree. In particular, the robber will almost surely eventually stay in one copy of the small tree which we denote by T_0.

Clearly it is enough to show that the cop will almost surely eventually stay in T_0 as well; in this case the game reduces to a game on a δ-regular tree which the cop almost surely wins by Lemma 4.1.

We start by showing that if the cop is located in a different tree than T_0, then they will almost surely get to the root of T_0 in B in finite time. To prove this, first note that $t_c \leq \frac{\delta}{4\delta-4}$ because $t_r \leq \frac{1}{2}$ and $\delta \geq 2$. Hence the probability that the cop moves towards the root of the δ-regular tree they are in (assuming that the robber is in a different copy) is at least $\frac{1}{2}$, and Theorem 2.6 implies that unless the cop and the robber eventually stay in the same copy of the δ-regular tree, the cop will almost surely return to the base tree B infinitely often. Again by Theorem 2.6, if the probability of the cop moving toward T_0 on a move in B is greater than or equal to the probability of moving further away from T_0, then the cop will almost surely reach T_0 eventually. This happens if and only if $1 - 2t_c + \frac{2t_c}{\Delta} \geq \frac{(\Delta-2)}{\Delta}2t_c$, or equivalently $t_c \leq \frac{\Delta}{4\Delta-6}$.

In order to show that the cop almost surely eventually remains in T_0, it suffices to show that every time the cop moves from B to the root of T_0, the probability of them never visiting B again is at least some $\epsilon > 0$.

We may condition on the event that the distance between $\mathbf{R_s}$ and B is at least $k \cdot \mathbf{s}$ for every k, since the probability of this event is positive by Theorem 2.9. In other words, without loss of generality we can assume that

$$\mathbb{P}[\forall \mathbf{s} > 0 : d(\mathbf{R_s}, B) > k \cdot \mathbf{s}] = 1.$$

Denote by $P_\mathbf{s}$ the unique path connecting the root of T_0 to $\mathbf{R_s}$, and denote by $\mathbf{C''_s}$ the vertex closest to $\mathbf{C_s}$ on this path. If $\mathbf{C_s} = \mathbf{C''_s}$, then the probability of increasing the distance between $\mathbf{C''_s}$ and B is $c + \frac{t_c}{\delta}$, the probability of decreasing this distance is $\frac{t_c}{\delta}$, and the probability of the distance remaining unchanged is $\frac{t_c(\delta-2)}{\delta}$. If $\mathbf{C_s} \neq \mathbf{C''_s}$ then the distance can only change if $\mathbf{R_s} = \mathbf{C''_s}$, in which case the distance is at least

$k \cdot \mathbf{s}$. Hence the distance between $\mathbf{C}'_{\mathbf{s}}$ and B is a random walk on $\mathbb{N}$ whose probability of moving away from 0 is strictly larger than its probability of moving towards 0, and thus it has a positive probability of never returning to 0. $\square$

In the remainder of this subsection we analyse RSB. The case $t_r > \frac{\delta}{4\delta-4}$ where the robber is too tipsy to get back to the base tree an infinite number of times is already covered in the previous theorem; we may thus focus on $t_r \leq \frac{\delta}{4\delta-4}$, and we first consider the case where this inequality is strict. We start by analyzing the number of steps in between consecutive cop and robber moves on the base tree.

Lemma 4.5. Consider the game on $X(\Delta, \delta)$. If $t_r < \frac{\delta}{4\delta-4}$, then with the above notation, we have

$$\mathbb{E}[T_r(n) - T_r(n-1)] \geq 2 \cdot \frac{1 - \left(\frac{4\delta-4}{\delta} - \frac{2}{\Delta}\right) t_r}{\left(1 - \frac{4\delta-4}{\delta} t_r\right)\left(1 - \frac{2}{\Delta} t_r\right)}$$

with equality if the robber plays RSB. Additionally, if the robber plays RSB, then $T_r(n) - T_r(n-1)$ has finite variance.

Similarly, if $t_c < \frac{\delta}{4\delta-4}$, then

$$\mathbb{E}[T_c(n) - T_c(n-1)] \geq 2 \cdot \frac{1 - \left(\frac{4\delta-4}{\delta} - \frac{2}{\Delta}\right) t_c}{\left(1 - \frac{4\delta-4}{\delta} t_c\right)\left(1 - \frac{2}{\Delta} t_c\right)}$$

with equality if the cop plays CSB. Additionally, if the cop plays CSB, then $T_c(n) - T_c(n-1)$ has finite variance.

Proof. We start by showing that RSB minimizes $\mathbb{E}[T_r(n) - T_r(n-1)]$. An easy coupling argument shows that RSB minimizes the time it takes the robber to return to the base tree whenever they leave the base tree: consider two games where in one the robber plays RSB and in the other they play an arbitrary strategy, and assume that every sober or tipsy robber move in one game is also a sober or tipsy robber move in the other. Then clearly the robber playing RSB is always at most as far away from B as the robber in the other game. Moreover, when the robber is at a vertex of the base tree, then RSB maximizes the probability that the next robber step will be along an edge of B.

For the remainder of the proof assume that the robber plays RSB. When the distance of the robber from the base tree is larger than 0, then the probability of the robber moving away from the base tree is $2\frac{t_r(\delta-1)}{\delta}$, when it is equal to 0, then the probability of the next move leading away from the base tree is $2\frac{t_r}{\Delta}$. The factor 2 in both of these probabilities is due to the fact that we have to condition on the event that the move is actually a robber move which has probability $\frac{1}{2}$. In particular, if $X_{\mathbf{s}}$ is a random walk on $\mathbb{N}$ started at $X_0 = 1$ with transition probabilities $p_{i,i+1} := 2\frac{t_r(\delta-1)}{\delta}$ for $i > 1$ and $p_{1,2} = 2\frac{t_r}{\Delta}$, then the random variable $U_1 := \min\{\mathbf{s} \mid X_{\mathbf{s}} = 0\}$ has the same distribution as the number of robber moves occurring between $T_r(n-1)$ and $T_r(n)$.

Since $2\frac{t_r(\delta-1)}{\delta} < \frac{1}{2}$, Proposition 2.8 implies that U_1 has finite expectation and variance, and hence the same is true for $T_r(n) - T_r(n-1)$. To determine $\mathbb{E}[U_1]$, note that like in the proof of Proposition 2.8

$$\mathbb{E}[U_1] = 1 + 2\frac{t_r}{\Delta}(\mathbb{E}[U_1] + \mathbb{E}[U_2]),$$

where U_2 is the number of steps until a random walk with the same transition probabilities as above started at $X_0 = 2$ reaches $X_{\mathbf{s}} = 1$ for the first time. By Theorem 2.6, we know that $\mathbb{E}[U_2] = \frac{1}{1-4\frac{t_r(\delta-1)}{\delta}}$, and solving the above equation for $\mathbb{E}[T_1]$ gives

$$\mathbb{E}[U_1] = \frac{1 - \left(\frac{4\delta-4}{\delta} - \frac{2}{\Delta}\right) t_r}{\left(1 - \frac{4\delta-4}{\delta} t_r\right)\left(1 - \frac{2}{\Delta} t_r\right)}.$$

Since on average every second move is a cop move, $\mathbb{E}[U_1] = \frac{1}{2}\mathbb{E}[T_r(n) - T_r(n-1)]$. This completes the proof of the first part.

For the second part note that if $t_c < \frac{\delta}{4\delta-4}$, then the exact same arguments apply for the cop. $\square$

Recall that $\mathbf{C}'_{\mathbf{s}}$ and $\mathbf{R}'_{\mathbf{s}}$ are projections of the cop's and robber's position to the base tree B. If the $\mathbf{C}'_{\mathbf{s}}$ is never equal to $\mathbf{R}'_{\mathbf{s}}$, then the cop obviously never reaches the same path as the robber on $X(\Delta, \delta)$, and thus they cannot win the game. The next result shows that this is essentially the only obstruction to CSB winning against RSB. We remark that a refinement of the argument given below shows that CS is also winning against RSB in the second case. Indeed, we conjecture that if the cop has a strategy to win the game on a tree, then they can always win the game by playing CS.

Theorem 4.6. Consider the game on $X(\Delta,\delta)$. Assume that $t_r \leq \frac{\delta}{4\delta-4}$ and $t_c \leq \min\left\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\right\}$. If

$$\frac{\left(1-\frac{4\delta-4}{\delta}t_r\right)\left(1-\frac{6}{\Delta}t_r\right)}{1-\left(\frac{4\delta-4}{\delta}-\frac{2}{\Delta}\right)t_r} - \frac{\left(1-\frac{4\delta-4}{\delta}t_c\right)\left(1-\frac{4\Delta-6}{\Delta}t_c\right)}{1-\left(\frac{4\delta-4}{\delta}-\frac{2}{\Delta}\right)t_c} > 0 \tag{4}$$

then the robber wins playing RSB against any cop strategy.

If the reverse inequality holds:

$$\frac{\left(1-\frac{4\delta-4}{\delta}t_r\right)\left(1-\frac{6}{\Delta}t_r\right)}{1-\left(\frac{4\delta-4}{\delta}-\frac{2}{\Delta}\right)t_r} - \frac{\left(1-\frac{4\delta-4}{\delta}t_c\right)\left(1-\frac{4\Delta-6}{\Delta}t_c\right)}{1-\left(\frac{4\delta-4}{\delta}-\frac{2}{\Delta}\right)t_c} < 0 \tag{5}$$

then CSB is winning against a robber playing RSB; in particular RSB does not win against every cop strategy.

Proof. We split the proof in the following cases:

(1) $t_r < \frac{\delta}{4\delta-4}$ and $t_c < \min\left\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\right\}$, and (4) holds,
(2) $t_r < \frac{\delta}{4\delta-4}$ and $t_c < \min\left\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\right\}$, and (5) holds,
(3) $t_r = \frac{\delta}{4\delta-4}$ and $t_c < \min\left\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\right\}$, in this case inequality (5) holds, and
(4) $t_r < \frac{\delta}{4\delta-4}$ and $t_c = \min\left\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\right\}$, in this case inequality (4) holds.

Note that we do not need to consider the case $t_r = \frac{\delta}{4\delta-4}$ and $t_c = \min\left\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\right\}$, because in this case neither of the two inequalities (4) and (5) is satisfied.

For the first two cases, we start by giving estimates on the number of summands in each of the sums in the definition of $Y_{\mathbf{s}}$ in Equation (2). For $t \in \{t_r, t_c\}$, let $\mu(t) := \frac{2\delta\Delta-(8\delta\Delta-4\delta-8\Delta)t}{\delta-(4(\delta-1)t)(\Delta-2t)}$. Since we may assume that the robber plays RSB, Lemma 4.5 implies that $T_r(n) - T_r(n-1)$ are i.i.d. random variables with finite mean $\mu(t_r)$ and finite variance. Thus by Theorem 2.5, for every $\epsilon > 0$ we can find a constant k such that

$$\mathbb{P}\left[n\left(\mu(t_r)-\epsilon\right) < T_r(n) < n\left(\mu(t_r)+\epsilon\right)\right] > 1 - e^{-kn}$$

for every large enough n. Substituting $n = \frac{\mathbf{s}}{\mu(t_r)\pm\epsilon}$ and using continuity of $f(x) = \frac{1}{x}$ at $x = \mu(t_r)$, this implies that for every $\epsilon > 0$ we can find a constant k (different to the constant k above) such that

$$\mathbb{P}\left[T_r\left(\mathbf{s}\left(\frac{1}{\mu(t_r)}-\epsilon\right)\right) < \mathbf{s} < T_r\left(\mathbf{s}\left(\frac{1}{\mu(t_r)}+\epsilon\right)\right)\right] > 1 - e^{-k\mathbf{s}}$$

for every large enough $\mathbf{s}$. In other words, the probability that number of summands in the first sum in the definition of $Y_{\mathbf{s}}$ deviates by more than $\mathbf{s}\epsilon$ from its expected value $\frac{\mathbf{s}}{\mu(t_r)}$ is less than $e^{-k\mathbf{s}}$.

If the cop plays CSB, then an analogous argument shows that for every ϵ there is a constant k such that

$$\mathbb{P}\left[T_c\left(\mathbf{s}\left(\frac{1}{\mu(t_c)}-\epsilon\right)\right) < \mathbf{s} < T_c\left(\left(\frac{1}{\mu(t_c)}+\epsilon\right)\right)\right] > 1 - e^{-k\mathbf{s}}.$$

If we allow the cop to play an arbitrary strategy, then the value $T_c(n) - T_c(n-1)$ will only increase compared to a cop playing CSB, whence we still obtain one side of the above concentration result:

$$\mathbb{P}\left[T_c\left(\mathbf{s}\left(\frac{1}{\mu(t_c)}-\epsilon\right)\right) < \mathbf{s}\right] > 1 - e^{-k\mathbf{s}},$$

in other words, the probability that the number of summands exceeds $\frac{\mathbf{s}}{\mu(t_c)}$ by more than $\mathbf{s}\epsilon$ is less than $e^{-k\mathbf{s}}$.

Now assume that we are in Case 1, and thus the bound in (4) holds. We claim that for any $\epsilon' > 0$ we can find an $\epsilon > 0$ such that the probability that the following inequality is violated for any given $\mathbf{s}$ is bounded above by $ae^{-b\mathbf{s}}$ for suitable $a, b > 0$. That is

$$
\begin{aligned}
\frac{Y_{\mathbf{s}}}{\mathbf{s}} &\geq \left(\sum_{n:T_r(n)<\mathbf{s}} \frac{\tilde{F}_r(n)}{\mathbf{s}}\right) + \left(\sum_{n:T_c(n)<\mathbf{s}} \frac{\tilde{F}_c(n)}{\mathbf{s}}\right) && (6)\\
&\geq \frac{\max\{n : T_r(n) \leq \mathbf{s}\}}{\mathbf{s}} \left(\mathbb{E}[\tilde{F}_r(1)] - \epsilon\right) + \frac{\max\{n : T_c(n) \leq \mathbf{s}\}}{\mathbf{s}} \left(\mathbb{E}[\tilde{F}_c(1)] - \epsilon\right) && (7)\\
&\geq \left(\frac{1}{\mu(t_r)} - \epsilon\right) \left(\mathbb{E}[\tilde{F}_r(1)] - \epsilon\right) + \left(\frac{1}{\mu(t_c)} + \epsilon\right) \left(\mathbb{E}[\tilde{F}_c(1)] - \epsilon\right) && (8)\\
&\geq \frac{\mathbb{E}[\tilde{F}_r(1)]}{\mu(t_r)} + \frac{\mathbb{E}[\tilde{F}_c(1)]}{\mu(t_c)} - \epsilon' && (9)\\
&= \frac{1}{2} \cdot \frac{\left(1 - \frac{4\delta-4}{\delta} t_r\right)\left(1 - \frac{6}{\Delta} t_r\right)}{1 - \left(\frac{4\delta-4}{\delta} - \frac{2}{\Delta}\right) t_r} - \frac{1}{2} \cdot \frac{\left(1 - \frac{4\delta-4}{\delta} t_c\right)\left(1 - \frac{4\Delta-6}{\Delta} t_c\right)}{1 - \left(\frac{4\delta-4}{\delta} - \frac{2}{\Delta}\right) t_c} - \epsilon'. && (10)
\end{aligned}
$$

Note that the inequality in (6) follows from the fact that $F_r(n) \geq \tilde{F}_r(n)$, and $F_c(n) \geq \tilde{F}_c(n)$, and the inequality in (7) follows from Theorem 2.5. For the inequality in (8) we note that $\mathbb{E}[\tilde{F}_r(1)] - \epsilon > 0$ whenever ϵ is small enough and that $\mathbb{E}[\tilde{F}_c(1)] - \epsilon < 0$ because $t_c < \frac{\Delta}{4\Delta-6}$; the inequality then follows from the bounds on the number of summands discussed above. For the last inequality in (9) we note that none of the coefficients are unbounded, and thus we can indeed find a suitable ϵ for any ϵ'.

If ϵ' is small enough, then the expression in (10) is larger than 0. Since $\sum_{\mathbf{s}\geq 1} ae^{-b\mathbf{s}} < \infty$, we conclude that there is some $\mathbf{s}_0$ such that the probability that $Y_{\mathbf{s}} > 0$ for every $\mathbf{s} > \mathbf{s}_0$ is positive. If the distance between $\mathbf{R}'_0$ and $\mathbf{C}'_0$ is larger than $\mathbf{s}_0$, then this implies that there is a positive probability that $\mathbf{R}'_{\mathbf{s}}$ and $\mathbf{C}'_{\mathbf{s}}$ never coincide, and thus the robber wins the game with positive probability.

Next consider Case 2. For $\mathbf{s} > 0$, we denote by $K(\mathbf{s})$ the number of $\mathbf{t} \leq \mathbf{s}$ such that $\mathbf{C}'_{\mathbf{t}} = \mathbf{R}'_{\mathbf{t}}$. If the robber plays RSB, then $F_r(n) > \tilde{F}_r(n)$ if and only if $\mathbf{C}'_{\mathbf{s}} = \mathbf{R}'_{\mathbf{s}}$ at time $\mathbf{s} = T_r(n)$, and in this case $F_r(n) - \tilde{F}_r(n) \leq 2$. Similarly, if the cop plays CSB, then $F_c(n) > \tilde{F}_c(n)$ if and only if $\mathbf{C}'_{\mathbf{s}} = \mathbf{R}'_{\mathbf{s}}$ at time $\mathbf{s} = T_c(n)$, and in this case $F_c(n) - \tilde{F}_c(n) \leq 2$. In particular,

$$
\frac{Y_{\mathbf{s}}}{\mathbf{s}} \leq \left(\sum_{n:T_r(n)<\mathbf{s}} \frac{\tilde{F}_r(n)}{\mathbf{s}}\right) + \left(\sum_{n:T_c(n)<\mathbf{s}} \frac{\tilde{F}_c(n)}{\mathbf{s}}\right) + \frac{2K(\mathbf{s})}{\mathbf{s}},
$$

and a similar calculation as above yields that the inequality

$$
\frac{Y_{\mathbf{s}}}{\mathbf{s}} \leq \frac{1}{2} \cdot \frac{\left(1 - \frac{4\delta-4}{\delta} t_r\right)\left(1 - \frac{6}{\Delta} t_r\right)}{1 - \left(\frac{4\delta-4}{\delta} - \frac{2}{\Delta}\right) t_r} - \frac{1}{2} \cdot \frac{\left(1 - \frac{4\delta-4}{\delta} t_c\right)\left(1 - \frac{4\Delta-6}{\Delta} t_c\right)}{1 - \left(\frac{4\delta-4}{\delta} - \frac{2}{\Delta}\right) t_c} + \epsilon' + \frac{2K(\mathbf{s})}{\mathbf{s}} \tag{11}
$$

is satisfied with probability $1 - ae^{-b\mathbf{s}}$. Since the distance between $\mathbf{C}'_{\mathbf{s}}$ and $\mathbf{R}'_{\mathbf{s}}$ is never negative, there is a constant lower bound on $Y_{\mathbf{s}}$. Hence choosing ϵ' sufficiently small and taking the limit $\mathbf{s} \to \infty$ in inequality (11), we conclude that almost surely

$$
\liminf_{\mathbf{s}\to\infty} \frac{K(\mathbf{s})}{\mathbf{s}} > 0.
$$

In fact, it is not hard to see that $\lim_{\mathbf{s}\to\infty} \frac{K(\mathbf{s})}{\mathbf{s}}$ exists, but we will not need this.

Let $L(n, \mathbf{s})$ denote the number of time steps $\mathbf{t} \leq \mathbf{s}$ for which $\mathbf{C}'_{\mathbf{t}} = \mathbf{R}'_{\mathbf{t}}$, and both $\mathbf{C}_{\mathbf{t}}$ and $\mathbf{R}_{\mathbf{t}}$ are at distance at most n from $\mathbf{C}'_{\mathbf{t}}$. We note that the distances between $\mathbf{R}_{\mathbf{s}}$ and $\mathbf{R}'_{\mathbf{s}}$ and between $\mathbf{C}_{\mathbf{s}}$ and $\mathbf{C}'_{\mathbf{s}}$ can be described by positive recurrent biased random walks on $\mathbb{N}_0$, respectively. By Theorem 2.10, and the above estimate on $\liminf_{\mathbf{s}\to\infty} \frac{K(\mathbf{s})}{\mathbf{s}}$ there is an $n_0 \in \mathbb{N}$ such that almost surely

$$
\liminf_{\mathbf{s}\to\infty} \frac{L(n_0, \mathbf{s})}{\mathbf{s}} > 0.
$$

This in particular implies that $L(n_0, \mathbf{s})$ tends to infinity as $\mathbf{s} \to \infty$, and therefore there are almost surely infinitely many time steps in which the distance between $\mathbf{C}_{\mathbf{s}}$ and $\mathbf{R}_{\mathbf{s}}$ is upper bounded by $2n_0$. Every time this happens, the probability that the cop will win in the subsequent $2n_0$ steps is larger than the probability that the cop wins by a sequence of random moves starting at distance $2n_0$ which in turn is lower bounded by $(\frac{t_c}{\Delta})^{2n_0}$. Therefore the cop will almost surely eventually win the game.

In Case 3, we first note that

$$\lim_{\mathbf{s}\to\infty}\left|\sum_{n:T_r(n)<\mathbf{s}}\frac{\tilde{F}_r(n)}{\mathbf{s}}\right|=0$$

because increasing t_r also increases $T_r(n)$. An analogous argument as in Case 2 shows that it will happen infinitely many times that $\mathbf{C_s}=\mathbf{R}'_\mathbf{s}$, but unfortunately we cannot bound the distance between $\mathbf{R_s}$ and $\mathbf{R}'_\mathbf{s}$ in the same way as in Case 2. However, we claim that $L(n_0,\mathbf{s})$ will still tend to infinity as $\mathbf{s}$ goes to infinity.

To prove this claim it is enough to show that if $d(\mathbf{C}_0,\mathbf{R}'_0)=0$, then there is a constant lower bound on the probability that there is some finite $\mathbf{s}$ such that $\mathbf{C}'_\mathbf{s}=\mathbf{R}'_\mathbf{s}$, and both $\mathbf{C_s}$ and $\mathbf{R_s}$ are at distance at most n from $\mathbf{C}'_\mathbf{s}$. Consider the two random processes $d(\mathbf{C_s},\mathbf{R}'_\mathbf{s})$ and $d(\mathbf{R_s},\mathbf{R}'_\mathbf{s})$. As long as $d(\mathbf{R_s},\mathbf{R}'_\mathbf{s})>0$, these random processes can be treated as independent random walks on $\mathbb{N}$. Since $d(\mathbf{C_s},\mathbf{R}'_\mathbf{s})$ is positive recurrent and $d(\mathbf{C}_0,\mathbf{R}'_0)=0$, the probability that $d(\mathbf{C_s},\mathbf{R}'_\mathbf{s})<n_0$ is bounded from below by some constant for every $\mathbf{s}>0$. This also holds at the (random) time step when $d(\mathbf{R_s},\mathbf{R}'_\mathbf{s})$ becomes smaller than n_0 for the first time which completes the proof of the claim. Hence $L(n_0,\mathbf{s})$ goes to infinity, and thus CSB is winning against RSB in this case.

Finally consider Case 4. Note that a cop with tipsiness $t_c=\min\{\frac{\delta}{4\delta-4},\frac{\Delta}{4\Delta-6}\}-\epsilon$ can pretend that $t_c=\min\{\frac{\delta}{4\delta-4},\frac{\Delta}{4\Delta-6}\}$ by playing a strategy in which they make a random move with probability ϵ. By Case 1, strategy RSB is winning against any cop strategy (and thus also against this specific strategy) if $t_c=\min\{\frac{\delta}{4\delta-4},\frac{\Delta}{4\Delta-6}\}-\epsilon$ for small enough ϵ. □

4.2. **Comparison of the two strategies.** The aim of this section is to compare the two strategies RSA and RSB against one another. Throughout this section we will assume that $t_r<\frac{1}{2}$ because for $t_r=\frac{1}{2}$ all robber moves are random, and thus there will not be any difference between the two robber strategies.

We first show that (apart from isolated cases) we may assume that $0<t_c<\min\{\frac{\delta}{4\delta-4},\frac{\Delta}{4\Delta-6}\}$ and that $0<t_r<\frac{\delta}{4\delta-4}$. Proposition 4.3 tells us that if $t_c>\min\{\frac{\delta}{4\delta-4},\frac{\Delta}{4\Delta-6}\}$ then the robber wins if they play each of the two strategies RSA or RSB. Our first result tells us that if $t_r>\frac{\delta}{4\delta-4}$ then the robber should always choose strategy RSA over strategy RSB.

Proposition 4.7. If $t_r>\frac{\delta}{4\delta-4}$ and the cop has a winning strategy against a robber playing RSA, then they also have a winning strategy against a robber playing RSB; in this case CS is winning against both RSA and RSB.

Proof. If the cop has a winning strategy against a robber playing RSA, then by Theorem 4.4 we know that $t_c\le\min\{\frac{t_r}{\delta-1},\frac{\Delta}{4\Delta-4}\}$, which means that CS is winning against RSB provided that $t_r>\frac{\delta}{4\delta-4}$. □

Hence we can assume that $0<t_c\le\min\{\frac{\delta}{4\delta-4},\frac{\Delta}{4\Delta-6}\}$, and that $0<t_r\le\frac{\delta}{4\delta-4}$. The next proposition deals with the cases where equality holds in one of the two bounds.

Proposition 4.8. Consider the game on $X(\Delta,\delta)$, and assume that $\Delta\ge 4$.

(1) If $t_c=0$, then CS is always a winning strategy for the cop.
(2) If $t_r=0$ and $t_c\neq 0$, then both RSA and RSB are winning robber strategies.
(3) If $t_r=\frac{\delta}{4\delta-4}$ and RSB is winning against every cop strategy, then so is RSA.

Proof. The first two statements follow from Proposition 2.7 by observing that the probability that the distance between the cop and the robber increases is at most $\frac{1}{2}$ in the first case, and strictly larger than $\frac{1}{2}$ in the second. For the third statement, recall that we are working under the assumption that $t_r<\frac{1}{2}$ and hence $\delta\ge 3$. If RSA is not winning, then $t_c\le\frac{\delta}{4(\delta-1)^2}$ by Theorem 4.4. In particular, $t_c<\frac{1}{4}\le\min\{\frac{\delta}{4\delta-4},\frac{\Delta}{4\Delta-6}\}$ and thus RSB is not winning against CSB by Theorem 4.6. □

The next lemma puts the winning conditions from Theorem 4.6 for a robber playing RSB in a more useful form.

Lemma 4.9. Let $\Delta\ge 4$. There is a function

$$f\colon\left[0,\frac{\delta}{4-4\delta}\right]\to\left[0,\min\left\{\frac{\delta}{4\delta-4},\frac{\Delta}{4\Delta-6}\right\}\right]$$

such that RSB is winning against any cop strategy if $f(t_r) < t_c \le \min\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\}$, and CSB is winning against RSB if $0 \le t_c < f(t_r)$.

This function is monotonically increasing and strictly convex, $f(0) = 0$, $f(\frac{\delta}{4-4\delta}) = \min\left\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\right\}$, and $f'(0) = \frac{2}{\Delta-1}$.

Proof. If we let $G(x) = \frac{(1-ax)(1-bx)}{(1-cx)}$ then $G'(x) = -\frac{b(1-ax)}{1-cx} - \frac{(a-c)(1-bx)}{(1-cx)^2}$ and $G''(x) = \frac{2(a-c)(b-c)}{(1-cx)^3}$ provided that $x \ne c$.

For fixed x, δ, and Δ this implies that the expression

$$F(x,y) = \frac{\left(1 - \frac{4\delta-4}{\delta}x\right)\left(1 - \frac{6}{\Delta}x\right)}{1 - \left(\frac{4\delta-4}{\delta} - \frac{2}{\Delta}\right)x} - \frac{\left(1 - \frac{4\delta-4}{\delta}y\right)\left(1 - \frac{4\Delta-6}{\Delta}y\right)}{1 - \left(\frac{4\delta-4}{\delta} - \frac{2}{\Delta}\right)y}$$

is monotonically increasing for y in the interval $\left[0, \min\left\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\right\}\right]$. It is not hard to check that $F(x,0) \le 0$ and $F\left(x, \min\left\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\right\}\right) \ge 0$ for $x \in \left[0, \frac{\delta}{4-4\delta}\right]$. Hence the equation $F(x, f(x)) = 0$ defines a unique function $f(x)$. By Theorem 4.6 this function has the property that if $f(t_r) < t_c \le \min\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\}$, then RSB is winning against any cop strategy, and if $0 \le t_c < f(t_r)$, then CSB is winning against RSB.

To see that f is monotonically increasing, recall that

$$f'(x_0) = -\frac{F_x}{F_y},$$

where the partial derivatives are evaluated at the point $(x_0, f(x_0))$. The partial derivatives have the form $\pm G'$ from above for suitable values of a, b, and c (depending on δ and Δ), and it is not hard to see that $F_x(x_0, y_0)$ is positive and $F_y(x_0, y_0)$ is negative for all pairs

$$(x_0, y_0) \in \left(0, \frac{\delta}{4\delta-4}\right) \times \left(0, \min\left\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\right\}\right).$$

To show convexity, recall that

$$f''(x_0) = \frac{-F_y^2 F_{xx} + 2F_x F_y F_{xy} - F_x^2 F_{yy}}{F_y^3}.$$

Note that the mixed second derivative F_{xy} is always zero and that the denominator F_y^3 is always negative. Moreover F_{xx} and F_{yy} have the form $\pm G''$ from above for suitable values of a, b, and c, and again it is not hard to see that for all values of (x_0, y_0) in our range we have $F_{xx}(x_0, y_0) \ge 0$ (with equality if $\delta = 2$ and $\Delta = 4$, here we are using the assumption that $\Delta \ge 4$), and $F_{yy}(x_0, y_0) > 0$.

Finally note that

$$F(0,0) = F\left(\frac{\delta}{4\delta-4}, \min\left\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\right\}\right) = 0,$$

and thus $f(0) = 0$, and $f(\frac{\delta}{4-4\delta}) = \min\left\{\frac{\delta}{4\delta-4}, \frac{\Delta}{4\Delta-6}\right\}$ as claimed. Moreover, evaluating $f'(x_0) = -\frac{F_x}{F_y}$ at $x_0 = f(x_0) = 0$ gives $\frac{2}{\Delta-1}$. ☐

We are now ready to prove Theorem 4.2. Before we do so, let us recall the statement of this theorem.

Theorem 4.2. Consider the game on $X(\Delta, \delta)$ with $\Delta \ge 4$. There is some $t_0 \in [0, \frac{\delta}{4(\delta-1)}]$ such that

(1) if $t_r \le t_0$ and the robber wins against every cop strategy by playing RSA for some given value t_c, then they also win by playing RSB, and
(2) if $t_r > t_0$ and the robber wins against every cop strategy by playing RSB for some given value t_c, then they also win by playing RSA.

If $\delta = 2$, then $t_0 = \frac{1}{2}$. If $2\delta \ge \Delta + 1$, then $t_0 = 0$. Otherwise the value of t_0 lies strictly between 0 and $\frac{1}{2}$ and can be determined by solving a quadratic equation.

Proof. First note that if $t_r = 0$, then by Proposition 4.8 the strategies RSA and RSB are both winning against any cop strategy for $t_c > 0$ and both losing against CS for $t_c = 0$. If $t_r > \frac{\delta}{4\delta-4}$ and RSB is winning for a given value t_c, then by Proposition 4.7 so is RSA. In particular, if the theorem is true for $t_r \in (0, \frac{\delta}{4\delta-4}]$, then it is true for $t_r \in [0, \frac{1}{2}]$.

Throughout the rest of this proof, assume that $0 < t_r \leq \frac{\delta}{4\delta-4}$. Denote by $f(x)$ the function provided by Lemma 4.9, and let $g(x) = \frac{x}{\delta-1}$. Observe that $f(0) = g(0) = 0$. By Theorem 4.4, RSA is winning against every possible cop strategy if and only if $t_c > g(t_r)$, otherwise it loses against CS. By Lemma 4.9, RSB is winning against every cop strategy if $t_c > f(t_r)$ and losing against CSB if $t_c < f(t_r)$.

If $\delta = 2$, then $\frac{\delta}{4\delta-4} = \frac{1}{2}$. By Lemma 4.9 we have that $f(\frac{1}{2}) = \min\{\frac{1}{2}, \frac{\Delta}{4\Delta-6}\}$ and since $\Delta \geq 4$ this is strictly less than $g(\frac{1}{2}) = \frac{1}{2}$. By convexity of f this implies that $f(x) < g(x)$ for $x \in (0, \frac{1}{2})$ and thus if RSA is winning against every cop strategy for a given value t_c, then so is RSB.

Now assume that $\delta \geq 3$. In this case $f(\frac{\delta}{4\delta-4}) > \frac{1}{4} \geq g(\frac{\delta}{4\delta-4})$. If $2\delta \geq \Delta + 1$, then $f'(0) \geq g'(0)$ and strict convexity of f implies that $f(x) > g(x)$ for every $x \in (0, \frac{\delta}{4\delta-4}]$. This in particular implies that if $t_r > 0$ and RSB is winning against every cop strategy for a given value t_c, then so is RSA. Finally, if $2\delta < \Delta + 1$, then $f'(0) < g'(0)$, and consequently there is a unique $t_0 \in (0, \frac{\delta}{4\delta-4}]$ such that $f(x) \leq g(x)$ for $x \leq t_0$ and $f(x) > g(x)$ for $x > t_0$. This implies that if $t_r \leq t_0$ and RSA is winning against every cop strategy for a given value t_c, then so is RSB, and if $t_r > t_0$ and RSB is winning against every cop strategy for a given value t_c, then so is RSA.

Recall that $f(x)$ is the unique solution to the equation $F(x, f(x)) = 0$, where

$$F(x,y) = \frac{\left(1 - \frac{4\delta-4}{\delta}x\right)\left(1 - \frac{6}{\Delta}x\right)}{1 - \left(\frac{4\delta-4}{\delta} - \frac{2}{\Delta}\right)x} - \frac{\left(1 - \frac{4\delta-4}{\delta}y\right)\left(1 - \frac{4\Delta-6}{\Delta}y\right)}{1 - \left(\frac{4\delta-4}{\delta} - \frac{2}{\Delta}\right)y}.$$

Since t_0 is a solution of $f(x) = g(x)$, we can find its value by solving the equation $F(x, \frac{x}{\delta-1}) = 0$ for x, which gives a quadratic equation in x as claimed. □

In Figure 3 we compare the efficacy of the robber's dualing strategies on $X(\Delta, \delta)$ for $\Delta = 7$ and $\delta \in \{2, 3, 4, 5\}$. The dashed lines are at $t_c = \frac{\delta}{4\delta-4}$ and $t_r = \frac{\delta}{4\delta-4}$, respectively; the dotted line is at $t_c = \frac{\Delta}{4\Delta-6}$. The red and blue lines are the graphs of the functions f and g from the proof of Theorem 4.2; if for $t_r < \frac{\delta}{4\delta-4}$ the point (t_r, t_c) lies above the blue line, then strategy RSA is winning against every cop strategy, and if it lies above the red line RSB is winning against every cop strategy.

5. Future Directions and Open Problems

It seems intuitively clear that on any tree, the cop's optimal strategy is to always move in the direction of the robber, but we were unable to provide a proof nor could we find a counterexample through the tools we used in our arguments. Hence we pose the following problems, which might be solved via a different perspective or with a different set of mathematical tools than those used in this manuscript.

Problem 5.1. Prove or provide a counterexample to the following statement: On any tree, the cop's optimal strategy is to always move in the direction of the robber.

Problem 5.2. Prove or provide a counterexample to the following statement: On any graph, if the cop has a winning strategy, then they also have a winning strategy which is "monotone" in the sense that the cop never purposefully increases the distance between them and the robber.

We conjecture that the statement in Problem 5.1 is true, and there is likely a counterexample to the statement in Problem 5.2.

On the grid graph, we have shown that the cop's choice of strategy has an impact on whether they win when $c > r$, but that if the robber employs any strategy that increases the distance between them and the cop, then they are equally likely to win when $r > c$. On the $X(\Delta, \delta)$ trees, the cop's intuitive strategy is to decrease the distance between them and the robber in each step, but the robber's strategy depends on conditions relating the proportion of sober cop and robber moves to Δ and δ. These examples lead us to ask, whether there any graphs that restrict the strategies of both players at the same time.

Problem 5.3. Are there graphs where the cop and robber each have multiple strategies they could employ where each of these strategies is clearly better depending on specific conditions on c and r?

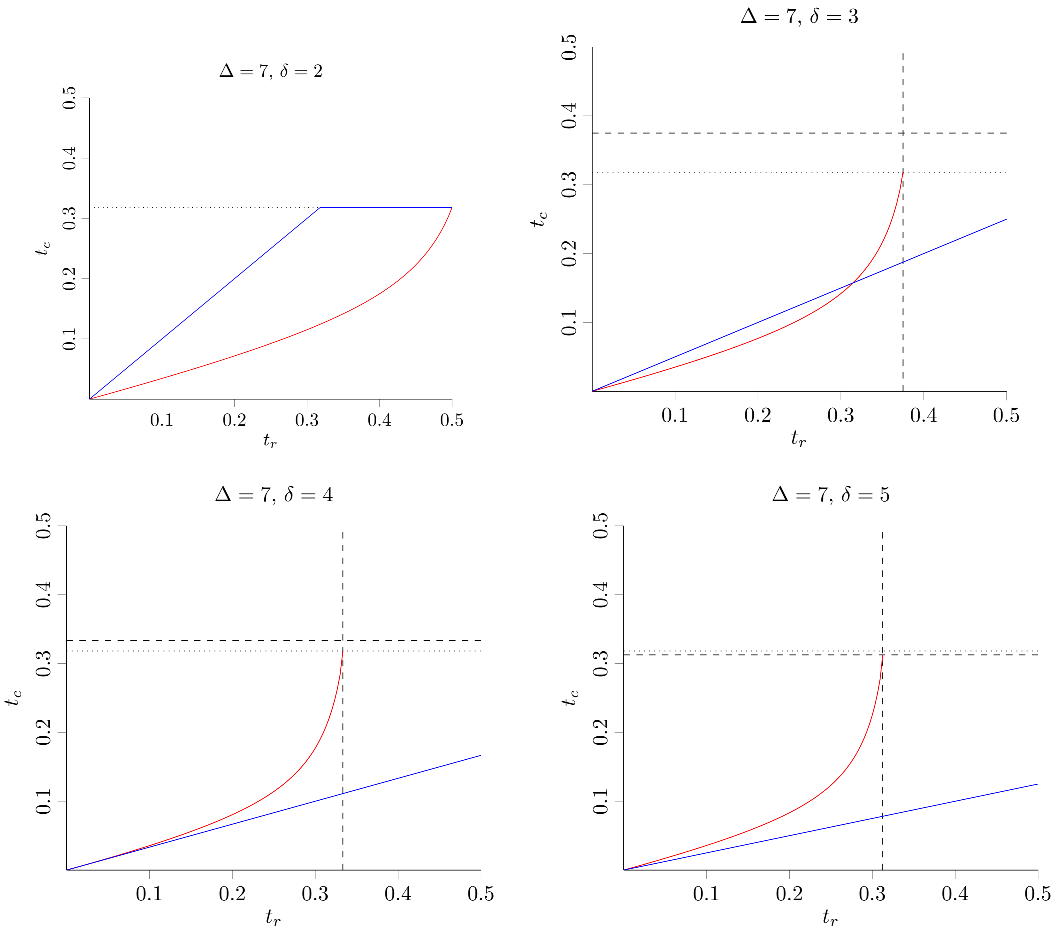

FIGURE 3. Comparing the efficacy of the robber's dualing strategies on $X(\Delta, \delta)$ for $\Delta = 7$ and $\delta \in \{2, 3, 4, 5\}$.

Recently, Songsuwan, Jiarasuksakun, Tanghanawatsakul, and Kaemawichanurat calculated the expected capture time of the game on the n-dimensional grid graph with the robber's tipsiness being 1 and the cop's tipsiness being 0 [19]. We think that it might be possible to refine our proof techniques to calculate the expected capture times for other tipsiness parameters as well.

Problem 5.4. Calculate the expected capture time on n-dimensional grid graphs for other tipsiness values.

In this paper, we have studied one of several interesting notions of tipsiness. A particularly interesting generalization is a model where the cop and the robber become 'more aware' of each others' presence when their distance is small, that is, they are more likely to play according to their strategies when the opposing player is nearby.

Problem 5.5. Study models where the tipsiness parameters increase with the distance between the cop and the robber. What can be said about optimal strategies and capture times? Do the results depend on how quickly the tipsiness parameters increase (compared to the dimension of the grid or to one another)?

In this manuscript we focus on infinite graphs, but there are also many interesting questions regarding the game on finite graphs. A very natural question to ask is the 'cost of drunkenness': how does the capture time depend on the tipsiness of either player?

Problem 5.6. On finite graphs, by how much can the expected capture time differ from the capture time in the deterministic game? What is the best strategy for a sober robber against a drunk cop on a finite tree, for instance a full binary tree of depth k? What about other tipsiness parameters for both players?

References

[1] V. Bardenova, V. Ciarcia, and E. Insko. Markov models for the tipsy cop and robber game on graphs. *Open J. Disc. Math.*, 11(3):61–93, 2021.

[2] M. T. Barlow, Y. Peres, and P. Sousi. Collisions of random walks. *Ann. Inst. Henri Poincaré Probab. Stat.*, 48(4):922–946, 2012.

[3] A. Bonato and R. J. Nowakowski. *The game of cops and robbers on graphs*, volume 61 of *Student Mathematical Library*. American Mathematical Society, Providence, RI, 2011.

[4] D. Chen, B. Wei, and F. Zhang. A note on the finite collision property of random walks. *Stat. Probab. Lett.*, 78(13):1742–1747, 2008.

[5] X. Chen. Gaussian bounds and collisions of variable speed random walks on lattices with power law conductances. *Stochastic Process. Appl.*, 126(10):3041–3064, 2016.

[6] X. Chen and D. Chen. Two random walks on the open cluster of $\mathbb{Z}^2$ meet infinitely often. *Sci. China, Math.*, 53(8):1971–1978, 2010.

[7] X. Chen and D. Chen. Some sufficient conditions for infinite collisions of simple random walks on a wedge comb. *Electron. J. Probab.*, 16:no. 49, 1341–1355, 2011.

[8] A. Devulder, N. Gantert, and F. Pène. Collisions of several walkers in recurrent random environments. *Electron. J. Probab.*, 23:Paper No. 90, 34, 2018.

[9] A. Devulder, N. Gantert, and F. Pène. Arbitrary many walkers meet infinitely often in a subballistic random environment. *Electron. J. Probab.*, 24:Paper No. 100, 25, 2019.

[10] P. E. Harris, E. Insko, A. Prieto-Langarica, R. Stoisavljevic, and S. Sullivan. Tipsy cop and drunken robber: a variant of the cop and robber game on graphs. *Adv. Appl. Discrete Math.*, 26, 2021.

[11] T. Hutchcroft and Y. Peres. Collisions of random walks in reversible random graphs. *Electron. Commun. Probab.*, 20:no. 63, 6, 2015.

[12] A. Kehagias and P. Prałat. Some remarks on cops and drunk robbers. *Theoret. Comput. Sci.*, 463:133–147, 2012.

[13] M. Krishnapur and Y. Peres. Recurrent graphs where two independent random walks collide finitely often. *Electron. Comm. Probab.*, 9:72–81, 2004.

[14] L. Lovász. Random walks on graphs: a survey. In *Combinatorics, Paul Erdős is eighty, Vol. 2 (Keszthely, 1993)*, volume 2 of *Bolyai Soc. Math. Stud.*, pages 353–397. János Bolyai Math. Soc., Budapest, 1996.

[15] R. Lyons and Y. Peres. *Probability on trees and networks*, volume 42 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, New York, 2016.

[16] R. Nowakowski and P. Winkler. Vertex-to-vertex pursuit in a graph. *Discrete Mathematics*, 43(2):235 – 239, 1983.

[17] G. Pólya. *Collected papers. Vol. IV*, volume 22 of *Mathematicians of Our Time*. MIT Press, Cambridge, MA, 1984. Probability; combinatorics; teaching and learning in mathematics, Edited by Gian-Carlo Rota, M. C. Reynolds and R. M. Shortt.

[18] A. Quilliot. Some results about pursuit games on metric spaces obtained through graph theory techniques. *European Journal of Combinatorics*, 7(1):55 – 66, 1986.

[19] N. Songsuwan, T. Jiarasuksakun, A. Tangthanawatsakul, and P. Kaemawichanurat. A cop and drunken robber game on n-dimensional infinite-grid graphs. *Mathematics*, 9(17), 2021.

[20] W. Woess. Random walks on infinite graphs and groups—a survey on selected topics. *Bull. London Math. Soc.*, 26(1):1–60, 1994.

[21] Wolfgang Woess. *Random walks on infinite graphs and groups*, volume 138 of *Camb. Tracts Math.* Cambridge: Cambridge University Press, paperback edition edition, 2008.

[22] Wolfgang Woess. *Denumerable Markov chains*. EMS Textbooks in Mathematics. European Mathematical Society (EMS), Zürich, 2009. Generating functions, boundary theory, random walks on trees.

(P. E. Harris) Department of Mathematics, University of Wisconsin - Milwaukee, United States
Email address: peharris@uwm.edu

(E. Insko) Department of Mathematics, Florida Gulf Coast University, United States
Email address: einsko@fgcu.edu

(F. Lehner) Department of Mathematics, University of Auckland, New Zealand
Email address: mail@florian-lehner.net