On-line size Ramsey number for monotone k-uniform ordered paths with uniform looseness

Xavier Pérez-Giménez* Paweł Prałat† Douglas B. West‡

July 20, 2020

Abstract

An *ordered hypergraph* is a hypergraph H with a specified linear ordering of the vertices, and the appearance of an ordered hypergraph G in H must respect the specified order on $V(G)$. In on-line Ramsey theory, Builder iteratively presents edges that Painter must immediately color. The *t-color on-line size Ramsey number* $\tilde{R}_t(G)$ of an ordered hypergraph G is the minimum number of edges Builder needs to play (on a large ordered set of vertices) to force Painter using t colors to produce a monochromatic copy of G. The *monotone tight path* $P_r^{(k)}$ is the ordered hypergraph with r vertices whose edges are all sets of k consecutive vertices.

We obtain good bounds on $\tilde{R}_t(P_r^{(k)})$. Letting $m = r - k + 1$ (the number of edges in $P_r^{(k)}$), we prove $m^{t-1}/(3\sqrt{t}) \le \tilde{R}_t(P_r^{(2)}) \le tm^{t+1}$. For general k, a trivial upper bound is $\binom{R}{k}$, where R is the least number of vertices in a k-uniform (ordered) hypergraph whose t-colorings all contain $P_r^{(k)}$ (and is a tower of height $k-2$). We prove $R/(k \lg R) \le \tilde{R}_t(P_r^{(k)}) \le R(\lg R)^{2+\varepsilon}$, where ε is a positive constant and tm is sufficiently large in terms of ε^{-1}. Our upper bounds improve prior results when t grows faster than $m/\log m$. We also generalize our results to ℓ-loose monotone paths, where each successive edge begins ℓ vertices after the previous edge.

1 Introduction

Ramsey theory studies the occurrence of forced patterns in colorings. We say that H *forces* G and write $H \to_t G$ when every t-coloring of the elements of H contains a monochromatic copy of G. In this paper H and G are k-uniform hypergraphs, we color the edges of H, and $t \ge 2$. Ramsey's Theorem [37] implies $K_n^{(k)} \to_t G$ when n is sufficiently large, where $K_n^{(k)}$ denotes the complete k-uniform hypergraph with n vertices. Our problem involves several variations on this.

For any monotone parameter, we can study its least value on the (hyper)graphs that force G. Besides the number of vertices (the classical problem), the most-studied parameter for this is the number of edges, yielding the *size Ramsey number* (proposed in [19], with early work surveyed in [21]). For example, Beck [3] solved a problem of Erdős by showing that the 2-color size Ramsey number of the path P_n is linear in n; after improvements in [6, 16, 28], the current best upper bound is $74n$ by Dudek and Prałat [17].

*Department of Mathematics, University of Nebraska–Lincoln, Lincoln, NE, USA, e-mail: `xperez@unl.edu`

†Department of Mathematics, Ryerson University, Toronto, ON, Canada, e-mail: `pralat@ryerson.ca`; Supported in part by NSERC and Ryerson University.

‡Departments of Mathematics, Zhejiang Normal University, Jinhua, China, and University of Illinois, Urbana, IL, USA, e-mail: `dwest@math.uiuc.edu`; Research supported by National Natural Science Foundation of China grants NSFC 11871439 and 11971439.

Another direction considers an ordered version of hypergraphs. An *ordered hypergraph* is a hypergraph on a linearly ordered vertex set. In the ordered sense, H is a *subhypergraph* of H' if H' contains a copy of H with the vertices appearing in the specified order. The ordered version of Ramsey's Theorem states that for an ordered k-uniform hypergraph G, there exist an ordered k-uniform hypergraph H such that $H \to_t G$ (meaning that every t-coloring of $E(H)$ contains a monochromatic copy of G in the ordered sense). This follows from Ramsey's Theorem because a complete ordered hypergraph with n vertices contains all ordered hypergraphs with n vertices; that is, it is enough to force $K^{(k)}_{|V(G)|}$ under any vertex ordering. Thus Ramsey numbers and size Ramsey numbers for ordered hypergraphs, being the least number of vertices or edges, respectively, in such an ordered hypergraph H, are also well-defined. Such problems have been studied in [2, 10, 12, 22, 29, 30, 31, 32].

In particular, let $[r]$ denote $\{1, \ldots, r\}$. The *monotone tight path* $P^{(k)}_r$ is the k-uniform ordered hypergraph with vertex set $[r]$ whose edges are all sets of k consecutive vertices. When G is understood to be an ordered hypergraph, we use $R_t(G)$ to denote the Ramsey number of G in the ordered sense. Thus $R_t(P^{(k)}_r)$ is the least n such that every t-coloring of the edges of the k-uniform complete ordered hypergraph with vertex set $[n]$ contains a monochromatic copy of $P^{(k)}_r$. This value was studied and applied in [18, 22, 29, 30].

An "on-line" version of Ramsey theory is a game between *Builder* and *Painter*, introduced by Beck [4] and by Kurek and Ruciński [27]. In each round, Builder presents an edge that Painter must color. Builder can force a monochromatic copy of G by presenting all edges of some H such that $H \to_t G$. However, Builder may be able to use Painter's choices to force G to appear sooner. On-line Ramsey problems have been studied for the number of edges [9, 13, 14, 23, 24, 27, 35, 36], the genus [23, 25, 34], and the maximum degree [8, 26, 39, 40]. For a monotone parameter φ, the *on-line φ Ramsey number* is the least t such that Builder can force Painter to produce a monochromatic G while keeping the value of φ to at most t on the presented object. The number of edges is the number of rounds (the length of the game) and hence is the natural parameter. It is so natural that the on-line size Ramsey number has confusingly also been called just the on-line Ramsey number, which term more properly is the minimum number of vertices needed. Easy arguments imply that the 2-color on-line size Ramsey number of the path P_n is at least $2n-3$ and at most $4n-7$ ([24]).

We combine these three variations on the Ramsey problem, studying the number of edges used in the on-line model to force an ordered hypergraph. In particular, we study the on-line size Ramsey number of monotone tight paths. Appropriate notation is needed for the resulting parameter. One common practice in Ramsey theory is to add a circumflex accent ($\hat{R}$) to indicate a size Ramsey number. Some recent papers use a tilde accent ($\tilde{R}$) to indicate the on-line version of the size Ramsey number (a circular accent $\mathring{R}$ has been used with on-line versions of other parameter Ramsey numbers). These choices free the subscript for the number of colors. For ordered Ramsey numbers, OR was used in [29], but now we follow [9, 31, 32] in using the same notation as in the classical problem when it is understood that the target and host are ordered hypergraphs. Thus we use $\tilde{R}_t(P^{(k)}_r)$ for the t-color on-line size Ramsey number of the monotone tight path $P^{(k)}_r$.

Our results and proofs are motivated by the characterization of the t-color off-line vertex Ramsey number of $P_r^{(k)}$ in terms of the size of an associated poset (partially ordered set), obtained by Moshkovitz and Shapira [30] (see [29] for an exposition and alternative presentation of the proof). Henceforth let m be the number of edges in $P_r^{(k)}$; note that $m = r - k + 1$. The arguments and bounds are stated more cleanly in terms of m. Let Q_1 be the poset consisting of t disjoint chains of size $m-1$. For $j > 1$, let Q_j be the poset consisting of all the down-sets in Q_{j-1}, ordered by inclusion. The value of $\tilde{R}_t(P_r^{(k)})$ is given in terms of $|Q_k|$, and the bounds on $|Q_k|$ follow inductively.

Theorem 1 (Moshkovitz and Shapira [30]). *$R_t(P_r^{(k)}) = |Q_k| + 1$. Furthermore,*

$$\mathrm{tow}_{k-2}(m^{t-1}/2\sqrt{t}) \le |Q_k| \le \mathrm{tow}_{k-2}(2m^{t-1}),$$

where $m = r - k + 1$ and $\mathrm{tow}_h(x)$ equals x when $h = 0$ and $2^{\mathrm{tow}_{h-1}(x)}$ when $h \ge 1$.

This result immediately implies $\tilde{R}_t(P_r^{(k)}) \le \binom{|Q_k|+1}{k}$, since $\binom{|Q_k|+1}{k}$ is the number of edges in $K^{(k)}_{|Q_k|+1}$. Building on ideas used in the exposition of this proof in [29], for sufficiently small positive ε we present a strategy for Builder that proves an upper bound of $|Q_k|(\lg |Q_k|)^{2+\varepsilon}$ when tm is sufficiently large in terms of ε^{-1}, where lg is the base-2 logarithm. Our Painter strategy for the lower bound yields roughly the same lower bound as in Theorem 1. Hence our upper and lower bounds on $\tilde{R}_t(P_r^{(k)})$ are towers of the same height.

The arguments for the upper and lower bound generalize trivially to the non-diagonal case $\tilde{R}_t(P_{r_1}^{(k)}, \ldots, P_{r_t}^{(k)})$, where Builder seeks to force a copy of $P_{r_i}^{(k)}$ in color i for some i. Simply let Q_1 be the disjoint union of t chains such that the ith chain has $r_i - k + 1$ elements.

Fox, Pach, Sudakov, and Suk [22] considered a game with a more restricted Builder, which was introduced by Conlon, Fox, and Sudakov [11]. Builder can only introduce a new vertex at the end of the ordering and can only present an edge joining the newest vertex to earlier vertices. Painter colors each presented edge immediately. Our Builder can simulate this game, so the optimal value $f_t(m)$ in their game is at least $\tilde{R}_t(P_{m+1}^{(2)})$. For constant t (and here $k = 2$), Fox et al. [22] proved

$$\frac{t-1-o(t)}{\log t} m^t \log m \le f_t(m) \le \left(1 + \frac{t-1}{\log(1+1/(t-1))}\log(m+1)\right)(m^t+1).$$

Since their Builder is weaker, their lower bound is naturally larger than ours; neither result implies the other. For large t (growing faster than $m/\log m$), our upper bound is smaller than theirs, but for constant t their upper bound is better.

They also studied the k-uniform version of their game, where their objective was to obtain an upper bound on the vertex Ramsey number of the monotone tight path in terms of the length of their game. Since $R_t(P_r^{(k)}) = |Q_k| + 1$ by [30], in their game also Builder must use more than $|Q_k|$ vertices to end the game.

Indeed, if Painter knows that the game is being played by their Builder, meaning that vertices will only be introduced from left to right, then Painter can use our strategy (in the general k-uniform case) with a supply of Q_k vertices (treating them as described in Section 3), achieving $|Q_k|/k$ as a lower bound against their Builder. Similarly, when the vertices are known initially (that is, in the off-line setting), our Painter strategy also implies that any hypergraph forcing $P_r^{(k)}$ has more than $|Q_k|$ vertices, thus yielding the lower bound $R_t(P_r^{(k)}) > |Q_k|$ in Theorem 1. A closer look at the upper bound strategy for Builder also yields the upper bound $R_t(P_r^{(k)}) \le |Q_k|+1$. The ideas in our proof are similar to the ideas in the proofs in [30] and [29].

Our proofs also generalize easily to describe the Ramsey number of the monotone *ℓ-loose k-uniform path* $P_r^{k,\ell}$ for $1 \le \ell \le k$. Here each edge consists of k consecutive vertices, and two consecutive edges have $k-\ell$ common vertices. (In particular, $r = k + \ell(m-1)$ when there are m edges.) Note that $P_r^{(k)} = P_r^{k,1}$, while $P_r^{k,k}$ is a k-uniform matching in which each edge ends before the next edge begins in the vertex ordering. Let $h = \lceil k/\ell \rceil$. Our arguments for the on-line version of the problem yield $R_t(P_r^{k,\ell}) = \ell|Q_h| + s$, where $s = k-(h-1)\ell$. This formula was obtained earlier by Cox and Stolee [12], expressed in different notation. They gave a separate argument for the case $\ell = k$ (matchings), though this formula applies to both.

In the last section we discuss an off-line version of this problem for directed graphs and hypergraphs, related to results of Ben-Eliezer, Krivelevich, and Sudakov [5].

2 On-line scenario: The graph case ($k = 2$)

The game ends when Builder forces Painter to produce a monochromatic monotone tight path with m edges. For clarity and because the numerical bounds are somewhat tighter in this case, we first consider the case $k = 2$. For the monotone path, $R_t(P_r^{(2)}) = m^t + 1$. The trivial upper bound is $\binom{R_t(P_r^{(2)})}{2}$, but our upper bound is not much larger than $R_t(P_r^{(2)})$. Like the result of [30], it is motivated by the short proof due to Seidenberg [41] of the Erdős–Szekeres Theorem [20] on monotone subsequences.

Theorem 2. *For $m = r-1$ with $r \ge 3$, always $m^{t-1}/(3\sqrt{t}) \le \tilde{R}_t(P_r^{(2)}) \le tm^{t+1}$.*

Proof. Let $M = \{0, 1, \dots, m-1\}$. Given $a = (a_1, \dots, a_t) \in M^t$, let $|a| = \sum a_i$.

Upper bound *(Builder strategy):* Builder uses $m^t + 1$ vertices, viewed as ordered from left to right. At any time, all vertices are labeled with vectors in M^t, where the ith coordinate of the label for v is the number of edges in the longest monotone path in color i that ends at v. All labels are initially the all-0 vector. Let $\mathbf{\Lambda}$ denote the "top" vector in M^t; its components all equal $m-1$.

Builder seeks to produce label $\mathbf{\Lambda}$ at one of the first m^t vertices, after which playing the edge from this vertex to the last (rightmost) vertex wins the game no matter what color Painter gives it. If no two vertices among the first m^t have the same label, then all labels occur, including $\mathbf{\Lambda}$.

Otherwise, some vertices u and v have the same label, say with u before v. These vertices cannot yet be adjacent, since their labels would then differ in the coordinate for the color of uv. Builder plays uv. The label for v increases in the coordinate for the color Painter uses on uv.

On each round, the label for the second vertex of the edge played increases by 1 in some coordinate. To avoid reaching $\mathbf{\Lambda}$ or reaching m in any coordinate, each label must increase fewer than $(m-1)t$ times. By the pigeonhole principle, within $m^t[(m-1)t-1]+1$ rounds some label reaches $\mathbf{\Lambda}$, and the next play wins. Note that $m^t[(m-1)t-1]+1 < tm^{t+1}$.

Lower bound *(Painter strategy):* Let $B = \{a \in M^t \colon\ |a| = \lfloor (m-1)t/2 \rfloor\}$. Until Builder uses more than $|B|$ vertices, Painter can assign different labels from B to all vertices used. These labels remain unchanged throughout the game. Let $a(v)$ denote the label assigned by Painter to v, with $a(v) = (a_1(v), \dots, a_t(v))$. When Builder plays an edge uv with u before v, Painter gives it a color i such that $a_i(v) > a_i(u)$. Such a coordinate exists, since $a(u) \neq a(v)$ and $|a(u)| = |a(v)|$.

Choosing colors in this way maintains for each vertex w the property that every monotone path in color i arriving at vertex w has at most $a_i(w)$ edges. This holds since along a monotone path in color i the ith coordinate of the label strictly increases with each step. Since $a(w) \in M^t$, no monochromatic monotone path has m edges. Since using more than $|B|$ vertices requires playing more than $|B|/2$ edges, Painter can survive at least $|B|/2$ rounds without creating a monochromatic monotone path with m edges.

The elements of M are the elements of Q_2, and B is a middle level. Using Chebyshev's Inequality and the pigeonhole principle, Moshkovitz and Shapira [30] showed $|B| \geq \frac{2}{3}m^{t-1}/\sqrt{t}$. □

Remark 3. It is well known by many arguments that B is a largest level in Q_2. (For example, the product of chains is a symmetric chain order, the convolution of symmetric log-concave sequences is symmetric and log-concave, explicit injections map one level to the next toward the middle, etc.) Since $|M^t| = m^t$ and there are $(m-1)t+1$ levels, we thus have $|B| > m^{t-1}/t$ by the pigeonhole principle alone.

Using the Chernoff bound instead of Chebyshev's Inequality in the argument in [30], we can improve the lower bound on $|B|$ to $0.7815987m^{t-1}/\sqrt{t}$. The value of $|B|$ was also studied by Alekseev [1]. A special case is that when $m \in o(e^t/\sqrt{t})$, the value of $|B|$ is asymptotic to $m^{t-1}/\sqrt{\pi t/6}$.

For the non-diagonal case, with m_i being the forbidden length in color i, the argument yields

$$\frac{\prod m_i}{2\sum m_i} \leq \tilde{R}_t(P^{(2)}_{r_1}, \dots, P^{(2)}_{r_t}) \leq \sum m_i \prod m_i.$$

Here the pigeonhole argument for the size of the largest antichain in Q_2 gives the lower bound on $|B|$. Again Chebyshev's Inequality can be used to improve it somewhat, but the resulting formula is more complicated.

Our lower bound remains valid against a stronger Builder. Suppose Builder can present any directed graph in seeking a monochromatic directed path, instead of only presenting edges directed from lower to higher vertices. The strategy for Painter establishes the same lower bound, where "an edge uv with u before v" becomes "an edge directed from u to v". This works because the labels for vertices are incomparable. We will return to the digraph problem in the last section.

3 On-line scenario: The hypergraph case

For the k-uniform monotone tight path, the flavor of the arguments extends that of the graph case, but the details are more delicate. As described in the introduction, let Q_1 be the poset consisting of t disjoint chains of $m-1$ elements each. The ith chain is associated with color i. For $j > 1$, the poset Q_j consists of the down-sets in Q_{j-1}, ordered by inclusion. The arguments are the same for the non-diagonal case, with the ith chain in Q_1 consisting of $m_i - 1$ elements, where $m_i = r_i - k + 1$.

We will first study the upper bound. Let G denote the current hypergraph of edges played by Builder and colored by Painter. In the strategy for Builder used to prove the upper bound, Builder will confine play to a fixed vertex set $[n]$, where $[n] = \{1, \ldots, n\}$, under the usual order on $\mathbb{N}$. Given a set $Y \subseteq [n]$, let Y^+ be the set obtained from Y by deleting the first vertex, and let Y^- be the set obtained from Y by deleting the last vertex. Let $\binom{[n]}{j}$ denote the family of j-element subsets of $[n]$. We recursively define functions $g_k, \ldots, g_1$ such that $g_j \colon \binom{[n]}{j} \to Q_{k-j+1}$, except that g_k is defined only on the k-sets that are actual edges of G. We also recursively define a notion of one j-set "following" another. Both definitions depend on G and pertain to the strategy that Builder will use to pick the next edge as long as G has no monochromatic $P_r^{(k)}$.

Definition 4. For $Y \in E(G)$, if Y has color i and the longest monochromatic tight path with last edge Y has p edges, then let $g_k(Y)$ be element p on the ith chain in Q_1. For $Y \in \binom{[n]}{j}$ with $j < k$, let $\overleftarrow{Y} = \{Z \in \binom{[n]}{j+1} \colon Z^+ = Y\}$; call the elements of $\overleftarrow{Y}$ the *precursors* of Y. Given that g_{j+1} has been defined, for $Y \in \binom{[n]}{j}$ define $g_j(Y)$ as follows:

$$g_j(Y) \text{ is the downset in } Q_{k-j} \text{ generated by } \{g_{j+1}(Z) \colon Z \in \overleftarrow{Y}\}.$$

Being a downset in Q_{k-j}, by definition $g_j(Y) \in Q_{k-j+1}$.

Definition 5. Given $Y_1, Y_2 \in \binom{[n]}{k}$, say that Y_2 *follows* Y_1 if $Y_1, Y_2 \in E(G)$ and $Y_1^+ = Y_2^-$. For $Y_1, Y_2 \in \binom{[n]}{j}$ with $j < k$, say that Y_2 *follows* Y_1 if

(A) $Y_1^+ = Y_2^-$ and

(B) for each maximal element w of $g_j(Y_1)$, the $(j+1)$-set $Y_1 \cup Y_2$ follows some precursor Z_1 of Y_1 such that $g_{j+1}(Z_1) = w$.

Note that (B) in Definition 5 holds trivially when $g_j(Y_1)$ is empty. Since a precursor Z_2 of Y_2 following a precursor Z_1 of Y_1 requires $Z_2^- = Z_1^+ = Y_1$, the set $Y_1 \cup Y_2$ is the only precursor of Y_2 that can follow a precursor of Y_1. When Y_2 follows Y_1, the set $Y_1 \cup Y_2$ is a set Z such that $Z^- = Y_2$ and $Z^+ = Y_2$.

Our strategy for Builder is based on a crucial property of g_j we prove next. We think of g_j as assigning a label in Q_{k-j+1} to a j-set Y.

Lemma 6. *If Y_2 follows Y_1 in $\binom{[n]}{j}$, then $g_j(Y_1) \not\geq g_j(Y_2)$ in Q_{k-j+1}.*

Proof. The proof is by induction on $k - j$. For $j = k$, if Y_2 follows Y_1 in $E(G)$, then either Y_1 and Y_2 have the same color, in which case $g_k(Y_2) > g_k(Y_1)$ in Q_1, or they have different colors, in which case $g_k(Y_1)$ and $g_k(Y_2)$ are incomparable in Q_1. In either case, $g_k(Y_1) \not\geq g_k(Y_2)$.

For $j < k$, suppose that the claim holds for $j + 1$. Given that Y_2 follows Y_1 in $\binom{[n]}{j}$, let $Z = Y_1 \cup Y_2 \in \binom{[n]}{j+1}$. If Y_1 has no precursors, then $g_j(Y_1)$ is empty; the claim is then trivially true, since Z is a precursor of Y_2 and thus $g_j(Y_2)$ is nonempty. Otherwise, let w be a maximal element of $g_j(Y_1)$. Since Y_2 follows Y_1, by definition Z follows some $Z_1 \in \overleftarrow{Y}_1$ with $g_{j+1}(Z_1) = w$. By the hypothesis for $j+1$, we have $w = g_{j+1}(Z_1) \not\geq g_{j+1}(Z)$ for all such Z_1. Since this holds for all w that are maximal in $g_j(Y_1)$, the label $g_{j+1}(Z)$ does not lie in the downset generated by the precursors of Y_1 (which by definition is $g_j(Y_1)$). However, since $Z \in \overleftarrow{Y}_2$, the label $g_{j+1}(Z)$ does lie in $g_j(Y_2)$. Hence as downsets in Q_{k-j}, the family $g_j(Y_2)$ is not contained in the family $g_j(Y_1)$, which means $g_j(Y_1) \not\geq g_j(Y_2)$ as elements of Q_{k-j+1}. □

The inductive definition of "follows" facilitates Lemma 6. To simplify the presentation of Builder's strategy, we provide a more explicit description of what "Y_2 follows Y_1" guarantees.

Definition 7. For a j-set Y with $j < k$ or an edge $Y \in E(G)$, we form a tree $U(Y)$. The nodes of the tree are elements of the posets $Q_{k-j+1}, \ldots, Q_1$. The root of $U(Y)$ is the label $g_j(Y) \in Q_{k-j+1}$. Any node z in $U(Y)$ that is an element of Q_i for $i > 1$ is also a downset in Q_{i-1}; the children of z in $U(Y)$ are the maximal elements of this downset in Q_{i-1}. The process iterates until we reach elements of Q_1 as the leaves of $U(Y)$. To avoid confusion, we will refer to the vertices of the tree $U(Y)$ as nodes and reserve the term "vertex" for the elements of $V(G)$.

An *instance* of $U(Y)$ for a j-set Y associates vertex sets to the nodes. Associated to the root of $U(Y)$, which is $g_j(Y) \in Q_{k-j+1}$, is the set Y. To a non-root node $w \in Q_i$ whose parent in $U(Y)$ is $z \in Q_{i+1}$ and has associated $(k-i)$-set Z, we associate a precursor Z' of Z such that $g_{k-i+1}(Z') = w$; note that Z' is a $(k-i+1)$-set. Iteratively, we choose associated sets moving away from the root. Since the leaves are in Q_1, their associated sets are k-sets: that is, edges.

We must confirm that the selection of labels is well-defined, so that every such tree $U(Y)$ has at least one instance of the form described. When w is a child in $U(Y)$ of the node $z \in Q_{i+1}$ with associated $(k-i)$-set Z, we have already chosen Z so that $g_{k-i}(Z) = z$. By definition, z is the downset in Q_i generated by $\{g_{k-i+1}(Z') \colon Z' \in \overleftarrow{Z}\}$, and w is a maximal element in that downset. Thus w must be the image under g_{k-i+1} of some precursor of Z.

Lemma 8. *A j-set Y_2 follows a j-set Y_1 if and only if $Y_1^+ = Y_2^-$ and there is an instance of $U(Y_1)$ such that for every node, deleting the first element of the associated set W and replacing it with the last element y of Y_2 yields a set Z following W. In particular, for every edge W associated with a leaf, replacing the first vertex of W with y yields an edge in G.*

Proof. First suppose that Y_2 follows Y_1. The tree $U(Y_1)$ is fixed, but the instance we construct depends on Y_2. We construct the needed instance of $U(Y_1)$ by finding sets to associate with nodes along each path from the root. The set Y_1 is associated with the root node.

Given that Y_2 follows Y_1, let $Z_1 = Y_1 \cup Y_2$. Note that Z_1 arises from Y_1 by adding the last vertex y from Y_2. By the definition of Y_2 following Y_1, the set Z_1 is required to be a $(j+1)$-set that, for each child w_1 of the root of $U(Y_1)$, follows some precursor W_1 of Y_1 that has label w_1. This selects W_1 as a $(j+1)$-set to associate with w_1 in the instance of $U(Y_1)$ we are building.

At the next level, descending from a child w_1, we have Z_1 following W_1. We let $Z_2 = W_1 \cup Z_1$. For each child w_2 of w_1, we apply the same argument to obtain the $(j+2)$-set W_2 to associate with w_2. Repeating this argument along any path from the root to a leaf of $U(Y_1)$, we obtain successively larger sets $Z_1, \ldots, Z_{k-j}$ that respectively follow sets $W_1, \ldots, W_{k-j}$ associated with the nodes along the path. Each Z_i is obtained by deleting the smallest element of W_i and adding y. Finally, Z_{k-j} is an edge following an edge W_{k-j} associated with the leaf at the end of the path. We obtain such an edge Z_{k-j} for each leaf. (It is just the edge Y_2 following Y_1 when $j = k$.)

The converse is the specialization of the assumed condition on $U(Y_1)$ to the root node. $\square$

Remark 9. For $Y \in \binom{[n]}{j}$, if no precursor of Y has a defined label, then the downset generated by $\overleftarrow{Y}$ is empty, and $g_j(Y)$ is the bottom element of Q_{k-j+1}. This occurs for any $(k-1)$-set whose precursors all are not edges of G and for any j-set with first vertex 1 (it has no precursors).

Each of $Q_2, \ldots, Q_k$ has one element of rank 0, which is the empty downset in the previous poset. Also each of $Q_3, \ldots, Q_k$ has one element of rank 1, which is the downset of size 1 consisting of the bottom element of the previous poset. Inductively, ranks 0 through $j-2$ of Q_j form a single chain with one element of each rank. For $0 \le i \le j-2$, let $\mathbf{V}_j^i$ be the element of rank i in Q_j.

With vertex set $[n]$ before any edges have been played, all k-sets have undefined labels. Hence the label of each $(k-1)$-set at the start is $\mathbf{V}_2^0$. A j-set with least element 1 has no precursor, so its label is $\mathbf{V}_{k-j+1}^0$. For $j < k-1$, a j-set Y with least element 2 has one precursor, which has label $\mathbf{V}_{k-j}^0$, so initially $g_j(Y) = \mathbf{V}_{k-j+1}^1$. Inductively, for $j < k$, a j-set Y with least element i has initial label $\mathbf{V}_{k-j+1}^{i-1}$ if $i \le k-j$ and label $\mathbf{V}_{k-j+1}^{k-j-1}$ if $i > k-j$. In particular, for the crucial case $j=1$, the initial label of the vertex i is $\mathbf{V}_k^{i-1}$ for $i \le k-1$ and $\mathbf{V}_k^{k-2}$ for $i > k-1$.

Our upper bound for general k is also valid for $k=2$, but in that case Theorem 2 provides a stronger bound. For $k=3$ our bound is a bit weaker than for larger k, which introduces some complication in the inductive proof. The combinatorial bound obtained first is valid for all k, m, t, but the bound in terms of $|Q_k|$ alone requires tm (or equivalently $|Q_1|$) to be sufficiently large.

Theorem 10. *For $k, m, t \in \mathbb{N}$ with $t, m \ge 2$ and $r = k+m-1$.*

$$\tilde{R}_t(P_r^{(k)}) \le |Q_k| \cdot |Q_{k-1}| \prod_{i=1}^{k-1} a_i,$$

where a_i is the size of the largest antichain in Q_i. Moreover, for any sufficiently small positive constant ε,

$$|Q_3| \cdot |Q_2| a_2 a_1 \le |Q_3| (\lg |Q_3|)^{2+\frac{1}{t-1}+\varepsilon} \qquad \textit{and} \qquad |Q_k| \cdot |Q_{k-1}| \prod_{i=1}^{k-1} a_i \le |Q_k| (\lg |Q_k|)^{2+\varepsilon} \quad (\textit{for } k \ge 4)$$

when tm is sufficiently large compared to ε^{-1}.

Proof. We give a strategy for Builder. Let $n = |Q_k| + 1$. Builder plays k-sets from the fixed ordered vertex set $[n]$, numbered from left to right. After each round the label functions $g_k, \ldots, g_1$ are defined as in Definition 4 for the hypergraph played so far. Let $\mathbf{\Lambda}_j$ be the unique top element in Q_j, for $2 \le j \le k$. Builder seeks a vertex z in $[n] - \{n\}$ with $g_1(z) = \mathbf{\Lambda}_k$. Since $\mathbf{\Lambda}_k$ is the downset in Q_{k-1} that is all of Q_{k-1}, this vertex z has a precursor $\{y, z\}$ with label $\mathbf{\Lambda}_{k-1}$. Iterating, some $(k-1)$-set Y ending at z has label $\mathbf{\Lambda}_2$. Since $\mathbf{\Lambda}_2 = (m-1, \ldots, m-1)$, in each color some precursor of Y is the last edge in a path of $m-1$ edges. Builder then plays the edge $Y \cup \{n\}$ to win.

Builder plays to force Painter to produce such a vertex z. Before any edges are played, the labels are as described in Remark 9. The labels of the first $k-1$ vertices never change (always $g_1(\{i\}) = \mathbf{V}_k^{i-1}$ for $i \le k-1$), since no edge can be played ending at one of those vertices. All vertices from $k-1$ to n initially have the same label $\mathbf{V}_k^{k-2}$ with rank $k-2$ in Q_k.

Playing an edge in the game creates a label for that edge. The label of an existing edge stays the same or moves upward on its chain, by the definition of g_k. For a j-set Y with $j < k$, by induction on $k - j$, the label $g_j(Y)$ stays the same or moves upward in Q_{k-j+1}, because the label is defined to be the downset generated by the labels of the precursors. The precursors remain the same (except that precursors can be added when $j = k - 1$). By the induction hypothesis, the labels of the precursors stay the same or move up. Hence the downset they generate stays the same or becomes larger, which means that $g_j(Y)$ stays the same or moves up.

After the first $k-2$ vertices and before the last, there are $|Q_k| - k + 2$ vertices, and their labels are initially (and hence always) above the bottom $k-2$ elements of Q_k. If $\mathbf{\Lambda}_k$ is not the label of any of them, then their labels are confined to a set of $|Q_k| - k + 1$ elements in $|Q_k|$. By the pigeonhole principle, two of these vertices have the same label. We claim that in this situation Builder can make a vertex label go up in Q_k.

Builder picks two vertices x and y having the same label, with x before y. Since x and y have the same label, Lemma 6 guarantees that y does not follow x. Builder plays edges to make y follow x. Since labels that change can only move up, Lemma 6 implies that playing edges to make y follow x causes the label of y to increase in Q_k.

In order to make y follow x, Builder uses an instance of $U(\{x\})$. For each leaf in $U(\{x\})$, the associated edge Z ends with x. By Lemma 8, y follows x if $Z^+ \cup \{y\}$ is an edge for each such Z. *Builder plays all such k-sets that are not already edges.*

The number of edges played by Builder to make y follow x is at most the number of leaves in $U(\{x\})$. Since the children in $U(\{x\})$ of each label in Q_j form an antichain in Q_{j-1} (because they are the maximal elements of a downset in Q_{j-1}), the number of leaves is bounded by $\prod_{i=1}^{k-1} a_i$, where a_i is the maximum size of an antichain in Q_i.

As long as no monotone tight path with m edges is created, the labels of the $|Q_k| - k + 1$ vertices we are considering can rise at most $|Q_{k-1}| - k$ times without reaching $\mathbf{\Lambda}_k$, since $\mathbf{\Lambda}_k$ is the full downset of size $|Q_{k-1}|$ in Q_{k-1}, and each of these labels initially is the unique downset of size $k - 1$. Hence

$$1 + [(|Q_k| - k + 1)(|Q_{k-1}| - k) + 1] \prod_{i=1}^{k-1} a_i$$

moves suffice for Builder to finish the game. Thus $\tilde{R}_t(P_r^{(k)}) \le |Q_k| \cdot |Q_{k-1}| \prod_{i=1}^{k-1} a_i$.

The remainder of the proof, obtaining an upper bound on $\tilde{R}_t(P_r^{(k)})$ in terms of $|Q_k|$ alone, is purely numerical. Let ε be a sufficiently small positive constant. We seek

$$|Q_2| a_2 a_1 \le (\lg |Q_3|)^{2+\frac{1}{t-1}+\varepsilon} \qquad \text{and} \qquad |Q_{k-1}| \prod_{i=1}^{k-1} a_i \le (\lg |Q_k|)^{2+\varepsilon} \quad (\text{for } k \ge 4). \tag{1}$$

We will find positive constants t_0 and m_0 in terms of ε such that (1) holds when $tm \ge t_0 m_0$. Note that we are not trying to optimize these bounds.

Let $q_i = |Q_i|$. The rank of an element of Q_i is its size as a downset in Q_{i-1}; hence Q_i has $|Q_{i-1}| + 1$ ranks. Since the minimal and maximal elements are unique, Q_i has a decomposition into the fewest chains such that no chain meets all ranks. Dilworth's Theorem [15] and the pigeonhole principle then yield $a_i \geq q_i/q_{i-1}$, and hence $a_i \leq q_i \leq a_i q_{i-1}$. Since the subsets of a largest antichain in Q_i generate distinct downsets, $q_{i+1} \geq 2^{a_i}$, so $a_i \leq \lg q_{i+1}$. To bound $q_{k-1} \prod_{i=1}^{k-1} a_i$ in terms of q_k, we need q_i to grow rapidly with i. Already we have $\lg q_{i+1} \geq q_i/q_{i-1}$ for $i \geq 2$, but we need better.

Consider first $k = 3$. The computation we use to prove the first part of (1) is

$$q_2 a_2 a_1 = tm^t a_2 \leq a_2 \left(\frac{m^{t-1}}{2\sqrt{t}} \right)^{t/(t-1)+\varepsilon} \leq (\lg q_3)^{2+\frac{1}{t-1}+\varepsilon}.$$

The first step follows from $a_1 = t$ and $q_2 = m^t$. For the rightmost inequality, we noted $a_2 \leq \lg q_3$ above, and Theorem 1 gives $m^{t-1}/2\sqrt{t} \leq \lg q_3$.

The middle inequality reduces to $t(2\sqrt{t})^{t/(t-1)+\varepsilon} \leq m^{\varepsilon(t-1)}$. When $m \geq 4^{1+2/\varepsilon}$, this holds for $t \geq 2$. When $(t-1)/\lg 4t \geq .5 + 2/\varepsilon$, from $t, m \geq 2$ we obtain

$$t(2\sqrt{t})^{t/(t-1)+\varepsilon} \leq t(2\sqrt{t})^{2+\varepsilon} \leq (4t)^{2+\varepsilon/2} = 2^{\varepsilon(2/\varepsilon+1/2)\lg 4t} \leq 2^{\varepsilon(t-1)} \leq m^{\varepsilon(t-1)}$$

Hence if we let $m_0 = 4^{1+2/\varepsilon}$ and let t_0 be the solution to $(t-1)/\lg 4t = .5+2/\varepsilon$, then the inequality will hold whenever $tm \geq t_0 m_0$, since that yields $t \geq t_0$ or $m \geq m_0$ when $t, m \geq 2$.

In order to prove the inequality of (1) for $k \geq 4$, it suffices to prove

$$\prod_{i=1}^{k-1} q_i \leq (\lg q_k)^{1+\varepsilon/2}, \tag{2}$$

because $a_i \leq q_i$ implies $q_{k-1} \prod_{i=1}^{k-1} a_i < (\prod_{i=1}^{k-1} q_i)^2$. In the induction step, we use $1 + \varepsilon/2 < 4$ to weaken the induction hypothesis, proving that $\prod_{i=1}^{k-2} q_i \leq (\lg q_{k-1})^4$ implies (2). As a base step to start the induction, we prove the weaker statement for $k = 3$. The computation for this is

$$q_2 q_1 = tm^{t+1} \leq (m^{t-1}/t)^4 = (q_2/q_1)^4 \leq a_2^4 \leq (\lg q_3)^4,$$

in which the only step needing further explanation is $tm^{t+1} \leq (m^{t-1}/t)^4$, which simplifies to $(mt)^5 \leq m^{3t}$. This holds when $t = 2$ and $m \geq 32$, or when $t \geq 3$ and $m \geq 4$. It does not hold when $t = m = 3$, but the desired inequality $tm^{t+1} \leq (q_2/q_1)^4$ does hold then. In any case, we obtain the desired inequality when $tm \geq 64$.

For the induction step, we first use $q_i \leq a_i q_{i-1}$, the induction hypothesis, and the fact that q_{k-1} (which exceeds $t(m-1)$) is sufficiently large to compute

$$\prod_{i=1}^{k-1} q_i \leq a_{k-1} q_{k-2} \prod_{i=1}^{k-2} q_i < a_{k-1} \left(\prod_{i=1}^{k-2} q_i \right)^2 \leq \lg q_k (\lg q_{k-1})^8 \leq q_{k-1}^{\varepsilon/3} \lg q_k.$$

Now let $\beta = \prod_{i=1}^{k-1} q_i$. We weaken $\beta \leq q_{k-1}^{\varepsilon/3} \lg q_k$ to $\beta \leq \beta^{\varepsilon/3} \lg q_k$. Rearranging to a bound on β now yields $\beta \leq (\lg q_k)^{1/(1-\varepsilon/3)} \leq (\lg q_k)^{1+\varepsilon/2}$, which completes the proof of (2) and the theorem. □

The argument for the lower bound, presented next, is easier.

Theorem 11. *With $r > k$ and $m = r - k + 1$, we have $\tilde{R}_t(P_r^{(k)}) \geq |Q_k|/(k \lg |Q_k|)$.*

Proof. With $Q_1, \ldots, Q_k$ defined as before, we give a strategy for Painter. Let A be a maximum-sized antichain in Q_k. We have noted that $|A| \geq |Q_k|/\lg|Q_k|$. When Builder uses new vertices, Painter gives them distinct unused elements of A as labels. Since each edge played has k vertices, Painter can make such assignments for at least $|A|/k$ edges.

Using these labels, Painter assigns labels to all j-sets in the set of vertices that have been used by Builder, for $1 \leq j \leq k$; these labels remain unchanged thereafter. The label $f_j(Y)$ assigned to a j-set Y is in Q_{k-j+1}. Thus the label of a k-set is in Q_1, and which chain it is on specifies the color to be used on the set if Builder plays it as an edge. Through the first $|A|/k$ rounds, this will enable Painter to avoid making a monochromatic monotone copy of $P_r^{(k)}$.

The property needed for the labels is that if Y_1 and Y_2 are j-sets such that $Y_1^+ = Y_2^-$ (or equivalently that $Y_1 = Y^-$ and $Y_2 = Y^+$ for some $(j+1)$-set Y), then $f_j(Y_1) \not\supseteq f_j(Y_2)$. For $j = 1$, the labels of vertices are chosen as incomparable elements in Q_k, so this holds by construction no matter what order Builder uses to introduce vertices.

For $1 \leq j \leq k-1$, we define f_{j+1} from f_j (Builder defined g_j from g_{j+1} in the proof of the upper bound). Given a $(j+1)$-set Y, consider Y^- and Y^+. Since $(Y^-)^+ = (Y^+)^-$, we are given f_j defined so that $f_j(Y^-) \not\supseteq f_j(Y^+)$. Hence some element of $f_j(Y^+)$ is not in $f_j(Y^-)$ (as downsets in Q_{k-j}). Painter chooses any such element as the label $f_{j+1}(Y)$.

Now consider $(j+1)$-sets Y_1 and Y_2 with $Y_1^+ = Y_2^-$. Both $f_j(Y_2^+)$ and $f_j(Y_2^-)$ are downsets in Q_{k-j}, and Painter chose $f_{j+1}(Y_2) \in f_j(Y_2^+) - f_j(Y_2^-)$. Hence the element $f_{j+1}(Y_2)$ is not below anything in the downset $f_j(Y_2^-)$, including $f_{j+1}(Y_1) \in f_j(Y_1^+) = f_j(Y_2^-)$. This means $f_{j+1}(Y_1) \not\geq f_{j+1}(Y_2)$, as needed for the process to continue.

We have now defined labels for all sets of at most k vertices among those used by Builder. The labels of k-sets lie in Q_1 and hence are colors with heights. When Builder plays a k-set, the color used by Painter is the color in its label. When edges Y_1 and Y_2 are consecutive in a monotone tight path in color i, so $Y_1^+ = Y_2^-$, the property $f_k(Y_1) \not\geq f_k(Y_2)$ implies that the height of the label in Q_1 strictly increases. Since the chains in Q_1 have only $m-1$ elements, no monochromatic monotone copy of $P_r^{(k)}$ occurs. □

We restrict vertex labels to an antichain in Q_k because Builder has the power to introduce new vertices between old vertices, and when vertex x is to the left of vertex y Painter needs to find an element in the label of y that is not in the label of x. If the vertices were known in advance, then the vertex Ramsey result $R_t(P_r^{(k)}) = |Q_k| + 1$ would already allow Painter to survive $|Q_k|/k$ edges in the on-line game. On the other hand, our arguments also yield this result.

Corollary 12 (Moshkovitz and Shapira [30])**.** *$R_t(P_r^{(k)}) = |Q_k| + 1$.*

Proof. When all vertices are known in advance, or when Builder is constrained to add vertices only at the high (i.e., right) end (as in the game studied by Fox et al. [22]), Painter can use all of Q_k as vertex labels, assigning them according to a linear extension, level by level. The initialization $f_1(\{x\}) \not\supseteq f_1(\{y\})$ for any vertices x and y with x before y then holds. The rest of the proof is exactly the same, yielding a lower bound of $|Q_k|/k$ for their game and requiring more than $|Q_k|$ vertices to be played to force a monochromatic copy of $P_r^{(k)}$.

Since the off-line situation is weaker for Builder, we must work harder for the upper bound. All the edges of $\binom{[n]}{k}$ will be played, with $n = |Q_k| + 1$. Painter knows that. If there is a t-coloring that avoids $P_r^{(k)}$, then Painter can prepare to play that coloring, no matter in what order we add the edges. We can allow the labels to be defined as in the on-line game as we add edges.

Initially, the labels are as at the start of the on-line game, as described in Remark 9. We imagine playing all the edges on the first $|Q_k|$ vertices first. If $\mathbf{\Lambda}_k$ appears as a label on a vertex, then as observed in the proof of Theorem 10 there is an edge using the last vertex that when added forces $P_r^{(k)}$. If $\mathbf{\Lambda}_k$ does not appear, then among the first $|Q_k|$ vertices there are vertices x and y (with y later than x) having the same labels. Lemma 6 as edges are processed maintains that two vertices cannot have the same label when one follows the other. Lemma 8 guarantees that when all the edges are processed, all the edges that need to be played to make y follow x have been played. Hence such x and y cannot exist, and $\mathbf{\Lambda}_k$ must occur as a label on a vertex. □

Generalizing these results to ℓ-loose k-uniform monotone paths is straightforward. The off-line value $R_t(P_r^{k,\ell})$ was obtained by Cox and Stolee [12]; they related the value for ℓ-loose paths to the value for tight paths. Translating their result into our notation, they proved

$$R(P_{r_1}^{k,\ell}, \ldots, P_{r_t}^{k,\ell}) = \ell R(P_{r_1}^{(h)}, \ldots, P_{r_t}^{(h)}) + k - \ell h,$$

specifying h to be the unique integer at least 2 such that $\frac{h-2}{h-1} < \frac{k-\ell}{k} \le \frac{h-1}{h}$. Their inequalities for h reduce to $h - 1 < k/\ell \le h$, which says simply $h = \lceil k/\ell \rceil$. With this translation, it becomes clear that our extension of the on-line result from tight paths to ℓ-loose paths yields their extension of the Moshkovitz–Shapira result in the same way that our result for tight paths yields the Moshkovitz–Shapira result (Corollary 12).

In generalizing our results to ℓ-loose paths, the key point is that edges whose last vertices differ by less than ℓ cannot belong to a common ℓ-loose k-uniform monotone path. Again we discuss only the situation where all r_i are equal. Note that explicit bounds on $|Q_h| \cdot |Q_{h-1}| \prod_{i=1}^{h-1} a_i$ in terms of $|Q_h|$ and $|Q_{h-1}|$ are given in Theorem 10 for $h \ge 3$, while $|Q_2| \cdot |Q_1| a_1 = t^2 m^{t+1}$ when $h = 2$.

Theorem 13. *Given $k, \ell, m, t \in \mathbb{N}$ with $t, m \ge 2$ and $\ell \in [k]$, let $r = k + \ell(m-1)$. Also let $h = \lceil k/\ell \rceil$ and $s = k - (h-1)\ell$. With Q_j defined in terms of k, r, t as in the introduction, $R_t(P_r^{k,\ell}) = \ell |Q_h| + s$. Moreover, if $\ell < k$ then $|Q_h|/k \lg |Q_h| \le \tilde{R}_t(P_r^{k,\ell}) \le |Q_h| \cdot |Q_{h-1}| \prod_{i=1}^{h-1} a_i$, where a_i denotes the size of a largest antichain in Q_i, while if $\ell = k$ then $|Q_1|/k \lg |Q_1| \le \tilde{R}_t(P_r^{k,\ell}) \le |Q_1| + 1$.*

Proof. (Sketch) The value ℓ is the *shift*; in an ℓ-loose k-uniform monotone path, it is the number of vertices at the beginning of an edge that are not included in the next edge.

Let Y^- and Y^+ be obtained from a set Y with $|Y| > \ell$ by deleting the last ℓ and the first ℓ elements, respectively. Note that s is the unique member of $[\ell]$ congruent to k modulo ℓ. Given j with $1 \le j \le h$, let $j' = k - (h-j)\ell$; the values of j' are $\{i \in [k] \colon i \equiv k \mod \ell\}$.

Lower Bound *(Painter strategy):* Painter will assign labels to subsets of the vertices whose size is congruent to k modulo ℓ. In particular, the label $f_j(Y)$ will be in Q_{h-j+1} for each j'-set Y of vertices. As noted earlier, in Q_h there is an antichain of size at least $|Q_h|/|Q_{h-1}|$. Painter initially fixes a largest antichain A in Q_h and uses distinct elements of A to name the vertices as they are introduced by Builder. The smallest sets given labels by Painter have size s. For each s-set Y, let $f_1(Y)$ be the element of A that Painter used to name its rightmost vertex.

For $1 \le j \le h$, again we need $f_j(Y_1) \not\ge f_j(Y_2)$ for j'-sets Y_1 and Y_2 such that there exists Y with $Y_1 = Y^-$ and $Y_2 = Y^+$. Note that such a set Y may be introduced after later moves by Builder's introduction of new vertices. However, if Y_1 and Y_2 have the same highest vertex, then this can never occur, and Painter can have the same label on Y_1 and Y_2.

For $1 \le j \le h-1$, define f_{j+1} from f_j by letting $f_{j+1}(Y)$ be any element of $f_j(Y^+)$ not in $f_j(Y^-)$. The inductive proof of the needed property $f_j(Y_1) \not\ge f_j(Y_2)$ is the same as in Theorem 11. The Painter strategy is as defined there: the resulting labels of k-sets under f_k lie in Q_1, and the color used by Painter on an edge played by Builder is the color of the chain containing its label. Since heights must strictly increase along ℓ-loose k-uniform paths, no monochromatic copy of $P_r^{k,\ell}$ occurs. Painter can survive any a_h/k edges, where $a_h = |A|$.

In a restricted version of the game where Builder must add vertices in order from low to high, or where the vertices are specified in advance, Painter can use all elements of Q_h as vertex names (in the order of a linear extension of Q_h). Furthermore, Painter can then use the same name on ℓ consecutive vertices, since edges whose highest vertices differ by less than ℓ cannot belong to the same copy of $P_r^{k,\ell}$, and no vertices will be inserted between two already having names. In addition, the first $s-1$ vertices receive no names from Q_h, since the smallest sets needing labels have size s. Again the process proceeds: s-sets receive as label the element of Q_h assigned to their highest vertex. Note that if $|\max Y_2 - \max Y_1| < \ell$, then Y_1 and Y_2 can never be extended leftward to edges in the same copy of $P_r^{k,\ell}$. In this way, Painter can survive $\ell|Q_h| + s - 1$ vertices. Hence $R_t(P_r^{k,\ell}) \ge \ell|Q_h| + s$, as in [12].

Upper Bound *(Builder strategy):* Builder uses $\ell|Q_h|+s$ vertices, assigning labels to sets whose size is congruent to k modulo ℓ, down to size s. Actually, Builder assigns labels only to sets whose last s vertices are consecutive, called *basic* sets; Builder also plays only basic edges. Henceforth consider only basic sets. Note that there are $\ell|Q_h| + 1$ basic sets of size s.

Builder assigns a label in Q_1 to edges and a label in Q_{h-j+1} to the sets of size j' for $h > j \ge 1$ (note that $j' = j$ when $\ell = 1$). For an edge Y with color i in G, the label $g_h(Y)$ is the element of height p on the ith chain in Q_1, where p is the number of edges in the longest ℓ-loose k-uniform monotone path with last edge Y in the current colored hypergraph. For $h > j \ge 1$, the *precursors* of a j'-set Y are the $(j'+\ell)$-sets obtained by adding ℓ elements to Y that are smaller than the least element of Y; that is, the precursors are the sets Z such that $Z^+ = Y$.

With these generalizations of earlier definitions, the definitions of g_j for $1 \le j < h$ and the relation of "follows" are the same as in Definitions 4 and 5. In particular, note that if Y_2 follows Y_1, then the rightmost element of Y_2 must be at least ℓ positions to the right of the rightmost element of Y_1. The statement and proof of Lemma 6 are the same, except that g_k and Q_{k-j+1} generalize to g_h and Q_{h-j+1}, and $\binom{[n]}{j}$ becomes $\binom{[n]}{j'}$. In Definition 7 and Lemma 8 we generalize j-set and $(j+1)$-set to basic j'-set and basic $(j'+\ell)$-set, and again k generalizes to h in various subscripts.

Now Remark 9 and Theorem 10 also generalize naturally to yield $\tilde{R}_t(P_r^{k,\ell}) \le |Q_h|{\cdot}|Q_{h-1}| \prod_{i=1}^{h-1} a_i$ for $\ell < k$ (or equivalently $h \ge 2$). Note that the labels $\mathbf{V}_h^0, \ldots, \mathbf{V}_h^{h-2}$ of the chain at the bottom of

Q_h are assigned to the first $(h-1)\ell$ basic sets of size s, where each label is used on s-sets starting at ℓ consecutive values. For $0 \le i \le (h-1)\ell - 1$, the set $[i+1, i+s] \in \binom{[n]}{s}$ is assigned label $\mathbf{V}_h^{\lfloor i/\ell \rfloor}$. These labels never change, since no edge can be played ending at one of these sets.

In particular, the s-set $[(h-2)\ell+1, (h-2)\ell+s]$ and all subsequent s-sets initially have the label $\mathbf{V}_h^{h-2}$. Since they never go down, they always avoid the bottom $h-2$ labels in Q_h and the top label $\mathbf{\Lambda}_h$, as long as none of them becomes $\mathbf{\Lambda}_h$. Among these, Builder considers *restricted basic s-sets*, having the form $[i\ell+1, i\ell+s]$ with $h-2 \le i \le |Q_h| - 1$. There are $|Q_h| - h + 2$ such sets, and their labels are confined to a set of $|Q_h| - h + 1$ labels in Q_h. Therefore, when Builder is ready to move the pigeonhole principle guarantees that some label in Q_h is assigned to at least two restricted basic s-sets. This guarantees the existence of two basic s-sets X and Y with the same label whose rightmost vertices differ by at least ℓ. By the generalization of Lemma 6, Y does not follow X. Builder can then play edges as guaranteed by the generalization of Lemma 8 to make Y follow X, which as in Theorem 10 makes the label of Y go up. A label can increase at most $|Q_{h-1}| - h$ times before reaching $\mathbf{\Lambda}_h$.

Hence Builder can play to force an s-set Z with label $\mathbf{\Lambda}_h$ whose highest element is at most $\ell(|Q_h|-1)+s$. As in Theorem 10, some $(k-\ell)$-set Y ending with Z will then have label $(m-1, \ldots, m-1)$, the top element of Q_2. Since $n = \ell|Q_h| + s$, there remain ℓ vertices beyond Z. By playing the k-set consisting of Y and the last ℓ vertices, Builder wins.

Since in fact the label of the leftmost restricted basic s-set never changes, the number of edges played is at most

$$1 + [(|Q_h| - h + 1)(|Q_{h-1}| - h) + 1] \prod_{i=1}^{h-1} a_i,$$

which for $h \ge 2$ is at most $|Q_h| \cdot |Q_{h-1}| \prod_{i=1}^{h-1} a_i$. Note, however, that since Builder used only $\ell|Q_h| + s$ vertices, we have $R_t(P_r^{k,\ell}) = \ell|Q_h| + s$. In the case $h = 1$ (that is, $\ell = s = k$), Builder simply plays the basic edges (intervals) $[ik+1, (i+1)k]$ for $0 \le i \le |Q_1|$. Since $[i'k+1, (i'+1)k]$ follows $[ik+1, (i+1)k]$ whenever $i < i'$, Painter is forced to use distinct colors on the edges and loses. This gives the desired upper bounds on $R_t(P_r^{k,\ell})$ and $\tilde{R}_t(P_r^{k,\ell})$ for $\ell = k$. □

4 Directed Graphs

The ordered Ramsey problem can be described using directed graphs and hypergraphs. An orientation of an edge is a permutation of its vertices. An ordered hypergraph can be viewed as a directed hypergraph in which the orientation of each edge is the permutation inherited from the vertex ordering. In particular, an ordered tight path is a directed hypergraph in which the edges are the k-sets of consecutive vertices, oriented in increasing order in each edge. In a general k-uniform directed hypergraph, k-sets may appear up to $k!$ times, once with each orientation.

When Builder has the power to play edges of a general directed hypergraph in seeking to force a monochromatic directed tight path, Painter can follow a strategy like that above, using an antichain in Q_k for vertex labels. All oriented j-tuples must be labeled, for $1 \le j \le k$, so the lower bound will be $|Q_k|/(k \lg |Q_k|)$.

Let us consider this problem in the off-line setting for $k = 2$. Hence we are seeking the size Ramsey number of the directed path P_{m+1} in the model where arbitrary host digraphs are allowed. The trivial upper bound is again $\binom{m^t+1}{2}$, achieved by playing increasing edges for all pairs on $R_t(P_{m+1})$ vertices in the ordered setting. For the off-line model, Builder is weaker, and we obtain a better lower bound than for the on-line game.

Theorem 14. *In the setting of directed graphs, $\hat{R}_t(P_{m+1}) \geq \binom{|B|+1}{2}$, where B is the family of elements in $\{0, \ldots, m-1\}^t$ with sum $\lfloor (m-1)t/2 \rfloor$. As noted in Theorem 2, $|B| \geq \frac{2}{3} m^{t-1}/\sqrt{t}$.*

Proof. A graph with fewer than $\binom{|B|+1}{2}$ edges is $(|B|-1)$-degenerate and hence $|B|$-colorable. Hence we may suppose that the underlying undirected graph of the host digraph is $|B|$-colorable. Painter specifies a proper vertex coloring whose colors correspond to the elements of B. Each vertex v has a label $a(v) \in B$, and adjacent vertices always have distinct labels. As in Theorem 2, Painter can choose for each (directed) edge uv a color i such that $a_i(v) > a_i(u)$. Again at every vertex w the length of any path in color i reaching w is at most $a_i(w)$, since the ith coordinate strictly increases along paths whose edges have color i. □

The off-line size Ramsey problem for paths in digraphs (with $t = 2$) was also studied by Ben-Eliezer, Krivelevich, and Sudakov [5]. They considered a model where Builder can present only oriented graphs (no 2-cycles) and a model where Builder can present any digraph, yielding size Ramsey numbers S_{ori} and S_{dir} respectively. Note that $S_{dir} \leq S_{ori}$ when the parameters are equal.

For the general digraph model, which we considered above, the arguments of [5] yield the following bounds:

$$\left(\frac{m+1}{3t-3}\right)^{2t-2} \leq S_{dir} \leq 4(m+1)^{2t-2}.$$

Since they focus on constant t, they state the result as $S_{dir} = \Theta(m^{2t-2})$. Since $|B| \geq \frac{2}{3} m^{t-1}/\sqrt{t}$, our lower bound strengthens theirs.

Their lower bound for S_{ori} is higher than their upper bound for S_{dir} (Bucic, Letzter, and Sudakov [7] improved their upper bound on S_{ori}). They prove

$$C_1(t)\frac{(m+1)^{2t-2}(\log(m+1))^{\frac{1}{t-1}}}{(\log\log(m+1))^{\frac{t+1}{t-1}}} \leq S_{ori} \leq C_2(m+1)^{2t-2}(\log(m+1))^2$$

where C_2 is an absolute constant, but $C_1(t)$ depends on t. They require

$$C_1(t) < \frac{C^{1/(t-1)}}{8(2t-2)^{t-1}(16(t-1)^2)^t}$$

for some absolute constant C. Therefore, their lower bound is at most

$$\frac{1}{(2t)^{3t}}\frac{(m+1)^{2t-2}(\log(m+1))^{\frac{1}{t-1}}}{(\log\log(m+1))^{\frac{t+1}{t-1}}},$$

which remains smaller than ours when t grows faster than $\sqrt{\log\log m}$.

References

[1] V. B. Alekseev, The number of monotone k-valued functions. (Russian) *Problemy Kibernet.* 28 (1974), 5–24, 278; correction, ibid. 29 (1974), 248.

[2] M. Balko, J. Cibulka, K. Král, and J. Kynčl, Ramsey numbers of ordered graphs, preprint available as arXiv:1310.7208.

[3] J. Beck, On size Ramsey number of paths, trees and cycles I. *J. Graph Th.* 7 (1983), 115–130.

[4] J. Beck, Achievement games and the probabilistic method. In *Combinatorics, Paul Erdős is eighty*, Vol. 1, 51–78, Bolyai Soc. Math. Stud., János Bolyai Math. Soc., Budapest, 1993.

[5] I. Ben-Eliezer, M. Krivelevich, and B. Sudakov, The size Ramsey number of a directed path. *J. Combin. Th. (B)* 102 (2012), 743–755.

[6] B. Bollobás, *Random graphs* (Second edition). Cambridge Studies in Advanced Mathematics, 73. Cambridge University Press, Cambridge, 2001.

[7] M. Bucić, S. Letzter, and B. Sudakov, Monochromatic paths in random tournaments, preprint.

[8] J. Butterfield, T. Grauman, W.B. Kinnersley, K.G. Milans, C. Stocker, and D.B. West, On-line Ramsey theory for bounded degree graphs. *Electron. J. Combin.* 18 (2011), no. 1, Paper 136, 17 pp.

[9] D. Conlon, On-line Ramsey numbers. *SIAM J. Discrete Math.* 23 (2009), 1954–1963.

[10] D. Conlon, J. Fox, C. Lee, and B. Sudakov, Ordered Ramsey numbers. *J. Combin. Th. (B)* 122 (2017), 353–383. preprint available as arXiv:1410.5292.

[11] D. Conlon, J. Fox, and B. Sudakov, Hypergraph Ramsey numbers. *J. Amer. Math. Soc.* 23 (2010), 247–266.

[12] C. Cox and D. Stolee, Ordered Ramsey numbers of loose paths and matchings. *Discrete Math.* 339 (2016), 499–505.

[13] J. Cyman and T. Dzido, A note on on-line Ramsey numbers for quadrilaterals. *Opuscula Math.* 34 (2014), 463–468.

[14] J. Cyman, T. Dzido, J. Lapinskas, and A. Lo, On-line Ramsey numbers of paths and cycles. *Electron. J. Combin.* 22 (2015), no. 1, Paper 1.15, 32 pp.

[15] R.P. Dilworth, A decomposition theorem for partially ordered sets. *Ann. of Math.* (2) 51, (1950), 161–166.

[16] A. Dudek and P. Prałat, An alternative proof of the linearity of the size-Ramsey number of paths, *Combin. Probab. Comput.*, 24(3) (2015), 551-555.

[17] A. Dudek and P. Prałat, On some multicolour Ramsey properties of random graphs, *SIAM J. Discrete Math.*, 31 (2017), no. 3, 2079–2092.

[18] D. Duffus, H. Lefmann, and V. Rödl, Shift graphs and lower bounds on Ramsey numbers $r_k(l;r)$. *Discrete Math.* 137 (1995), 177–187.

[19] P. Erdős, R. J. Faudree, C. C. Rousseau, and R. H. Schelp, The size Ramsey number, *Periodica Mathematica Hungarica* 9 (1978), 145–161.

[20] P. Erdős and G. Szekeres, A combinatorial problem in geometry, *Compositio Math.* 2 (1935), 463–470.

[21] R. J. Faudree and R. H. Schelp, A survey of results on the size Ramsey number, Paul Erdős and his mathematics, II (Budapest, 1999), *Bolyai Soc. Math. Stud.* 11 (János Bolyai Math. Soc., Budapest, 2002), 291–309.

[22] J. Fox, J. Pach, B. Sudakov, and A. Suk, Erdős-Szekeres-type theorems for monotone paths and convex bodies, *Proc. London Math. Soc.* 105 (2012), 953–982.

[23] J. A. Grytczuk, M. Hałuszczak, and H. A. Kierstead, On-line Ramsey Theory. *Electron. J. Combin.* 11 (2004), Paper R57, 10 pp.

[24] J. A. Grytczuk, H. A. Kierstead, and P. Pralat, On-line Ramsey Numbers for Paths and Stars. *Discrete Math. Theor. Comp. Science* 10 (2008), 63–74.

[25] H. A. Kierstead and G. Konjevod, Coloring number and on-line Ramsey theory for graphs and hypergraphs. *Combinatorica*, 29 (2009), 49–64.

[26] W.B. Kinnersley and D.B. West, Multicolor on-line degree Ramsey numbers of trees. *J. Comb.* 3 (2012), 91–100.

[27] A. Kurek and A. Ruciński, Two variants of the size Ramsey number, *Discuss. Math. Graph Theory* 25 (2005), 141–149.

[28] S. Letzter, Path Ramsey number for random graphs. *Combin. Probab. Comput.* 25 (2016), no. 4, 612–622.

[29] K.G. Milans, D. Stolee, and D.B. West, Ordered Ramsey theory and track representations of graphs. *J. Comb.* 6 (2015), 445–456.

[30] G. Moshkovitz and A. Shapira, Ramsey Theory, Integer Partitions and a New Proof of the Erdős-Szekeres Theorem, *Advances in Math.* 262 (2014), 1107–1129.

[31] D. Mubayi, Variants of the Erdős–Szekeres and Erdős–Hajnal Ramsey problems, *European J. Combin.* 62 (2017), 197–205.

[32] D. Mubayi and A. Suk, Off-diagonal hypergraph Ramsey numbers, *J. Combin. Th. Ser. B* 125 (2017), 168–177.

[33] J. Nešetřil and V. Rödl, The Ramsey property for graphs with forbidden complete subgraphs. *J. Combin. Th. Ser. B* 20 (1976), no. 3, 243–249.

[34] Š. Petříčková, Online Ramsey theory for planar graphs. *Electron. J. Combin.* 21 (2014), no. 1, Paper 1.64, 14 pp.

[35] P. Prałat, A note on small on-line Ramsey numbers for paths and their generalization, *Australas. J. Combin.* 40 (2008), 27–36.

[36] P. Prałat, A note on off-diagonal small on-line Ramsey numbers for paths, *Ars Combin.* 107 (2012), 295–306.

[37] F. P. Ramsey, On a problem of formal logic, *Proc. Lond. Math. Soc.* 30 (1930), 264–286.

[38] D. Reimer, The Ramsey size number of dipaths, Discrete Math. 257 (2002), 173–175.

[39] D. Rolnick, Trees with an on-line degree Ramsey number of four. *Electron. J. Combin.* 18 (2011), no. 1, Paper 173, 21 pp.

[40] D. Rolnick, The on-line degree Ramsey number of cycles. *Discrete Math.* 313 (2013), 2084–2093.

[41] A. Seidenberg, A simple proof of a theorem of Erdos and Szekeres, *J. London Math. Soc.* 34 (1959), 352.

[42] V.G. Vizing, On an estimate of the chromatic class of a p-graph. *Diskret. Anal.* 3 (1964), pp. 25–30.