


Charles University in Prague
Faculty of Mathematics and Physics

DOCTORAL THESIS



Wilf-Type Classifications

**Extremal and Enumerative Theory
of Ordered Structures**

Vít Jelínek

Department of Applied Mathematics

Supervised by doc. RNDr. Martin Klazar, Dr.
Branch: Discrete models and algorithms (I4)

2008

Děkuji docentu Martinu Klazarovi za užitečné rady a za jeho trpělivost a podporu při sepisování této práce.

I declare that I wrote this thesis myself, and all the sources of information I used are properly referenced. I agree with the publication of this thesis.

Prague, 18. 9. 2008

Vít Jelínek

Contents

Introduction: Basic ordered structures	7
Permutations	8
Words	8
Matrices	9
Fillings of diagrams	9
Ordered graphs	11
Set partitions	13
Summary	14
1 Wilf classes of permutations	18
1.1 Symmetries of permutations	18
1.2 Non-trivial Wilf equivalences	18
1.3 The shape-Wilf equivalence	19
2 Identities involving diagonal patterns	24
2.1 Known results on diagonal patterns	24
2.2 Constrained rectangular fillings	26
3 Wilf order	34
3.1 Shape-Wilf order of small patterns	35
3.2 Skew order	35
4 Involutions	40
5 Words	45
5.1 Basic terminology	45
5.2 Compositions	46
5.3 Strong equivalence of words	46
5.4 Semi-standard fillings of Ferrers shapes	46
5.5 Patterns equivalent to 12^k	47
5.6 Patterns equivalent to 1232^k	49
6 Partitions	52
6.1 Basic facts and previous results	53
6.2 General classes of equivalent patterns	54
6.3 Patterns equivalent to $12 \cdots m(m+1)$	63
6.4 Binary patterns	64
6.5 Sporadic equivalences	75
6.6 The pattern 12112	78

Conclusion: Beyond Wilf-type classifications	88
Hereditary classes	88
Growth rates	88
Speeds of hereditary classes	89
Ordered graphs	90
Atomic and molecular relational structures	91
Further reading	92
Acknowledgement	92
Bibliography	93
A Notation	97

Název práce: Extremální a enumerativní teorie uspořádaných struktur

Autor: Vít Jelínek

Katedra: Katedra aplikované matematiky

Vedoucí doktorské práce: doc. RNDr. Martin Klazar, Dr.

e-mail vedoucího: klazar@kam.mff.cuni.cz

Abstrakt: Počítání permutací, které neobsahují danou zakázanou podpermutaci, patří mezi základní témata enumerativní kombinatoriky. V této práci ukážeme, že mnohé výsledky tohoto výzkumu lze aplikovat i na obecnější struktury, než jsou permutace, například na uspořádané grafy, na slova nad uspořádanou abecedou, nebo na množinové rozklady.

Hlavní část této práce souvisí s *Wilfovou ekvivalencí*, definovanou následovně: dvě permutace σ a τ jsou Wilf-ekvivalentní, pokud pro každé n platí, že počet permutací řádu n neobsahujících σ je stejný jako počet permutací řádu n neobsahujících τ . Obdobně lze definovat ekvivalence na obecnějších objektech, než jsou permutace. V této práci systematicky zkoumáme analogie Wilfovy ekvivalence u obecnějších struktur, jako jsou například matice, diagramy, permutace multimnožin, nebo množinové rozklady. Najdeme nové třídy Wilf-ekvivalentních objektů a ověříme, že mezi permutacemi multimnožin a mezi rozklady malé délky žádná další Wilf-ekvivalentní dvojice neexistuje.

Abychom začlenili tyto výsledky do kontextu, tak v závěru práce krátce zmíníme několik souvisejících témat enumerativní kombinatoriky, jako například výzkum rychlosti růstu dědičných tříd permutací a uspořádaných grafů, nebo výzkum atomických tříd relačních struktur.

Klíčová slova: Wilfova ekvivalence, množinové rozklady, relační struktury

Title: Extremal and enumerative theory of ordered structures

Author: Vít Jelínek

Department: Department of Applied Mathematics

Supervisor: doc. RNDr. Martin Klazar, Dr.

Supervisor's e-mail address: klazar@kam.mff.cuni.cz

Abstract: The study of permutations that avoid a given pattern is among the main topics of enumerative combinatorics. In this thesis, we show that many results of this study may be extended to more general structures, such as ordered graphs, words over a linearly ordered alphabet, or set partitions.

The main part of this thesis is related to *Wilf equivalence*, defined as follows: two permutations σ and τ are Wilf equivalent if, for every n , the number of permutations of order n that avoid σ is the same as the number of permutations of order n that avoid τ . We may define analogous equivalence relations on more general objects than permutations. In this thesis, we investigate the Wilf-type equivalences of more general structures, such as matrices, diagrams, multiset permutations, or set partitions. We present new examples of Wilf equivalent structures, and we show that among multiset permutations and set partitions of small size there are no other equivalent pairs.

To put these results in context, we briefly mention, at the end of this thesis, several related topics of enumerative combinatorics, such as the research of growth rates of hereditary classes of permutations and ordered graphs, or the research of atomic classes of relational structures.

Keywords: Wilf equivalence, set partitions, relational structures

Introduction: Basic ordered structures

In this thesis, we use the term ‘ordered structure’ to refer to several types of combinatorial objects that can be regarded as generalizations of permutations. We will be interested in pattern avoidance of these structures, as an extension of the intensively studied concept of pattern avoidance of permutations.

For the study of pattern-avoiding permutations, one of the central notions is the Wilf equivalence. Two permutations σ and τ are said to be Wilf equivalent, if for every n the number of permutations of order n that avoid τ is the same as the number of permutations of order n that avoid σ .

Equivalence relations analogous to Wilf equivalence can be studied for any family of combinatorial objects for which there is a well-defined concept of pattern avoidance. These Wilf-type equivalences are the topic of this thesis.

In the research of Wilf equivalence (and its analogs for other types of structures) the natural ultimate goal would be to find the complete classification of all the equivalent pairs of objects. Unfortunately, this goal seems far out of reach, mostly because our understanding of larger patterns is very limited. In this situation, it is natural to first focus on the Wilf-type classification of small patterns, where we can use computer-generated enumeration data to find all possible candidates for Wilf-type equivalence, and then try to prove the equivalence of these candidates. Thus, a more realistic goal of the study of pattern avoidance is the Wilf-classification of all the patterns that are within reach of computerized enumeration, with emphasis put on criteria that can be generalized to larger patterns as well.

In the study of pattern-avoiding permutations, this approach has provided the full classification of permutation patterns of size at most seven, which appears to be the bound of computerized enumeration. Many techniques and results developed in the course of this classification are also applicable to larger patterns.

In this thesis, we apply a similar approach to other types of pattern-avoiding structures, most of which can be regarded as generalizations of permutations. Among other results, we present the classification of words (which may also be regarded as multiset permutations) of length at most six, and of set partitions of length at most seven. Most of these results have been previously published as joint papers with several coauthors. Our presentation in this thesis often closely follows the journal version, with one important exception: in this thesis, we concentrate on the combinatorial arguments, and completely ignore the accompanying computerized enumeration. The reader which is interested in this aspect of the work may consult the original sources given in the references.

In the rest of this chapter, we give an overview of the main ordered structures, and explain the relationships between them. The main part of the thesis starts after this chapter. Since almost all the thesis is devoted to a very narrow topic of Wilf-type classifications, we have decided to add, after the main part of the thesis, a brief concluding chapter which aims to provide an overview of alternative approaches to the enumeration of pattern-avoiding classes.

For reader’s convenience, we summarized the notation we use in Appendix A.

Permutations

Permutations are the prototypical class of ordered structures. We define a permutation π of order n as a sequence $\pi_1\pi_2\cdots\pi_n$ in which each number from the set $[n] = \{1, 2, \dots, n\}$ appears exactly once. We let \mathcal{S}_n denote the set of all permutations of order n . An *involution* is a permutation $\pi = \pi_1\pi_2\cdots\pi_n$ which satisfies the equivalence $\pi_i = j \iff \pi_j = i$ for every $i, j \in [n]$.

Let $\sigma \in \mathcal{S}_k$ and $\tau \in \mathcal{S}_n$ be two permutations, with $k \leq n$. We say that τ *contains* σ if there are k indices $i(1), i(2), \dots, i(k)$ with $1 \leq i(1) < i(2) < \cdots < i(k) \leq n$, such that $\sigma_a < \sigma_b$ if and only if $\tau_{i(a)} < \tau_{i(b)}$, for any $a, b \in [k]$. If τ does not contain σ , we say that τ *avoids* σ .

We let $\mathcal{S}_n(\sigma)$ denote the set of permutations of order n that avoid a permutation σ . More generally, if \mathcal{F} is a set of permutations, then $\mathcal{S}_n(\mathcal{F})$ is the set of all the permutations of order n that avoid all the elements of \mathcal{F} .

By a result of Marcus and Tardos [50], it is known that for every permutation σ , the limit $\lim_{n \rightarrow \infty} \sqrt[n]{|\mathcal{S}_n(\sigma)|}$ exists and is finite. This limit is known as the *Stanley–Wilf limit* of σ .

Although we chose to define a permutation as a sequence of integers, there are other ways to represent a permutation, e.g., a permutation matrix, a permutation graph, or a permutation matching. Each of these representations offers a natural way to embed the containment order of permutations into the containment order of more general structures. In the remaining sections of this brief introductory chapter, we will introduce these structures, and explain their relationship with permutations.

Words

Words over a linearly ordered alphabet $[k] = \{1, 2, \dots, k\}$ are a natural encoding for permutations of multisets with at most k distinct elements. Thus, the notion of pattern avoidance of words is a natural extension of pattern avoidance of permutations. Let us not define this notion formally.

Let A be an arbitrary set, called *the alphabet*. Let A^n be the set of all the sequences of length n whose elements belong to A . Such sequences are called *words of length n over A* . A *subword* of a word w is a (not necessarily contiguous) subsequence of the symbols of w .

In this thesis, we will assume that the alphabet is a subset of \mathbb{N} , unless otherwise noted. Words over the alphabet $[k]$ will be called *k -ary words*. Thus, a permutation of order n may be regarded as a special case of an n -ary word of length n .

Let $v = v_1v_2\cdots v_n$ and $w = w_1w_2\cdots w_n$ be two words of length n over the alphabet \mathbb{N} . We say that v and w are *order-isomorphic* if for every pair of indices $i, j \in [n]$ we have the equivalence $v_i < v_j \iff w_i < w_j$. Note that if v and w are order-isomorphic words, then $v_i = v_j$ if and only if $w_i = w_j$. We will say that a word w *contains a copy of v* , or simply w *contains v* , if w has a subword which is order-isomorphic to v . If w does not contain any copy of v , we say that w *avoids v* . We let $A^n(v)$ denote the set of words from A^n that avoid v . It is clear that the containment relation of permutations defined in the previous section is a special case of the containment relation of words.

We remark that our notion of containment of words is substantially based on the fact the underlying alphabet is linearly ordered. We should mention that there are other (perhaps more natural) notions of word-containment, which do not refer to any ordering of the alphabet. In this thesis, we will not consider these alternative notions, since it would make us drift too far away from our main topic. An interested reader may find more information in the work of Klazar [42, 44].

Let w be a word over the alphabet \mathbb{N} , and assume that k is the largest integer that appears as a symbol in w . The word w is called *reduced* if every symbol from the set $[k]$ appears in w at least once. It is easy to see that every word w over \mathbb{N} that contains k distinct symbols is order-isomorphic to a unique reduced word y , where y is a k -ary word, which we will call *the reduction of w* . Of course, a word x contains a word w if and only if x contains the reduction of w . Thus, when we study pattern-avoiding classes of words, we may restrict our attention to the situation when the avoided pattern is a reduced word.

Matrices

Let X be a set of integers. We let $X^{k \times \ell}$ denote the set of matrices with k rows and ℓ columns, whose elements belong to X . We will always use the ‘cartesian’ numbering of rows and columns, i.e., we will assume that columns are numbered from left to right, and rows are numbered from bottom to top. An intersection of a row and a column will be called *a cell* of the matrix. In a matrix $M \in \mathbb{N}^{k \times \ell}$, we let M_{ij} denote the cell in row i and column j . A *01-matrix* is a matrix whose cells are equal to 0 or 1; in such case, we will speak of *0-cells* and *1-cells*, respectively. A *submatrix of a matrix M* is obtained from M by erasing some of its rows or columns.

Let $P \in \mathbb{N}^{k \times \ell}$ and $M \in \mathbb{N}^{m \times n}$ be two matrices. We say that M *contains a copy of P* , or briefly M *contains P* , if M has a submatrix M' with k rows and ℓ columns, such that for every $i \in [k]$ and $j \in [\ell]$ we have the inequality $P_{ij} \leq M'_{ij}$. In this thesis, we almost always restrict ourselves to situations when the pattern P is a 01-matrix.

A *permutation matrix of order n* is a 01-matrix M with n rows and n columns, with the property that every row and every column of M has exactly one 1-cell. We will assume that a permutation matrix M of order n represents the permutation $\tau = \tau_1 \cdots \tau_n \in \mathcal{S}_n$ defined by the relation $\tau_j = i$ if and only if $M_{ij} = 1$. This correspondence provides a bijection between the set \mathcal{S}_n and the set of permutation matrices of order n . Note that the symmetric permutation matrices correspond precisely to involutions.

If σ and τ are two permutations, and M_σ and M_τ their corresponding permutation matrices, then it is not difficult to see that τ contains σ if and only if M_τ contains M_σ , which happens if and only if M_τ has M_σ as a submatrix. This shows that the containment relation of permutations can be viewed as a special case of the containment relation of matrices.

In fact, this reasoning may be extended to words over the alphabet \mathbb{N} . Let $w = w_1 w_2 \cdots w_n \in \mathbb{N}^n$ be a word, and let $m \in \mathbb{N}$ be the largest symbol appearing in w . We may represent w by a 01-matrix $M \in \{0, 1\}^{m \times n}$ where the j -th column of M has a 1-cell in row w_j and all the remaining cells in this column are equal to zero. Notice that if w is in fact a permutation, then M is its corresponding permutation matrix. If x is a reduced word and y an arbitrary word over \mathbb{N} , then y contains x if and only if the matrix representing y contains the matrix representing x .

Fillings of diagrams

The notion of a matrix can be further generalized, by relaxing the assumption that all the rows and all the columns have the same length. This idea leads to the concept of a filling of a diagram. In full generality, a *diagram* is a finite set of cells in the plane, where each cell is a square of unit size whose vertices have integer coordinates.

We will assume that the rows of the diagram are numbered from bottom to top, and the columns are numbered from left to right. The numbering is fixed in such a way that




Figure 1: Examples of special types of diagrams. From left to right: a Ferrers shape, a skew shape, a stack shape, and a moon shape.

the first numbered column contains the leftmost cell of the diagram and the last numbered column contains its rightmost cell; the numbering of rows is fixed analogously. We say that the i -th row of the diagram *intersects* the j -th column, if the diagram contains the cell that belongs to the i -th row and j -th column. For a diagram D , we let $r(D)$ and $c(D)$ denote, respectively, the number of rows and the number of columns of D .

For our purposes, we will mostly use four special types of diagrams: the Ferrers diagrams, the skew diagrams, the stack diagrams and moon diagrams (see Fig. 1). A *Ferrers diagram* (also known as a *Ferrers shape*) is a diagram satisfying the following conditions:

- The rows of the diagram are contiguous and left-justified, i.e., if a row has exactly k -cells, then these cells appear in the columns $1, 2, \dots, k$.
- For every $i \geq 1$, the length of row i is greater than or equal to the length of the row $i + 1$.

A *skew diagram* (or *skew shape*) is a diagram that can be obtained as a difference of two Ferrers diagrams that share a common bottom-left corner. Formally, a skew shape is a diagram with the property that the vertical coordinates of the bottom cells of its columns form a nonincreasing sequence, and the vertical coordinates of the top cells of its columns form a nonincreasing sequence as well.

A *stack diagram* (also known as a *stack polyomino* or *stack shape*) is a diagram with the following properties:

- Each row is contiguous, i.e., if two cells in the same row belong to the diagram, then all the cells between these two also belong to the diagram.
- If a column intersects row i , then the column intersects all the rows $1, 2, \dots, i$.

A stack polyomino can also be regarded as a diagram obtained by gluing a copy of a Ferrers shape reflected along a vertical axis and glued to another (non-reflected) Ferrers shape.

A *moon diagram* (or *moon polyomino*) is a diagram with the following properties:

- Both the rows and the columns of the diagram are contiguous.
- Each two rows are *comparable*, which means that the set of columns intersected by a row i is either a subset or a superset of the set of columns intersected by a row j . (Notice that this condition is equivalent to saying that each two columns are comparable.)

Clearly, every Ferrers shape is also a stack polyomino and a skew shape, and every stack shape is also a moon shape.

A *filling* of a diagram is a mapping which assigns to each cell of the diagram an integer. A matrix may be viewed as a filling of a rectangular diagram. We will now define

a containment relation of fillings which extends the containment of matrices defined in the previous section. Let D be a filling of a diagram. Let P be another filling (P stands for ‘pattern’), and let $r = r(P)$ and $c = c(P)$. We say that D contains a copy of P if, in the filling D , we may choose r row indices $i_1 < i_2 < \dots < i_r$ and c column indices $j_1 < j_2 < \dots < j_c$ such that the following conditions are satisfied:

- For every $k \in [r]$ and $\ell \in [c]$, the k -th row of P intersects the ℓ -th column of P if and only if, in the diagram D , the row i_k intersects the column j_ℓ . In other words, the rows $i_1 < i_2 < \dots < i_r$ and the columns $j_1 < j_2 < \dots < j_c$ induce in D a subdiagram with the same shape as P .
- If, in the filling P , the k -th row intersects the ℓ -th column, then the cell that corresponds to this intersection is filled with a number that is less than or equal to the number in the intersection of row i_k and column j_ℓ in D .

The *transpose* of a diagram F , denoted by F^T , is the diagram obtained by flipping F along the main diagonal; in other words, F^T contains the cell (i, j) if and only if F contains the cell (j, i) . The transpose of a filling is defined analogously. A diagram or a filling is called *symmetric* if it is equal to its transpose. Note that while a symmetric diagram may have a non-symmetric filling, any filling of a non-symmetric diagram is necessarily non-symmetric.

Note that the transpose of a Ferrers shape is also a Ferrers shape, the transpose of a skew shape is a skew shape, and the transpose of a moon shape is a moon shape.

Let us now define several special types of fillings, which will be later useful. A 01-filling is a filling that only uses the numbers 0 and 1. A 01-filling is called *semi-standard* if each column has exactly one 1-cell. A *transversal* (also called a *standard filling*) is a 01-filling which contains exactly one 1-cell in every row as well as in every column. Notice that transversals of rectangular shapes correspond exactly to permutation matrices. A *zero row* (or *zero column*) is a row (or column) in a filling that only contains zeros. A filling is called *dense* if it has no zero rows and no zero columns. A 01-filling is called *sparse* if every row and every column contains at most one 1-cell. A 01-filling is called *semi-sparse* if every column has at most one 1-cell.

In this thesis, we follow the convention that in figures of fillings or matrices, all the zeros are omitted, i.e., the 0-cells are represented as empty boxes. This makes the figures less cluttered.

Ordered graphs

An *ordered graph* $G = (V, E, \prec)$ is a graph with vertex set V and edge set E , whose vertices are linearly ordered by the relation \prec . An intuitive way to represent an ordered graph is to draw its vertices as a sequence of points on a horizontal line, where the left-to-right ordering of the points corresponds to the linear order \prec ; the edges are then represented as circular arcs connecting the corresponding pair of vertices (see Fig. 2 for an example). Thus, we will often say, e.g., that a vertex v is to the left of a vertex w , which means that v is smaller than w in the ordering \prec . Most of the time, we will work with ordered graphs whose vertices are integers. In such situation, we always assume that the ordering of vertices corresponds to the usual ordering of integers, and we write $G = (V, E)$ instead of $G = (V, E, \prec)$.

We say that two ordered graphs $G = (V, E, \prec)$ and $H = (W, F, \triangleleft)$ are *isomorphic* if there is a bijection $\phi: V \rightarrow W$, with the property that $\{u, v\}$ is an edge of G if and only if $\{\phi(u), \phi(v)\}$ is an edge of H , and $u \prec v$ if and only if $\phi(u) \triangleleft \phi(v)$. Since an isomorphism




Figure 2: An example of an IM-free ordered graph with its adjacency filling. Note that the empty boxes of the filling represent 0-cells.

of ordered graphs must preserve the order of the vertices, it follows that no two distinct graphs on the same ordered vertex set can be isomorphic.

We say that an ordered graph $G = (V, E, \prec)$ is a *subgraph* of $H = (W, F, \triangleleft)$ if V is a subset of W , E is a subset of F , and the two ordering relations \prec and \triangleleft agree on V . Induced ordered subgraphs are defined analogously.

We say that a vertex v of an ordered graph is a *left-vertex*, or briefly L-vertex, if all the neighbours of v are to the right of v . Similarly, a *right-vertex*, or R-vertex, is a vertex that is to the right of all its neighbors. Thus, an isolated vertex is both an L-vertex and an R-vertex. A middle-vertex, or M-vertex, is a vertex that is neither left nor right. An ordered graph is called *M-free* if it has no M-vertex. It is called *IM-free* if it has no M-vertex and no isolated vertex.

Let us now describe a bijection between IM-free ordered graphs and dense 01-fillings of Ferrers shapes. Let $G = (V, E, \prec)$ be an IM-free ordered graph with m L-vertices and n R-vertices. Let $l_1 \prec l_2 \prec \dots \prec l_m$ be the sequence of its L-vertices, ordered from left to right, and let $r_1 \succ r_2 \succ \dots \succ r_n$ be its R-vertices, ordered from right to left. Let D be a diagram with m columns and n rows, with the property that the i -th row and j -th column intersect in D if and only if the vertex r_i is to the right of the vertex l_j (see again Figure 2 for an example). Observe that the diagram D is a Ferrers diagram, and the shape of D uniquely determines the linear order \prec of the vertices of G . We now fill the diagram D with zeros and ones in such a way, that the cell in row i and column j is a 1-cell if and only if the vertex r_i is connected to the vertex l_j by an edge of G . The filling obtained by this procedure will be called *the adjacency filling* of the graph G .

It is not difficult to see that every dense 01-filling of a Ferrers shape is the adjacency filling of a unique IM-free ordered graph (up to isomorphism). We thus have a bijection between dense 01-fillings of Ferrers shapes and IM-free ordered graphs. Moreover, this bijection preserves the containment relations defined on the two classes of objects. Indeed, if G and H are two IM-free ordered graphs with adjacency fillings F_G and F_H , it follows easily from the definitions that G has a (not necessarily induced) subgraph isomorphic to H if and only if the filling F_G contains F_H . This fact provides a connection between pattern avoidance in fillings and pattern avoidance in graphs which we will often exploit in this thesis.

Let us now mention several special classes of fillings, together with their corresponding classes of graphs. First of all, notice that the transversals of Ferrers diagrams correspond exactly to ordered graphs with all degrees equal to one (of course, any such graph is IM-free). These graphs will be called *ordered matchings*. More specifically, the permutation matrices, considered as adjacency fillings, correspond precisely to ordered matchings in which every L-vertex is to the left of any R-vertex. Matchings with this property will be called *permutation matchings*. The permutation matrix that represents a permutation $\tau = \tau_1 \tau_2 \dots \tau_n$ is the adjacency filling of the ordered matching on the vertex set $[2n]$, where an L-vertex $j \in [n]$ is connected to the R-vertex $2n + 1 - \tau_j$.

Another useful class of ordered graphs are the so-called sprinkler graphs, introduced

(with a different terminology) by de Mier [20]. An ordered graph is called *sprinkler graph* if it is M-free and each of its R-vertices has degree one. Thus, every connected component of a sprinkler graph is a star, with the center of the star being the leftmost vertex of the component. A dense filling of a Ferrers diagram is an adjacency filling of a sprinkler graph if and only if it is a transpose of a semi-standard filling.

Apart from permutation matchings, there is another way to represent a permutation by an ordered graph. Let $\tau = \tau_1\tau_2 \cdots \tau_n$ be a permutation of order n . Let us define an ordered graph G on the vertex set $[n]$ by the following rule: for every $i, j \in [n]$, with $i < j$, the graph G contains the edge ij if and only if $\tau_i > \tau_j$. Those ordered graphs G that represent a permutation in this way are called *permutation graphs*. Each permutation graph represents a unique permutation. A permutation τ contains a permutation σ if and only if the permutation graph representing τ contains the graph representing σ as an induced subgraph.

Set partitions

A *set partition* of order n is a collection of nonempty disjoint sets B_1, B_2, \dots, B_k , called *blocks*, whose union is the set $[n]$. We always order the blocks in the increasing order of their minimal elements, i.e., we have $\min B_1 < \min B_2 < \dots < \min B_k$.

There are several ways to encode a set partition, and several corresponding notions of partition containment. Let us first mention an approach of Chen et al. [17, 18]. This approach is based on the notion of direct successor. For two numbers $i, j \in [n]$ and a partition $\Pi = (B_1, \dots, B_k)$ of order n , we say that j is a *direct successor* of i in Π , if $i < j$, i and j belong to the same block of Π , and no number that is larger than i and smaller than j belongs to the same block as i and j . Clearly every number that is not the largest element of its block has a unique direct successor. We may represent a partition Π of order n by an ordered graph on the vertex set $[n]$, in which two vertices are connected by an edge if and only if one of them is the direct successor of the other in Π . The ordered graph defined in this way is a vertex-disjoint union of monotone paths, where each path corresponds to a block of the original partition. We will call this graph *the path-representation of Π* . Note that every ordered matching is a path representation of a partition.

Another way to represent set partitions is to use sprinkler graphs, defined in the previous section. Again, a partition $\Pi = (B_1, \dots, B_k)$ of order n is represented by an ordered graph, but this time two vertices $i, j \in [n]$ are connected by an edge if i belongs to the same block as j and i is the smallest element of its block. Clearly, this yields a sprinkler graph, which we will call *the sprinkler representation of Π* . Every sprinkler graph represents a unique set partition.

Another encoding of set partitions was considered by Sagan [58] and later by Jelínek and Mansour [37]. This encoding is based on the concept of canonical sequence. Let $\Pi = (B_1, \dots, B_k)$ be a partition of $[n]$ with k blocks. We will represent Π by a k -ary word $\pi = \pi_1\pi_2 \cdots \pi_n \in [k]^n$, where $\pi_j = i$ if and only if $j \in B_i$. The sequence π will be called the *canonical sequence* of Π . Note that π has the following two properties:

- Every number $i \in [k]$ appears in π at least once (i.e., π is a reduced word).
- For every $i \in [k - 1]$, the first occurrence of i in π comes before the first occurrence of $i + 1$.

Every sequence that satisfies, for some value of k , the two properties above is a canonical

sequence of a unique set partition. The sequences of this form are also known as *restricted-growth functions*.

The three representations of set partitions described above suggest (at least) three different possibilities to define containment relation of set partitions. We may either view the containment relation of partitions as a special case of the subgraph (or induced subgraph) relation of ordered graphs, or alternatively, we may define partition containment as a special case of word containment, with a partition being represented by its canonical sequence. In this thesis, we will mostly be interested in the last option. Thus, we say that a partition Π *contains* a partition Σ if the canonical sequence of Π contains a subsequence order-isomorphic to the canonical sequence of Σ . Let us point out the containment order of partitions defined in this way generalizes the containment of permutations, in the following sense: if $\pi = \pi_1\pi_2 \cdots \pi_n \in \mathfrak{S}_n$ and $\sigma = \sigma_1\sigma_2 \cdots \sigma_k \in \mathfrak{S}_k$ are two permutations, then π contains σ if and only if the canonical sequence $123 \cdots n\pi_1\pi_2 \cdots \pi_n$ contains the canonical sequence $123 \cdots k\sigma_1\sigma_2 \cdots \sigma_k$.

Summary

In Figure 3, we summarize the main classes of ordered structures defined so far, and outline their relationships.

In the thesis, we will deal with several of these structures in greater detail, aiming to find common features in the pattern avoidance behaviour of these classes. We will be mostly interested in identities between the sizes of pattern-avoiding classes. To easily describe such identities, we will use the following terminology: we will say that two objects σ and τ are *equirestrictive* in a class of objects \mathcal{C} , if for every n , the number of σ -avoiding objects of size n in \mathcal{C} is the same as the number of τ -avoiding objects of size n in \mathcal{C} . Similarly, we will say that σ is *more restrictive than* τ (in a class \mathcal{C}) if, for every n , the number of σ -avoiding elements of size n in \mathcal{C} does not exceed the number its τ -avoiding elements.

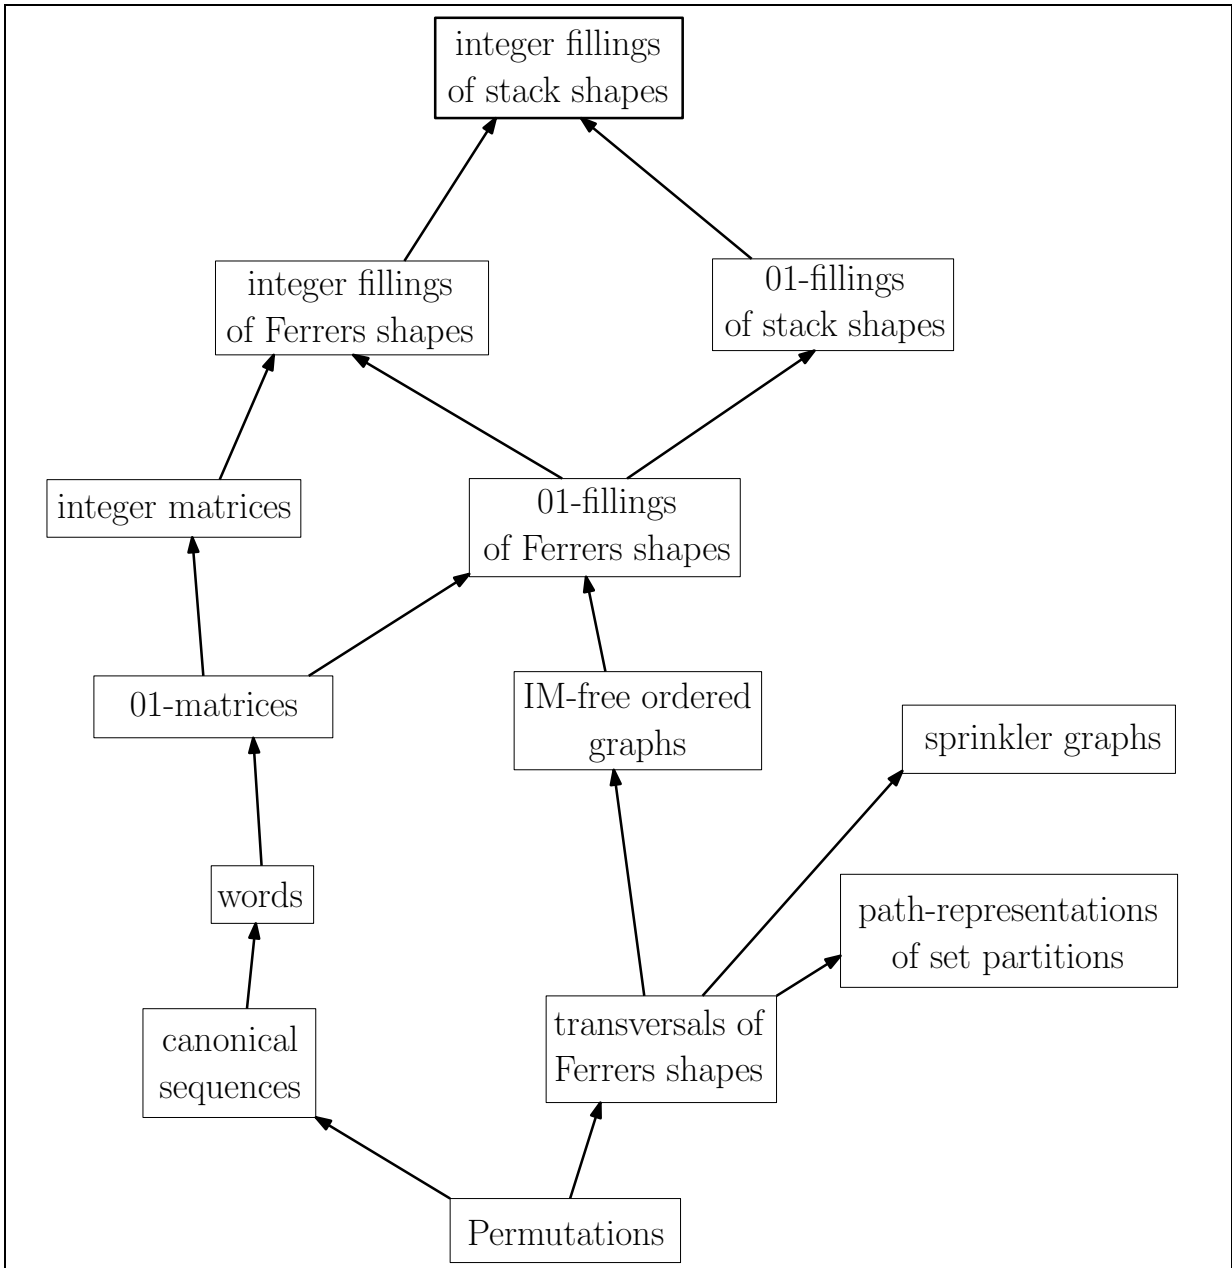


Figure 3: The ZOO of ordered structures. This figure presents an overview of the main classes of ordered structures considered in this thesis. Arrows indicate that the objects of the bottom class can be naturally represented by objects of the upper class, while preserving the containment relation.

Wilf-type classifications

As we announced in the introduction, this thesis is mostly devoted to the variations on the theme of Wilf equivalence.

Recall that $\mathcal{S}_n(\tau)$ denotes the set of permutations of order n that avoid τ . Two permutations σ and τ are called *Wilf equivalent*, denoted by $\sigma \stackrel{w}{\sim} \tau$, if for every n we have the equality $|\mathcal{S}_n(\sigma)| = |\mathcal{S}_n(\tau)|$. Clearly, two Wilf equivalent permutations have the same order. The equivalence classes of the Wilf equivalence are known as *Wilf classes*.

Naturally, the concept of Wilf equivalence can be easily extended to other ordered structures. Thus, every class of ordered structures with a corresponding containment relation gives rise to a Wilf-type equivalence relation.

A large part of the first three chapters is devoted to an overview of previous results related to Wilf equivalence. We usually present these previous results without proof, unless the method of the proof is necessary for the understanding of our own results presented in the rest of the thesis.

Chapter 1 contains an overview of known results related to the Wilf-classification of permutations. Apart from these results, mostly presented without proofs, we also state and prove several lemmas that were proven in the context of Wilf-equivalence, but whose ideas can easily be generalized to other ordered structures.

In the second chapter, we will investigate in greater detail the topic of diagonal patterns in fillings of diagrams. We will be particularly interested in the theorems of Backelin, West and Xin [6], of Krattenthaler [48], and of Rubey [57]. These theorems will play a significant part in the remaining chapters, since they have important consequences in the study of ordered structures. Part of the second chapter is also devoted to the author's own result related to diagonal fillings, which does not seem to have as far-reaching consequences as the above-mentioned theorems, but it deals with similar topic.

In Chapter 3, we will mention the topic *Wilf order*, which is a quasi-order relation \preceq defined on the set of permutations by writing $\sigma \preceq \tau$ if and only if $|\mathcal{S}_n(\sigma)| \leq |\mathcal{S}_n(\tau)|$ for every n . The first part of this chapter is again devoted to the overview of previous results, and in the second part, we explore a connection between Wilf order and fillings of skew shapes, which yields a new family of Wilf-comparable permutations.

In the remaining chapters of the main part, we will study the Wilf-type classification of involutions, words and set partitions. Most of the results presented in these chapters have been published in a series of papers [22, 37, 38] as the joint work of the author with Dukes, Mansour and Reifegerste.

We conclude this thesis by an overview of several promising alternative directions of research related to the topic of hereditary permutation classes.

Chapter 1

Wilf classes of permutations

1.1 Symmetries of permutations

Before we deal with the main results on Wilf equivalence, let us introduce some more terminology. Let $\pi = \pi_1\pi_2 \cdots \pi_n \in \mathfrak{S}_n$ be a permutation. The *reversal* of π is a permutation $\sigma = \sigma_1\sigma_2 \cdots \sigma_n$ obtained by writing π backwards, i.e., $\sigma_i = \pi_{n-i+1}$. The *inverse* of a permutation π is a permutation $\rho = \rho_1\rho_2 \cdots \rho_n$ defined by the equivalence $\rho_i = j \iff \pi_j = i$, for every $i, j \in [n]$. We will denote the reversal of π by $\bar{\pi}$ and the inverse by π^{-1} . The two operations can be easily visualised, when we represent the permutations by their permutation matrices. The permutation matrix representing $\bar{\pi}$ is obtained from the matrix representing π by reversing the order of columns, while the matrix representing π^{-1} is the transpose of the matrix representing π .

If a permutation σ can be obtained from a permutation π by a sequence of reversals and inverses, we say that σ is *symmetric to* π . The *symmetry class* of π is the set of all the permutations that are symmetric to π . The symmetry class may have up to eight elements.

Permutation containment is preserved by both the reversal and the inverse, in the following sense: a permutation π contains a permutation σ if and only if $\bar{\pi}$ contains $\bar{\sigma}$, which happens if and only if π^{-1} contains σ^{-1} . It is thus clear that each permutation is Wilf equivalent to its reversal and to its inverse, and in particular, every symmetry class is a subset of the corresponding Wilf class.

1.2 Non-trivial Wilf equivalences

Results related to Wilf equivalence can be traced back to 1973, when Knuth [46, 47] showed that for any permutation τ of order three, the cardinality of $\mathfrak{S}_n(\tau)$ is equal to the n -th Catalan number $C_n = \frac{1}{n+1} \binom{2n}{n}$ (sequence A000108 in OEIS [68]). In particular, all the permutations of order three are Wilf equivalent. Note that the permutations of order three fall into two symmetry classes ($\{123, 321\}$ and $\{132, 213, 231, 312\}$), so this example demonstrates that the symmetry classes are a strict refinement of Wilf classes.

Let us mention that the Wilf classification of patterns of size three was completed by Simion and Schmidt [62], who determined the cardinality of $\mathfrak{S}_n(\mathcal{F})$ for any set $\mathcal{F} \subseteq \mathfrak{S}_3$.

To determine the Wilf classes of patterns of size four took a lot more effort. The 24 permutations of order four fall into seven symmetry classes, for which we may choose the representatives 1234, 1243, 1324, 1342, 1432, 2143 and 2413 (see Fig. 1.1).

It has been determined that these seven patterns fall into the following three Wilf classes:




Figure 1.1: The seven pairwise non-symmetric permutations of order four, represented as matrices. For clarity, we omit the 0-cells.

- 1324;
- $1342 \stackrel{w}{\sim} 2413$;
- $1234 \stackrel{w}{\sim} 1243 \stackrel{w}{\sim} 1432 \stackrel{w}{\sim} 2143$.

The first step towards this classification was made in 1994 by Stankova [63], who showed that 1342 is Wilf equivalent with 2413. This seems to be a sporadic case of equivalence—so far, no one has managed to interpret this result as a special case of a more general identity. In contrast, the Wilf equivalence of the four patterns 1234, 1243, 1432, and 2143 follows from a more general result, which will be presented in the next section.

1.3 The shape-Wilf equivalence

Apart from the sporadic pair $1342 \stackrel{w}{\sim} 2413$, all known pairs of non-symmetric Wilf equivalent permutations are described by two general results by Stankova and West [65], and by Backelin, West and Xin [6]. Both these results are based on the concept of shape-Wilf equivalence, which is an analogue of Wilf equivalence for transversals of Ferrers diagrams. Recall that a 01-filling of a diagram is called transversal, if every row and every column of the filling has exactly one 1-cell. Throughout the rest of this chapter, we assume that every filling we mention is a filling of a Ferrers diagram, unless otherwise noted.

For a Ferrers diagram F , let \mathcal{T}_F denote the set of all the transversals of the shape F . Let $\mathcal{T}_F(\sigma)$ denote the set of the transversals of F that avoid the pattern σ . We say that two transversals σ and τ are *shape-Wilf equivalent*, denoted by $\sigma \stackrel{LW}{\sim} \tau$, if for every Ferrers diagram F , the set $\mathcal{T}_F(\sigma)$ has the same cardinality as $\mathcal{T}_F(\tau)$.

A permutation, represented by its permutation matrix, is a transversal of a square shape. With a slight abuse of terminology, we will omit the distinction between a permutation and its permutation matrix, and we will say that two permutations are shape-Wilf equivalent if their permutation matrices, treated as transversal fillings of a square shape, are shape-Wilf equivalent. If there is no risk of confusion, we will freely switch between the two possible representations of a permutation.

Observe that if two permutations are shape-Wilf equivalent, then they are also Wilf equivalent. To see this, let F denote the square diagram with n rows and n columns, and assume that σ and τ are shape-Wilf equivalent permutations. The shape-Wilf equivalence implies the equality $|\mathcal{T}_F(\sigma)| = |\mathcal{T}_F(\tau)|$. However, the σ -avoiding transversals of F are precisely the permutation matrices representing the σ -avoiding permutations of order n . Thus, the equality $|\mathcal{T}_F(\sigma)| = |\mathcal{T}_F(\tau)|$ implies $|\mathcal{S}_n(\sigma)| = |\mathcal{S}_n(\tau)|$. Since this argument works for any $n \in \mathbb{N}$, we see that σ and τ are indeed Wilf equivalent.

Unlike Wilf equivalence, the shape-Wilf equivalence is not necessarily preserved by reversal or by inverse. In fact, among the three symmetric permutations 132, 312, and 213, no two are shape-Wilf equivalent.

Let us adopt the following intuitive notation: if $A \in \mathbb{N}^{n \times n}$ and $B \in \mathbb{N}^{m \times m}$ are two square matrices, we let $\begin{pmatrix} 0 & B \\ A & 0 \end{pmatrix}$ denote the matrix with $m+n$ rows and $m+n$ columns,

whose bottom-left corner contains a copy of A , its top-right corner contains a copy of B , and the remaining cells are equal to zero. We are now ready to state the following proposition, which is due to Backelin et al. [6]. Later on, slightly modified versions of this proposition were applied in more general settings [19, 37].

Proposition 1 (Proposition 2.3 from Backelin et al. [6]). *Let A and B be two shape-Wilf equivalent permutations of order n , and let C be an arbitrary permutation of order m . Then the permutations $\begin{pmatrix} 0 & C \\ A & 0 \end{pmatrix}$ and $\begin{pmatrix} 0 & C \\ B & 0 \end{pmatrix}$ are shape-Wilf equivalent (and hence also Wilf equivalent).*

Before we present the proof of this proposition, let us make several remarks related to its statement. First of all, the assumption that A and B are shape-Wilf equivalent is essential, and it is not enough to just assume that the two permutations are Wilf equivalent. For example, consider $A = 132$, $B = 123$, and $C = 1$: we know that $A \stackrel{w}{\sim} B$, but the two permutations $\begin{pmatrix} 0 & C \\ A & 0 \end{pmatrix} = 1324$ and $\begin{pmatrix} 0 & C \\ B & 0 \end{pmatrix} = 1234$ are not even Wilf equivalent. Let us also point out, that the two matrices $\begin{pmatrix} 0 & C \\ A & 0 \end{pmatrix}$ and $\begin{pmatrix} 0 & C \\ B & 0 \end{pmatrix}$ in the conclusion of the proposition cannot be replaced by $\begin{pmatrix} 0 & A \\ C & 0 \end{pmatrix}$ and $\begin{pmatrix} 0 & B \\ C & 0 \end{pmatrix}$. For example, consider $A = 213$, $B = 123$, and $C = 1$. Then A and B are shape-Wilf equivalent (as we will soon see), but $\begin{pmatrix} 0 & A \\ C & 0 \end{pmatrix} = 1324$ and $\begin{pmatrix} 0 & B \\ C & 0 \end{pmatrix} = 1234$ are again not even Wilf equivalent.

Before proving Proposition 1, we first prove a simple lemma, which shows that a bijection between pattern-avoiding transversals may be extended into a bijection between pattern-avoiding sparse fillings. Recall that a 01-filling is sparse, if every row and every column contains at most one 1-cell.

Lemma 2. *Let A and B be shape-Wilf equivalent permutations. For any Ferrers shape F , there is a bijection ϕ between the set of sparse A -avoiding 01-fillings of F and the set of sparse B -avoiding 01-fillings of F . Moreover, ϕ preserves the position of zero rows and zero columns.*

Proof. Since $A \stackrel{w}{\sim} B$, there is an invertible mapping ϕ_0 which transforms A -avoiding transversals into B -avoiding transversals of the same shape. The basic idea of the proof is straightforward: we simply extend ϕ_0 to a bijection which operates on sparse fillings, by ignoring the zero rows and columns, and applying ϕ_0 on the nonzero rows and columns.

Let us now explain the argument more formally (see Figure 1.2). Let Φ be an A -avoiding sparse filling of a Ferrers shape F . By removing from Φ all the zero rows and zero columns, we obtain an A -avoiding transversal Φ^- of a Ferrers subdiagram F^- of F . Let $\Psi^- = \phi_0(\Phi^-)$. By inserting into Ψ^- the zero rows and zero columns whose positions correspond to the position of the zero rows and columns of Φ , we extend Ψ^- into a sparse B -avoiding filling Ψ of the shape F . It is easy to see that the transform $\phi: \Phi \mapsto \Psi$ has all the required properties. Note that the insertion of zero row or column cannot create an occurrence of a forbidden pattern, because the patterns A and B themselves have no zero row or column. \square




Figure 1.2: Illustration of Lemma 2: a shape-preserving bijection between transversals can be extended into a shape-preserving bijection between sparse fillings.

Let us now focus on the proof of Proposition 1. The main trick involved in the proof is not too difficult, but it is extremely useful and will reappear in this thesis in several different contexts. For this reason, we will now present the proof in full detail. We will refer to the trick as *the red-green argument*.

Proof of Proposition 1. Choose A , B and C as in the statement of Proposition 1. By assumption, $A \stackrel{\text{lw}}{\sim} B$. In particular, there is a bijection ϕ that maps A -avoiding transversals to B -avoiding transversals, while preserving the underlying shape. Let us fix an arbitrary Ferrers shape F . Our aim is to describe a bijection that maps the $\begin{pmatrix} 0 & C \\ A & 0 \end{pmatrix}$ -avoiding transversals of F to the $\begin{pmatrix} 0 & C \\ B & 0 \end{pmatrix}$ -avoiding transversals of F .

Let (i, j) denote the cell of F in row i and column j . We will say that a cell (i', j') of F is *top-right* of a cell (i, j) if $i > i'$ and $j > j'$. Note that for a fixed cell (i, j) , the cells that are top-right from (i, j) form a Ferrers subshape of F (see Fig. 1.3).




Figure 1.3: The cells that are north-east from the cell (i, j) are shaded. They form a Ferrers subdiagram.

Let Φ be an arbitrary $\begin{pmatrix} 0 & C \\ A & 0 \end{pmatrix}$ -avoiding transversal of F . We will distinguish in Φ two types of cells: a cell (i, j) is *green* (with respect to the filling Φ) if the subdiagram formed by the cells that are top-right from (i, j) contains the pattern C ; otherwise, the cell (i, j) is *red*. Let Φ_G and Φ_R denote the subfillings of Φ formed by the green cells and the red cells of Φ , respectively. Observe that Φ_G is a sparse filling of a Ferrers diagram, while Φ_R is a sparse filling of a skew shape.

Let us now make several observations about the properties of this two-coloring. First of all, we claim that if (i, j) is a green cell, then the red filling Φ_R contains a copy of C which is situated to the top-right of (i, j) . To see this, choose a copy C' of the matrix C in Φ , such that C' is situated to the top-right of (i, j) , and the bottom row of C' is as far to the top as possible. We know that such a submatrix C' exists, otherwise (i, j) would not be green. We claim that all the cells of C' are red. For this, it is sufficient to show that the bottom-left corner of C' is red. However, if the bottom-left corner of C' were green, then Φ would contain a copy of C to the top-right of this corner, contradicting the choice of C' . This shows that a cell of Φ is green if and only if it has a red copy of C situated to its top-right.

Next, we observe that Φ_G avoids the matrix A . Indeed, if Φ_G contained A , then Φ would contain $\begin{pmatrix} 0 & C \\ A & 0 \end{pmatrix}$, which would contradict our assumptions. By Lemma 2, there is a shape-preserving bijection ϕ that transforms A -avoiding sparse fillings into B -avoiding sparse fillings, while preserving the zero rows and columns. Let us replace the filling Φ_G by the filling $\phi(\Phi_G)$ inside the diagram F , while the filling Φ_R remains unchanged. Let Ψ be the resulting filling of F . Let Ψ_G be the subfilling of Ψ obtained as the image of Φ_G by the bijection ϕ .

Let τ be the transform $\Phi \mapsto \Psi$ described above. We claim that τ is a bijection between $(\begin{smallmatrix} 0 & C \\ A & 0 \end{smallmatrix})$ -avoiding and $(\begin{smallmatrix} 0 & C \\ B & 0 \end{smallmatrix})$ -avoiding fillings of the diagram F . To see this, we need to check that Ψ is a $(\begin{smallmatrix} 0 & C \\ B & 0 \end{smallmatrix})$ -avoiding transversal of F , and that the mapping $\Phi \mapsto \Psi$ can be inverted. To see that Ψ is a transversal, it suffices to recall that ϕ preserves the position of the zero rows and zero columns in Φ_G . The zero rows (or columns) of Φ_G are exactly the rows (or columns) that contain a red 1-cell in Φ_R .

Let us argue that Ψ avoids $(\begin{smallmatrix} 0 & C \\ B & 0 \end{smallmatrix})$. For contradiction, assume that Ψ contains a copy B' of the matrix B , as well as a copy C' of the matrix C , and that C' is situated top-right from B' . We know that Ψ_G avoids B . Thus, at least one cell of B' must be red. This implies that all the cells of C' are red. This in turn means that all the cells that are bottom-left from C' must be green, which is impossible, because B' has at least one red cell.

It remains to show that τ is invertible. Since we already know that the transformation ϕ of the green cells is invertible, it suffices to notice that the cells of the filling Ψ_G are exactly the cells that are green with respect to the filling Ψ . Earlier, we have pointed out that a cell (i, j) is green with respect to Φ if and only if there is a copy C' of the matrix C which is situated to the top-right of (i, j) and consists entirely of cells that are red with respect to Φ . Since this red copy of C is not affected by the transform τ , we know that such a cell (i, j) is also green with respect to Ψ . By the same argument, we see that a cell that is red with respect to Φ is also red with respect to Ψ . It is now clear that the mapping τ is invertible. \square

Proposition 1 allows us to construct infinite families of Wilf equivalent permutations from a single pair of shape-Wilf equivalent patterns. At first glance, it is not even clear whether there are any distinct shape-Wilf equivalent patterns at all. So far, there are two results involving shape-Wilf equivalence of permutation patterns. The first is the following theorem of Stankova and West, which yields a single pair of shape-Wilf equivalent permutations.

Theorem 3 (Stankova–West [65]). *The two permutations 312 and 231 are shape-Wilf equivalent.*

There are nowadays two known proofs of Theorem 3. The original argument of Stankova and West was based on an inductive construction of 312-avoiding transversals. Another argument, due to Jelínek [34], shows that for every shape F there is a bijection between the set $\mathcal{T}_F(312)$ and the set $\mathcal{T}_F(X)$, where X is an infinite set of symmetric transversals of Ferrers shapes. Since the permutation matrix representing 231 is the transpose of the matrix of 312, and all the elements of X are symmetric, it follows that 312 and 231 are shape-Wilf equivalent. Both known proofs of Theorem 3 are rather lengthy, and we omit them for the sake of brevity.

Unfortunately, Theorem 3 has no known generalization. Not only are we not able to extend the theorem to provide more pairs of shape-Wilf equivalent patterns, but we are also not aware of any result that would extend the equivalence of 312 and 231 to a more general type of fillings other than transversals of Ferrers shapes.

The other known result on shape-Wilf equivalence is the following Theorem due to Backelin, West and Xin.

Theorem 4 (Backelin–West–Xin [6]). *For any k , the identity permutation $12 \cdots k$ is shape-Wilf equivalent to the anti-identity permutation $k(k-1) \cdots 1$.*

Unlike Theorem 3, Theorem 4 has several interesting generalizations which concern more general fillings as well as more general diagrams. We will deal with these generalizations in the next chapter.

It has been verified by computer enumeration [6] that all the Wilf classes of permutations of order at most seven can be described by the results we mentioned in this chapter. For larger patterns, computer enumeration quickly becomes infeasible, and a full Wilf classification seems out of reach. Likewise, apart from the work of Simion and Schmidt [62], who have enumerated permutations avoiding an arbitrary set of patterns of size 3, there is only limited understanding of permutations simultaneously avoiding multiple patterns.

Chapter 2

Identities involving diagonal patterns

In this chapter, we will review several recent results that involve identities between fillings of diagrams that avoid increasing chains of k positive elements and those that avoid decreasing chains of k positive elements. Results of this type seem to be a recurrent topic in the study of pattern avoidance. They seem to point towards a more general combinatorial phenomenon, which is not yet fully understood.

The main reason we devote a special chapter to this type of results is that they seem to be a natural generalization of various identities obtained in the study of other combinatorial structures, such as graphs [19, 20], words [37], or set partitions [17, 18, 37]. Indeed, many of the results mentioned in the later chapters of this thesis are based on the identities between diagonal-avoiding fillings mentioned in this chapter.

In Section 2.1 of this chapter, we will present the most important previous results related to diagonal-avoiding fillings. The theorems presented here do not convey the full strength of the results obtained in this field of study. For instance, we do not deal with simultaneous avoidance of increasing and decreasing chains in diagrams, since this would require introducing new terminology that would not be useful in the rest of the thesis. The interested reader may consult the papers of Krattenthaler [48], de Mier [19, 20] or Rubey [57], from which these results originate.

In Section 2.2, the author will offer his own contribution to the topic of this chapter, by proving an identity involving diagonal-free fillings of rectangular shapes.

2.1 Known results on diagonal patterns

Let I_k denote the identity matrix of order k , i.e., the matrix representing the permutation $12 \cdots k$. Let J_k be the anti-identity matrix, i.e., the matrix representing $k(k-1) \cdots 1$. As we already stated in Theorem 4, the two matrices I_k and J_k are shape-Wilf equivalent. We also mentioned that this equivalence can be generalized to more general fillings and more general shapes. We will now present an overview of these generalizations.

The original proof Theorem 4, due to Backelin et al., was first published in 2001. The proof was based on an elementary argument, providing an explicit bijection between I_k -avoiding and J_k -avoiding transversals. In 2004, Bousquet-Mélou and Steingrímsson [14] have shown that the bijection of Backelin et al. maps symmetric fillings to symmetric fillings, thus obtaining the following result (recall that a filling is symmetric, if it is equal to its transpose).

Theorem 5 (Bousquet-Mélou–Steingrímsson [14]). *For any $k \in \mathbb{N}$ and any symmetric Ferrers shape F , the two patterns I_k and J_k are equirestrictive in the class of the symmetric transversals of F . In particular, the two patterns are equirestrictive in the class of involutions.*

Even before Theorem 5 was proved, Jaggard [32] has applied a modification of the red-green argument of Proposition 1, to show that the theorem would imply the following corollary.

Corollary 6. *Let k and n be natural numbers, let A be any permutation matrix. The number of involutions of order n that avoid $\begin{pmatrix} 0 & A \\ I_k & 0 \end{pmatrix}$ is the same as the number of those that avoid $\begin{pmatrix} 0 & A \\ J_k & 0 \end{pmatrix}$.*

In 2006, Krattenthaler [48] has shown a different way to obtain bijections between I_k -avoiding and J_k -avoiding fillings of a given shape. His approach is based on the theory of growth diagrams, and seems even more powerful than the original approach of Backelin, West and Xin. Here is a simplified version of one of Krattenthaler's main results.

Theorem 7 (Krattenthaler [48]). *Let k be an integer, let F be a Ferrers shape. Let \mathcal{F} be the set of all the integer fillings of F . There is a bijection $\kappa: \mathcal{F} \rightarrow \mathcal{F}$ with the following properties:*

- *A filling $\Phi \in \mathcal{F}$ contains I_k if and only if $\kappa(\Phi)$ contains J_k .*
- *For any i , the sum of the elements of the i -th row of Φ is equal to the sum of the elements of the i -th row of $\kappa(\Phi)$. Similarly, the sum of the elements of the i -th column of Φ is equal to the sum of the elements of the i -th column of $\kappa(\Phi)$.*

Note that unlike the previously mentioned results, Theorem 7 speaks of fillings by arbitrary integers, rather than 01-fillings. There can be no bijection satisfying the conditions of Theorem 7 that would map 01-fillings to 01-fillings.

Again, as was pointed out by de Mier [19], we may apply a modification of the red-green argument to obtain the following result.

Corollary 8. *Let k be an integer, let A be a matrix, and let F be a Ferrers shape. The two matrices $\begin{pmatrix} 0 & A \\ I_k & 0 \end{pmatrix}$ and $\begin{pmatrix} 0 & A \\ J_k & 0 \end{pmatrix}$ are equirestrictive among the nonnegative fillings of F . Furthermore, the bijection between the two pattern-avoiding classes preserves the sum of the entries in each row and each column of the filling.*

Rubey [57] has managed to generalize these results to fillings of moon diagrams. The *content* of a diagram is the multiset of the lengths of its columns. Here are the results obtained by Rubey:

Theorem 9 (Rubey [57]). *Let k be an integer. Let M and M' be two moon polyominoes with the same content. There is a bijection between I_k -avoiding 01-fillings of M and I_k -avoiding 01-fillings of M' . Furthermore, if M' is obtained from M by a permutation of its columns (i.e., without altering the vertical position of the columns) then the bijection preserves the number of 1-cells in each row.*

Notice that Theorem 9 implies that I_k and J_k are equirestrictive in the set of 01-fillings of a moon polyomino. This is because the J_k -avoiding fillings of a moon polyomino M correspond to the I_k -avoiding fillings of M' , where M' is the mirror image of M . By Theorem 9, the I_k -avoiding fillings of M' are then in bijection with the I_k -avoiding fillings of M .

For fillings by arbitrary nonnegative integers, Rubey obtains an analogous result:

Theorem 10 (Rubey [57]). *Let k be an integer. Let M and M' be two moon polyominoes with the same content. There is a bijection between I_k -avoiding nonnegative fillings of M and I_k -avoiding nonnegative fillings of M' . Furthermore, if M' is obtained from M by a permutation of its columns, then the bijection preserves the sum of the entries in each row.*

2.2 Constrained rectangular fillings

In this section, we take the opportunity to present the author's own contribution to the rich family of identities involving I_k - and J_k -avoiding fillings. The contents of this section are based on our contribution presented at FPSAC 2007 [35].

We will consider pattern avoidance in rectangular tables with prescribed row- and column-sums. Let us start with basic definitions. A *constrained table of shape* $r \times s$ is an empty table with r rows and s columns, together with two sequences of nonnegative integers: the *row constraints* (x_1, \dots, x_r) , and the *column constraints* (y_1, \dots, y_s) , satisfying

$$\sum_{i=1}^r x_i = \sum_{j=1}^s y_j.$$

A *filling* of the constrained table is a nonnegative integer matrix $M = (M_{ij})$ with r rows and s columns, such that the sum of the entries in the i -th row is equal to x_i , and the sum of the entries in the j -th column is equal to y_j , formally:

$$\begin{aligned} \forall i \in [r]: \quad & \sum_{j=1}^s M_{ij} = x_i \\ \forall j \in [s]: \quad & \sum_{i=1}^r M_{ij} = y_j. \end{aligned}$$

For two sequences $x = (x_1, \dots, x_r)$ and $y = (y_1, \dots, y_s)$ of nonnegative integers, we let $T[x \times y]$ denote the constrained table with row-constraints x and column-constraints y , and we let $f(x \times y)$ denote the total number of fillings of $T[x \times y]$. The unordered multiset $\{x_1, x_2, \dots, x_r, y_1, y_2, \dots, y_s\}$ will be called *the scoreline* of the table $T[x \times y]$. For a sequence $x = (x_1, \dots, x_r)$ and a permutation $\pi \in S_r$ we write $\pi(x)$ for the sequence $(x_{\pi(1)}, x_{\pi(2)}, \dots, x_{\pi(r)})$.

Note that if T is a table of shape $r \times r$ with all the row- and column constraints equal to 1, then the fillings of T are exactly the permutation matrices of order r . Furthermore, if P is itself a permutation matrix of a permutation $\pi \in S_n$, then the P -avoiding fillings of T are precisely the permutation matrices corresponding to the π -avoiding permutations. Thus, the concept of pattern avoidance in rectangular fillings is a generalization of pattern avoidance in permutations.

Notice that for any permutation π of appropriate order, we have the identity $f(x \times y) = f(\pi(x) \times y)$. This is because every filling M of $T[x \times y]$ can be transformed into a filling of $T[\pi(x) \times y]$ by permuting the rows of M according to the permutation π . Of course, this simple bijection in general does not preserve pattern avoidance. However, if P is a permutation matrix of order at most three, not only do we have the identity $f(x \times y; P) = f(\pi(x) \times \rho(y); P)$ for any π and ρ , but in fact, we can prove a stronger identity, stated in the following theorem, which is the main result of this section.

Theorem 11 (J. [35]). *Let $T[x \times y]$ be a constrained table, let P be a permutation matrix of order at most three. Then $f(x \times y; P)$ is uniquely determined by the scoreline of $T[x \times y]$ and the order of P .*

For example, consider the two tables $T = T[(2, 2) \times (1, 1, 1, 1)]$ and $T' = T[(2, 1, 1) \times (2, 1, 1)]$. Both these tables have the same scoreline $\{2, 2, 1, 1, 1, 1\}$. Theorem 11 implies that they must have the same number of P -avoiding fillings for any permutation matrix P of order at most three. Indeed, if P has order two, then both tables admit exactly one P -avoiding filling, and if P has order three, then all the six fillings of T are P -avoiding,




Figure 2.1: The corner flip operation.

and T' also has six P -avoiding fillings (as well as one filling containing P). This example also shows that Theorem 11 cannot be extended to permutation patterns of order greater than three, since all the six possible fillings of T as well as all the seven possible fillings of T' clearly avoid any pattern larger than three.

Before we present the proof of Theorem 11, we collect some simple observations that deal with permutation patterns of order at most two. We then show that previous results on fillings of Ferrers shapes imply that $f(x \times y; P) = f(x \times y; Q)$ for any two permutation matrices P, Q of order three. These arguments show that it is sufficient to prove Theorem 11 for the case when $P = J_3$.

Next, we will use the RSK algorithm together with basic results on Young tableaux to prove that if $P = J_k$ for some k , then $f(x \times y; P) = f(\pi(y) \times \rho(y); P)$, where π and ρ are arbitrary permutations of appropriate order.

As the last step of the proof, we introduce an operation called *corner flip*, defined as follows: let $T[x \times y]$ be a constrained table of shape $r \times s$. Assume that for some $t \leq r$ and $u \leq s$ we have

$$\sum_{i=1}^t x_i = \sum_{j=1}^u y_j.$$

A corner flip is an operation that transforms the table $T[x \times y]$ into a table $T[x' \times y']$ of shape $(t + s - u) \times (u + r - t)$, where $x' = (x_1, x_2, \dots, x_t, y_{u+1}, y_{u+2}, \dots, y_s)$ and $y' = (y_1, y_2, \dots, y_u, x_{t+1}, x_{t+2}, \dots, x_r)$ (see Fig. 2.1).

We will show that corner flips preserve the number of J_3 -avoiding fillings. It is easy to see that any two tables with the same scoreline can be transformed to each other by a sequence of row permutations, column permutations and corner flips. Combining these facts, we obtain the proof of Theorem 11.

After we prove Theorem 11, we will present some remarks on the connection between the fillings of rectangular shapes, the pattern avoidance in permutations, and other related concepts.

Let us first deal with the values of $f(x \times y; P)$ when P is a permutation matrix of order at most two. In the trivial case when P has order one, we see that $f(x \times y; P) = 0$ unless all the components of x and y are zero, in which case $f(x \times y; P) = 1$. Let us now turn to the slightly less trivial case of permutation matrices of order two:

Lemma 12. *If $P = I_2$ or $P = J_2$, and if $T = T[x \times y]$ is any constrained table, then $f(x \times y; P) = 1$.*

Proof. It suffices to prove the lemma for $P = J_2$, the other case is analogous. Let $x =$

(x_1, \dots, x_r) and $y = (y_1, \dots, y_s)$. We proceed by induction on $r + s$. If $r = 1$ or $s = 1$, the claim is clear.

Otherwise, let $k = \min\{x_r, y_s\}$. Observe that if $M = (M_{ij})$ is a J_2 -avoiding filling of $T[x \times y]$, then $M_{rs} = k$, otherwise both the last row and the last column of M would contain a positive entry other than M_{rs} , and these two entries would form the forbidden pattern P . Assume now that $k = x_r$ (the case $k = y_s$ is symmetric). For any J_2 -avoiding filling M of T , the last row of M is equal to $(0, \dots, 0, k)$. Furthermore, the remaining rows of M form a J_2 -avoiding filling of $T' = T[(x_1, \dots, x_{r-1}) \times (y_1, \dots, y_{s-1}, y_s - k)]$. By the induction hypothesis, there is exactly one J_2 -avoiding filling of T' , and adding a row $(0, \dots, 0, k)$ to the top of this filling produces a J_2 -avoiding filling of T . \square

It remains to prove Theorem 11 for permutation patterns of order three. Using Corollary 8, we may easily conclude that $f(x \times y; P) = f(x \times y; Q)$ for any two patterns P, Q of order three.

Lemma 13. *For any constrained rectangular table $T[x \times y]$ and any two permutation matrices P, Q of order three, we have $f(x \times y; P) = f(x \times y; Q)$.*

Proof. For P, Q chosen among I_3, J_3 and $\begin{pmatrix} 0 & I_1 \\ J_2 & 0 \end{pmatrix}$, the claim is a special case of Corollary 8. For the other cases, we can easily establish the required identity by exploiting the symmetries of the rectangle; take, e.g., $P = \begin{pmatrix} I_2 & 0 \\ 0 & I_1 \end{pmatrix}$ and $Q = \begin{pmatrix} 0 & I_1 \\ J_2 & 0 \end{pmatrix}$: let us write $x = (x_1, \dots, x_r)$ and let us define $\bar{x} = (x_r, x_{r-1}, \dots, x_1)$. Clearly,

$$f(x \times y; P) = f(\bar{x} \times y; Q) = f(\bar{x} \times y; I_3) = f(x \times y; J_3) = f(x \times y; Q).$$

The remaining cases are settled similarly. \square

As the next step towards the proof of Theorem 11, we prove the following result:

Proposition 14. *Let $T[x \times y]$ be a constrained table of shape $r \times s$. For every $\pi \in S_r$ and $\rho \in S_s$, and for every positive integer n , we have $f(x \times y; J_n) = f(\pi(x) \times \rho(y); J_n)$.*

Of course, in the statement of the proposition, we could have used any other pattern of the form $\begin{pmatrix} 0 & I_{n-k} \\ J_k & 0 \end{pmatrix}$ instead of J_n . Our choice of J_n is purely a matter of convenience.

Proposition 14 is an easy consequence of known results on Young tableaux and the Robinson–Schensted–Knuth (or RSK) algorithm. We will now state the necessary results without proof; a useful presentation of several variants of the RSK algorithm and their relation to pattern avoidance in fillings can be found in Krattenthaler’s paper [48]. The proofs of the basic properties of Young tableaux and the RSK correspondence can be found in textbooks of combinatorics, such as [27] or [66].

We first state the necessary definitions:

A *partition* (also known as *integer partition*, not to be confused with set partitions defined earlier) of size n and length r is a nonincreasing sequence $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_r$ of r positive integers whose sum is n .

A *Young tableau*, or more verbosely, a *column-strict semi-standard Young tableau*, is a filling of a Ferrers shape such that the elements of every row form a weakly increasing sequence and the elements of every column form a strictly increasing sequence. If P is a Young tableau with r rows, and the i -th row of P has length λ_i , then the sequence $\lambda = (\lambda_1, \dots, \lambda_r)$ is a partition, which we will simply call *the shape* of P .

The *content* of a Young tableau P is a sequence $\mu = (\mu_1, \dots, \mu_k)$ where μ_i is the number of cells of P that contain the number i . The number of Young tableaux of shape λ and content μ is known as the Kostka number, denoted $K_{\lambda, \mu}$.

The proof of the following standard fact can be found e.g. in [27]:

Fact 15. Let λ be a partition of n and let $\mu = (\mu_1, \dots, \mu_k)$ be a sequence of nonnegative numbers whose sum is n , let π be a permutation of order k . For $\mu' = \pi(\mu) = (\mu_{\pi(1)}, \dots, \mu_{\pi(k)})$, we have the identity $K_{\lambda\mu} = K_{\lambda\mu'}$.

Let g_π be a bijection that transforms a Young tableau P of content μ to a Young tableau $g_\pi(P)$ of the same shape and of content $\pi(\mu)$.

We now summarize the properties of the RSK algorithm which we will use in our proof:

Fact 16. The RSK algorithm provides a bijection between fillings of $T[x \times y]$ and ordered pairs of Young tableaux (P, Q) such that P and Q have the same shape, P has content x and Q has content y . Furthermore, the filling avoids J_n if and only if P and Q have less than n rows.

These facts immediately imply Proposition 14:

Proof of Proposition 14. Let x, y, π, ρ be as in Proposition 14. The J_n -avoiding fillings of $T[x \times y]$ are mapped by the RSK algorithm to pairs of Young tableaux (P, Q) of the same shape λ with at most $n - 1$ rows, where the content of P is x and the content of Q is y . This pair may be transformed into a pair of tableaux $(g_\pi(P), g_\rho(Q))$ of shape λ and content $\pi(x)$ and $\rho(y)$. By the RSK algorithm, such pairs correspond to J_n -avoiding fillings of $T[\pi(x) \times \rho(y)]$. \square

We remark that the bijection established above does not, in general, preserve the multiset of the entries used in the corresponding fillings. In particular, it does not send 01-fillings onto 01-fillings. This cannot be avoided because, for example, $T[(2, 1, 1) \times (2, 1, 1)]$ has no J_2 -avoiding 01-filling, while $T[(1, 2, 1) \times (2, 1, 1)]$ has one such filling.

The last ingredient of our proof is the operation called corner flip, illustrated on Fig. 2.1. Let us fix $x = (x_1, \dots, x_r)$ and $y = (y_1, \dots, y_s)$ such that $\sum_{i=1}^r x_i = \sum_{j=1}^s y_j$. Let us also fix $t \leq r$ and $u \leq s$ such that $\sum_{i=1}^t x_i = \sum_{j=1}^u y_j$. Recall that a corner flip is an operation that transforms a table $T[x \times y]$ into a table $T[x' \times y']$, where $x' = (x_1, x_2, \dots, x_t, y_{u+1}, y_{u+2}, \dots, y_s)$ and $y' = (y_1, y_2, \dots, y_u, x_{t+1}, x_{t+2}, \dots, x_r)$. We prove the following proposition:

Proposition 17. With the notation as above, $f(x \times y; J_3) = f(x' \times y'; J_3)$.

We introduce the following terminology: let M be a matrix with at least t rows and at least u columns. The *south-west corner* of M , denoted by M_{SW} , is the submatrix of M formed by the intersection of the first t rows with the first u columns. Similarly, M_{SE} denotes the *south-east corner* of M , which is the intersection of the first t rows of M with the columns of index greater than u . The north-east and north-west corners of M are defined analogously. Thus, a matrix M of shape $r \times s$ can be expressed as $M = \begin{pmatrix} M_{\text{NW}} & M_{\text{NE}} \\ M_{\text{SW}} & M_{\text{SE}} \end{pmatrix}$.

Notice that if M is a filling of $T[x \times y]$, then the sum of the entries of M_{SE} is equal to the sum of the entries of M_{NW} (recall that we assume that the first t rows have the same sum as the first u columns). The rows of M with indices $1, \dots, t$ are called *the southern rows*, the rows with indices greater than t are *the northern rows*, and similarly for the eastern and western columns.

Let (X, Y) be a pair of matrices. We say that a matrix M *completes X and Y inside $T[x \times y]$* if M is a J_3 -avoiding filling of $T[x \times y]$ with $M_{\text{SW}} = X$ and $M_{\text{NE}} = Y$. The following two lemmas immediately imply Proposition 17.

Lemma 18. For any pair of matrices (X, Y) , there is at most one M that completes (X, Y) inside $T[x \times y]$.

Lemma 19. *A pair of matrices (X, Y) can be completed inside $T[x \times y]$ if and only if the pair (X, Y^T) can be completed inside $T[x' \times y']$, where Y^T denotes the transpose of Y .*

By these lemmas, there is a bijection ϕ that maps a J_3 -avoiding filling M of $T[x \times y]$ to the J_3 -avoiding filling $\phi(M) = M'$ of $T[x' \times y']$ uniquely determined by the condition $M'_{\text{SW}} = M_{\text{SW}}$ and $M'_{\text{NE}} = M_{\text{NE}}^T$. The existence of such a bijection implies Proposition 17. It remains to prove the two lemmas.

Proof of Lemma 18. It is enough to prove that if M is a J_3 -avoiding filling of $T[x \times y]$ then both M_{NW} and M_{SE} avoid J_2 . By Lemma 12, a J_2 -avoiding matrix is uniquely determined by its row sums and column sums; in particular, M_{SE} and M_{NW} are determined by x , y and the two matrices $X = M_{\text{SW}}$ and $Y = M_{\text{NE}}$.

Assume that M is a J_3 -avoiding filling of $T[x \times y]$ and M_{NW} contains J_2 . Since the sum of entries of M_{NW} is equal to the sum of the entries of M_{SE} , we know that M_{SE} contains at least one positive entry. This positive entry and the occurrence of J_2 inside M_{NW} form the forbidden pattern J_3 , which is a contradiction, showing that M_{NW} avoids J_2 . By the same argument, we obtain that M_{SE} avoids J_2 as well. \square

Before we present the proof of Lemma 19, we state and prove a lemma that characterizes the pairs (X, Y) that can be completed inside $T[x \times y]$. We will say that a pair of matrices (X, Y) is *plausible* for $T[x \times y]$, if X and Y both avoid J_3 , X has shape $t \times u$, Y has shape $(r - t) \times (s - u)$, and the row sums and column sums of the matrix $\begin{pmatrix} 0 & Y \\ X & 0 \end{pmatrix}$ do not exceed the corresponding constraints x and y .

Lemma 20. *Let (X, Y) be a pair of matrices, let $M_0 = \begin{pmatrix} 0 & Y \\ X & 0 \end{pmatrix}$. Let \bar{x}_i be the sum of the i -th row of M_0 and \bar{y}_j the sum of its j -th column. We say that the i -th row (or j -th column) is saturated if $x_i = \bar{x}_i$ (or $y_j = \bar{y}_j$). The pair (X, Y) can be completed inside $T[x \times y]$ if and only if the following conditions are satisfied:*

- (a) (X, Y) is plausible.
- (b) $\sum_{i=1}^t (x_i - \bar{x}_i) = \sum_{j=u+1}^s (y_j - \bar{y}_j)$ (which is equivalent to $\sum_{j=1}^u (y_j - \bar{y}_j) = \sum_{i=t+1}^r (x_i - \bar{x}_i)$.)
- (c) Let i_S be the largest index of a southern row of M_0 such that for every $i < i_S$, the i -th row is saturated (in other words, i_S is the first unsaturated row, or $i_S = t$ if all southern rows are saturated, see Figure 2.2). Similarly, let j_W be the largest index of a western column such that for every $j < j_W$, the j -th column is saturated. The submatrix of M_0 induced by the rows $\{i_S + 1, \dots, t\}$ and columns $\{j_W + 1, \dots, u\}$ has all entries equal to 0.
- (d) With i_S and j_W as above, the submatrix of M_0 induced by the rows $\{1, \dots, i_S\}$ and columns $\{j_W + 1, \dots, u\}$ avoids J_2 . The submatrix induced by the rows $\{i_S + 1, \dots, t\}$ and columns $\{1, \dots, j_W\}$ avoids J_2 as well.
- (e) Let i_N be the smallest row-index of a northern row such that for every $i > i_N$, the i -th row is saturated. Similarly, let j_E be the smallest column index of an eastern column such that for every $j > j_E$, the j -th column is saturated. The submatrix of M_0 induced by the rows $\{t + 1, \dots, i_N - 1\}$ and columns $\{u + 1, \dots, j_E - 1\}$ has all entries equal to 0.
- (f) With i_N and j_E as above, the submatrix of M_0 induced by the rows $\{t + 1, \dots, i_N - 1\}$ and columns $\{j_E, \dots, s\}$ avoids J_2 . The submatrix induced by the rows $\{i_N, \dots, r\}$ and columns $\{u + 1, \dots, j_E - 1\}$ avoids J_2 as well.

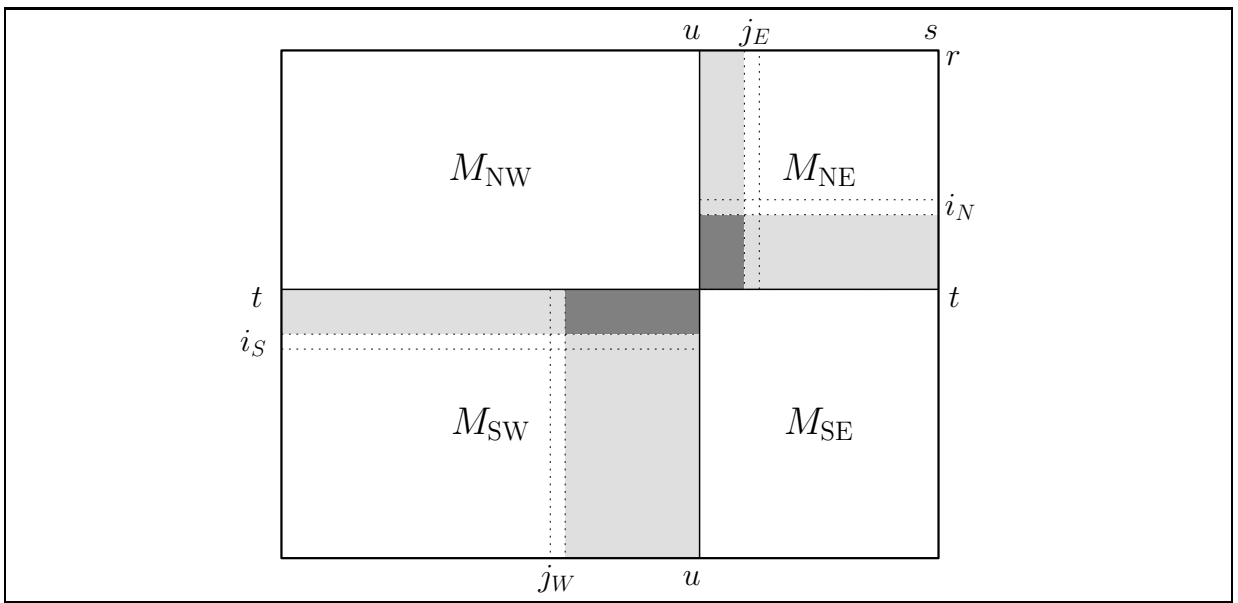


Figure 2.2: Illustration of the four conditions (c), (d), (e) and (f) from Lemma 20. The dark gray rectangles correspond to the submatrices with all entries equal to zero by conditions (c) and (e). The light gray rectangles correspond to submatrices avoiding J_2 by conditions (d) and (f).

Proof. We first show that the conditions are necessary. This is obvious in the case of (a) and (b). Assume that M completes X and Y in $T[x \times y]$. Assume, for contradiction, that condition (c) does not hold. Then M has a positive entry $M_{ij} > 0$ with $i_S < i \leq t$ and $j_W < j \leq u$. Since i_S is smaller than t , it is unsaturated, otherwise we would get a contradiction with i_S 's maximality. Thus, M has at least one positive entry in row i_S and an eastern column. Similarly, M has a positive entry in column j_W and a northern row. These three positive entries form the forbidden pattern J_3 .

Assume now, that condition (d) fails. If the submatrix induced by the rows $1, \dots, i_S$ and columns $\{j_W + 1, \dots, u\}$ contains J_2 , it means that $j_W < u$ and j_W is unsaturated. Hence, M contains a positive entry in column j_W and a northern row, creating the forbidden J_3 . By an analogous argument, there is no J_2 in the submatrix formed by rows $\{i_S + 1, \dots, t\}$ and columns $\{1, \dots, j_W\}$.

The arguments for the necessity of (e) and (f) are symmetric to the arguments given for the necessity of (c) and (d), respectively.

It remains to show that the conditions (a) to (f) are sufficient. Assume that X and Y satisfy these conditions. Fix J_2 -avoiding matrices M_{SE} and M_{NW} in such a way that $M = \begin{pmatrix} M_{NW} & Y \\ X & M_{SE} \end{pmatrix}$ is a filling of $T[x \times y]$ (we do not know yet that M avoids J_3). By condition (b) and by Lemma 12, we know that such M_{NE} and M_{SW} exist and are uniquely determined. By the proof of Lemma 18, we know that M is the only candidate for a completion of (X, Y) inside $T[x \times y]$.

It remains to show that M avoids J_3 . For contradiction, assume that M contains J_3 . Fix three positive cells in M forming J_3 . Assume that these cells appear in rows $i_1 < i_2 < i_3$ and columns $j_1 > j_2 > j_3$. At most one of the three cells is in M_{NW} and at most one is in M_{SE} , because these two corners avoid J_2 by construction. It follows that the cell (i_2, j_2) is either inside X or inside Y . Assume that it is inside X (the other case is symmetric). Thus, we have $i_2 \leq t$ and $j_2 \leq u$. However, it is not possible to have the complete copy of J_3 inside X (because (X, Y) is plausible and thus X avoids J_3), so we may assume, losing no generality, that the nonzero cell (i_1, j_1) is in M_{SE} . It follows that i_1 is not saturated, which means that $i_S \leq i_1 < i_2$.

If the 1-cell (i_3, j_3) is in M_{NW} , we similarly obtain $j_W \leq j_3 < j_2$ contradicting condition (c). On the other hand, if this cell is inside X , then we have a contradiction with condition (c) or (d). \square

With the characterization of the matrix pairs (X, Y) that can be completed inside $T[x \times y]$, the proof of Lemma 19 is easy:

Proof of Lemma 19. It suffices to check that a pair (X, Y) satisfies the conditions of Lemma 20 with respect to $T[x \times y]$ if and only if the pair (X, Y^T) satisfies these conditions with respect to $T[x' \times y']$. This is obvious for conditions (a) and (b). For the remaining four conditions, we may observe that a saturated southern row or western column of $\begin{pmatrix} 0 & Y \\ X & 0 \end{pmatrix}$ remains saturated in $\begin{pmatrix} 0 & Y^T \\ X & 0 \end{pmatrix}$. Similarly, a saturated northern row of index $t + i$ in $\begin{pmatrix} 0 & Y \\ X & 0 \end{pmatrix}$ corresponds to a saturated eastern column of index $u + i$ in $\begin{pmatrix} 0 & Y^T \\ X & 0 \end{pmatrix}$ and vice versa. Combining this with the observation that transposition preserves copies of J_2 , we see that the last four conditions of Lemma 20 are unaffected by the transition from $\begin{pmatrix} 0 & Y \\ X & 0 \end{pmatrix}$ to $\begin{pmatrix} 0 & Y^T \\ X & 0 \end{pmatrix}$ and from $T[x \times y]$ to $T[x' \times y']$. \square

Let us now assemble these pieces into the proof of Theorem 11.

Proof of Theorem 11. We have observed earlier that the result is easy for matrices of order at most two (see Lemma 12 and the preceding discussion). Thanks to Lemma 13, we only need to prove the theorem for a single permutation matrix P of order three. Our matrix of choice is J_3 . By Propositions 14 and 17, all we have to do is notice that for any two tables $T = T[x \times y]$ and $T' = T[x' \times y']$ with the same scoreline, we may transform T into T' by a sequence of permutations and corner flips, which is indeed easily seen. \square

We conclude this section with some remarks and examples that put fillings of rectangular shapes into a broader context of pattern avoidance in fillings.

Let us only consider patterns that are permutation matrices, and let us make no distinction between a permutation and its matrix.

Recall that two permutations π, σ are *Wilf equivalent* (denoted by $\pi \overset{w}{\sim} \sigma$) if they are equirestrictive in the set of all permutations. In the notation of rectangle fillings, this may be written as $f(1^n \times 1^n; \pi) = f(1^n \times 1^n; \sigma)$, where 1^n is the sequence of n ones.

Allowing arbitrary constraints, we write $\pi \overset{g}{\sim} \sigma$ if for every constrained table $T[x \times y]$ we have $f(x \times y; \pi) = f(x \times y; \sigma)$ (the letter ‘g’ stands for ‘general’ fillings, as opposed to the transversal fillings considered in Wilf and shape-Wilf equivalences).

As we have already seen, the integer fillings of rectangular shapes naturally generalize to integer fillings of Ferrers shapes. Let $T_\lambda[x \times y]$ denote the Ferrers diagram of shape λ with row constraints x and column constraints y . Let $f_\lambda(x \times y; \pi)$ be the number of fillings of $T_\lambda[x \times y]$ that avoid a pattern π . Finally, let us write $\pi \overset{lg}{\sim} \sigma$ if the identity $f_\lambda(x \times y; \pi) = f_\lambda(x \times y; \sigma)$ holds for any constrained Ferrers shape $T_\lambda[x \times y]$. The equivalence $\overset{lg}{\sim}$ is to the shape-Wilf equivalence $\overset{lw}{\sim}$, what $\overset{g}{\sim}$ is to the Wilf equivalence $\overset{w}{\sim}$.

In general, $\overset{lg}{\sim}$ is different from $\overset{lw}{\sim}$; for example, for $\lambda = (4, 4, 4, 3)$ we have

$$18 = f_\lambda((1, 1, 2, 1) \times (2, 1, 1, 1); \begin{pmatrix} I_1 & 0 \\ 0 & I_2 \end{pmatrix}) \neq f_\lambda((1, 1, 2, 1) \times (2, 1, 1, 1); \begin{pmatrix} I_2 & 0 \\ 0 & I_1 \end{pmatrix}) = 17,$$

even though $\begin{pmatrix} I_1 & 0 \\ 0 & I_2 \end{pmatrix} \overset{lw}{\sim} \begin{pmatrix} I_2 & 0 \\ 0 & I_1 \end{pmatrix}$. Thus, $\overset{lg}{\sim}$ is a proper refinement of $\overset{lw}{\sim}$.

As we have seen, all the permutations of order three are $\overset{g}{\sim}$ -equivalent, which shows that $\overset{g}{\sim}$ is different from $\overset{lw}{\sim}$ and $\overset{lg}{\sim}$. To see that $\overset{g}{\sim}$ is also different from $\overset{w}{\sim}$, consider the

following two patterns:

$$P = \begin{pmatrix} & & 1 \\ & 1 & \\ 1 & & \\ & & 1 \end{pmatrix} \quad \text{and} \quad \overline{P} = \begin{pmatrix} & & 1 \\ 1 & & \\ & 1 & \\ & & 1 \end{pmatrix}.$$

Clearly $P \stackrel{w}{\sim} \overline{P}$, since \overline{P} is symmetric to P ; on the other hand, for $x = (1, 1, 1, 2, 1)$ we have

$$165 = f(x \times x; P) \neq f(x \times x; \overline{P}) = 166,$$

which shows that P and \overline{P} are not $\stackrel{g}{\sim}$ -equivalent. This example can also be interpreted as $f(x \times x; P) \neq f(\overline{x} \times x; P)$, where \overline{x} is the sequence x written backwards. This shows that Proposition 14 does not generalize to all forbidden patterns.

We may apply the red-green argument to obtain further examples of $\stackrel{g}{\sim}$ -equivalent patterns. In general, it is not true that $A \stackrel{g}{\sim} A'$ implies $\begin{pmatrix} 0 & B \\ A & 0 \end{pmatrix} \stackrel{g}{\sim} \begin{pmatrix} 0 & B \\ A' & 0 \end{pmatrix}$ (compare the example P above with I_4). On the other hand, using Corollary 8, a single pair of $\stackrel{lg}{\sim}$ -equivalent patterns can be turned into a family of $\stackrel{g}{\sim}$ -equivalent patterns, by a red-green argument similar to Proposition 1. In particular, if $A \stackrel{lg}{\sim} A'$ and $B \stackrel{lg}{\sim} B'$ then $\begin{pmatrix} 0 & \overrightarrow{B} \\ A & 0 \end{pmatrix} \stackrel{g}{\sim} \begin{pmatrix} 0 & \overrightarrow{B'} \\ A' & 0 \end{pmatrix}$, where \overrightarrow{B} and $\overrightarrow{B'}$ are the matrices obtained from B and B' by the rotation of 180 degrees (similar arguments can be made for other symmetries of the square). For instance, we may conclude that $I_4 \stackrel{g}{\sim} \begin{pmatrix} 0 & J_2 \\ J_2 & 0 \end{pmatrix}$.

It is not known whether there are any examples of $\stackrel{g}{\sim}$ -equivalent or even $\stackrel{lg}{\sim}$ -equivalent patterns, apart from those that can be deduced from the equivalence of I_k and J_k using obvious symmetries and the red-green argument outlined above.

Chapter 3

Wilf order

Currently, it seems difficult to proceed with the task of Wilf classification of permutations. The Wilf equivalences of small patterns are well understood, while for larger patterns, it is difficult to use computer-assisted enumeration to generate new conjectures. In this situation, the topic of Wilf order starts receiving more attention as a promising tool to gain more insight into Wilf classification of permutations.

For two permutations σ and τ , let us write $\sigma \preceq \tau$ if σ is more restrictive than τ in the class of permutations, i.e., the number of σ -avoiding permutations of a given order is smaller than or equal to the number of τ -avoiding permutations. The relation \preceq defines a quasi-order of permutations.

Analogously, we may define the shape-Wilf order, denoted by \preceq_{sw} , where $\sigma \preceq_{\text{sw}} \tau$ means that for every Ferrers shape F , the number of σ -avoiding transversals of F does not exceed the number of its τ -avoiding transversals.

In this chapter, we will often deal with shapes and their fillings. It is thus convenient to keep the convention that a permutation is represented by a permutation matrix, and to ignore the distinction between a permutation and its matrix.

Let $\text{diag}(A_1, A_2, \dots, A_k)$ denote the block-diagonal matrix whose blocks are the matrices A_1, \dots, A_k , in left-to-right order. Formally, we may define $\text{diag}(A_1, A_2, \dots, A_k)$ inductively, by saying that $\text{diag}(A_1) = A_1$, while for $k > 1$, we have

$$\text{diag}(A_1, A_2, \dots, A_k) = \begin{pmatrix} 0 & \text{diag}(A_2, A_3, \dots, A_k) \\ A_1 & 0 \end{pmatrix}.$$

Note that a block-diagonal matrix whose every block is a permutation matrix is itself a permutation matrix. A block-diagonal permutation whose every block is a diagonal matrix (i.e., either I_k or J_k for some k) is known as *layered permutation*.

Similarly to Wilf equivalence, new results involving the Wilf order may be deduced from results involving the shape-Wilf order. Indeed, all the known results on Wilf order are corollaries to the following proposition.

Proposition 21. *Let A , B and C be three permutations. If $A \preceq_{\text{sw}} B$ then $\begin{pmatrix} 0 & C \\ A & 0 \end{pmatrix} \preceq_{\text{sw}} \begin{pmatrix} 0 & C \\ B & 0 \end{pmatrix}$, and hence also $\begin{pmatrix} 0 & C \\ A & 0 \end{pmatrix} \preceq \begin{pmatrix} 0 & C \\ B & 0 \end{pmatrix}$.*

The proof of this proposition is based on a red-green argument completely analogous to the proof of Proposition 1, and we omit it.

In the rest of this chapter we review the known results on shape-Wilf order, and then present several conjectures related to the Wilf and shape-Wilf order relations.

3.1 Shape-Wilf order of small patterns

Since the two permutations in \mathcal{S}_2 are shape-Wilf equivalent, the first non-trivial results related to the shape-Wilf order involve patterns of size three. As we have seen, \mathcal{S}_3 is partitioned into three shape-Wilf classes:

- $312 \stackrel{\text{w}}{\sim} 231$,
- $123 \stackrel{\text{w}}{\sim} 213 \stackrel{\text{w}}{\sim} 321$,
- 132 .

The first result involving the shape-Wilf order has been obtained by the author [34]. It deals with the first two of the three shape-Wilf classes above. A different proof of the same result has been also obtained by Stankova [64]. The statement of the result is simple:

Theorem 22 (J. [34]). $231 \stackrel{\text{w}}{\preceq} 123$.

The other known result on shape-Wilf order, which finishes the classification of patterns of size three, is due to Stankova [64], and its statement is equally simple:

Theorem 23 (Stankova [64]). $123 \stackrel{\text{w}}{\preceq} 132$.

Stankova has pointed out that the shape-Wilf ordering of patterns of size three, together with Proposition 21, makes it possible to deduce the Wilf ordering of all the permutations in \mathcal{S}_4 . Recall that \mathcal{S}_4 has three Wilf classes, which may be represented by the patterns 2314, 1234, and 1324. From Proposition 21 and Theorems 22 and 23, we obtain the chain

$$2314 \stackrel{\text{w}}{\preceq} 1234 \stackrel{\text{w}}{\preceq} 1324.$$

3.2 Skew order

Recall that a skew shape is a shape obtained as the difference of two Ferrers shapes sharing a common bottom-left corner. It turns out that pattern avoidance among transversals of skew shapes has interesting consequences for shape-Wilf order and hence also for the Wilf order.

For a pair σ, τ of permutations, let us write $\sigma \stackrel{\text{sk}}{\preceq} \tau$, if for every skew shape S the number of σ -avoiding transversals of S does not exceed the number of its τ -avoiding transversals. The relation $\stackrel{\text{sk}}{\preceq}$ will be called *the skew order*. Every Ferrers shape is a skew shape, which means that $\sigma \stackrel{\text{sk}}{\preceq} \tau$ implies $\sigma \stackrel{\text{w}}{\preceq} \tau$.

At this point, the reader might wonder why we have not yet introduced the ‘skew equivalence’, as the natural refinement of the shape-Wilf equivalence. The reason is, that we are not aware of any pair of distinct patterns that would be equirestrictive with respect to transversals of skew shapes.

Let us present several new definitions related to skew shapes. Recall that $r(S)$ and $c(S)$ denote respectively the number of rows and columns of a diagram S . We say that a skew shape S is *proper* if $r(S) = c(S)$, and moreover, every row and every column of S contains at least one cell of S . A proper skew shape S is called *permissible* if it has at least one transversal. The following simple observation characterizes the permissible skew shapes.

Observation 24. A proper skew shape S with n rows is permissible if and only if for every $i \in [n]$ the shape S contains the cell in row i and column $n - i + 1$.

Proof. If S contains all the cells $\{(i, n - i + 1), i \in [n]\}$, then the filling which assigns the value 1 to these cells and the value 0 to all the other cells is a transversal, hence S is permissible.

Assume now that for some $i \in [n]$ the cell $(i, n - i + 1)$ does not belong to S . Without loss of generality, assume that the cell $(i, n - i + 1)$ is to the left of all the cells of S in row i . The definition of skew shape then implies that all the cells in rows $1, 2, \dots, i$ are strictly to the right of the column $n - i + 1$. Thus, the bottom i rows of S intersect at most $i - 1$ columns of S , which implies that S has no transversal. \square

Our motivation for the study of the skew order is based on the following proposition, which allows us to construct a family of \succsim -comparable patterns from a single pair of $\stackrel{\text{sk}}{\succsim}$ -comparable patterns.

Proposition 25. Let A, B and C be three permutations. If $A \stackrel{\text{sk}}{\succsim} B$, then $\begin{pmatrix} 0 & A \\ C & 0 \end{pmatrix} \stackrel{\text{sk}}{\succsim} \begin{pmatrix} 0 & B \\ C & 0 \end{pmatrix}$, and hence also $\begin{pmatrix} 0 & A \\ C & 0 \end{pmatrix} \stackrel{\text{sk}}{\prec} \begin{pmatrix} 0 & B \\ C & 0 \end{pmatrix}$.

Proof. The proof is very similar to the proof of Proposition 1, and uses an analogous reg-green argument. Instead of repeating the whole proof again, we content ourselves with sketching the main points.

Let S be a skew shape. Our aim is to present an injective mapping that transforms a $\begin{pmatrix} 0 & A \\ C & 0 \end{pmatrix}$ -avoiding transversal of S into a $\begin{pmatrix} 0 & B \\ C & 0 \end{pmatrix}$ -avoiding transversal of S . Let Φ be a $\begin{pmatrix} 0 & A \\ C & 0 \end{pmatrix}$ -avoiding transversal of S . We color the cells of Φ red and green: a cell (i, j) is green with respect to Φ if the subfilling of Φ to the bottom-left of (i, j) (i.e., the subfilling formed by the intersection of the bottommost $i - 1$ rows with the leftmost $j - 1$ columns) contains the pattern C . A cell is red if it is not green.

Let Φ_G be the subfilling of Φ formed by the green cells. Clearly, Φ_G is a sparse A -avoiding skew filling. Since $A \stackrel{\text{sk}}{\succsim} B$, we know that there is an injective mapping ϕ_0 that transforms A -avoiding skew transversals into B -avoiding skew transversals of the same shape. By an argument analogous to Lemma 2, we may extend ϕ_0 into a shape-preserving mapping ϕ that transforms A -avoiding sparse skew fillings injectively into B -avoiding sparse fillings. Furthermore, ϕ preserves the zero rows and zero columns of the filling.

We apply the mapping ϕ to the filling Φ_G , transforming it into a B -avoiding filling Ψ_G , while all the red cells of Φ remain unchanged. We thus obtain a $\begin{pmatrix} 0 & B \\ C & 0 \end{pmatrix}$ -avoiding transversal Ψ of the shape S . As in the proof of Proposition 1, we again see that a cell (i, j) is red with respect to the filling Φ if and only if it is red with respect to Ψ . This implies that the transformation we described here is indeed an injection. \square

To make Proposition 25 useful, we need to find some $\stackrel{\text{sk}}{\succsim}$ -comparable patterns. The most natural candidates are the diagonal patterns I_n and J_n . After extensive computer enumeration, we are confident enough to make the following conjecture.

Conjecture 26. For every $k \in \mathbb{N}$, the following holds:

1. $I_k \stackrel{\text{sk}}{\succsim} J_k$.
2. For any permutation C , we have $\begin{pmatrix} 0 & I_k \\ C & 0 \end{pmatrix} \stackrel{\text{sk}}{\prec} \begin{pmatrix} 0 & J_k \\ C & 0 \end{pmatrix}$.
3. For any two permutations C and D , we have $\text{diag}(C, I_k, D) \stackrel{\text{sk}}{\preceq} \text{diag}(C, J_k, D)$.
4. If A is a layered permutation of order k , then $I_k \stackrel{\text{sk}}{\preceq} A$.




Figure 3.1: The pattern P used in the proof of Theorem 28.

The four statements of this conjecture are listed in the order of decreasing strength, i.e., each statement is a consequence of the previous one. Indeed, the second statement follows from the first by Proposition 25, the third from the second by Proposition 21, while the fourth follows from the third by a simple observation.

With regards to the fourth claim of Conjecture 26, it is noteworthy that Bóna [13] has proved the following asymptotic version of the claim.

Theorem 27 (Bóna [13]). *If A is a layered permutation of order k , then the Stanley–Wilf limit of A is greater than or equal to the Stanley–Wilf limit of I_k .*

Let us mention that the Stanley–Wilf limit of I_k is equal to $(k-1)^2$. In fact, Regev [55] has found an explicit formula for the number of I_k -avoiding permutations of order n .

We have so far been unable to prove Conjecture 26 in full generality. We are only able to verify the conjecture for the smallest nontrivial case, i.e., $k = 2$.

Theorem 28. $I_2 \stackrel{sk}{\preceq} J_2$. *In particular, for any permutation C we have $\begin{pmatrix} 0 & I_2 \\ C & 0 \end{pmatrix} \stackrel{\preceq}{\asymp} \begin{pmatrix} 0 & J_2 \\ C & 0 \end{pmatrix}$, and for any pair of permutations C and D we have $\text{diag}(C, I_2, D) \preceq \text{diag}(C, J_2, D)$.*

Notice that even this simplest case of Conjecture 26 stated in Theorem 28 is already more general than Theorem 23.

To prove Theorem 28, we first show that every permissible skew shape has exactly one I_2 -avoiding transversal. To complete the proof, it then suffices to show that every permissible skew shape has at least one J_2 -avoiding transversal. We will in fact prove a stronger statement, by showing that every permissible skew shape has exactly one transversal that simultaneously avoids J_2 and the skew pattern P in Figure 3.1. Since the number of $\{J_2, P\}$ -avoiding transversals cannot be greater than the number of J_2 -avoiding transversals, the proof of Theorem 28 will follow easily.

Lemma 29. *Every permissible skew shape has exactly one I_2 -avoiding transversal.*

Proof. Recall from Observation 24 that every permissible skew shape S with n rows admits the ‘antidiagonal’ transversal, whose 1-cells are exactly the cells of the form $(i, n - i + 1)$, for $i \in [n]$. Clearly, this transversal avoids I_2 .

We claim that any other transversal of S contains I_2 . To see this, represent a transversal of S by a permutation $\tau = \tau_1 \tau_2 \cdots \tau_n$ where τ_i is the index of the column containing the 1-cell in row i . If the transversal is different from the antidiagonal transversal described above, then τ is different from the anti-identity permutation J_n . Thus, τ must contain 12 as a subpermutation, and consequently, the transversal must contain I_2 . \square

Let us now concentrate on the pattern J_2 . Note that a skew shape may admit more than one J_2 -avoiding transversal. For example, the pattern P from Figure 3.1, as well as its transpose P^T are two J_2 -avoiding transversals of the same shape.

As we said above, our aim is to prove that every permissible skew shape has exactly one transversal that avoids both J_2 and P . We will again represent a transversal of a shape S by the permutation τ , defined as in the proof of the previous lemma.

To compare two transversals of a given shape, we will use the standard notion of lexicographic ordering: for two different sequences of integers $\tau = \tau_1\tau_2\cdots\tau_n$ and $\sigma = \sigma_1\sigma_2\cdots\sigma_n$, we say that τ is *lexicographically smaller* than σ if $\tau_i < \sigma_i$, where i is the smallest index where the two sequences differ.

Notice that the unique I_2 -avoiding transversal of a given skew shape S (which is represented by the anti-identity permutation) is the lexicographically largest of all transversals of S . For the unique $\{J_2, P\}$ -avoiding transversal, we have the opposite characterisation.

Lemma 30. *Let S be a permissible skew shape. Let τ be the lexicographically smallest transversal of S . Then τ is the unique transversal of S that avoids both patterns J_2 and P .*

Proof. Let $\tau = \tau_1\tau_2\cdots\tau_n$ be the lexicographically smallest transversal of a skew shape S . Let us first prove that τ avoids both forbidden patterns. Notice that neither of the two forbidden patterns is itself the lexicographically smallest transversal of its underlying shape. Thus, if τ contains a copy of J_2 or P , we may modify the filling of the subshape that contains this copy by replacing the forbidden pattern with a lexicographically smaller transversal of the same shape. This modification transforms τ into a lexicographically smaller transversal, contradicting its minimality.

Assume now, for the sake of contradiction, that $\sigma = \sigma_1\sigma_2\cdots\sigma_n$ is a transversal that is different from τ and also avoids J_2 and P . Let us refer to the 1-cells of τ as τ -cells, while the 1-cells of σ will be called σ -cells. Let i be the smallest index where σ and τ differ (see Fig. 3.2). By the minimality of τ , we have $\tau_i < \sigma_i$. Let j be the index of the highest row that intersects the column σ_i . Let k be the row-index of the σ -cell that appears in column τ_i (i.e., k satisfies the equality $\sigma_k = \tau_i$).

The index k cannot be smaller than i , since the σ -cells below row i coincide with the τ -cells. We also know that $k \neq i$ since $\tau_i \neq \sigma_i$. This leaves us with $k > i$. If $k \leq j$, then the two rows i and k and the two columns σ_k and σ_i induce a copy of J_2 in σ .

Assume now that $k > j$. Let us say that a cell is *high* if its row index is greater than j , and a cell is *low* otherwise. Clearly, there are exactly j low σ -cells and j low τ -cells. In column σ_k , there is a low τ -cell in row i together with a high σ -cell in row k . Since the number of low σ -cells equals the number of low τ -cells, there must also be a column that contains a low σ -cell and a high τ -cell. Let c be such a column, and let ℓ be the row index of the σ -cell in column c . Note that $c < \sigma_i$, otherwise c would not contain any high cells. If the column c intersects row i , then the two rows i and ℓ with the two columns c and σ_i induce a copy of J_2 in σ . On the other hand, if the column c does not intersect row i , this means that $c < \tau_i = \sigma_k$, and the three rows i, ℓ, k together with the columns c, σ_k, σ_i induce in σ a copy of P .

In any case, we get a contradiction. □

As we already explained, Theorem 28 is a direct consequence of the two lemmas above.




Figure 3.2: Illustration of the proof of Lemma 30.

Chapter 4

Involutions

Let us now turn our attention to the topic of pattern avoidance in involutions. Recall that an involution is a permutation whose matrix is symmetric. Let \mathcal{J}_n denote the set of involutions of order n , and let $\mathcal{J}_n(\sigma)$ be the set of the involutions of order n that avoid σ . For two permutation patterns σ and τ , we write $\sigma \stackrel{I}{\sim} \tau$ if the two patterns are equirestrictive in the set of involutions, i.e., if $|\mathcal{J}_n(\sigma)| = |\mathcal{J}_n(\tau)|$ for each n . We will call the relation $\stackrel{I}{\sim}$ *the I-Wilf equivalence*.

Note that a permutation is not necessarily I-Wilf equivalent to its reversal. However, it is clear that a permutation is I-Wilf equivalent to its inverse, as well as to the permutation obtained by reflecting the permutation matrix along the decreasing diagonal.

Pattern avoidance of involutions was already studied by Simion and Schmidt [62], who classified patterns of size three with respect to I-Wilf equivalence. They have shown that for any pattern $\tau \in \{123, 213, 132, 321\}$, there are exactly $\binom{n}{\lfloor n/2 \rfloor}$ τ -avoiding involutions of order n , while for $\tau \in \{231, 312\}$, there are 2^{n-1} such τ -avoiding involutions.

From known symmetries of the RSK algorithm (which may be found, e.g., in Fulton's book [27]), it is easy to see that $I_k \stackrel{I}{\sim} J_k$ for any k , where I_k and J_k denote the identity and the anti-identity matrix of order k .

Guibert [29] has shown that $3412 \stackrel{I}{\sim} 4321$ and $2143 \stackrel{I}{\sim} 1243$. He conjectured that both 2143 and 1432 are actually I-Wilf equivalent to 4321 (and hence also to 1234). The first part of this conjecture was proved by Guibert, Pergola and Pinzani [30], who proved the equivalence $1234 \stackrel{I}{\sim} 2143$.

Jaggard [32] generalized these results by proving that for any permutation X , we have $\begin{pmatrix} 0 & X \\ I_2 & 0 \end{pmatrix} \stackrel{I}{\sim} \begin{pmatrix} 0 & X \\ J_2 & 0 \end{pmatrix}$ and $\begin{pmatrix} 0 & X \\ I_3 & 0 \end{pmatrix} \stackrel{I}{\sim} \begin{pmatrix} 0 & X \\ J_3 & 0 \end{pmatrix}$.

Furthermore, Jaggard made the following conjectures:

1. $\begin{pmatrix} 0 & X \\ I_k & 0 \end{pmatrix} \stackrel{I}{\sim} \begin{pmatrix} 0 & X \\ J_k & 0 \end{pmatrix}$ for any $k \geq 1$ and any permutation matrix X ,
2. $12345 \stackrel{I}{\sim} 45312$,
3. $123456 \stackrel{I}{\sim} 456123 \stackrel{I}{\sim} 564312$.

The first of these conjectures was settled by Bousquet-Mélou and Steingrímsson [14], whose result was presented here as Corollary 6 in Chapter 2. From this corollary, it is possible to deduce the following equivalence, valid for any permutation matrix X and any two integers $k, l \geq 0$:

$$\begin{pmatrix} & & I_l \\ & X & \\ I_k & & \end{pmatrix} \stackrel{I}{\sim} \begin{pmatrix} & & J_l \\ & X & \\ J_k & & \end{pmatrix}.$$




Figure 4.1: An example of an F^\perp -shape. The shaded cells are the corners.

The other two conjectures of Jaggard are both consequences of a more general result proved by Dukes, Jelínek, Mansour and Reifegerste [22]. The proof of this result is the main topic of this chapter.

The proof we present here is taken from the above-mentioned paper of Dukes et al. [22]. Let us remark that the paper in fact considers a more general setting of *signed permutations*, which may be represented by matrices with entries 0, 1 and -1 , with the property that each row and each column has exactly one nonzero entry. In this thesis, we will restrict ourselves to the less general setting of ordinary permutations, to avoid unnecessary technical complications.

We also point out that, independently of Dukes et al. [22], Jaggard and Marincel [33] have proved that for every $k \geq 5$, the permutation J_k is I-Wilf equivalent to $(k-1)k(k-2)(k-3)\cdots 4312$. The proof of Jaggard and Marincel uses a different method than the proof of Dukes et al. that we will show in this thesis.

The proof we are about to present is based on previously known results on shape-Wilf equivalence, combined with a suitably adapted version of the red-green argument. For technical reasons, apart from using usual Ferrers diagrams (which are bottom-left aligned shapes), we will also need to refer to the bottom-right aligned corner shapes, i.e., the shapes obtained from Ferrers diagrams by reflection along a vertical axis. We will call these shapes *the F^\perp -shapes*. We will write $\sigma \overset{sw}{\sim} \tau$ if σ and τ are equirestrictive with respect to transversals of F^\perp -shapes. Of course, $\sigma \overset{sw}{\sim} \tau$ is merely a shorthand for saying that the reversal of σ is shape-Wilf equivalent to the reversal of τ .

Here are the results we are about to prove.

Theorem 31 (Dukes, J., Mansour, Reifegerste [22]). *If A and B are $\overset{sw}{\sim}$ -equivalent matrices and X is any permutation matrix then the following equivalences hold:*

$$\begin{pmatrix} A^T & & & \\ & X^T & & \\ & & X & \\ & & & A \end{pmatrix} \overset{I}{\sim} \begin{pmatrix} B^T & & & \\ & X^T & & \\ & & X & \\ & & & B \end{pmatrix} \quad (4.1)$$

$$\begin{pmatrix} A^T & & & \\ & X^T & & \\ & & 1 & \\ & & & X \\ & & & & A \end{pmatrix} \overset{I}{\sim} \begin{pmatrix} B^T & & & \\ & X^T & & \\ & & 1 & \\ & & & X \\ & & & & B \end{pmatrix} \quad (4.2)$$

As we already mentioned, the proof of Theorem 31 uses a suitably adapted version of the red-green argument. Before we state argument precisely, let us make the following definition: a cell of an F^\perp -shape is called a *corner* if it is the leftmost cell of its row and also the topmost cell of its column. See Figure 4.1 for an example.

Proposition 32. *Let F be an F^\perp -shape, and let A, B, C be permutations, such that A and B are $\overset{w}{\sim}$ -equivalent. We set*

$$X = \begin{pmatrix} C & 0 \\ 0 & A \end{pmatrix} \text{ and } Y = \begin{pmatrix} C & 0 \\ 0 & B \end{pmatrix}.$$

There is a bijection between X -avoiding and Y -avoiding sparse fillings of F . This bijection preserves the number of nonzero entries in each row and column. In particular, X and Y are $\overset{w}{\sim}$ -equivalent. Furthermore, if C is nonempty, the bijection preserves the filling in the corners of F .

The proof of Proposition 32 is analogous to the proof of Proposition 1, and we omit it.

Note that Proposition 32 yields some information even when C is the empty matrix. In such situation, the proposition shows that a bijection between pattern-avoiding transversals can be extended to a bijection between pattern-avoiding sparse fillings, by simply ignoring the rows and columns with no nonzero entries.

We will now show how the results on shape Wilf equivalence may be applied to obtain new classes of I-Wilf equivalent patterns. Let us first give the necessary definitions. For an $n \times n$ matrix M , let M^- denote the subfilling of M formed by the cells of M which are strictly below the main diagonal, and let M_0^- denote the subfilling formed by the cells on the main diagonal and below it. See Figure 4.2 for an example. Note that both M^- and M_0^- are fillings of F^\perp -shapes, and if M is a permutation matrix, then the two fillings are sparse.




Figure 4.2: The fillings M^- and M_0^- .

The rows and columns of M^- and M_0^- will have their numbering inherited from the matrix M . In particular, this means that the leftmost column of M^- does not contain any cell at all, while the only cell in the second column of M^- will be referred to as cell $(1, 2)$.

Analogously, we define M^+ to be the filled shape corresponding to the entries strictly above the main diagonal of M .

Clearly, a symmetric matrix M is completely determined by M_0^- . Observe that a symmetric 01-matrix M is an involution if and only if, for every $i = 1, \dots, n$, the filling M_0^- has exactly one 1-cell in the union of all cells of the i -th row and i -th column.

The filling M^- does not determine a symmetric matrix M uniquely, since it does not carry any information about the diagonal cells. However, if we further assume that M is an involution, then it is again easy to see that M is determined by M^- . Note that in such case, the filling M^- has the property that the union of i -th row and i -th column has at most one 1-cell. Conversely, it is not difficult to see that a filling M^- that satisfies the property above identifies a unique involution M .

For a permutation P , let P' denote the involution $\begin{pmatrix} P^T & 0 \\ 0 & P \end{pmatrix}$. We are now ready to state and prove the first part of our first result on I-Wilf equivalence, which corresponds to the equation (4.1) of Theorem 31.

Proposition 33. *If A and B are two $\overset{w}{\sim}$ -equivalent permutation matrices, then $A' \overset{I}{\sim} B'$. Moreover, the bijection between $\mathcal{J}_n(A')$ and $\mathcal{J}_n(B')$ preserves fixed points.*

Proof. Let $M \in \mathcal{J}_n$ be an involution. We claim that M avoids A' if and only if M^- avoids A . To see this, notice that any occurrence of A' in M can be restricted either to an occurrence of A in M^- or an occurrence of A^T in M^+ . However, since M^+ is the transpose of M^- , we know that M^+ contains A^T if and only if M^- contains A . It follows that if M contains A' then M^- contains A . The converse is even easier to see.

Let us choose $M \in \mathcal{J}_n(A')$. Since M^- is a sparse A -avoiding filling, we may apply the bijection from Proposition 32 to M^- , to obtain a B -avoiding sparse filling Ψ of the same shape.

The new filling Ψ has a nonzero entry in a row i (or column i) whenever M^- has a nonzero entry in the same row (or column, respectively). In particular, the filling Ψ has the property that the union of the i -th row and i -th column has at most one 1-cell. This implies that $\Psi = N^-$ for an involution $N \in I_n$.

Furthermore, the fixed points of N are in the same position as the fixed points of M , because the position of the fixed points is determined by the zero rows and columns, which are preserved by the bijection from Proposition 32.

Clearly, since N^- avoids B , we know that B avoids B' . Each step of this construction can be inverted which proves the bijectivity. The bijection preserves fixed points by construction. \square

By a similar reasoning, we obtain an analogous result for patterns of odd size. For a permutation M , let M'' denote the involution

$$\begin{pmatrix} M^T & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & M \end{pmatrix},$$

and let M^* denote the permutation $\begin{pmatrix} 1 & 0 \\ 0 & M \end{pmatrix}$.

We are now ready to state a proposition which proves the equation (4.2) from Theorem 31.

Proposition 34. *If A and B are $\overset{w}{\sim}$ -equivalent, then $A'' \overset{I}{\sim} B''$. Moreover, the bijection between $I_n(A'')$ and $I_n(B'')$ preserves fixed points.*

Proof. By an argument analogous to the proof of Proposition 33, we may observe that an involution M avoids A'' if and only if M_0^- avoids the pattern A^* . By Proposition 32, the two patterns A^* and B^* are $\overset{w}{\sim}$ -equivalent and furthermore, the bijection realizing this equivalence preserves the filling of the corners of the shape. Note the corners of M_0^- correspond exactly to the diagonal cells of the original permutation matrix M .

Now we consider M_0^- for an involution $M \in I_n(A'')$. The bijection of Proposition 32 maps M_0^- to a B^* -avoiding filling Ψ . Since the bijection preserves the number of nonzero entries in each row and each column of M_0^- , and since it also preserves the entries on the intersection of i -th row and i -th column (these are precisely the corners), we know that the bijection preserves, for each i , the number of nonzero entries in the union of the i -th row and i -th column. In particular, the filling Ψ has exactly one nonzero entry in the union of i -th row and i -th column, which guarantees that there is a unique involution N satisfying $N_0^- = \Psi$. Since Ψ avoids B^* , we know that N avoids B'' .

Because the bijection preserves the entries in the diagonal cells (i, i) , $i = 1, \dots, n$, the permutations M and N have the same fixed points. This provides the required bijection. \square

The proof of the main result now follows easily from the previous propositions.

Proof of Theorem 31. Let A and B be $\overset{\text{w}}{\sim}$ -equivalent patterns. Let C be an arbitrary pattern. By Proposition 32, the patterns $\begin{pmatrix} C & 0 \\ 0 & A \end{pmatrix}$ and $\begin{pmatrix} C & 0 \\ 0 & B \end{pmatrix}$ are $\overset{\text{w}}{\sim}$ -equivalent as well. By applying Propositions 33 and 34 to these two patterns, we obtain directly the two equations of Theorem 31. \square

By recalling that the two diagonal patterns I_k and J_k are shape-Wilf equivalent, and hence also $\overset{\text{w}}{\sim}$ -equivalent, we may now easily see that the two remaining conjectures of Jaggard were correct.

Corollary 35. *We have $54321 \overset{L}{\sim} 45312$ and $654321 \overset{L}{\sim} 456123 \overset{L}{\sim} 564312$.*

Let us mention that a computer enumeration performed by Mansour [22] has verified that among permutations of order at most 7, there are no pairs of I-Wilf equivalent patterns, apart from those that are covered by presently known results.

Chapter 5

Words

In this chapter, we will investigate a very natural generalization of the concept of permutation, namely the k -ary words, also referred to as multiset permutations.

The results presented in this sections are based on a forthcoming article by Jelínek and Mansour [38]. These results extend previous work of Burstein [15], who described the equivalence classes of k -ary words of length at most 3, and of Savage and Wilf [60], who dealt with integer compositions (as well as words) avoiding patterns of length at most 3. We will present several new bijective arguments that extend these results to larger patterns.

5.1 Basic terminology

Let us first recall the terminology and notation related to pattern avoidance of k -ary words.

Let $[k] = \{1, 2, \dots, k\}$ be a linearly ordered alphabet of k letters. We let $[k]^n$ denote the set of words of length n over this alphabet.

Consider two words, $\sigma \in [k]^n$ and $\tau \in [\ell]^m$. Assume additionally that τ contains all letters 1 through ℓ —a word with this property is called *reduced*. We say that σ contains an *occurrence* of τ , or simply that σ *contains* τ , if σ has a subsequence *order-isomorphic* to τ , i.e., if there exist $1 \leq i_1 < \dots < i_m \leq n$ such that, for any two indices $1 \leq a, b \leq m$, $\sigma_{i_a} < \sigma_{i_b}$ if and only if $\tau_a < \tau_b$. If σ contains no occurrences of τ , we say that σ *avoids* τ .

In this chapter, the term *pattern* will refer to an arbitrary reduced word. For a pattern τ , let $[k]^n(\tau)$ denote the set of k -ary words of length n which avoid the pattern τ . We say that two patterns τ and τ' are *word-equivalent* (or, more briefly, *w-equivalent*), if for all values of k and n , we have the identity $|[k]^n(\tau)| = |[k]^n(\tau')|$.

There are two operations on words which trivially preserve the w-equivalence, called the reversal and the complement. The *reversal* of a word $\tau \in [k]^m$, denoted by $\bar{\tau}$, is obtained by writing the letters of τ in the reverse order, i.e., the i -th letter of $\bar{\tau}$ is equal to the $(m - i + 1)$ -th letter of τ . The *complement* of a word τ , denoted by τ^C , is obtained by turning τ “upside-down”, i.e., a letter j is replaced by the letter $\ell - j + 1$, where ℓ is the largest letter of τ . For example, $\overline{1232} = 2321$, $1232^C = 3212$, and $\overline{1232^C} = \overline{1232}^C = 2123$.

Several authors have previously considered pattern avoidance in words [1, 4, 15, 16, 56]. In 1998, Burstein [15] proved that 123 and 132 are w-equivalent. In 2002, Burstein and Mansour [16] proved the w-equivalence of 121 and 112. By these two results we obtain that there are three w-equivalence classes of word patterns of length three:

- 111,
- 112, 121, 122, 211, 212, 221,

- 123, 132, 213, 231, 312, 321.

5.2 Compositions

A *composition* $\sigma = \sigma_1\sigma_2\cdots\sigma_m$ of $n \in \mathbb{N}$ is an ordered collection of one or more positive integers whose sum is n . The numbers $\sigma_1, \dots, \sigma_m$ are called *parts* of the composition. We let \mathcal{C}_n denote set of all compositions of n . A composition may thus be regarded as a k -ary word.

We again say that the composition $\sigma \in \mathcal{C}_n$ *contains* a pattern $\tau \in [\ell]^s$, if σ contains a subsequence order-isomorphic to τ . Let $\mathcal{C}_n(\tau)$ denote the set of all the compositions in \mathcal{C}_n that avoid τ . We say that two patterns τ and τ' are *c-equivalent*, if for all values of n , we have $|\mathcal{C}_n(\tau)| = |\mathcal{C}_n(\tau')|$. It is easy to see that every pattern is c-equivalent to its reversal. However, a pattern does not need to be c-equivalent to its complement.

Savage and Wilf [60] considered pattern avoidance in compositions for a single pattern $\tau \in \mathfrak{S}_3$ (recall that \mathfrak{S}_3 is the set of the permutations on three letters), and showed that the number of compositions of n avoiding $\tau \in \mathfrak{S}_3$ is independent of τ , that is, the three patterns 123, 213 and 132 are all c-equivalent. Recently, Heubach, Mansour and Munagi [26] showed that 112 is c-equivalent to 121, and 122 is c-equivalent to 212. These two results complete the classification of patterns of length three in compositions, and show they form exactly four c-equivalence classes:

- 123, 213, 132, 231, 312, 321,
- 112, 121, 211,
- 122, 212, 221,
- 111.

5.3 Strong equivalence of words

We now introduce an equivalence relation on words, which refines both the w-equivalence and the c-equivalence. For a word σ of length n , the *content* of σ is the unordered multiset of the n letters appearing in σ . In particular, two words have the same content, if one can be obtained from the other by a suitable rearrangement of letters.

We say that two patterns τ, τ' are *strongly equivalent*, denoted by $\tau \overset{\dagger}{\sim} \tau'$, if for every k, n there is a bijection f between $[k]^n(\tau)$ and $[k]^n(\tau')$ with the property that for every $\sigma \in [k]^n(\tau)$, the word $f(\sigma)$ has the same content as σ . Clearly, if two patterns are strongly equivalent, then they are also w-equivalent and c-equivalent. Each pattern is strongly equivalent to its reversal, and if two patterns τ and σ are strongly equivalent, then their complements τ^C and σ^C are strongly equivalent as well. Strong equivalence has already been considered (under different terminology) by Savage and Wilf [60], who proved that all permutation patterns of length 3 are strongly equivalent.

5.4 Semi-standard fillings of Ferrers shapes

Several families of strongly equivalent words may be deduced from known results on fillings of diagrams.

To establish the link between words and fillings, we will represent k -ary words of length n as 01-matrices with k rows and n columns and exactly one 1-cell in each column. For a

word σ of length n over the alphabet $[k]$, we let $M(\sigma, k)$ be the $k \times n$ matrix with a 1-cell in row i and column j if and only if the j -th letter of σ is equal to i .

With this representation, we may use known bijections on fillings of diagrams to obtain directly new equivalences among words. Recall that a semi-standard filling of a Ferrers shape is a 01-filling in which every column has exactly one 1-cell. We will say that two matrices M and M' are *Ferrers equivalent*, denoted by $M \stackrel{\text{fs}}{\sim} M'$, if for every Ferrers shape F the number of M -avoiding semi-standard fillings is equal to the number of M' -avoiding semi-standard fillings. We say that M and M' are *strongly Ferrers equivalent* if for every Ferrers shape F there is a bijection between M -avoiding and M' -avoiding semi-standard fillings of F that preserves the number of 1-cells in each row.

The following lemma allows us to translate results about fillings of Ferrers shapes into results about words. The lemma is based on the red-green argument that we have already encountered several times.

For a word $\rho \in [\ell]^n$ and an integer k , we let $\rho + k$ denote the word obtained by increasing each letter of ρ by k .

Lemma 36. *Let τ and τ' be two patterns with k letters, let ρ be a pattern with ℓ letters. If $M(\tau, k)$ and $M(\tau', k)$ are strongly Ferrers equivalent matrices then the two $(k + \ell)$ -letter patterns $\tau(\rho + k)$ and $\tau'(\rho + k)$ are strongly equivalent words. (Here $\tau(\rho + k)$ denotes the concatenation of τ and $\rho + k$.)*

Proof. Let us write $\sigma = \tau(\rho + k)$ and $\sigma' = \tau'(\rho + k)$. For a given m and n , choose a word $x \in [m]^n(\sigma)$, and let $M = M(x, m)$ be its corresponding matrix. Note that M avoids the matrix $M(\sigma, k + \ell)$.

Color the cells of M red and green, where a cell c is green if and only if the submatrix of M strictly to the right and strictly to the top of c contains $M(\rho, \ell)$, otherwise the cell is red. Note that the green cells form a Ferrers diagram and that the nonzero columns of this diagram induce an $M(\tau, k)$ -avoiding semi-standard filling. Using the strong Ferrers equivalence of $M(\tau, k)$ and $M(\tau', k)$, we may transform this filling into a $M(\tau', k)$ -avoiding filling. This operation transforms M into a matrix M' representing a σ' -avoiding word x' with the same content as x .

To see that this operation can be inverted, observe that the operation has only modified the filling of the green cells of M . Observe also that for every green cell c of M , there is a copy of $M(\rho, \ell)$ strictly to the right and strictly above c which only consists of red cells. Thus the red cells of M coincide with the red cells of M' .

We thus have a bijection showing that $\sigma \stackrel{\text{fs}}{\sim} \sigma'$. □

By Theorem 7, we know that the matrices I_k and J_k are strongly Ferrers equivalent. Applying Lemma 36, we thus obtain the following result.

Theorem 37 (J., Mansour [38]). *For any pattern ρ and any integer $k \in \mathbb{N}$, the word $12 \cdots k(\rho + k)$ is strongly equivalent to $k(k - 1) \cdots 1(\rho + k)$.*

5.5 Patterns equivalent to 12^k

From now on, we will often use the shorthand notation n^k to denote the word consisting of the symbol n repeated k times.

In this section, we will deal with a family of patterns that are strongly equivalent to the pattern 12^k . Our aim is to prove the following result.

Theorem 38 (J., Mansour [38]). *For any two integers i and j the matrix $M(2^i 12^j, 2)$ is strongly Ferrers equivalent to $M(12^{i+j}, 2)$.*

In view of Lemma 36, the theorem directly yields the following result.

Corollary 39. *For any pattern ρ , the words $2^i 12^j(\rho + 2)$ and $12^{i+j}(\rho + 2)$ are strongly equivalent.*

Rather than proving Theorem 38 directly, we shall prove a more refined result, which will become useful later. To state the refinement, we need additional terminology.

First of all, we will now work in the more general setting of stack shapes, instead of the Ferrers shapes considered above. The notions of Ferrers equivalence and strong Ferrers equivalence can be naturally extended to semi-standard fillings of stack shapes: we will say that two matrices M and M' are *stack equivalent*, denoted by $M \stackrel{s}{\sim} M'$, if they are equirestrictive with respect to semi-standard fillings of every stack polyomino. We will say that they are *strongly stack equivalent* if they are stack equivalent and the corresponding bijection preserves the number of 1-cells in each row.

Let Φ be a filling of a stack polyomino and let $t \geq 1$ be an integer. A sequence c_1, c_2, \dots, c_t of 1-cells in Φ is called a *decreasing chain* (or *increasing chain*) if for every $i \in [t - 1]$ the column containing c_i is to the left of the column containing c_{i+1} and the row containing c_i is above the row of c_{i+1} (or below the row of c_{i+1} , respectively).

A filling is *t-falling* (or *t-rising*) if it has at least t rows, and in its bottom t rows, the leftmost 1-cells of the nonzero rows form a decreasing chain (or increasing chain, respectively).

In the rest of this section, we let S_q^p denote the sequence $2^p 12^q$, where p, q are nonnegative integers.

Here is the promised refinement of Theorem 38.

Lemma 40. *For every $p, q \geq 0$, the matrix $M(S_q^p, 2)$ is strongly stack equivalent to the matrix $M(S_0^{p+q}, 2)$. Furthermore, if $p \geq 1$, then for every stack polyomino P , there is a bijection f between the $M(S_q^p, 2)$ -avoiding and $M(S_0^{p+q}, 2)$ -avoiding semi-standard fillings of P with the following properties.*

- *The bijection f preserves the number of 1-cells in every row.*
- *Both f and f^{-1} map t -falling fillings to t -falling fillings, for every $t \geq 1$.*

Proof. Let $M = M(S_q^p, 2)$ and $M' = M(S_0^{p+q}, 2)$, for some $p, q \geq 0$. Let P be a stack polyomino. We will proceed by induction over the number of rows of P . If P has only one row, then a constant mapping is the required bijection. Assume now that P has $r \geq 2$ rows, and assume that we are presented with a semi-standard filling Φ of P . Let P^- be the diagram obtained from P by erasing the r -th row as well as every column that contains a 1-cell of Φ in the r -th row. The filling Φ induces in P^- a semi-standard filling Φ^- .

We claim that for every $p, q \geq 0$, the filling Φ avoids M if and only if the following two conditions are satisfied.

- (a) The filling Φ^- avoids M .
- (b) If the r -th row of Φ contains m 1-cells in columns $c_1 < c_2 < \dots < c_m$ and if $m \geq p + q$, then for every i such that $p \leq i \leq m - q$, the column c_i is either the rightmost column of the r -th row of Π , or it is directly adjacent to the column c_{i+1} , i.e., $c_i + 1 = c_{i+1}$ (see Figure 5.1).

Clearly, the two conditions are necessary. We now show that they are sufficient. The first condition guarantees that Φ does not contain any copy of M that would be confined to the first $r - 1$ rows. The second condition guarantees that Φ has no copy of M that would intersect the r -th row.

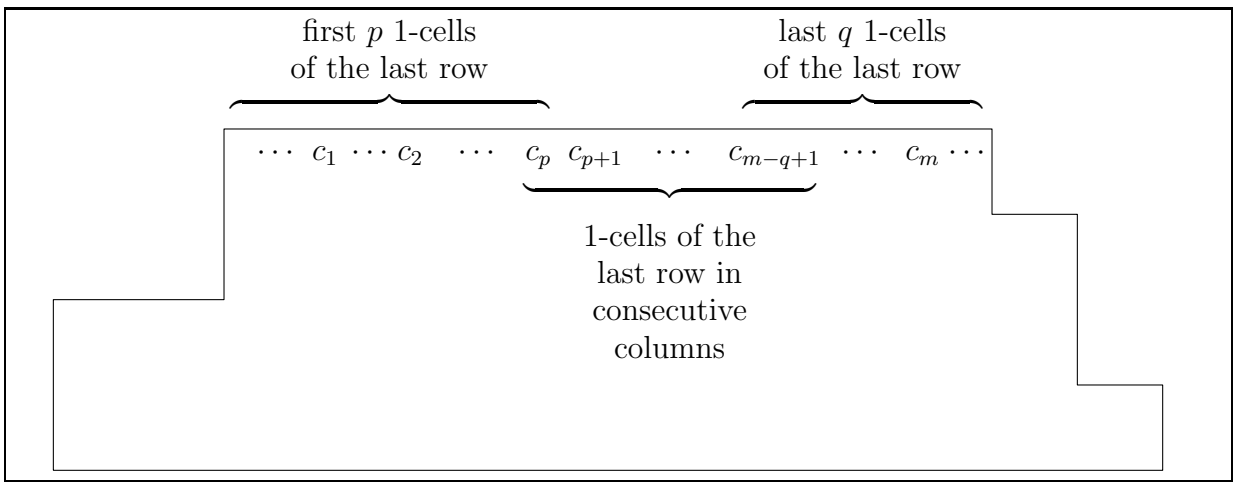


Figure 5.1: Illustration of condition (b) in the proof of Lemma 40.

We now define recursively the required bijection between M -avoiding and M' -avoiding fillings. Let Φ be an M -avoiding filling of P , let Φ^- and c_1, \dots, c_m be as above. By the induction hypothesis, we already have a bijection between M -avoiding and M' -avoiding fillings of the shape P^- . This bijection maps Φ^- to a filling Ψ^- of P^- . Let Ψ be the filling of P that has the same values as Φ in the r -th row, and the columns not containing a 1-cell in the r -th row are filled according to Ψ^- . Note that Ψ contains no copy of M' in its first $r - 1$ rows and it contains no copy of M that would intersect the r -th row.

If Ψ has fewer than $p + q$ 1-cells in the r -th row, we define $f(\Phi) = \Psi$, otherwise we modify Ψ in the following way. For every $i = 1, \dots, q$, we consider the columns with indices strictly between c_{m-q+i} and $c_{m-q+i+1}$ (if $i = q$, we take all columns to the right of c_m that intersect the last row). We remove these columns from Ψ and re-insert them between the columns c_{p+i-1} and c_{p+i} . Note that these transformations preserve the relative left-to-right order of all the columns that do not contain a 1-cell in their r -th row. In particular, the resulting filling still has no copy of M' in the first $r - 1$ rows. By construction, the filling also satisfies condition (b) for the values $p' = p + q$ and $q' = 0$ used instead of the original p and q . Hence, it is a M' -avoiding filling. This construction provides a bijection f between M -avoiding and M' -avoiding fillings.

It is clear that f preserves the number of 1-cells in each row. It remains to check that if $p \geq 1$, then f preserves the t -falling property. Let us fix t , and let r be the number of rows of P . If $r < t$ then no filling of P is t -falling. If $r = t$, then Φ is t -falling if and only if Φ^- is $(t - 1)$ -falling and the r -th row is either empty or has a 1-cell in the leftmost column of P . These conditions are preserved by f and f^{-1} , provided $p \geq 1$. Finally, if $r > t$, then Φ is t -falling if and only if Φ^- is t -falling. We now obtain the required result from the induction hypothesis and from the fact that the relative position of the 1-cells of the first $r - 1$ rows does not change when we transform Ψ into $f(\Phi)$. \square

Lemma 40 immediately implies Theorem 38, which in turn implies Corollary 39.

5.6 Patterns equivalent to 1232^k

Corollary 39 shows that all the words of length k that have a single symbol '1' and $k - 1$ symbols '2' are strongly equivalent. In other words, any pattern that can be obtained from 12^{k-1} by rearranging its symbols is strongly equivalent to it. We will now show that a similar property also holds for the pattern $12^{k-2}3$. This time, the proof is quite different, and does not use the notion of fillings.

Theorem 41 (J., Mansour [38]). *Let $k \geq 3$ be an integer. All the patterns of length k that consist of a single symbol ‘1’, a single symbol ‘3’ and $k - 2$ symbols ‘2’ are strongly equivalent.*

Proof. Let k be fixed. Let $\tau(i, j)$ denote the word of length k whose i -th symbol is ‘1’, the j -th symbol is ‘3’ and the remaining symbols are equal to ‘2’. Our aim is to show that all the patterns in the set $\{\tau(i, j), i \neq j, 1 \leq i, j \leq k\}$ are strongly equivalent. Since each word is strongly equivalent to its reversal, we only need to deal with the words $\tau(i, j)$ with $i < j$. From Corollary 39, we deduce that the words $\{\tau(i, k), i = 1, \dots, k - 1\}$ are all strongly equivalent, and hence the words $\{\tau(1, j), j = 2, \dots, k\}$ are all strongly equivalent as well.

To prove the theorem, it thus suffices to show that for every $i < j < k$, the word $\tau(i, j)$ is strongly equivalent to the word $\tau(i + 1, j + 1)$. Let m be an integer. We will say that a word σ *contains* $\tau(i, j)$ *at level* m if there is a pair of symbols ℓ, h such that $\ell < m < h$, and such that the word σ contains a subword over the alphabet $\{\ell, m, h\}$ which is order-isomorphic to $\tau(i, j)$. For example, the word 132342 contains the pattern 1223 at level 3 (due to the subword 1334), while it avoids 1223 at level 2 (since it does not contain the subword $\ell 22 h$ for any values of $\ell < 2 < h$).

Assume now that we are given a fixed pair of indices i, j , with $i < j < k$, and we want to provide a content-preserving bijection between $\tau(i, j)$ -avoiding and $\tau(i + 1, j + 1)$ -avoiding words of length n . We will say that a word σ is an m -*hybrid* if for every $\bar{m} < m$, the word σ avoids $\tau(i, j)$ at level \bar{m} , while for every $\tilde{m} \geq m$, σ avoids $\tau(i + 1, j + 1)$ at level \tilde{m} .

We will present, for any $m \geq 1$, a content-preserving bijection between m -hybrids and $(m + 1)$ -hybrids. By composing these bijections, we obtain the required bijection between $\tau(i, j)$ -avoiding and $\tau(i + 1, j + 1)$ -avoiding words.

Let $m \geq 1$ be fixed. Let σ be an arbitrary word. A letter of σ is called *low* if it is smaller than m , and a letter is called *high* if it is greater than m . A *low cluster* of σ is a maximal block of consecutive low symbols of σ . A *high cluster* is defined analogously. Thus, every symbol of σ different from m belongs to a unique cluster. The *landscape* of σ is a word over the alphabet $\{L, m, H\}$ obtained by replacing every low cluster of σ by a single symbol L, and every high cluster of σ by a single symbol H. For example, if $m = 3$, the landscape of the word 133212443 is the word L33LH3.

Note that σ contains $\tau(i, j)$ at level m if and only if the landscape of σ contains the subsequence $m^{i-1}Lm^{j-i-1}Hm^{k-j}$.

We will now describe the bijection between m -hybrids and $(m + 1)$ -hybrids. Let σ be an m -hybrid word, let X be its landscape. We split X into three parts $X = PmS$, where P is the prefix of X formed by all the symbols of X that appear before the first occurrence of m in X , and S is the suffix of all the symbols that appear after the first occurrence of m . Let us define a word X' by $X' = SmP$. Note that X' contains a subsequence $m^{i-1}Lm^{j-i-1}Hm^{k-j}$ if and only if X contains a subsequence $m^iLm^{j-i-1}Hm^{k-j-1}$. Thus, since X is a landscape of a word that avoids $\tau(i + 1, j + 1)$ at level m , we know that any word with landscape X' must avoid $\tau(i, j)$ at level m .

Let us define a word σ' by the following three rules.

1. The word σ' has landscape X' .
2. For any p , the p -th low cluster of σ' consists of the same sequence of symbols as the p -th low cluster of σ .
3. For any q , the q -th high cluster of σ' consists of the same sequence of symbols as the q -th high cluster of σ .

Clearly, there is a unique word σ' satisfying these properties. Note that the subsequence of all the low symbols of σ is the same as the subsequence of all the low symbols of σ' , and these sequences are partitioned into low clusters in the same way. An analogous property holds for the high symbols too.

We claim that σ' is an $(m + 1)$ -hybrid. We have already pointed out that σ' avoids $\tau(i, j)$ at level m . Let us now argue that σ' avoids $\tau(i, j)$ at level \bar{m} , for every $\bar{m} < m$. For contradiction, assume that σ' contains a subsequence $T = \bar{m}^{i-1}\ell\bar{m}^{j-i-1}h\bar{m}^{k-j}$, for some $\ell < \bar{m} < h$. If $h < m$, then all the symbols of T are low, and since σ has the same subsequence of low symbols as σ' , we know that σ also contains T as a subsequence, contradicting the assumption that σ is an m -hybrid.

Assume now that $h \geq m$. Let x and y be the two symbols adjacent to h in the sequence T (note that h is not the last symbol of T , so x and y are well defined). Both x and y are low, and they belong to distinct low clusters of σ' , because the symbol h is not low. Since the low symbols of σ are the same as the low symbols of σ' , and they are partitioned into clusters in the same way, we know that σ contains a subsequence $\bar{m}^{i-1}\ell\bar{m}^{j-i-1}h'\bar{m}^{k-j}$, where h' is a non-low symbol. This shows that σ contains $\tau(i, j)$ at level \bar{m} , which is impossible, because σ is an m -hybrid.

By an analogous argument, we may show that σ' avoids $\tau(i + 1, j + 1)$ at any level $\tilde{m} > m$. We conclude that the mapping described above transforms an m -hybrid σ into an $(m + 1)$ -hybrid σ' . It is clear that the mapping is reversible and provides the required bijection between m -hybrids and $(m + 1)$ -hybrids. \square

By computer enumeration [38], it has been verified that all the w-equivalence classes of patterns of length at most six and all the c-equivalence classes of patterns of length at most five can be described using the criteria given in this chapter.

Chapter 6

Partitions

In this final chapter of the main part of the thesis, we will deal with pattern avoidance of set partitions. Let us first recall the main notions related to this topic. A *partition of size n* is a collection B_1, B_2, \dots, B_d of nonempty disjoint sets, called *blocks*, whose union is the set $[n] = \{1, 2, \dots, n\}$. We will assume that B_1, B_2, \dots, B_d are listed in increasing order of their minimum elements, that is, $\min B_1 < \min B_2 < \dots < \min B_d$.

There are several possibilities to represent a set partition. For our purposes, we chose to represent a partition of size n by its *canonical sequence*, which is an integer sequence $\pi = \pi_1\pi_2\cdots\pi_n$ such that $\pi_i = k$ if and only if $i \in B_k$. For instance, 1231242 is the canonical sequence of the partition of $\{1, 2, \dots, 7\}$ with the four blocks $\{1, 4\}$, $\{2, 5, 7\}$, $\{3\}$ and $\{6\}$.

Note that a sequence π over the alphabet $[d]$ represents a partition with d blocks if and only if it has the following properties.

- Each number from the set $[d]$ appears at least once in π .
- For each i, j such that $1 \leq i < j \leq d$, the first occurrence of i precedes the first occurrence of j .

We remark that sequences satisfying these properties are also known as *restricted growth functions*. The idea of representing a set partition by a restricted growth function was first suggested by Hutchinson [31], as a basis for an efficient algorithm to generate all set partitions. The algorithmic aspects of restricted growth functions were later investigated by Williamson [72], and by Savage [59]. Simion [61, Section 3.4] mentions the connection between restricted growth functions and various combinatorial statistics of set partitions. Milne [51, 52, 53] and Wachs [71] used restricted growth functions as a tool in the study of combinatorial identities.

Throughout this chapter, we will identify a set partition with its corresponding canonical sequence. In this representation, the containment relation of set partitions can be regarded as a special case of the containment relation of k -ary words, which we considered in Chapter 5.

Let \mathcal{P}_n denote the set of all the partitions of $[n]$, let $\mathcal{P}_n(\sigma)$ denote the set of all partitions of $[n]$ that avoid σ , and let p_n and $p_n(\sigma)$ denote the cardinality of \mathcal{P}_n and $\mathcal{P}_n(\sigma)$, respectively. We say that two partitions σ and σ' are *equivalent*, denoted by $\sigma \approx \sigma'$, if $p_n(\sigma) = p_n(\sigma')$ for each n .

The concept of pattern avoidance based on restricted growth functions has been introduced by Sagan [58], who considered, among other topics, the enumeration of partitions avoiding patterns of size three. In this thesis, we extend this study to larger patterns. In particular, we will present recent results of Jelínek and Mansour [37], which yield new families of equivalent partition patterns. By computer enumeration, it has been verified

that the criteria described here cover all the equivalence classes of patterns of size $n \leq 7$. As usual, we omit the enumeration data in this thesis; however, they are available in the original paper [37] and references therein.

Let us remark, that there are several alternative ways to define containment for set partitions. For instance, Chen et al. [17, 18] have used the path-representation of set partitions. In this setting, they have obtained, among other results, an identity between k -noncrossing and k -nonnesting partitions. It was later pointed out by Krattenthaler [48] that this identity is a consequence of more general identities between diagonal-avoiding fillings of diagrams.

Other possible representations of set partitions have been considered by Klazar [42, 44], and by Goyt [28]. However, we are not aware of any attempt at systematic Wilf-type classification in these settings.

6.1 Basic facts and previous results

Let us first introduce several notational conventions that will be applied throughout the rest of this thesis. For a finite sequence $S = s_1 s_2 \cdots s_p$ and an integer k , we let $S + k$ denote the sequence $(s_1 + k)(s_2 + k) \cdots (s_p + k)$. For a symbol k and an integer d , the constant sequence (k, k, \dots, k) of length d is denoted by k^d . To prevent confusion, we will use capital letters S, T, \dots to denote arbitrary sequences of positive integers, and we will use lowercase greek symbols $(\pi, \sigma, \tau, \dots)$ to denote canonical sequences representing partitions.

Occasionally, it will be convenient to represent an infinite sequence $(a_n)_{n=0}^\infty$ by its *exponential generating function* (or EGF for short), which is the formal power series $F(x) = \sum_{n \geq 0} \frac{a_n x^n}{n!}$. We will deal with the generating functions of the sequences of the form $(p_n(\pi))_{n \geq 0}$, where π is a given pattern. We simply call such a generating function *the EGF of the pattern π* .

Let us summarize previous results relevant to our topic. Let $\exp(x) = \sum_{n \geq 0} \frac{x^n}{n!}$ and $\exp_{<k}(x) = \sum_{n=0}^{k-1} \frac{x^n}{n!}$. We first state two simple propositions, which already appear in Sagan's paper on pattern-avoiding partitions [58].

Proposition 42. *A partition avoids the pattern 1^k if and only if each of its blocks has size less than k . The EGF of the pattern 1^k is equal to*

$$\exp(\exp_{<k}(x) - 1). \tag{6.1}$$

Proposition 43. *A partition avoids the pattern $12 \cdots k$ if and only if it has fewer than k blocks. The corresponding EGF is equal to*

$$\exp_{<k}(\exp(x) - 1). \tag{6.2}$$

We omit the proofs of these two propositions. Let us just remark that the formulas given above are obtained by standard manipulation of EGFs. A common generalization of these formulas can be found, e.g., in Flajolet and Sedgewick's book [24, Proposition II.2].

The enumeration of partitions with fewer than k blocks is closely related to the Stirling numbers of the second kind $S(n, m)$, defined as the number of partitions of $[n]$ with exactly m blocks (see sequence A008277 in the OEIS [68]).

Sagan [58] has described and enumerated the pattern-avoiding classes $\mathcal{P}_n(\pi)$ for the five patterns π of length three. We summarize the relevant results in Table 6.1. We again omit the proofs.

τ	$p_n(\tau)$
111	sequence A000085 in [68]
112, 121, 122, 123	2^{n-1}

Table 6.1: Number of partitions in $\mathcal{P}_n(\tau)$, where $\tau \in \mathcal{P}_3$.

6.2 General classes of equivalent patterns

Most of our results on \approx -equivalent partitions yield infinite families of \approx -equivalent pairs of patterns. We thus begin by presenting these ‘general’ results and then, in Sections 6.5 and 6.6, we deal with two ‘sporadic’ cases of equivalent pairs, which are necessary to complete our classification of small patterns.

Pattern-avoiding fillings of diagrams. Since pattern avoidance in partitions is just a special case of pattern avoidance in words, it should be no surprise that many of our results may be reduced to results on pattern-avoiding fillings of shapes. Let us thus begin by explaining the relationship between fillings and canonical sequences.

Recall that an F^\perp -shape is a bottom-right aligned reflected copy of a Ferrers shape. We will say that two 01-matrices M and M' are F^\perp -equivalent, denoted by $M \overset{s}{\sim} M'$, if they are equirestrictive with respect to semi-standard fillings of F^\perp -shapes. Let us also recall that $r(F)$ and $c(F)$ denote, respectively, the number of rows and columns of a diagram F .

Since our next arguments mostly deal with semi-standard fillings, we will drop the adjective ‘semi-standard’ and simply use the term ‘filling’, when there is no risk of ambiguity.

The following argument, which is similar to Lemma 2, explains the close link between semi-standard fillings and semi-sparse fillings. We state it here as a remark, so that we may refer to it later. The proof is essentially the same as the proof of Lemma 2, and we only sketch it here.

Remark 44. Let M and M' be two F^\perp -equivalent 01-matrices with a 1-cell in every column, and let f be a bijection between M -avoiding and M' -avoiding semi-standard fillings of F^\perp -shapes. There is a natural way to extend f into a bijection between M -avoiding and M' -avoiding sparse fillings of F^\perp -shapes. Assume that Φ is a sparse M -avoiding filling of an F^\perp -shape F . The non-zero columns of Φ form a semi-standard filling of a (not necessarily contiguous) subdiagram of F . We apply f to this subfilling to transform Φ into a sparse M' -avoiding filling of F .

A completely analogous argument can be made for stack polyominoes instead of F^\perp -shapes.

Let $S = s_1 s_2 \cdots s_m$ be a sequence of positive integers, and let $k \geq \max\{s_i : i \in [m]\}$ be an integer. Recall that $M(S, k)$ denotes the 01-matrix with k rows and m columns which has a 1-cell in row i and column j if and only if $s_j = i$.

We now describe the correspondence between partitions and fillings of F^\perp -shapes (recall that $\tau + k$ denotes the sequence obtained from τ by adding k to every element). Although the correspondence is based on a routine red-green argument, there are several technical difficulties that need to be addressed. For this reason, we state the following lemma with full proof.

Lemma 45. *Let S and S' be two nonempty sequences over the alphabet $[k]$, let τ be an arbitrary partition. If $M(S, k)$ is F^\perp -equivalent to $M(S', k)$ then the partition pattern*

$\sigma = 12 \cdots k(\tau + k)S$ is \approx -equivalent to $\sigma' = 12 \cdots k(\tau + k)S'$.

Proof. Let π be a partition of $[n]$ with m blocks. Let M denote the matrix $M(\pi, m)$. Fix a partition τ with t blocks, and let T denote the matrix $M(\tau, t)$. We will color the cells of M red and green. If τ is nonempty, then the cell in row i and column j is colored green if and only if the submatrix of M induced by the rows $i + 1, \dots, m$ and columns $1, \dots, j - 1$ contains T . If τ is empty, then the cell in row i and column j is green if and only if row i has at least one 1-cell strictly to the left of column j . A cell is red if it is not green.

Note that the green cells form an F^\perp -shape, and the entries of the matrix M form a semi-sparse filling Φ_G of this shape. Also, note that the leftmost 1-cell of each row is always red, and any 0-cell of the same row to the left of the leftmost 1-cell is red too.

It is not difficult to see that the partition π avoids σ if and only if the filling Φ_G of the ‘green’ diagram avoids $M(S, k)$, and π avoids σ' if and only if Φ_G avoids $M(S', k)$. Since $M(S, k) \stackrel{s}{\approx} M(S', k)$, there is a bijection f that maps $M(S, k)$ -avoiding fillings of an F^\perp -shape onto $M(S', k)$ -avoiding fillings of the same shape. By Remark 44, f can be extended to semi-sparse fillings. Using this extension of f , we construct the following bijection between $\mathcal{P}_n(\sigma)$ and $\mathcal{P}_n(\sigma')$: for a partition $\pi \in \mathcal{P}_n(\sigma)$ with m blocks, we take M and Φ_G as above. By assumption, Φ_G is $M(S, k)$ -avoiding. Using the bijection f and Remark 44, we transform Φ_G into an $M(S', k)$ -avoiding semi-sparse filling $f(\Phi_G) = \Psi_G$, while the filling of the red cells of M remains the same. We thus obtain a new matrix M' .

Note that if we color the cells of M' red and green using the criterion described in the first paragraph of this proof, then each cell of M' will receive the same color as the corresponding cell of M , even though the occurrences of T in M' need not correspond exactly to the occurrences of T in M . Indeed, if τ is nonempty, then for each green cell g of M , there is an occurrence of T to the left and above g consisting entirely of red cells. This occurrence is contained in M' as well, which guarantees that the cell g remains green in M' . A similar argument can be made if τ is empty.

By construction, M' has exactly one 1-cell in each column, hence there is a sequence π' over the alphabet $[m]$ such that $M' = M(\pi', m)$. We claim that π' is a canonical sequence of a partition. To see this, note that for every $i \in [m]$, the leftmost 1-cell of M in row i is red and the preceding 0-cells in row i are red too. It follows that the leftmost 1-cell of row i in M is also the leftmost 1-cell of row i in M' . Thus, the first occurrence of the symbol i in π appears at the same place as the first occurrence of i in π' , hence π' is indeed a canonical sequence. The green cells of M' avoid $M(S', k)$, so π' avoids σ' . Obviously, the transform $\pi \mapsto \pi'$ is invertible and provides a bijection between $\mathcal{P}_n(\sigma)$ and $\mathcal{P}_n(\sigma')$. \square

In general, the relation $12 \dots kS \approx 12 \dots kS'$ does not imply that $M(S, k)$ and $M(S', k)$ are F^\perp -equivalent. In Section 6.6, we will prove that $12112 \approx 12212$, even though $M(112, 2)$ is not F^\perp -equivalent to $M(212, 2)$.

On the other hand, the relation $12 \dots kS \approx 12 \dots kS'$ allows us to establish a somewhat weaker equivalence between pattern-avoiding fillings, using the following lemma.

Lemma 46. *Let S be a nonempty sequence over the alphabet $[k]$, and let $\tau = 12 \dots kS$. For every n and m , there is a bijection f that maps the set of τ -avoiding partitions of $[n]$ with m blocks onto the set of the $M(S, k)$ -avoiding semi-standard fillings Φ of F^\perp -shapes with $n - m$ columns and at most m rows.*

Proof. Let π be a τ -avoiding partition of $[n]$ with m blocks. Let $M = M(\pi, m)$, and let us consider the same red and green coloring of M as in the proof of Lemma 45, i.e., the green cells of a row i are precisely the cells that are strictly to the right of the leftmost 1-cell in row i .

Note that M has exactly m red 1-cells, and each 1-cell is red if and only if it is the leftmost 1-cell of its row. Note also that if c_i is the column containing the red 1-cell in row i , then either c_i is the rightmost column of M , or column $c_i + 1$ is the leftmost column of M with exactly i green cells.

Let Φ_G be the filling formed by the green cells. As was pointed out in the previous proof, the filling Φ_G is a semi-sparse $M(S, k)$ -avoiding filling of an F^\perp -shape. Note that for each $i = 1, \dots, m - 1$, the filling Φ_G has exactly one zero column of height i , and this column, which corresponds to c_{i+1} , is the rightmost of all the columns of Φ_G with height at most i .

Let Φ_G^- be the subfilling of Φ_G induced by all the nonzero columns of Φ_G . Observe that Φ_G^- is a semi-standard $M(S, k)$ -avoiding filling of an F^\perp -shape with exactly $n - m$ columns and at most m rows; we thus define $f(\pi) = \Phi_G^-$.

Let us now show that the mapping f defined above can be inverted. Let Ψ be an $M(S, k)$ -avoiding filling of an F^\perp -shape with $n - m$ columns and at most m rows. We insert $m - 1$ zero columns c_2, c_3, \dots, c_m into the filling Ψ as follows: each column c_i has height $i - 1$, and it is inserted immediately after the rightmost column of $\Psi \cup \{c_2, \dots, c_{i-1}\}$ that has height at most $i - 1$. Note that the filling obtained by this operation corresponds to the green cells of the original matrix M . Let us call this semi-sparse filling Ψ_G .

We now add a new 1-cell on top of each zero column of Ψ_G , and we add a new 1-cell in front of the bottom row, to obtain a semi-standard filling of a diagram with n columns and m rows. The diagram can be completed into a matrix $M = M(\pi, m)$, where π is easily seen to be a canonical sequence of a τ -avoiding partition. \square

Lemma 45 provides a tool to deal with partition patterns of the form $12 \cdots k(\tau + k)S$ where S is a sequence over $[k]$ and τ is a partition. We now describe a similar correspondence between partitions and fillings of stack polyominoes, which will be useful for dealing with patterns of the form $12 \cdots kS(\tau + k)$. We use a similar argument as in the proof of Lemma 45.

Lemma 47. *If τ is a partition, and S and S' are two nonempty sequences over the alphabet $[k]$ such that $M(S, k) \stackrel{\Delta^s}{\sim} M(S', k)$, then the partition $\sigma = 12 \cdots kS(\tau + k)$ is equivalent to the partition $\sigma' = 12 \cdots kS'(\tau + k)$.*

Proof. Fix a partition τ with t blocks. Let π be any partition of $[n]$ with m blocks, let $M = M(\pi, m)$. We will color the cells of M red and green. A cell of M in row i and column j is green, if it satisfies the following conditions.

- (a) The submatrix of M formed by the intersection of the top $m - i$ rows and the rightmost $n - j$ columns contains $M(\tau, t)$.
- (b) The matrix M has at least one 1-cell in row i appearing strictly to the left of column j .

A cell is red if it is not green. Note that the green cells form a stack polyomino and the matrix M induces a semi-sparse filling Φ_G of this polyomino.

As in Lemma 45, it is easy to verify that the partition π avoids the pattern σ if and only if the filling Φ_G avoids $M(S, k)$, and π avoids σ' if and only if Φ_G avoids $M(S', k)$.

The rest of the argument is analogous to the proof of Lemma 45. Assume that $M(S, k)$ and $M(S', k)$ are stack equivalent via a bijection f . By Remark 44, we extend f to a bijection between $M(S, k)$ -avoiding and $M(S', k)$ -avoiding semi-sparse fillings of a given stack polyomino. Consider a partition $\pi \in \mathcal{P}_n(\sigma)$ with m blocks, and define M and Φ_G as above. Apply f to the filling Φ_G to obtain an $M(S', k)$ -avoiding filling Ψ_G ; the filling

of the red cells of M remains the same. This yields a matrix M' and a sequence π' such that $M' = M(\pi', k)$. We may easily check that the green cells of M' are the same as the green cells of M . By rule (b) above, the leftmost 1-cell of each row of M is unaffected by this transform. It follows that the first occurrence of i in π' is at the same place as the first occurrence of i in π , and in particular, π' is a partition. By the observation of the previous paragraph, π' avoids σ' and the transform $\pi \mapsto \pi'$ is a bijection from $\mathcal{P}_n(\sigma)$ to $\mathcal{P}_n(\sigma')$. \square

The following simple result about pattern avoidance in fillings will turn out to be useful in the analysis of pattern avoidance in partitions.

Proposition 48. *If S is a nonempty sequence over the alphabet $[k - 1]$, then $M(S, k)$ is stack equivalent to $M(S + 1, k)$. If S and S' are two sequences over $[k - 1]$ such that $M(S, k - 1) \stackrel{\text{js}}{\approx} M(S', k - 1)$ then $M(S, k) \stackrel{\text{js}}{\approx} M(S', k)$, and if $M(S, k - 1) \stackrel{\Delta^s}{\approx} M(S', k - 1)$ then $M(S, k) \stackrel{\Delta^s}{\approx} M(S', k)$.*

Proof. To prove the first part, let us define $M = M(S, k)$, $M^- = M(S, k - 1)$, and $M' = M(S + 1, k)$. Notice that a filling Φ of a stack polyomino P avoids M if and only if the filling obtained by erasing the topmost cell of every column of Φ avoids M^- . Similarly, Φ avoids M' , if and only if the filling obtained by erasing the bottom row of Φ avoids M^- . We will now describe a bijection between M -avoiding and M' -avoiding fillings. Fix an M -avoiding filling Φ . In every column of this filling, move the topmost element into the bottom row, and move every other element into the row directly above it. This yields an M' -avoiding filling. The second claim of the theorem is proved analogously. \square

Note that a sequence S over the alphabet $[k - 1]$ does not necessarily contain all the symbols $\{1, \dots, k - 1\}$. In particular, every sequence over $[k - 2]$ is also a sequence over $[k - 1]$. Thus, if S is a sequence over $[k - 2]$, we may use Proposition 48 to deduce $M(S, k) \stackrel{\Delta^s}{\approx} M(S + 1, k) \stackrel{\Delta^s}{\approx} M(S + 2, k)$.

For convenience, we translate the first part of Proposition 48 into the language of pattern-avoiding partitions, using Lemma 45 and Lemma 47. We omit the straightforward proof.

Corollary 49. *If S is a nonempty sequence over $[k - 1]$ and τ is an arbitrary partition, then*

$$12 \cdots k(\tau + k)S \approx 12 \cdots k(\tau + k)(S + 1) \text{ and } 12 \cdots kS(\tau + k) \approx 12 \cdots k(S + 1)(\tau + k). \quad \square$$

We now state another result related to pattern avoidance in F^\perp -shapes, which has important consequences for our study of partitions. Recall that for two matrices A and B , let $\begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix}$ denote the matrix with $r(A) + r(B)$ rows and $c(A) + c(B)$ columns with a copy of A in the top left corner and a copy of B in the bottom right corner.

The following lemma is analogous to Proposition 1. We omit the straightforward proof.

Lemma 50. *If A and A' are two F^\perp -equivalent matrices, and if B is an arbitrary matrix, then $\begin{pmatrix} B & 0 \\ 0 & A \end{pmatrix} \stackrel{\text{js}}{\approx} \begin{pmatrix} B & 0 \\ 0 & A' \end{pmatrix}$.*

With the help of Lemma 50, we may easily prove the following proposition.

Proposition 51. *Let $s_1 > s_2 > \cdots > s_m$ and $t_1 > t_2 > \cdots > t_m$ be two strictly decreasing sequences over the alphabet $[k]$, let r_1, \dots, r_m be positive integers. Define weakly decreasing sequences $S = s_1^{r_1} s_2^{r_2} \cdots s_m^{r_m}$ and $T = t_1^{r_1} t_2^{r_2} \cdots t_m^{r_m}$. We have $M(S, k) \stackrel{\text{js}}{\approx} M(T, k)$, and in particular, if τ an arbitrary partition, then $12 \cdots k(\tau + k)S \approx 12 \cdots k(\tau + k)T$.*

Proof. We proceed by induction over minimum j such that $s_i = t_i$ for each $i \leq m - j$. For $j = 0$, we have $S = T$ and the result is clear. If $j > 0$, assume without loss of generality that $s_{m-j+1} - t_{m-j+1} = d > 0$. Consider the sequence $t'_1 > t'_2 > \dots > t'_m$ such that $t'_i = t_i$ for every $i \leq m - j$ and $t'_i = t_i + d$ for every $i > m - j$.

The sequence $(t'_i)_{i=1}^m$ is strictly decreasing, and its first $m - j + 1$ terms are equal to the corresponding term of $(s_i)_{i=1}^m$. Define $T' = (t'_1)^{r_1}(t'_2)^{r_2} \dots (t'_m)^{r_m}$. By induction, $M(S, k) \stackrel{\mathcal{S}}{\sim} M(T', k)$. To prove that $M(T, k) \stackrel{\mathcal{S}}{\sim} M(T', k)$, first write $T = T_0T_1$, where T_0 is the prefix of T containing all the symbols of T greater than t_{m-j+1} and T_1 is the suffix of the remaining symbols. Notice that $T' = T_0(T_1 + d)$. We may write $M(T, k) = \begin{pmatrix} B & 0 \\ 0 & A \end{pmatrix}$ and $M(T', k) = \begin{pmatrix} B & 0 \\ 0 & A' \end{pmatrix}$, where $A = M(T_1, t_{m-j} - 1)$ and $A' = M(T_1 + d, t_{m-j} - 1)$. By Proposition 48, $A \stackrel{\mathcal{S}}{\sim} A'$, and by Lemma 50, $M(T, k) \stackrel{\mathcal{S}}{\sim} M(T', k)$, as claimed. The last claim of the proposition follows from Lemma 45. \square

Notice that Proposition 51 implies that for any $m \in \mathbb{N}$ and any decreasing sequence $m \geq a_1 > a_2 > \dots > a_k \geq 1$, the partition $12 \dots ma_1 a_2 \dots a_k$ is equivalent to $12 \dots mk(k - 1) \dots 1$. In particular, there is a set of (at least) $\binom{m}{k}$ equivalent patterns of size $m + k$. By choosing $k = \lfloor m/2 \rfloor$, we obtain an exponentially large class of equivalent patterns. In the whole realm of pattern avoidance of ordered structures, we are not aware of any other exponentially large Wilf-type equivalence class.

Non-crossing and non-nesting partitions. The main application of the framework we have developed above is the identity between non-crossing and non-nesting partitions. This identity is a natural consequence of the often-used identity between I_k -avoiding and J_k -avoiding fillings.

We define non-crossing and non-nesting partitions in the following way.

Definition 52. A partition is *k-noncrossing* if it avoids the pattern $12 \dots k12 \dots k$, and it is *k-nonnesting* if it avoids the pattern $12 \dots kk(k - 1) \dots 1$.

Let us point out that there are several different concepts of ‘crossings’ and ‘nestings’ of set partitions used in the literature: for example, Klazar [42] has considered two blocks X, Y of a partition to be crossing (or nesting) if there are four elements $x_1 < y_1 < x_2 < y_2$ (or $x_1 < y_1 < y_2 < x_2$, respectively) such that $x_1, x_2 \in X$ and $y_1, y_2 \in Y$, and similarly for k -crossings and k -nestings. Unlike our approach, Klazar’s definition makes no assumption about the relative order of the minimal elements of X and Y , which allows more general configurations to be considered as crossing or nesting. Thus, Klazar’s k -noncrossing and k -nonnesting partitions are a proper subset of our k -noncrossing and k -nonnesting partitions, (except for 2-noncrossing partitions where the two concepts coincide).

Another approach to crossings in partitions has been studied by Chen et al. [17, 18]. This approach uses the path-representation of a partition, where a partition of $[n]$ with blocks B_1, B_2, \dots, B_k is represented by a graph on the vertex set $[n]$, with $a, b \in [n]$ connected by an edge if they belong to the same block and there is no other element of this block between them. In this terminology, a partition is k -crossing (or k -nesting) if the representing graph contains k edges which are pairwise crossing (or nesting), where two edges $e_1 = \{a < b\}$ and $e_2 = \{a' < b'\}$ are crossing (or nesting) if $a < a' < b < b'$ (or $a < a' < b' < b$, respectively). Let us call such partitions graph- k -crossing and graph- k -nesting, to avoid confusion with our own terminology of Definition 52. It is not difficult to see that a partition is graph-2-noncrossing if and only if it is 2-noncrossing, but for nestings and for k -crossings with $k > 2$, the two concepts are incomparable. For instance the partition 12121 is graph-2-nonnesting but it contains 1221, while 12112 is graph-2-nesting and avoids 1221. Similarly, 1213123 has no graph-3-crossing and contains 123123, while 1232132 has a graph-3-crossing and avoids 123123 (see Fig. 6.1).




Figure 6.1: Comparison of path-representation and canonical function of a partition.

Chen et al. [18] have shown that the number of graph- k -noncrossing and graph- k -nonnesting partitions of $[n]$ is equal. Below, we prove that the same is also true for k -noncrossing and k -nonnesting partitions. It is interesting to note that the proofs of both these results are based on a reduction to theorems on pattern avoidance in the fillings of Ferrers diagrams, in particular Theorem 7 of Krattenthaler (this is only implicit in [18], a direct construction is given by Krattenthaler [48]), although the constructions employed in the proofs of these results are different.

From Theorem 7, we may easily deduce that $I_k \overset{s}{\approx} J_k$. Furthermore, Theorem 9 of Rubey implies that $I_k \overset{\Delta s}{\approx} J_k$. This is not quite as straightforward, since Rubey's theorem deals with integer fillings of moon polyominoes with prescribed row-sums. However, since a transposed copy of a stack polyomino is a special case of a moon polyomino, Rubey's general result applies to fillings of stack polyominoes with prescribed column-sums as well. In particular, it yields a bijection between I_k -avoiding and J_k -avoiding semi-standard fillings of an arbitrary stack polyomino. Combining these results with Lemma 47, we obtain the following result.

Theorem 53 (J., Mansour [37]). *Let τ be a partition and let k be an integer. We have the following identities:*

$$12 \cdots k(\tau + k)12 \cdots k \approx 12 \cdots k(\tau + k)k(k - 1) \cdots 21$$

and

$$12 \cdots k12 \cdots k(\tau + k) \approx 12 \cdots kk(k - 1) \cdots 21(\tau + k).$$

In particular, the number of non-crossing partitions of size n is equal to the number of non-nesting partitions of size n .

The patterns $12 \cdots k(k+1)12 \cdots k$ and $12 \cdots k12 \cdots k(k+1)$. There is one more result on pattern-avoiding partitions which can be proved using the identities between diagonal-avoiding polyomino fillings—it is the equivalence of the pattern $12 \cdots k(k+1)12 \cdots k$ and the pattern $12 \cdots k12 \cdots k(k+1)$. This time, however, the reduction is more involved than the routine arguments we used to prove the identity between non-crossing and non-nesting partitions. Before we prove this identity, we need some preparation.

Let P be a stack polyomino. Recall that the *content* of P is the multiset of column-heights of P . We will represent the content by the sequence of the column heights of P listed in nondecreasing order.

The key ingredient of our proof is the following result of Rubey [57].

Theorem 54 (Rubey [57], adapted). *Let P and P' be two stack polyominoes with the same content, and let $k \geq 1$ be an integer. There is a bijection between the I_k -avoiding semi-standard fillings of P and the I_k -avoiding semi-standard fillings of P' .*

The theorem above is essentially a special case of Proposition 5.3 from Rubey's paper [57]. The only complication is that Rubey's proposition deals with arbitrary non-negative integer fillings, rather than semi-standard fillings. However, as was pointed out

in the last paragraph of Section 4 in [57], it is easy to see that Rubey's bijection maps semi-standard fillings to semi-standard fillings.

Let us now analyze in more detail the partitions avoiding $12 \cdots k(k+1)12 \cdots k$.

Definition 55. Let $\pi = \pi_1 \cdots \pi_n$ be a partition. We say that an element π_i is *left-dominating* if $\pi_i \geq \pi_j$ for each $j < i$. We say that a left-dominating element π_i *left-dominates* an element π_j , if $\pi_i > \pi_j$, $i < j$, and π_i is the rightmost left-dominating element with these two properties. Clearly, if π_j not left-dominating, then it is left-dominated by a unique left-dominating element. On the other hand, a left-dominating element is not left-dominated by any other element. If an element is not left-dominating, we call it simply *left-dominated*.

The *left shadow* of π is the sequence $\bar{\pi}$ obtained by replacing each left-dominated element by the symbol '*'. We will say that a non-star symbol i *left-dominates* an occurrence of a star, if i is the rightmost non-star to the left of the star.

For example, if $\pi = 123232144$, the left shadow of π is the sequence $\bar{\pi} = 123*3**44$. In $\bar{\pi}$, the leftmost occurrence of '3' left-dominates a single star, while the second occurrence of '3' left-dominates two stars.

It is not difficult to see that a sequence $\bar{\pi}$ over the alphabet $\{1, 2, \dots, m, *\}$ is a left shadow of a partition with m blocks if and only if it satisfies the following conditions.

- The non-star symbols of $\bar{\pi}$ form a non-decreasing sequence.
- Each of the symbols $1, 2, \dots, m$ appears at least once.
- No occurrence of the symbol 1 may left-dominate an occurrence of *. Any other non-star symbol may left-dominate any number of stars, and each star is dominated by a non-star.

Any sequence that satisfies these three conditions will be called a *left-shadow sequence*. Note that a left-shadow sequence is uniquely determined by the multiplicities of its non-star symbols and by the number of stars dominated by each non-star.

Definition 56. Let $\pi = \pi_1 \cdots \pi_n$ be a partition, let $\Phi = \Phi(\pi)$ be the semi-standard filling of an F^\perp -shape defined by the following conditions.

1. The columns of Φ correspond to the left-dominated elements of π . The i -th column of Φ has height j if the i -th left-dominated element of π is dominated by an occurrence of $j + 1$.
2. The i -th column of Φ has a 1-cell in row j if the i -th left-dominated element of π is equal to j .

Note that the shape of the underlying diagram of $\Phi(\pi)$ is determined by the left shadow of π . More precisely, the number of columns of height h in Φ is equal to the number of stars in the left shadow which are dominated by an occurrence of $h + 1$. It is easy to see that the left shadow $\bar{\pi}$ and the filling $\Phi(\pi)$ together uniquely determine the partition π . In fact, for every semi-standard filling Φ' with the same shape as $\Phi(\pi)$, there is a (unique) partition π' with the same left-shadow as π , and with $\Phi(\pi') = \Phi'$.

The following observation is a straightforward application of the terminology introduced above. We omit its proof.

Observation 57. A partition π avoids the pattern $12 \cdots k(k+1)12 \cdots k$ if and only if the filling $\Phi(\pi)$ avoids I_k . □

We now focus on the partitions that avoid the pattern $12 \cdots k 12 \cdots k(k+1)$.

Definition 58. Let $\pi = \pi_1 \cdots \pi_n$ be a partition. We say that an element π_i is *right-dominating* if either $\pi_i \geq \pi_j$ for each $j > i$ or $\pi_i > \pi_j$ for each $j < i$. If π_i is not right-dominating, we say that it is *right-dominated*. We say that π_i *right-dominates* π_j if π_i is the leftmost right-dominating element appearing to the right of π_j , and π_j itself is not right-dominating.

The *right shadow* $\tilde{\pi}$ of a partition π is obtained by replacing each right-dominated element of π by a star.

For example, the right shadow of the partition $\pi = 12213423312$ is the sequence $12 * * 34 * 33 * 2$. A sequence $\tilde{\pi}$ over the alphabet $\{1, 2, \dots, m, *\}$ is the right shadow of a partition with m blocks if and only if it satisfies the following conditions.

- The non-star symbols of $\tilde{\pi}$ form a sequence $(1, 2, \dots, m, s_1, s_2, \dots, s_p)$ where the sequence $s_1 s_2 \cdots s_p$ is nonincreasing.
- No occurrence of the symbol 1 may right-dominate an occurrence of *. Any other non-star symbol may right-dominate any number of stars, and each star is right-dominated by a non-star.

Any sequence that satisfies these two conditions will be called a *right-shadow sequence*. A right-shadow sequence is uniquely determined by the multiplicities of its non-star symbols and by the number of stars right-dominated by each non-star.

Definition 59. Let $\pi = \pi_1 \cdots \pi_n$ be a partition. Let $\Psi = \Psi(\pi)$ be the semi-standard filling of a stack polyomino defined by the following conditions.

1. The columns of Ψ correspond to the right-dominated elements of π . The i -th column of Ψ has height j if the i -th right-dominated element of π is dominated by an occurrence of $j+1$.
2. The i -th column of Ψ has a 1-cell in row j if the i -th right-dominated element of π is equal to j .

Let S be the underlying diagram of $\Psi(\pi)$. Notice that S is uniquely determined by the right shadow $\tilde{\pi}$ of the partition π , although there may be different right shadows corresponding to the same shape S . The sequence $\tilde{\pi}$ and the filling $\Psi(\pi)$ together determine the partition π . For a fixed $\tilde{\pi}$, the mapping $\pi \mapsto \Psi(\pi)$ gives a bijection between partitions with right shadow $\tilde{\pi}$ and fillings of S .

The proof of the following observation is again straightforward and we omit it.

Observation 60. *A partition π avoids the pattern $12 \cdots k 12 \cdots k(k+1)$ if and only if the filling $\Psi(\pi)$ avoids I_k . \square*

We are now ready to prove the main result of this section.

Theorem 61 (J., Mansour [37]). *For any $k \geq 1$, the patterns $12 \cdots k(k+1)12 \cdots k$ and $12 \cdots k 12 \cdots k(k+1)$ are \approx -equivalent.*

Proof. We will describe a bijection between the two pattern-avoiding classes. Let π be a partition with m blocks that avoids $12 \cdots k(k+1)12 \cdots k$. Let $\bar{\pi}$ be its left shadow, and let $\Phi(\pi)$ be the filling from Definition 56. Let F denote the underlying shape of $\Phi(\pi)$. By Observation 57, $\Phi(\pi)$ avoids I_k .

Let $\tilde{\sigma}$ be the right-shadow sequence determined by the following two conditions.

1. For each symbol $i \in [m]$, the number of occurrences of i in $\bar{\pi}$ is equal to the number of its occurrences in $\tilde{\sigma}$.
2. For any i and j , the number of stars left-dominated by the j -th occurrence of i in $\bar{\pi}$ is equal to the number of stars right-dominated by the j -th occurrence of i in $\tilde{\sigma}$.

Note that these conditions determine $\tilde{\sigma}$ uniquely. As an example, consider the left-shadow sequence $\bar{\pi} = 123 * 3 * *44*$. In $\tilde{\sigma}$, the non-star elements form the subsequence 123443. The first occurrence of 3 in $\bar{\pi}$ left-dominates a single star, the second occurrence of 3 left-dominates two stars, and the second occurrence of 4 left-dominates one star. Hence, $\tilde{\sigma}$ is the sequence $12 * 34 * 4 * *3$.

Next, let S be the stack polyomino whose columns correspond to the stars of $\tilde{\sigma}$, where the i -th column has height h if the i -th star of $\tilde{\sigma}$ is right-dominated by $h + 1$. In the example above, if $\tilde{\sigma} = 12 * 34 * 4 * *3$, then S has four columns of heights $(2, 3, 2, 2)$. Clearly, S has the same content as F . By Theorem 54, there is a bijection f between the I_k -avoiding fillings of F and the I_k -avoiding fillings of S . This bijection transforms $\Phi(\pi)$ into a filling Ψ of S . Define a partition σ by replacing the i -th star in $\tilde{\sigma}$ by the row-index of the 1-cell in the i -th column of Ψ . By construction, σ is a partition with right shadow $\tilde{\sigma}$, and $\Psi(\sigma) = \Psi$. By Observation 60, σ avoids $12 \cdots k12 \cdots k(k+1)$.

This transformation, which is easily seen to be invertible, provides the required bijection. This completes the proof. \square

Patterns of the form $1(\tau + 1)$. We will now establish a general relationship between the partitions that avoid a pattern τ and the partitions that avoid the pattern $1(\tau + 1)$. The key result is the following theorem.

Theorem 62 (J., Mansour [37]). *Let τ be an arbitrary pattern, and let $F(x)$ be its corresponding EGF. Let $\sigma = 1(\tau + 1)$, and let $G(x)$ be its EGF. For every $n \geq 1$, the following holds:*

$$p_n(\sigma) = \sum_{i=0}^{n-1} \binom{n-1}{i} p_i(\tau). \quad (6.3)$$

In terms of generating functions, this is equivalent to

$$G(x) = 1 + \int_0^x F(t)e^t dt. \quad (6.4)$$

Proof. Fix σ and τ as in the statement of the theorem. Let π be an arbitrary partition, and let π^- denote the partition obtained from π by erasing every occurrence of the symbol 1, and decreasing every other symbol by 1; in other words, π^- represents the partition obtained by removing the first block from the partition π . Clearly, a partition π avoids σ if and only if π^- avoids τ . Thus, for every σ -avoiding partition $\pi \in \mathcal{P}_n(\sigma)$ there is a unique τ -avoiding partition $\rho \in \cup_{i=0}^{n-1} \mathcal{P}_i(\tau)$ satisfying $\pi^- = \rho$. On the other hand, for a fixed $\rho \in \mathcal{P}_i(\tau)$, there are $\binom{n-1}{i}$ partitions $\pi \in \mathcal{P}_n(\sigma)$ such that $\pi^- = \rho$. This gives equation (6.3).

To get equation (6.4), we multiply both sides of (6.3) by $\frac{x^n}{n!}$ and sum for all $n \geq 1$. This yields

$$\begin{aligned}
G(x) - 1 &= \sum_{n \geq 1} \frac{x^n}{n!} \sum_{i=0}^{n-1} \binom{n-1}{i} p_i(\tau) = \int_0^x \sum_{n \geq 1} \frac{t^{n-1}}{(n-1)!} \sum_{i=0}^{n-1} \binom{n-1}{i} p_i(\tau) dt \\
&= \int_0^x \sum_{n \geq 0} \frac{t^n}{n!} \sum_{i=0}^n \binom{n}{i} p_i(\tau) dt = \int_0^x \sum_{n \geq 0} \sum_{i=0}^n \frac{t^i}{i!} p_i(\tau) \frac{t^{n-i}}{(n-i)!} dt \\
&= \int_0^x \left(\sum_{i \geq 0} \frac{t^i}{i!} p_i(\tau) \right) \left(\sum_{k \geq 0} \frac{t^k}{k!} \right) dt = \int_0^x F(t) e^t dt,
\end{aligned}$$

which is equivalent to equation (6.4). \square

The following result is an immediate consequence of Theorem 62.

Corollary 63. *If $\tau \approx \tau'$ then $1(\tau + 1) \approx 1(\tau' + 1)$, and more generally, $12 \cdots k(\tau + k) \approx 12 \cdots k(\tau' + k)$. In particular, since $123 \approx 122 \approx 112 \approx 121$, we see that for every $m \geq 2$ the patterns $12 \cdots (m-1)m(m+1)$, $12 \cdots (m-1)mm$, $12 \cdots (m-1)(m-1)m$ and $12 \cdots (m-1)m(m-1)$ are equivalent. Conversely, if $1(\tau + 1) \approx 1(\tau' + 1)$, then $\tau \approx \tau'$.*

Proof. To prove the last claim, notice that equation (6.3) can be inverted to obtain

$$p_{n-1}(\tau) = \sum_{i=0}^{n-1} (-1)^i \binom{n-1}{i} p_{n-i}(\sigma).$$

The other claims follow directly from Theorem 62. \square

6.3 Patterns equivalent to $12 \cdots m(m+1)$

The partitions that avoid $12 \cdots m(m+1)$, or equivalently, the partitions with at most m blocks, are a very natural pattern-avoiding class of partitions. Their number may be expressed by $p_n(12 \cdots (m+1)) = \sum_{i=0}^m S(n, i)$, where $S(n, i)$ is the Stirling number of the second kind, which is equal to the number of partitions of $[n]$ with exactly i blocks.

As an application of the previous results, we will now present two classes of patterns that are equivalent to the pattern $12 \cdots (m+1)$. From this result, we obtain an alternative combinatorial interpretation of the Stirling numbers $S(n, i)$.

Our result is summarized in the following theorem.

Theorem 64 (J., Mansour [37]). *For every $m \geq 2$, the following patterns are equivalent:*

- (a) $12 \cdots (m-1)m(m+1)$,
- (b) $12 \cdots (m-1)md$, where d is any number from the set $[m]$,
- (c) $12 \cdots (m-1)dm$, where d is any number from the set $[m-1]$.

Proof. From Corollary 63, we get the equivalences

$$12 \cdots m(m+1) \approx 12 \cdots (m-1)mm \approx 12 \cdots (m-1)(m-1)m.$$

The equivalences

$$12 \cdots (m-1)mm \approx 12 \cdots (m-1)md \text{ and } 12 \cdots (m-1)(m-1)m \approx 12 \cdots (m-1)dm$$

are obtained by a repeated application of Corollary 49. \square

6.4 Binary patterns

Let us now focus on the avoidance of *binary* patterns, i.e., the patterns that only contain the symbols 1 and 2.

We will first consider the forbidden patterns of the form $1^k 2 1^\ell$. We have already seen that $112 \sim 121$. The following theorem offers a generalization.

Theorem 65 (J., Mansour [37]). *For any three integers j, k, m satisfying $1 \leq j, k \leq m$, the pattern $1^j 2 1^{m-j}$ is equivalent to the pattern $1^k 2 1^{m-k}$.*

Proof. It is sufficient to prove that the equivalence $1^j 2 1^{m-j} \approx 1^m 2$ holds for every m and j . We will use Lemma 40 from page 48 to provide a bijection between $\mathcal{P}_n(1^j 2 1^{m-j})$ and $\mathcal{P}_n(1^m 2)$. Fix $\tau \in \mathcal{P}_n(1^j 2 1^{m-j})$, and assume that τ has b blocks. Represent τ by a matrix $M = M(\tau, b)$. Since τ was a canonical sequence, the leftmost 1-cells of the b rows of M form an increasing chain.

Let us turn the matrix M upside down, to obtain a matrix \overline{M} . Clearly, \overline{M} avoids $M(2^j 1 2^{m-j}, 2)$. Also, \overline{M} is b -falling. We can treat \overline{M} as a b -falling filling of a stack polyomino, and apply Lemma 40 to transform it into a b -falling $M(2^m 1, 2)$ -avoiding filling \overline{M}' . Turning \overline{M}' upside down again, we obtain a semi-standard b -raising matrix M' , which corresponds to a $1^m 2$ -avoiding partition.

This transformation is the required bijection. \square

Using our results on fillings, we can add another pattern to the equivalence class covered by Theorem 65.

Theorem 66 (J., Mansour [37]). *For every $m \geq 1$, the pattern 12^m is equivalent to the pattern 121^{m-1} .*

Proof. This is just Corollary 49 with $k = 2$ and $S = 1^{m-1}$. \square

Corollary 67. *Let m be a positive integer, let τ be any pattern from the set*

$$T = \{1^k 2 1^{m-k} : 1 \leq k \leq m\} \cup \{12^m\}.$$

The EGF $F(x)$ of a pattern $\tau \in T$ is given by

$$F(x) = 1 + \int_0^x \exp\left(t + \sum_{i=1}^{m-1} \frac{t^i}{i!}\right) dt.$$

Proof. Theorems 65 and 66 show that all the patterns from the set T are equivalent, so we will compute the EGF of $\tau = 12^m$. The formula for $F(x)$ follows directly from equation (6.1) on page 53 and Theorem 62. \square

We now turn to another type of binary patterns, namely the patterns of the form $12^k 12^{m-k}$ with $1 \leq k \leq m$. In the rest of this section, S_q^p denotes the sequence $2^p 12^q$ and \overline{S}_q^p denotes the sequence $1^p 2 1^q$, where p, q are nonnegative integers. Our first aim is to show that for fixed m and arbitrary $k \in [m]$, all the pattern of the form $12^k 12^{m-k}$ belong to the same \approx -equivalence class. In fact, we are able to prove a more general result.

Theorem 68 (J., Mansour [37]). *For any partition τ , for any $k \geq 2$, and for any $p, q \geq 0$, we have the following equivalences:*

$$12 \cdots k(\tau + k) S_q^p \approx 12 \cdots k(\tau + k) S_0^{p+q}$$

and

$$12 \cdots k S_q^p(\tau + k) \approx 12 \cdots k S_0^{p+q}(\tau + k).$$

Proof. By Lemma 40, the two matrices $M(S_q^p, 2)$ and $M(S_0^{p+q}, 2)$ are stack equivalent, and hence also F^\perp -equivalent. By Proposition 48, this implies that $M(S_q^p, k) \stackrel{\Delta^s}{\approx} M(S_0^{p+q}, k)$ for any $k \geq 2$. Lemma 45 then gives the first equivalence, and Lemma 47 gives the second. \square

Next, we present two theorems that make use of the full strength of Lemma 40, including the preservation of the t -falling property. Recall that $\overline{S}_q^p = 1^p 21^q$.

Theorem 69 (J., Mansour [37]). *Let τ be any partition with k blocks, let $p \geq 1$ and $q \geq 0$. The pattern $\sigma = \tau(\overline{S}_q^p + k)$ is equivalent to $\sigma' = \tau(\overline{S}_0^{p+q} + k)$.*

Proof. Let π be a partition of $[n]$ with m blocks, let $M = M(\pi, m)$. We color the cells of M red and green, where a cell in row i and column j is green if and only if the submatrix of M formed by the intersection of the first $i - 1$ rows and $j - 1$ columns of M contains $M(\tau, k)$. It is not difficult to see that for each green cell (i, j) there is an occurrence of $M(\tau, k)$ which appears in the first $i - 1$ rows and the first $j - 1$ columns and which consists entirely of red cells. Thus, for any matrix M' obtained from M by modifying the filling of M 's green cells, a cell is green with respect to M' if and only if it is green with respect to M .

Let G be the diagram formed by the green cells of M , and let Φ be the filling of G by the values from M . Note that G is an upside-down copy of an F^\perp -shape. It is easy to see that the partition π avoids σ if and only if Φ avoids $M(\overline{S}_q^p, 2)$, and π avoids σ' if and only if Φ avoids $M(\overline{S}_0^{p+q}, 2)$.

Let us now assume that π is σ -avoiding. We now describe a procedure to transform π into a σ' -avoiding partition π' (see Figure 6.2). We first turn the filling Φ and the diagram G upside down, which transforms G into an F^\perp -shape \overline{G} , and it also transforms the $M(\overline{S}_q^p, 2)$ -avoiding filling Φ into an $M(S_q^p, 2)$ -avoiding filling $\overline{\Phi}$ of \overline{G} . Then we apply the bijection f of Lemma 40 to $\overline{\Phi}$, ignoring the zero columns. We thus obtain a filling $\overline{\Psi} = f(\overline{\Phi})$ which avoids $M(S_0^{p+q}, 2)$. We turn this filling upside down, obtaining a $M(\overline{S}_0^{p+q}, 2)$ -avoiding filling Ψ of G . We then fill the green cells of M with the values of Ψ while the filling of the red cells remains the same. We thus obtain a matrix M' . The matrix M' has exactly one 1-cell in each column, so there is a sequence π' over the alphabet $[m]$ such that $M' = M(\pi', m)$.

By construction, the sequence π' has no subsequence order-isomorphic to σ' . We now need to show that π' is a restricted-growth sequence. For this, we will use the preservation of the t -falling property. Let c_i be the leftmost 1-cell of the i -th row of M , let c'_i be the leftmost 1-cell of the i -th row of M' . We know that the cells c_1, \dots, c_m form an increasing chain, because π was a restricted-growth sequence. We want to show that the cells c'_1, \dots, c'_m form an increasing chain as well.

Let s be the largest index such that the cell c_s is red in M . We set $s = 0$ if no such cell exists. Note that the cells c_1, \dots, c_s are red and the cells c_{s+1}, \dots, c_m are green in M . We have $c_i = c'_i$ for every $i \leq s$. If $s > 0$, we also see that all the green 1-cells of M are in the columns to the right of c_s . This means that even in the matrix M' all the green 1-cells are to the right of c_s , because the zero columns of Φ must remain zero in Ψ . In particular, all the cells c'_{s+1}, \dots, c'_m appear to the right of c'_s .

It remains to show that c'_{s+1}, \dots, c'_m form an increasing chain. We know that the cells c_{s+1}, \dots, c_m form an increasing chain in M and in Φ . When G is turned upside down, this chain becomes a decreasing chain $\overline{c_{s+1}}, \dots, \overline{c_m}$ in $\overline{\Phi}$. This chain shows that $\overline{\Phi}$ is $(m - s)$ -falling. By Lemma 40, $\overline{\Psi}$ must be $(m - s)$ -falling as well, hence it contains a decreasing chain $\overline{c'_{s+1}}, \dots, \overline{c'_m}$ in its bottom $m - s$ rows. This decreasing chain corresponds to an




Figure 6.2: Illustration of the proof of Theorem 69.

increasing chain c'_{s+1}, \dots, c'_m in M' , showing that π' is a restricted-growth function, as claimed.

It is obvious that the above construction can be reversed, which shows that it is indeed a bijection between $\mathcal{P}_n(\sigma)$ and $\mathcal{P}_n(\sigma')$. \square

The following result is proved by a similar approach, but the argument is slightly more technical.

Theorem 70 (J., Mansour [37]). *Let T be an arbitrary sequence over the alphabet $[k]$, let $p \geq 1$ and $q \geq 0$. The partition $\sigma = 12 \cdots k(\overline{S}_q^p + k)T$ is equivalent to $\sigma' = 12 \cdots k(\overline{S}_0^{p+q} + k)T$.*

Proof. Let π be a partition of $[n]$ with m blocks, let $M = M(\pi, m)$. As in the previous proof, we color the cells of M red and green. A cell in row i and column j will be green if the submatrix of M formed by rows $1, \dots, i-1$ and columns $j+1, \dots, n$ contains $M(T, k)$.

Let G be the diagram formed by the green cells and Φ its filling inherited from M . Let r be the number of rows of G . The partition π contains σ if and only if Φ contains $M(\overline{S}_q^p, 2)$. Note that the diagram G is an upside-down copy of a Ferrers shape.

We apply the same construction as in the previous proof. Let $\overline{\Phi}$ be the upside down copy of Φ . The filling $\overline{\Phi}$ is r -falling and it avoids $M(S_q^p, 2)$. We apply the mapping f from Lemma 40 to transform $\overline{\Phi}$ into an r -falling semi-sparse filling $\overline{\Psi}$. We then turn $\overline{\Psi}$ upside down again and reinsert it into the green cells of the original matrix. This yields a matrix M' with exactly one 1-cell in each column. Hence, there exists a sequence π' , such that $M' = M(\pi', m)$. The sequence π' has no subsequence order-isomorphic to σ' .

We need to prove that π' is a restricted-growth sequence. Let c_i be the leftmost 1-cell in row i of M and let c'_i be the leftmost 1-cell in row i of M' . To prove that π' is a partition, we want to show that c'_1, \dots, c'_m form an increasing chain in M' .

Let us fix two row indices $i < j$. We claim that c'_i is left of c'_j . If both c'_i and c'_j are green, then the claim follows from the preservation of the r -falling property. If both c'_i and c'_j are red, then $c'_i = c_i$ and $c'_j = c_j$. The claim then follows from the fact that c_1, \dots, c_m

is an increasing chain. If c'_j is red and c'_i is green, the claim holds as well, because $c_j = c'_j$, and all the green cells below row j must appear to the left of the column of c_j .

Finally, assume that c'_j is green and c'_i is red. We have $c'_i = c_i$. All the 1-cells of Φ that are to the left of c_i are also below row i . Let x be the number of such 1-cells. Then x is equal to the number of nonzero columns of Φ that are to the left of c_i . Since the number of these nonzero columns is preserved by the mapping f , we see that Ψ also has x 1-cells left of c_i . Since f preserves the number of 1-cells in each row, both Φ and Ψ have exactly x 1-cells below row i . All the 1-cells of Ψ below row i must appear to the left of c_i , and since there are only x 1-cells of Ψ to the left of c_i , they must all appear below row i . Hence, all the green 1-cells above row i (including the cell c'_j) appear to the right of c_i . \square

Patterns equivalent to 12^k13 . We will now focus on the following sets of patterns:

$$\begin{aligned}\Sigma_t^+ &= \{12^{p+1}12^q32^r : p, q, r \geq 0, p + q + r = t\} \\ \Sigma_t^- &= \{12^{p+1}32^q12^r : p, q, r \geq 0, p + q + r = t\} \\ \Sigma_t &= \Sigma_t^+ \cup \Sigma_t^-\end{aligned}$$

Our aim is to show that all the patterns in Σ_t are equivalent. Throughout this section, we will assume that t is arbitrary but fixed. We will write Σ^+, Σ^- and Σ instead of Σ_t^+, Σ_t^- and Σ_t , if there is no risk of ambiguity.

The approach we will use is similar to the idea we employed to prove Theorem 41, but with considerably more technical details to take care of.

We will use the following definition.

Definition 71. Let σ be a pattern over the alphabet $\{1, 2, 3\}$, let π be a partition with m blocks, and let $k \leq m$ be an integer. We say that π *contains σ at level k* , if there are symbols $\ell, h \in [m]$ such that $\ell < k < h$, and the partition π contains a subsequence S made of the symbols $\{\ell, k, h\}$ which is order-isomorphic to σ .

For example, the partition $\pi = 1231323142221$ contains $\sigma = 121223$ at level 3, because π contains the subsequence 131334, but π avoids σ at level 2, because π has no subsequence of the form $\ell 2 \ell 2 2 h$ with $\ell < 2 < h$.

Our plan is to show, for suitable pairs $\sigma, \sigma' \in \Sigma$, that for every k there is a bijection f_k that maps the partitions avoiding σ at level k to the partitions avoiding σ' at level k , while preserving σ' -avoidance at all levels $j < k$ and preserving σ -avoidance at all levels $j > k + 1$. Composing the maps f_k for $k = 2, \dots, n - 1$, we will obtain a bijection between $\mathcal{P}_n(\sigma)$ and $\mathcal{P}_n(\sigma')$.

To formalize this idea we will need more definitions.

Definition 72. Consider a partition π , and fix a level $k \geq 2$. A symbol of π is called *k-low* if it is smaller than k and *k-high* if it is greater than k . A *k-low cluster* (or *k-high cluster*) is a maximal consecutive sequence of k -low symbols (or k -high symbols, respectively) in π . The *k-landscape* of π is a word over the alphabet $\{L, k, H\}$ obtained from π by replacing each k -low cluster with a single symbol L and each k -high cluster with a single symbol H.

A word W over the alphabet $\{L, k, H\}$ is called a *k-landscape word* if it satisfies the following conditions.

- The first symbol of W is L, the second symbol of W is k .
- No two symbols L are consecutive in W , no two symbols H are consecutive in W .

Clearly, the landscape of a partition is a landscape word.

Two k -landscape words W and W' are said to be *compatible*, if each of the three symbols $\{L, k, H\}$ has the same number of occurrences in W as in W' .

We will often drop the prefix k from these terms, if the value of k is clear from the context.

To give an example, consider $\pi = 1231323142221$: it has five 3-low clusters, namely 12, 1, 2, 1 and 2221, it has one 3-high cluster 4, and its 3-landscape is L3L3L3LHL.

If W and W' are two compatible k -landscape words, we have a natural bijection between partitions with landscape W and partitions with landscape W' . If π has landscape W , we map π to the partition π' of landscape W' which has the same k -low clusters and k -high clusters as π , and moreover, the k -low clusters appear in the same order in π as in π' , and also the k -high clusters appear in the same order in π as in π' . It is not difficult to check that these rules define a unique sequence π' and this sequence is indeed a partition. This provides a bijection between partitions of landscape W and partitions of landscape W' which will be called *the k -shuffle from W to W'* .

The key property of shuffles is established by the next lemma.

Lemma 73. *Let W and W' be two compatible k -landscape words. Let π be a partition with k -landscape W and let π' be the partition obtained from π by the shuffle from W to W' . Let σ be a pattern from Σ , and let j be an integer. The following holds.*

1. *If σ does not end with the symbol 1 and $j > k$, then π' contains σ at level j if and only if π contains σ at level j .*
2. *If σ does not end with the symbol 3 and $j < k$, then π' contains σ at level j if and only if π contains σ at level j .*

Proof. We begin with the first claim of the lemma. Let $\sigma = 12^{p+1}32^q12^r$ be an arbitrary pattern from Σ^- (the case $\sigma \in \Sigma^+$ is analogous). By assumption, we have $r > 0$. Assume that π contains σ at a level $j > k$. In particular, π has a subsequence $S = \ell j^{p+1} h j^q \ell j^r$, with $\ell < j < h$.

If $k < \ell$, then all the symbols of S are k -high. Since the shuffle preserves the relative order of high symbols, π' contains the subsequence S as well.

If $k \geq \ell$, then the shuffle preserves the relative order of the symbols j and h , which are all high. Let x and y be the two symbols of S directly adjacent to the second occurrence of ℓ in S (if $q > 0$, both these symbols are equal to j , otherwise one of them is h and the other j). The two symbols are both high, but they must appear in different k -high clusters. After the shuffle, the two symbols x and y will again be in different clusters, separated by a non-high symbol $\ell' \leq k$, and since the first occurrence of ℓ' in π' precedes any occurrence of j , the partition π' will contain a subsequence $\ell' j^{p+1} h j^q \ell' j^r$, which is order-isomorphic to σ .

We see that the shuffle preserves the occurrence of σ at level j . Since the inverse of the shuffle from W to W' is the shuffle from W' to W , we see that the inverse of a shuffle preserves the occurrence of σ at level j as well.

The second claim of the lemma is proved by a similar argument. Assume that π contains σ at a level $j < k$. Thus, π contains a subsequence S over the alphabet $\{\ell < j < h\}$, which is order-isomorphic to σ . If $h < k$, then the symbols of S are low and hence preserved by the shuffle. If $h \geq k$, let x and y be the two symbols of S adjacent to the symbol of h . Recall that σ does not end with the symbol 3, so x and y are both well defined. The symbols x and y must appear in two distinct low clusters. After the shuffle is performed there will be a non-low symbol h' between x and y . Hence, π' will contain a subsequence order isomorphic to σ . \square

We will use shuffles as basic building blocks for our bijections. The first example is the following lemma.

Lemma 74. *For every $p, q, r \geq 0$, the pattern $\sigma = 12^{p+1}12^q32^r$ is equivalent to the pattern $\sigma' = 12^{p+1}32^q12^r$.*

Proof. Let us fix $p, q, r \geq 0$ with $t = p + q + r$. For a given k , a partition π of $[n]$ is called a k -hybrid if π avoids σ' at every level $j < k$ and π avoids σ at every level $j \geq k$. We will show that for every $k \in \{2, \dots, n-1\}$ there is a bijection f_k between k -hybrids and $(k+1)$ -hybrids. Since 2-hybrids are precisely the σ -avoiding partitions of $[n]$ and n -hybrids are precisely the σ' -avoiding partitions of $[n]$, this gives the required result.

Let us fix k . Note that a partition π contains σ at level k if and only if its k -landscape W contains a subsequence $k^{p+1}Lk^qHk^r$. Similarly, π contains σ' at level k if and only if W contains a subsequence $k^{p+1}Hk^qLk^r$.

Let π be a k -hybrid with landscape W . If π has fewer than $t+1$ occurrences of k , then it is also a $(k+1)$ -hybrid and we put $f_k(\pi) = \pi$. Otherwise, we write $W = XYZ$, where X is the shortest prefix of W that has $p+1$ symbols k and Z is the shortest suffix of W that has r symbols k . By assumption, X and Z do not overlap (although they may be adjacent if $q=0$). Let \bar{Y} be the word obtained by reversing the order of the letters of Y , and let us define $W' = X\bar{Y}Z$. Note that W' is a landscape word compatible with W , and that W avoids $k^{p+1}Lk^qHk^r$ if and only if W' avoids $k^{p+1}Hk^qLk^r$. We apply to π the shuffle from W to W' which transforms it into a partition $\pi' = f_k(\pi)$.

Lemma 73 implies that π' is a $(k+1)$ -hybrid. Hence, f_k is the required bijection. \square

Another result in the same spirit is the following lemma.

Lemma 75. *For every $p, q, r \geq 0$, the pattern $\sigma = 12^{p+2}12^q32^r$ is equivalent to the pattern $\sigma' = 12^{p+1}12^q32^{r+1}$.*

Proof. We follow a similar argument as in the proof of Lemma 74. As before, a k -hybrid is a partition that avoids σ' at every level $j < k$ and that avoids σ at every level $j \geq k$. We will present a bijection f_k between k -hybrids and $(k+1)$ -hybrids. Note that π avoids σ at level k if and only if its landscape W avoids $k^{p+2}Lk^qHk^r$.

Fix a k -hybrid π with landscape W . If π has fewer than $p+2+q+r$ occurrences of k , then it is also a $(k+1)$ -hybrid and we define $f_k(\pi) = \pi$; otherwise, we write $W = XSYZ$ where X is the shortest prefix of W that has $p+1$ occurrences of k , Z is the shortest suffix with r occurrences of k , S is the subword that starts just after the $(p+1)$ -th occurrence of k and ends immediately after the $(p+2)$ -th occurrence of k . We define $W' = XY\bar{S}Z$, where \bar{S} is the reversal of S .

Note that in the definition of W' , we need to take $W' = XY\bar{S}Z$ instead of the seemingly more natural definition $W' = XYSZ$. This is because in general, the string $XYSZ$ need not be a landscape word, since it may contain two consecutive occurrences of either L or H. Our definition guarantees that W' is a correct landscape word, and that W' avoids $k^{p+1}Lk^qHk^{r+1}$ if and only if W avoids $k^{p+2}Lk^qHk^r$ (which is if and only if Y avoids Lk^qH).

The rest of the argument is the same as in the previous lemma. \square

We may now state and prove the main result of this paragraph.

Theorem 76 (J., Mansour [37]). *For every t , the patterns in the set Σ_t are equivalent.*

Proof. By Theorem 68, we already know that for any $p, q \geq 0$, the pattern $12^{p+1}12^q3$ is equivalent to the pattern $12^{p+q+1}13$. This, together with the two previous lemmas gives the required result. \square

More ‘landscape’ patterns. We will show that with a little bit of additional effort, the previous arguments involving landscapes can be adapted to prove, for every $p, q \geq 0$, the following equivalences:

- $1232^p412^q \approx 1232^p42^q1$
- $1232^p142^q \approx 12312^p42^q$
- $123^{p+1}143^q \approx 123^{p+1}13^q4$
- $123^{p+1}413^q \approx 12343^p13^q$

Throughout this paragraph, we will say that τ is a *1-2-4 pattern* if τ has the form $123S$ where S is a sequence that satisfies these two conditions:

- S has exactly one occurrence of the symbol ‘1’, exactly one occurrence of the symbol ‘4’, and all its remaining symbols are equal to ‘2’.
- Neither the first nor the last symbol of S is equal to ‘4’.

Similarly, a *1-3-4 pattern* is a pattern of the form $123S$ where S satisfies these two conditions:

- S has one occurrence of ‘1’ and one occurrence of ‘4’, and all its remaining symbols are equal to ‘3’.
- The last symbol of S not equal to ‘1’.

We decided to exclude the patterns of the form 1232^p12^q4 , 12342^p12^q and 1233^p43^q1 from the set of 1-2-4 and 1-3-4 patterns defined above, because some of the arguments we will need in the following discussion (namely in Lemma 77) would become more complicated if these special types of patterns were allowed. We need not be too concerned about this constraint, because we have already dealt with the patterns of the three excluded types in Theorem 68 and Theorem 70. From Theorem 68, we obtain the equivalences $1232^p12^q4 \approx 1232^{p+q}14$ and $12342^p12^q \approx 12342^{p+q}1$, while from Theorem 70, we obtain $1233^p43^q1 \approx 1233^{p+q}41$.

For our arguments, we need to extend some of the terminology of the previous section to cover the new family of patterns. Let τ be a 1-2-4 pattern, k be a natural number, and π be a partition. We say that π *contains* τ *at level* k , if π has a subsequence T order-isomorphic to τ such that the occurrences of the symbol ‘2’ in τ correspond to the occurrences of the symbol k in T . Similarly, if τ is a 1-3-4 pattern, we say that a partition π *contains* τ *at level* k if π has a subsequence T order-isomorphic to τ with the symbol k in T corresponding to the symbol ‘3’ in τ .

Our aim is to prove an analogue of Lemma 73 for 1-2-4 and 1-3-4 patterns. Unfortunately, general k -shuffles may behave badly with respect to the avoidance of these patterns. However, we will define special types of k -shuffles that have the properties we need. We first introduce some new definitions.

Let W be a k -landscape word. We say that two occurrences of the symbol H in W are *separated* if there is at least one occurrence of L between them. Similarly, two symbols L are separated if there is at least one H between them. As an example, consider the k -landscape word $W = LkLkHkkHLkH$. In W , neither the first two occurrences of L nor the first two occurrences of H are separated, while the second and third occurrence of H, as well as the second and third occurrence of L are separated. We also say that two

clusters of a partition are separated if the corresponding symbols of the landscape word are separated.

Let W and W' be two k -landscape words. We say that W and W' are *H-compatible* if they are compatible, and for any i, j , the i -th and j -th occurrence of H in W are separated if and only if the i -th and j -th occurrence of H in W' are separated. An L-compatible pair of words is defined analogously.

For example, the two compatible words $W = LkHkkHL$ and $W' = LkHkLHk$ are L-compatible (since the two occurrences of L are separated in both words) but they are not H-compatible (the two symbols H are not separated in W but they are separated in W').

The following lemma explains the relevance of these concepts.

Lemma 77. *Let k be an integer. The following holds.*

- (1) *Let W and W' be two L-compatible k -landscape words, and let τ be a 1-2-4 pattern. Let π be an arbitrary partition, and let π' be the partition obtained from π by the k -shuffle from W to W' . For every $j < k$, π contains τ at level j if and only if π' contains τ at level j . Moreover, if the last symbol of τ is equal to 2, then the previous equivalence also holds for every $j > k$.*
- (2) *Let W and W' be two H-compatible k -landscape words, and let τ be a 1-3-4 pattern. Let π be an arbitrary partition, and let π' be the partition obtained from π by the k -shuffle from W to W' . For every $j > k$, π contains τ at level j if and only if π' contains τ at level j . Moreover, if the last symbol of τ is equal to 3, then the previous equivalence also holds for every $j < k$.*

Proof. We first prove (1). Assume that π contains a 1-2-4 pattern τ at level j . If $j > k$, we may use the same argument as in the proof of the first part of Lemma 73 to see that the occurrence of τ is preserved by the shuffle as long as τ does not end with a 1.

Assume now that $j < k$. Let us write $\tau = 1232^p42^q12^r$ (the case when τ has the form $1232^p12^q42^r$ is analogous). Note that our definition of 1-2-4 pattern implies that $p \neq 0$.

By assumption, π contains a subsequence T order-isomorphic to τ , with the symbol 2 of τ corresponding to the symbol j in T . We label from left-to-right the $1 + p + q + r$ occurrences of j in T by $j_0, j_1, \dots, j_{p+q+r}$. Let $a < b < c$ denote the symbols of T that correspond respectively to the symbols 1, 3 and 4 in τ ; we label the two occurrences of a in T by a_0 and a_1 . With this notation, we may write T as follows:

$$T = a_0 j_0 b j_1 \cdots j_p c j_{p+1} \cdots j_{p+q} a_1 j_{p+q+1} \cdots j_{p+q+r}.$$

We distinguish several cases, based on the relative order of b, c and k . If $c < k$, then all the symbols of T are k -low and their relative position is preserved by the shuffle, which means that T is also a subsequence of π' .

If $c \geq k$ and $b < k$, then the symbols $a < j < b$ are k -low. Let x and y be the two symbols adjacent to c in T . Typically $x = j_p$ and $y = j_{p+1}$, unless q is zero, in which case $y = a_1$. Recall that c cannot directly follow b and it cannot be the last element of T by the definition of 1-2-4 pattern. The elements x and y are low and they appear in two distinct low clusters. After the shuffle, the occurrences of a, b and j in T have the same relative order, and the elements x and y still belong to different clusters. Thus, π' contains a symbol greater than b between x and y . This shows that π' has a subsequence order-isomorphic to τ .

It remains to consider the most complicated case, when $c > k$ and $b \geq k$. This is when we first use the L-compatibility assumption. Let x and y be again the two symbols adjacent to c in T . By the definition of 1-2-4 patterns, x and y are both k -low. Since b is not k -low and c is high, the partition π has the following properties.

1. The symbol j_1 does not belong to the leftmost low cluster.
2. The two symbols x and y belong to two separated low clusters.

The two properties are preserved by the shuffle. In particular, in π' , the symbol j_1 does not belong to the leftmost low cluster, which means that there is at least one non-low symbol appearing in π' before j_1 . Since π' is a partition in its canonical sequential form, this implies that all the symbols $1, 2, \dots, k$ appear in π' in this order before j_1 . Let a', j' and k' denote respectively the leftmost occurrences of a, j and k in π' . We also know, from the L-compatibility of W and W' , that in π' the two symbols x and y appear in distinct and separated low clusters. In particular, π' contains a k -high symbol c' between x and y . Putting it all together, we see that π' contains the subsequence

$$T' = a'j'k'j_1 \cdots j_p c' j_{p+1} \cdots j_{p+q} a_1 j_{p+q+1} \cdots j_{p+q+r},$$

which is order isomorphic to τ .

Thus, π contains a 1-2-4 pattern τ at level j , if and only if π' contains τ at level j . This completes the proof of (1).

Claim (2) is proved by a similar argument. Let τ be a 1-3-4 pattern of the form $123^{p+1}13^q43^r$ (the case when $\tau = 123^{p+1}43^q13^r$ is analogous and easier). Assume that π contains τ at level j , represented by a sequence T of the form

$$T = a_0 b j_0 j_1 \cdots j_p a_1 j_{p+1} \cdots j_{p+q} c j_{p+q+1} \cdots j_{p+q+r},$$

with $a < b < j < c$.

If $j < k$, we apply the same argument as in the proof of the second claim of Lemma 73 to prove that if τ does not end with 4, then the occurrence of τ is preserved by the shuffle.

Next, we assume that $j > k$ and we distinguish several cases based on the relative order of a, b and k .

If $a > k$, then all the symbols of T are k -high and their order is preserved by the shuffle.

If $a \leq k$, and $b > k$, we let x and y denote the two symbols adjacent to a_1 in T , and we observe that π' has a non-high element a' between x and y . The first occurrence of a' in π' must appear to the left of any k -high symbol, hence π' contains a subsequence $a'bj^{p+1}a'j^q c j^r$ order-isomorphic to τ .

If $a < k$ and $b \leq k$, we define x and y as in the previous paragraph. This time, x and y belong to two separated high clusters, so π' has a k -low element a' between x and y , and in particular, π' contains the subsequence $a'kj^{p+1}a'j^q c j^r$. \square

With the help of Lemma 77, we may prove all the equivalence relations announced at the beginning of this paragraph. We split the proofs into four lemmas and then summarize the results in a theorem.

Lemma 78. *Let $p, q \geq 1$. The pattern $\tau = 1232^p412^q$ is equivalent to $\tau' = 1232^p42^q1$.*

Proof. For an integer k we say that a partition π is a k -hybrid if π avoids τ' at level j for every $j < k$ and it avoids τ at level j for every $j \geq k$. To prove the claim, it is enough to establish a bijection f_k between k -hybrids and $(k+1)$ -hybrids.

We say that a k -high cluster of π is *extra-high* if it contains a symbol greater than $k+1$. We claim that π contains τ at level k if and only if by scanning the k -landscape W of π from left to right we may find (not necessarily consecutively) the leftmost high cluster, followed by p occurrences of the symbol k , followed by an extra-high cluster, followed by a low cluster, followed by q occurrences of k . To see this, it suffices to notice that the

leftmost high cluster contains the symbol $k + 1$, and to the left of this cluster we may always find all the symbols $12 \cdots k$ in increasing order.

By a similar argument, we see that π contains τ' at level k if and only if it contains, left-to-right, the leftmost high-cluster, p occurrences of k , an extra-high cluster, q occurrences of k and a low cluster.

Now assume that π is a k -hybrid partition. Let H' be the leftmost extra-high cluster of π such that between H' and the leftmost high cluster of π there are at least p occurrences of k . If no such cluster exists, or if π has fewer than q symbols equal to k to the right of H' , then π avoids both τ and τ' at level k , and we define $f_k(\pi) = \pi$.

Otherwise, let W be the k -landscape of π . We will decompose W as

$$W = XH'Yk_qS_1k_{q-1}S_2 \cdots k_1S_q,$$

where H' represents the extra-high cluster defined above, and k_i represents the i -th symbol k in π , counted from the right. The symbols X, Y and S_1, \dots, S_q above refer to the corresponding subwords of W appearing between these symbols.

By construction, none of the S_i 's contains the symbol k , so each of them is an alternating sequence over the alphabet $\{L, H\}$, possibly empty. Since π avoids τ at level k , the subword Y does not contain the symbol L .

We decompose S_1 into two parts $S_1 = H^*S_1^-$ in the following way: if the first letter of S_1 is H , then we put $H^* = H$ and S_1^- is equal to S_1 with the first letter removed. If S_1 does not start with H , then H^* is the empty string and $S_1^- = S_1$.

Now, we define the word W' by

$$W' = XH'S_1^-k_1S_2k_2S_3k_3 \cdots k_{q-1}S_qk_qH^*Y.$$

It is not difficult to check that W' is a landscape word (note that neither Y nor S_1^- can start with the symbol H), and that W' is L -compatible with W (recall that Y contains no L).

Let π' be the partition obtained from π by the shuffle from W to W' . Note that the prefix of π through the cluster H' is not affected by the shuffle, because the words W and W' share the same prefix up to the symbol H' . In particular, the shuffle preserves the property that H' is the leftmost extra-high cluster with at least p symbols k between H' and the leftmost high cluster of π' . It is routine to check that π' avoids τ' at level k . By Lemma 77, π' is a $(k + 1)$ -hybrid partition. It is easy to see that for any given $(k + 1)$ -hybrid partition π' , we may uniquely invert the procedure above and obtain a k -hybrid partition π .

Defining $f_k(\pi) = \pi'$, we obtain the required bijection between k -hybrids and $(k + 1)$ -hybrids. \square

The proofs of the following three lemmas follow the same basic argument as the proof of Lemma 78 above. The only difference is in the decompositions of the corresponding landscape words W and W' . We omit the common parts of the arguments and concentrate on pointing out the differences.

Lemma 79. *Let $p, q \geq 1$. The pattern $\tau = 1232^p142^q$ is equivalent to $\tau' = 12312^p42^q$.*

Proof. A partition π contains τ at level k if and only if it contains, from left to right, the leftmost high cluster, p copies of k , a low cluster, an extra-high cluster, and q copies of k . Similar characterization applies to τ' .

Let H_1 denote the leftmost high cluster of π , let H' denote the rightmost extra-high cluster of π that has the property that there are at least q occurrences of k to the right

of H' . If H' does not exist, or if there are fewer than p occurrences of k between H_1 and H' , then π contains neither τ nor τ' at level k and we put $f_k(\pi) = \pi$. Otherwise, let W be the landscape of π , and let us write

$$W = XH_1S_1k_1S_2k_2 \cdots S_pk_pYH'Z,$$

where none of the S_i contains k , and Y avoids L . Define S_p^- and H^* by writing $S_p = S_p^-H^*$ where S_p^- does not end with the letter H and H^* is equal either to H or to the empty string, depending on whether S_p ends with H or not.

Now we write

$$W' = XH_1\bar{Y}k_1H^*S_1k_2S_2 \cdots k_pS_p^-H'Z,$$

where \bar{Y} is the reversal of Y . The rest of the proof is analogous to Lemma 78. \square

We now apply the same arguments to 1-3-4 patterns.

Lemma 80. *For any $p \geq 0$ and $q \geq 1$, the pattern $\tau = 123^{p+1}13^q4$ is equivalent to the pattern $\tau' = 123^{p+1}143^q$.*

Proof. As usual, a k -hybrid is a partition that avoids τ at every level $j \geq k$ and that avoids τ' at every level below k .

Let us say that a k -cluster of a partition π is *extra-low* if it contains a symbol smaller than $k - 1$. A partition contains τ at level k if and only if it has $p + 1$ occurrences of k followed by an extra-low cluster, followed q occurrences of k , followed by a high cluster. Similarly, a partition contains τ' at level k if and only if it has $p + 1$ copies of k , followed by an extra-low cluster, followed by a high cluster, followed by q copies of k .

Assume π is a k -hybrid partition. Let L' denote the leftmost extra-low cluster of π that has at least $p + 1$ copies of k to its left. If L' does not exist, or if it has fewer than q copies of k to its right, we put $f_k(\pi) = \pi$. Otherwise, we decompose the landscape word W of π as

$$W = XL'S_1k_1S_2k_2 \cdots S_{q-1}k_{q-1}S_qk_qY,$$

where the S_i do not contain k . By assumption, Y avoids H . Next, we write $Y = L^*Y^-$ where L^* is an empty string or a single symbol L , and Y^- does not start with L . We define W' by

$$W' = XL'Y^-k_1L^*S_1k_2 \cdots S_{q-1}k_qS_q.$$

The words W and W' are H -compatible. We define the bijection between k -hybrids and $(k + 1)$ -hybrids in the usual way. \square

Lemma 81. *For every $p \geq 0$ and $q \geq 1$, the pattern $\tau = 123^{p+1}413^q$ is equivalent to the pattern $\tau' = 12343^p13^q$.*

Proof. As before, take π to be a k -hybrid partition. Let L' be the rightmost extra-low cluster that has at least q copies of k to its right. If L' has at least $p + 1$ copies of k to its left, we decompose the landscape W of π as

$$W = Lk_1S_1k_2S_2 \cdots k_pS_pk_{p+1}YL'Z.$$

Next, we write $S_p = S_p^-L^*$ with the usual meaning and define

$$W' = Lk_1L^*\bar{Y}k_2S_1k_3S_2 \cdots S_{p-1}k_{p+1}S_p^-L'Z.$$

The rest is the same as before. \square

We now summarize our results.

Theorem 82 (J., Mansour [37]). *For every $p, q \geq 0$, we have the following equivalences:*

1. $1232^p 412^q \approx 1232^p 42^q 1$
2. $1232^p 142^q \approx 12312^p 42^q$
3. $123^{p+1} 143^q \approx 123^{p+1} 13^q 4$
4. $123^{p+1} 413^q \approx 12343^p 13^q$

Proof. If p and q are both positive, the results follow directly from the four preceding lemmas.

If $p = 0$, the second and the fourth claim are trivial, the first one is a special case of Theorem 68, and the third is covered by Lemma 80.

If $q = 0$, the first and the third claim are trivial, the second is a special case of Theorem 68, and the fourth follows from Theorem 70. \square

6.5 Sporadic equivalences

The results that we have presented so far have always yielded infinite families of equivalent pairs of patterns. However, the computer enumeration of small patterns undertaken by Jelínek and Mansour [37] has revealed two likely pairs of equivalent patterns which are not covered by any of the previous general classes. The two equivalences suggested by the enumerative data are $1123 \approx 1212$ and $12112 \approx 12212$. To complete the classification of small partition patterns, we will show that the two pairs of patterns are indeed equivalent. We are not able to generalize any of these equivalences to a more general family of equivalent patterns. For this reason, we call them the ‘sporadic’ pairs.

Enumeration of 1123-avoiding partitions. Let us first deal with the equivalence $1212 \approx 1123$. Unlike in the previous arguments, we do not present a direct bijection between pattern-avoiding classes. Instead, we prove that $p_n(1123)$ is equal to the n -th Catalan number, i.e., $p_n(1123) = \frac{1}{n+1} \binom{2n}{n}$. Since it is well known, at least since 1970’s [49], that noncrossing partitions are enumerated by the Catalan numbers, this will yield the desired equivalence.

We achieve our goal by proving that $p_n(1123)$ is equal to the number of Dyck paths of semilength n . A *Dyck path of semilength n* is a nonnegative path on the two-dimensional integer lattice from $(0, 0)$ to $(2n, 0)$ composed of *up-steps* connecting (x, y) to $(x + 1, y + 1)$ and *down-steps* connecting (x, y) to $(x + 1, y - 1)$. It is well known that these paths are enumerated by Catalan numbers (for a survey of the many combinatorial structures enumerated by the Catalan numbers, see the Catalan Addendum of Stanley [67]).

Let $D(n, k)$ be the set of Dyck paths of semilength n whose last up-step is followed by exactly k down-steps. Let $d(n, k)$ be the cardinality of $D(n, k)$. Additional combinatorial interpretations of $d(n, k)$ can be found in the OEIS [68, sequence A033184].

Lemma 83. *The numbers $d(n, k)$ are determined by the following set of recurrences:*

$$d(1, 1) = 1 \tag{6.5}$$

$$d(n, k) = 0 \quad \text{if } k < 1 \quad \text{or } k > n \tag{6.6}$$

$$d(n, k) = \sum_{j=k-1}^{n-1} d(n-1, j) \quad \text{for } n \geq 2, n \geq k \geq 1. \tag{6.7}$$

Proof. Only the third recurrence is nontrivial. We prove it by presenting a bijection between $D(n, k)$ and the disjoint union $\bigcup_{j=k-1}^{n-1} D(n-1, j)$. Assume that k and n are fixed, with $n \geq 2$ and $k \leq n$. Take a Dyck path $P \in D(n, k)$. By erasing the last up-step and the last down-step of D , we get a Dyck path $P' \in D(n-1, j)$, where $j \geq k-1$. Conversely, given a Dyck path $P' \in D(n-1, j)$ with $j \geq k-1$, we insert a down-step at the end of D' , and then insert an up-step into the resulting path immediately before its last k down-steps. This inverts the mapping above. \square

We now focus on 1123-avoiding partitions. First of all, we will present a correspondence between 1123-avoiding partitions and 123-avoiding words. A 123-avoiding word is a sequence s_1, s_2, \dots, s_ℓ of positive integers, such that there are no three indices $i < j < k$ that would satisfy $s_i < s_j < s_k$. We define the *rank* of a word to be equal to $\ell + m - 1$, where ℓ is the length of the word and $m = \max\{s_i, i = 1, \dots, \ell\}$ is the largest symbol of the word.

For example, there are five 123-avoiding words of rank 3: 111, 12, 21, 22, and 3. There are fourteen 123-avoiding words of rank 4: 1111, 112, 121, 122, 211, 212, 221, 222, 13, 23, 31, 32, 33, and 4.

Claim 84. *A 1123-avoiding partition π of $[n]$ with m blocks has the following form:*

$$\pi = 123 \cdots (m-2)(m-1)S \tag{6.8}$$

where S is a 123-avoiding word of rank n , with maximum element m . Conversely, if S is any 123-avoiding word of rank n with maximum element m then π defined by the formula (6.8) is a canonical sequence of a 1123-avoiding partition of $[n]$.

In particular, the number of 123-avoiding words of rank n with last element k is equal to the number of 1123-avoiding partitions of size n with last element k .

Proof. Let $\pi = \pi_1 \cdots \pi_n$ be a 1123-avoiding partition with m blocks, with $\pi_n = k$. Observe that for every $i \in [m-1]$, the symbol π_i is equal to i , otherwise π would contain the forbidden pattern. It follows that π can be decomposed as $\pi = 123 \cdots (m-2)(m-1)S$, where the word S has length $\ell = n - m + 1$ and maximum element equal to m . In particular, S has rank n and its last element is equal to k .

We now check that S is 123-avoiding. If S contained a subsequence xyz for $x < y < z$ then the original partition would contain a subsequence xyz , which is forbidden. It follows that S obtained from a 1123-avoiding partition π has all the required properties.

Conversely, if S is a 123-avoiding sequence of rank n and maximum element m , then it is routine to verify that $\pi = 12 \cdots (m-1)S$ is a 1123-avoiding partition of size n with m blocks. Clearly, the last element of π is equal to the last element of S . \square

Let $T(n, k)$ be the set of 123-avoiding words of rank n with last element equal to k . Let $t(n, k)$ be the cardinality of $T(n, k)$. By the previous claim, $t(n, k)$ is also equal to the number of 1123-avoiding partitions of size n with last element equal to k . To show that 1123-avoiding partitions of size n have the same enumeration as Dyck paths of semilength n , it suffices to show that $d(n, k) = t(n, k)$ for each n, k . To show this, we will prove that $t(n, k)$ is determined by the same set of recurrences as $d(n, k)$.

Claim 85. *The numbers $t(n, k)$ satisfy the following set of recurrences:*

$$t(1, 1) = 1 \tag{6.9}$$

$$t(n, k) = 0 \quad \text{if } k < 1 \quad \text{or } k > n \tag{6.10}$$

$$t(n, k) = \sum_{i=k-1}^{n-1} t(n-1, i) \quad \text{for } n \geq 2, n \geq k \geq 1 \tag{6.11}$$

Proof. Only the recurrence (6.11) is nontrivial. Let us fix $n \geq 2$ and $k \leq n$. To prove the recurrence, we need a bijection from $T(n, k)$ to $\cup_{i=k-1}^{n-1} T(n-1, i)$.

Let us first consider the case $k = 1$. A word $S \in T(n, 1)$ can be transformed into a word $S' \in \cup_{i=0}^{n-1} T(n-1, i)$, by simply erasing the last element of S . This provides a bijection between $T(n, 1)$ and $\cup_{i=0}^{n-1} T(n-1, i)$.

In the rest of the proof, we deal with the case $k > 1$. Let $S \in T(n, k)$ be a 123-avoiding word of length ℓ . The word S can be uniquely expressed as $S = S_0 1^b k$, where S_0 is the (possibly empty) longest proper prefix of S whose last element is different from 1. If S_0 is nonempty, let j be the last element of S_0 .

Let us decompose $T(n, k)$ into a disjoint union of two sets T_1 and T_2 defined by

$$\begin{aligned} T_1 &= \{S \in T(n, k): S_0 \text{ is nonempty, and } j \geq k\} \\ T_2 &= \{S \in T(n, k): S_0 \text{ is empty, or } j < k\}. \end{aligned}$$

Note that if S belongs to T_2 and S_0 is nonempty, then all the elements of S_0 are greater than or equal to j . Indeed, if S_0 contained an element i smaller than j , then S would contain a subsequence ijk , which would create a copy of 123 in S .

Let S' be a word from $\cup_{j=k-1}^{n-1} T(n-1, j)$. S' may be uniquely expressed as $S' = S'_0 (k-1)^c$, where $c \geq 0$ and S'_0 is the (possibly empty) longest prefix of S' whose last element is different from $k-1$. Note that if the last element of S' is greater than $k-1$ then $S' = S'_0$. If S'_0 is nonempty, let j' be the last element of S'_0 .

We decompose $\cup_{j=k-1}^{n-1} T(n-1, j)$ into a disjoint union of two sets T'_1 and T'_2 , where

$$\begin{aligned} T'_1 &= \{S' \in \cup_{i=k-1}^{n-1} T(n-1, i): S'_0 \text{ is nonempty, and } j' \geq k\} \\ T'_2 &= \{S' \in \cup_{i=k-1}^{n-1} T(n-1, i): S'_0 \text{ is empty, or } j' < k-1\}. \end{aligned}$$

Since j' is never equal to $k-1$, the two sets T'_1 and T'_2 form a disjoint partition of $\cup_{i=k-1}^{n-1} T(n-1, i)$. Note that T'_2 is in fact a subset of $T(n-1, k-1)$.

To prove the claim, it suffices to give a bijection f_1 between T_1 and T'_1 , and a bijection f_2 between T_2 and T'_2 .

We first construct f_1 . Choose $S \in T_1$ and write $S = S_0 1^b k$ as above. Let j be the last element of S_0 . Define $S' = f_1(S) = S_0 (k-1)^b$. Let us check that S' belongs to T'_1 . It is easy to see that S' avoids 123. The length of S' is one less than the length of S , and the maximum of S' is equal to the maximum of S , hence S' has rank $n-1$. We know that $j \geq k$. In particular $j \neq k-1$, and hence S_0 is the longest prefix of S' whose last element is different from $k-1$. This shows that $S' \in T'_1$.

It is routine to check that f_1 can be inverted.

Let us now construct f_2 . Choose $S \in T_2$, and write $S = S_0 1^b k$ as above. If S_0 is nonempty, let j be the last element of S_0 . Recall that no element of S_0 is smaller than j , and that j , if defined, is greater than 1 by definition of S_0 . In particular, $S_0 - 1$ is a (possibly empty) sequence of positive numbers. Define $S' = f_2(S) = (S_0 - 1)(k-1)^{b+1}$. The length of S' is equal to the length of S , and the maximum of S' is one less than the maximum of S , hence S' has rank $n-1$. It may be routinely checked that S' avoids 123. Note that the last element of $S_0 - 1$ is smaller than $k-1$, and hence S' belongs to T'_2 .

The inverse of f_2 is easy to obtain. Choose $S' \in T'_2$, with $S' = S'_0 (k-1)^b$, where S'_0 is the longest prefix of S' not ending with $k-1$. As we pointed out earlier, S' must end with the symbol $k-1$, hence $b \geq 1$. Define $S = (S_0 + 1) 1^{b-1} k$. It may be routinely checked that S belongs to T_2 . \square

The following results are direct consequences of Claim 84 and Claim 85. We omit their proofs.

Theorem 86 (J., Mansour [37]). *The number of 1123-avoiding matchings of size n with last element equal to k is equal to the number of Dyck paths of semilength n whose last up-step is followed by k down-steps.* \square

Corollary 87. *The number of 1123-avoiding matchings of size n is $C_n = \frac{1}{n+1} \binom{2n}{n}$. In particular, 1123 is \approx -equivalent to 1212 and to 1221.* \square

From Theorem 86 we may derive the closed-form expression for $t(n, k)$. Since the number of Dyck paths that end with an up-step followed by k down-steps is equal to the number of non-negative lattice paths from $(0, 0)$ to $(2n - k - 1, k - 1)$, we may apply standard arguments for the enumeration of non-negative lattice paths to obtain the formula

$$t(n, k) = \frac{k}{n} \binom{2n - k - 1}{n - 1}.$$

We omit the details of the argument.

Classification of patterns of size 4. Theorem 86 and the general results presented in the previous sections allow us to fully classify patterns of length four by their equivalence classes (see Table 6.2).

τ	$p_n(\tau)$
1213, 1223, 1231, 1232, 1233, 1234	[68, Sequence A007051] (see Equation (6.2))
1123, 1212, 1221	$\frac{1}{n+1} \binom{2n}{n}$ [68, Sequence A000108] (see Theorem 86)
1122	1, 1, 2, 5, 14, 42, 133, 441, ...
1112, 1121, 1211, 1222	[68, Sequence A005425] (see Corollary 67)
1111	[68, Sequence A001680] (see Equation (6.1))

Table 6.2: The numbers $p_n(\tau)$ for $\tau \in \mathcal{P}_4$.

6.6 The pattern 12112

For a full characterization of the equivalence of patterns up to size seven, we need to consider one more sporadic case, namely the pattern 12112. Our aim is to show that this pattern is equivalent to the three patterns 12221, 12212, and 12122. The latter three patterns are all equivalent by Theorem 68. It is thus sufficient to show that $12112 \approx 12212$.

We remark that the proof involving the pattern 12112 does not use the notion of F^\perp -equivalence. In fact, the matrix $M(112, 2)$ is not F^\perp -equivalent to the three F^\perp -equivalent matrices $M(221, 2)$, $M(212, 2)$ and $M(122, 2)$.

The basic idea. The bijection between $\mathcal{P}_n(12112)$ and $\mathcal{P}_n(12212)$, which we are about to construct, is probably the most complicated construction of this thesis. Before we deal with the technical details, we first introduce the basic terminology and notation that we will use throughout the proof, and then outline the key idea of the bijection.

Let $S = s_1 s_2 \cdots s_n$ be a word of length n over the alphabet $[m]$, such that every symbol of $[m]$ appears in S at least once. For $i \in [m]$ let f_i and ℓ_i denote the index of the first and the last symbol of S that is equal to i . Formally, $f_i = \min\{j: s_j = i\}$ and $\ell_i = \max\{j: s_j = i\}$.

Definition 88. For $k \in [m]$, we say that the word S is a *k -semicanonical sequence* (*k -sequence* for short), if S has the following properties.

- For every i and j such that $1 \leq i < k$ and $i < j$, we have $f_i < f_j$.
- For every i and j such that $k \leq i < j \leq m$, we have $l_i < l_j$.

Note that m -semicanonical sequences are precisely the canonical sequences of partitions of $[n]$ with m blocks (i.e., the sequences satisfying $f_i < f_{i+1}$ for $i \in [m-1]$), while the 1-semicanonical sequences are precisely the sequences satisfying $l_i < l_{i+1}$ for $i \in [m-1]$.

Note that for every fixed $k \in [m]$ and a fixed partition $\pi = \pi_1 \cdots \pi_n$ with m blocks, there is exactly one k -sequence $S = s_1 \cdots s_n$ with the property $s_i = s_j \iff \pi_i = \pi_j$. To construct such a k -sequence for a given partition π , we consider the m blocks of the partition, and arrange them into a sequence B_1, B_2, \dots, B_m by the following rules.

- The first $k-1$ blocks B_1, \dots, B_{k-1} are ordered in the increasing order of their minimum elements, in the same way as in the usual canonical representation that we have used so far. The minimum elements of these $k-1$ blocks are smaller than the minimum elements of the remaining blocks.
- The blocks B_k, B_{k+1}, \dots, B_m are ordered in the increasing order of their maximum elements.

A partition of $[n]$ with m blocks B_1, B_2, \dots, B_m ordered by the previous two rules can then be represented by a k -sequence $s_1 s_2 \cdots s_n$ where $s_i = j$ if $i \in B_j$.

In particular, assuming n and m are fixed, the number of k -sequences is independent of k , and each partition of $[n]$ with m blocks is represented by a unique k -sequence. To prove the equivalence $12112 \approx 12212$, we will exploit a remarkable property of the pattern 12112, described by the following key lemma.

Lemma 89 (Key Lemma). *For every fixed n and m , the number of 12112-avoiding k -sequences is independent of k . Thus, for every $k \in [m]$, the number of 12112-avoiding k -sequences of length n with m symbols is equal to the number of 12112-avoiding partitions of n with m blocks.*

Before we prove Lemma 89, let us explain how it implies $12112 \approx 12212$.

Theorem 90 (J., Mansour [37]). *The pattern 12112 is equivalent to 12212. In fact, for every m and n , there is a bijection between 12112-avoiding partitions of $[n]$ with m blocks and 12212-avoiding partitions of $[n]$ with m blocks.*

Proof. Fix m and n . We know that the 12112-avoiding partitions of $[n]$ with m blocks are precisely the m -semicanonical sequences over $[m]$ of length n , and by Lemma 89, these sequences are in bijection with 1-semicanonical 12112-avoiding sequences of the same length and alphabet. It remains to provide a bijection between the 12112-avoiding 1-sequences and the 12212-avoiding partitions.

Take a 1-semicanonical 12112-avoiding sequence S with m symbols and length n , reverse the order of letters in S , and then replace each symbol i of the reverted sequence by the symbol $m-i+1$ (intuitively, we take the sequence S , represented by the matrix $M(S, m)$, and rotate it by 180 degrees). It is easy to check that this transform is an involution which maps 12112-avoiding 1-sequences onto 12212-avoiding m -sequences, which are precisely the 12212-avoiding partitions of $[n]$ with m blocks. \square

It now remains to prove Lemma 89. For the rest of the proof, unless otherwise noted, we will assume that m and n are fixed, and that each sequence we consider has length n and m distinct symbols.

In the following arguments, it is often convenient to represent a sequence $S = s_1 \cdots s_n$ by the matrix $M(S, m)$. Recall that $M(S, m)$ is the 01-matrix with m rows and n columns, with a 1-cell in row i and column j if and only if $s_j = i$. A matrix representing a k -sequence will be called *k-semicanonical matrix* (or just *k-matrix*), and a matrix representing a 12112-avoiding sequence will be simply called *12112-avoiding matrix*. In accordance with earlier terminology, we will use the term *semi-sparse matrix* for a 01-matrix with at most one 1-cell in each column, and we will use the term *semi-standard matrix* for a 01-matrix with exactly one 1-cell in each column. For a 01-matrix M , we let $f_i(M)$ and $\ell_i(M)$ denote the column-index of the first and the last 1-cell in the i -th row of M . We will write f_i and ℓ_i instead of $f_i(M)$ and $\ell_i(M)$ if there is no risk of confusion.

Before we formulate the proof of Lemma 89, let us present a brief sketch of the main idea. We will first build a bijection that transforms a $(k + 1)$ -matrix M into a k -matrix, ignoring 12112-avoidance for a while. Let the last 1-cell in row k of M be in column c , let us call the row k *the key row of M* . If the last 1-cell in row $k + 1$ appears to the right of column c , then M is already a k -matrix and we are done. On the other hand, if row $k + 1$ has no 1-cell to the right of c , we swap the key row k with the row $k + 1$, to obtain a new matrix M' whose key row is now the row $k + 1$. We again check whether the row directly above the key row has a 1-cell to the right of column c , and if not, we swap the rows $k + 1$ and $k + 2$.

We repeat this procedure until we reach the situation when the key row is either the topmost row of the matrix, or the row above the key row has a 1-cell to the right of column c . This procedure transforms the original $k + 1$ matrix into a k -matrix. Also, the procedure is invertible (note that the first 1-cell of the key row is always to the left of any other 1-cell in the rows $k, k + 1, \dots, m$).

Unfortunately, this simplistic approach does not preserve 12112-avoidance. However, we will present an algorithm which follows the same basic structure as the procedure above, but instead of merely swapping the key row with the row above it, it performs a more complicated step, which preserves 12112-avoidance of the matrix. The description of this step is the main ingredient of our proof.

To formalize our argument, we need to introduce more definitions. Let M be a 01-matrix with exactly one 1-cell in each column and at least one 1-cell in each row, and let us write $f_i = f_i(M)$ and $\ell_i = \ell_i(M)$. Let k, p and q be three row-indices of M , with $k \leq p \leq q$. We will say that M is a (k, p, q) -matrix, if M satisfies the following conditions.

- The matrix obtained from M by erasing row p is a k -semicanonical matrix with $m - 1$ rows.
- For each $i < k$, we have $f_i < f_p$. For every $j \geq k, j \neq p$, we have $f_p < f_j$.
- The number q is determined by the relation $q = \max\{j: \ell_j \leq \ell_p\}$. Thus, the first condition implies that $\ell_j \leq \ell_p$ for every $j \in \{k, k + 1, \dots, q\}$.

In a (k, p, q) -matrix, row p will be called *the key row*.

Intuitively, a (k, p, q) -matrix is an intermediate stage of the above-described procedure which transforms a $(k + 1)$ -matrix into a k -matrix by moving the key row towards the top. The number p is the index of the key row in a given step of the procedure, while the number q is the topmost row that needs to be swapped with the key row to produce the required k -matrix. In particular, a matrix M is $(k + 1)$ -semicanonical if and only if it is a (k, k, q) -matrix for some value of q , and M is k -semicanonical if and only if it is a (k, q, q) -matrix for some q .

As an example, consider the sequence $S = 1331232431$ with $n = 10$ and $m = 4$. This sequence corresponds to the following matrix $M = M(S, 4)$.

$$M = \begin{pmatrix} 0000000100 \\ 0110010010 \\ 0000101000 \\ 1001000001 \end{pmatrix} \qquad M' = \begin{pmatrix} 0110010010 \\ 0000000100 \\ 0000101000 \\ 1001000001 \end{pmatrix}$$

The matrix M is a $(2, 3, 4)$ -matrix. If we exchange the third row (which acts as the key row) with the fourth row, we obtain a $(2, 4, 4)$ -matrix M' representing the 2-sequence $S' = 1441242341$. The matrix M' can also be regarded as a $(1, 1, 4)$ -matrix, with the key row at the bottom.

Observe that the following lemma implies Lemma 89.

Lemma 91. *For arbitrary $k \leq p < q$, there is a bijection ϕ between 12112-avoiding (k, p, q) -matrices and 12112-avoiding $(k, p + 1, q)$ -matrices.*

Thus, all we need to do to prove the Key Lemma, and hence also Theorem 90, is to prove Lemma 91.

Before we construct the bijection ϕ , we need to prove several basic properties of the 12112-avoiding (k, p, q) -matrices.

Tools of the proof. Let us introduce some more terminology. If $x \in [m]$ is a row of a matrix M , then an x -column is a column of M that has a 1-cell in row x . Similarly, if $X \subseteq [m]$ is a set of rows of M , we will say that a column j is an X -column if j has a 1-cell in a row belonging to X .

If x, y is a pair of rows of M with $x < y$, we will say that M contains 12112 in (x, y) if the submatrix of M induced by the pair of rows x, y contains 12112. If X and Y are two sets of rows, we will say that M contains 12112 in (X, Y) if there is an $x \in X$ and $y \in Y$ such that $x < y$ and M contains 12112 in (x, y) .

Throughout this paragraph, we will assume that k, p, q are fixed, and that $k \leq p < q$.

We now state a pair of simple but useful observations. Their proofs are straightforward, and we omit them.

Observation 92. *Let M be a semi-sparse 01-matrix, and let $x < y$ be two rows of M , such that $f_x < f_y$. The matrix M avoids 12112 in (x, y) if and only if M has at most one x -column s satisfying $f_y < s < \ell_y$. If such a unique column s exists, we will say that s separates row y . The y -columns that are to the left of the separating column s will be called front y -columns (with respect to row x) and their 1-cells will be called front 1-cells. Similarly, the y -columns to the right of s will be called rear y -columns and their 1-cells are rear 1-cells. If there is no such separating column, then we will assume that all the y -columns and their 1-cells are front. \square*

Observation 93. *Let M be a semi-sparse 01-matrix, and let $x < y$ be a pair of rows such that $\ell_x < \ell_y$. Let t be the number of 1-cells in row x , and let c_i be the i -th x -column, i.e., $f_x = c_1 < c_2 < \dots < c_t = \ell_x$. The matrix M avoids 12112 in (x, y) , if and only if every y -column appears either to the left of column c_1 , or between the columns c_{t-1} and c_t , or to the right of column c_t . These three types of y -columns (and their 1-cells) will be called left, middle, and right y -columns (or 1-cells) with respect to row x . \square*

The following lemma provides a criterion for avoidance of the pattern 12112, which will be useful later in the proof.

Lemma 94. *Let M be a 12112-avoiding (k, p, q) -matrix, and let j be a row of M with $k \leq j \leq p$. Let M' be a semi-sparse 01-matrix of the same size as M , with the property that for every $i \notin \{j, j+1, \dots, q\}$, the i -th row of M is equal to the i -th row of M' . If M' has a copy of the pattern 12112 in a pair of rows $x < y$, then $j \leq x \leq q$.*

Proof. Let M and M' be as above. We will call the rows $\{j, j+1, \dots, q\}$ *mutable*, and the remaining rows will be called *constant*.

Assume that M' has a copy of 12112 in the rows $x < y$. Clearly, at least one of the two rows x, y must be mutable, and in particular, we must have $x \leq q$. The lemma claims that x must be mutable. For contradiction, assume that $x < j$. This implies that y is mutable. We distinguish two possibilities; either $x < k$ or $k \leq x < j$.

Assume that $x < k$. From the definition of the (k, p, q) -matrix, we obtain that all the columns of M to the left of $f_p(M)$ and to the right of $\ell_p(M)$ contain a 1-cell in one of the constant rows. Since M' is semi-sparse, we conclude that in M' , all the 1-cells in the mutable rows can only appear in the columns i such that $f_p(M) \leq i \leq \ell_p(M)$.

Now, we apply Observation 92 to the rows x and p in the matrix M , and conclude that M (and hence also M') has at most one x -column s such that $f_p(M) \leq s \leq \ell_p(M)$. Therefore M' also has at most one x -column between $f_y(M')$ and $\ell_y(M')$. By Observation 92, this shows that x cannot form the pattern 12112 with any of the mutable rows y of M' .

Assume now that $k \leq x < j$. As before, we have $y \in \{j, \dots, q\}$. Let $c_1 < c_2 < \dots < c_t$ be the x -columns of M (and hence of M' as well, since x is constant). For any mutable row i , we have $\ell_x(M) < \ell_i(M)$ by the definition of (k, p, q) -matrix. By Observation 93, all of the i -columns of M appear either to the left of c_1 or to the right of c_{t-1} . In particular, all the 1-cells between the columns c_1 and c_{t-1} belong to the constant rows. This implies that M' can have no occurrence of 12112 in the two rows $x < y$. \square

We will now describe a simple operation, called *pseudoswap*, on 12112-avoiding pairs of rows.

Assume that M is a semi-sparse matrix with a pair of adjacent rows x and $y = x + 1$ that avoids 12112 in (x, y) . Assume furthermore that $f_x < f_y \leq \ell_y < \ell_x$. The pseudoswap of the two rows is performed as follows.

Easy case. If the row y is not separated by an x -column (in the sense of Observation 92), or if M has at most one rear y -column with respect to row x , the pseudoswap is performed by simply swapping the two rows of M .

Hard case. Assume M has an x -column s separating y , and that it has $r > 1$ rear y -columns $c_1 < c_2 < \dots < c_r$ (see Figure 6.3). In this case, the pseudoswap preserves the position of all the 1-cells in columns c_1, \dots, c_{r-1} (i.e., the 1-cells in these columns remain in row y), and all the other 1-cells in rows x, y are moved from x to y and vice versa. Note that after the pseudoswap is performed, the columns $s < c_1 < c_2 < \dots < c_{r-1}$ all contain a 1-cell in row y , and these r 1-cells are precisely the middle 1-cells of y with respect to x (in the sense of Observation 93).

Let M' be the matrix obtained from M by the pseudoswap. It can be routinely checked that M' avoids 12112 in (x, y) . Let us write f'_i for $f_i(M')$ and ℓ'_i for $\ell_i(M')$. Clearly, $f'_x = f_y$ and $f'_y = f_x$, and also $\ell'_x = \ell_y$ and $\ell'_y = \ell_x$. Also, if row y of M has $r \geq 0$ rear cells with respect to row x , then in M' , row y has r middle cells with respect to x .

It is not difficult to see that the pseudoswap can be inverted. Let M' be a sparse matrix avoiding 12112 in two adjacent rows $x < y$, such that $f'_y < f'_x \leq \ell'_x < \ell'_y$. If M' has fewer than two middle y -columns, we invert the easy case of the pseudoswap by exchanging the




Figure 6.3: The illustration of the hard case of the pseudoswap operation.

two rows. On the other hand, if M' has $r > 1$ middle y -columns $m_1 < \dots < m_r$, we invert the hard case by preserving the position of the 1-cells in columns m_2, m_3, \dots, m_r and inverting all the other $\{x, y\}$ -columns.

We will be mostly interested in the situation when the pseudoswap is applied to the pair of rows $(p, p + 1)$ in a (k, p, q) -matrix with $p < q$. It is not hard to see that this operation yields a $(k, p + 1, q)$ -matrix. Unfortunately, under some circumstances, the hard case of the pseudoswap may create a copy of the pattern 12112 in the remaining rows of the matrix. Thus, the pseudoswap alone is not sufficient to provide the required bijection between 12112-avoiding (k, p, q) -matrices and 12112-avoiding $(k, p + 1, q)$ -matrices.

Let us now look in more detail at the hard case of the pseudoswap. Recall that if X and Y are two sets of rows of M , we say that M avoids 12112 in (X, Y) , if there is no $x \in X$ and $y \in Y$ such that $x < y$ and the two rows x, y contain a copy of 12112.

The following lemma is illustrated in Figure 6.4.

Lemma 95. (a) Let M be a (k, p, q) -matrix that avoids 12112 in $(p, p + 1)$. Let $f_p(M) = b_1 < b_2 < \dots < b_t = \ell_p(M)$ be the p -columns of M . Assume that the row $p + 1$ is separated by the column b_i , and that it has $r \geq 2$ rear 1-cells. Let $c_1 < c_2 < \dots < c_s$ be the front $(p + 1)$ -columns and let $d_1 < d_2 < \dots < d_r$ be the rear $(p + 1)$ -columns. By Observation 92, we have the inequalities

$$b_1 < \dots < b_{i-1} < c_1 < \dots < c_s < b_i < d_1 < \dots < d_r < b_{i+1} < \dots < b_t.$$

Let $X = \{p, p + 1\}$ and let Y be the set of all the rows above $p + 1$ that contain at least one 1-cell to the left of the column d_{r-1} ; formally,

$$Y = \{y > p + 1 : f_y(M) < d_{r-1}\}.$$

The matrix M avoids 12112 in (X, Y) if and only if each Y -column y satisfies one of the following three inequalities:

1. $b_{i-1} < y < c_1 = f_{p+1}$
2. $d_{r-1} < y < d_r$
3. $d_r < y < b_{i+1}$

The rows in Y are precisely the rows above $p + 1$ that are separated by the p -column b_i .

(b) Let M' be a $(k, p + 1, q)$ -matrix that avoids 12112 in $(p, p + 1)$. Let $\alpha_1 < \dots < \alpha_u < \beta_1 < \dots < \beta_r < \gamma_1 < \dots < \gamma_v$ be the $(p + 1)$ -columns of M' , where the α_i , β_i and γ_i denote respectively the left, middle and right $(p + 1)$ -columns with respect to row p . Assume that there are at least two middle 1-cells. Let $\delta_1 < \dots < \delta_w$ be the p -columns of M' . By Observation 93, we have the inequalities

$$\alpha_1 < \dots < \alpha_u < \delta_1 < \dots < \delta_{w-1} < \beta_1 < \dots < \beta_r < \delta_w < \gamma_1 < \dots < \gamma_v.$$

Let $X = \{p, p+1\}$ and let Y' be the set of all the rows above $p+1$ that contain at least one 1-cell to the left of column β_r . The matrix M' avoids 12112 in (X, Y') if and only if each Y' -column y satisfies one of the following three inequalities:

1. $\beta_{r-1} < y < \beta_r$
2. $\beta_r < y < \delta_w$
3. $\delta_w < y < \gamma_1$

The rows in Y' are precisely the rows above $p+1$ that are separated by the $(p+1)$ -column β_r .




Figure 6.4: Illustration of Lemma 95: part (a) is above, part (b) below. The black dots correspond to 1-cells in rows p and $p+1$, and the grey rectangles correspond to possible positions of the 1-cells in the rows of Y or Y' .

Proof. Let us consider part (a). Fix a row $y \in Y$. By the definition of a (k, p, q) -matrix, we have $d_r < \ell_y$. By Observation 93, we see that M avoids 12112 in $(p+1, y)$ if and only if every y -column j satisfies either $j < c_1 = f_{p+1}$, $d_{r-1} < j < d_r$, or $j > d_r$. The first y -column satisfies $f_y < d_{r-1}$ by the definition of Y , and hence $f_y < c_1 = f_{p+1} < b_i$. Since $\ell_y > \ell_{p+1} = d_r > b_i$, we see that if the pair of rows $(p+1, y)$ avoids 12112, then y is separated by b_i and the two rows (p, y) avoid 12112 if and only if $b_{i-1} < f_y < \ell_y < b_{i+1}$. This proves part (a) of the lemma.

The proof of part (b) is analogous and we omit it. \square

The bijection. We are now ready to present the bijection ϕ , promised in Lemma 91. Let M be a 12112-avoiding (k, p, q) -matrix with $p < q$, and let us write f_i and ℓ_i for $f_i(M)$ and $\ell_i(M)$. By the definition of (k, p, q) -matrix and by the assumption $p < q$, we know that $f_p < f_{p+1} \leq \ell_{p+1} < \ell_p$, so we may perform the pseudoswap of the rows p and $p+1$ in M . Let M' be the $m \times n$ matrix obtained from M by this pseudoswap. Let $f'_i = f_i(M')$ and $\ell'_i = \ell_i(M')$. Note that $f'_i = f_i$ and $\ell'_i = \ell_i$ for every $i \notin \{p, p+1\}$.

We already know that M' is a $(k, p+1, q)$ -matrix. We now distinguish two cases, depending on whether the pseudoswap we performed was easy or hard.

Easy case. If the row $p+1$ of M has at most one rear 1-cell with respect to row p , then M' is 12112-avoiding, and we may define $\phi(M) = M'$. Indeed, from the definition of the pseudoswap we know that M' cannot contain a copy of 12112 in the rows $(p, p+1)$,

and since we are performing the easy case of the pseudoswap, we cannot create any new copy of the forbidden pattern that would intersect the remaining $m - 2$ rows.

Hard case. Assume that the row $p + 1$ of M has $r > 1$ rear 1-cells. Let $b_1 < \dots < b_t$, $c_1 < \dots < c_s$, $d_1 < d_2 < \dots < d_r$, and Y have the same meaning as in part (a) of Lemma 95. Let Y_1 , Y_2 and Y_3 denote, respectively, the Y -columns that lie between b_{i-1} and c_1 , between d_{r-1} and d_r , and between d_r and b_{i+1} .

The bijection ϕ is now constructed in two steps. In the first step, we perform the pseudoswap of the rows p and $p + 1$. Let M' be the result of this first step. Let us now apply the notation of part (b) of Lemma 95 to the matrix M' (see Figure 6.4). Note that $d_{r-1} = \beta_r$, and hence $Y = Y'$. Part (b) of Lemma 95 requires that all the Y' -columns of a 12112-avoiding $(k, p + 1, q)$ -matrix fall into one of the three groups:

- columns between $\delta_w < y < \gamma_1$. In M' , we have $\delta_w = d_r$ and $\gamma_1 = b_{i+1}$, so these columns are precisely the columns in Y_3 .
- columns between $\beta_r < y < \delta_w$. In M' , these are precisely the columns in Y_2 .
- columns between $\beta_{r-1} < y < \beta_r$. In M' , there are no Y -columns in this range.

On the other hand, if Y_1 is nonempty, then these columns violate the inequalities of part (b) in Lemma 95, showing that M' is not 12112-avoiding. To correct this, we apply the second step of the bijection ϕ . Consider the submatrix of M' induced by the columns Y_1 and the columns $Z = \{\delta_1 < \dots < \delta_{w-1} < \beta_1 < \dots < \beta_{r-1}\}$. Note that the columns Y_1 are to the left of any column of Z . Now we rearrange the columns inside this submatrix, so that all the columns in Y_1 appear after the columns in Z , keeping the relative order of the columns in Y_1 , as well as those in Z . This transforms M' into a matrix M'' . We define $\phi(M) = M''$.

Since M'' is clearly a $(k, p + 1, q)$ -matrix, it remains to check that M'' avoids 12112. Let $x < y$ be a pair of rows of M'' . We want to check that M'' avoids 12112 in these two rows. Let us consider the following cases separately.

The case $x < p$. The rows below row p are unaffected by ϕ . The rows above row q are preserved as well, because any row $z \in Y$ must satisfy $\ell_z < b_{i+1} \leq \ell_p$, so no row above q belongs to Y . Thus, we may apply Lemma 94, to see that M'' avoids 12112 in the rows (x, y) .

The case $x = p, y = p + 1$. The properties of pseudoswap guarantee that M'' avoids 12112 in these two rows.

The case $x \in X = \{p, p + 1\}$ and $y \in Y'$. By construction, M'' satisfies the inequalities of part (b) of Lemma 95, and thus it avoids 12112 in (X, Y) .

The case $x \in X = \{p, p + 1\}$, $y \notin Y'$ and $y > p + 1$. By the definition of Y' , we have $f_y(M'') = f_y(M) > d_{r-1} = \beta_r$. In any column to the right of β_r the mapping ϕ acts by exchanging the rows p and $p + 1$. It is easy to check that this action cannot create a copy of 12112 in (x, y) (note that in any of the three matrices M , M' and M'' , both the rows p and $p + 1$ have a 1-cell to the left of β_r).

The case $y > x > p + 1$. The submatrix of M'' induced by the rows above $p + 1$ only differs from the corresponding submatrix of M by the position of the zero columns. Thus, it cannot contain any copy of 12112.

This shows that $\phi(M)$ is indeed a 12112-avoiding $(k, p + 1, q)$ -matrix.

It is routine to check that the mapping ϕ can be inverted, and by a case analysis similar to the arguments above, it turns out that the inverse of ϕ preserves 12112-avoidance. This shows that ϕ is indeed the required bijection.

This completes the proof of Lemma 91, from which, and as we explained before, Lemma 89 and Theorem 90 follow directly.




Figure 6.5: The ordered graphs G_1 and G_2 corresponding to the filling patterns $M(2, 112)$ and $M(2, 212)$.

Consequences. Theorem 90 has several consequences for pattern-avoiding fillings of F^\perp -shapes and for pattern-avoiding ordered graphs.

By Lemma 46, there is a bijection between 12112-avoiding partitions of $[n]$ with m blocks and $M(2, 112)$ -avoiding semi-standard fillings of F^\perp -shapes with $n - m$ columns and at most m rows. Similarly, there is an analogous bijection between 12212-avoiding partitions and $M(2, 212)$ -avoiding fillings of F^\perp -shapes. Thus, we obtain the following direct consequence of Theorem 90.

Corollary 96. *For every r and c , there is a bijection between the $M(2, 112)$ -avoiding semi-standard fillings of all the F^\perp -shapes with r rows and c columns and the $M(2, 212)$ -avoiding semi-standard fillings of all the F^\perp -shapes with r rows and c columns.*

It would be tempting to assume that for a given F^\perp -shape F , the $M(2, 112)$ -avoiding semi-standard fillings of F are in bijection with the $M(2, 212)$ -avoiding semi-standard fillings of F , i.e., that the two matrices $M(2, 112)$ and $M(2, 212)$ are F^\perp -equivalent. However, as we already mentioned in the introduction of Section 6.6, this is not the case. For instance, the F^\perp -shape F with five columns of height 4 and one column of height 2 has 866 $M(2, 112)$ -avoiding fillings but only 865 $M(2, 212)$ -avoiding fillings. Thus, the bijection of Corollary 96 in general cannot preserve the shape of the underlying diagram.

In the introduction of this thesis, we described a one-to-one correspondence between dense fillings of Ferrers shapes and IM-free ordered graph. By applying the same idea, we may obtain a one-to-one correspondence between semi-standard fillings of F^\perp -shapes and sprinkler graphs. For convenience, let us describe the correspondence here.

Recall that a sprinkler graph is an ordered graph, in which every vertex has either exactly one neighbor to its left, or an arbitrary number (possibly zero) of neighbors to its right.

Every semi-standard filling Φ of an F^\perp -shape with c columns and r rows can be represented by an ordered graph with $c + r$ linearly ordered vertices. The graph has two kinds of vertices: the *right vertices* r_1, \dots, r_c , which have degree one, and are to the right of their neighbors, and the *left vertices* ℓ_1, \dots, ℓ_r , which may have arbitrary degree one and are to the left of all their neighbors.

The i -th column of Φ is associated with the i -th right vertex r_i , and the j -th row of Φ is associated with the j -th left vertex ℓ_j . All the vertices are linearly ordered by a left-to-right relation $<$ with the properties $r_1 < \dots < r_c$, $\ell_1 < \ell_2 < \dots < \ell_r$, and furthermore, $\ell_j < r_i$ if and only if row j intersects column i inside Φ . A 1-cell in row j and column i corresponds to an edge between ℓ_j and r_i . Note that if ℓ_j and r_i are connected by an edge, then $\ell_j < r_i$.

In this representation, the semi-standard fillings of F^\perp -shapes correspond precisely to the sprinkler graphs. The $M(112, 2)$ avoiding fillings of F correspond precisely to ordered graphs which avoid a subgraph G_1 with five vertices $\ell_a < \ell_b < r_x < r_y < r_z$ and three

edges $\ell_a r_x$, $\ell_a r_y$, and $\ell_b r_z$. Similarly, the fillings avoiding $M(212, 2)$ correspond to graphs avoiding the subgraph G_2 with vertices $\ell_a < \ell_b < r_x < r_y < r_z$ and edges $\ell_a r_y$, $\ell_b r_x$, and $\ell_b r_z$ (see Figure 6.5).

Theorem 90 then immediately yields the following result.

Corollary 97. *There is a bijection between G_1 -avoiding sprinkler graphs and G_2 -avoiding sprinkler graphs that preserves the number of left vertices and right vertices.*

Whether this result can be extended to more general types of graphs or more general pairs of patterns is an open problem.

Conclusion: Beyond Wilf-type classifications

Wilf-type classification has been one of the main topics of the study of pattern-avoiding ordered structures. However, it would be unwise to over-emphasize its importance and ignore alternative approaches. In fact, it seems that lately these alternative approaches are gaining increasing amount of attention.

For this reason, we will devote the concluding chapter of this thesis to a brief overview of several promising approaches which are currently being actively pursued, and are likely to bring new results in near future.

Hereditary classes

Throughout this thesis, we have mostly considered classes of structures that avoided a single forbidden pattern. This makes good sense in the context of Wilf-type classification, but it appears to be a somewhat artificial restriction once we begin to study other combinatorial aspects. It often turns out that a more natural concept is the concept of hereditary classes. A class \mathcal{C} of permutations is called *hereditary*, if for every $\tau \in \mathcal{C}$, the class \mathcal{C} contains all subpermutations of τ . A hereditary class is *proper* if it does not contain all permutations.

Of course, we may analogously define hereditary classes of other structures than permutations. The notion of hereditary class makes sense for any family of structures ordered by a containment relation.

Recall that \mathcal{S}_n is the set of all the permutations of order n , and $\mathcal{S}_n(T)$ is the set of permutations of order n that avoid all the patterns from the set T . Let us write

$$\mathcal{S} = \bigcup_{n=1}^{\infty} \mathcal{S}_n \quad \text{and} \quad \mathcal{S}(T) = \bigcup_{n=1}^{\infty} \mathcal{S}_n(T).$$

For any set of patterns T , the class $\mathcal{S}(T)$ is clearly hereditary. Conversely, any hereditary class $\mathcal{C} \subseteq \mathcal{S}$ can be expressed as $\mathcal{C} = \mathcal{S}(T)$, where T is the (possibly infinite) set of minimal permutations that do not belong to \mathcal{C} . The set T is called *the basis* of \mathcal{C} . For a hereditary class \mathcal{C} , we let \mathcal{C}_n denote the set $\mathcal{C} \cap \mathcal{S}_n$. The function $n \mapsto |\mathcal{C}_n|$ is known as the *speed* of the class \mathcal{C} .

Growth Rates

The most natural task in the study of pattern avoidance is to determine how many permutations of a given order a pattern-avoiding class \mathcal{C} contains.

One of the most significant achievements in this line of research is the following result due to Marcus and Tardos [50].

Theorem 98 (Marcus–Tardos Theorem (formerly Stanley–Wilf Conjecture)). *For a proper hereditary class of permutations \mathcal{C} there is a constant c such that $|\mathcal{C}_n| \leq c^n$.*

By Theorem 98, we may define, for a proper hereditary class \mathcal{C} , its *upper growth rate* $\overline{\text{gr}}(\mathcal{C})$ as $\limsup_{n \rightarrow \infty} \sqrt[n]{|\mathcal{C}_n|}$, and its *lower growth rate* $\underline{\text{gr}}(\mathcal{C})$ as $\liminf_{n \rightarrow \infty} \sqrt[n]{|\mathcal{C}_n|}$. If the upper and lower growth rates are the same, we speak simply of the *growth rate* $\text{gr}(\mathcal{C})$.

Arratia [5] has shown, by a superadditivity argument, that if the basis of \mathcal{C} consists of a single element τ , then $\overline{\text{gr}}(\mathcal{C}) = \underline{\text{gr}}(\mathcal{C})$; in such case, the growth rate is often referred to as the Stanley–Wilf limit of τ . It is an open question whether the growth rate exists for any hereditary class \mathcal{C} of permutations.

There are several results related to the evaluation of the Stanley–Wilf limits of a specific pattern. We have already mentioned that Regev [55] gave an asymptotic formula for the number of I_k -avoiding permutation, which implies that the Stanley–Wilf limit of $\mathcal{S}(I_k)$ is equal to $(k-1)^2$.

In 1997, Bóna [11] has found a formula for $|\mathcal{S}_n(1342)|$, and showed that the Stanley–Wilf limit of 1342 is equal to 8. In 2005, Bóna [13] has shown that the Stanley–Wilf limit of 12453 is equal to $9 + 4\sqrt{2}$. This was the first example of a pattern whose Stanley–Wilf limit is not an integer. It is conjectured [70] that the growth rate of a finitely-based permutation class is either an integer or an algebraic irrational.

In general, the problem of finding the Stanley–Wilf limit of a given pattern seems rather difficult. For instance, the Stanley–Wilf limit of 1324 is still not known. The best known lower bound is due to Albert et al. [2] who showed that $\text{gr}(\mathcal{S}(1324)) \geq 9.35$. This result shows that the Stanley–Wilf limit of 1324 is the largest among the permutations of order four.

A related question is to estimate the largest and the smallest Stanley–Wilf limit of a pattern $\tau \in \mathcal{S}_k$, as a function of k . The proof of the Marcus–Tardos theorem gives an upper bound which is superexponential in terms of k , but this is believed to be far from optimal. The largest known Stanley–Wilf limit is due to Bóna [12], who constructed, for any k , a pattern of size $3k+1$ whose Stanley–Wilf limit is equal to $k^2 \text{gr}(\mathcal{S}(1324))$. For the minimum possible Stanley–Wilf limit, an argument of Valtr [39] shows that any pattern of size k has the Stanley–Wilf limit at least $(1 + o(1))(k-1)^2/e^3$. The smallest known Stanley–Wilf limit was again obtained by Bóna [12], whose method yields a pattern of length $3k+1$ and Stanley–Wilf limit $8k^2$.

Speeds of hereditary classes

Recently, several researchers have considered the question of determining general criteria for the functions that can be obtained as speeds of hereditary classes.

For hereditary permutation classes, the papers by Kaiser and Klazar [39], Albert and Linton [3] and Vatter [69] have the yielded the following results:

- If the speed of a hereditary class \mathcal{C} is asymptotically smaller than $(1 + \epsilon)^n$ for every $\epsilon > 0$, then the speed is eventually equal to a polynomial [39].
- If the growth rate of a hereditary permutation class is less than 2, then the growth rate is equal to the positive root of $1 + x + x^2 + \dots + x^{k-1} - x^k$, for some integer k [39].
- There is a constant $\kappa \approx 2.20557$ (the unique positive root of $1 + 2x^2 - x^3$) such that there are only countably many hereditary classes of growth rate smaller than κ , while there are uncountably many classes of growth rate κ [43, 70].

- There is a constant $\lambda \approx 2.48187$ (the unique real root of $x^5 - 2x^4 - 2x^2 - 2x - 1$) such that for every real number $c \geq \lambda$ there is a hereditary permutation class of growth rate c [3, 69].

Notice that these results yield a dichotomy between hereditary classes of polynomial speed and hereditary classes of exponential speed. For growth rates up to κ , there is still a discrete hierarchy of possible speeds. However, after a certain threshold, there are no more gaps between the possible growth rates, and any growth rate is allowed. This cancels any hope of a ‘nice’ hierarchy of all the permutation growth rates. For this reason, it has been suggested to study more restricted classes of permutations, e.g., the hereditary classes with finite base. So far, however, we are not aware of any substantial result in this restricted setting.

Ordered graphs

The results on the speeds of permutation classes can be, to a great extent, generalized to hereditary properties of ordered graphs. A *property* of graphs is a class of graphs that is closed under isomorphism. A property of ordered graph is *hereditary* if it is closed under taking induced ordered subgraphs. For a property \mathcal{C} of ordered graphs, let \mathcal{C}_n be the number of graphs from \mathcal{C} on the vertex set $[n]$. The *speed* of a property \mathcal{C} is the function $n \mapsto |\mathcal{C}_n|$. Note that no two distinct ordered graphs on the vertex set $[n]$ are isomorphic, so the speed in fact counts the number of isomorphism classes of order n in \mathcal{C} .

Since a permutation $\tau \in \mathfrak{S}_n$ can be represented by its permutation graph G_τ , which is an ordered graph on the vertex set $[n]$, and since any induced subgraph of G_τ is (up to isomorphism) a permutation graph that represents a subpermutation of τ , we see that a hereditary class \mathcal{C} of permutations can be represented by a hereditary property $G_{\mathcal{C}} = \{G_\tau : \tau \in \mathcal{C}\}$, which has the same speed as \mathcal{C} . This means that the possible speeds of hereditary properties of ordered graphs are a superset of the possible speeds of hereditary permutation classes.

The study of possible speeds of hereditary properties of ordered graphs was initiated by Balogh, Bollobás and Morris. They were motivated by results on permutation speeds mentioned above, as well as previous results on speeds of classes of (unordered) labelled graphs, obtained (among others) by Balogh, Bollobás and Weinreich [8, 9].

In the setting of ordered graphs, Balogh, Bollobás and Morris have obtained the following result.

Theorem 99 (Balogh et al. [7]). *Let \mathcal{C} be a hereditary property of ordered graphs, and let $f(n) = |\mathcal{C}_n|$ be the speed of \mathcal{C} . One of the following conditions holds:*

- There are constants n_0 and K such that $f(n) = K$ for every $n \geq n_0$.
- There are integers $n_0, k \geq 1, a_0, \dots, a_k$ such that for every $n \geq n_0$, we have $f(n) = \sum_{i=0}^k a_i \binom{n}{i}$.
- There is an integer $k \geq 2$ and a polynomial p , such that $F_n^{(k)} \leq f(n) \leq p(n)F_n^{(k)}$. Here $F_n^{(k)}$ are the k -Fibonacci numbers, defined by the recurrence $F_n^{(k)} = \sum_{i=1}^k F_{n-i}^{(k)}$ for $n \geq 1$, with the initial conditions $F_0^{(k)} = 1$, and $F_n^{(k)} = 0$ for $n < 0$.
- The inequality $f(n) \geq 2^{n-1}$ holds for every $n \in \mathbb{N}$.

We remark that for any k , the growth rate of the k -Fibonacci sequence $\lim_{n \rightarrow \infty} \sqrt[n]{F_n^{(k)}}$ is the positive root of $1 + x + x^2 + \dots + x^{k-1} - x^k$, and is strictly smaller than 2.

Atomic and molecular relational structures

The attempts to generalize hereditary permutation classes need not stop with ordered graphs. A promising approach to the understanding of ordered structures is based on the theory of hereditary and atomic classes of relational structures.

Since we have not mentioned relational structures before, let us briefly introduce the necessary terminology. Let $\Sigma = (a(1), a(2), \dots, a(k))$ be a finite sequence of positive integers and let V be a (possibly infinite) set. A *relational structure with signature Σ on the vertex set V* is a $(k + 1)$ -tuple $R = (V, E_1, \dots, E_k)$, where $E_i \subseteq V^{a(i)}$. The sets E_i are called *the relations* of R and the integer $a(i)$ is called *the arity* of E_i . The size of the vertex set V is referred to as *the order* of R .

A relational structure whose all relations have arity 2 is called *binary relational structure*.

Let $R = (V, E_1, \dots, E_k)$ and $R' = (W, F_1, \dots, F_k)$ be two relational structures of the same signature. We say that R and R' are *isomorphic*, if there is a bijection $\phi: V \rightarrow W$ such that for every i and every $a(i)$ -tuple $(v_1, \dots, v_{a(i)}) \in V^{a(i)}$ we have the equivalence $(v_1, \dots, v_{a(i)}) \in E_i \iff (\phi(v_1), \dots, \phi(v_{a(i)})) \in F_i$. We say that R' is a substructure of R if $W \subset V$ and for every i , $F_i = E_i \cap W^{a(i)}$. We say that R *contains* R' if R' is isomorphic to a substructure of R . For a relational structure R , the *age* of R is the set of all the finite relational structures that are contained in R .

A *property of relational structures* is a class of finite relational structures of the same signature that is closed under isomorphism. A property is hereditary, if it closed under taking substructures. A property is *atomic* if it is equal to the age of a single (possibly infinite) relational structure. If a property is the union of finitely many atomic properties, we call it *molecular*.

Atomic classes have been introduced by Fraïssé [25], who has also shown that a hereditary property \mathcal{C} of relational structures is atomic if and only if each two elements R and S of \mathcal{C} are *jointly embeddable*, which means that there is an element $T \in \mathcal{C}$ that contains both R and S . Another equivalent definition of an atomic property states that a hereditary property \mathcal{C} is atomic if and only if it cannot be expressed as the union of two proper hereditary properties different from \mathcal{C} .

For molecular properties, we are able to give a similar characterization in terms of joint embeddability [36]. For an integer k , a hereditary property \mathcal{C} cannot be expressed as the union of k atomic properties if and only if it contains a set of $k + 1$ structures, no two of which are jointly embeddable in \mathcal{C} . A property is not molecular, if and only if it contains an infinite such subset.

The *speed* of a hereditary class \mathcal{C} is the function that assigns to an integer n the number of nonisomorphic structures of order n in \mathcal{C} .

Relational structures are very general concept; in particular, hereditary classes of relational structures generalize hereditary classes of ordered graphs, hereditary classes of set partitions (with various containment relations), as well as many other structures. It thus makes good sense to study their possible speeds.

The research related to relational structures seems to have proceeded independently of the research of the research of hereditary classes of permutations and graphs. In fact, in the context of relational structures, people have more often considered the speeds of atomic classes, rather than general hereditary classes. We are not aware of any results that would deal with the full generality of hereditary properties of relational structures.

Nevertheless, several results on the speeds of atomic classes of relational structures exist [54], and they appear similar to the results related to speeds of hereditary classes of more specific objects. Thus, it seems plausible that there might be a common generaliza-

tion of these lines of research.

It is also remarkable that the classes of ordered graphs or speed smaller than 2^n that occur in the classification of Theorem 99, as well as the permutation classes of growth rate smaller than the countability threshold κ , are in fact molecular classes of structures. This suggests that the speeds of molecular classes of structures might be easier to handle and their growth rates might be more constrained than the growth rates of general hereditary classes.

Further reading

Obviously, we only provided a sketchy and incomplete overview of topics related to hereditary classes of permutations and other structures. For the benefit of an interested reader, we provide several references to more thorough surveys of these topics.

The survey of Kitaev and Mansour [40] deals mostly with the topic of pattern avoidance in permutations and words. It also deals with several alternative notions of pattern avoidance in permutations.

A more general (and also more recent) survey by Klazar [45] deals, among other topics, with growth rates of hereditary properties. Another survey, by Bollobás [10], is exclusively devoted to growth rates of hereditary and monotone properties of combinatorial structures.

The survey of Pouzet [54] deals with atomic classes of relational structures, including an overview of the main algebraic and order-theoretic tools used in their study.

Acknowledgement

I am grateful to Martin Klazar for his patient and inspiring guidance and a lot of valuable advice throughout my PhD study and during my work on this thesis.

Bibliography

- [1] M. ALBERT, R. ALDRED, M.D. ATKINSON, C. HANDLEY, AND D. HOLTON, Permutations of a multiset avoiding permutations of length 3, *Eur. J. Combin.* **22** (2001), 1021–1031.
- [2] M. ALBERT, M. ELDER, A. RECHNITZER, P. WESTCOTT, AND M. ZABROCKI, A lower bound on the growth rate of the class of 4231 avoiding permutations, *Adv. Appl. Math.* (in press).
- [3] M. ALBERT AND S. LINTON, Growing at a perfect speed, *preprint*.
- [4] N. ALON AND E. FRIEDGUT, On the number of permutations avoiding a given pattern, *J. Comb. Theory A* **89** (2000), 133–140.
- [5] R. ARRATIA, On the Stanley–Wilf conjecture for the number of permutations avoiding a given pattern, *Elect. J. Combin.*, **6** (1999), #N1.
- [6] J. BACKELIN, J. WEST, AND G. XIN, Wilf-equivalence for singleton classes, *Adv. Appl. Math.* **32:2** (2007), 133–148.
- [7] J. BALOGH, B. BOLLOBÁS, AND R. MORRIS, Hereditary properties of ordered graphs, in: M. Klazar, J. Kratochvíl, M. Loeb, J. Matoušek, R. Thomas, and P. Valtr (eds.) *Topics in Discrete Mathematics*, Springer, Heidelberg (2006), pp. 179–213.
- [8] J. BALOGH, B. BOLLOBÁS, AND D. WEINREICH, The penultimate rate of growth for graph properties, *Eur. J. Combin.* **22** (2001), 277–289.
- [9] J. BALOGH, B. BOLLOBÁS, AND D. WEINREICH, The speed of hereditary properties of graphs, *J. Comb. Theory B* **79** (2000), 131–156.
- [10] B. BOLLOBÁS, Hereditary and monotone properties of combinatorial structures, in: A. Hilton and J. Talbot (eds.) *Surveys in Combinatorics 2007*, Cambridge University Press (2007), pp. 1–40.
- [11] M. BÓNA, Exact enumeration of 1342-avoiding permutations: A close link with labeled trees and planar maps, *J. Comb. Theory A* **80** (1997), 257–272.
- [12] M. BÓNA, New records in Stanley–Wilf limits, *Eur. J. Combin.* **28** (2007), 75–85.
- [13] M. BÓNA, The limit of a Stanley–Wilf sequence is not always rational, and layered patterns beat monotone patterns, *J. Comb. Theory A* **110(2)** (2005), 223–235.
- [14] M. BOUSQUET-MÉLOU AND E. STEINGRÍMSSON, Decreasing subsequences in permutations and Wilf equivalence for involutions, *J. Algebr. Comb.* **22** (2005), 383–409.

- [15] A. BURSTEIN, Enumeration of words with forbidden patterns, Ph.D. thesis, University of Pennsylvania, 1998.
- [16] A. BURSTEIN AND T. MANSOUR, Words restricted by patterns with at most 2 distinct letters, *Electron. J. Comb.* **9:2** (2002), #R3.
- [17] W. Y. C. CHEN, E. Y. P. DENG, AND R. R. X. DU, Reduction of m -regular noncrossing partitions, *Eur. J. Combin.* **26:2** (2005), 237–243.
- [18] W. Y. C. CHEN, E. Y. P. DENG, R. R. X. DU, R. P. STANLEY, AND C. H. YAN, Crossings and nestings of matchings and partitions, *T. Am. Math. Soc.* **359(4)** (2007), 1555–1575.
- [19] A. DE MIER, k -noncrossing and k -nonnesting graphs and fillings of Ferrers diagrams, *Combinatorica* **27(6)** (2007), 699–720.
- [20] A. DE MIER, On the symmetry of the distribution of k -crossings and k -nestings in graphs, *Electron. J. Comb.* **13** (2006), #N21.
- [21] W. M. B. DUKES, T. MANSOUR, AND A. REIFEGERSTE, Wilf classification of three and four letter signed patterns, *Discrete Math.* **308(15)** (2008), 3125–3133.
- [22] W. M. B. DUKES, V. JELÍNEK, T. MANSOUR, AND A. REIFEGERSTE, New equivalences for pattern avoiding involutions, *P. Am. Math. Soc.* (in press).
- [23] S.-P. EU, T.-S. FU, AND C.-J. LAI, On the enumeration of parking functions by leading terms, *Adv. Appl. Math.* **35(4)** (2005), 392–406.
- [24] P. FLAJOLET AND R. SEDGEWICK, Analytic Combinatorics, Part A, Chapter II; available at <http://algo.inria.fr/flajolet/Publications/books.html>, 2006.
- [25] R. FRAÏSSÉ, Sur l’extension aux relations de quelques propriétés des ordres, *Ann. Sci. Ecole Norm. Sup.* **71(3)** (1954), 363–388.
- [26] S. HEUBACH, T. MANSOUR, AND A. MUNAGI, Avoiding permutation patterns of type $(2, 1)$ in compositions, *preprint*.
- [27] W. FULTON, *Young Tableaux*, London Mathematical Society Student Texts 35; Cambridge University Press (1997).
- [28] A. M. GOYT, Avoidance of partitions of a three-element set, [arXiv:math.CO/0603481](https://arxiv.org/abs/math/0603481).
- [29] O. GUIBERT, Combinatoire des permutations à motifs exclus en liaison avec mots, cartes planaires et tableaux de Young, Ph.D. thesis, University Bordeaux 1 (1995).
- [30] O. GUIBERT, E. PERGOLA, AND R. PINZANI, Vexillary involutions are enumerated by Motzkin numbers, *Ann. Comb.* **5(2)** (2001), 153–174.
- [31] G. HUTCHINSON, Partitioning algorithms for finite sets, *Comm. ACM* **6** (1963), 613–614.
- [32] A. D. JAGGARD, Prefix exchanging and pattern avoidance by involutions, *Electron. J. Comb.* **9** (2003), #R16.
- [33] A. D. JAGGARD AND J. J. MARINCEL, Generating tree isomorphisms for pattern-avoiding involutions, www.ams.org/amsmtgs/2098_abstracts/1023-05-1618.pdf.

- [34] V. JELÍNEK, Dyck paths and pattern-avoiding matchings, *Eur. J. Combin.* **28(1)** (2007), 202–213.
- [35] V. JELÍNEK, Pattern-avoiding fillings of rectangular shapes, *Proceedings of FPSAC 2007* (2007).
- [36] V. JELÍNEK AND M. KLAZAR, Duality for partitions of bounded number of blocks, *unpublished manuscript*.
- [37] V. JELÍNEK AND T. MANSOUR, On pattern-avoiding partitions, *Electron. J. Comb.* **15** (2008), #R39.
- [38] V. JELÍNEK AND T. MANSOUR, Wilf equivalence on k -ary words, *preprint*.
- [39] T. KAISER AND M. KLAZAR, On growth rates of hereditary permutation classes, *Electron. J. Comb.* **9(2)** (2003), #R10.
- [40] S. KITAEV AND T. MANSOUR, A survey on certain pattern problems, *manuscript*.
- [41] M. KLAZAR, Extremal problems for ordered (hyper)graphs: applications of Davenport–Schinzel sequences, *Eur. J. Combin.* **25** (2004), 125–140.
- [42] M. KLAZAR, On abab-free and abba-free set partitions, *Eur. J. Combin.* **17** (1996), 53–68.
- [43] M. KLAZAR, On the least exponential growth admitting uncountably many closed permutation classes, *Theor. Comput. Sci.* **321(2-3)** (2004), 271–281.
- [44] M. KLAZAR, On trees and noncrossing partitions, *Discrete Appl. Math.* **82** (1998), 263–269.
- [45] M. KLAZAR, Overview of some general results in combinatorial enumeration, [arXiv:0803.4292v1](https://arxiv.org/abs/0803.4292v1).
- [46] D. E. KNUTH, *The Art of Computer Programming, Vol. 1, Fundamental Algorithms*, Addison Wesley, Reading, Massachusetts, 1973.
- [47] D. E. KNUTH, *The Art of Computer Programming, Vol. 3, Sorting and Searching*, Addison Wesley, Reading, Massachusetts, 1973.
- [48] C. KRATTENTHALER, Growth diagrams, and increasing and decreasing chains in fillings of Ferrers shapes, *Adv. Appl. Math.* **37(3)** (2006), 404–431.
- [49] G. KREWERAS, Sur les partitions non croisées d’un cycle, *Discrete Math.* **1(4)** (1972), 333–350.
- [50] A. MARCUS AND G. TARDOS, Excluded permutation matrices and the Stanley–Wilf conjecture, *J. Comb. Theory A* **107(1)** (2004), 153–160.
- [51] S. C. MILNE, Restricted growth functions and incidence relations of the lattice of partitions of an n -set, *Adv. in Math.* **26** (1977), 290–305.
- [52] S. C. MILNE, A q -analog of restricted growth functions, Dobinski’s equality, and Charlier polynomials, *Trans. Amer. Math. Soc.* **245** (1978), 89–118.
- [53] S. C. MILNE, Restricted growth functions, rank row matchings of partition lattices, and q -Stirling numbers, *Adv. in Math.* **43** (1982), 173–196.

- [54] M. POUZET, The profile of relations, *Glob. J. Pure Appl. Math.*, **2(3)** (2006), 237–272.
- [55] A. REGEV, Asymptotic values for degrees associated with strips of Young diagrams, *Adv. Math.* **41** (1981), 115–136.
- [56] A. REGEV, Asymptotics of the number of k -words with an ℓ -descent, *Electron. J. Comb.* **5** (1998), #R15.
- [57] M. RUBEY, Increasing and decreasing sequences in fillings of moon polyominoes, [arXiv:math.CO/0604140](https://arxiv.org/abs/math/0604140).
- [58] B. SAGAN, Pattern avoidance in set partitions, *DIMACS TR: 2006-19*, [arXiv:math.CO/0604292](https://arxiv.org/abs/math/0604292).
- [59] C. D. SAVAGE, A survey of combinatorial Gray codes, *SIAM Rev.* **39** (1997), 605–629.
- [60] C. D. SAVAGE AND H. S. WILF, Pattern avoidance in compositions and multiset permutations, *Adv. Appl. Math.* **36(2)** (2006), 194–201.
- [61] R. SIMION, Noncrossing partitions, *Discrete Math.* **217** (2000), 367–409.
- [62] R. SIMION AND F. SCHMIDT, Restricted permutations, *Eur. J. Combin.* **6** (1985), 383–406.
- [63] Z. STANKOVA, Forbidden subsequences, *Discrete Math.*, **132** (1994), 291–316.
- [64] Z. STANKOVA, Shape-Wilf ordering on permutations of length 3, *Electron. J. Comb.* **14** (2007), #R56.
- [65] Z. STANKOVA AND J. WEST, A new class of Wilf-equivalent permutations, *J. Algebr. Comb.* **15** (2002), 271–290.
- [66] R. P. STANLEY, *Enumerative Combinatorics* (vol. 2), Cambridge Studies in Advanced Mathematics 62, Cambridge University Press (1999).
- [67] R. P. STANLEY, *The Catalan Addendum*, www-math.mit.edu/~rstan/ec/catadd.pdf.
- [68] N. J. A. SLOANE (ED.), The On-Line Encyclopedia of Integer Sequences, <http://www.research.att.com/~njas/sequences/>
- [69] V. VATTER, Permutation classes of every growth rate (a.k.a. Stanley–Wilf limit) above 2.48187, [arXiv:0807.2815v1](https://arxiv.org/abs/0807.2815v1).
- [70] V. VATTER, Small permutation classes, [arXiv:0712.4006v2](https://arxiv.org/abs/0712.4006v2).
- [71] M. L. WACHS, σ -restricted growth functions and p, q -Stirling numbers, *J. Comb. Theory A* **68** (1994), 470–480.
- [72] S. G. WILLIAMSON, Ranking algorithms for lists of partitions, *SIAM J. Comput.* **5** (1976), 602–617.

Appendix A

Notation

The following table summarizes the main notation used in this thesis. The numbers in brackets indicate the page where the corresponding notion is defined.

\mathbb{N}	the set of positive integers $\{1, 2, \dots\}$
\mathbb{Z}	the set of integers
$[n]$	the set $\{1, 2, \dots, n\}$
\mathfrak{S}_n	the set of permutations of order n
$\mathfrak{S}_n(\sigma)$	the set permutations of order n avoiding the pattern σ
\mathcal{I}_n	the set of involutions of order n
$\mathcal{I}_n(\sigma)$	the set of involutions of order n avoiding the pattern σ
A^n	the set of words of length n over the alphabet A
$A^n(w)$	the set of words of length n over the alphabet A avoiding the pattern w
F^T	the transpose of F
\mathcal{P}_n	the set of partitions of $[n]$
$\mathcal{P}_n(\sigma)$	the set of partitions of $[n]$ that avoid σ
$p_n(\sigma)$	the cardinality of $\mathcal{P}_n(\sigma)$
τ^C	the complement of a word τ (45)
\simeq^w	the Wilf equivalence (17)
\preceq	the Wilf order (34)
\simeq^{sh}	the shape-Wilf equivalence (19)
\preceq^{sh}	the shape-Wilf order (34)
\preceq^{sk}	the skew order (35)
\simeq^{I}	the I-Wilf equivalence (40)
$\simeq^{\text{--}}$	the strong equivalence of words (46)
\simeq^{reg}	equivalence with respect to general fillings of rectangles (32)
\simeq^{reg}	equivalence with respect to general fillings of Ferrers shapes (32)
\simeq^{f}	Ferrers equivalence (47)
\triangleleft^{s}	stack equivalence (48)
\simeq	equivalence of partitions (52)
\simeq^{JW}	equivalence with respect to transversals of F^{\perp} -shapes (41)
\simeq^{J}	F^{\perp} -equivalence (54)