

A general theory of Wilf-equivalence for Catalan structures

Michael Albert

Department of Computer Science
University of Otago, New Zealand

`malbert@cs.otago.ac.nz`

Mathilde Bouvel

Institute of Mathematics
University of Zurich, Switzerland

`mathilde.bouvel@math.uzh.ch`

Submitted: Oct 13, 2015; Accepted: Dec 12, 2015; Published: Dec 23, 2015

Mathematics Subject Classifications: 05A05, 05A15, 05A16, 05A19

Abstract

The existence of apparently coincidental equalities (also called Wilf-equivalences) between the enumeration sequences or generating functions of various hereditary classes of combinatorial structures has attracted significant interest. We investigate such coincidences among non-crossing matchings and a variety of other Catalan structures including Dyck paths, 231-avoiding permutations and plane forests. In particular we consider principal subclasses defined by not containing an occurrence of a single given structure. An easily computed equivalence relation among structures is described such that if two structures are equivalent then the associated principal subclasses have the same enumeration sequence. We give an asymptotic estimate of the number of equivalence classes of this relation among structures of size n and show that it is exponentially smaller than the n^{th} Catalan number. In other words these “coincidental” equalities are in fact very common among principal subclasses. Our results also allow us to prove in a unified and bijective manner several known Wilf-equivalences from the literature.

1 Introduction

The Catalan numbers are renowned for their ubiquity in problems of combinatorial enumeration. A few of the many contexts in which they arise are: plane forests (counted by number of nodes), non-crossing matchings or arch systems (counted by number of matched pairs or arches), Dyck paths, and 231-avoiding permutations. Many others are considered in the book [18] based on the “Catalan Addendum” [19]. The contexts mentioned above share the additional property – to be detailed in Section 2 – that each admits a natural substructure relation and that there are bijections between them which preserve that relationship. So, one can further consider those structures of each type which do not contain some designated substructure(s). As part of a previous work (see an extended

abstract [4], or [5]) the present authors considered certain coincidences of enumeration (often called *Wilf-equivalences*) between such classes of Catalan structures avoiding a given substructure (in our case, permutations avoiding 231 and π). Using a non-standard bijection we were able to explain some of those coincidences. However, when we turned to the more general question:

How many distinct enumeration sequences are there for classes of 231-avoiding permutations defined by a single additional restriction?

we were struck by the difference between the computed numbers and any known general equivalences. Specifically it seemed that there were many more such coincidences (and so fewer enumeration sequences) than one might have expected. This phenomenon will be explained in the current paper. We will show in Section 5 that although there are $\text{Cat}_n = \binom{2n}{n}/(n+1) \sim (1/\sqrt{\pi})n^{-3/2}4^n$ distinct classes of permutations avoiding 231 and an additional permutation of size n these classes have asymptotically at most $cn^{-3/2}\gamma^n$ distinct enumeration sequences where $c \approx 1.13$ and $\gamma \approx 2.4975$ (these are approximate values only).

A particularly wide collection of such classes share generating functions derived from the continued fraction representation of $C(t) = \sum \text{Cat}_n t^n$, the generating function of the Catalan numbers. Since $C = 1/(1 - tC)$ it follows that:

$$C = \frac{1}{1 - \frac{t}{1 - \frac{t}{1 - \frac{t}{1 - \dots}}}}$$

This fraction can be truncated after n levels producing a sequence of generating functions:

$$C^{(0)} = 1$$

$$C^{(n)} = \frac{1}{1 - tC^{(n-1)}} \quad \text{for } n \geq 1$$

The functions $C^{(n)}$ enumerate many specific subclasses of the Catalan classes above – for instance the 231-avoiding permutations that also avoid a descending permutation of size n , or the Dyck paths of height at most n . Other examples can be found in [4, 16]. Previously these enumeration coincidences were understood on an analytic (or perhaps more properly arithmetic) level only. We can explain them and many others bijectively – among other things we can show, combining Propositions 13, 14 and 17:

The number of 231-avoiding permutations, π , of size n for which the generating function of the class of permutations avoiding both 231 and π is $C^{(n)}(t)$ is the n^{th} Motzkin number.

The proof of this fact also describes (at least in principle) bijections between any two such classes. Furthermore, we show that for any other 231-avoiding permutation θ of size n the generating function for 231 and θ -avoiding permutations is dominated (term by term and eventually strictly) by $C^{(n)}(t)$.

The main tool in producing these results is a binary relation on Catalan structures defined purely intrinsically by four very simple rules in Section 4. This relation induces an equivalence relation \sim on these Catalan structures whose equivalence classes are the connected components of the binary relation. Remarkably, if $A \sim B$ then the collection of structures not containing A has the same generating function as the collection of structures not containing B so that one generating function may be associated with each equivalence class of \sim . For convenience in the description and proofs we will work mostly in the domain of arch systems, but of course all the results translate to the other domains directly using the natural bijections of Section 2. We have been able to verify that through size 15 (where \sim has only 16,709 equivalence classes on the 9,694,845 Catalan structures) that A and B are \sim -equivalent if and only if the corresponding generating functions are the same. So we have:

Conjecture 1. The equivalence relation \sim coincides with Wilf-equivalence.

In the final section we discuss this conjecture and further open problems.

In the next section we introduce our basic terminology and notation and consider in more detail a particular quartet of Catalan structures: arch systems, Dyck paths, plane forests, and 231-avoiding permutations. This is followed by some preparatory results before we introduce the relation \sim and prove its main property, namely that it refines Wilf-equivalence in Theorem 7. We can represent the collection of all \sim -equivalence classes, which we call *cohorts*, as a slight modification of the family of non-plane forests and this also permits us to determine the number of cohorts in structures of size n both through a functional equation or recurrence and asymptotically. We then consider further relationships between the cohorts (more precisely, between the generating functions of the corresponding avoidance classes) and the properties of the special main cohort mentioned above – which is maximal in terms of the associated generating function. We also discuss cardinalities of cohorts, asking in particular whether or not the main cohort is also the largest in terms of cardinality. Finally we consider some open problems that arise from this work.

2 A quartet of well-behaved Catalan structures

2.1 Arch systems, Dyck paths, plane forests, and 231-avoiding permutations

Among the most well-known Catalan structures are certainly the Dyck paths. A *Dyck path* of semi-length n is a path in the positive quarter-plane taking steps $u = (1, 1)$ and $d = (1, -1)$, starting at $(0, 0)$ and ending at $(2n, 0)$. Steps u and d of a Dyck path may be paired by associating to each u step the first d step on its right at the same ordinate. These pairs (u, d) may also be seen as pairs of opening and closing parentheses and under

this correspondence Dyck paths correspond to parentheses words where parentheses are properly matched. A subpath of a Dyck path is defined by the deletion of some pairs of steps (u, d) (or equivalently of matched parentheses). The deletion here is intended as a *contraction* of the segment of each deleted step into a point so that deleting k pairs of steps in a Dyck path of semi-length n provides a Dyck path of semi-length $n - k$. In contrast to [7] we require more of a subpath than that it simply occur as a subword over the alphabet $\{u, d\}$. We also require that the “matched pairs” of u ’s and d ’s should be the same in both paths.

Another natural way of representing proper parentheses words is as non-crossing matchings or arch systems. These form a second family of Catalan structures and will be essential in the presentation of our results. An *arch system* of size n is a set of n arches connecting $2n$ points arranged along a baseline such that all arches are above the baseline and no pair of arches cross. The left end of each arch encodes an opening parenthesis and its right end the corresponding closing parenthesis. A subsystem of an arch system can be obtained simply by deleting some of the original system’s arches.

We can concatenate arch systems A and B in the obvious way – just draw the arch system B strictly to the right of A on the same baseline. The resulting arch system will be denoted AB .

Definition 2. An *atom* is a non empty arch system that cannot be written as the concatenation of two non empty arch systems, i.e., one that has a single outermost arch. Atoms will generally be denoted by lower case letters. The *contents* of an atom a are the unique arch system, A , such that a is obtained by adding a single arch outside all of A and we write $a = \overline{A}$.

Since every non empty arch system is a unique concatenation of atoms we see immediately that the generating function, $A(t)$, for arch systems according to the number of arches satisfies:

$$\begin{aligned} A(t) &= 1 + t A(t) + (t A(t))^2 + (t A(t))^3 + \cdots \\ &= \frac{1}{1 - t A(t)} \end{aligned}$$

proving that – and this should be no surprise – arch systems are enumerated by the Catalan numbers.

There is a bijection between arch systems with n arches and non-empty plane forests with n nodes obtained simply by mapping each arch to a node in such a way that if one arch lies within another then its node is a descendant of the other, and if it lies to the left of another then its node does so too. Equivalently, describing this recursively: take an arch system A , write it as a concatenation of atoms $A = a_1 a_2 \cdots a_m$ and associate to it a forest of m trees whose roots, r_i , correspond to the outermost arches of the a_i (and are arranged from left to right for i from 1 through m) and such that the tree rooted at r_i is (up to the addition of the root r_i) the forest of the contents of a_i . This bijection also preserves the “substructure” relationship provided that in the case of forests we maintain ancestry in substructures (i.e., if a child, x , of a node, y , is deleted then all the children of x




Figure 1: The Dyck path, arch system and plane forest corresponding to the 231-avoiding permutation 41327658. Reading u steps (or arch beginnings) as push operations and d steps (or arch ends) as pop operations on a stack the Dyck path (or arch system) successfully sorts this permutation i.e., the output sequence would be 12345678.

remaining become children of y preserving their left to right order both among themselves and with respect to their new siblings).

Finally, we can consider 231-avoiding permutations of $\{1, 2, \dots, n\}$. These are those permutations π which, when written in one line notation, contain no subsequence bca with $a < b < c$. Here the substructure relationship (known as the *pattern* relationship among permutations) involves deleting some symbols and then relabeling the remaining ones to form a permutation of $\{1, 2, \dots, m\}$ for some $m < n$ while maintaining their relative order (e.g., if we delete 2 from 31254 we obtain 2143). It is perhaps not immediately clear that these are also in bijection with Dyck paths, arch systems or plane forests. However, these permutations are precisely those that can be sorted by a single pass through a stack [14] and we can form a Dyck path by adding a step u whenever pushing an element on to the stack and a step d whenever popping one from the stack. Since the sequence of push and pop operations to sort a permutation is easily seen to be unique and every sequence of operations sorts *some* permutation this is clearly a bijection. Moreover, it respects the substructure relationship since, when deleting an element, we just delete the pair of matched steps, or equivalently the arch in the corresponding arch system which corresponds to push and pop operations that affect that element. This bijection can also be realized intrinsically. The n arches are labeled with the integers from 1 through n according to the following rules: if two arches are nested then the outer arch has a greater label than the inner one and if two arches are not nested the arch to the left has a lesser label than the arch to the right. The permutation is then read by reading the labels of the arches in order of their leftmost endpoints. This means that the left to right maxima of the permutation (i.e., the elements that have no greater element to their left) correspond to outermost arches and within them an arch system is constructed using the same principle recursively on the following lesser elements. An example of these correspondences is given in Figure 1.

With a slightly different view point the main idea of the above discussion can be rephrased by stating that the four following posets are isomorphic: those of Dyck paths, arch systems, plane forests and 231-avoiding permutations (where, in all four cases, the order is that of (induced) substructure). In addition the bijections described above are order-preserving.

2.2 A note about other Catalan structures

Of course there are other classical combinatorial structures that are enumerated by the Catalan numbers: plane binary trees, 321-avoiding permutations, or non-crossing partitions, just to name a few. And of course there are also classical (although less canonical) bijections between them and Dyck paths, plane forests or 231-avoiding permutations. However, it is deliberate that we restrict our attention in this work to the quartet of Catalan objects presented above. The reason is that the substructure relation on Dyck paths, plane forests or 231-avoiding permutations does not translate naturally to the contexts of binary trees, 321-avoiding permutations, nor non-crossing partitions. A consequence of this is that the results we prove below for this particular quartet of Catalan structures do not apply in other Catalan contexts.

In the case of binary trees, it is not even clear how to define a natural notion of subtree induced by a subset of the vertices while preserving the binary condition. This fact somehow explains why the link between 231-avoiding permutations and binary trees with respect to pattern avoidance is not as natural as one might hope for – see [10, Section 6].

The situation is different for 321-avoiding permutations, or non-crossing partitions. Indeed both have a natural notion of induced substructure but it does not mimic the substructure order on Catalan structures in our quartet. Specifically, the corresponding posets are not isomorphic to the Catalan poset that we are considering. This is shown for instance on Figure 2 in the case of 321- and 231-avoiding permutations.




Figure 2: The first few levels of the posets for 231-avoiding permutations and 321-avoiding permutations the order being that of pattern, i.e., (induced) substructure.

2.3 Excluding substructures in our quartet

Restricting our attention to our quartet of preferred Catalan structures we are interested in considering the problem:

Given a single structure, A , what is the generating function of the collection of structures that do not have A as a substructure?

Going back to some examples discussed in the introduction note that Dyck paths of height at most n corresponds to Dyck paths that do not have $u^n d^n$ as a subpath. Under the correspondences we have described these correspond to: arch systems that do not have $N_n = \overline{\dots \cap \dots}$, the nested arch system with n arches, as a subsystem, plane forests of depth at most n , and 231-avoiding permutations with no $n(n-1) \dots 21$ pattern.

Structures that do not have A as a substructure are said to *avoid* A and we will denote the set (or *class*) of them by $\text{Av}(A)$. If a structure does not avoid A it is said to *involve* or *contain* A . In this paper we will only be considering the avoidance of a single structure – but of course in general we could consider any collection of structures closed downwards under the substructure relation.

Definition 3. The classes $\text{Av}(A)$ and $\text{Av}(B)$ are *Wilf-equivalent*, written $\text{Av}(A) \simeq \text{Av}(B)$, if there is a bijection between them that preserves the size of each structure. Equivalently, the generating functions F_A of $\text{Av}(A)$ and F_B of $\text{Av}(B)$ are equal.

When $\text{Av}(A)$ and $\text{Av}(B)$ are Wilf-equivalent it need not, and generally will not, be the case that the classes $\text{Av}(A)$ and $\text{Av}(B)$ are isomorphic as posets. Sometimes *par abus de langage* we may say that A and B are Wilf-equivalent when we mean that $\text{Av}(A)$ and $\text{Av}(B)$ are. If A and B are of different sizes, then they cannot possibly be Wilf-equivalent so effectively Wilf-equivalence is an equivalence relation on structures of size n for each n . As such, the n^{th} Catalan number is an upper bound for the number of its equivalence classes there, but we shall see that this is far from the actual number of equivalence classes.

3 Arch systems containing and avoiding subsystems

If an arch system X contains some arch system P then there is a *leftmost* occurrence of P in X (which we often denote P_L) by which we mean the occurrence of P whose rightmost point (i.e. the point of X that corresponds to the final point of P in this occurrence) is as far left as possible. If there are two such occurrences with the same rightmost point we designate as P_L the one whose second rightmost point is as far left as possible etc. There is also a corresponding notion of *rightmost* occurrence.

One advantage of working with arch systems is that it is clear that when searching for a substructure of X equal to some given arch system we may proceed in a greedy fashion. That is:

Observation 4. Suppose that P , Q and X are arch systems and that PQ is a substructure of X . Then, in witnessing this we may use the leftmost occurrence, P_L , of P in X .

We will use this observation (and some obvious generalizations) repeatedly without further comment. Note however that we do not suggest that X must factor into a part containing P and a part containing Q . For example the system $\overline{P \cap} \cap$ has $P \cap \cap$ as a substructure but no such factorization.

For any arch system, A , let F_A denote the generating function of $\text{Av}(A)$. It is a result of [15] (expressed in somewhat different terms of course) that F_A is necessarily a rational

function. In fact, given a factorization of A into atoms we can write down a system of equations that allow for the recursive computation of F_A (again, this is already done in [15] and, in somewhat more general terms, in [2]). The following proposition simply translates that result into the current context.

Proposition 5. *Let A be an arch system, with $A = a_1 a_2 \cdots a_m$ its factorization into atoms, and $a_1 = \overline{A_1}$. Then the generating function of $\text{Av}(A)$ is*

$$F_A = 1 + tF_{A_1}F_A + t(F_{a_1} - F_{A_1})F_{a_2 \dots a_m} + t \sum_{k=2}^m (F_{a_1 \dots a_k} - F_{a_1 \dots a_{k-1}})F_{a_k \dots a_m}.$$

In particular, F_A is rational.

Fundamentally the first part of the proposition is proved simply by partitioning A -avoiding arch systems according to “how much of A ” can be found within the first arch and the conclusion of the second part follows by an easy inductive argument.

4 A refinement of Wilf-equivalence

In this section we introduce an equivalence relation, \sim , on the collection of arch systems. We will then establish that this relation refines Wilf-equivalence, i.e., that $A \sim B$ implies $\text{Av}(A) \simeq \text{Av}(B)$. So, without further ado:

Definition 6. The binary relation, \sim , on arch systems is the finest equivalence relation that satisfies:

$$A \sim B \implies \overline{A} \sim \overline{B} \tag{1}$$

$$a \sim b \implies PaQ \sim PbQ \tag{2}$$

$$PabQ \sim PbaQ \tag{3}$$

$$a\overline{bc} \sim \overline{ab}c. \tag{4}$$

where A , B , P and Q denote arbitrary arch systems; and a , b and c denote arbitrary atoms or empty arch systems. The equivalence classes of \sim will be called *cohorts*.

Note that if $A \sim B$ then A and B have the same number of arches. Note also that $A \sim B \Leftrightarrow \overline{A} \sim \overline{B}$ since (non trivial) equivalences between atoms may only be produced by rule (1).

The main result is:

Theorem 7. *If A and B are arch systems and $A \sim B$ then $\text{Av}(A) \simeq \text{Av}(B)$.*

This will be proved by a series of arguments in the following subsections.

Interestingly, another equivalence relation (say, \equiv) on Catalan structures has been defined in a similar fashion by Rudolph [17]. She proves in this paper that two \equiv -equivalent 132-avoiding permutations π and τ are *equipopular*, that is: for any n , the total

number of occurrences of π and τ in 132-avoiding permutations of size n are equal. In other words, \equiv refines equipopularity, and the analogy with \sim refining Wilf-equivalence is clear. What is further interesting in the case of \equiv , is that it *coincides* with equipopularity as shown in [9]. As a consequence the number of equivalence classes for equipopularity among permutations of size n is given by the number of partitions of n .

We give bijective proofs for all parts of Theorem 7 including some for cases where it is also possible (and sometimes easier) to obtain the corresponding result by manipulation of generating functions. We call proofs of the latter type *analytic* since they form a large part of the field of analytic combinatorics, though we will not in general be making use of any analytic properties of the corresponding generating functions. One reason to provide bijective proofs is that they can frequently be refined to allow for term by term comparisons between the generating functions for inequivalent cohorts while this is not so easily accomplished when only analytic proofs are available. A second reason is that these bijective proofs are needed for proving our claim of the introduction: that we are able (at least in principle) to provide bijections between any two classes of permutations $\text{Av}(231, \pi)$ and $\text{Av}(231, \tau)$ for π and τ of size n whose generating function is $C^{(n)}$.

To prove Theorem 7 it is sufficient to show that its conclusion holds for each of the four cases arising in Definition 6. The proof is therefore subdivided into such cases. For compactness of notation we have found it convenient to denote functional application in exponential form, i.e., the image of an arch system X under a map τ will be denoted X^τ .

4.1 Proof of Theorem 7

Proof of case (1): $\text{Av}(A) \simeq \text{Av}(B) \implies \text{Av}(\overline{A}) \simeq \text{Av}(\overline{B})$. Let A and B be given and suppose that $\text{Av}(A) \simeq \text{Av}(B)$. We may further assume that A and B are not empty or the result holds trivially. Take σ to be any size-preserving bijection between $\text{Av}(A)$ and $\text{Av}(B)$. Define a map τ on atoms $x = \overline{X}$ belonging to $\text{Av}(\overline{A})$ by $x^\tau = \overline{X^\sigma}$. This is possible since $x \in \text{Av}(\overline{A})$ if and only if $X \in \text{Av}(A)$. Now extend τ to concatenations of atoms in the obvious way, i.e., $(x_1 x_2 \dots x_m)^\tau = x_1^\tau x_2^\tau \dots x_m^\tau$. Since $\text{Av}(\overline{A})$ consists exactly of arch systems which are concatenations of atoms whose contents belong to $\text{Av}(A)$ (and correspondingly $\text{Av}(\overline{B})$ consists exactly of arch systems which are concatenations of atoms whose contents belong to $\text{Av}(B)$) $\tau : \text{Av}(\overline{A}) \rightarrow \text{Av}(\overline{B})$ is a size preserving bijection. \square

Proof of case (2): $\text{Av}(a) \simeq \text{Av}(b) \implies \text{Av}(PaQ) \simeq \text{Av}(PbQ)$. Let P and Q be arbitrary arch systems and a and b be atoms such that $a \sim b$. Assume that a and b are not empty (or the result holds trivially) and let $\sigma : \text{Av}(a) \rightarrow \text{Av}(b)$ be a size preserving bijection. We will define a size preserving bijection $\tau : \text{Av}(PaQ) \rightarrow \text{Av}(PbQ)$.

Suppose that $X \in \text{Av}(PaQ)$. If $X \in \text{Av}(PQ)$ we define $X^\tau = X$. Otherwise take the leftmost copy, P_L , of P in X and the rightmost copy, Q_R , of Q . The arches that begin before the end of P_L but end after it, and those that end after the beginning of Q_R but begin before it divide the segment between the end of P_L and the beginning of Q_R into intervals. This is illustrated in Figure 3. Since a is an atom, any occurrence of a between the end of P_L and the beginning of Q_R would have to be entirely contained in one of the




Figure 3: The situation arising in the proof of case (2). In an arch system X involving PQ but avoiding PaQ the leftmost copy of P , denoted P_L , and the rightmost copy of Q , denoted Q_R are designated. Arches with one endpoint inside and one endpoint outside the interval between P_L and Q_R create a sequence of subintervals (I_1 through I_4 here) that must avoid a . To produce a PbQ avoiding arch system, a bijection mapping a -avoiding systems to b -avoiding systems is applied to the I_i and the remainder of the system is left unchanged.

intervals. So, each of these intervals contains an arch system that avoids a and conversely, if we are given an arch system with this property, it avoids PaQ . Define X^τ by applying σ to each of the intervals while retaining the structure of X up to the end of P_L and from the beginning of Q_R (including the arches that define the intervals). It is immediate to check that this defines a bijection from $\text{Av}(PaQ)$ to $\text{Av}(PbQ)$. \square

Proof of case (3): $\text{Av}(PabQ) \simeq \text{Av}(PbaQ)$. The claim is trivial when a or b is empty. For the non-trivial case let a and b be non empty arbitrary atoms and P and Q arbitrary arch systems. We wish to construct a bijection $\tau : \text{Av}(PabQ) \rightarrow \text{Av}(PbaQ)$. It will be helpful in what follows for the reader to refer to Figure 4. As in the previous case consider an arch system X . If X avoids PaQ then define $X^\tau = X$. Otherwise take P_L to be the leftmost P , a_L the leftmost atom involving a following P_L and Q_R the rightmost Q in X . Furthermore, denote by C the contents of a_L . As in the previous proof the interval between P_L and Q_R is subdivided by those arches that have only one endpoint in this interval, say there are i (resp. j) such arches with only their right (resp. left) endpoint between P_L and Q_R . But now also one of those intervals (the one containing a_L) is further subdivided before and after a_L by a_L itself and any arches nested over a_L . Denote by k the number of such arches (including the outermost arch of a_L). All the designated subintervals to the left of a_L must avoid a (since a_L was leftmost) while those to the right of it must avoid b (since X avoids $PabQ$). To define X^τ simply reverse the order of these subintervals (keeping the arch systems within them fixed i.e., the contents of a subinterval are not changed only its position between P_L and Q_R). The structure of the arch system outside these intervals is unchanged, that is: the arch system before P_L and after Q_R is not modified and there are still k arches on top of C and i (resp. j) arches with only their right (resp. left) endpoint between P_L and Q_R . In the resulting arch system X^τ , P_L and Q_R are still the leftmost copies of P and the rightmost copies of Q respectively (since nothing before the end of P_L or after the start of Q_R has been changed). Between these,




Figure 4: A typical situation arising in the proof of case (3). In the top diagram the original $PabQ$ avoiding arch system is shown. Each interval A_i must avoid a and each interval B_j must avoid b . In the bottom diagram its image is shown – the atom a_L and the nest of arches around it are moved to the right to allow copies of the B_j to be placed on the left and copies of A_i on the right as seen in the middle two diagrams.

the atom a_L has become the rightmost atom involving a . Since all of the intervals before it but following P_L avoid b , X^τ avoids $PbaQ$. Moreover it is clear that we can reverse this construction, so $\tau : \text{Av}(PabQ) \rightarrow \text{Av}(PbaQ)$ is a size preserving bijection as claimed. \square

In the proof of case (3) we have chosen to reverse $A_1 \dots A_{i+k}CB_1 \dots B_{j+k}$ in X^τ . But many variants of τ could have been defined by choosing any other permutation of the A_ℓ , B_m and C that respects that all the B_m are to the left of C and all the A_ℓ to its right.

Turning now to case (4) we will give an analytic proof below but here give a bijective proof – albeit one that it is not nearly as concrete as the previous cases. This proof, with minor variations, is due to Jonathan Bloom (the present authors had previously only been able to prove the result bijectively for the case where one of a , b or c was empty.)

We first need a preparatory lemma which is a very simple combinatorial fact in general

(basically: from a bijection from X to Y and a second one from a subset of X to a subset of Y we can construct one that agrees with the second one on the designated subset.)

Lemma 8. *Let a, b and c be atoms (or empty). There is a bijection $\tau : \text{Av}(abc) \rightarrow \text{Av}(bac)$ which has the property that for all $A \in \text{Av}(abc) \setminus \text{Av}(ab)$, $A^\tau \in \text{Av}(bac) \setminus \text{Av}(ba)$. That is, if X avoids abc but involves ab , then X^τ avoids bac but involves ba .*

Proof. Let $\sigma : \text{Av}(abc) \rightarrow \text{Av}(bac)$ be any bijection (such as the one constructed in the proof of the preceding case) and let $\nu : \text{Av}(ba) \rightarrow \text{Av}(ab)$ be any bijection. Note that $\text{Av}(ab) \subseteq \text{Av}(abc)$ and $\text{Av}(ba) \subseteq \text{Av}(bac)$. For $X \in \text{Av}(ab)$ define $X^\tau = X^{\nu^{-1}}$. For $X \in \text{Av}(abc) \setminus \text{Av}(ab)$ define a sequence $Y_1 = X^\sigma$. Now if Y_i is defined and belongs to $\text{Av}(ba)$ define $Y_{i+1} = Y_i^{\nu\sigma}$. Since σ and ν were bijections, it must be the case that for some i , $Y_i \in \text{Av}(bac) \setminus \text{Av}(ba)$ and we define $X^\tau = Y_i$. \square

Proof of case (4): $\text{Av}(a\overline{b}c) \simeq \text{Av}(b\overline{a}c)$. Let $X \in \text{Av}(a\overline{b}c)$ be given. We wish to define its image X^τ , and will assume that Y^τ has already been defined for all Y of smaller size with the property that if Y contains the pattern ab , then Y^τ contains the pattern ba . We consider two cases depending on whether X is in $\text{Av}(abc)$ or in $\text{Av}(a\overline{b}c) \setminus \text{Av}(abc)$.

If $X \in \text{Av}(abc)$, then we can bijectively map X to $X^\tau \in \text{Av}(bac)$ using the preceding lemma.

Now assume $X \in \text{Av}(a\overline{b}c) \setminus \text{Av}(abc)$ and let $\overline{C_R}$ be the rightmost occurrence of c in X (in particular, the arch on the outside of $\overline{C_R}$ matches the arch on the outside of the atom c). Consider the intervals defined by the nest of arches (if any) over $\overline{C_R}$. The only condition on the intervals to the right of $\overline{C_R}$ is that they avoid c . We label these C_i . Now assume there are $n + m + 1$ intervals to the left of $\overline{C_R}$. These intervals come in three flavors. The n leftmost intervals A_i avoid a , the $(n + 1)$ st interval Z contains a , and the remaining m intervals B_i must avoid b . So we have

$$X = A_1 \left[\dots \left[A_m \left[Z \left[B_1 \left[\dots \left[B_n \overline{C_R} C_{m+n+1} \right] \dots \right] C_{m+2} \right] C_{m+1} \right] C_m \right] \dots \right] C_1.$$

As X is assumed to contain the pattern abc , it is now clear from the above decomposition that Z must actually contain the pattern ab . By our induction hypothesis we know that Z^τ contains the pattern ba . Lastly, we define

$$X^\tau = B_1 \left[\dots \left[B_n \left[Z^\tau \left[A_1 \left[\dots \left[A_m \overline{C_R} C_{m+n+1} \right] \dots \right] C_{n+2} \right] C_{n+1} \right] C_n \right] \dots \right] C_1.$$

Observe that X^τ contains the pattern bac , so that $X^\tau \in \text{Av}(b\overline{a}c) \setminus \text{Av}(bac)$ as needed. In particular, as X contains the pattern ab we also have that X^τ contains ba . That τ is a bijection in this case follows by induction and the fact that $\overline{C_R}$ is still the rightmost occurrence of c in X^τ . \square

As in the proof of case (3) we can again define many variants of the bijection $\tau : \text{Av}(a\overline{b}c) \rightarrow \text{Av}(b\overline{a}c)$ by replacing in X^τ the sequence $y_1, y_2 \dots$ by any permutation of the y_i , where $y = A, B$ or C .

4.2 An analytic proof

The bijection provided in case (4) above is a little unsatisfactory since it is in part implicit, based on bijections from $\text{Av}(abc)$ to $\text{Av}(bac)$ and $\text{Av}(ab)$ to $\text{Av}(ba)$. In the case where one of a , b or c is empty, it can be made explicit. Here, however, we consider an analytic proof of the same result.

Proof of case (4). Let $a = \overline{[A]}$, $b = \overline{[B]}$ and $c = \overline{[C]}$. For an arch system X let F_X be the generating function of $\text{Av}(X)$. Using the general technique described in Proposition 5 we can compute the generating function $F_{a\overline{[bc]}}$ in terms of F_A , F_B and F_C .

$$\begin{aligned} F_{a\overline{[bc]}} &= 1 + tF_A F_{a\overline{[bc]}} + t(F_{a\overline{[bc]}} - F_A)F_{\overline{[bc]}} \\ F_{\overline{[bc]}} &= 1 + tF_{bc}F_{\overline{[bc]}} \\ F_{bc} &= 1 + tF_B F_{bc} + t(F_{bc} - F_B)F_c \\ F_c &= 1 + tF_C F_c \end{aligned}$$

Solving the system¹ for $F_{a\overline{[bc]}}$ in terms of F_A , F_B and F_C gives a terrible mess which is nevertheless symmetric in F_A , F_B and F_C . In fact the solution is tidier if written in terms of F_a , F_b and F_c (recall that $F_a = 1/(1 - tF_A)$, i.e., $F_A = (F_a - 1)/(tF_a)$ etc.):

$$F_{a\overline{[bc]}} = \frac{1 - t(F_a F_b + F_b F_c + F_c F_a - F_a F_b F_c)}{1 - t(F_a + F_b + F_c - F_a F_b F_c)}$$

Accordingly, $F_{a\overline{[bc]}}$ is symmetric in F_a , F_b and F_c . This proves that $\text{Av}(a\overline{[bc]}) \simeq \text{Av}(c\overline{[ab]})$. Now use case (3) to reach the desired conclusion. \square

We have seen in the above proof that, for any atom $a = \overline{[A]}$, F_a completely determines F_A and conversely via the relations $F_a = 1/(1 - tF_A)$ and $F_A = (F_a - 1)/(tF_a)$. This simple fact also provides an analytic proof that:

Observation 9. For any atoms $a = \overline{[A]}$ and $b = \overline{[B]}$, $\text{Av}(A) \simeq \text{Av}(B)$ if and only if $\text{Av}(a) \simeq \text{Av}(b)$.

5 The combinatorial class of cohorts

From Theorem 7 it follows that the number of different generating functions of classes of arch systems avoiding an arch system with n arches (or equivalently the number of Wilf-equivalence classes of permutation classes $\text{Av}(231, \pi)$ for π of size n avoiding 231) is at most the number of cohorts (i.e., equivalence classes of \sim) for n element structures. In Conjecture 1 we suggest that these numbers may actually be equal explaining our interest in the enumeration of cohorts. In any case the number of cohorts certainly provides an upper bound for the number of such Wilf-equivalence classes. Towards the goal of enumerating cohorts we first associate with each cohort a *single* structure and

¹Or rather, having *Mathematica* solve it for you.

then enumerate such structures. These structures that represent cohorts may be seen as choosing one representative in the set of all structures (e.g., all arch systems) that form a cohort. Alternatively – and it is rather this point of view we choose – we can think of the structure representing a cohort as an abstract structure from which all structures in the cohort may be recovered.

5.1 The structure of a cohort

It is easiest to describe the single (abstract) structure associated with a cohort in the context of plane forests. In fact our first result shows that we could just as easily drop the qualifier ‘plane’:

Proposition 10. *If two plane forests A and B are isomorphic as non-plane forests, then $A \sim B$.*

Proof. This follows directly by induction from rules (1), (2) and (3). Specifically, suppose that plane forests A and B which are isomorphic as non-plane forests are given and that the result holds for all plane forests of lesser size. If A and B are trees (corresponding to atoms in the context of arch systems), then the result applies to the forests obtained by deleting their roots (i.e., the contents of these atoms) and hence by rule (1) to A and B . Otherwise, each of A and B is the concatenation of the same number of trees (i.e. atoms), say m . First, using rule (3) we can find $A' \sim A$ so that $A' = a_1 a_2 \dots a_m$, $B = b_1 b_2 \dots b_m$, and each tree a_i is isomorphic to b_i . Then using rule (2) we are done. \square

We note that this proposition already establishes that there are no more cohorts for n element structures than there are rooted non-plane forest with n nodes, or equivalently rooted non-plane *trees* with $n + 1$ nodes. As the asymptotic enumeration of these (see for example [11, Proposition VII.5 and note VII.21]) has exponential growth rate approximately 2.956 we already see exponentially fewer Wilf-equivalence classes than there are structures of size n . However, the final rule provides a further reduction.

Let us focus our attention on \sim -equivalences *between atoms (or trees) only* that may be derived from rule (4). In this context an equivalent form of this rule is $\overline{a(bc)} \sim \overline{ab}c$. So in terms of trees rule (4) allows us to rotate subtrees at binary branches. Furthermore, it also allows unary nodes to be lifted through binary ones (from the case when c is empty) via $\overline{a(b)} \sim \overline{ab}$. Finally, in the case where b and c are empty, rule (4) rewrites as $\overline{a\cap} \sim \overline{a}$, allowing to transform a leaf hanging below a binary node x into a unary node between x and its other child. These operations on trees are shown in Figure 5.

Consider any subtree of a plane forest that has a binary root. In this tree replace any subtree whose root has three or more children by a symbol representing that atom (and temporarily call such atoms, *large*). As a result we obtain a tree, T , all of whose internal nodes have one or two children and where the leaves are either large atoms or bare nodes. As shown in Figure 5(ii) and (iii) we can lift the unary nodes and bare nodes through the binary ones to obtain a \sim -equivalent tree T' with a chain of unary nodes running from the root connected to a full binary tree all of whose leaves are labeled with




Figure 5: \sim -equivalences on trees that are derived from rule (4).

large atoms. Finally, we can rotate the large atoms (see Figure 5(i)), permute them (from $PabQ \sim PbaQ$), and replace them by equivalent large atoms (from $a \sim b \Rightarrow PaQ \sim PbQ$). So we see that two such full binary trees (with leaves that are large atoms) are \sim -equivalent if and only if they have the same number of nodes (and hence leaves) and there is a bijection between their sets of leaves such that items in correspondence in these sets are \sim -equivalent large atoms. More properly, note that these “sets” of leaves are actually multisets since repetitions are allowed.

For ease of explanation, in the rest of this section we will focus on *atomic cohorts*, i.e., cohorts that contain at least one atom (or tree). Note that this is not an actual restriction: atomic cohorts for $(n + 1)$ element structures are in bijective correspondence with cohorts for n element structures, since $A \sim B \Leftrightarrow \overline{[A]} \sim \overline{[B]}$.

The above discussion leads to a recursive description of (representatives for) atomic cohorts. Consider the recursive specification of a variety, \mathcal{A} , of non-plane tree-like structures:

$$\begin{aligned} \mathcal{A} &= \bullet + \bullet \times \mathcal{A} + \mathcal{B} + \sum_{k \geq 2} \Delta_{k-1} \times \text{MSET}_k(\mathcal{B}) \\ \mathcal{B} &= \bullet \times \text{MSET}_{\geq 3}(\mathcal{A}), \end{aligned} \tag{5}$$

where \bullet refers to a class with a single object of size 1, \times denotes ordered pairs, Δ_m denotes a class with a single object of size m , $+$ denotes disjoint union, and MSET denotes the multiset construction, with the subscript denoting the number of elements in the multiset. Equivalently, as non-plane trees:

$$\mathcal{A} = \bullet + \begin{array}{c} \bullet \\ | \\ \mathcal{A} \end{array} + \mathcal{B} + \sum_{k \geq 2} \underbrace{\begin{array}{c} \Delta_{k-1} \\ \swarrow \quad \searrow \\ \mathcal{B} \cdots \mathcal{B} \end{array}}_{k \text{ children}} \quad \text{and} \quad \mathcal{B} = \underbrace{\begin{array}{c} \bullet \\ \swarrow \quad \searrow \\ \mathcal{A} \cdots \mathcal{A} \end{array}}_{\geq 3 \text{ children}}$$

Proposition 11. *There is a size-preserving bijection between atomic cohorts and \mathcal{A} .*

Proof. This is basically simply a direct translation of the preceding discussion, where we have unraveled all possible equivalences following from rules (1) to (4). The class \mathcal{B} represents “large atoms”. Then the elements of \mathcal{A} are described in order as: a single node, a root with one child, a large atom, or an atom corresponding to a full binary tree with k leaves labeled by large atoms. \square

We shall use this description to refine the asymptotic enumeration of the number of cohorts.

5.2 The number of cohorts

The first 15 values of the number of cohorts of arch systems of size $n \geq 1$ are given by:

1, 1, 2, 4, 8, 16, 32, 67, 142, 307, 669, 1 478, 3 290, 7 390, 16 709, ...

Furthermore, for each cohort of size up to 15, we can produce a representative arch system X for that cohort and check that the generating functions of the classes $\text{Av}(X)$ are all distinct. With Theorem 7 this ensures that the above also shows the first few terms of the sequence enumerating Wilf-equivalence classes of classes $\text{Av}(A)$ for A of size n . Notice that more terms of the enumeration sequence of cohorts may be obtained from Equation (6) below – the next few terms are 38 027, 86 993, 200 018, 461 847, 1 070 675. From Theorem 7 these are upper bounds on the number of Wilf-equivalence classes of $\text{Av}(A)$ but we cannot ensure that they are equal (although we suspect they are). In the following, we therefore study the asymptotic behavior of the number of cohorts of arch systems of size n .

As already noted, the number of cohorts of arch systems of size n equals the number of *atomic* cohorts of arch systems of size $n + 1$. Here we can make profitable use of (5) to provide a functional equation for the generating function $A(t) = \sum a_n t^n$ counting atomic cohorts which is susceptible to asymptotic analysis using the techniques of Section VII.5 of [11], or with minor variations of the twenty steps of [12]. Specifically we obtain:

$$\begin{aligned} A &= t + tA + \frac{1}{t}M_{\geq 2}(tB) + B \\ B &= tM_{\geq 3}(A) \end{aligned} \tag{6}$$

where

$$\begin{aligned} M(Z) &= \exp\left(\frac{Z(t)}{1} + \frac{Z(t^2)}{2} + \frac{Z(t^3)}{3} + \frac{Z(t^4)}{4} + \dots\right) \\ M_{\geq 2}(Z) &= M(Z) - 1 - Z(t) \\ M_{\geq 3}(Z) &= M(Z) - 1 - Z(t) - \frac{1}{2}(Z(t)^2 + Z(t^2)) \end{aligned}$$

are operators representing the generating functions that enumerate multisets of objects (and respectively such multisets of size at least 2 or 3) counted by the generating function Z .

Clearly the power series A dominates $t + tM_{\geq 3}(A)$ term by term and so a_n is at least the number of non-plane trees with n nodes in which each internal node has at least 3 children. This trivial estimate suffices to show that the radius of convergence, ρ_A , of A is less than 1 (and hence so is that of B). Now observe that in general

$$M(Z) = \exp(Z) \exp(W) \tag{7}$$

where

$$W = \frac{Z(t^2)}{2} + \frac{Z(t^3)}{3} + \frac{Z(t^4)}{4} + \dots$$

If the radius of convergence of Z is $r < 1$, then the radius of convergence of W is easily seen to be at least $\sqrt{r} > r$. This suggests that when analyzing the radius of convergence of generating functions defined by functional equations involving the M operator we treat these as implicit definitions of the desired function in terms of “known” analytic functions which, while related to the function we are analyzing, are analytic in a disc around the origin strictly containing the radius of convergence of the function we seek. Effectively these are the first five steps of [12]. So to proceed we view (6) as an implicit definition of A in terms of these “known” functions after having eliminated B entirely and noting also that the terms corresponding to $Z(t^2)$ in any occurrences of $M_{\geq 3}$ should also be treated as “known”. Thus we aim to find the radius of convergence of the solution to $F(t, y) = 0$ where:

$$F(t, y) = -y + t + ty + tM_{\geq 3}(y) + M_{\geq 2}(t^2M_{\geq 3}(y))/t$$

In this expression we replace the subscripted M operators by their definitions above and then on the remaining occurrences of M use the form given by (7) to replace the definition of F by one involving y , t and some functions of t known to be analytic on the domain of interest. Continuing with the steps of [12] as we know already that the solution y is a generating function we can find its radius of convergence ρ_A by determining the smallest positive root of the equation $F_y(t, y) = 0$ (where F_y is the derivative of F with respect to y).

Of course in finding this root we first take the derivative formally and then replace y and all the related “known” functions by polynomial approximations of some degree, denoted n , obtained by using equation (6) as a recurrence for generating terms of A . The results of these approximations for various values of n are as follows:

n	ρ_A	$1/\rho_A$
50	0.4069	2.4575
100	0.4083	2.4763
200	0.4022	2.4863
400	0.4014	2.4916
800	0.4009	2.4943

These values agree well with the numerical estimates obtained by simply looking at computed coefficients of A and fitting an asymptotic expression of the form $a_n \sim \alpha n^{-3/2} \gamma^n$. Note however that the apparent accuracy is significantly less than that given in examples VII.21 and VII.22 of [11]. We suspect that this arises due to the iterated application of M and the correction terms that are part of the definitions of $M_{\geq 2}$ and $M_{\geq 3}$. Another possible reason is that we also truncate the “known” parts at degree n . Approximate values of α and γ are $\alpha \approx 0.454$ and $\gamma \approx 2.4975$.

To justify the asymptotic form used above thereby reaching step 14 of the 20 steps of [12] (which is where we intend to stop) requires checking that $F_{yy}(\rho_A, A(\rho_A)) \neq 0$. Fortunately, we can compute F_{yy} modulo some “known” functions (in the usual sense) as:

$$F_{yy} = e^{p(y)} t (t^2(1 + y - e^y c_1(t))^2 + e^y c_1(t) - 1), \text{ where} \\ p(y) = (-1/2)t^2(2 + 2y + y^2 + c_2(t) - 2e^y c_3(t)) + c_4(t),$$

with c_2 , c_3 and c_4 analytic and real at ρ_A . Further $c_1(t) = \exp w(t)$ where $w(t)$ is a series with positive coefficients. So $e^y c_1 - 1 > 0$ at $t = \rho_A$ and thus $F_{yy}(\rho_A, A(\rho_A)) \neq 0$.

Recall that atomic cohorts of arch systems with $n + 1$ arches are in bijection with cohorts of arch systems with n arches so to obtain the general asymptotics we multiply the constant term from the atomic asymptotics by γ yielding:

Theorem 12. *The number of cohorts of arch systems with n arches behaves asymptotically as $cn^{-3/2}\gamma^n$ for some positive constants c and γ .*

As noted previously, we estimate that $c \approx 1.13$ and $\gamma \approx 2.4975$ and could certainly establish rigorously based on the previous arguments that $2.49 < \gamma < 2.5$.

6 The main cohort, and comparing generating functions associated with cohorts

We start this section by defining a special cohort of arch systems of any size n and studying its properties. We specifically deal with the generating function of any avoidance class $\text{Av}(X)$ for an arch system X in this cohort and with the number of arch systems contained in this cohort. This will complete the proofs of our claims of the introduction. This special cohort is called the main cohort.

We shall see that the main cohort is the largest in terms of generating function associated to an avoidance class. This follows from the results reported in this section about the comparison between cohorts (of structures of the same size) with respect to the (common) generating functions of the classes $\text{Av}(X)$ they represent. Specifically, we provide some rules on arch systems that allow the comparison between the generating functions of their cohorts and consequently show that the main cohort is largest in the sense that its generating function dominates that of any other cohort.

6.1 The main cohort

Following the discussion of Subsection 5.1, for each n there is a unique cohort of structures of size n that arises from all unary-binary plane forests (i.e., no large atoms are involved) – by definition such forests consist of at most two trees which are themselves unary-binary trees. We call this the *main cohort* for structures of size n and denote it by \mathcal{M}_n . A representative of this cohort is the system N_n of n nested arches whose corresponding forest is a chain of n nodes. But from its description in terms of forests it is clear that the main cohort also includes all the arch systems of size n that can be built using the following operations, and only these: concatenate two atoms that belong to \mathcal{M}_j and \mathcal{M}_k for $j + k = n$, or place an arch over an arch system of \mathcal{M}_{n-1} . For the same reason, if we let M_n denote the number of atoms (i.e., trees) of size n in the cohort \mathcal{M}_n , it is immediate that the generating function $M(t) = \sum M_n t^n$ satisfies:

$$M = t + tM + tM^2.$$

This identifies (M_n) as the sequence of Motzkin numbers (offset by 1):

$$M_{n+1} = \text{Motz}_n = \sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2k} \text{Cat}_k.$$

Recalling that the number of atoms in the main cohort for structures of size $n + 1$ is equal to the total number of arch systems in the main cohort for structures of size n , we obtain:

Proposition 13. *The size of the main cohort for structures of size n is the n -th Motzkin number: $|\mathcal{M}_n| = \text{Motz}_n$.*

Furthermore, to \mathcal{M}_n corresponds one generating function: that of any $\text{Av}(X)$ for $X \in \mathcal{M}_n$. Taking $X = N_n$, where N_n is the nest of n arches, these generating functions $C^{(n)}$ are easily seen to satisfy

$$C^{(n)} = \frac{1}{1 - tC^{(n-1)}} \quad (\text{with initial condition } C^{(0)} = 1),$$

giving that:

Proposition 14. *For any structure X in \mathcal{M}_n , the generating function of $\text{Av}(X)$ is $C^{(n)}$.*

This justifies the remarks concerning the sequence of generating functions $(C^{(n)})$ made in the introduction.

Note that Proposition 14 provides an alternative proof of the enumeration of classes $\text{Av}(231, \pi)$ (by $C^{(n)}$ for $n = |\pi|$) for several families of patterns π that appear in the literature: namely decreasing patterns and patterns of the form $1n(n-1) \dots 32$ [8], reverse of 2-layered permutations and 132-avoiding wedge-patterns of [15, 16], and patterns $\lambda_k \oplus \lambda_{n-k}$ of [4]. Indeed, all such patterns belong to the main cohort of the corresponding size.

For any pair of structures A and B in \mathcal{M}_n we can find a chain of \sim -equivalences between them. Therefore, the bijective proofs of Subsection 4.1 provide a bijection between $\text{Av}(A)$ and $\text{Av}(B)$. A special case of this statement answers a question raised in [16] about the description of a bijection between $\text{Av}(132, \pi)$ and $\text{Av}(132, \tau)$, for π any 2-layered pattern and τ any 132-avoiding wedge-pattern.

The name *main* cohort has been chosen because, as already announced, this cohort is the largest for the generating function associated to its avoidance classes. We prove below that $C^{(n)}$ dominates (term by term) the generating function F_X of $\text{Av}(X)$ for any arch system X of size n . Moreover, unless $X \in \mathcal{M}_n$, eventually $C^{(n)}$ dominates F_X strictly.

We suspect that the main cohort might also be the largest in an other sense, namely in terms of how many elements a cohort of structures of a given size can contain. We shall discuss the question of maximality of the main cohort in terms of cardinality in Section 7

6.2 Comparing avoidance classes between cohorts

One (maybe the most important) purpose of this subsection is to prove that the main cohort is the largest in terms of the generating function associated with $\text{Av}(X)$, for X in this cohort. This claim is proved as a consequence of more general statements, that allow the comparison of such generating functions associated with various cohorts.

Let us start by introducing some notation. For any cohort \mathcal{C} , and any A and B in \mathcal{C} , we know from Theorem 7 that $\text{Av}(A)$ and $\text{Av}(B)$ have the same generating function. We may therefore associate this generating function with \mathcal{C} and, when doing so, we denote it $F_{\mathcal{C}}$. For two cohorts \mathcal{C} and \mathcal{D} , with generating functions $F_{\mathcal{C}} = \sum c_n t^n$ and $F_{\mathcal{D}} = \sum d_n t^n$, we write $\mathcal{C} \leq \mathcal{D}$ when for all n , $c_n \leq d_n$. We also write $\mathcal{C} < \mathcal{D}$ when $\mathcal{C} \leq \mathcal{D}$ and there exists n_0 such that for all $n \geq n_0$ $c_n < d_n$. Finally, for any arch system A , we denote by \mathcal{C}_A the cohort containing A , that is to say the equivalence class of A for \sim .

Variations on the bijective proofs of cases (1), (2) and case (4) of Theorem 7 allow us to provide some recursive rules for the comparison of cohorts \mathcal{C}_A for this order \leq of comparison of their generating functions.

Proposition 15. *For any arch systems A and B , if $\mathcal{C}_A \leq \mathcal{C}_B$ then $\mathcal{C}_{\overline{A}} \leq \mathcal{C}_{\overline{B}}$, and if $\mathcal{C}_A < \mathcal{C}_B$ then $\mathcal{C}_{\overline{A}} < \mathcal{C}_{\overline{B}}$.*

Proof. To prove that $\mathcal{C}_{\overline{A}} \leq \mathcal{C}_{\overline{B}}$ (resp. $\mathcal{C}_{\overline{A}} < \mathcal{C}_{\overline{B}}$) we should compare (term by term) the enumeration sequences of $\text{Av}(\overline{A})$ and $\text{Av}(\overline{B})$ proving that the latter is weakly (resp. eventually strictly) larger. To do that it is enough to give a size-preserving injection (resp. size-preserving injection which fails to be surjective in any size from some n_0) from $\text{Av}(\overline{A})$ to $\text{Av}(\overline{B})$ given one from $\text{Av}(A)$ to $\text{Av}(B)$. This follows immediately from the same arguments used in the proof of case (1) of Theorem 7, essentially by replacing “bijection” wherever it occurs by “injection” (resp. “injection which is not surjective in any size from some n'_0 ” – observe that $n_0 = n'_0 + 1$). \square

Proposition 16. *For any arch system A and any atom b , if $\mathcal{C}_A \leq \mathcal{C}_b$ then $\mathcal{C}_{PAQ} \leq \mathcal{C}_{PbQ}$, and unless $A = a$ is an atom such that $a \sim b$, $\mathcal{C}_{PAQ} < \mathcal{C}_{PbQ}$. Moreover, if $\mathcal{C}_A < \mathcal{C}_b$ then $\mathcal{C}_{PAQ} < \mathcal{C}_{PbQ}$.*

Proof. To prove $\mathcal{C}_{PAQ} \leq \mathcal{C}_{PbQ}$ we describe a size-preserving injection from $\text{Av}(PAQ)$ to $\text{Av}(PbQ)$ based on one from $\text{Av}(A)$ to $\text{Av}(b)$.

With the same decomposition used in the proof of case (2) of Theorem 7 we see that, given an injection from $\text{Av}(A)$ to $\text{Av}(b)$, an injection from $\text{Av}(PAQ)$ to $\text{Av}(PbQ)$ can be constructed. This uses the fact that if a concatenation $I_1 I_2 \dots I_k$ of arch systems avoids A , then each arch system I_i must avoid A .

If $\mathcal{C}_A < \mathcal{C}_b$, this injection cannot possibly be a bijection (except for the first few sizes $n \leq \text{some } n_0$). Indeed, it is easy to construct elements of any size $n + |P| + |Q|$ of $\text{Av}(PbQ)$ that do not lie in its image from elements of $\text{Av}(b)$ of size n that do not lie in the image of the original injection. In fact, for this injection to be a bijection, we need two conditions. The first one is that a concatenation of arch systems should avoid A if and only if each arch system in this sequence avoids A : this happens exactly when A is an atom. The

second condition is that the injection from $\text{Av}(A)$ to $\text{Av}(b)$ needs to be a bijection, i.e., that $A \sim b$. \square

Propositions 15 and 16 are enough to prove that the main cohorts $\mathcal{M}_n = \mathcal{C}_{N_n}$ are the largest in the sense that their generating functions $F_{\mathcal{M}_n}$ eventually dominate the generating functions of any other cohort of arch systems of size n . Recall that N_n is the arch system consisting of n nested arches.

Proposition 17. *For every arch system A of size n , either A is in the cohort of N_n or $\mathcal{C}_A < \mathcal{C}_{N_n}$.*

Proof. The proof is by induction. The base case ($n = 1$) is clear. So assume that $n \geq 2$ and that the statement holds for all $n' < n$. Consider an arch system A of size n . Either $A = (\overline{X})$ or $A = Xa$ where a is an atom and X a non empty arch system.

In the first case, by induction, we know that exactly one of the following holds:

- X is in the cohort of N_{n-1} ; and then A is in the cohort of N_n by rule (1).
- $\mathcal{C}_X < \mathcal{C}_{N_{n-1}}$; but then Proposition 15 ensures that $\mathcal{C}_A < \mathcal{C}_{N_n}$.

In the second case, denoting the size of X by j , we know that either X is in the cohort of N_j or $\mathcal{C}_X < \mathcal{C}_{N_j}$.

Assume first that $X \sim N_j$. If X is an atom, then $Xa \sim N_ja$ by rule (2). Now either $a \sim N_{n-j}$, in which case $N_ja \sim N_jN_{n-j} \sim N_n$ so that $A = Xa$ is in the cohort of N_n ; or $\mathcal{C}_a < \mathcal{C}_{N_{n-j}}$, and Proposition 16 ensures that $\mathcal{C}_A = \mathcal{C}_{Xa} < \mathcal{C}_{XN_{n-j}} \leq \mathcal{C}_{N_jN_{n-j}}$ (using Proposition 16 again, since $\mathcal{C}_X \leq \mathcal{C}_{N_j}$ by induction). We conclude using $\mathcal{C}_{N_jN_{n-j}} = \mathcal{C}_{N_n}$.

If X is not an atom, we deduce from $X \sim N_j$ that $\mathcal{C}_X \leq \mathcal{C}_{N_j}$ and Proposition 16 (applied twice) and induction ensure that $\mathcal{C}_{Xa} < \mathcal{C}_{N_ja} \leq \mathcal{C}_{N_jN_{n-j}} = \mathcal{C}_{N_n}$.

The last case is $\mathcal{C}_X < \mathcal{C}_{N_j}$, in which case Proposition 16 gives $\mathcal{C}_{Xa} < \mathcal{C}_{N_ja} \leq \mathcal{C}_{N_n}$ (as before). \square

Finally, the bijective proof of the specialization of case (4) of Theorem 7 (where one of a , b or c is empty) can also be adapted to the comparison of cohorts.

Proposition 18. *For any arch system A , and any arch system b which is an atom or empty, $\mathcal{C}_{A(\overline{b})} \leq \mathcal{C}_{(\overline{b}A)}$. Moreover, unless A is an atom, $\mathcal{C}_{A(\overline{b})} < \mathcal{C}_{(\overline{b}A)}$.*

The arguments are very similar to the above, and we do not give them here as we have not been explicit about the nature of the bijection for that specialized case (and the result *does not* apply to the full case.)

7 Cardinalities of cohorts: some hints about extremal cases

Many questions may be asked about cohorts. In Section 5 we have described them as a combinatorial class and obtained results on their enumeration most importantly asymptotically. Section 6 has addressed the comparison between cohorts in terms of the generating

function of their corresponding avoidance classes. But cohorts may be compared with respect to other criteria, the most natural one being their cardinalities. In the following, we give preliminary results (and open questions) regarding the cardinalities of cohorts. We focus on extremal cases, namely on candidates to being the cohort with maximal cardinality (among cohorts of structures of a given size) and on singleton cohorts.

7.1 Is the main cohort the largest?

We have seen in Subsection 6.1 that the main cohort \mathcal{M}_n for arch systems of size n (of which a representative is the system N_n of n nested arches) contains Motz_n arch systems. Recall that Motzkin numbers grow exponentially like 3^n .

As far as we can observe on the data we have, the main cohort is always (and by far) the one with largest cardinality among all cohorts of structures of a given size. And it is tempting to formulate the conjecture that this generalizes to any size, i.e., that for every positive integer $n \geq 3$ the size of \mathcal{M}_n is greater than the size of any other cohort of arch systems of size n . If this would hold, the main cohort would then be the largest according to two criteria: cardinality, and generating function of the corresponding avoidance class.

In favor of this conjecture, notice that the main cohort is constructed using the smallest building blocks, i.e., any other cohort must involve somehow one or more atoms consisting of at least four arches (such as $(\cap\cap\cap)$). Our attempts to formalize this idea into a proof that the main cohort is the one with largest cardinality have however been unsuccessful.

We are now inclined to believe that cohorts that are larger than the main cohort may exist even if this does not happen before a very large value of n . We have tried the route that consists in building such a cohort for very large n using our asymptotic estimate of the number of cohorts. Even if we were not successful in this direction as well the method we have tried seems worthy of being described. The idea is to try and build an arch system $A = a_1 a_2 \dots a_k$ of size n whose cohort contains α^n elements, for some $\alpha > 3$. To maximize the cardinality of the cohort of A , the a_i would be atoms taken in distinct cohorts of the same size s , these cohorts being the k largest cohorts of size- s structures, in terms of their cardinality. Assuming that there exist constants b and c such that there are at least b^s cohorts of cardinality at least c^s among structures of size s , with $bc > 3$, we would be able to prove that the main cohort of structures of size n is not the largest, at least for sufficiently large n .

Several factors contribute to making the cardinality of a cohort large. One is minimizing the size of the atoms used in building a representative of this cohort and in this regard the main cohort seems a good candidate to being a large cohort. Another is permuting atoms that belong to different cohorts a feature that does not arise in the main cohort. It is hard to foresee how these two factors interplay so we will not state any bold conjecture but simply ask the question: is the main cohort \mathcal{M}_n the one with largest cardinality among all cohorts of structures of size n ?

7.2 Singleton cohorts

At the other end of the chain we consider the cohorts that contain one single arch system. Modulo Conjecture 1 these correspond to the only arch systems A that can be recognized directly from the generating function of $\text{Av}(A)$.

Proposition 19. *The cohort of a (non empty) arch system A is a singleton if and only if:*

- $A = b^k$ where $k \geq 3$ and b is an atom which is the only atom in its cohort² or
- $A = a^2$ where a is an atom whose contents are some b^k as in the first condition or
- A is an atom whose contents are either empty or some b^k as in the first condition.

Moreover the atoms which are the only atoms in their cohort are: \cap and the atoms whose (non empty) contents belong to a singleton cohort.

Proof. Suppose first that an arch system A is a concatenation of two or more atoms. For such arch systems rule (3) would yield more than one element in A 's cohort unless these atoms were all identical. Further rule (2) would do likewise if that atom were not the only atom in its cohort.

On the other hand if these conditions are met and A is a concatenation of at least three atoms then rules (1) and (4) cannot be applied so such A are indeed arch systems whose cohort is a singleton.

If the cohort of $A = a^2$ is a singleton and $a = \overline{[X]}$ then clearly the cohort of X must be a singleton (else rule (1) would apply). Furthermore, X must be the concatenation of at least three atoms, or else rule (4) could be applied in A . Conversely, if X satisfies these conditions then none of the rules can be applied to yield any other element of A 's cohort.

If $A = \overline{[X]}$ is an atom that forms a singleton cohort, then its contents X (if not empty) must belong to a singleton cohort (else rule (1) would apply). X cannot be an atom since $\overline{[Y]} \sim \cap \overline{[Y]}$ (from rule (4) with $c = \overline{[Y]}$ and a and b empty). Similarly, X cannot be the concatenation of two atoms, since $\overline{[ab]} = a\overline{[b]}$ (from rule (4) with c empty). So X must satisfy the first condition. Conversely if the contents X of A do satisfy this condition then the cohort of A will be a singleton: indeed, the only rules allowing one to find a \sim -equivalent of an atom are rule (1) and the special cases of rule (4) – which do not apply here since X is the concatenation of at least three atoms.

If an atom is the only atom in its cohort, then obviously its contents are either empty or belong to a singleton cohort. Conversely, consider an atom that is either \cap or $\overline{[X]}$ where the cohort of X is a singleton. Certainly, \cap is the only atom in its cohort (which is indeed a singleton here). We claim that for any arch system X whose cohort is a singleton, $\overline{[X]}$ is the only atom in its cohort. Such X satisfies one of the conditions of Proposition 19. If $X = b^k$ as in the first condition, then none of the rules (1) to (4) apply to $\overline{[X]}$ – note

² Note that this condition is less restrictive than the cohort of b being a singleton.

that here the cohort of $\overline{[X]}$ is actually a singleton from the third condition. If $X = a^2$ as in the second condition, then only special cases of rule (4) apply to $\overline{[X]} = \overline{[aa]}$, producing two \sim -equivalent to $\overline{[X]}$, namely $\overline{[a]a}$ and $a\overline{[a]}$. If $X = \overline{[Y]}$ is an atom as in the third condition, then only special cases of rule (4) apply to $\overline{[X]} = \overline{[\overline{[Y]}]}$ producing two (one if Y is empty) \sim -equivalent to $\overline{[X]}$, namely $\overline{[Y]} \cap$ and $\cap \overline{[Y]}$. In all cases, we observe that $\overline{[X]}$ is indeed the only atom in its cohort. \square

We may wonder how many singleton cohorts there are. In order to translate the conditions of Proposition 19 into recurrences allowing to count singleton cohorts we introduce several auxiliary functions: $S_1(n)$ counts the atomic singleton cohorts, $S_2(n)$ counts the singleton cohorts of the form a^2 , and $S_{\geq 3}(n)$ counts the singleton cohorts of the form b^k for $k \geq 3$. Also $A(n)$ counts the number of cohorts that contain a single atom. Then we obtain as recursive conditions:

$$\begin{aligned} S_1(n) &= S_{\geq 3}(n-1), \\ S_2(n) &= \begin{cases} 0 & n \text{ odd} \\ S_{\geq 3}(n/2-1) & n \text{ even} \end{cases}, \\ S_{\geq 3}(n) &= \sum_{k \geq 3, k|n} A(n/k), \\ A(n) &= S_1(n-1) + S_2(n-1) + S_{\geq 3}(n-1). \end{aligned}$$

These together with appropriate boundary conditions determine all the functions and hence the total number $S(n)$ of singleton cohorts, $S(n) = S_1(n) + S_2(n) + S_{\geq 3}(n)$. Note that the actual recurrences really just involve $S_{\geq 3}$ and A as follows:

$$\begin{aligned} S_{\geq 3}(n) &= \sum_{k \geq 3, k|n} A(n/k), \\ A(n) &= \begin{cases} S_{\geq 3}(n-1) + S_{\geq 3}(n-2) + S_{\geq 3}((n-3)/2) & n \text{ odd} \\ S_{\geq 3}(n-1) + S_{\geq 3}(n-2) & n \text{ even} \end{cases}. \end{aligned}$$

It might be possible to derive from the above some information on the “average” behavior of $S(n)$, the number of singleton cohorts of n -arch systems. But this would likely involve tricky computations with number theoretic arguments, that we leave aside for the moment.

8 Conclusions and open problems

8.1 Conjectures and questions left open

Several questions are left open in this work. An important one is certainly to provide a completely explicit bijective proof of our main result (Theorem 7), that is: providing an explicit bijection for case 4 of this theorem. Even a sensible combinatorial explanation of the rather tidy expression for $F_{a(\overline{bc})}$ in terms of F_a , F_b and F_c would represent progress in

this direction. Another problem is to prove or disprove that the main cohort is the largest in terms of number of elements it contains.

But the most intriguing problem is certainly to prove a converse statement to our main theorem: that not only does \sim refine Wilf-equivalence but also coincides with it. This is stated as Conjecture 1 at the beginning of our paper, and we offer a stronger version of this conjecture, by way of conclusion.

Conjecture 20. For any two arch systems A and B , both with n arches, either A and B are in the same cohort (i.e. $A \sim B$), or the enumeration sequences of $\text{Av}(A)$ and $\text{Av}(B)$ differ at the latest at size $2n - 2$.

We have been able to check that this stronger conjecture holds up to arch systems A and B of size 15. We further know that the size $2n - 2$ is the smallest one for which such a conjecture could be true. Indeed, we have identified families of arch systems A_n and B_n of any size $n \geq 4$ such that the enumeration sequences of $\text{Av}(A_n)$ and $\text{Av}(B_n)$ coincide up to size $2n - 3$ but differ at $2n - 2$. These are described below.

Let \cap^k denote the concatenation of k empty arches. Now, for any $n \geq 4$, set $C_n = \overline{\cap^{n-4}}$, $A_n = \overline{C_n \cap \cap}$, and $B_n = \cap \cap \overline{C_n}$. We claim that there is a size preserving bijection between $\text{Av}(A_n)$ and $\text{Av}(B_n)$ restricted to arch systems with at most $2n - 3$ arches, but that there are more arch systems of size $2n - 2$ avoiding A_n than B_n .

Observe that $A_n = \overline{bA}$ and $B_n = A\overline{b}$ for $b = C_n$ and $A = \cap \cap$. So the proof of Proposition 18 provides an injection φ from $\text{Av}(B_n)$ to $\text{Av}(A_n)$. It is relatively easy to see that φ is actually a bijection when restricted to arch systems with at most $2n - 3$ arches. This essentially amounts to examining where these at most $2n - 3$ arches can be in arch systems containing C_n but avoiding A_n . It is also not hard to see that the arch system $\overline{C_n \cap} \overline{C_n \cap}$ of size $2n - 2$ avoids A_n but is not in the image of φ .

8.2 Extensions of our work in the Catalan world(s)

So far, we have studied Wilf-equivalences among classes of the form $\text{Av}(231, \pi)$. Equivalently, these classes are described as subposets (also called order ideals) defined by the avoidance of *one* substructure in the Catalan poset of 231-avoiding permutations (or Dyck paths, or plane forests, or arch systems). A natural continuation is then to consider such subposets defined by the avoidance of *several* substructures, starting with classes $\text{Av}(231, \pi, \tau)$.

Because of symmetries, we don't expect to find as many different enumeration sequences for classes of the form $\text{Av}(231, \pi, \tau)$ as $\text{Cat}_n \times \text{Cat}_m$, where $n = |\pi|$ and $m = |\tau|$. The naive upper bound on the number of Wilf-equivalence classes in this case is rather the number of symmetry classes of avoidance classes $\text{Av}(231, \pi, \tau)$. For small sizes of π and τ , Table 1 shows some experimental data comparing these upper bounds with what we suspect is the actual number of Wilf-equivalence classes. These estimations are actually lower bounds, based on experimental data only: fixing n and m , they count the number of different sequences $(|\text{Av}_k(231, \pi, \tau)|)_{1 \leq k \leq N}$ for some fixed value of N (at least 15 in our experiments). In support of our guess that these lower bounds are actually tight, let us

mention that we have chosen N great enough so that a stabilization of the number of different sequences occur. Namely, for any choice of n and m in Table 1, the number of different sequences $(|\text{Av}_k(231, \pi, \tau)|)_{1 \leq k \leq K}$ is the same for $K = N - 2, N - 1$ and N .

n	m	Number of symmetry classes	Estimated number of Wilf-classes
3	3	7	3
3	4	25	7
3	5	43	10
3	6	82	16
4	4	55	16
4	5	257	65
5	5	461	166

Table 1: Number of Wilf-equivalence classes of avoidance classes of the form $\text{Av}(231, \pi, \tau)$ where $n = |\pi|$ and $m = |\tau|$: the naive upper bound, and an estimation of the actual value, obtained from data up to permutations of size 15 or more.

There is of course also the question of whether our binary relation \sim could be used to explain Wilf-equivalences between classes of the form $\text{Av}(231, \pi, \tau)$. It should be noticed that $\pi \sim \pi'$ and $\tau \sim \tau'$ is not a sufficient condition for $\text{Av}(231, \pi, \tau)$ and $\text{Av}(231, \pi', \tau')$ to be Wilf-equivalent. For instance, although $4321 \sim 4312$, the classes $\text{Av}(231, 4123, 4321)$ and $\text{Av}(231, 4123, 4312)$ are not enumerated by the same sequence. However, we suspect that a stronger condition on π, π', τ and τ' could guarantee that $\text{Av}(231, \pi, \tau)$ and $\text{Av}(231, \pi', \tau')$ are Wilf-equivalent. Informally, this condition would state that $\pi \sim \pi'$, $\tau \sim \tau'$ and that the sequence of \sim -equivalence to go from π to π' also allows to go from τ to τ' . This condition is still to be expressed properly, but we can illustrate it by a (very simple) example: taking $\pi = 41235$, $\pi' = 15234$, $\tau = 43215$ and $\tau' = 15432$, the equivalences $\pi \sim \pi'$ and $\tau \sim \tau'$ both follow by rule (3), permuting atoms of size 4 and 1, and the enumerations of $\text{Av}(231, \pi, \tau)$ and $\text{Av}(231, \pi', \tau')$ agree up to size 17.

While staying in the Catalan world, examining enumeration of non-principal subclasses of $\text{Av}(231)$ as explained above is not the only question to investigate. Indeed, as explained in Subsection 2.2, other families of Catalan objects have a natural substructure order, which give rise to posets that are different from the one of $\text{Av}(231)$. In these posets, the problem of classifying Wilf-equivalences between principal (resp. general) subposets – defined by the avoidance of one (resp. several) substructures – is as natural as in the case of $\text{Av}(231)$.

In particular consider the poset associated with $\text{Av}(321)$. Preliminary experiments on classes $\text{Av}(321, \pi)$ have shown that it seems possible to define some rules in the same flavor as those defining \sim in such a way that if π and τ are equivalent for those new rules then $\text{Av}(321, \pi)$ and $\text{Av}(321, \tau)$ are Wilf-equivalent. Our investigations on this topic are just at the beginning and should be continued in future work. As in the case of $\text{Av}(231)$, we can hope for a set of rules that captures Wilf-equivalence completely (parallel to Conjecture 1), and for an asymptotic enumeration of equivalence classes for the analogue

of \sim in this context (parallel to Theorem 12). Again we expect a collapse in comparison with the Catalan numbers but which we believe will be less dramatic than in the $\text{Av}(231)$ case.

8.3 Further extensions

To the best of our knowledge, this work is the first global approach to the study of Wilf-equivalences, a popular topic of research in the field of permutation patterns from its early days until now – and arguably so in the wider context of hereditary classes of combinatorial structures. It is performed in the context of Catalan structures, and more precisely in that of permutations avoiding 231 and another pattern π – which we could call *principal subclasses* of $\text{Av}(231)$. We believe that similar investigations aiming at classifying all Wilf-equivalences between principal subclasses of (well-behaved) permutation classes should be carried out. One promising example being considered by the first author, Cheyne Homberger and Jay Pantone is the class of separable permutations, $\text{Av}(2413, 3142)$. This comment is motivated in part by the results of [3] which provide a partial parallel of Proposition 5 and those of [6] extending the work of [9, 17] to this context, but more generally because the separable permutations permit several other “well-structured” representations.

We can even hope to extend our ideas further, to a partial classification of Wilf-equivalences between principal permutation classes, i.e., classes of permutations defined by the avoidance of a single pattern. The framework of matchings with excluded sub-matchings, as defined in [13], could provide a good tool for that. Matchings are similar to arch systems, but arches are allowed to cross. Namely, a matching of size n is a set of n arches connecting $2n$ points arranged along a baseline, with all arches above the baseline. Obviously, our families $\text{Av}(A)$ of arch systems avoiding a given arch system A can be seen as matchings with excluded sub-matchings: namely, those avoiding \cap and A . But (principal) permutation classes $\text{Av}(\pi)$ can also be represented as matchings with excluded sub-matchings. Indeed, permutations are in immediate correspondence with matchings having all their arches opened before any arch is closed, or equivalently with matchings avoiding $\cap\cap$. Under this correspondence, a permutation class $\text{Av}(\pi)$ is simply the class of matchings avoiding $\cap\cap$ and the matching encoding π . If it were possible to adapt our work to such cases, and in particular to provide an upper bound on the asymptotic number of Wilf-equivalence classes of principal permutation classes, this would be a major achievement in the field.

More directly, the first author and Jonathan Bloom have recently considered further Wilf-equivalences that arise from arch systems with the addition of a new singleton object (which could be thought of as an arch whose endpoints coincide), or where the notion of avoidance is with respect to the analogous concept to factor order in words (where the present paper concerns the concept analogous to subword order). These results will appear in a future paper.

Acknowledgments

Much of the work in this paper was supported by the software suite *PermLab* [1], and extensions of it. Generating function computations were carried out in *Mathematica* [20].

We are grateful to Cheyne Homberger for pointing out references [9] and [17] to our attention, and to Vit Jelinek for the idea of extending our work to other Catalan posets. We are especially grateful to Jonathan Bloom for finding the full bijective proof of case (4) of Theorem 7.

Mathilde Bouvel would like to thank the Department of Computer Science of the University of Otago for their hospitality and support in January and February of 2014 when most of this work was carried out.

References

- [1] Michael Albert. *PermLab: Software for Permutation Patterns*, 2012. Software package available at: <http://www.cs.otago.ac.nz/PermLab>.
- [2] Michael Albert and Michael Atkinson. Simple permutations and pattern restricted permutations. *Discrete Math.* 300 (1-3):1–15, 2005.
- [3] Michael Albert, Michael Atkinson, and Vincent Vatter. Subclasses of the separable permutations. *Bull. Lond. Math. Soc.* 43 (5):859–870, 2011.
- [4] Michael Albert and Mathilde Bouvel. Operators of equivalent sorting power and related Wilf-equivalences. *DMTCS Proceedings (FPSAC)* AS:671–682, 2013.
- [5] Michael Albert and Mathilde Bouvel. Operators of equivalent sorting power and related Wilf-equivalences. *Elec. J. Combin.* 21(4), 2014. #P4.11.
- [6] Michael Albert, Cheyne Homberger, and Jay Pantone. Equipopularity Classes in the Separable Permutations, 2014. Available as [arXiv:1410.7312v1](https://arxiv.org/abs/1410.7312v1).
- [7] Axel Bacher, Antonio Bernini, Luca Ferrari, Benjamin Gunby, Renzo Pinzani, and Julian West. The Dyck pattern poset. *Discrete Math.* 321:12–23, 2014.
- [8] Timothy Chow and Julian West. Forbidden subsequences and Chebyshev polynomials. *Discrete Math.* 204:119–128, 1999.
- [9] Lynn Chua and Krishanu Roy Sankar. Equipopularity Classes of 132-Avoiding Permutations. *Elec. J. Combin.* 21(1), 2014. #P1.59.
- [10] Michael Dairyko, Lara Pudwell, Samantha Tyner, and Casey Wynn. Non-Contiguous Pattern Avoidance in Binary Trees. *Elec. J. Combin.* 19(3), 2012. #P22.
- [11] Philippe Flajolet and Robert Sedgewick. *Analytic combinatorics*. Cambridge University Press, 2009.
- [12] Frank Harary, Robert W. Robinson, and Allen J. Schwenk. Twenty-step algorithm for determining the asymptotic number of trees of various species. *J. Austral. Math. Soc. Ser. A* 20 (4):483–503, 1975.

- [13] Vit Jelinek. Dyck paths and pattern-avoiding matchings. *Europ. J. Combin.* 28:202–213, 2007.
- [14] Donald E. Knuth. *The Art of Computer Programming*. Addison-Wesley Publishing Co., Reading, Mass.-London-Amsterdam, Second edition, 1975. Volume 1: Fundamental algorithms, Addison-Wesley Series in Computer Science and Information Processing.
- [15] Toufik Mansour and Alek Vainshtein. Restricted 132-avoiding permutations. *Adv. in Applied Math.* 26:258–269, 2001.
- [16] ———. Restricted permutations and Chebyshev polynomials. *Séminaire Lotharingien de Combinatoire* 47, 2002. Article B47c.
- [17] Kate Rudolph. Pattern Popularity in 132-avoiding Permutations. *Elec. J. Combin.* 20(1), 2013. #P8.
- [18] Richard Stanley. *Catalan Numbers*. Cambridge University Press, 2015.
- [19] ———. *Catalan Addendum*. <http://www-math.mit.edu/~rstan/ec/catadd.pdf> (version of 25 May 2013).
- [20] Wolfram Research, Inc. *Mathematica, Version 9.0*. Wolfram Research, Inc., Version 9.0, 2012.