Counting Lines: How Geometry Becomes Combinatorics

Ryan Bansal

Abstract

Two of the best-known multiplication rules in algebraic combinatorics read like the instructions of a game. Pieri's rule adds boxes to a Young diagram, never two in the same column; Monk's rule swaps two entries of a permutation across a marked position, provided the number of inversions rises by exactly one. The rules are easy to apply and, as usually stated, silent about where they come from. This exposition develops the single geometric mechanism that produces both: a flag is a measuring device, and as a fixed flag grows one dimension at a time, a moving subspace becomes visible in stages. In a Grassmannian, this measurement is a staircase, recorded by a partition and realized geometrically by a Schubert variety; Pieri's rule then resolves the classical count of the two lines meeting four general lines in $\mathbb{P}^3$. In the complete flag variety, the same measurement becomes a two-dimensional rank table, recorded by a permutation and organized by Bruhat order. Monk's rule emerges as the precise statement of the metaphor: one simple incidence condition moves a Schubert condition along exactly the Bruhat covers crossing a chosen position. The two rules are thus one geometric event seen at two resolutions. The elementary linear-algebraic geometry is developed directly, while the deeper intersection-theoretic inputs are isolated and stated explicitly as standard facts.

1 Introduction

A classical theorem of enumerative geometry says that four sufficiently general lines in $\mathbb{P}^3$ have exactly two common transversals: there are two lines meeting all four [8]. At first this sounds like a problem about solving equations for an unknown line. Schubert calculus changes the viewpoint. Instead of studying one line at a time, it places all possible lines into a single geometric space and studies the conditions cut out inside that space. By the middle of this paper, the count of two will fall out of a short calculation with Young diagrams.

The calculation that produces the count is itself a small surprise. It uses Pieri's rule, which says: to impose the simplest incidence condition, add one box to a Young diagram in every way that leaves a Young diagram. Deeper in the subject sits Monk's rule, which says: to impose the simplest condition on a complete flag, swap two entries of a permutation across a marked position, provided the number of inversions rises by exactly one. Stated this way, the rules are precise, easy to teach, and opaque. Nothing in "add a box" or "swap across the mark" mentions a line in space. Why should geometric problems be governed by bookkeeping rules about diagrams and permutations—and why these rules rather than others?

This paper is organized around one mechanism that answers the question. A condition such as "meet a fixed line" is first turned into a condition on a Grassmannian, the space of all subspaces of a fixed dimension in $\mathbb{C}^n$. Then a fixed flag is used as a measuring device: as the flag grows one dimension at a time, a moving subspace becomes visible in stages. Every combinatorial object in the story records this measurement. In a Grassmannian, the record is a staircase, and a partition lists how many steps early each jump of the staircase occurs: a box is one step of early visibility. In a complete flag variety, the record is a two-dimensional rank table, and a permutation is its compact label; Bruhat order tracks degeneration, the one-way motion by which measurements become more special in limits. In this language the rules lose their arbitrariness. Imposing one simple incidence condition forces the measurement to become special by exactly one step, and the terms of each rule list the legal steps: in the simplest case of Pieri's rule, the new box is a jump of the staircase arriving one step early where it does not collide with its neighbor, and a Monk swap is a minimal specialization of the rank table crossing the position the condition watches. The two rules are the same geometric event, read at two resolutions.

The subject has excellent expositions, from Fulton's textbook development [3] to Sottile's notes [8] and Gillespie's wide-ranging tour [4]. Those works aim at coverage: the full combinatorial toolkit, or several families of varieties at once. This paper aims at one mechanism, carried as far as it will go by elementary means. Everything before cohomology is proved directly with linear algebra, at the level of echelon forms and dimension counts. The genuinely deep inputs—the existence of Schubert classes and their intersection behavior, the flag-variety cell decomposition, the coinvariant presentation, and one controlled degeneration—are not smuggled in: each is stated precisely where it is used and marked as a standard fact, so the reader always knows which parts of the story are elementary and which are machinery. A reader with linear algebra and no algebraic geometry can follow the whole path, and will meet each combinatorial object at the moment the geometry forces it to appear.

Grassmannians also reach far beyond this classical story: a k-dimensional principal-component model is a point of a Grassmannian [2], subspaces serve as codewords in network coding [5], and Grassmannians organize the lines lying on a variety. These applications are not pursued here, but they are one more reason the basic dictionary of this paper—subspaces as points, incidence conditions as subvarieties—is worth learning well.

Throughout the paper, all vector spaces and varieties are over $\mathbb{C}$, and all dimensions are complex dimensions. The word *general* means that the chosen objects avoid a proper algebraic exceptional set. In the examples below, this is exactly what rules out accidental coincidences and makes intersections have the expected dimension. Thus the count of two transversals is a count of complex projective lines; a real configuration may have fewer real solutions.

We begin with echelon forms on Grassmannians, then explain how flags and partitions encode Schubert conditions. Pieri's rule gives a first model for how multiplying Schubert classes corresponds to imposing incidence conditions, and settles the four-lines problem. We then pass to complete flag varieties, where rank tables, Bruhat order, and Schubert polynomials lead to Monk's rule. The conclusion returns to the opening question with the whole dictionary in hand.

2 Grassmannians

A Grassmannian is a space whose points are subspaces. For the four-lines problem and other similar examples, it helps to work with projective lines. Projective space itself is a way to forget the length of a vector and remember only its direction. It is obtained from $\mathbb{C}^n \setminus \{0\}$ by identifying nonzero scalar multiples:

$$[x_1 : \cdots : x_n] = [\lambda x_1 : \cdots : \lambda x_n] \qquad (\lambda \neq 0).$$

Thus $(2, 0, 1, 0)$, $(4, 0, 2, 0)$, and $(-2, 0, -1, 0)$ all determine the same point $[2 : 0 : 1 : 0]$ in $\mathbb{P}^3$. A point of projective space is a whole line through the origin, collapsed to one object.

The same idea explains why the Grassmannian appears in the problem about lines. Just as a projective point corresponds to a line through the origin, a projective line in $\mathbb{P}^3$ corresponds to a two-dimensional vector plane through the origin in $\mathbb{C}^4$. Start with the coordinate plane

$$\{(a, b, 0, 0) : a, b \in \mathbb{C}\} \subset \mathbb{C}^4.$$

This plane is spanned by $(1,0,0,0)$ and $(0,1,0,0)$, because every vector in it is built from those two directions. Passing to projective space turns a nonzero vector in this plane into

$$[a : b : 0 : 0].$$

Multiplying both a and b by the same nonzero number gives the same projective point, so only the direction inside the coordinate plane matters. For instance, when $a \neq 0$,

$$[a : b : 0 : 0] = [1 : b/a : 0 : 0].$$

Thus this whole part is traced by the single number b/a. The one remaining case is $a = 0$, which gives the point $[0 : 1 : 0 : 0]$. Together these are the projective points coming from directions in the coordinate plane, and they form one projective line in $\mathbb{P}^3$. Conversely, every projective line in $\mathbb{P}^3$ comes from exactly one two-dimensional subspace of $\mathbb{C}^4$.

Definition 2.1. The Grassmannian $\mathrm{Gr}(k,n)$ is the set of all k-dimensional subspaces of $\mathbb{C}^n$.

Thus

$$\mathrm{Gr}(2,4) \leftrightarrow \{\text{two-dimensional subspaces of } \mathbb{C}^4\} \leftrightarrow \{\text{projective lines in } \mathbb{P}^3\}.$$

More generally, $\mathrm{Gr}(k,n)$ parametrizes projective $(k-1)$-planes in $\mathbb{P}^{n-1}$.

To analyze a point in $\mathrm{Gr}(2,4)$, start with two independent vectors spanning a two-dimensional subspace of $\mathbb{C}^4$ and place them as the rows of a matrix. Row operations change the chosen basis, but not the plane it spans, so we put the matrix in a right-to-left echelon form. In this form, the leading 1's move from right to left as the rows go down. Each leading 1 has zeros in the rest of its column, and all entries to its left in the same row are zero. The entries to its right are flexible complex numbers.

$$\begin{bmatrix} 1 & 3 & 1 & 6 \\ 1 & 3 & 2 & 8 \end{bmatrix} \longrightarrow \begin{bmatrix} 0 & 0 & 1 & 2 \\ 1 & 3 & 0 & 4 \end{bmatrix} \qquad \begin{bmatrix} 1 & 1 & 6 & 8 \\ 1 & 2 & 8 & 11 \end{bmatrix} \longrightarrow \begin{bmatrix} 0 & 1 & 2 & 3 \\ 1 & 0 & 4 & 5 \end{bmatrix}$$

Figure 1: Two basis matrices reduced to right-to-left echelon form.

The four columns are labeled e_1, e_2, e_3, e_4 from left to right. In the first example, the leading 1's are in the e_3 and e_1 columns. In the second example, they are in the e_2 and e_1 columns. In general, each point of the Grassmannian is the row span of a unique full-rank $k \times n$ matrix in this right-to-left echelon form.

3 Flags and Partitions

A flag is a reference object that lets us "measure" a subspace. It is a sequence of reference subspaces that grows one dimension at a time.

Definition 3.1. A flag in $\mathbb{C}^n$ is a chain of subspaces

$$0 = F_0 \subset F_1 \subset F_2 \subset \cdots \subset F_n = \mathbb{C}^n, \qquad \dim F_q = q.$$

In this paper, the coordinate flag is the one obtained by adding the coordinate directions from right to left:

$$F_q = \mathrm{span}(e_n, e_{n-1}, \ldots, e_{n-q+1}).$$

Equivalently, $F_1 = \mathrm{span}(e_n)$, then $F_2 = \mathrm{span}(e_n, e_{n-1})$, and so on.

In $\mathbb{C}^4$, the coordinate flag looks like

$$F_1 = (0,0,0,*), \qquad F_2 = (0,0,*,*), \qquad F_3 = (0,*,*,*), \qquad F_4 = \mathbb{C}^4.$$

Here $*$ denotes an arbitrary complex number. As the flag grows, a general plane becomes more "visible" in the flag over time. For example, take the plane from the left echelon matrix in Figure 1:

$$V = \operatorname{span}\big((0,0,1,2),(1,3,0,4)\big).$$

A vector in V has the form

$$a(0,0,1,2) + b(1,3,0,4) = (b,\ 3b,\ a,\ 2a+4b).$$

We focus on the intersection of V with the coordinate flag. Inside F_1, the first three coordinates must vanish, so $a = b = 0$. Thus there is no nonzero vector of V inside F_1. Inside F_2, the first two coordinates must vanish, so $b = 0$, leaving one visible direction $(0,0,1,2)$.

Inside F_3, the first coordinate still forces $b = 0$, so no new direction appears. Finally, inside $F_4 = \mathbb{C}^4$, the whole plane is visible. The dimension of the intersection is better seen as a staircase:

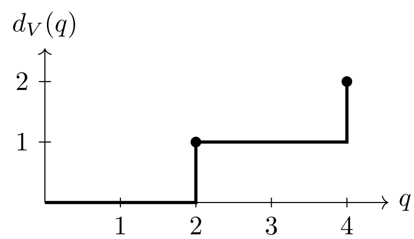

Figure 2: A visibility staircase in $\mathrm{Gr}(2,4)$, with jumps at 2 and 4.

The leading 1's are in the e_3 and e_1 columns. Since the flag reveals the coordinate directions in the order e_4, e_3, e_2, e_1, these become jumps at $q = 2$ and $q = 4$.

For a general $V \in \mathrm{Gr}(k,n)$, the same measurement is written

$$d_V(q) = \dim(V \cap F_q).$$

It counts how many independent directions of V have become visible by the time the flag reaches F_q.

Lemma 3.2. *The sequence $d_V(0), d_V(1), \ldots, d_V(n)$ starts at 0, ends at k, and each step changes by either 0 or 1.*

Proof. Let $W_q = V \cap F_q$ be the intersection. Since $F_q \subset F_{q+1}$, we have $W_q \subset W_{q+1}$, so the dimensions cannot decrease. The only question is how much larger W_{q+1} can be.

Both W_{q+1} and F_q are subspaces of F_{q+1}. Their intersection is

$$W_{q+1} \cap F_q = (V \cap F_{q+1}) \cap F_q = V \cap F_q = W_q.$$

Using the dimension formula for a sum of subspaces,

$$\dim(W_{q+1} + F_q) = \dim W_{q+1} + \dim F_q - \dim W_q.$$

But $W_{q+1} + F_q$ is still contained in F_{q+1}, whose dimension is $q+1$, while $\dim F_q = q$. Therefore

$$\dim W_{q+1} + q - \dim W_q \le q + 1.$$

Rearranging gives

$$0 \leq \dim W_{q+1} - \dim W_q \leq 1.$$

Thus each step changes by 0 or 1. Finally, $W_0 = 0$ and $W_n = V$, so the sequence starts at 0 and ends at k. □

We can now classify the possible staircases. Let

$$q_1 < q_2 < \cdots < q_k$$

be the stages at which the k jumps occur. The ith jump cannot occur before stage i, because each stage reveals at most one new direction. It also cannot occur after stage $n - k + i$, because the remaining $k - i$ jumps still need room to occur. Thus

$$i \leq q_i \leq n - k + i.$$

The upper bound is the latest possible position of the ith jump. We record how far the jump has moved from that latest position by setting

$$\lambda_i = n - k + i - q_i.$$

These numbers automatically satisfy

$$0 \leq \lambda_i \leq n - k, \qquad \lambda_i - \lambda_{i+1} = q_{i+1} - q_i - 1 \geq 0 \quad (1 \leq i < k).$$

The second inequality is the algebraic form of the no-collision rule: two jumps cannot occur at the same stage.

Definition 3.3. A *partition* $\lambda = (\lambda_1, \ldots, \lambda_k)$ lists the row lengths of a left-justified Young diagram. It fits inside the $k \times (n - k)$ rectangle if

$$n - k \geq \lambda_1 \geq \cdots \geq \lambda_k \geq 0.$$

The corresponding jump positions are

$$q_i(\lambda) = n - k + i - \lambda_i.$$

Thus partitions in the rectangle and possible visibility staircases are two descriptions of exactly the same data. The number λ_i says how many stages early the ith direction becomes visible.

The same dictionary is visible in an echelon matrix. Begin with the right-to-left row-reduced matrix and shade the triangular collection of entries forced by the latest possible pivots. After removing this common triangle and the zeros in the other pivot columns, the remaining rectangle has k rows and $n - k$ columns. A box records one step by which a pivot, and therefore a jump, has moved early. The unfilled boxes below and to the right of the boundary path are the free parameters of the corresponding cell [3, Chs. 9–10].

For example, in $\mathrm{Gr}(2, 4)$ the latest jumps are $(3, 4)$. Moving the first jump one stage earlier gives $(2, 4)$, so the diagram has one box in its first row and the echelon matrix has one fewer free parameter.

For the example in Figure 3, the jumps are $(2, 4)$, so the partition is $(1, 0)$, which we abbreviate to (1). By contrast, the impossible sequence $(0, 1)$ would ask for two jumps at $q = 3$. It is excluded both by the staircase lemma and by the requirement that partition rows weakly decrease.

For $\mathrm{Gr}(2, 4)$, the following diagrams show all possible jump positions:

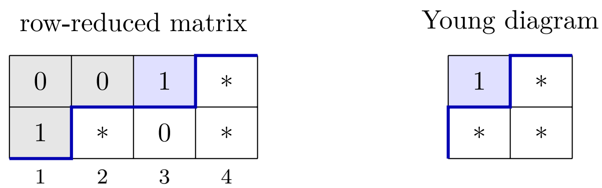

Figure 3: The echelon and Young diagrams for jumps at $q = 2$ and $q = 4$.

The empty diagram is the generic case with jumps as late as possible. The full rectangle forces both jumps as early as possible, so the plane is already fully visible inside F_2. We can characterize an element of the Grassmannian by the partition recording its jumps, together with the free entries of the corresponding echelon matrix.

4 Schubert Cells and Schubert Varieties

Partitions record visibility patterns, but we wish to turn these back into geometric conditions. What exactly does it mean for a subspace to have a specific corresponding partition? Schubert cells and Schubert varieties make that conversion.

Definition 4.1. The Schubert cell associated to λ is

$$C_\lambda = \{V \in \mathrm{Gr}(k, n) : d_V \text{ jumps exactly at } q_1(\lambda), \dots, q_k(\lambda)\}.$$

In other words, the Schubert cell contains all the subspaces in the Grassmannian with a specific corresponding partition. By Lemma 3.2, every subspace produces a staircase, and every staircase has exactly k jumps, so every V lies in exactly one Schubert cell. We can use this definition to create a looser condition.

Definition 4.2. The Schubert variety associated to λ is the union of cells

$$\Omega_\lambda = \bigcup_{\lambda' \supseteq \lambda} C_{\lambda'},$$

where $\lambda' \supseteq \lambda$ means $\lambda'_i \geq \lambda_i$ for $1 \leq i \leq k$.

Ω_λ collects the subspaces whose jumps in visibility occur no later than the corresponding jumps for the Schubert cell C_λ.

To count how large each cell is, we return to the echelon matrix and adapt it to the jumps. We place the direction that becomes visible at stage q_i in row i. This direction lies in F_{q_i}, so it is supported in the last q_i coordinates, and its leading 1 sits in column $n + 1 - q_i$. As in Section 2, we clear the rest of every pivot column, so the entries of row i that remain free are those to the right of its leading 1 and outside the other pivot columns.

Proposition 4.3 (Dimension of Schubert cells)**.** *For $\lambda \subseteq k \times (n - k)$, the free entries of the adapted echelon form identify the Schubert cell C_λ with a copy of $\mathbb{C}^{k(n-k)-|\lambda|}$. In particular,*

$$\dim(C_\lambda) = k(n - k) - |\lambda|, \qquad |\lambda| = \lambda_1 + \cdots + \lambda_k.$$

Proof. Take $q_i = q_i(\lambda)$. A plane in C_λ has exactly one echelon basis of the form just described, with the leading 1 of row i in column $n+1-q_i$. We count the free entries of row i. There are $q_i - 1$ columns to the right of its leading 1, and $i-1$ of them are pivot columns of earlier rows: for $j < i$ we have $q_j < q_i$, which places the pivot of row j in column $n+1-q_j > n+1-q_i$. Row i therefore has

$$(q_i - 1) - (i-1) = q_i - i = n - k - \lambda_i$$

free entries, and the whole cell has

$$\sum_{i=1}^{k}(n-k-\lambda_i) = k(n-k) - |\lambda|.$$

Setting $\lambda = \varnothing$ gives the open cell, with $k(n-k)$ free parameters. Thus

$$\dim \mathrm{Gr}(k,n) = k(n-k)$$

is recovered from the star-counting picture. □

We now move our attention to the Schubert variety. At first glance, it may look like nothing more than a bookkeeping tool that groups together cells. Why is this looser set the right geometric object? The answer appears when we let a plane move.

We can watch a family of planes tilt. For $t \in \mathbb{C}$, consider the plane

$$V_t = \mathrm{span}\left((0,t,1,a),\ (1,b,0,c)\right).$$

For every $t \neq 0$, the first row has its leftmost nonzero entry in column 2 (scaling by $1/t$ makes it a leading 1), so the first direction becomes visible at $q_1 = 3$. This means the plane V_t lies in $C_\varnothing$.

At $t = 0$ that entry vanishes, the leading 1 moves to column 3, and the first direction is already visible at stage 2. Thus the plane lies in $C_{(1)}$. A family can therefore sit inside one cell for every $t \neq 0$ and land in a more special cell at $t = 0$.

The reverse jump is impossible: a special fiber cannot suddenly become less visible. To see why, we turn the visibility conditions into equations. On any matrix chart for the Grassmannian, choose a basis matrix for V and stack its k rows on top of q rows spanning F_q. The row span of the stacked matrix is $V + F_q$, so the dimension formula gives

$$\mathrm{rank}(\text{stacked matrix}) = \dim(V + F_q) = k + q - \dim(V \cap F_q).$$

Thus $\dim(V \cap F_q) \geq i$ is equivalent to saying that all minors of size $k+q-i+1$ vanish. If these polynomial equations hold for every $t \neq 0$ in a family, they also hold at $t = 0$. Therefore $\dim(V_t \cap F_q)$ may jump up in a special fiber, as it did in the example above, but it cannot jump down.

In algebraic geometry, a set cut out by polynomial equations is called *closed*. The *closure* $\overline{S}$ is the smallest closed set containing S. Any limit of a one-parameter family inside S belongs to $\overline{S}$. For the Schubert cells, the explicit families below will also show that every point added by the closure occurs through exactly this kind of specialization.

Proposition 4.4 (Two descriptions of the Schubert variety)**.** *For every* $\lambda \subseteq k \times (n-k)$*,*

$$\Omega_\lambda = \left\{V \in \mathrm{Gr}(k,n) : \dim(V \cap F_{q_i(\lambda)}) \geq i \text{ for } 1 \leq i \leq k\right\} \qquad \text{and} \qquad \overline{C_\lambda} = \Omega_\lambda.$$

Proof. We first check the inequality description. Suppose $V \in C_{\lambda'}$ with $\lambda' \supseteq \lambda$. The ith jump of V occurs at $q_i(\lambda') \le q_i(\lambda)$, so at least i directions are visible by stage $q_i(\lambda)$; that is, $\dim(V \cap F_{q_i(\lambda)}) \ge i$. Conversely, suppose the inequalities hold, and let λ' be the partition of V. The ith inequality says that at least i jumps occur by stage $q_i(\lambda)$, so $q_i(\lambda') \le q_i(\lambda)$, which is $\lambda'_i \ge \lambda_i$. Thus V lies in a cell with $\lambda' \supseteq \lambda$.

For the closure statement, we already turned each visibility inequality into polynomial equations on matrix charts, so Ω_λ is closed. Since it contains C_λ, we get $\overline{C_\lambda} \subseteq \Omega_\lambda$.

For the reverse containment, it is enough to move one box at a time. Let μ be obtained from ν by adding one box in row i, and take $V \in C_\mu$. The ith jump of V occurs one stage earlier than the corresponding jump for C_ν, while all its other jumps agree. In the adapted echelon basis $v_1, \dots, v_k$ for V, replace row i by

$$v_i(t) = v_i + te_{n+1-q_i(\nu)}$$

and leave the other rows unchanged. The new coordinate lies one column to the left of the old pivot of row i. Because both μ and ν are partitions, that column is not another pivot column. For $t \neq 0$, it is therefore the new pivot of row i; clearing its entries from later rows does not change their pivots. The resulting plane has the jump positions of ν, so it lies in C_ν. At $t = 0$ it returns to V.

Every larger partition $\lambda' \supseteq \lambda$ can be reached from λ by a sequence of legal one-box additions. Repeating the preceding specialization shows that every point of $C_{\lambda'}$ lies in $\overline{C_\lambda}$. Since Ω_λ is the union of these cells, we obtain $\Omega_\lambda \subseteq \overline{C_\lambda}$. This proves the reverse containment; see Fulton [3, Ch. 9] for the general matrix description of the same argument. □

Together, the dimension and closure statements give

$$\dim \Omega_\lambda = k(n-k) - |\lambda|, \qquad \operatorname{codim} \Omega_\lambda = |\lambda|.$$

Containment of Young diagrams records the one-way motion of jumps. For example,

$$\Omega_{(1)} = C_{(1)} \sqcup C_{(2)} \sqcup C_{(1,1)} \sqcup C_{(2,1)} \sqcup C_{(2,2)}.$$

The open part $C_{(1)}$ consists of planes whose first jump is exactly at F_2. The other pieces are the more special limits where some jump has moved farther left.

We can now ask the pivotal question: what does $\Omega_{(1)}$ require of a point $V \in \mathrm{Gr}(2,4)$? Both V and F_2 are two-dimensional vector subspaces, so projectivizing turns them into lines in $\mathbb{P}^3$. Write $\ell = \mathbb{P}(V)$ for the moving line and $L = \mathbb{P}(F_2)$ for the fixed line.

The condition $V \in \Omega_{(1)}$ is

$$\dim(V \cap F_2) \ge 1.$$

In other words, $V \cap F_2$ contains a nonzero vector. Its one-dimensional span becomes a projective point belonging to both ℓ and L. Hence

$$V \in \Omega_{(1)} \iff \ell \cap L \neq \varnothing.$$

The condition "meet a fixed line" is therefore a Schubert variety.

The same condition is also a polynomial equation. We take the chart of $\mathrm{Gr}(2,4)$ with row matrix

$$\begin{pmatrix} a & b & 1 & 0 \\ c & d & 0 & 1 \end{pmatrix}.$$

Since F_2 consists of the vectors whose first two coordinates vanish, the plane meets F_2 when some nonzero combination of the rows has its first two coordinates equal to zero. We write out the

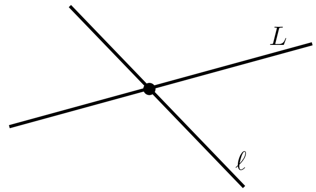

Figure 4: The moving line ℓ meets the fixed line L.

combination $x(a, b, 1, 0) + y(c, d, 0, 1)$. Its first two coordinates are $xa + yc$ and $xb + yd$, and a nonzero solution (x, y) exists when

$$\det \begin{pmatrix} a & b \\ c & d \end{pmatrix} = ad - bc = 0.$$

Thus, on this chart, "meet a fixed line" is the single determinant equation $ad - bc = 0$. The calculation is the smallest example of the minors used above to express all Schubert incidence conditions.

The same dictionary gives a concrete meaning to every shape in the 2×2 rectangle. Write

$$p = \mathbb{P}(F_1), \qquad L = \mathbb{P}(F_2), \qquad P = \mathbb{P}(F_3).$$

Then the Schubert varieties in $\mathrm{Gr}(2, 4)$ are summarized by

λ	(q_1, q_2)	condition on the projective line $\ell = \mathbb{P}(V)$
$\varnothing$	$(3, 4)$	no condition
(1)	$(2, 4)$	ℓ meets L
(2)	$(1, 4)$	ℓ passes through p
$(1, 1)$	$(2, 3)$	ℓ lies in P
$(2, 1)$	$(1, 3)$	ℓ passes through p and lies in P
$(2, 2)$	$(1, 2)$	$\ell = L$.

This table will let us read the first products geometrically before using the general multiplication rule.

5 Pieri's Rule

We can now express a single incidence condition as a Schubert variety. The four-lines problem, however, imposes four conditions at once, and each condition "meet L_i" uses a different line. In our language, each condition is a copy of $\Omega_{(1)}$ built from a different flag. So we ask: how do we intersect Schubert varieties defined by different flags, and how do we count the result?

Our first key observation is that the choice of flag should not matter. To see why, we drop one size down, to $\mathbb{C}^2$, where a flag is a single line through the origin. Take the two flags

$$\ell = \mathrm{span}(e_2), \qquad \ell' = \mathrm{span}(e_1 + e_2).$$

The invertible matrix

$$M = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

carries ℓ to ℓ', since $Me_2 = e_1 + e_2$. Now we slide M back to the identity:

$$M_t = \begin{pmatrix} 1 & 1-t \\ 0 & 1 \end{pmatrix}, \qquad M_0 = M, \quad M_1 = I.$$

The moving line $M_t\,\ell = \operatorname{span}\big((1-t,\,1)\big)$ tilts continuously from ℓ' back to ℓ. The slide never breaks down, because $\det M_t = 1$ for every t.

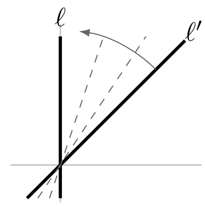

The same works in general. An invertible matrix carries any flag in $\mathbb{C}^n$ to any other flag, and the space of invertible complex matrices is connected, so a path of such matrices slides the first flag into the second. Along that path, a Schubert variety of shape λ deforms continuously through Schubert varieties of the same shape.

This observation suggests that the flag should be forgotten and that all Schubert varieties of shape λ should receive one shared name, σ_λ. The path alone, however, does not yet prove that counts stay fixed. In a special configuration, solutions may collide, acquire multiplicity, or even form a positive-dimensional family. The principle of *conservation of number* says that the total count remains stable once these effects are recorded correctly. Cohomology will provide the precise form of that principle below.

Now we impose two conditions at once: which planes in $\mathrm{Gr}(2,4)$ meet both a line L and a second line L'? We deform, as in Section 4. Slide L' until it touches L. In the limit configuration, the two lines share a point p and span a plane P. Take any line ℓ meeting both. If ℓ meets them in two distinct points, then ℓ has two points inside P, so ℓ lies inside P. Otherwise ℓ meets them in one common point, and the only common point is p.

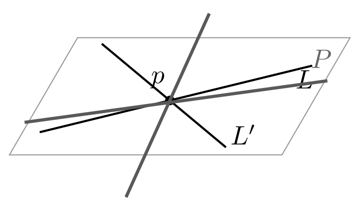

Set-theoretically, the limiting condition has split into two familiar shapes. "Pass through p" makes the direction of p visible immediately; its shape is (2). "Lie inside P" makes everything visible by stage 3; its shape is $(1,1)$. Both are Schubert conditions for a flag built from p, L, and P. This picture predicts the identity

$$\sigma_{(1)}\,\sigma_{(1)} = \sigma_{(2)} + \sigma_{(1,1)}.$$

What the picture does not prove by itself is that the two components occur with multiplicity one. That will follow from the intersection-theoretic statement below and, more specifically, from Pieri's rule.

More generally, we would like multiplication to follow the same recipe: impose Ω_λ and Ω_μ for flags in general position, slide the flags into a useful special position, and list the Schubert pieces of the limit:

$$\sigma_\lambda\,\sigma_\mu = \sum_\nu c_\nu\,\sigma_\nu,$$

where c_ν should record the multiplicity with which shape ν appears.

For this bookkeeping to make sense, two things must be true. The list of pieces must not depend on the choices made along the way, and the names must be independent, so that the list can be read off the product. The following theorem guarantees both. Proving it requires singular cohomology, Poincaré duality, and transversality. See Fulton [3, Chs. 9–10] or Sottile [8, §§1–3].

Theorem 5.1 (Schubert classes and generic intersections; standard input)**.** *The integral cohomology of* $\mathrm{Gr}(k,n)$ *has an additive basis*

$$\sigma_\lambda = \mathrm{PD}([\Omega_\lambda]_{\mathrm{fund}}) \in H^{2|\lambda|}(\mathrm{Gr}(k,n),\mathbb{Z}), \qquad \lambda \subseteq k \times (n-k).$$

If a collection of Schubert varieties is moved into general relative position, its intersection has the cohomology class given by the cup product of their classes. In complementary total codimension, the coefficient of the point class counts the resulting finite intersection, with multiplicities.

In other words, the theorem certifies our bookkeeping. The fundamental class $[\Omega_\lambda]_{\mathrm{fund}}$ records the subvariety itself in homology, and Poincaré duality, denoted PD, turns it into the cohomology class σ_λ of the condition it imposes. An *additive basis* means that every class is a unique integer combination of these names, while the *cup product* is the product of conditions. We will not need a construction of cohomology; only three of its consequences matter here: classes add when components appear with multiplicity, multiply when conditions are imposed in general position, and count points in top degree. The group $H^{2|\lambda|}$ is the formal home for a codimension-$|\lambda|$ condition. The factor of 2 appears because topology counts real parameters, and one complex parameter is two real parameters.

Counting is now a special case of decomposing. Suppose generally positioned conditions have total codimension $k(n-k)$, the whole dimension of the Grassmannian. Then their joint condition has dimension zero, so it is a finite set of planes. Only one shape has dimension zero: the full $k \times (n-k)$ rectangle, which by Proposition 4.3 leaves no free parameters. The decomposition must therefore read $N \cdot \sigma_{\text{full rectangle}}$, and N is the number of solutions. This is the theorem's last sentence. "With multiplicities" allows a degenerate configuration to count a solution more than once. For flags in general position, transversality makes each solution count once.

We can now ask the section's real question. Suppose a plane already satisfies the condition Ω_λ, and we impose one more "meet" condition. Which shapes appear in the decomposition? The extra condition is the one-box condition $\sigma_{(1)}$. In $\mathrm{Gr}(k,n)$, its variety $\Omega_{(1)}$ asks that V meet the subspace F_{n-k}, exactly as "meet a fixed line" did in $\mathrm{Gr}(2,4)$.

Theorem 5.2 (Pieri's rule, one-box case; standard input)**.**

$$\sigma_\lambda\, \sigma_{(1)} = \sum_\mu \sigma_\mu,$$

where the sum is over the partitions $\mu \subseteq k \times (n-k)$ *obtained from* λ *by adding one box.*

This is the one-box case of the classical Pieri rule; see Fulton [3, Ch. 5] or Sottile [8, §3].

The rule says exactly what one would hope: multiplying by the one-box condition adds one box in every legal place. In $\mathrm{Gr}(2,5)$ the rectangle is 2×3, and from the shape $(2,1)$ there are only two possible one-box moves.

Thus

$$\sigma_{(2,1)}\sigma_{(1)} = \sigma_{(3,1)} + \sigma_{(2,2)}.$$

The equation says that after imposing one more simple incidence condition, the old visibility pattern can become more special in exactly the two ways shown in Figure 5.

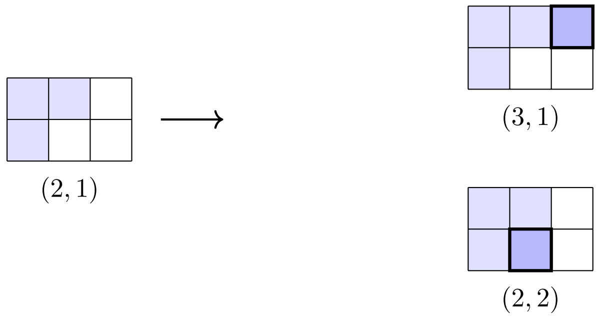

Figure 5: Multiplication by $\sigma_{(1)}$ adds one visible box.

The staircase picture explains why these are the natural candidate terms. One more incidence condition has codimension one, so the visibility pattern must become more special by one step: one jump arrives one stage early. The legal positions are exactly those where the moved jump does not collide with its neighbor, the same constraint that made partitions weakly decreasing in Section 3. What Pieri's rule supplies beyond this local picture is the global guarantee: every legal shape appears, no other shape appears, and every coefficient is one.

Remark 5.3. The general Pieri rule describes multiplication by $\sigma_{(r)}$, the name of the one-row shape with r boxes: the sum runs over the ways to add r boxes to λ with no two in the same column. We only need the one-box case; for the general rule see Fulton [3, Ch. 5] or Sottile [8, §3].

We can now count the four-lines problem, one condition at a time. Each "meet L_i" contributes a factor of $\sigma_{(1)}$, so we compute $\sigma_{(1)}^4$ in $\mathrm{Gr}(2, 4)$, where the rectangle is 2×2.

The first condition gives $\sigma_{(1)}$ itself. The second condition adds a box to (1) in the two legal ways, reproducing the two-lines decomposition we found geometrically:

$$\sigma_{(1)}^2 = \sigma_{(2)} + \sigma_{(1,1)}.$$

The third condition adds a box to each term. Inside the 2×2 rectangle, both (2) and $(1, 1)$ admit only the move to $(2, 1)$, so the two branches merge:

$$\sigma_{(1)}^3 = 2\,\sigma_{(2,1)}.$$

The fourth condition completes the rectangle:

$$\sigma_{(1)}^4 = 2\,\sigma_{(2,2)}.$$

Figure 6 shows the whole computation at once.

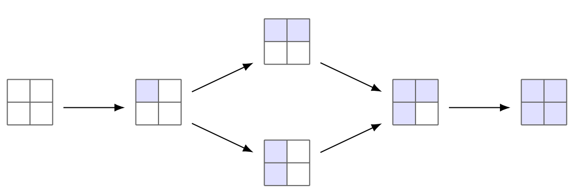

Figure 6: The four box additions of $\sigma_{(1)}^4$ in $\mathrm{Gr}(2, 4)$. Two monotone paths lead from the empty diagram to the full rectangle, matching the coefficient 2.

Finally, we identify $\sigma_{(2,2)}$. The full rectangle forces every direction to be visible as early as possible, and by Proposition 4.3 its cell has dimension zero: the only plane in it is F_2 itself. Thus $\sigma_{(2,2)}$ is the class of a single point, and

$$\sigma_{(1)}^4 = 2\,\sigma_{(2,2)}$$

says that four general lines in $\mathbb{P}^3$ have exactly two common transversals. The coefficient 2 counts the two chains of box additions in Figure 6. This settles the count promised in the introduction.

6 Flag Varieties and Rank Tables

Pieri's rule tells us what happens when one more condition is imposed on a single moving subspace. Monk's rule, the second rule of the introduction, speaks about permutations instead of partitions. Where do the permutations come from? They appear when the moving object is upgraded from one subspace to a whole chain of subspaces, measured against the fixed flag all at once.

Definition 6.1. The complete flag variety Fl_n is the space of all chains

$$V_1 \subset V_2 \subset \cdots \subset V_n = \mathbb{C}^n, \qquad \dim V_p = p.$$

In other words, a point of Fl_n is itself a flag: it chooses a line inside a plane inside a three-dimensional space, and so on. To keep the two flags apart, we call $V_\bullet$ the moving flag and $F_\bullet$ the fixed one.

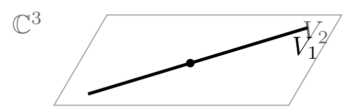

Before measuring, it is useful to know the size of the space.

Proposition 6.2. *The complete flag variety* Fl_n *has complex dimension*

$$\dim \mathrm{Fl}_n = \binom{n}{2}.$$

Proof. Project a flag to its first member V_1. The possible lines V_1 form $\mathrm{Gr}(1,n) = \mathbb{P}^{n-1}$, a space of dimension $n-1$. Once V_1 is fixed, the remaining chain

$$V_2/V_1 \subset \cdots \subset V_n/V_1 = \mathbb{C}^n/V_1$$

is a complete flag in the $(n-1)$-dimensional quotient. Thus the fibers of this projection are copies of Fl_{n-1}. Locally the projection looks like a product, so the dimensions of the base and fiber add. If $d_n = \dim \mathrm{Fl}_n$, then

$$d_n = (n-1) + d_{n-1}, \qquad d_1 = 0.$$

So $d_n = (n-1) + (n-2) + \cdots + 1 = \binom{n}{2}$. ☐

Intuitively, when the pth new direction is chosen, there are $n-p$ remaining directions into which it can "tilt".

For one moving subspace we recorded $\dim(V \cap F_q)$. For a moving flag we record every comparison at once:

$$\dim(V_p \cap F_q), \qquad 1 \le p, q \le n.$$

For example, in $\mathbb{C}^3$, a fixed flag is $F_1 = \operatorname{span}(e_3) \subset F_2 = \operatorname{span}(e_3, e_2) \subset F_3 = \mathbb{C}^3$. Take the moving flag

$$V_1 = \operatorname{span}(e_2 + ae_3), \qquad V_2 = \operatorname{span}(e_2 + ae_3,\ e_1 + be_3), \qquad V_3 = \mathbb{C}^3,$$

with $a, b \in \mathbb{C}$. A vector of V_1 never lies in F_1, and it always lies in F_2. So the first row of measurements is $0, 1, 1$. A vector of V_2 has the form $y\, e_1 + x\, e_2 + (ax + by)e_3$, so it lies in F_1 only when $x = y = 0$, and it lies in F_2 only when $y = 0$, which leaves the line V_1. So the second row is $0, 1, 2$. The third row is $1, 2, 3$, since V_3 is everything. Arranged with rows indexed by p and columns by q, the measurements form the rank table

$$\begin{pmatrix} 0 & 1 & 1 \\ 0 & 1 & 2 \\ 1 & 2 & 3 \end{pmatrix}.$$

Notice that the parameters a and b do not affect the table. In particular, setting $a = b = 0$ gives the coordinate flag $\operatorname{span}(e_2) \subset \operatorname{span}(e_2, e_1) \subset \mathbb{C}^3$ with the same measurements. Flags built from coordinate vectors are the models for all possible tables, and the order in which the coordinate vectors appear is a permutation. We write permutations in one-line notation, listing the values $w(1)\, w(2) \cdots w(n)$: the permutations of $\{1, \dots, n\}$ form the symmetric group S_n, and $w = 231$ means $w(1) = 2$, $w(2) = 3$, $w(3) = 1$. For $w \in S_n$, define the coordinate flag

$$E_w : \qquad \operatorname{span}(e_{n+1-w(1)}) \subset \operatorname{span}(e_{n+1-w(1)}, e_{n+1-w(2)}) \subset \cdots \subset \mathbb{C}^n.$$

The number $w(p)$ records the stage at which the pth direction of E_w becomes visible to the fixed flag. In the example above, e_2 is revealed at stage 2, then e_1 at stage 3, then e_3 at stage 1, so the coordinate flag is E_w for $w = 231$.

Definition 6.3. For $w \in S_n$, define its rank table by

$$r_w(p, q) = \#\{i \le p : w(i) \le q\}.$$

Equivalently, draw the permutation matrix of w. Then $r_w(p, q)$ is the number of dots in the northwest $p \times q$ rectangle.

By construction,

$$\dim\big((E_w)_p \cap F_q\big) = r_w(p, q) :$$

by stage q, the flag has revealed the directions with $w(i) \le q$, and $r_w(p, q)$ counts how many of the first p of them have appeared.

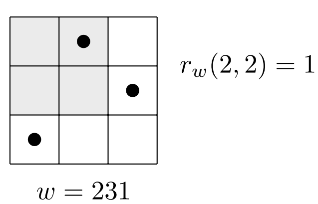

Figure 7: A rank-table entry counts dots in a northwest rectangle.

Figure 7 shows the permutation 231. The shaded rectangle has one dot, so $r_{231}(2, 2) = 1$, matching the table we computed. In other words, the rank table is the two-dimensional version of the staircase: instead of recording when one subspace becomes visible, it records how much of each V_p is visible inside each F_q.

We now repeat the pattern of Section 4: first the exact conditions, then the looser ones.

Definition 6.4. The Schubert cell of $w \in S_n$ is

$$C_w = \{V_\bullet \in \mathrm{Fl}_n : \dim(V_p \cap F_q) = r_w(p,q) \text{ for all } p, q\}\,.$$

In other words, C_w contains the moving flags whose visibility table is exactly that of the model flag E_w. Our computed example says that the flags $V_1 = \mathrm{span}(e_2 + ae_3)$ and $V_2 = \mathrm{span}(e_2 + ae_3,\ e_1 + be_3)$ all lie in C_{231}. Conversely, the adapted echelon form of any flag with this rank table has exactly these two free entries, so this family fills the cell. The parameters a, b are the two tilts that do not change the table.

Definition 6.5. The Schubert variety of $u \in S_n$ is the union of cells

$$X_u(F_\bullet) = \bigcup_v C_v, \qquad \text{over the } v \text{ with } r_v(p,q) \geq r_u(p,q) \text{ for all } p, q.$$

In other words, X_u collects the moving flags that are at least as visible as the model E_u at every stage. Once we know that every moving flag lies in some cell, the argument of Proposition 4.4 identifies X_u with the locus of inequalities $\dim(V_p \cap F_q) \geq r_u(p,q)$. That fact, together with the analogues of the two propositions we proved by hand for Grassmannians, is the standard geometric input of this section; see Fulton [3, Ch. 10]. For the combinatorics of Bruhat order, see Björner–Brenti [1, Chs. 2–3].

Theorem 6.6 (Flag Schubert cells; standard input)**.** *Every flag in Fl_n lies in exactly one cell C_w. The cell C_w is a copy of affine space $\mathbb{C}^{\ell(w)}$, where $\ell(w)$ is the number of inversions $\#\{i < j : w(i) > w(j)\}$. The closure of C_w is $X_w(F_\bullet)$, and the containments*

$$X_u(F_\bullet) \subseteq X_v(F_\bullet) \quad \iff \quad u \leq v$$

define a partial order on S_n, the Bruhat order.

The geometric reason for the order is the same as before: under degeneration, intersection dimensions can increase, but they cannot decrease. In Fl_3, the six cell dimensions are

w	123	213	132	231	312	321
$\ell(w)$	0	1	1	2	2	3

Our example matches: $\ell(231) = 2$, one inversion for each free parameter of C_{231}.

For $n = 3$, the whole order is small enough to see at once. We display each rank table as a 3×3 matrix, with rows indexed by p and columns by q:

$$r_{123} = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 2 & 2 \\ 1 & 2 & 3 \end{pmatrix} \qquad r_{213} = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 2 & 2 \\ 1 & 2 & 3 \end{pmatrix} \qquad r_{132} = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 2 \\ 1 & 2 & 3 \end{pmatrix}$$
$$r_{231} = \begin{pmatrix} 0 & 1 & 1 \\ 0 & 1 & 2 \\ 1 & 2 & 3 \end{pmatrix} \qquad r_{312} = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 1 & 2 \\ 1 & 2 & 3 \end{pmatrix} \qquad r_{321} = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 2 \\ 1 & 2 & 3 \end{pmatrix}.$$

We read Figure 8 from bottom to top: one edge means one minimal one-step specialization, called a *cover*. Every edge joins some w to wt_{ab}, where t_{ab} swaps the entries of w in positions a and b, and raises the inversion count by exactly one. Covers have a useful elementary test.

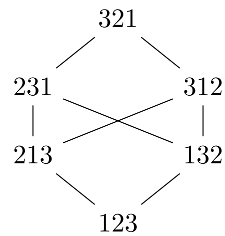

Figure 8: The Bruhat order on S_3.

Lemma 6.7 (Transposition cover criterion)**.** *Let $a < b$. Then*

$$\ell(wt_{ab}) = \ell(w) + 1$$

if and only if $w(a) < w(b)$ and there is no c with $a < c < b$ satisfying

$$w(a) < w(c) < w(b).$$

Proof. Set $x = w(a)$ and $y = w(b)$. If $x > y$, swapping the entries removes the inversion (a, b) rather than adding one, so the length cannot increase by one. Suppose $x < y$. The pair (a, b) contributes one new inversion after the swap. For each $a < c < b$, compare the two pairs involving c. If $x < w(c) < y$, both pairs become inversions after the swap whereas neither was an inversion before; this contributes two more. If $w(c)$ lies outside the open interval (x, y), the total contribution of those two pairs is unchanged. Hence

$$\ell(wt_{ab}) - \ell(w) = 1 + 2\#\{c : a < c < b,\ x < w(c) < y\}.$$

The stated criterion follows. $\square$

For $w = 132$, both swaps involving the first position pass the test:

$$132\, t_{12} = 312, \qquad 132\, t_{13} = 231.$$

Both results sit one level above 132 in Figure 8, as the lemma predicts.

7 Schubert Polynomials

Each Grassmannian condition had an algebraic name σ_λ. What are the names for the flag conditions X_u? Two bookkeeping choices come first.

The first choice is indexing by codimension. A name should record how many constraints its condition imposes, and constraints are counted by codimension, not dimension. The variety $X_u(F_\bullet)$ has *dimension* $\ell(u)$. To flip to codimension, let $w_0 = n\, n{-}1 \cdots 1$ be the permutation that reverses the order. Here $w_0 w$ means the composition $w_0 \circ w$. Composing on the left with w_0 replaces each value $w(i)$ by $n + 1 - w(i)$, so a pair of positions is inverted in $w_0 w$ exactly when it is not inverted in w, and

$$\ell(w_0 w) = \tbinom{n}{2} - \ell(w).$$

The variety $X_{w_0 w}$ therefore has codimension $\ell(w)$, and we name conditions the way we named them for the Grassmannian:

$$Y_w := X_{w_0 w}(F_\bullet), \qquad \sigma_w := \mathrm{PD}([Y_w]_{\mathrm{fund}}).$$

For example, in Fl_3 take $w = 213$. Then $w_0 w = 231$, so $Y_{213} = X_{231}$ has dimension $\ell(231) = 2$ and codimension $\ell(213) = 1$: the name 213 records one imposed condition. In this convention, longer permutations name more restrictive conditions.

The second choice is a ring in which the names can be multiplied. For Grassmannians we used cohomology as a black box; for flag varieties the black box has a concrete polynomial form. The theorem below says: work with polynomials in $x_1, \ldots, x_n$, and declare two polynomials equal when their difference belongs to the ideal generated by symmetric polynomials with zero constant term.

Theorem 7.1 (Coinvariant presentation; standard input)**.** *There is a graded-ring presentation*

$$H^*(\mathrm{Fl}_n, \mathbb{Z}) \simeq \frac{\mathbb{Z}[x_1, \ldots, x_n]}{\langle f \in \mathbb{Z}[x_1, \ldots, x_n]^{S_n} : f(0, \ldots, 0) = 0 \rangle}, \qquad \deg x_i = 2.$$

The Schubert classes σ_w, $w \in S_n$, form an additive basis. They are represented in this quotient by the Schubert polynomials $\mathfrak{S}_w(x)$.

This presentation, including the identification with Schubert classes, is standard; see Macdonald [7, §§1–4] or Fulton [3, Ch. 10]. The Schubert polynomials are the distinguished representatives that keep the geometry visible after passing to the quotient. We construct them by a recursion that walks down the Bruhat order.

The recursion uses one operator and one family of permutations. Let s_i denote the adjacent transposition $t_{i\,i+1}$, so that ws_i swaps the entries of w in positions i and $i+1$; the swap creates or removes exactly one inversion, so $\ell(ws_i) = \ell(w) \pm 1$. Acting on polynomials, let s_i interchange the variables x_i and x_{i+1}, and define

$$\partial_i f = \frac{f - s_i f}{x_i - x_{i+1}}.$$

The numerator vanishes when $x_i = x_{i+1}$, so the quotient is again a polynomial, one ordinary polynomial degree lower. For example,

$$\partial_1(x_1^2) = \frac{x_1^2 - x_2^2}{x_1 - x_2} = x_1 + x_2.$$

Definition 7.2 (Schubert polynomials)**.** Start with

$$\mathfrak{S}_{w_0} = x_1^{n-1} x_2^{n-2} \cdots x_{n-1}.$$

Whenever $\ell(ws_i) < \ell(w)$, define recursively

$$\mathfrak{S}_{ws_i} = \partial_i \mathfrak{S}_w.$$

Different length-decreasing paths from w_0 down to the same permutation give the same polynomial—the operators ∂_i obey the relations that make the recursion consistent (see Macdonald [7, §2])—so $\mathfrak{S}_w$ is well defined. In particular,

$$\mathfrak{S}_{s_k} = x_1 + \cdots + x_k.$$

For $n = 3$, we can compute every Schubert polynomial by hand. The longest permutation is 321, so

$$\mathfrak{S}_{321} = x_1^2 x_2.$$

Applying divided differences along length-decreasing edges gives

w	one divided difference step	$\mathfrak{S}_w$
321	–	$x_1^2x_2$
312	$\partial_2(x_1^2x_2)$	x_1^2
231	$\partial_1(x_1^2x_2)$	x_1x_2
132	$\partial_1(x_1^2)$	x_1+x_2
213	$\partial_2(x_1x_2)$	x_1
123	$\partial_1(x_1)$	1

$$\begin{tikzcd}
& 321 \arrow[dl, "\partial_1"'] \arrow[dr, "\partial_2"] & \\
231 \arrow[d, "\partial_2"'] & & 312 \arrow[d, "\partial_1"] \\
213 \arrow[dr, "\partial_1"'] & & 132 \arrow[dl, "\partial_2"] \\
& 123 &
\end{tikzcd}$$

The ordinary polynomial degree of $\mathfrak{S}_w$ is $\ell(w)$, exactly the complex codimension of Y_w. In the cohomological grading of Theorem 7.1, each variable has degree 2, so the same class lies in degree $2\,\ell(w)$. These are two scales for the same size: complex codimension in the geometry and real degree in cohomology.

8 Monk's Rule

We can finally explain the second rule of the introduction. The simplest condition to impose on a moving flag is the one we have used all along: a chosen member V_k must meet a fixed subspace nontrivially. Its class is σ_{s_k}, represented by the polynomial

$$\mathfrak{S}_{s_k} = x_1 + \cdots + x_k.$$

Let

$$\pi_k : \mathrm{Fl}_n \longrightarrow \mathrm{Gr}(k,n), \qquad (V_1 \subset \cdots \subset V_n) \longmapsto V_k$$

be the map that keeps the kth member of the moving flag and forgets the rest. The flags whose V_k satisfies the Grassmannian condition $\Omega_{(1)}$ form the preimage

$$D_k = \pi_k^{-1}(\Omega_{(1)}) = \{V_\bullet : V_k \cap F_{n-k} \neq 0\}.$$

Like $\Omega_{(1)}$, the condition D_k has codimension one; codimension-one subvarieties are called *divisors*. Its class is σ_{s_k}, represented by $x_1 + \cdots + x_k$; we take this identification as a standard fact alongside Theorem 7.1 (see Fulton [3, Ch. 10]). Thus multiplying by σ_{s_k}—or by its representative $x_1+\cdots+x_k$ in the coinvariant presentation—means imposing one simple incidence condition on V_k.

For Fl_3, the two simple divisors are especially easy to picture:

$$D_1 = \{V_\bullet : V_1 \subset F_2\}, \qquad D_2 = \{V_\bullet : F_1 \subset V_2\}.$$

The first condition measures the line V_1; the second measures the plane V_2. In both cases, one member of the moving flag becomes visible one step earlier than it generically would.

Write t_{ab} for the transposition that swaps *positions* a and b in one-line notation; thus wt_{ab} is obtained from w by swapping its entries in positions a and b.

Theorem 8.1 (Monk's rule in cohomology)**.** *For $w \in S_n$ and $1 \le k < n$,*

$$\sigma_w \sigma_{s_k} = \sum_{\substack{1 \le a \le k < b \le n \\ \ell(wt_{ab}) = \ell(w)+1}} \sigma_{wt_{ab}} \quad \textit{in } H^*(\mathrm{Fl}_n, \mathbb{Z}).$$

The Schubert polynomials are unchanged when a permutation is enlarged by appending fixed points. They therefore form a stable family indexed by S_∞, the group of permutations of the positive integers that move only finitely many entries; see Macdonald [7, §4]. In this setting Monk's rule also has a literal polynomial version. For $w \in S_\infty$, Monk's rule says

$$\mathfrak{S}_w(x)(x_1 + \cdots + x_k) = \sum_{\substack{1 \le a \le k < b \\ \ell(wt_{ab}) = \ell(w)+1}} \mathfrak{S}_{wt_{ab}}(x),$$

where b is not bounded by a fixed n. Only finitely many transpositions pass the cover test, so the sum is still finite; see Lascoux–Schützenberger [6] or Macdonald [7, §4] for the polynomial rule. The distinction matters at the edge of S_n. For example, in S_2 the class σ_{21}^2 is zero because it lies above the top cohomological degree, whereas the polynomial product is

$$\mathfrak{S}_{21}^2 = x_1^2 = \mathfrak{S}_{312}.$$

The permutation 312 lies in S_3, not S_2. Thus the finite formula is an equality of cohomology classes, or equivalently of polynomials modulo the coinvariant ideal, while the stable formula is an equality of the polynomials themselves. In the examples below no extra stable terms occur, so the two formulas can be read term by term.

There are two filters in both versions of the sum. The length condition says that wt_{ab} is a Bruhat cover of w. The inequality $a \le k < b$ says that the swap crosses the position selected by the divisor D_k. A single incidence condition does not see every Bruhat cover; it sees only the covers crossing its cut.

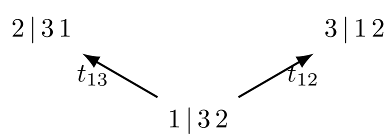

Figure 9: For $w = 132$ and $k = 1$, the bar marks the cut that the swap must cross.

Why are the crossing covers exactly the terms of the product? The geometric input is a degeneration. We define the divisor with a second flag in general position, slide that flag into the fixed one along a one-parameter family, and watch what the intersection becomes. The following statement is the one intersection-theoretic degeneration fact that we use. It is often called the geometric Chevalley formula or the moving-flag proof of Monk's rule; see Fulton [3, Ch. 10] or Sottile [8, §4].

Theorem 8.2 (Moving-divisor degeneration; standard input)**.** *Let $D_k(E_\bullet)$ be the divisor*

$$\{V_\bullet : V_k \cap E_{n-k} \neq 0\}$$

defined using a flag $E_\bullet$ in general position relative to $F_\bullet$. As $E_\bullet$ specializes to $F_\bullet$ along a standard one-parameter family, the intersection of Y_w with $D_k(E_\bullet)$ has a flat limit equal to the reduced union

$$\bigcup_{\substack{1\leq a\leq k<b\leq n\\ \ell(wt_{ab})=\ell(w)+1}} Y_{wt_{ab}}.$$

Every displayed component occurs with multiplicity one.

Three phrases in the statement deserve plain words. A *one-parameter family* means that the flag $E_\bullet$ depends on a parameter t and reaches $F_\bullet$ at $t = 0$, just like the families of Section 4. The *flat limit* is obtained by following the intersections for $t \neq 0$ and taking their limiting position in a way that preserves dimension and class. It need not equal the naive set-theoretic intersection $Y_w \cap D_k(F_\bullet)$, which may acquire extra pieces in the special position. Finally, *reduced*, with *multiplicity one*, means that each listed component appears exactly once, with no hidden doubling. In other words, the correctly measured limit breaks cleanly into the displayed Schubert varieties.

The rank conditions explain why the listed components are the only candidates. Imposing D_k changes visibility at the kth stage of the moving flag, so the relevant one-step changes must cross position k. The degeneration theorem supplies the two global facts that are not visible from rank tables alone: no other components appear, and each allowed component appears once.

Derivation of Monk's rule from the degeneration. In plain terms: the product is represented by an intersection with a generically placed divisor, the degeneration slides that divisor into special position without changing its class, and the theorem identifies the result. The class of $D_k(E_\bullet)$ is independent of the general flag $E_\bullet$ and equals σ_{s_k}. By the generic-intersection principle, the product $\sigma_w\sigma_{s_k}$ represents the general intersection. Its class does not change along the flat family, and Theorem 8.2 identifies the limit as the union of the displayed Schubert varieties, each appearing once. Adding up their names gives the cohomology formula. The stable polynomial identity is the algebraic version of the same rule, with the finite-versus-stable distinction explained above. □

The S_3 case in Figure 9 now has both a geometric and an algebraic reading. The allowed transpositions are exactly the swaps across the displayed bar:

$$132\,t_{12} = 312, \qquad 132\,t_{13} = 231.$$

Monk's rule predicts

$$\sigma_{132}\sigma_{213} = \sigma_{312} + \sigma_{231}.$$

The polynomial calculation makes the same two terms appear:

$$\underbrace{(x_1 + x_2)}_{\mathfrak{S}_{132}}\underbrace{x_1}_{\mathfrak{S}_{213}} = x_1^2 + x_1x_2 = \underbrace{\mathfrak{S}_{312}}_{x_1^2} + \underbrace{\mathfrak{S}_{231}}_{x_1x_2}.$$

In Figure 8, 132 has two covers, but the important point is why they appear in this product: both covers cross position 1, the position selected by the simple incidence divisor D_1.

The first multiplication step in the four-lines problem fits this frame. In Fl_4 with $k = 2$, the divisor D_2 is the flag version of "meet a fixed line," and its class is $\sigma_{1324} = \sigma_{s_2}$, with $\mathfrak{S}_{1324} = x_1 + x_2$.

Imposing the condition twice means squaring this class. Monk's rule allows exactly the two swaps across position 2 that raise the length by one:

$$1324\, t_{13} = 2314, \qquad 1324\, t_{24} = 1423,$$

so $\sigma_{1324}^2 = \sigma_{2314} + \sigma_{1423}$. A divided-difference computation gives $\mathfrak{S}_{2314} = x_1 x_2$ and $\mathfrak{S}_{1423} = x_1^2 + x_1 x_2 + x_2^2$, and indeed

$$(x_1 + x_2)^2 = x_1 x_2 + (x_1^2 + x_1 x_2 + x_2^2).$$

The permutations 2314 and 1423 are the flag labels of the partitions $(1,1)$ and (2). Thus this is exactly the identity $\sigma_{(1)}^2 = \sigma_{(1,1)} + \sigma_{(2)}$ from Section 5, read at the finer resolution of the flag variety.

We now make this dictionary general. Let $\lambda = (\lambda_1, \ldots, \lambda_k)$ fit inside the $k \times (n-k)$ rectangle. Define $w_\lambda \in S_n$ by prescribing its first k entries:

$$w_\lambda(i) = i + \lambda_{k+1-i} \qquad (1 \le i \le k),$$

and then placing the unused values in increasing order. The first k entries increase because the rows of λ weakly decrease, and they are at most n because the diagram fits in the rectangle. The remaining entries also increase, so the only possible descent of w_λ is at position k. Such a permutation is called *k-Grassmannian*.

The boxes of λ are already visible in its inversions. Among the values smaller than $w_\lambda(i)$, exactly $i-1$ occur before position i. Hence $w_\lambda(i) - i = \lambda_{k+1-i}$ smaller values occur after it. Summing over the first block gives

$$\ell(w_\lambda) = \sum_{i=1}^{k} \lambda_{k+1-i} = |\lambda|.$$

Thus the boxes of the partition and the inversions of its permutation measure the same codimension.

The projection π_k makes the geometric connection equally precise. To see this directly, set $u = w_0 w_\lambda$. The first k entries of u are

$$q_k(\lambda), q_{k-1}(\lambda), \ldots, q_1(\lambda).$$

It follows that $r_u(k, q_i(\lambda)) = i$, so the kth row of the rank conditions for $X_u = Y_{w_\lambda}$ is exactly $\dim(V_k \cap F_{q_i(\lambda)}) \ge i$. The other rows add no new condition on V_k. If $p < k$, passing from V_k to its subspace V_p can remove at most $k-p$ visible directions; if $p > k$, the inclusion $V_k \subset V_p$ and the dimension formula give the remaining automatic lower bounds. Thus

$$\pi_k^{-1}(\Omega_\lambda) = Y_{w_\lambda}, \qquad \pi_k^*(\sigma_\lambda) = \sigma_{w_\lambda}.$$

Completing a k-plane to a full flag adds finer measurements, but it does not alter the Schubert condition already imposed on that plane. The pullback π_k^* is injective—a standard consequence of the flag-bundle description of π_k—so no Grassmannian class is lost by making this completion; see Fulton [3, Ch. 10].

Finally, apply the transposition cover criterion to w_λ. Suppose a box can be added in row r, and put $a = k + 1 - r$. Adding that box increases $w_\lambda(a)$ by one. Because the box is legal, the next value $w_\lambda(a) + 1$ does not occur among the first k entries and does not exceed n; it therefore occurs at a unique position $b > k$. Swapping positions a and b produces w_μ, where μ is the enlarged partition. The two swapped values are consecutive, so no intermediate value can block the swap, and Lemma 6.7 says that it is a Bruhat cover.

Conversely, suppose a cover of w_λ crosses position k. If the swapped values were $x < y$ with $y > x + 1$, then the value $x + 1$ would occur between their positions: either later in the first increasing block or earlier in the second increasing block. The cover criterion rules this out. Thus $y = x + 1$, and the swap increases exactly one part of λ by one. It is therefore a legal one-box addition. Consequently,

$$\sigma_{w_\lambda}\sigma_{s_k} = \sum_\mu \sigma_{w_\mu},$$

where μ runs over the legal one-box additions to λ. Under π_k^*, this is precisely Pieri's formula

$$\sigma_\lambda\sigma_{(1)} = \sum_\mu \sigma_\mu.$$

Pieri and Monk are therefore not merely similar recipes. The one-box Pieri rule is the part of Monk's rule seen after forgetting every member of a flag except V_k. The swap must cross position k because that is the stage watched by the divisor D_k, and its length must rise by one because a single incidence condition has codimension one. In the Grassmannian, this controlled change is drawn as one new box; in the full rank table, it is recorded as one crossing Bruhat cover.

9 Conclusion

The objects in this paper can look unrelated on first encounter. A partition is a Young diagram, a permutation is a string of numbers, a Schubert polynomial is an algebraic expression, and Bruhat order can appear to be a purely combinatorial relation. Their common origin is geometric. A flag measures a subspace by revealing the ambient space one dimension at a time. For a Grassmannian this creates a staircase, and its early jumps are recorded by a partition; for a complete flag variety it creates a rank table, recorded by a permutation. Allowing a measurement to become more special gives closures and Bruhat order, and Schubert classes encode the resulting conditions algebraically.

Seen this way, the question of the introduction—why these rules rather than others?—has a short answer: the two multiplication rules are the same event at two resolutions. In the Grassmannian, one extra incidence condition adds a box to a partition; iterating this is what counted the two transversals to four general lines in $\mathbb{P}^3$. In the flag variety, the richer rank table resolves the same event into Bruhat covers crossing a chosen position, and that is Monk's rule. The projection $\pi_k : \mathrm{Fl}_n \to \mathrm{Gr}(k, n)$ makes the connection literal: the partition λ becomes the k-Grassmannian permutation w_λ, and its addable boxes become the covers crossing position k. In both cases the algebra is remembering a single geometric fact: how subspaces can become visible. "Add a box" says that one direction has become visible a step early; "swap across the mark" says that a one-step specialization has crossed the watched position. The recipes were never arbitrary.

Monk's rule is also a doorway rather than an endpoint. Multiplying by a general Schubert class in a Grassmannian is governed by the Littlewood–Richardson rule, and Monk's rule is the simplest case of the Chevalley formula, which holds for the flag varieties of any semisimple group. The visibility metaphor developed here—flags as measuring devices, degeneration as one-way motion toward more special patterns—is a reliable guide into that larger landscape; see Fulton [3] and Sottile [8] for the next steps.

10 Acknowledgements

The author would like to thank Liam Spilker for his guidance and mentorship throughout the process of writing this paper. The author would also like to thank Simon Rubinstein-Salzedo for his lectures on mathematical writing and for creating an environment that encouraged the author to pursue mathematical exposition.

References

[1] A. Bjoŕner and F. Brenti, *Combinatorics of Coxeter Groups*, Graduate Texts in Mathematics 231, Springer, 2005.

[2] A. Edelman, T. A. Arias, and S. T. Smith, The geometry of algorithms with orthogonality constraints, *SIAM Journal on Matrix Analysis and Applications* **20** (1998), no. 2, 303–353. doi:10.1137/S0895479895290954.

[3] W. Fulton, *Young Tableaux: With Applications to Representation Theory and Geometry*, London Mathematical Society Student Texts 35, Cambridge University Press, 1997.

[4] M. Gillespie, Variations on a theme of Schubert calculus, in *Recent Trends in Algebraic Combinatorics*, Association for Women in Mathematics Series 16, Springer, 2019, 115–158. arXiv:1804.08164.

[5] R. Kötter and F. R. Kschischang, Coding for errors and erasures in random network coding, *IEEE Transactions on Information Theory* **54** (2008), no. 8, 3579–3591. doi:10.1109/TIT.2008.926449.

[6] A. Lascoux and M.-P. Schützenberger, Polynômes de Schubert, *C. R. Acad. Sci. Paris Sér. I Math.* **294** (1982), 447–450.

[7] I. G. Macdonald, *Notes on Schubert Polynomials*, Publications du LACIM 6, Université du Québec à Montréal, 1991.

[8] F. Sottile, Elementary Schubert calculus, arXiv:math/0102047, 2001.