Splitting a Polynomial into Linear Factors after an Injective Ring Extension

GPT-5.5, with prompting by Darij Grinberg

July 17, 2026

Abstract. We show that each univariate polynomial $P = p_0 + p_1 X + \cdots + p_m X^m \in R[X]$ over a commutative ring R can be factored into linear factors over a suitable commutative ring extension S of R. The proof proceeds by universal construction; the injectivity of the structure homomorphism $R \to S$ is shown using a combinatorial study of multisymmetric polynomials in $\mathbb{Z}[a_1, b_1, a_2, b_2, \ldots, a_m, b_m]$. In the process, a homogeneous variant of the Garsia–Stanton basis is constructed.

Manifest. This is an answer to https://mathoverflow.net/questions/429134, generated by GPT-5.5 with significant shepherding and editing. I still intend to write up this argument in my own words and possibly extend it to an expository paper due to its many connections to algebraic combinatorics.[1]

All rings in this note are commutative and unital, and all ring homomorphisms preserve the unity.

A small notational warning: if $x_1, \ldots, x_N$ are indeterminates, then $\mathbb{Z}[x_1, \ldots, x_N]$ denotes a polynomial ring. If $u_1, \ldots, u_N$ are elements of some ring B, then $\mathbb{Z}[u_1, \ldots, u_N] \subseteq B$ denotes the subalgebra generated by these elements. Such a subalgebra need not be a polynomial ring unless this is proved.

1 The universal splitting map is split

Fix a positive integer m. Set

$$B_m = \mathbb{Z}[a_1, b_1, a_2, b_2, \ldots, a_m, b_m],$$

the polynomial ring in the $2m$ indeterminates $a_1, b_1, \ldots, a_m, b_m$ over $\mathbb{Z}$. For each subset $I \subseteq [m] := \{1, 2, \ldots, m\}$, define the monomial

$$z_I = \prod_{i \in I} a_i \prod_{i \notin I} b_i \in B_m.$$

I thank Victor Reiner and Brendon Rhoades for help with finding precedents in the literature. –Darij

For $0 \le r \le m$, define

$$E_r = \sum_{\substack{I \subseteq [m];\\ |I|=r}} z_I.$$

For example,

$$\begin{aligned}
E_0 &= z_\varnothing = b_1 b_2 \cdots b_m;\\
E_1 &= \sum_{i=1}^{m} z_{\{i\}} = a_1 b_2 b_3 \cdots b_m + b_1 a_2 b_3 \cdots b_m + b_1 b_2 a_3 \cdots b_m + \cdots + b_1 b_2 \cdots b_{m-1} a_m;\\
E_2 &= \sum_{1 \le i < j \le m} z_{\{i,j\}} = \sum_{1 \le i < j \le m} b_1 b_2 \cdots b_{i-1} a_i b_{i+1} b_{i+2} \cdots b_{j-1} a_j b_{j+1} b_{j+2} \cdots b_m;\\
&\cdots;\\
E_m &= z_{[m]} = a_1 a_2 \cdots a_m.
\end{aligned}$$

We observe the following:

Proposition 1.1. *We have*

$$\sum_{j=0}^{m} E_j X^j = \prod_{r=1}^{m} (a_r X + b_r) \qquad \text{in the univariate polynomial ring } B_m[X]. \tag{1}$$

Proof. This is a homogeneous form of Viète's theorem. The proof is easy: Multiply the right hand side and collect terms with like powers of X. □

Let

$$C_m = \mathbb{Z}[E_0, E_1, \ldots, E_m] \subseteq B_m.$$

At this point, C_m means the subalgebra of B_m generated by $E_0, E_1, \ldots, E_m$. Lemma 1.3 will show that this subalgebra is actually a polynomial ring over $\mathbb{Z}$ on these generators.

The main claim is the following.

Theorem 1.2 (splitting theorem)**.** *The inclusion*

$$C_m \hookrightarrow B_m$$

is split as a C_m-module homomorphism. Equivalently, there is a C_m-linear map

$$\rho : B_m \to C_m$$

whose restriction to C_m is the identity.

The proof of this theorem will be delivered at the end of Subsection 1.5.

1.1 Algebraic independence of the E_r's

Lemma 1.3. *The elements $E_0, E_1, \ldots, E_m$ are algebraically independent over $\mathbb{Z}$. Hence*

$$C_m = \mathbb{Z}[E_0, E_1, \ldots, E_m]$$

is a polynomial ring over $\mathbb{Z}$ on the displayed generators.

Proof. We will use the theory of monomial orders ([4, §2.2]). If a monomial order has been chosen in a polynomial ring, then $\mathrm{LT}(f)$ denotes the leading monomial of a nonzero polynomial f, not including its coefficient. For a monomial M and a variable x, we write $\nu_x(M)$ for the exponent of x in M.

Choose any monomial order on B_m. For each r, write

$$\mathrm{LT}(E_r) = z_{I_r}$$

for some r-element subset $I_r \subseteq [m]$.

We claim that

$$I_0 \subset I_1 \subset \cdots \subset I_m.$$

Indeed, suppose that $I_r \not\subseteq I_{r+1}$ for some $0 \leq r < m$. Choose

$$x \in I_r \setminus I_{r+1}, \qquad y \in I_{r+1} \setminus I_r.$$

Put

$$I_r' = (I_r \setminus \{x\}) \cup \{y\}, \qquad I_{r+1}' = (I_{r+1} \setminus \{y\}) \cup \{x\}.$$

Then

$$z_{I_r} z_{I_{r+1}} = z_{I_r'} z_{I_{r+1}'}.$$

But $z_{I_r} > z_{I_r'}$ and $z_{I_{r+1}} > z_{I_{r+1}'}$ by the definition of leading term. Multiplicativity of the monomial order gives

$$z_{I_r} z_{I_{r+1}} > z_{I_r'} z_{I_{r+1}'},$$

a contradiction to the previous equality. This proves the claim.

A chain with one r-element set for each r is obtained by adding the elements of $[m]$ in some order. Relabeling the indices $1, 2, \ldots, m$, we may therefore assume

$$I_r = \{1, 2, \ldots, r\} \qquad \text{for all } 0 \leq r \leq m.$$

Hence

$$\mathrm{LT}(E_r) = M_r := a_1 a_2 \cdots a_r b_{r+1} b_{r+2} \cdots b_m.$$

For $\mathbf{k} = (k_0, k_1, \ldots, k_m) \in \mathbb{N}^{m+1}$, the leading monomial of

$$E^{\mathbf{k}} := E_0^{k_0} E_1^{k_1} \cdots E_m^{k_m}$$

is therefore

$$M^{\mathbf{k}} := M_0^{k_0} M_1^{k_1} \cdots M_m^{k_m}.$$

The exponents of the variables in $M^{\mathbf{k}}$ determine $\mathbf{k}$: namely,

$$k_0 = \nu_{b_1}(M^{\mathbf{k}}),$$

$$k_i = \nu_{a_i}(M^{\mathbf{k}}) - \nu_{a_{i+1}}(M^{\mathbf{k}}) \qquad (1 \le i < m),$$

and

$$k_m = \nu_{a_m}(M^{\mathbf{k}}).$$

Thus the monomials $M^{\mathbf{k}}$ are pairwise distinct. Therefore the elements $E^{\mathbf{k}}$ are $\mathbb{Z}$-linearly independent, since each has a distinct leading monomial with coefficient 1. This proves algebraic independence. $\square$

Remark 1.4. *The proof of Lemma 1.3 was written so that it works with any monomial order. If one only wants a quick proof of algebraic independence, one can instead choose a specific lexicographic order, for example*

$$a_1 > a_2 > \cdots > a_m > b_1 > b_2 > \cdots > b_m.$$

Then

$$\mathrm{LT}(E_r) = a_1 a_2 \cdots a_r b_{r+1} b_{r+2} \cdots b_m$$

immediately, and the distinct-leading-monomial argument for the products $E_0^{k_0} \cdots E_m^{k_m}$ *goes through without the preliminary chain argument.*

1.2 The diagonal subring and chain monomials

Give B_m the $\mathbb{Z}^m$-grading determined by

$$\deg a_i = \deg b_i = \varepsilon_i,$$

where $\varepsilon_1, \dots, \varepsilon_m$ is the standard basis of $\mathbb{Z}^m$. Let

$$\delta = (1, 1, \dots, 1) \in \mathbb{Z}^m.$$

Each z_I and each E_r has multidegree δ. Define the diagonal subring

$$A_m = \bigoplus_{n \ge 0} (B_m)_{n\delta},$$

where $(B_m)_\gamma$ denotes the γ-th degree component of B_m. This is a subring of B_m, because the product of two homogeneous elements of multidegrees $n\delta$ and $n'\delta$ is homogeneous of multidegree $(n + n')\delta$.

A monomial of B_m is called *diagonal of diagonal degree* n if it has multidegree $n\delta$. Equivalently, it has the form

$$\prod_{i=1}^{m} a_i^{h_i} b_i^{n-h_i}$$

for a uniquely determined vector

$$h = (h_1, \dots, h_m) \in \{0, 1, \dots, n\}^m.$$

This vector h is called the *height vector* of the diagonal monomial, and its entries h_i are called the *heights* of the respective labels i. Thus

$$h_i = \nu_{a_i}(M)$$

for a diagonal monomial M of diagonal degree n (where $\nu_x(M)$ denotes the exponent of a variable x in M); equivalently, the exponent of b_i in M is $n - h_i$. Clearly, the diagonal monomials (of all diagonal degrees n) form a basis of the $\mathbb{Z}$-module A_m.

A *chain monomial* will mean a product

$$z_{I_1} z_{I_2} \cdots z_{I_n}$$

where the subsets $I_1, I_2, \ldots, I_n$ form a multichain under inclusion (i.e., any i and j satisfy $I_i \subseteq I_j$ or $I_j \subseteq I_i$). Since the ring is commutative, the order of the factors is immaterial. When convenient, we write such a monomial in the form $z_{I_1} z_{I_2} \cdots z_{I_n}$ with

$$I_n \subseteq I_{n-1} \subseteq \cdots \subseteq I_1.$$

Note that $n = 0$ is allowed.

Lemma 1.5. *The subring A_m is generated, as a $\mathbb{Z}$-algebra, by the 2^m monomials*

$$z_I \qquad (I \subseteq [m]).$$

Moreover, the chain monomials are precisely the ordinary monomial basis of A_m written in this generating set. In particular, they form a $\mathbb{Z}$-basis of A_m.

Proof. Each z_I has multidegree δ, so the subring generated by the z_I's is contained in A_m.

Conversely, consider a monomial of multidegree $n\delta$. Such a monomial has the form

$$\prod_{i=1}^{m} a_i^{h_i} b_i^{n-h_i}$$

for a unique vector

$$h = (h_1, \ldots, h_m) \in \{0, 1, \ldots, n\}^m.$$

For each $s \in \{1, 2, \ldots, n\}$, set

$$I_s(h) = \{i \in [m] \mid h_i \geq s\}.$$

Then

$$I_n(h) \subseteq I_{n-1}(h) \subseteq \cdots \subseteq I_1(h)$$

and

$$\prod_{s=1}^{n} z_{I_s(h)} = \prod_{i=1}^{m} a_i^{h_i} b_i^{n-h_i}.$$

Thus every diagonal monomial is a chain monomial in the z_I's. Conversely, every chain monomial is visibly a diagonal monomial.[2] The ordinary monomials in the a_i's and b_i's form a $\mathbb{Z}$-basis of B_m, so the diagonal monomials form a $\mathbb{Z}$-basis of A_m. Hence the chain monomials form a $\mathbb{Z}$-basis of A_m. $\square$

We shall also use the elementary straightening identity

$$z_I z_J = z_{I\cap J} z_{I\cup J} \qquad (I, J \subseteq [m]), \tag{2}$$

which holds directly in A_m, since both sides have the same exponent of each variable a_i and b_i.

Remark 1.6 (Hibi-ring context, not used in the proof)**.** *One standard way to package Lemma 1.5 and (2) is by introducing the polynomial ring*

$$T_m = \mathbb{Z}[y_I \mid I \subseteq [m]]$$

and the ideal J_m generated by

$$y_I y_J - y_{I\cap J} y_{I\cup J} \qquad (I, J \subseteq [m]).$$

Then the map

$$T_m/J_m \longrightarrow A_m, \qquad y_I \longmapsto z_I,$$

is an isomorphism. This is the Hibi ring of the Boolean lattice $2^{[m]}$; see Hibi [7]. We shall not need this presentation below.

1.3 A bijection between monomials and pairs of permutations and tuples

Recall that the symmetric group S_m consists of the permutations of $[m] := \{1, 2, \ldots, m\}$. We write permutations in one-line notation; for example, 31524 or $(3, 1, 5, 2, 4)$ is a permutation in S_5. (We shall occasionally omit commas and parentheses in this way.) In particular, $12\cdots m$ is the identity permutation in S_m.

For a permutation $\pi = (\pi_1, \ldots, \pi_m) \in S_m$, set

$$\operatorname{Des}(\pi) = \{s \in \{1, \ldots, m-1\} \mid \pi_s > \pi_{s+1}\}, \qquad \operatorname{des}(\pi) = |\operatorname{Des}(\pi)|.$$

For $0 \le j \le m$, write

$$P_j^\pi = \{\pi_1, \pi_2, \ldots, \pi_j\},$$

so that $P_0^\pi = \varnothing$ and $P_m^\pi = [m]$. Define the Garsia–Stanton descent monomial

$$u_\pi = \prod_{s\in\operatorname{Des}(\pi)} z_{P_s^\pi} \in A_m.$$

[2] And it is easy to see that distinct chains $(I_1, I_2, \ldots, I_n)$ lead to different chain monomials $z_{I_1} z_{I_2} \cdots z_{I_n}$: indeed, the length n of the chain is the diagonal degree of the chain monomial, and each exponent $h_i = \nu_{a_i}(z_{I_1} z_{I_2} \cdots z_{I_n})$ is the number of subsets $I_1, I_2, \ldots, I_n$ that contain i, which (because of the chain condition $I_n \subseteq I_{n-1} \subseteq \cdots \subseteq I_1$) entails that the first h_i sets $I_1, I_2, \ldots, I_{h_i}$ contain i while the remaining $n - h_i$ sets $I_{h_i+1}, I_{h_i+2}, \ldots, I_n$ don't. Thus, from the monomial $z_{I_1} z_{I_2} \cdots z_{I_n}$, we can tell which i's lie in which I_j's.

Thus $u_{12\cdots m} = 1$.

Given $\pi \in S_m$ and $\mathbf{i} = (i_0, i_1, \ldots, i_m) \in \mathbb{N}^{m+1}$, put

$$M_{\pi,\mathbf{i}} = u_\pi \prod_{j=0}^{m} \left(z_{P_j^\pi} \right)^{i_j}. \tag{3}$$

This is a chain monomial, of diagonal degree

$$\operatorname{des}(\pi) + i_0 + i_1 + \cdots + i_m.$$

If $h = (h_1, \ldots, h_m) \in \mathbb{Z}^m$, then the *sorting permutation* of h, denoted $\operatorname{spr}(h)$, will mean the permutation in S_m obtained by listing the labels $1, 2, \ldots, m$ in decreasing order of their heights h_i, breaking ties in increasing order of the labels. That is, $\operatorname{spr}(h)$ is the unique permutation $\pi \in S_m$ satisfying

$$h_{\pi_1} \geq h_{\pi_2} \geq \cdots \geq h_{\pi_m} \tag{4}$$

and

$$\text{if } h_{\pi_s} = h_{\pi_{s+1}}, \text{ then } \pi_s < \pi_{s+1}. \tag{5}$$

(This is related to the notion of standardization in combinatorics, but not literally the same; we use the name "sorting permutation" to avoid confusion.)

Lemma 1.7 (height vector of the distinguished chain monomial)**.** *Let $\pi \in S_m$ and $\mathbf{i} = (i_0, i_1, \ldots, i_m) \in \mathbb{N}^{m+1}$. Let h be the height vector of the chain monomial $M_{\pi,\mathbf{i}}$. Then*

$$\operatorname{spr}(h) = \pi.$$

Proof. We analyze the height vector $h = (h_1, h_2, \ldots, h_m)$ of $M_{\pi,\mathbf{i}}$. For $1 \leq s \leq m$, put

$$d_s^\pi = \#\{q \in \operatorname{Des}(\pi) \mid q \geq s\}, \qquad g_s = i_s + i_{s+1} + \cdots + i_m.$$

The factor u_π contributes d_s^π to the height of the label π_s, while the prefix factors

$$\prod_{j=0}^{m} \left(z_{P_j^\pi} \right)^{i_j}$$

contribute g_s to the height of the same label. Hence

$$h_{\pi_s} = d_s^\pi + g_s \qquad (1 \leq s \leq m). \tag{6}$$

The sequence $g_1, g_2, \ldots, g_m$ is weakly decreasing. The sequence $d_1^\pi, d_2^\pi, \ldots, d_m^\pi$ is also weakly decreasing; more precisely,

$$d_s^\pi = d_{s+1}^\pi + 1 \quad \text{if } s \in \operatorname{Des}(\pi), \qquad d_s^\pi = d_{s+1}^\pi \quad \text{if } s \notin \operatorname{Des}(\pi).$$

Thus (6) gives

$$h_{\pi_1} \geq h_{\pi_2} \geq \cdots \geq h_{\pi_m}.$$

Moreover, if $s \in \mathrm{Des}(\pi)$, then $d^\pi_s = d^\pi_{s+1} + 1$, whence

$$h_{\pi_s} > h_{\pi_{s+1}}.$$

If equality $h_{\pi_s} = h_{\pi_{s+1}}$ occurs, then therefore $s \notin \mathrm{Des}(\pi)$, so $\pi_s < \pi_{s+1}$. This is exactly the rule defining the sorting permutation: list the labels in decreasing height, and break equal-height ties increasingly. Hence $\mathrm{spr}(h) = \pi$. □

Lemma 1.8 (labelled-antichain bijection)**.** *The map*

$$(\pi, \mathbf{i}) \longmapsto M_{\pi,\mathbf{i}}$$

is a bijection from

$$\{(\pi, \mathbf{i}) \mid \pi \in S_m,\ \mathbf{i} \in \mathbb{N}^{m+1}\}$$

to the set of all chain monomials in A_m. More precisely, for each $n \geq 0$, it restricts to a bijection from the pairs satisfying

$$\mathrm{des}(\pi) + i_0 + i_1 + \cdots + i_m = n$$

to the chain monomials of diagonal degree n.

Proof. Well-definedness. For a fixed pair $(\pi, \mathbf{i})$, the sets

$$P^\pi_0 \subseteq P^\pi_1 \subseteq \cdots \subseteq P^\pi_m$$

form a chain. The additional factors appearing in u_π are among the same sets P^π_s. Hence $M_{\pi,\mathbf{i}}$ is a chain monomial. Its diagonal degree is

$$\mathrm{des}(\pi) + i_0 + i_1 + \cdots + i_m,$$

since each z_I has diagonal degree 1.

Surjectivity. Let a chain monomial of diagonal degree n be given. By Lemma 1.5, it has a unique height vector

$$h = (h_1, \ldots, h_m) \in \{0, 1, \ldots, n\}^m,$$

where the monomial is

$$\prod_{i=1}^m a_i^{h_i} b_i^{n-h_i}.$$

Let

$$\pi = \mathrm{spr}(h)$$

be the sorting permutation of h. Thus π lists the labels in decreasing order of their heights, breaking ties increasingly:

$$h_{\pi_1} \geq h_{\pi_2} \geq \cdots \geq h_{\pi_m},$$

and if $h_{\pi_s} = h_{\pi_{s+1}}$, then $\pi_s < \pi_{s+1}$.

For $1 \leq s \leq m$, put

$$d^\pi_s = \#\{q \in \mathrm{Des}(\pi) \mid q \geq s\}.$$

This sequence $(d_1^\pi, d_2^\pi, \dots, d_m^\pi)$ is the "descent staircase" of π: it drops by 1 exactly at those s that are descents of π. Now define

$$g_s = h_{\pi_s} - d_s^\pi \qquad (1 \le s \le m).$$

We claim that

$$g_1 \ge g_2 \ge \cdots \ge g_m \ge 0. \tag{7}$$

Indeed, if $s \notin \mathrm{Des}(\pi)$, then $d_s^\pi = d_{s+1}^\pi$, and the inequality $g_s \ge g_{s+1}$ follows from $h_{\pi_s} \ge h_{\pi_{s+1}}$. If $s \in \mathrm{Des}(\pi)$, then $\pi_s > \pi_{s+1}$; by the tie-breaking rule this forces $h_{\pi_s} > h_{\pi_{s+1}}$, whence $h_{\pi_s} \ge h_{\pi_{s+1}} + 1$. Since $d_s^\pi = d_{s+1}^\pi + 1$ in this case, we again get $g_s \ge g_{s+1}$. Finally, from $d_m^\pi = 0$, we obtain $g_m = h_{\pi_m} \ge 0$. Thus, (7) is proved.

Set

$$i_s = g_s - g_{s+1} \qquad (1 \le s < m),$$
$$i_m = g_m,$$

and

$$i_0 = n - \mathrm{des}(\pi) - g_1.$$

The last integer i_0 is nonnegative, since $g_1 = h_{\pi_1} - \mathrm{des}(\pi)$ and $h_{\pi_1} \le n$. These formulas give a vector $\mathbf{i} = (i_0, i_1, \dots, i_m) \in \mathbb{N}^{m+1}$, since (7) yields that $i_1, i_2, \dots, i_m$ are nonnegative as well. Moreover, $i_0 + i_1 + \cdots + i_m = n - \mathrm{des}(\pi)$, whence $\mathrm{des}(\pi) + i_0 + i_1 + \cdots + i_m = n$; in other words, the monomial $M_{\pi,\mathbf{i}}$ has diagonal degree n.

For $1 \le s \le m$, the height of $M_{\pi,\mathbf{i}}$ at the label π_s is

$$d_s^\pi + i_s + i_{s+1} + \cdots + i_m.$$

By the definition of the i's, this equals $d_s^\pi + g_s = h_{\pi_s}$. Hence $M_{\pi,\mathbf{i}}$ has the same height vector as the original chain monomial, and it also has diagonal degree n. These data determine a diagonal monomial, so $M_{\pi,\mathbf{i}}$ is the original chain monomial. So our original chain monomial has the form $M_{\pi,\mathbf{i}}$ for some $\pi \in S_m$ and $\mathbf{i} \in \mathbb{N}^{m+1}$. This proves the surjectivity of the map in the lemma.

Injectivity. We show that the pair $(\pi, \mathbf{i})$ can be recovered from the monomial $M_{\pi,\mathbf{i}}$. Let h be its height vector, and let n be its diagonal degree. By Lemma 1.7, the sorting permutation of h is π. Thus π is recovered from the monomial.

Once π is known, its descent staircase d_s^π is known. Define

$$g_s = h_{\pi_s} - d_s^\pi \qquad (1 \le s \le m).$$

For the monomial $M_{\pi,\mathbf{i}}$, equation (6) says precisely that

$$g_s = i_s + i_{s+1} + \cdots + i_m.$$

Hence the successive differences recover

$$i_s = g_s - g_{s+1} \qquad (1 \le s < m), \qquad i_m = g_m.$$

Finally,

$$i_0 = n - \mathrm{des}(\pi) - g_1,$$

because n is the diagonal degree of $M_{\pi,\mathbf{i}}$ and $\operatorname{des}(\pi) + g_1 = \operatorname{des}(\pi) + i_1 + \cdots + i_m$ is the contribution of all factors other than $(z_{P_0^\pi})^{i_0}$. Thus all entries of $\mathbf{i}$ are recovered. Therefore two different pairs $(\pi, \mathbf{i})$ cannot give the same chain monomial. This proves the injectivity of the map in the lemma. $\square$

Remark 1.9 (Stanley's fundamental theorem on P-partitions)**.** *Lemma 1.8 is precisely the antichain case of Stanley's fundamental theorem on P-partitions [8, Lemma 3.15.3]. A function $h : [m] \to \{0, 1, \ldots, n\}$ is a P-partition for the labelled antichain on $[m]$, since there are no order relations. Sorting h gives a unique permutation π; the tie-breaking rule produces the strict inequalities at descents of π. Subtracting the descent staircase d_s^π leaves an ordinary partition $g_1 \geq \cdots \geq g_m \geq 0$, whose successive differences are the numbers $i_1, \ldots, i_m$; the number i_0 records the unused vertical room up to n.*

Remark 1.10 (Eulerian identity from the bijection)**.** *Comparing cardinalities in Lemma 1.8 for fixed diagonal degree n gives Worpitzky's identity*

$$(n+1)^m = \sum_{\pi \in S_m} \binom{n - \operatorname{des}(\pi) + m}{m},$$

where the binomial coefficient is interpreted as 0 when $n - \operatorname{des}(\pi) < 0$. Indeed, the left hand side counts height vectors $h \in \{0, 1, \ldots, n\}^m$, while for a fixed π the number of $(i_0, \ldots, i_m) \in \mathbb{N}^{m+1}$ satisfying $i_0 + \cdots + i_m = n - \operatorname{des}(\pi)$ is the displayed binomial coefficient. Equivalently,

$$\sum_{n \geq 0} (n+1)^m t^n = \frac{\sum_{\pi \in S_m} t^{\operatorname{des}(\pi)}}{(1-t)^{m+1}}.$$

Conversely, if one already knows this identity, then in Lemma 1.8 it would be enough to prove either surjectivity or injectivity for each fixed n; the equality of the finite cardinalities would force the other property by the pigeonhole principle.

1.4 Triangularity and freeness

Our next goal is to show that the chain monomials $M_{\pi,\mathbf{i}}$ are the leading terms of the polynomials $u_\pi E_{\mathbf{i}}$ with respect to a certain partial order on diagonal monomials. We shall now define this order. It is useful to keep in mind that this order is **not a monomial order**. It is only an order on the basis of diagonal chain monomials of a fixed diagonal degree; it does not respect multiplication.

The *decreasing rearrangement* of any m-tuple $h = (h_1, \ldots, h_m) \in \mathbb{Z}^m$ is defined to be the m-tuple $h^\downarrow = (h_1^\downarrow \geq h_2^\downarrow \geq \cdots \geq h_m^\downarrow)$ that contains the same entries as h but sorted in weakly decreasing order. For instance, $(2, 5, 2, 7)^\downarrow = (7, 5, 2, 2)$. We note the following simple fact:

Lemma 1.11. *Let $h = (h_1, \ldots, h_m)$ and $k = (k_1, \ldots, k_m)$ be two m-tuples in $\mathbb{Z}^m$. If $\operatorname{spr}(h) = \operatorname{spr}(k)$ and $h^\downarrow = k^\downarrow$, then $h = k$.*

Proof. Let $\pi = \operatorname{spr}(h) = \operatorname{spr}(k)$. Then, from (4), we see that $h^\downarrow = (h_{\pi_1}, h_{\pi_2}, \ldots, h_{\pi_m})$. Likewise, $k^\downarrow = (k_{\pi_1}, k_{\pi_2}, \ldots, k_{\pi_m})$. Thus, from $h^\downarrow = k^\downarrow$, we obtain $(h_{\pi_1}, h_{\pi_2}, \ldots, h_{\pi_m}) = (k_{\pi_1}, k_{\pi_2}, \ldots, k_{\pi_m})$. Since π is a permutation of $[m]$, this gives $h_i = k_i$ for every label i. Thus $h = k$. $\square$

Let $h = (h_1, \dots, h_m)$ and $k = (k_1, \dots, k_m)$ be two height vectors with entries in $\{0, 1, \dots, n\}$. Let

$$h^\downarrow = (h_1^\downarrow \geq h_2^\downarrow \geq \cdots \geq h_m^\downarrow) \qquad \text{and} \qquad k^\downarrow = (k_1^\downarrow \geq k_2^\downarrow \geq \cdots \geq k_m^\downarrow)$$

be their decreasing rearrangements. We say that $h^\downarrow$ is *dominated by* $k^\downarrow$ if

$$h_1^\downarrow + \cdots + h_q^\downarrow \leq k_1^\downarrow + \cdots + k_q^\downarrow \qquad \text{for every } q \in \{1, 2, \dots, m\}.$$

We define the *dominance–lex order* on diagonal monomials of diagonal degree n as follows. First compare their decreasingly sorted height vectors by dominance. If these sorted height vectors are unequal, then the monomial whose sorted height vector strictly dominates the other is declared to be larger. If the decreasingly sorted height vectors are equal, break ties by the sorting permutation: the height vector whose sorting permutation is lexicographically smaller is declared to be larger. Explicitly, two distinct diagonal monomials $\mathbf{m}$ and $\mathbf{n}$, with respective height vectors h and k, satisfy $\mathbf{m} > \mathbf{n}$ if and only if either

$$h^\downarrow \text{ strictly dominates } k^\downarrow \text{ (i.e., dominates it and is not equal to it),}$$

or else

$$h^\downarrow = k^\downarrow \quad \text{and} \quad \mathrm{spr}(h) < \mathrm{spr}(k)$$

(where the $<$ sign refers to lexicographic order of one-line notations). This is a partial order[3], because dominance is only a partial order. Nevertheless, the leading term statements below make sense: a monomial is a leading monomial of a polynomial if it is the unique maximal monomial among the monomials appearing in that polynomial. Not every polynomial has a leading monomial, but if it does, then this monomial is unique.

For $\mathbf{i} = (i_0, i_1, \dots, i_m) \in \mathbb{N}^{m+1}$, put

$$E_{\mathbf{i}} = E_0^{i_0} E_1^{i_1} \cdots E_m^{i_m}.$$

Also put

$$g_s = i_s + i_{s+1} + \cdots + i_m \qquad (1 \leq s \leq m).$$

Thus

$$g_1 \geq g_2 \geq \cdots \geq g_m \geq 0.$$

Lemma 1.12 (dominance for terms of $E_{\mathbf{i}}$)**.** *Let $w = (w_1, \dots, w_m)$ be the height vector of a diagonal monomial appearing in $E_{\mathbf{i}}$. Then*

$$w^\downarrow \textit{ is dominated by } (g_1, g_2, \dots, g_m).$$

Moreover, for any permutation $\rho \in S_m$, the term

$$\prod_{j=0}^{m} \left(z_{P_j^\rho} \right)^{i_j}$$

of $E_{\mathbf{i}}$ has height vector w given by

$$w_{\rho_s} = g_s \qquad (1 \leq s \leq m).$$

[3]It is easily seen to be asymmetric and transitive.

Proof. Recall that $E_{\mathbf{i}} = E_0^{i_0} E_1^{i_1} \cdots E_m^{i_m}$, where each E_j is the sum of the z_J's over all j-element subsets $J \subseteq [m]$. Thus, expanding this product, we obtain a sum of terms, each of which is a product of such z_J's. To be more specific: A term of $E_{\mathbf{i}}$ is obtained by choosing, for each of the i_j copies of E_j, a j-element subset $J \subseteq [m]$, and then multiplying the z_J's corresponding to all the chosen J's. In the height vector of this term, the height w_x is then the total number of chosen subsets J that contain x.

Let S be any q-element subset of $[m]$. Each chosen j-element subset meets S in at most $\min(q, j)$ elements. Hence

$$\sum_{x \in S} w_x \leq \sum_{j=0}^{m} i_j \min(q, j) = g_1 + g_2 + \cdots + g_q.$$

Taking the maximum over all q-element subsets S gives the desired dominance inequality.

The final assertion follows directly from the definition of the prefixes $P_j^\rho = \{\rho_1, \ldots, \rho_j\}$: the label ρ_s lies in P_j^ρ exactly when $j \geq s$, so it is counted $i_s + i_{s+1} + \cdots + i_m = g_s$ times. $\square$

Remark 1.13. *Lemma 1.12 says, in particular, that the leading term of $E_{\mathbf{i}}$ itself is*

$$\prod_{j=0}^{m} \left(z_{\{1,2,\ldots,j\}} \right)^{i_j},$$

with respect to the dominance–lex order.[4] *Indeed, the maximal sorted height vector is $(g_1, \ldots, g_m)$, and among all terms with this sorted height vector the lexicographically smallest sorting permutation is $12 \cdots m$. This is the usual triangularity underlying the fundamental theorem on symmetric polynomials (see, e.g., [3, §4.3, Proposition 5]), in homogenized form. The point of the next lemmas is that, after multiplying by u_π, the relevant prefix order becomes $\pi_1, \pi_2, \ldots, \pi_m$ rather than $1, 2, \ldots, m$.*

Lemma 1.14 (dominance after multiplying by u_π)**.** *Let $\pi \in S_m$ and $\mathbf{i} = (i_0, \ldots, i_m) \in \mathbb{N}^{m+1}$. Let h^* be the height vector of*

$$M_{\pi,\mathbf{i}} = u_\pi \prod_{j=0}^{m} \left(z_{P_j^\pi} \right)^{i_j}.$$

Let h be the height vector of any diagonal monomial appearing in

$$u_\pi E_{\mathbf{i}}.$$

Then $h^\downarrow$ is dominated by $(h^*)^\downarrow$.

Proof. Put

$$d_s = d_s^\pi = \#\{q \in \mathrm{Des}(\pi) \mid q \geq s\}, \qquad g_s = i_s + i_{s+1} + \cdots + i_m \qquad (1 \leq s \leq m).$$

The factor u_π contributes height d_s to the label π_s. Hence, the height vector h^* satisfies

$$h^*_{\pi_s} = d_s + g_s \qquad (1 \leq s \leq m),$$

[4] We have not nailed down the leading coefficient yet. (It is 1, and this is a particular case of Lemma 1.16 below.)

and $\operatorname{spr}(h^*) = \pi$ by Lemma 1.7. In particular,

$$(h^*)^{\downarrow} = (d_1 + g_1, d_2 + g_2, \dots, d_m + g_m)$$

because $d_1 \geq d_2 \geq \cdots \geq d_m$ and $g_1 \geq g_2 \geq \cdots \geq g_m$.

Now consider the diagonal monomial in $u_\pi E_{\mathbf{i}}$ whose height vector is h. Let $w = (w_1, w_2, \dots, w_m)$ be the height vector of the corresponding monomial in $E_{\mathbf{i}}$. Thus,

$$h_x = d_x^{\pi,\mathrm{lab}} + w_x, \qquad \text{where } d_{\pi_s}^{\pi,\mathrm{lab}} = d_s$$

(since $(d_1^{\pi,\mathrm{lab}}, d_2^{\pi,\mathrm{lab}}, \dots, d_m^{\pi,\mathrm{lab}})$ is the height vector of u_π). Let S be a q-element subset of $[m]$. Since $d_1 \geq d_2 \geq \cdots \geq d_m$, the largest possible sum of the d-contributions over q labels is $d_1 + \cdots + d_q$. Hence

$$\sum_{x \in S} d_x^{\pi,\mathrm{lab}} \leq d_1 + \cdots + d_q. \tag{8}$$

By Lemma 1.12, applied in its proof to the same subset S, we also have

$$\sum_{x \in S} w_x \leq g_1 + \cdots + g_q. \tag{9}$$

Adding these inequalities gives

$$\sum_{x \in S} h_x \leq \sum_{s=1}^{q} (d_s + g_s) = \sum_{s=1}^{q} (h^*)^{\downarrow}_s. \tag{10}$$

Taking the maximum over all q-element subsets S proves the dominance claim. $\square$

Lemma 1.15 (the equality case)**.** *Keep the notation of Lemma 1.14. Assume that $h^{\downarrow} = (h^*)^{\downarrow}$. Then*

$$\operatorname{spr}(h) \geq_{\mathrm{lex}} \pi.$$

Moreover, if $\operatorname{spr}(h) = \pi$, then $h = h^*$.

Proof. Let

$$\sigma = \operatorname{spr}(h).$$

For $q \in \{1, 2, \dots, m\}$, set

$$S_q = \{\sigma_1, \dots, \sigma_q\}.$$

This is a q-element subset on which the sum of the entries of h is maximal. Since $h^{\downarrow} = (h^*)^{\downarrow}$, the dominance inequality (10) in Lemma 1.14 is an equality for $S = S_q$. The two inequalities (8) and (9) that were added in the proof of that lemma therefore must both be equalities. In particular,

$$\sum_{x \in S_q} d_x^{\pi,\mathrm{lab}} = d_1 + d_2 + \cdots + d_q \qquad (1 \leq q \leq m). \tag{11}$$

We now interpret this condition. Decompose the word $\pi_1, \dots, \pi_m$ into its increasing runs[5]:

$$R_1 \mid R_2 \mid \cdots \mid R_t.$$

The cuts occur exactly at the descents of π. The sequence $d_1, d_2, \dots, d_m$ is constant on each run and strictly decreases when we move from one run to the next. Thus the labels in R_1 have the largest d-value, the labels in R_2 have the next largest d-value, and so on. Equation (11), for all q, says exactly that the first q entries of σ always choose q labels with largest possible d-sum. Consequently, σ must list all labels of R_1 before all labels of R_2, all labels of R_2 before all labels of R_3, and so on. Indeed, if a label from a later run occurred before a label from an earlier run, then for a suitable q the set S_q would contain the former but not the latter; replacing the former by the latter would strictly increase the left hand side of (11), a contradiction.

Thus σ is obtained from π by permuting entries only inside the increasing runs of π. Each such run is already written in increasing order, and increasing order is lexicographically smallest among all orderings of the same set of labels. Therefore

$$\sigma \geq_{\text{lex}} \pi.$$

Finally, if $\operatorname{spr}(h) = \pi$, then h and h^* have the same decreasing rearrangement and the same sorting permutation. Thus, Lemma 1.11 yields $h = h^*$. □

Lemma 1.16 (coefficient of the top contribution)**.** *Let* $\pi \in S_m$ *and* $\mathbf{i} \in \mathbb{N}^{m+1}$*. In the expansion of* $u_\pi E_{\mathbf{i}}$*, the monomial* $M_{\pi,\mathbf{i}}$ *appears with coefficient* 1*.*

Proof. We must show that the coefficient of the monomial

$$M_{\pi,\mathbf{i}} = u_\pi \prod_{j=0}^{m} \left(z_{P_j^\pi}\right)^{i_j}$$

in $u_\pi E_{\mathbf{i}}$ is 1. Since multiplication by the monomial u_π sends distinct monomials to distinct monomials, it is enough to show that the coefficient of

$$\prod_{j=0}^{m} \left(z_{P_j^\pi}\right)^{i_j}$$

in $E_{\mathbf{i}}$ is 1.

For this proof only, we equip the polynomial ring B_m with the lexicographic monomial order corresponding to

$$a_{\pi_1} > a_{\pi_2} > \cdots > a_{\pi_m} > b_{\pi_1} > b_{\pi_2} > \cdots > b_{\pi_m}.$$

We claim that, for each j, the leading monomial of E_j is $z_{P_j^\pi}$, with coefficient 1. Indeed, among all j-element subsets $J \subseteq [m]$, the subset $P_j^\pi = \{\pi_1, \dots, \pi_j\}$ is lexicographically first in the following sense: if $J \neq P_j^\pi$, and t is the first index such that π_t belongs to exactly one

[5] The *increasing runs* of a word are its inclusion-maximal increasing contiguous subsequences. For instance, the increasing runs of the word 425816397 are 4, 258, 16, 39 and 7.

of J and P_j^π, then $t \le j$, and $\pi_t \in P_j^\pi \setminus J$. Thus $z_{P_j^\pi}$ has exponent 1 on the variable a_{π_t}, whereas z_J has exponent 0 on this variable. Hence

$$z_{P_j^\pi} > z_J$$

in our lexicographic order.

Therefore the polynomial E_j has leading monomial

$$\mathrm{LT}(E_j) = z_{P_j^\pi} \qquad (0 \le j \le m),$$

with leading coefficient 1. Since leading monomials and leading coefficients multiply in a monomial order, the product

$$E_{\mathbf{i}} = E_0^{i_0} E_1^{i_1} \cdots E_m^{i_m}$$

has leading monomial

$$\prod_{j=0}^{m} \left(z_{P_j^\pi} \right)^{i_j}$$

with coefficient 1.

Thus, as we have explained above, the coefficient of $M_{\pi,\mathbf{i}}$ in $u_\pi E_{\mathbf{i}}$ is also 1. $\square$

Lemma 1.17 (antichain triangularity)**.** *Let $\pi \in S_m$ and $\mathbf{i} = (i_0, \ldots, i_m) \in \mathbb{N}^{m+1}$. Then the polynomial*

$$F_{\pi,\mathbf{i}} := u_\pi E_{\mathbf{i}} = u_\pi E_0^{i_0} E_1^{i_1} \cdots E_m^{i_m} \in A_m$$

has leading chain monomial, with respect to the dominance–lex order, equal to

$$M_{\pi,\mathbf{i}} = u_\pi \prod_{j=0}^{m} \left(z_{P_j^\pi} \right)^{i_j}.$$

The coefficient of this leading chain monomial is 1*.*

Proof. Let h^* be the height vector of $M_{\pi,\mathbf{i}}$, and let h be the height vector of any monomial M' appearing in $F_{\pi,\mathbf{i}}$. Our first goal is to prove that $M_{\pi,\mathbf{i}} \ge M'$ in the dominance-lex order. By Lemma 1.14, the sorted height vector $h^\downarrow$ is dominated by $(h^*)^\downarrow$. If this dominance is strict, then $M_{\pi,\mathbf{i}} > M'$ follows immediately.

It remains to consider the case $h^\downarrow = (h^*)^\downarrow$. By Lemma 1.15, we then have $\mathrm{spr}(h) \ge_{\mathrm{lex}} \pi = \mathrm{spr}(h^*)$ (since $\mathrm{spr}(h^*) = \pi$ by Lemma 1.7). If this inequality is strict, then the tie-breaker in the dominance–lex order again ensures $M_{\pi,\mathbf{i}} > M'$.

The remaining case is when $h^\downarrow = (h^*)^\downarrow$ and $\mathrm{spr}(h) = \mathrm{spr}(h^*)$. But in this case, Lemma 1.15 gives $h = h^*$, which forces $M_{\pi,\mathbf{i}} = M'$: the two monomials have the same height vector and the same diagonal degree, and these data determine a diagonal monomial.

So we have shown that the monomial $M_{\pi,\mathbf{i}}$ is the leading chain monomial of $F_{\pi,\mathbf{i}}$. It remains to show that its coefficient is 1. But this follows from Lemma 1.16. $\square$

Remark 1.18. *Lemma 1.17 is not a Gröbner-basis statement about the elements $E_0, E_1, \ldots, E_m$. The dominance–lex order is a partial order on the diagonal monomial basis, not a multiplicative monomial order on B_m.*

1.5 Proof of Theorem 1.2

We now recall once again the three nested rings we are working with:

$$\begin{aligned} & C_m = \mathbb{Z}[E_0, E_1, \ldots, E_m] \\ \subseteq\ & A_m = (\text{diagonal subring of } B_m) = \bigoplus_{n \geq 0} (B_m)_{n\delta} \\ \subseteq\ & B_m = \mathbb{Z}[a_1, b_1, a_2, b_2, \ldots, a_m, b_m]. \end{aligned}$$

Proposition 1.19. *The C_m-module A_m is free with basis*

$$\{u_\pi \mid \pi \in S_m\}.$$

In particular, since $u_{12\cdots m} = 1$, the inclusion $C_m \hookrightarrow A_m$ has a C_m-linear retraction.

Proof. Since C_m is the polynomial ring $\mathbb{Z}[E_0, \ldots, E_m]$, it has a basis consisting of all monomials $E_{\mathbf{i}} = E_0^{i_0} \cdots E_m^{i_m}$ with $\mathbf{i} = (i_0, i_1, \ldots, i_m) \in \mathbb{N}^{m+1}$. The products

$$F_{\pi,\mathbf{i}} = u_\pi E_0^{i_0} \cdots E_m^{i_m} \qquad (\pi \in S_m,\ \mathbf{i} \in \mathbb{N}^{m+1})$$

are precisely the elements obtained by multiplying the proposed basis elements u_π by this monomial basis of C_m.

By Lemma 1.17, each $F_{\pi,\mathbf{i}}$ has leading chain monomial $M_{\pi,\mathbf{i}}$ with coefficient 1. By Lemma 1.8, the monomials $M_{\pi,\mathbf{i}}$ are pairwise distinct and, in fact, run over the entire chain-monomial basis of A_m.

Therefore, in each diagonal degree, the transition matrix from the family

$$\{F_{\pi,\mathbf{i}}\}$$

to the chain-monomial basis of A_m is unitriangular after choosing any linear extension of the partial order used in Lemma 1.17. Hence the $F_{\pi,\mathbf{i}}$ form a $\mathbb{Z}$-basis of A_m.

Equivalently,

$$A_m = \bigoplus_{\pi \in S_m} C_m u_\pi$$

as a C_m-module. Since the identity permutation $12 \cdots m$ has no descents, $u_{12\cdots m} = 1$, so the summand for the identity permutation is exactly C_m. Projection onto this summand is a C_m-linear retraction $A_m \to C_m$. □

Remark 1.20 (What is free after base change)**.** *The freeness part of Proposition 1.19 is stronger than the retraction part: it says that the diagonal subring A_m is a free C_m-module of rank $m!$, with a basis containing 1. Consequently, for every C_m-algebra R, the base change*

$$R \otimes_{C_m} A_m$$

is a free R-module of rank $m!$, with basis $1 \otimes u_\pi$ for $\pi \in S_m$.

The analogous freeness statement for the whole ring B_m is false for $m \geq 2$, even if one allows an infinite basis. In fact, B_m is not flat over C_m for $m \geq 2$. See Corollary 1.25 below for a proof of this.

Thus the freeness used in this paper is precisely the freeness of the diagonal summand A_m, not of all of B_m. What we use for B_m is only that A_m is a C_m-direct summand of B_m via the multidegree projection. Hence, after any base change $C_m \to R$, the R-algebra

$$R \otimes_{C_m} B_m$$

contains the finite free R-module $R \otimes_{C_m} A_m$ as an R-module direct summand, and in particular contains the copy of $R = R \otimes_{C_m} C_m$ as a direct summand. This is what we will apply below.

Remark 1.21 (Hilbert-series consistency)**.** *Taking Hilbert series in the free decomposition of Proposition 1.19 gives again the generating-function identity of Remark 1.10. Thus the algebraic basis theorem and the labelled-antichain P-partition count are two explanations of the same numerical identity for Eulerian numbers.*

Remark 1.22 (Garsia–Stanton context)**.** *The monomials u_π are variants of the Garsia–Stanton descent monomials ([6, (7.21)], [5, (6.7)], [2, (2)]). The idea goes back to Garsia's and Stanton's work on group actions on Stanley–Reisner rings and invariant rings [5], [6]. In the classical coinvariant algebra of type A, the corresponding monomials form what is now usually called the Garsia–Stanton descent basis; see, for example, Adin–Brenti–Roichman [1]. The proof above can be viewed as the labelled-antichain, diagonal-monomial shadow of this construction. A close precursor to our proof, but situated in the ordinary polynomial ring $\mathbb{Z}[x_1, x_2, \ldots, x_m]$ instead of A_m, appears in Allen's [2].*

Remark 1.23 (Cohen–Macaulay and quotient-ring context)**.** *Another standard route is to identify A_m with the homogeneous coordinate ring of the Segre embedding of $(\mathbb{P}^1)^m$, or equivalently with the Hibi ring of the Boolean lattice. This ring is Cohen–Macaulay; the elements $E_0, \ldots, E_m$ form a homogeneous system of parameters; hence they form a regular sequence. This gives the Hilbert series of $A_m/(E_0, \ldots, E_m)A_m$, after which one can lift a basis of the quotient. The present proof avoids that quotient and proves the free C_m-basis directly by triangularity. For background on Stanley–Reisner rings and the Cohen–Macaulay viewpoint, see [9].*

Proof of Theorem 1.2. By Proposition 1.19, the inclusion $C_m \hookrightarrow A_m$ has a C_m-linear retraction $\rho : A_m \to C_m$. The decomposition of B_m into multihomogeneous components gives a direct sum decomposition into C_m-submodules according to the cosets of $\mathbb{Z}\delta$ in $\mathbb{Z}^m$, because multiplication by any E_r shifts multidegree by δ. The summand corresponding to the coset of 0 is A_m. Hence the multihomogeneous projection

$$B_m \to A_m$$

is C_m-linear (even A_m-linear) and restricts to the identity on A_m. Composing this projection with the retraction $\rho : A_m \to C_m$ gives a C_m-linear retraction

$$B_m \to A_m \to C_m.$$

This proves the theorem. □

1.6 Non-freeness of the polynomial ring over the diagonal subring

This subsection is meant to establish the non-freeness (and non-flatness) of B_m as a C_m-module, which was claimed in Remark 1.20. Other than this, it is inconsequential.

Proposition 1.24 (A root syzygy among the coefficients)**.** *For each $i \in [m]$, we have the identity*

$$\sum_{j=0}^{m} (-1)^j a_i^{m-j} b_i^j E_j = 0 \qquad \text{in the polynomial ring } B_m.$$

Proof. Fix $i \in [m]$. Recall from Proposition 1.1 that

$$\sum_{j=0}^{m} E_j X^j = \prod_{r=1}^{m} (a_r X + b_r) \qquad \text{in the univariate polynomial ring } B_m[X].$$

Homogenizing this equality (by substituting X/Y for X and multiplying the result by Y^m), we obtain

$$\sum_{j=0}^{m} E_j X^j Y^{m-j} = \prod_{r=1}^{m} (a_r X + b_r Y) \qquad \text{in the bivariate polynomial ring } B_m[X, Y].$$

Now, substitute $X = -b_i$ and $Y = a_i$ here. The right hand side becomes

$$\prod_{r=1}^{m} (a_r(-b_i) + b_r a_i),$$

which is 0, since its i-th factor is $a_i(-b_i) + b_i a_i = 0$. Thus the left hand side is 0 as well; that is,

$$\sum_{j=0}^{m} E_j (-b_i)^j a_i^{m-j} = 0.$$

In other words,

$$\sum_{j=0}^{m} (-1)^j a_i^{m-j} b_i^j E_j = 0.$$

This proves the proposition. ∎

Corollary 1.25. *Assume that $m \geq 2$. Then, the C_m-module B_m is not flat.*

Proof. Taking $i = 1$ in Proposition 1.24, we find

$$a_1^m E_0 - a_1^{m-1} b_1 E_1 + a_1^{m-2} b_1^2 E_2 - \cdots + (-1)^m b_1^m E_m = 0.$$

Hence

$$b_1^m E_m \in (E_0, E_1, \ldots, E_{m-1}) B_m.$$

However,

$$b_1^m \notin (E_0, E_1, \ldots, E_{m-1}) B_m.$$

Indeed, specialize $b_2 = b_3 = \cdots = b_m = 0$. Then

$$\prod_{r=1}^{m}(a_r X + b_r) \mapsto (a_1 X + b_1) a_2 X \cdots a_m X,$$

so the ideal $(E_0, E_1, \ldots, E_{m-1})B_m$ maps into the principal ideal

$$(b_1 a_2 a_3 \cdots a_m),$$

which does not contain b_1^m when $m \geq 2$. Thus E_m is a zero-divisor modulo $(E_0, E_1, \ldots, E_{m-1})B_m$. Hence

$$E_0, E_1, \ldots, E_m$$

is not a regular sequence in B_m. But $E_0, E_1, \ldots, E_m$ is a regular sequence in the polynomial ring $C_m = \mathbb{Z}[E_0, E_1, \ldots, E_m]$. Since flat base change preserves regular sequences, B_m cannot be flat over C_m. □

In particular, B_m is not free over C_m for $m \geq 2$. The case $m = 1$ is exceptional and trivial: then $C_1 = B_1$.

2 Application to splitting arbitrary polynomials

We now deduce some consequences for univariate polynomials.

Theorem 2.1. *Let R be a commutative ring, and let m be a positive integer. Let $P = p_0 + p_1 X + \cdots + p_m X^m \in R[X]$ (with $p_0, p_1, \ldots, p_m \in R$) be a univariate polynomial of degree $\leq m$. Then there exists a commutative ring S and an injective ring homomorphism $R \hookrightarrow S$ such that P factors in $S[X]$ as a product of m polynomials of degree at most 1.*

Proof. Let A_m, B_m, C_m, and $E_0, E_1, \ldots, E_m$ be as in Section 1. By Lemma 1.3, $C_m = \mathbb{Z}[E_0, E_1, \ldots, E_m]$ is a polynomial ring over $\mathbb{Z}$ on the generators $E_0, E_1, \ldots, E_m$. Hence there is a unique ring homomorphism

$$\varphi : C_m \to R, \qquad E_i \longmapsto p_i.$$

Define

$$S = R \otimes_{C_m} B_m,$$

where R is regarded as a C_m-algebra via φ.

The inclusion $C_m \hookrightarrow B_m$ is split as a C_m-module map by Theorem 1.2. Tensoring this split injection with R over C_m yields a split injection of R-modules

$$R = R \otimes_{C_m} C_m \hookrightarrow R \otimes_{C_m} B_m = S \tag{12}$$

(since split injections remain split injections under base change). In particular, the natural ring homomorphism $R \to S$ is injective. We identify R with its image in S, equating each $r \in R$ with $r \otimes_{C_m} 1 \in S$.

More precisely, Proposition 1.19 yields that

$$R \otimes_{C_m} A_m$$

is a free R-module of rank $m!$, with basis $1 \otimes u_\pi$ for $\pi \in S_m$, and the copy of R generated by $1 \otimes u_{12\cdots m} = 1$ is a direct summand of this free module. The splitting algebra $S = R \otimes_{C_m} B_m$ contains this module as an R-module direct summand. Thus the construction gives not merely injectivity of $R \to S$, but a split injection of underlying R-modules.

The canonical ring morphism $\xi : B_m \to R \otimes_{C_m} B_m = S$ (sending each $u \in B_m$ to $1 \otimes u \in S$) sends each $E_j \in B_m$ to $1 \otimes_{C_m} E_j = \varphi(E_j) \otimes_{C_m} 1 = p_j \otimes_{C_m} 1 = p_j \in S$. Let $\alpha_i, \beta_i \in S$ be the images of $a_i, b_i \in B_m$, respectively, under this morphism ξ. In $B_m[X]$ we have the universal identity

$$\prod_{i=1}^{m} (a_i X + b_i) = \sum_{j=0}^{m} E_j X^j$$

(a restatement of (1)). After tensoring with R (that is, applying ξ), this becomes

$$\prod_{i=1}^{m} (\alpha_i X + \beta_i) = \sum_{j=0}^{m} p_j X^j = P$$

in $S[X]$. Hence P splits in $S[X]$ as a product of m degree-≤ 1 polynomials. $\square$

Proposition 2.2. *The ring S constructed in the above proof of Theorem 2.1 can be written explicitly as*

$$S \cong R[a_1, b_1, \ldots, a_m, b_m] \Big/ (E_0 - p_0, E_1 - p_1, \ldots, E_m - p_m),$$

where now the same formula for E_r is interpreted inside the polynomial ring $R[a_1, b_1, \ldots, a_m, b_m]$:

$$E_r = \sum_{\substack{I \subseteq [m];\\ |I| = r}} \prod_{i \in I} a_i \prod_{i \notin I} b_i.$$

Proof. Let $a'_1, b'_1, a'_2, b'_2, \ldots, a'_m, b'_m$ be $2m$ further indeterminates, and let

$$E'_r = \sum_{\substack{I \subseteq [m];\\ |I| = r}} \prod_{i \in I} a'_i \prod_{i \notin I} b'_i$$

over the ring $\mathbb{Z}$. Then, we claim that

$$B_m \cong C_m [a'_1, b'_1, \ldots, a'_m, b'_m] \Big/ (E'_0 - E_0, E'_1 - E_1, \ldots, E'_m - E_m) \tag{13}$$

as C_m-algebras. Indeed, Lemma 1.3 shows that C_m is the polynomial ring $\mathbb{Z}[E_0, E_1, \ldots, E_m]$; thus,

$$\begin{aligned}
& C_m [a'_1, b'_1, \ldots, a'_m, b'_m] \Big/ (E'_0 - E_0, E'_1 - E_1, \ldots, E'_m - E_m) \\
& \cong (\mathbb{Z}[E_0, E_1, \ldots, E_m]) [a'_1, b'_1, \ldots, a'_m, b'_m] \Big/ (E'_0 - E_0, E'_1 - E_1, \ldots, E'_m - E_m) \\
& \cong (\mathbb{Z} [a'_1, b'_1, \ldots, a'_m, b'_m]) [E_0, E_1, \ldots, E_m] \Big/ (E'_0 - E_0, E'_1 - E_1, \ldots, E'_m - E_m) \\
& \cong \mathbb{Z} [a'_1, b'_1, \ldots, a'_m, b'_m],
\end{aligned}$$

since any ring P and any elements $u_0, u_1, \ldots, u_m \in P$ satisfy the P-algebra isomorphism

$$P\left[E_0, E_1, \ldots, E_m\right] \Big/ \left(u_0 - E_0, u_1 - E_1, \ldots, u_m - E_m\right) \cong P$$

(because adjoining a variable to a ring P and then setting this variable equal to an element of P gives back P). Hence, we obtain an R-algebra isomorphism

$$\begin{aligned} & C_m\left[a_1', b_1', \ldots, a_m', b_m'\right] \Big/ \left(E_0' - E_0, E_1' - E_1, \ldots, E_m' - E_m\right) \\ & \cong \mathbb{Z}\left[a_1', b_1', \ldots, a_m', b_m'\right] \cong \mathbb{Z}\left[a_1, b_1, \ldots, a_m, b_m\right] = B_m. \end{aligned}$$

This isomorphism is furthermore a C_m-algebra isomorphism, since it is easily seen to send each E_i to E_i (via E_i'). This proves (13).

Now,

$$\begin{aligned} S &= R \otimes_{C_m} B_m \\ &\cong R \otimes_{C_m} \left(C_m\left[a_1', b_1', \ldots, a_m', b_m'\right] \Big/ \left(E_0' - E_0, E_1' - E_1, \ldots, E_m' - E_m\right)\right) \\ &\qquad \text{(by (13))} \\ &\cong \left(R \otimes_{C_m} C_m\left[a_1', b_1', \ldots, a_m', b_m'\right]\right) \Big/ \left(E_0' - E_0, E_1' - E_1, \ldots, E_m' - E_m\right) \\ &\qquad \text{(since quotienting by an ideal commutes with base change)} \\ &\cong R\left[a_1', b_1', \ldots, a_m', b_m'\right] \Big/ \left(E_0' - p_0, E_1' - p_1, \ldots, E_m' - p_m\right) \\ &\qquad \left(\begin{array}{c} \text{since } R \otimes_{C_m} C_m\left[a_1', b_1', \ldots, a_m', b_m'\right] \cong R\left[a_1', b_1', \ldots, a_m', b_m'\right] \text{ as } R\text{-algebras,} \\ \text{and the isomorphism sends each } E_j \text{ to } p_j \end{array}\right) \\ &\cong R[a_1, b_1, \ldots, a_m, b_m] \Big/ \left(E_0 - p_0, E_1 - p_1, \ldots, E_m - p_m\right). \end{aligned}$$

This proves the proposition. □

Note that Proposition 2.2 gives what is probably the simplest construction of S. The point of Theorem 1.2 is that the natural map $R \to S$ is injective for every specialization $E_i \mapsto p_i$.

Theorem 2.3 (Universal property of the splitting algebra)**.** *Let R, m, P and $p_0, \ldots, p_m$ be as in Theorem 2.1. Let*

$$S \cong R[a_1, b_1, \ldots, a_m, b_m] \Big/ \left(E_0 - p_0, E_1 - p_1, \ldots, E_m - p_m\right)$$

be the R-algebra of Proposition 2.2, and let $\alpha_i, \beta_i \in S$ denote the images of a_i, b_i in this algebra S.

For every R-algebra T, the map

$$\operatorname{Hom}_{R\text{-alg}}(S, T) \longrightarrow \left\{ (c_1, d_1, \ldots, c_m, d_m) \in T^{2m} \;\middle|\; P = \prod_{i=1}^{m} (c_i X + d_i) \text{ in } T[X] \right\}$$

that sends an R-algebra homomorphism $f : S \to T$ to

$$(f(\alpha_1), f(\beta_1), \ldots, f(\alpha_m), f(\beta_m))$$

is a bijection. Equivalently, S is universal for ordered factorizations of P into exactly m linear-or-constant factors.

Proof. Let T be an R-algebra. Since S is presented as a quotient of the polynomial R-algebra

$$R[a_1, b_1, \ldots, a_m, b_m],$$

an R-algebra homomorphism $f : S \to T$ is the same as a choice of images

$$c_i = f(\alpha_i), \qquad d_i = f(\beta_i) \qquad (1 \le i \le m)$$

for the variables a_i, b_i satisfying the defining relations of the quotient. These relations are

$$E_j(c_1, d_1, \ldots, c_m, d_m) = p_j \qquad (0 \le j \le m),$$

where

$$E_j(c_1, d_1, \ldots, c_m, d_m) = \sum_{\substack{I \subseteq [m]; \\ |I|=j}} \prod_{i \in I} c_i \prod_{i \notin I} d_i.$$

But these are exactly the coefficient identities obtained by expanding

$$\prod_{i=1}^{m} (c_i X + d_i) = \sum_{j=0}^{m} E_j(c_1, d_1, \ldots, c_m, d_m) X^j.$$

Thus the defining relations hold if and only if

$$\prod_{i=1}^{m} (c_i X + d_i) = \sum_{j=0}^{m} p_j X^j = P \qquad \text{in } T[X].$$

This proves that the displayed map is well-defined and surjective.

It remains only to note uniqueness. Once the tuple $(c_1, d_1, \ldots, c_m, d_m)$ is specified, there is a unique R-algebra map

$$R[a_1, b_1, \ldots, a_m, b_m] \to T$$

sending $a_i \mapsto c_i$ and $b_i \mapsto d_i$. If the tuple gives a factorization of P, then the relations $E_j - p_j$ vanish under this map, so it factors uniquely through S. Hence two homomorphisms $S \to T$ with the same tuple are equal. This proves injectivity of the displayed map and completes the proof. □

A factorization with fewer than m factors can be regarded as one with exactly m factors by adjoining extra factors $0X + 1$.

References

[1] R. M. Adin, F. Brenti, and Y. Roichman, *Descent representations and multivariate statistics*, Transactions of the American Mathematical Society **357** (2005), no. 8, 3051–3082.

[2] E. E. Allen, *The Descent Monomials and a Basis for the Diagonally Symmetric Polynomials*, Journal of Algebraic Combinatorics **3** (1994), 5–16.

[3] Siegfried Bosch, *Algebra: From the Viewpoint of Galois Theory*, Springer 2018.

[4] David A. Cox, John Little, Donal O'Shea, *Ideals, Varieties, and Algorithms*, Undergraduate Texts in Mathematics, 5th edition, Springer 2025.

[5] Adriano M. Garsia, *Combinatorial Methods in the Theory of Cohen–Macaulay Rings*, Advances in Mathematics **38** (1980), 229–266.

[6] A. M. Garsia and D. Stanton, *Group actions on Stanley–Reisner rings and invariants of permutation groups*, Advances in Mathematics **51** (1984), no. 2, 107–201.

[7] T. Hibi, *Distributive lattices, affine semigroup rings and algebras with straightening laws*, in *Commutative Algebra and Combinatorics*, Advanced Studies in Pure Mathematics **11**, North-Holland, Amsterdam, 1987, pp. 93–109.

[8] R. P. Stanley, *Enumerative Combinatorics, Volume 1*, second edition, Cambridge Studies in Advanced Mathematics **49**, Cambridge University Press, 2012.

[9] R. P. Stanley, *Combinatorics and Commutative Algebra*, second edition, Progress in Mathematics **41**, Birkhäuser, Boston, 1996.