Discrete Faà di Bruno via Möbius Inversion

Heinrich Hartmann*

heinrichhartmann.com/math/2026-Moebius-Faa-di-Bruno

Abstract

We approach discrete and differential Faà di Bruno formulas from a Möbius inversion angle. On the Boolean cube, Newton's discrete Taylor formula and the definition of iterated forward differences form a zeta–Möbius dual pair, and composing two Taylor expansions and inverting once yields a closed discrete Faà di Bruno formula at a fixed basepoint: for arbitrary maps f, g between abelian groups,

$$\Delta(f \circ g;\, x;\, u_1, \ldots, u_k) = \sum_{H \in \mathrm{Cov}(k)} \Delta(f;\, g(x);\, (\Delta(g; x; u_T))_{T \in H}),$$

where $\mathrm{Cov}(k)$ denotes the coverings of $[k]$ by nonempty subsets. Grouping repeated directions gives binomial versions on multi-index grids, and iterating gives formulas for m-fold composites, with integer covering coefficients governed by explicit cross and level recursions, a discrete analogue of the Constantine–Savits formulas.

The relationship between coverings and partitions appearing in classical Faà di Bruno formulas, is exhibited in an algebraic setting. The discrete formulas are Taylor expansions over the function algebra of the Boolean cube, $B_k = \mathbf{k}[\delta_1, \ldots, \delta_k]/(\delta_i^2 - \delta_i)$, whose idempotent generators absorb overlapping products; in the differential analogue $A_k = \mathbf{k}[\varepsilon_1, \ldots, \varepsilon_k]/(\varepsilon_i^2)$, nilpotent generators annihilate overlaps and only partitions remain. Both algebras are fibers of the flat deformation $C_k = \mathbf{k}[t][x_1, \ldots, x_k]/(x_i^2 - tx_i)$, over which a single weighted covering formula interpolates: its coefficients are difference quotients, non-partition coverings carry positive powers of t, and evaluation at $t = 0$ yields the classical partition-indexed Faà di Bruno formula.

We demonstrate how these algebraic identities can be lifted to the analytical setting of C^n maps between Banach spaces, recovering the multivariate Faà di Bruno formula of Constantine–Savits and extending it to composites of several maps. Boolean finite differences, binomial grid formulas, infinitesimal Taylor algebras, and Fréchet derivatives thus appear as four realizations of one Möbius-dual Faà di Bruno formula, connected by a flat family.

Outline

1 Introduction **2**
1.1 Related work . . . 5
2 Möbius inversion **6**
3 Discrete Möbius calculus **7**
3.1 Taylor duality . . . 7
3.2 Faà di Bruno duality . . . 8
3.3 Iterated Faà di Bruno . . . 9
4 Polynomial Möbius calculus **12**
4.1 Taylor algebras . . . 12
4.2 Taylor–Möbius duality . . . 12
4.3 Faà di Bruno duality . . . 13
4.4 Iterated Faà di Bruno . . . 14
4.5 Faà di Bruno deformation . . . 15
5 Fréchet Möbius calculus **16**

*Hartmann IT GmbH / heinrichhartmann.com/math / Heinrich@HeinrichHartmann.com

6 Applications **18**
6.1 Self-enumeration . . . 18
6.2 Degree bound for polynomial maps . . . 18
6.3 Newton–Mahler composition . . . 19
6.4 Discrete jet composition . . . 19
6.5 Product rules . . . 19
7 Acknowledgements **20**

1 Introduction

We approach discrete and differential calculus from the point of view of Möbius inversion. The motivating observation is elementary: the discrete Taylor formula is the Möbius inverse of the definition of iterated forward differences. If $g : X \to Y$ is a map between abelian groups, then the values $g(x + \sum_{i \in S} u_i)$ for $S \subseteq [k]$ form a Boolean cube, and the forward difference

$$\Delta(g; x; u_1, \ldots, u_k) = \sum_{S \subseteq [k]} (-1)^{k-|S|}\, g(x + \textstyle\sum_{i \in S} u_i)$$

is exactly the Möbius transform of this cube. Newton's finite Taylor formula is the inverse zeta expansion.

(1) Remark (Notation)**.** We prefer explicit arguments over subscript notation, writing $\Delta(g; x; u)$ instead of $\Delta_u g(x)$. Semicolons and commas are both argument separators; we use semicolons as a visual hint when arguments are of different kinds. Dropping trailing arguments denotes currying: $\Delta(g; x)$ is the map $u \mapsto \Delta(g; x; u)$.

This observation extends without changing its nature from Boolean cubes to multi-index grids. Repeating directions replaces subsets by slot sets and gives a binomial form of Möbius inversion. Thus the discrete Taylor formula and its higher multi-index versions are two instances of the same zeta/Möbius duality.

The same principle carries over to the differential setting once one works over an infinitesimal Taylor algebra. The discrete calculus is algebra over the Boolean cube algebra $B_k = \Bbbk[\delta_1, \ldots, \delta_k]/(\delta_i^2 - \delta_i)$, whose spectrum is the cube $\{0, 1\}^k$, or over the grid algebra $B_k^\nu = \Bbbk[\delta_1, \ldots, \delta_k]/((\delta_i)_{\nu_i+1})$ with falling-factorial relations, whose natural module basis is the Newton basis $(\delta_i)_j$. The infinitesimal analogues replace idempotent relations by nilpotent ones: the Taylor algebra $A_k = \Bbbk[\varepsilon_1, \ldots, \varepsilon_k]/(\varepsilon_i^2)$ and its truncation $A_k^\nu = \Bbbk[\varepsilon_1, \ldots, \varepsilon_k]/(\varepsilon_i^{\nu_i+1})$. We will see below that these are naturally expressed as two fibers of the same flat deformation.

Working in A_k^ν, the Möbius duality transfers exactly from finite grids to infinitesimal grids. Evaluating a polynomial map on the infinitesimal grid $x + \sum_i \beta_i \varepsilon_i v_i$ for $0 \le \beta \le \nu$ gives a Taylor expansion on the zeta side, and the same alternating sieve extracts the differential monomials: in A_k^γ,

$$D^\gamma(p; x; v_\bullet)\, \varepsilon^\gamma = \sum_{\beta \le \gamma} (-1)^{|\gamma - \beta|} \binom{\gamma}{\beta}\, p(x + \textstyle\sum_i \beta_i \varepsilon_i v_i).$$

In this sense, the finite-difference and differential Taylor formulas are parallel Möbius-dual constructions: finite grids give forward differences, infinitesimal grids give differentials.

This low-tech viewpoint becomes especially useful for composition. Our first result is a discrete Faà di Bruno formula indexed by coverings, closed at the original basepoint.

(2) Theorem (Discrete Faà di Bruno)**.** *Let X, Y, Z be abelian groups, $g : X \to Y$, $f : Y \to Z$ arbitrary maps, $x \in X$, and $u_1, \ldots, u_k \in X$. Then*

$$\Delta(f \circ g; x; u_1, \ldots, u_k) = \sum_{H \in \mathrm{Cov}(k)} \Delta(f; g(x); (\Delta(g; x; u_T))_{T \in H}),$$

where $\mathrm{Cov}(k)$ *denotes the set of coverings of* $[k]$ *by nonempty subsets.*

The formula is exact and integral, closed at the original basepoint $g(x)$, and uses only the original direction increments $\Delta(g; x; u_T)$. It follows by composing two discrete Taylor expansions and applying Boolean Möbius inversion: the zeta side of the composite is trivially the composite of the zeta sides, and the covering formula is its Möbius inverse.

For $k = 2$ there are five coverings of $\{1, 2\}$. Writing $y = g(x)$ and $g_T = \Delta(g; x; u_T)$,

$$\Delta(f \circ g; x; u_1, u_2) = \underbrace{\Delta(f; y; g_1, g_2) + \Delta(f; y; g_{12})}_{\text{partitions, weight 2}} + \underbrace{\Delta(f; y; g_1, g_{12}) + \Delta(f; y; g_2, g_{12})}_{\text{weight 3}} + \underbrace{\Delta(f; y; g_1, g_2, g_{12})}_{\text{weight 4}},$$

where the *weight* of a covering is the total size $\mathrm{wt}(H) = \sum_{T \in H} |T|$ of its blocks. Under the scaling $u_i = tv_i$, a term of weight w is of order t^w: dividing by t^2 and letting $t \to 0$, the three overlapping coverings vanish and the two partition terms become the classical second-order chain rule $Df \cdot D^2 g(v_1, v_2) + D^2 f(Dg\, v_1, Dg\, v_2)$. In general the covering counts $1, 5, 109, 32297, \ldots$ A003465 replace the Bell numbers $1, 2, 5, 15, \ldots$ A000110 of the smooth theory.

Traditionally, this identity is grouped differently. In the example above, every covering except the partition $\{\{1\}, \{2\}\}$ contains the block $\{1, 2\}$; the corresponding four terms constitute the finite Taylor expansion of a single difference of f at the shifted basepoint $y + g_1 + g_2$, and are commonly contracted into it:

$$\Delta(f \circ g;\, x;\, u_1, u_2) = \underbrace{\Delta(f;\, y;\, g_1, g_2)}_{\{\{1\},\{2\}\}} + \underbrace{\Delta(f;\, y + g_1 + g_2;\, g_{12})}_{\{\{1,2\}\}},$$

two terms instead of five, indexed by the partitions of $\{1, 2\}$. Duarte and Torres [DT12] show that such a contraction is always possible: the finite-difference chain rule can be written as a partition-indexed sum with recursively shifted basepoints and corrected directions, and this is the direction in which the theory has developed. The contracted formulas are compact and of high practical utility, but they are not closed forms, and the basepoint shifts break the symmetry in the direction arguments, which makes higher-composition questions for $f_m \circ \cdots \circ f_1$ very difficult to approach. We argue that the expanded fixed-basepoint covering form is the natural way to view the formula: it is closed and symmetric in the directions, it carries the weight grading that governs the collapse to the classical chain rule, and it generalizes directly.

The same construction iterates through composites of several maps and refines to multi-index grids, and both coefficient systems of the resulting duality are explicit.

(3) Theorem (Iterated Faà di Bruno duality)**.** *Let* $X_0 \xrightarrow{f_1} X_1 \xrightarrow{f_2} \cdots \xrightarrow{f_m} X_m$ *be arbitrary maps between abelian groups,* $x \in X_0$*, directions* $u_s \in X_0$ *for* $s \in S$*,* $\gamma \in \mathbb{N}_0^S$*, and* $z = (f_m \circ \cdots \circ f_1)(x)$*. Then*

$$\Delta(f_m \circ \cdots \circ f_1;\, x;\, u_\bullet) = \sum_{K \in \mathrm{Cov}_m(S)} \Delta^K(f_1, \ldots, f_m;\, x;\, u_\bullet),$$

$$\Delta(f_m \circ \cdots \circ f_1; x; u_\bullet^\gamma) = \sum_{\kappa \in \mathcal{M}_+^m(S)} \mathrm{Cov}_m(\gamma, \kappa)\, \Delta^\kappa(f_1, \ldots, f_m; x; u_\bullet),$$

$$(f_m \circ \cdots \circ f_1)(x + \textstyle\sum_s \gamma_s u_s) = z + \displaystyle\sum_{\kappa \in \mathcal{M}_+^m(S)} \mathrm{Pow}_m(\gamma, \kappa)\, \Delta^\kappa(f_1, \ldots, f_m; x; u_\bullet),$$

where $\mathrm{Cov}_m(S)$ *is the set of* m*-fold iterated coverings of* S*, the binomial sums range over the* m*-fold iterated multi-index sets* $\mathcal{M}_+^m(S) = \mathcal{M}_+(\mathcal{M}_+^{m-1}(S))$*, the iterated increments* Δ^K *and* Δ^κ *are formed by nesting inner differences as directions of outer differences, and* $\mathrm{Pow}_m(\gamma, \kappa)$ *and* $\mathrm{Cov}_m(\gamma, \kappa)$ *count* m*-fold iterated subsets, respectively coverings, of the slot set* $S(\gamma)$ *with profile* κ*. The two coefficient systems are Möbius inverses of each other in* γ*,*

$$\mathrm{Cov}_m(\gamma,\kappa)=\sum_{\beta\le\gamma}(-1)^{\mathrm{wt}(\gamma-\beta)}\binom{\gamma}{\beta}\,\mathrm{Pow}_m(\beta,\kappa),$$

and both satisfy explicit level recursions in m.

The coefficients are nonnegative integers, and no factorial denominators appear anywhere in the discrete theory: all identities hold over arbitrary abelian groups, in any characteristic.

On the differential side, the same Taylor-composition argument produces the classical partition-indexed Faà di Bruno formulas. The reason coverings become partitions is visible already in the multiplication laws of the cube algebras. Idempotent generators absorb overlaps, nilpotent generators annihilate them:

$$\delta^T\delta^U=\delta^{T\cup U}\quad\text{in } B_k,\qquad \varepsilon^T\varepsilon^U=\varepsilon^{T\cup U}\cdot[T\cap U=\emptyset]\quad\text{in } A_k.$$

A product of faces contributes to the top face exactly when the exponents cover $[k]$; in A_k only the disjoint coverings contribute. Covering logic becomes partition logic.

The two multiplication laws are fibers of one family. Over $\Bbbk[t]$, the deformation cube algebra

$$C_k=\Bbbk[t][x_1,\dots,x_k]/(x_i^2-tx_i),\qquad x^Sx^T=t^{|S\cap T|}\,x^{S\cup T},$$

is free as a $\Bbbk[t]$-module; its fiber at $t=0$ is the Taylor algebra A_k, and after inverting t the rescaling $x_i=t\delta_i$ identifies it with the Boolean cube algebra. Expanding $p(x+\sum_i v_ix_i)$ in the basis x^S produces coefficients $c_S(t)\in\Bbbk[t]$ that interpolate the two calculi:

$$t^{|S|}\,c_S(t)=\Delta(p;x;(tv_i)_{i\in S}),\qquad c_S(0)=D(p;x;v_S).$$

The two readings of $c_S(t)$ explain each other. Since C_k is free over $\Bbbk[t]$ with basis x^S, the coefficients $c_S(t)$ are polynomials in t from the outset. Evaluating the expansion at the vertices of the rescaled cube $x_i=t\delta_i$, $\delta\in\{0,1\}^k$, and applying Boolean Möbius inversion identifies $t^{|S|}c_S(t)$ with the forward difference $\Delta(p;x;(tv_i)_{i\in S})$. Combining the two statements: the forward difference is divisible by $t^{|S|}$ in $\Bbbk[t]$, the difference quotient $c_S(t)$ is itself a polynomial in t, and its value at $t=0$, computed in the fiber A_k, is the derivative. The passage from difference quotients to derivatives is thus a divisibility statement about polynomials in t, rather than an analytic limit. Composing two such expansions yields a single Faà di Bruno identity over $\Bbbk[t]$, indexed by coverings and weighted by overlap:

$$[x^{[k]}]\,(f\circ g)(x+\textstyle\sum_i v_ix_i)=\displaystyle\sum_{H\in\mathrm{Cov}(k)}t^{\,\mathrm{wt}(H)-k}\,\mathrm{FdB}_H(t),$$

where $\mathrm{FdB}_H(t)$ is the covering term formed from difference quotients. At $t=1$ this is the discrete covering formula; at $t=0$ every overlapping covering vanishes and the classical partition formula remains, as in the $k=2$ example above. The covering and partition Faà di Bruno formulas are one identity, evaluated at two fibers. The grid algebras deform in the same way, via $C_k^\nu=\Bbbk[t][x_1,\dots,x_k]/(\prod_{j=0}^{\nu_i}(x_i-jt):1\le i\le k)$ with its Newton basis: in this family, Newton interpolation degenerates literally to Taylor expansion.

For polynomial maps $q:\Bbbk^e\to\Bbbk^d$, $p:\Bbbk^d\to\Bbbk$, the special fiber gives the classical partition form [Lev06]

$$D(p\circ q;x;v_{[k]})=\sum_{\pi\in\mathrm{Part}([k])}D(p;q(x);(D(q;x;v_B))_{B\in\pi}).$$

Grouping repeated directions gives the Constantine–Savits multi-index formula [CS96]. Iterating the construction gives an iterated Constantine–Savits formula for m-fold composites with recursive partition coefficients $\mathrm{Part}_m(\gamma,\kappa)$.

These polynomial identities lift to the Banach setting. For C^n maps between Banach spaces, the Taylor jet of a composite agrees up to order n with the composite of the individual Taylor jets:

$$T^n(f_m \circ \cdots \circ f_1;\, x) = \pi_{\leq n}(T^n(f_m;\, x_{m-1}) \circ \cdots \circ T^n(f_1;\, x_0)).$$

Since the partition Faà di Bruno formulas are polynomial coefficient identities, they apply directly to these Taylor jets and therefore to any C^n composite. In the polynomial setting, the deformation is precisely the passage to the limit of difference quotients: scaling the increments by t, the non-partition terms of the discrete covering formula carry positive powers of t and vanish as $t \to 0$, while the partition terms specialize to their smooth counterparts (35).

Putting things together, we obtain the following generalization of Constantine–Savits to Banach spaces and iterated multi-indices.

(4) Theorem (Fréchet iterated Faà di Bruno / recursive Constantine–Savits)**.** *Let $X_0, \ldots, X_m$ be Banach spaces, let $f_r : X_{r-1} \to X_r$ be C^n near x_{r-1}, and put $x_0 = x$, $x_r = f_r(x_{r-1})$. Fix directions $v_1, \ldots, v_k \in X_0$ and let $\gamma \in \mathbb{N}_0^k$ with $1 \leq |\gamma| \leq n$. Then*

$$D^\gamma(f_m \circ \cdots \circ f_1;\, x;\, v_\bullet) = \sum_{\kappa \in \mathcal{M}_+^m(k)} \mathrm{Part}_m(\gamma, \kappa)\, D^\kappa(f_1, \ldots, f_m;\, x;\, v_\bullet),$$

where $\mathrm{Part}_m(\gamma, \kappa)$ counts m-fold partitions of the slot set $S(\gamma)$ with profile κ. Writing $\mathrm{leaf}(\lambda) \in \mathbb{N}_0^k$ for the leaf multi-index of λ, obtained by recursively summing the leaves of λ with multiplicity, the coefficients vanish unless $\sum_\lambda \kappa(\lambda)\, \mathrm{leaf}(\lambda) = \gamma$, and satisfy the closed recursion $\mathrm{Part}_1(\gamma, \alpha) = \delta_{\gamma,\alpha}$ and, for $m \geq 2$,

$$\mathrm{Part}_m(\gamma, \kappa) = \frac{\gamma!}{\kappa!} \prod_{\substack{\beta \in \mathcal{M}_+^{m-1}(S) \\ \alpha = \mathrm{leaf}(\beta)}} \left(\frac{1}{\alpha!} \cdot \mathrm{Part}_{m-1}(\alpha, \beta) \right)^{\kappa(\beta)}.$$

For $m = 2$ this is the Constantine–Savits coefficient $\mathrm{Part}_2(\gamma, \kappa) = \gamma!/(\kappa! \prod_\beta (\beta!)^{\kappa(\beta)})$.

A final section collects five short applications of the discrete calculus, among them discrete and Fréchet product rules, governed by ordered coverings and ordered partitions respectively, and an easy proof that polynomial maps between abelian groups are closed under composition, with $\deg(f \circ g) \leq \deg f \cdot \deg g$. The paper is accompanied by a software package, available on GitHub, containing symbolic expansions and validations of the covering formula (degree 4, ca. $2 \cdot 10^6$ terms) and of the Duarte–Torres recursion (degree 10, ca. 115,000 terms).

Taken together: Boolean finite differences, binomial grid formulas, infinitesimal Taylor algebras, and Fréchet derivatives are four fibers of one Möbius-dual Taylor calculus.

1.1 Related work

The closest discrete predecessor is the formula of Duarte and Torres [DT12], and the present work grew out of an attempt to understand their dense and insightful paper properly, and to repackage its results in a more functorial form. Their discrete Faà di Bruno formula is indexed by partitions and uses recursively shifted basepoints and corrected directions; the covering formula in this paper writes all terms at the original basepoint $g(x)$ and is indexed by coverings of the direction set. The covering formula is fully expanded at the original basepoint $g(x)$; the Duarte–Torres formula is best understood as a grouping of the same terms by shifted basepoint, collecting all covering contributions that share a common recursively shifted evaluation point into a single partition-indexed term. Thus the two approaches describe the same finite-difference chain rule in different coordinates: Duarte–Torres gives a compact recursive partition formula with shifted basepoints and corrected directions, whereas the present paper gives a closed fixed-basepoint covering formula.

A second neighboring literature is the divided-difference form of Faà di Bruno. Floater and Lyche [FL07] derive chain rules for divided differences $[t_0, \ldots, t_n]$ of univariate functions, organized by the linear order of interpolation nodes; related divided-difference forms were obtained by Wang and Wang [WW06] and extended by Xu and Wang [XW10]. Our finite-difference formula is cubical rather than linearly ordered:

it concerns $\Delta(g; x; u_1, \ldots, u_k)$ at one basepoint with independent directions, and its natural combinatorics is that of coverings of $[k]$. In the univariate equal-step case, normalized forward differences are divided differences on a uniform grid, so the two settings meet after specialization. The deformation cube makes this connection explicit: the coefficients $c_S(t)$ of the family C_k are precisely difference quotients that extend through $t = 0$ to derivatives. A general calculus built directly on such scaled difference-quotient maps has been developed by Bertram, Glöckner, and Neeb [BGN04]; our flat family can be read as a polynomial, higher-order companion to their first-order framework.

On the smooth side, the multi-index Faà di Bruno formula of Constantine and Savits [CS96] is one of the main targets recovered here. Our infinitesimal Taylor algebra A_k^ν gives a direct Möbius-dual derivation of their coefficients from the square-free partition formula by grouping repeated directions; closely related bookkeeping between multi-indices and partitions of multisets appears in Hardy [Har06]. The slot-set construction $S(\gamma)$ provides a uniform way to pass between partition formulas and multi-index formulas. Iterated one-dimensional Faà di Bruno formulas have been studied through higher-order Bell polynomials by Natalini and Ricci [NR04]; the iterated formula in this paper gives a multivariate version for m-fold composites with recursively defined partition coefficients.

The finite-difference operators used here also belong to the older theory of cross-effects and polynomial maps. The alternating expression $\Delta(g; x; u_1, \ldots, u_k)$ is the k-fold deviation or cross-effect of a map between abelian groups, in the sense of Eilenberg and Mac Lane [EL54]. Vanishing of sufficiently high differences is the classical finite-difference characterization of polynomial operations, going back to Mazur and Orlicz [MO34] and appearing in later work such as Leibman's theory of polynomial mappings of groups [Lei02]; in the different setting of functors between abelian categories, Bauer, Johnson, Osborne, Riehl, and Tebbe [BJO+18] prove a categorified higher-order chain rule for cross-effects, where equalities are replaced by chain homotopy equivalences and similar combinatorics governs the resulting Faà di Bruno decomposition. The present paper works entirely at the level of maps between abelian groups, where the covering Faà di Bruno formula is an explicit, exact chain rule for cross-effects; as a corollary, polynomial maps are closed under composition with $\deg(f \circ g) \leq \deg f \cdot \deg g$ (42).

Finally, Faà di Bruno formulas have a rich algebraic life in coalgebras, Hopf algebras, and incidence algebras, beginning with work of Joni and Rota [JR79] and surveyed from several viewpoints by Frabetti and Manchon [FM14]. The present paper does not develop the Hopf-algebraic formalism. Instead, it uses ordinary Möbius inversion on Boolean and multi-index grids, together with the cube algebras B_k, A_k, and their deformation C_k.

2 Möbius inversion

Möbius inversion generalizes the elementary principle that individual terms of a sum can be extracted via alternating differences. In its general form, it applies to any locally finite partially ordered set (see [Sta11] for the general theory). In this text we need it for Boolean lattices of subsets, and for the componentwise partial order on multi-indices.

We write $[n] = \{1, \ldots, n\}$, $\mathcal{P}(S)$ for the power set of a finite set S, $\mathcal{P}_+(S) = \mathcal{P}(S) \setminus \{\emptyset\}$ for the set of nonempty subsets, and $\mathcal{P}(n)$ for $\mathcal{P}([n])$. Note $|\mathcal{P}(k)| = 2^k$. We think of maps $a : \mathcal{P}(S) \to X$ as *cubes* in X with $2^{|S|}$ vertices $a(T)$ for $T \subseteq S$ and legs $a(\{i\})$.

(5) Proposition (Boolean Möbius inversion)**.** *Let G be an abelian group. For a cube $a : \mathcal{P}(k) \to G$, set*

$$\zeta(a; S) := \sum_{T \subseteq S} a(T), \qquad \mu(a; S) := \sum_{T \subseteq S} (-1)^{|S|-|T|}\, a(T).$$

Then ζ and μ define inverse bijections on $\mathrm{Map}(\mathcal{P}(k), G)$.

Proof. Exchange summation order and use $(1-1)^{|S|} = \sum_{T \subseteq S} (-1)^{|S|-|T|} = \delta_{S,\emptyset}$.

□

(6) Definition (Multi-indices and slot sets)**.** • A *multi-index* on a finite set S is a map $\alpha : S \to \mathbb{N}_0$, i.e. an element of $\mathbb{N}_0^S$. The componentwise partial order is $\beta \le \alpha$ iff $\beta_s \le \alpha_s$ for all s.

- The *weight* is $|\alpha| = \mathrm{wt}(\alpha) = \sum_{s\in S} \alpha_s$, the *height* is $\mathrm{ht}(\alpha) = \max_{s\in S} \alpha_s$.
- The *factorial* is $\alpha! = \prod_{s\in S} \alpha_s!$, the *falling factorial* is $(\alpha)_\beta = \prod_{s\in S}(\alpha_s)_{\beta_s}$ where $(n)_r = n(n-1)\cdots(n-r+1)$.
- A *Boolean realization* of $\alpha \in \mathbb{N}_0^S$ is a set $S(\alpha) = \{(s,i) : s \in S,\ 0 < i \le \alpha_s\}$ with canonical projection $\pi : S(\alpha) \to S$.
- The *fiber measure* of a map $q : S' \to S$ of finite sets is $\nu(q) \in \mathbb{N}_0^S$, $\nu(q)_s = |q^{-1}(s)|$. In particular $\nu(\pi) = \alpha$ and $|S'| = \mathrm{wt}(\nu(q))$.
- We embed $\mathcal{P}(S) \hookrightarrow \mathbb{N}_0^S$ by $U \mapsto 1_U$; the image consists of multi-indices of height ≤ 1, and $S(1_U) \cong U$.

We think of maps $A : \mathbb{N}_0^S \to X$ as *grids* in X. The points $A(1_T)$ for $T \subseteq S$ form the coordinate cube; the remaining points $A(\alpha)$ for $\mathrm{ht}(\alpha) > 1$ fill the interior of the grid.

We define binomial Möbius inversion via the Boolean realization of multi-indices.

(7) Proposition (Binomial Möbius inversion)**.** *Let G be an abelian group. For a map $A : \mathbb{N}_0^S \to G$ and $\alpha \in \mathbb{N}_0^S$, let $A_\alpha = A \circ \nu : \mathcal{P}(S(\alpha)) \xrightarrow{\nu} \mathbb{N}_0^S \xrightarrow{A} G$. Set*

$$\zeta(A;\alpha) := \zeta(A_\alpha; S(\alpha)) = \sum_{\beta\le\alpha} \frac{(\alpha)_\beta}{\beta!}\, A(\beta), \qquad \mu(A;\alpha) := \mu(A_\alpha; S(\alpha)) = \sum_{\beta\le\alpha} (-1)^{\mathrm{wt}(\alpha-\beta)}\, \frac{(\alpha)_\beta}{\beta!}\, A(\beta).$$

Then ζ and μ define inverse bijections on $\mathrm{Map}(\mathbb{N}_0^S, G)$, extending the Boolean case.

Proof. The number of subsets $T \subseteq S(\alpha)$ with $\nu(T) = \beta$ is $\frac{(\alpha)_\beta}{\beta!}$, giving the explicit sums. For the inversion, observe that the cube $A_\alpha = A \circ \nu$ satisfies $(\zeta A)_\alpha = \zeta(A_\alpha)$: for $T \subseteq S(\alpha)$ with $\nu(T) = \beta$, the restriction of A_α to $\mathcal{P}(T)$ is the realization of A_β on T, hence $\zeta(A_\alpha; T) = \sum_{R\subseteq T} A(\nu(R)) = \zeta(A;\beta)$. Applying Boolean Möbius inversion (5) to A_α on $\mathcal{P}(S(\alpha))$: $\mu(\zeta A;\alpha) = \mu((\zeta A)_\alpha; S(\alpha)) = \mu(\zeta(A_\alpha); S(\alpha)) = A_\alpha(S(\alpha)) = A(\alpha)$, and symmetrically for $\zeta\mu$. Alternatively, exchange summation order, use $\binom{\alpha}{\beta}\binom{\beta}{\gamma} = \binom{\alpha}{\gamma}\binom{\alpha-\gamma}{\beta-\gamma}$ to decouple, and apply the multi-index binomial theorem $(1-1)^\alpha = \sum_{\beta\le\alpha}(-1)^{\mathrm{wt}(\beta)}\binom{\alpha}{\beta} = \delta_{\alpha,0}$. □

3 Discrete Möbius calculus

This section develops the composition calculus of forward differences for arbitrary maps between abelian groups: Taylor duality, the covering Faà di Bruno formula, and their iterated and binomial refinements with the coefficient systems Pow_m and Cov_m. All identities are exact with integer coefficients. Every formula evaluates maps only at the points $x + \sum_{i\in S} u_i$, so the domain may in fact be any commutative monoid; we state results for abelian groups and use monoid domains such as $\mathbb{N}_0$ in the applications.

3.1 Taylor duality

(8) Definition (Forward differences and translations)**.** Let X, Y be abelian groups, $g : X \to Y$, $x \in X$, and directions $u_\bullet = (u_1, \dots, u_k) \in X^k$. The *iterated forward difference* and the *iterated translation* are:

$$\Delta(g;x;u_\bullet) := \sum_{S\subseteq[k]} (-1)^{k-|S|}\, g(x + \textstyle\sum_{i\in S} u_i), \qquad T(g;x;u_\bullet) := g(x + \textstyle\sum_{i=1}^k u_i).$$

For $S \subseteq [k]$, write $u_S = (u_i)_{i\in S}$. As Δ is invariant under permutation of directions, we can define:

$$\Delta^S(g;x;u_\bullet) := \Delta(g;x;u_S), \qquad T^S(g;x;u_\bullet) := g(x + \textstyle\sum_{i\in S} u_i).$$

For $\alpha \in \mathbb{N}_0^k$, write $u_\bullet^\alpha = (u_1^{\times\alpha_1}, \ldots, u_k^{\times\alpha_k})$ for the direction list with u_i repeated α_i times. Then:

$$\Delta^\alpha(g; x; u_\bullet) := \Delta(g; x; u_\bullet^\alpha) = \sum_{\beta\leq\alpha} (-1)^{|\alpha|-|\beta|} \frac{(\alpha)_\beta}{\beta!} g(x + \textstyle\sum_i \beta_i u_i), \qquad T^\alpha(g; x; u_\bullet) := g(x + \textstyle\sum_{i=1}^k \alpha_i u_i),$$

where we grouped the $2^{|\alpha|}$ Boolean terms of $\Delta(g; x; u_\bullet^\alpha)$ by profile.

(9) Proposition (Taylor duality)**.** *Let X, Y be abelian groups, $g : X \to Y$, $x \in X$, $u_\bullet = (u_1, \ldots, u_k) \in X^k$, and $y = g(x)$.*

1) Boolean Taylor duality.

$$g(x + \textstyle\sum_{i=1}^k u_i) = y + \displaystyle\sum_{\emptyset\neq S\subseteq[k]} \Delta(g; x; u_S), \qquad \Delta(g; x; u_\bullet) = \sum_{S\subseteq[k]} (-1)^{k-|S|} g(x + \textstyle\sum_{i\in S} u_i).$$

2) Binomial Taylor duality. *For $\alpha \in \mathbb{N}_0^k$:*

$$g(x + \textstyle\sum_{i=1}^k \alpha_i u_i) = y + \displaystyle\sum_{0<\beta\leq\alpha} \frac{(\alpha)_\beta}{\beta!} \Delta(g; x; u_\bullet^\beta), \qquad \Delta(g; x; u_\bullet^\alpha) = \sum_{\beta\leq\alpha} (-1)^{|\alpha|-|\beta|} \frac{(\alpha)_\beta}{\beta!} g(x + \textstyle\sum_{i=1}^k \beta_i u_i).$$

Proof. 1) The second identity is the definition of Δ. The first is its Boolean Möbius inverse (5).

2) The second identity is the grouped form of Δ^α (8). The first is its Binomial Möbius inverse (7).

□

3.2 Faà di Bruno duality

Write $\mathcal{P}_+(k)$ for the set of nonempty subsets of $[k]$, $\mathcal{P}_+^2(k) = \mathcal{P}_+(\mathcal{P}_+(k))$ for the set of nonempty families of such subsets, and $\mathrm{Cov}(k) = \{H \in \mathcal{P}_+^2(k) \mid \bigcup H = [k]\}$ for the set of all coverings of $[k]$ by nonempty subsets.

(10) Theorem (Discrete Faà di Bruno Duality)**.** *Let X, Y, Z be abelian groups, $g : X \to Y$, $f : Y \to Z$ arbitrary maps, $x \in X$, $u_1, \ldots, u_k \in X$. Write $y = g(x)$ and $z = f(y)$. Then:*

1) Boolean Faà di Bruno *($k \geq 1$)*:

$$\Delta(f \circ g; x; u_\bullet) = \sum_{H\in\mathrm{Cov}(k)} \Delta(f; y; (\Delta(g; x; u_T))_{T\in H}).$$

2) Boolean Taylor composition:

$$(f \circ g)(x + \textstyle\sum_{i=1}^k u_i) = z + \displaystyle\sum_{H\in\mathcal{P}_+^2(k)} \Delta(f; y; (\Delta(g; x; u_T))_{T\in H}).$$

Proof. 2) Write $g_T := \Delta(g; x; u_T)$. By Taylor duality (9), $T(g; x; u_S) = y + \sum_{\emptyset\neq T\subseteq S} g_T$. Applying Taylor duality to f at y in directions $(g_T)_{T\in\mathcal{P}_+(S)}$ gives $T(f \circ g; x; u_S) = z + \sum_{H\in\mathcal{P}_+^2(S)} \Delta(f; y; (g_T)_{T\in H})$.

3) Define $\varphi(S) := \sum_{H\in\mathrm{Cov}(S)} \Delta(f; y; (g_T)_{T\in H})$. By Möbius inversion (5), it suffices to show $\zeta(\varphi; S) = T(f\circ g; x; u_S) - z$. Indeed, $\zeta(\varphi; S) = \sum_{R\subseteq S} \sum_{H\in\mathrm{Cov}(R)} \Delta(f; y; (g_T)_{T\in H}) = \sum_{H\in\mathcal{P}_+^2(S)} \Delta(f; y; (g_T)_{T\in H})$, since each $H \in \mathcal{P}_+^2(S)$ appears in exactly one summand, namely the one for $R = \bigcup H$. Now apply part 2.

□

(11) Remark (Binomial Faà di Bruno)**.** The Boolean Faà di Bruno formula extends naturally to a binomial version in terms of multi-indices $\gamma \in \mathbb{N}_0^S$, by applying the Boolean formula to the slot set $S(\gamma)$ and grouping by profile. The resulting expression resembles the Constantine–Savits form [CS96]; the precise coefficients and a recursive version are given in (18) below.

3.3 Iterated Faà di Bruno

The Boolean Faà di Bruno formula extends to m-fold compositions, indexed by $\mathrm{Cov}_m(k)$. We first set up the combinatorial apparatus: higher power sets, higher multi-indices, and higher partitions.

(12) Definition (Higher power sets and coverings)**.** Recall that $\mathcal{P}(S)$ is the power set and $\mathcal{P}_+(S) = \mathcal{P}(S) \setminus \{\emptyset\}$ the set of nonempty subsets.

- The *higher power sets* are $\mathcal{P}_+^0(S) = S$, $\mathcal{P}_+^1(S) = \mathcal{P}_+(S)$, and $\mathcal{P}_+^r(S) := \mathcal{P}_+(\mathcal{P}_+^{r-1}(S))$ for $r \geq 2$.
- For $K \in \mathcal{P}_+^r(S)$, the *leaf support* $\mathrm{leaf}(K) \subseteq S$ is defined recursively: $\mathrm{leaf}(K) := K$ for $r = 1$ and $\mathrm{leaf}(K) := \bigcup_{L \in K} \mathrm{leaf}(L)$ for $r \geq 2$.
- We say K *covers* S if $\mathrm{leaf}(K) = S$, and write $\mathrm{Cov}_r(S) := \{K \in \mathcal{P}_+^r(S) \mid \mathrm{leaf}(K) = S\}$ for the set of *r-fold coverings.*

For $S = [k]$, we write $\mathcal{P}_+^r(k)$ and $\mathrm{Cov}_r(k)$.

(13) Definition (Higher multi-indices)**.** For a set S, let $\mathcal{M}(S)$ be the set of finitely supported maps $S \to \mathbb{N}_0$, and $\mathcal{M}_+(S) = \mathcal{M}(S) \setminus \{0\}$. For finite S, $\mathcal{M}(S) = \mathbb{N}_0^S$.

- The *iterated multiset sets* are $\mathcal{M}_+^0(S) = S$, $\mathcal{M}_+^1(S) = \mathcal{M}_+(S)$, and $\mathcal{M}_+^{r+1}(S) := \mathcal{M}_+(\mathcal{M}_+^r(S))$ for $r \geq 1$.
- For $\kappa \in \mathcal{M}_+^r(S)$ with $r \geq 1$, the *support* is $\mathrm{Supp}(\kappa) := \{\lambda \in \mathcal{M}_+^{r-1}(S) : \kappa(\lambda) > 0\}$.
- The *leaf multi-index* $\mathrm{leaf}(\kappa) \in \mathbb{N}_0^S$ is defined recursively: $\mathrm{leaf}(\alpha) := \alpha$ for $r = 1$ and $\mathrm{leaf}(\kappa) := \sum_{\lambda \in \mathrm{Supp}(\kappa)} \kappa(\lambda)\, \mathrm{leaf}(\lambda)$ for $r \geq 2$.

The *higher profile map* $\nu : \mathcal{P}_+^m(S(\gamma)) \to \mathcal{M}_+^m(S)$ is defined recursively: $\nu(T) \in \mathbb{N}_0^S$ for $m = 1$ as in (6), and $\nu(K)(\lambda) := \#\{L \in K : \nu(L) = \lambda\}$ for $m \geq 2$.

The embedding $\mathcal{P}_+(S) \hookrightarrow \mathcal{M}_+(S)$ via $T \mapsto 1_T$ extends to $\mathcal{P}_+^r(S) \hookrightarrow \mathcal{M}_+^r(S)$ at every level: the Boolean case is the height-one restriction of the multi-index iteration.

(14) Definition (Higher partitions)**.**
- A *partition* of a finite set S is a set $\pi = \{B_1, \ldots, B_r\}$ of nonempty pairwise disjoint subsets with $B_1 \sqcup \cdots \sqcup B_r = S$. Write $\mathrm{Part}(S)$ for the set of partitions and $\mathrm{Part}(k)$ for $\mathrm{Part}([k])$.
- *Higher partitions* are defined recursively: $\mathrm{Part}_1(S) = \{S\}$, and for $m \geq 1$, $\mathrm{Part}_{m+1}(S) := \{\{H_B\}_{B \in \pi} \mid \pi \in \mathrm{Part}(S),\ H_B \in \mathrm{Part}_m(B)\}$. For $m = 2$, $\mathrm{Part}_2(S) = \mathrm{Part}(S)$.
- The *weight* of $H \in \mathcal{P}_+^r(S)$ is $\mathrm{wt}(H) := |H|$ for $r = 1$ and $\mathrm{wt}(H) := \sum_{K \in H} \mathrm{wt}(K)$ for $r \geq 2$.
- A *multi-index partition* of $\gamma \in \mathbb{N}_0^S$ is an iterated multi-index $\kappa \in \mathcal{M}_+^m(S)$ with $\mathrm{leaf}(\kappa) = \gamma$; we write $\kappa \vdash \gamma$. For $m = 2$ this is a multiset κ of nonzero multi-indices with $\sum_\beta \kappa(\beta)\, \beta = \gamma$.

(15) Lemma (Weight bound)**.** *For any covering $H \in \mathrm{Cov}_r(S)$ with $r \geq 1$, $\mathrm{wt}(H) \geq |S|$, with equality if and only if $H \in \mathrm{Part}_r(S)$.*

Proof. For $r = 1$: $\mathrm{Cov}_1(S) = \mathrm{Part}_1(S) = \{S\}$ and $\mathrm{wt}(S) = |S|$. For $r = 2$: $\mathrm{wt}(H) = \sum_{T \in H} |T| \geq |\bigcup H| = |S|$, with equality iff the blocks are pairwise disjoint. For $r \geq 3$: by induction, $\mathrm{wt}(H) = \sum_{K \in H} \mathrm{wt}(K) \geq \sum_{K \in H} |\mathrm{leaf}(K)| \geq |S|$, with equality iff each $K \in \mathrm{Part}_{r-1}(\mathrm{leaf}(K))$ and the leaf supports partition S.

□

(16) Definition (Iterated increments)**.** Let $X_0 \xrightarrow{f_1} X_1 \xrightarrow{f_2} \cdots \xrightarrow{f_m} X_m$ be maps between abelian groups, $x \in X_0$, $u : S \to X_0$ a direction family, and $x_r = (f_r \circ \cdots \circ f_1)(x)$.

1) *Boolean iterated increments.* For $r = 1$ and $T \in \mathcal{P}_+(S)$, set $\Delta^T(f_1; x; u) := \Delta(f_1; x; u_T)$. For $r \geq 2$ and $K \in \mathcal{P}_+^r(S)$, define recursively:

$$\Delta^K(f_1,\dots,f_r;\,x;\,u) := \Delta(f_r;\,x_{r-1};\,(\Delta^L(f_1,\dots,f_{r-1};\,x;\,u))_{L\in K}).$$

2) *Binomial iterated increments.* For $r=1$ and $\alpha\in\mathcal{M}_+(S)$, set $\Delta^\alpha(f_1;\,x;\,u) := \Delta(f_1;\,x;\,u^\alpha)$ (8). For $r\ge 2$ and $\kappa\in\mathcal{M}_+^r(S)$, define recursively:

$$\Delta^\kappa(f_1,\dots,f_r;\,x;\,u) := \Delta(f_r;\,x_{r-1};\,(\Delta^\lambda(f_1,\dots,f_{r-1};\,x;\,u)^{\times\kappa(\lambda)})_{\lambda\in\mathrm{Supp}(\kappa)}),$$

where for each λ with $\kappa(\lambda)>0$, the direction $\Delta^\lambda(\dots)$ appears $\kappa(\lambda)$ times. The Boolean case is recovered by restriction to $\mathcal{P}_+^r(S)\subset\mathcal{M}_+^r(S)$.

(17) Lemma (Profile invariance)**.** *Let $\gamma\in\mathbb{N}_0^S$, let $u:S\to X_0$ be a direction family, and equip $S(\gamma)$ with the slot directions $u\circ\pi$. Then for every $K\in\mathcal{P}_+^m(S(\gamma))$,*

$$\Delta^K(f_1,\dots,f_m;\,x;\,u\circ\pi) = \Delta^{\nu(K)}(f_1,\dots,f_m;\,x;\,u).$$

Proof. For $m=1$ and $T\subseteq S(\gamma)$, the direction list $(u_{\pi(t)})_{t\in T}$ is a permutation of $u^{\nu(T)}$, so the claim is permutation invariance of Δ (8). For $m\ge 2$, apply the induction hypothesis to each $L\in K$: $\Delta^L=\Delta^{\nu(L)}$. In the outer difference, $\kappa(\lambda)$ of the directions equal Δ^λ, which is the defining expression for $\Delta^{\nu(K)}$.

□

(18) Theorem (Iterated Faà di Bruno)**.** *Let $X_0 \xrightarrow{f_1} X_1 \xrightarrow{f_2} \cdots \xrightarrow{f_m} X_m$ be arbitrary maps between abelian groups, $x\in X_0$, and $z=(f_m\circ\cdots\circ f_1)(x)$. Let S be a finite set, $u:S\to X_0$ a family of directions, and $\gamma\in\mathbb{N}_0^S$; write $u_S=(u_s)_{s\in S}$, and u^γ for the family with each u_s repeated γ_s times. For $S=[k]$ this recovers the tuple notation $u_\bullet=(u_1,\dots,u_k)$.*

1) Boolean Faà di Bruno *($S\neq\emptyset$)*:

$$\Delta(f_m\circ\cdots\circ f_1;\,x;\,u_S) = \sum_{K\in\mathrm{Cov}_m(S)} \Delta^K(f_1,\dots,f_m;\,x;\,u).$$

2) Binomial Faà di Bruno:

$$\Delta(f_m\circ\cdots\circ f_1;x;u^\gamma) = \sum_{\kappa\in\mathcal{M}_+^m(S)} \mathrm{Cov}_m(\gamma,\kappa)\,\Delta^\kappa(f_1,\dots,f_m;x;u),$$

where $\mathrm{Cov}_m(\gamma,\kappa)=\#\{K\in\mathrm{Cov}_m(S(\gamma)):\nu(K)=\kappa\}$ *counts m-fold coverings of $S(\gamma)$ with profile κ.*

3) Boolean Taylor composition:

$$(f_m\circ\cdots\circ f_1)(x+\textstyle\sum_{s\in S}u_s) = z + \displaystyle\sum_{K\in\mathcal{P}_+^m(S)} \Delta^K(f_1,\dots,f_m;\,x;\,u).$$

4) Binomial Taylor composition:

$$(f_m\circ\cdots\circ f_1)(x+\textstyle\sum_s\gamma_s u_s) = z + \displaystyle\sum_{\kappa\in\mathcal{M}_+^m(S)} \mathrm{Pow}_m(\gamma,\kappa)\,\Delta^\kappa(f_1,\dots,f_m;x;u),$$

where $\mathrm{Pow}_m(\gamma,\kappa)=\#\{K\in\mathcal{P}_+^m(S(\gamma)):\nu(K)=\kappa\}$ *counts m-fold iterated subsets of $S(\gamma)$ with profile κ.*

Proof. 3) By induction on m. The cases $m=1$ and $m=2$ are (9) and (10). Write $F=f_m\circ\cdots\circ f_1$, $F'=f_{m-1}\circ\cdots\circ f_1$, $x_r=(f_r\circ\cdots\circ f_1)(x)$, and $d_K:=\Delta^K(f_1,\dots,f_{m-1};\,x;\,u)$ for $K\in\mathcal{P}_+^{m-1}(S)$. The induction hypothesis gives $F'(x+\sum_{s\in S}u_s)=x_{m-1}+\sum_{K\in\mathcal{P}_+^{m-1}(S)}d_K$. Applying Taylor duality (9) to f_m at x_{m-1} in directions d_K:

$$F(x+\textstyle\sum_{s\in S} u_s) = x_m + \displaystyle\sum_{H\in\mathcal{P}_+(\mathcal{P}_+^{m-1}(S))} \Delta(f_m;\, x_{m-1};\, (d_K)_{K\in H}).$$

Since $\mathcal{P}_+(\mathcal{P}_+^{m-1}(S)) = \mathcal{P}_+^m(S)$ and $\Delta(f_m;\, x_{m-1};\, (d_K)_{K\in H}) = \Delta^H(f_1,\dots,f_m;\, x;\, u)$ by (16), this completes the induction.

1) Define $\varphi(R) := \sum_{H\in\mathrm{Cov}_m(R)} \Delta^H$ for $R\subseteq S$. By Möbius inversion, it suffices to show $\zeta(\varphi; R) = T(F; x; u_R) - z$. Indeed, $\zeta(\varphi; R) = \sum_{Q\subseteq R}\sum_{H\in\mathrm{Cov}_m(Q)} \Delta^H = \sum_{H\in\mathcal{P}_+^m(R)} \Delta^H$, since each $H\in\mathcal{P}_+^m(R)$ covers exactly $Q = \mathrm{leaf}(H)$. Now apply part 3.

2,4) Apply parts 1,3 to $S(\gamma)$. By profile invariance (17), Δ^K depends only on the profile $\kappa = \nu(K)$, so grouping gives $\sum_\kappa \mathrm{Cov}_m(\gamma,\kappa)\,\Delta^\kappa$ for the FdB side and $\sum_\kappa \mathrm{Pow}_m(\gamma,\kappa)\,\Delta^\kappa$ for the Taylor side.

□

(19) Proposition (Newton coefficient recursion)**.** *For $\alpha\in\mathcal{M}_+(S)$ and $\kappa\in\mathcal{M}_+^m(S)$ with $m\geq 2$:*

$$\mathrm{Pow}_1(\gamma,\alpha) = \frac{(\gamma)_\alpha}{\alpha!}, \qquad \mathrm{Pow}_2(\gamma,\kappa) = \frac{1}{\kappa!}\prod_{\alpha\in\mathrm{Supp}(\kappa)} \left(\frac{(\gamma)_\alpha}{\alpha!}\right)_{\kappa(\alpha)},$$

$$\mathrm{Pow}_m(\gamma,\kappa) = \frac{1}{\kappa!}\prod_{\lambda\in\mathrm{Supp}(\kappa)} (\mathrm{Pow}_{m-1}(\gamma,\lambda))_{\kappa(\lambda)}.$$

Proof. For $m=1$: $\mathrm{Pow}_1(\gamma,\alpha) = \#\{T\subseteq S(\gamma) : \nu(T)=\alpha\} = \frac{(\gamma)_\alpha}{\alpha!}$, counting subsets with profile α.

For $m\geq 2$: an element $K\in\mathcal{P}_+^m(S(\gamma))$ is a nonempty set of elements of $\mathcal{P}_+^{m-1}(S(\gamma))$. The profile $\kappa\in\mathcal{M}_+^m(S)$ records $\kappa(\lambda) = \#\{\text{elements of } K \text{ with profile } \lambda\}$. Since the elements are chosen as a set (no repeats) from $\mathrm{Pow}_{m-1}(\gamma,\lambda)$ available elements with profile λ, the count is $\prod_\lambda (\mathrm{Pow}_{m-1}(\gamma,\lambda))_{\kappa(\lambda)}/\kappa(\lambda)!$.

□

(20) Proposition (Covering coefficient recursion)**.**

$$\mathrm{Cov}_1(\gamma,\alpha) = \delta_{\gamma,\alpha}, \qquad \mathrm{Cov}_2(\gamma,\kappa) = \sum_{\beta\leq\gamma} (-1)^{\mathrm{wt}(\gamma-\beta)} \frac{(\gamma)_\beta}{\beta!}\frac{1}{\kappa!}\prod_{\alpha\in\mathrm{Supp}(\kappa)} \left(\frac{(\beta)_\alpha}{\alpha!}\right)_{\kappa(\alpha)},$$

and generally, for $m\geq 2$:

1) Cross recursion *(Cov_m from Pow_m):*

$$\mathrm{Cov}_m(\gamma,\kappa) = \sum_{\beta\leq\gamma} (-1)^{\mathrm{wt}(\gamma-\beta)} \frac{(\gamma)_\beta}{\beta!}\,\mathrm{Pow}_m(\beta,\kappa).$$

2) Level recursion *(Cov_m from Cov_{m-1}):*

$$\mathrm{Cov}_m(\gamma,\kappa) = \sum_{\beta\leq\gamma} (-1)^{\mathrm{wt}(\gamma-\beta)} \frac{(\gamma)_\beta}{\beta!}\frac{1}{\kappa!}\prod_{\lambda\in\mathrm{Supp}(\kappa)} \left(\sum_{\alpha\leq\beta}\frac{(\beta)_\alpha}{\alpha!}\,\mathrm{Cov}_{m-1}(\alpha,\lambda)\right)_{\kappa(\lambda)}.$$

Proof. 1) An element $K\in\mathcal{P}_+^m(S(\gamma))$ with profile κ has leaf support $\mathrm{leaf}(K)\subseteq S(\gamma)$. Each subset $R\subseteq S(\gamma)$ with $\nu(R)=\beta$ supports $\mathrm{Cov}_m(\beta,\kappa)$ coverings with profile κ, and there are $\frac{(\gamma)_\beta}{\beta!}$ such subsets. Hence $\mathrm{Pow}_m(\gamma,\kappa) = \sum_{\beta\leq\gamma}\frac{(\gamma)_\beta}{\beta!}\mathrm{Cov}_m(\beta,\kappa)$. Binomial Möbius inversion (7) gives $\mathrm{Cov}_m(\gamma,\kappa) = \sum_{\beta\leq\gamma}(-1)^{\mathrm{wt}(\gamma-\beta)}\frac{(\gamma)_\beta}{\beta!}\mathrm{Pow}_m(\beta,\kappa)$.

2) Substitute the Newton recursion (19) $\mathrm{Pow}_m(\beta,\kappa) = \frac{1}{\kappa!}\prod_\lambda(\mathrm{Pow}_{m-1}(\beta,\lambda))_{\kappa(\lambda)}$ into 1), then replace $\mathrm{Pow}_{m-1}(\beta,\lambda) = \sum_{\alpha\leq\beta}\frac{(\beta)_\alpha}{\alpha!}\mathrm{Cov}_{m-1}(\alpha,\lambda)$, which is the zeta inverse of 1) at level $m-1$.

□

4 Polynomial Möbius calculus

This section passes from the discrete setting to the differential setting on polynomial rings. The discrete formulas of the previous section are algebra over the Boolean cube algebra B_k; the differential formulas are algebra over the infinitesimal Taylor algebra A_k. A flat deformation family C_k interpolates between them. The key objects (Taylor–Möbius duality, Faà di Bruno by composition, and the Constantine–Savits coefficients) all transfer from B_k to A_k with Δ replaced by D and coverings replaced by partitions.

Throughout this section, let $\Bbbk$ be a commutative ring. The square-free results ((32) at $\nu = \mathbf{1}$, (33), (25), (35)) hold over any $\Bbbk$; the binomial refinements require $\mathbb{Q} \subseteq \Bbbk$ for the factorial denominators.

4.1 Taylor algebras

(21) Definition (Taylor algebras)**.** For $\nu \in \mathbb{N}_0^k$, define:

$$B_k^\nu := \Bbbk[\delta_1, \ldots, \delta_k]/(\delta_i)_{\nu_i+1}, \qquad A_k^\nu := \Bbbk[\varepsilon_1, \ldots, \varepsilon_k]/(\varepsilon_i^{\nu_i+1}).$$

Here $(\delta_i)_{\nu_i+1} = \delta_i(\delta_i - 1)\cdots(\delta_i - \nu_i)$ is the falling factorial, and relations carrying a free index i generate one relation for each $1 \le i \le k$. We call B_k^ν the *grid algebra* (for $\mathbb{Q} \subseteq \Bbbk$, the function algebra on the finite grid $\prod_i \{0, \ldots, \nu_i\}$; for $\nu = \mathbf{1}$ this holds over any $\Bbbk$) and A_k^ν the *Taylor algebra* (its infinitesimal version).

(22) Example ($\nu = \mathbf{1}$)**.** For $\nu = (1, \ldots, 1)$, the relations simplify to $\delta_i^2 = \delta_i$ and $\varepsilon_i^2 = 0$. Write $B_k = B_k^{\mathbf{1}}$ and $A_k = A_k^{\mathbf{1}}$. Both are free of rank 2^k with basis $\delta^S := \prod_{i\in S} \delta_i$ resp. $\varepsilon^S := \prod_{i \in S} \varepsilon_i$ for $S \subseteq [k]$. The multiplication rules are:

- $\delta^S \delta^T = \delta^{S\cup T}$ in B_k (idempotent: overlaps absorbed).
- $\varepsilon^S \varepsilon^T = \varepsilon^{S\cup T}$ if $S \cap T = \emptyset$, and $\varepsilon^S \varepsilon^T = 0$ otherwise, in A_k (nilpotent: overlapping products vanish).

A product $\delta^{T_1} \cdots \delta^{T_p} = \delta^{T_1 \cup \cdots \cup T_p}$ contributes to the top face $\delta^{[k]}$ exactly when $T_1 \cup \cdots \cup T_p = [k]$, i.e. a covering. A product $\varepsilon^{T_1} \cdots \varepsilon^{T_p}$ is nonzero only when the T_i are pairwise disjoint, i.e. a partition.

4.2 Taylor–Möbius duality

(23) Proposition (Boolean Taylor–Möbius duality)**.** *Let $p \in \Bbbk[y_1, \ldots, y_d]$, $x \in \Bbbk^d$, and $v_1, \ldots, v_k \in \Bbbk^d$. In B_k:*

$$p(x + \textstyle\sum_i \delta_i v_i) = \displaystyle\sum_{S\subseteq[k]} \Delta(p; x; v_S)\, \delta^S, \qquad \Delta(p; x; v_S)\, \delta^S = \sum_{T \subseteq S} (-1)^{|S|-|T|}\, p(x + \textstyle\sum_{i\in T} \delta_i v_i).$$

In A_k:

$$p(x + \textstyle\sum_i \varepsilon_i v_i) = \displaystyle\sum_{S\subseteq[k]} D(p; x; v_S)\, \varepsilon^S, \qquad \varepsilon^S\, D(p; x; v_S) = \sum_{T \subseteq S} (-1)^{|S|-|T|}\, p(x + \textstyle\sum_{i\in T} \varepsilon_i v_i).$$

Here $D(p; x; v_S)$ denotes the ε^S-coefficient of $p(x + \sum_i v_i \varepsilon_i)$ in A_k; this agrees with the iterated directional derivative when $\Bbbk \supseteq \mathbb{Q}$.

Proof. In B_k: expand $p(x + \sum_i \delta_i v_i) = \sum_S c_S\, \delta^S$ in the free basis. The branch maps $\rho_T : \delta_i \mapsto 1_{i\in T}$ give $p(x + \sum_{i\in T} v_i) = \sum_{S\subseteq T} c_S$; Boolean Möbius inversion (5) gives $c_S = \Delta(p; x; v_S)$. The Möbius identity is the same inversion applied inside B_k: the right side is $\sum_{T\subseteq S} (-1)^{|S|-|T|}\, p(x + \sum_{i\in T} \delta_i v_i) = \sum_{T\subseteq S} (-1)^{|S|-|T|} \sum_{R\subseteq T} c_R\, \delta^R = c_S\, \delta^S$. In A_k: the ε^S-coefficient is $D(p; x; v_S)$ by the multinomial theorem; the Möbius identity says the same alternating sieve recovers differentials from infinitesimal grid evaluations.

□

(24) Proposition (Binomial Taylor–Möbius duality)**.** *Let $p \in \Bbbk[y_1, \ldots, y_d]$, $x \in \Bbbk^d$, and $v_1, \ldots, v_k \in \Bbbk^d$. In B_k^ν:*

$$p(x + \textstyle\sum_i \delta_i v_i) = \displaystyle\sum_{0 \le \alpha \le \nu} \frac{1}{\alpha!} \Delta^\alpha(p; x; v_\bullet)\,(\delta)_\alpha, \qquad \Delta^\alpha(p; x; v_\bullet) = \sum_{\beta \le \alpha} (-1)^{\mathrm{wt}(\alpha-\beta)} \frac{(\alpha)_\beta}{\beta!}\, p(x + \textstyle\sum_i \beta_i v_i).$$

In A_k^ν:

$$p(x + \textstyle\sum_i \varepsilon_i v_i) = \displaystyle\sum_{0 \le \alpha \le \nu} \frac{1}{\alpha!} D^\alpha(p; x; v_\bullet)\,\varepsilon^\alpha, \qquad \varepsilon^\nu\, D^\nu(p; x; v_\bullet) = \sum_{\beta \le \nu} (-1)^{\mathrm{wt}(\nu-\beta)} \frac{(\nu)_\beta}{\beta!}\, p(x + \textstyle\sum_i \beta_i \varepsilon_i v_i).$$

Here $D^\alpha(p; x; v_\bullet) := D(p; x; v_\bullet^{\times\alpha})$, the square-free D applied to the slot realization (8). By the multinomial theorem, $\alpha!\,[\varepsilon^\alpha]\,p(x + \sum_i \varepsilon_i v_i) = D(p; x; v_\bullet^{\times\alpha})$; when $\mathbb{Q} \subseteq \Bbbk$, both sides equal the iterated directional derivative $(D_{v_1}^{\alpha_1} \cdots D_{v_k}^{\alpha_k} p)(x)$.

Proof. In B_k^ν: the branch maps $\rho_\beta : B_k^\nu \to \Bbbk$ send $\delta_i \mapsto \beta_i$ for $0 \le \beta \le \nu$. Since $\rho_\beta((\delta)_\alpha) = (\beta)_\alpha$, the expansion coefficients are determined by binomial Möbius inversion (7), giving $\Delta^\alpha/\alpha!$. The Möbius identity is the definition of Δ^α.

In A_k^ν: the expansion gives $D^\alpha/\alpha!$ as the ε^α-coefficient by the multinomial theorem. For the Möbius identity: the coefficient of $D^\alpha \varepsilon^\alpha/\alpha!$ in the alternating sum is $\prod_i \Delta^{\nu_i}(x^{\alpha_i})(0)$, which annihilates $\alpha_i < \nu_i$ and gives $\nu_i!$ at $\alpha_i = \nu_i$.

□

4.3 Faà di Bruno duality

(25) Theorem (Boolean Faà di Bruno)**.** *Let $q : \Bbbk^e \to \Bbbk^d$ and $p : \Bbbk^d \to \Bbbk$ be polynomials, $x \in \Bbbk^e$, $v_1, \ldots, v_k \in \Bbbk^e$, and $y = q(x)$.*

$$\Delta(p \circ q; x; v_{[k]}) = \sum_{\substack{H \subseteq \mathcal{P}_+([k]) \\ \bigcup H = [k]}} \Delta(p; y; (\Delta(q; x; v_T))_{T \in H}),$$

$$D(p \circ q; x; v_{[k]}) = \sum_{\pi \in \mathrm{Part}([k])} D(p; y; (D(q; x; v_B))_{B \in \pi}).$$

Proof. Write $q_T := \Delta(q; x; v_T)$ for $\emptyset \ne T \subseteq [k]$. By cubical Taylor (23), $q(x + \sum_i \delta_i v_i) = y + \sum_{\emptyset \ne T} q_T\, \delta^T$ in B_k. Since $(\delta^T)^2 = \delta^T$, the map $\delta_T \mapsto \delta^T$ defines an algebra map $B_l \to B_k$ where $l = 2^k - 1$. Applying cubical Taylor to p at y in the l directions q_T via this map:

$$(p \circ q)(x + \textstyle\sum_i \delta_i v_i) = p(y + \textstyle\sum_T q_T \delta^T) = \displaystyle\sum_{S \subseteq [k]} \sum_{\substack{H \subseteq \mathcal{P}_+(S) \\ \bigcup H = S}} \Delta(p; y; (q_T)_{T \in H})\, \delta^S.$$

Extracting $[\delta^{[k]}]$ gives the B_k formula: the covering condition $\bigcup H = [k]$ comes from $\delta^{T_1} \cdots \delta^{T_r} = \delta^{T_1 \cup \cdots \cup T_r}$.

The same argument in A_k: since $(\varepsilon^T)^2 = 0$, the map $\varepsilon_T \mapsto \varepsilon^T$ defines $A_l \to A_k$. Now $\varepsilon^{T_1} \cdots \varepsilon^{T_r} = 0$ unless the T_i are pairwise disjoint, so coverings reduce to partitions.

□

(26) Remark. The B_k formula is the algebraic form of the discrete Faà di Bruno (10). Since $B_k \otimes Y \cong \mathrm{Map}(\{0,1\}^k, Y)$ and arbitrary maps act vertexwise, the B_k proof above is the same proof as (10) expressed in the coordinate ring of the cube.

4.4 Iterated Faà di Bruno

(27) Proposition (Binomial Faà di Bruno / Constantine–Savits). *Let $q : \Bbbk^e \to \Bbbk^d$ and $p : \Bbbk^d \to \Bbbk$ be polynomials, $x \in \Bbbk^e$, $v_1, \ldots, v_k \in \Bbbk^e$, $y = q(x)$, and $\gamma \in \mathbb{N}_0^k$:*

$$D(p \circ q;\, x;\, v_\bullet^{\times\gamma}) = \sum_{\kappa \vdash \gamma} \frac{\gamma!}{\kappa! \prod_\alpha (\alpha!)^{\kappa(\alpha)}} D(p;\, y;\, (D(q;\, x;\, v_\bullet^{\times\alpha}))_\alpha^{\times\kappa(\alpha)}).$$

Proof. Apply the A_k Faà di Bruno (25) to $S(\gamma)$ with $|\gamma|$ directions where v_i is repeated γ_i times. Since interchanging repeated copies of v_i does not change D^H, we group partitions by their multi-index profile κ. The number of partitions with profile κ is $\mathrm{Part}_2(\gamma, \kappa) = \gamma!/(\kappa! \prod_\alpha (\alpha!)^{\kappa(\alpha)})$.

□

(28) Definition (Iterated differentials). For polynomial maps $f_1, \ldots, f_m$ with $f_r : \Bbbk^{d_{r-1}} \to \Bbbk^{d_r}$, $x \in \Bbbk^{d_0}$, $v_1, \ldots, v_k \in \Bbbk^{d_0}$, and $x_r = (f_r \circ \cdots \circ f_1)(x)$, define the *iterated differential* for $\kappa \in \mathcal{M}_+^m(k)$: $D^\alpha(f_1;\, x;\, v_\bullet) = D(f_1;\, x;\, v_\bullet^{\times\alpha})$ for $m = 1$, and for $m \geq 2$:

$$D^\kappa(f_1, \ldots, f_m;\, x;\, v_\bullet) := D(f_m;\, x_{m-1};\, (D^\lambda(f_1, \ldots, f_{m-1};\, x;\, v_\bullet)^{\times\kappa(\lambda)})_{\lambda \in \mathrm{Supp}(\kappa)}).$$

For $H \in \mathrm{Part}_m(k)$ with $m \geq 2$, set $D^H(f_1, \ldots, f_m;\, x;\, v_\bullet) := D(f_m;\, x_{m-1};\, (D^L(f_1, \ldots, f_{m-1};\, x;\, v_\bullet))_{L \in H})$.

(29) Proposition (Iterated infinitesimal Faà di Bruno). *Let $f_1, \ldots, f_m$ be polynomial maps as above, $z = x_m$, and $\gamma \in \mathbb{N}_0^k$.*

1) Boolean Faà di Bruno:

$$D(f_m \circ \cdots \circ f_1;\, x;\, v_\bullet) = \sum_{H \in \mathrm{Part}_m(k)} D^H(f_1, \ldots, f_m;\, x;\, v_\bullet).$$

2) Binomial Faà di Bruno:

$$D^\gamma(f_m \circ \cdots \circ f_1;\, x;\, v_\bullet) = \sum_{\kappa \in \mathcal{M}_+^m(k)} \mathrm{Part}_m(\gamma, \kappa)\, D^\kappa(f_1, \ldots, f_m;\, x;\, v_\bullet),$$

where $\mathrm{Part}_m(\gamma, \kappa) := \#\{H \in \mathrm{Part}_m(S(\gamma)) : \nu(H) = \kappa\}$ *is the partition grouping coefficient.*

Proof. 1) Evaluate $f_m \circ \cdots \circ f_1$ on the infinitesimal cube in A_k by composing m Taylor–Möbius expansions. The nilpotent multiplication $\varepsilon^{T_1} \cdots \varepsilon^{T_p} = 0$ unless the T_i are pairwise disjoint forces partition logic at each level, so only higher partitions $H \in \mathrm{Part}_m(k)$ contribute. Reading off the $\varepsilon^{[k]}$-component gives the scalar identity.

2) Apply part 1 to $S(\gamma)$ with $|\gamma|$ directions where v_i is repeated γ_i times. Since D^H depends only on the profile $\kappa = \nu(H)$, we group the m-fold partitions of $S(\gamma)$ by profile. The number of m-fold partitions with profile κ is $\mathrm{Part}_m(\gamma, \kappa)$.

□

(30) Proposition (Partition coefficient recursion)**.** $\mathrm{Part}_m(\gamma,\kappa) = 0$ *unless* $\kappa \vdash \gamma$, *and* $\mathrm{Part}_1(\gamma,\alpha) = \delta_{\gamma,\alpha}$. *For* $m \geq 2$ *and* $\kappa \vdash \gamma$, *the coefficients satisfy the level recursion*

$$\mathrm{Part}_m(\gamma,\kappa) = \frac{\gamma!}{\kappa!} \prod_{\substack{\beta \in \mathcal{M}_+^{m-1}(S) \\ \alpha = \mathrm{leaf}(\beta)}} \left(\frac{1}{\alpha!} \cdot \mathrm{Part}_{m-1}(\alpha,\beta) \right)^{\kappa(\beta)}.$$

The first closed cases are

$$\mathrm{Part}_2(\gamma,\kappa) = \frac{\gamma!}{\kappa!} \prod_{\alpha \in \mathcal{M}_+(S)} (\alpha!)^{-\kappa(\alpha)}, \qquad \mathrm{Part}_3(\gamma,\kappa) = \frac{\gamma!}{\kappa!} \prod_{\beta \in \mathcal{M}_+^2(S)} \left(\frac{1}{\beta!} \prod_{\alpha \in \mathcal{M}_+(S)} (\alpha!)^{-\beta(\alpha)} \right)^{\kappa(\beta)}.$$

All products here and below are finite, since only factors indexed by the iterated supports of κ *differ from* 1.

Proof. The leaves of an m-fold partition $H \in \mathrm{Part}_m(B)$ of a block $B \subseteq S(\gamma)$ partition B, so $\mathrm{leaf}(\nu(H)) = \nu(B)$ by induction on m. In particular, $\mathrm{Part}_{m-1}(\mathrm{leaf}(\beta),\beta)$ is the only nonvanishing evaluation of $\mathrm{Part}_{m-1}(\cdot,\beta)$, and the profile of an m-fold partition of $S(\gamma)$ is a multi-index partition of γ.

For the recursion, an m-fold partition with profile κ is built in two independent steps. First, partition $S(\gamma)$ into blocks matched with the elements of the multiset κ, where a block matched with β has type $\mathrm{leaf}(\beta)$; there are $\gamma!/(\kappa! \prod_\beta (\mathrm{leaf}(\beta)!)^{\kappa(\beta)})$ such matched partitions. Second, equip each block matched with β with an $(m-1)$-fold partition of profile β, in $\mathrm{Part}_{m-1}(\mathrm{leaf}(\beta),\beta)$ ways.

□

4.5 Faà di Bruno deformation

The two Taylor algebras B_k^ν and A_k^ν are fibers of a flat deformation over $\Bbbk[t]$. The deformation construction carries over mutatis mutandis to the iterated and binomial settings via the grid algebra C_k^ν; we present the Boolean two-map case here, since it is the least notation-heavy and already contains the essential mechanism.

(31) Definition (Deformation cube algebra)**.** For $\nu \in \mathbb{N}_0^k$, define:

$$C_k^\nu := \Bbbk[t][x_1,\dots,x_k] \,/\, \left(\textstyle\prod_{j=0}^{\nu_i} (x_i - jt) \,:\, 1 \leq i \leq k \right).$$

For $\nu = \mathbf{1}$, write $C_k = C_k^{\mathbf{1}}$; the relation $x_i^2 = t x_i$ gives $x^S x^T = t^{|S \cap T|}\, x^{S \cup T}$. At $t = 1$ all coverings contribute; at $t = 0$ only partitions remain.

(32) Lemma (Flatness and fibers)**.** C_k^ν *is a free* $\Bbbk[t]$*-module of rank* $\prod_i (\nu_i + 1)$. *In particular,* C_k^ν *is flat over* $\Bbbk[t]$, *with fibers*

$$C_k^\nu/(t) \cong A_k^\nu, \qquad C_k^\nu/(t-1) \cong B_k^\nu.$$

Proof. The relations $\prod_{j=0}^{\nu_i} (x_i - jt) = 0$ are monic of degree $\nu_i + 1$ in x_i, so C_k^ν is free over $\Bbbk[t]$ with monomial basis $\{x^\alpha : 0 \leq \alpha \leq \nu\}$. The fiber identifications follow by substituting $t = 0$ and $t = 1$.

□

(33) Proposition (Taylor–Möbius duality in C_k)**.** *Let* $p \in \Bbbk[y_1,\dots,y_d]$, $x \in \Bbbk^d$, *and* $v_1,\dots,v_k \in \Bbbk^d$. *Expand* $p(x + \sum_i v_i x_i) = \sum_{S \subseteq [k]} c_S(t)\, x^S$ *in* C_k. *Then*

$$c_S(t) = \frac{1}{t^{|S|}}\, \Delta(p; x; (tv_i)_{i \in S}), \qquad c_S(0) = D(p; x; v_S).$$

The coefficients $c_S(t)$ lie in $\Bbbk[t]$; in particular, $\Delta(p;x;tv_1,\dots,tv_k)$ is divisible by t^k in $\Bbbk[t]$.

Proof. Write $p(x+\sum_i v_ix_i)=\sum_S c_S\,x^S$ in C_k (free over $\Bbbk[t]$ with basis x^S). In the generic fiber $C_k[t^{-1}]$, set $\delta_i=t^{-1}x_i$, so $\delta_i^2=\delta_i$. The branch maps $\rho_T:\delta_i\mapsto 1_{i\in T}$ give $\rho_T(p(x+\sum_i v_ix_i))=p(x+\sum_{i\in T}tv_i)$. Since $\rho_T(x^S)=t^{|S|}\cdot 1_{S\subseteq T}$, Boolean Möbius inversion (5) gives $t^{|S|}c_S(t)=\sum_{T\subseteq S}(-1)^{|S|-|T|}p(x+\sum_{i\in T}tv_i)=\Delta(p;x;(tv_i)_{i\in S})$, hence $c_S(t)=t^{-|S|}\,\Delta(p;x;(tv_i)_{i\in S})$. Since $c_S\in\Bbbk[t]$ (freeness), the divisibility follows, and $c_S(0)=D(p;x;v_S)$ by definition.

□

(34) Lemma (Multilinear part of the mixed difference)**.** *For a polynomial p, a point y, and $w_1,\dots,w_r\in\Bbbk^d$:*

$$\Delta(p;\,y;\,\varepsilon_1w_1,\dots,\varepsilon_rw_r)=D(p;\,y;\,w_1,\dots,w_r)\,\varepsilon^{[r]}\qquad \textit{in } A_r.$$

In particular, the multilinear part of $\Delta(p;y;w_\bullet)$ in $(w_1,\dots,w_r)$ equals $D(p;y;w_\bullet)$.

Proof. This is the A_r Möbius display of (23) applied at y with directions $w_\bullet$: the left side is the alternating sum $\sum_{T\subseteq[r]}(-1)^{r-|T|}\,p(y+\sum_{j\in T}\varepsilon_jw_j)$, and the $\varepsilon^{[r]}$-coefficient of $p(y+\sum_j\varepsilon_jw_j)$ is $D(p;y;w_\bullet)$ by definition. Substituting $w_j\mapsto\varepsilon_jw_j$ retains exactly the multilinear monomials of $\Delta(p;y;w_\bullet)$, since $\varepsilon_j^2=0$.

□

(35) Theorem (Faà di Bruno deformation)**.** *Let $q:\Bbbk^e\to\Bbbk^d$ and $p:\Bbbk^d\to\Bbbk$ be polynomials, $x\in\Bbbk^e$, $v_1,\dots,v_k\in\Bbbk^e$, and $y=q(x)$. Write $c_S(t)$ for the x^S-coefficient of $(p\circ q)(x+\sum_i v_ix_i)\in C_k$. Then*

$$c_S(t)=\frac{1}{t^{|S|}}\sum_{\substack{H\subseteq\mathcal{P}_+(S)\\ \bigcup H=S}}\Delta(p;\,y;\,(\Delta(q;\,x;\,(tv_i)_{i\in T}))_{T\in H}),\qquad c_S(0)=\sum_{\pi\in\mathrm{Part}(S)}D(p;\,y;\,(D(q;\,x;\,v_B))_{B\in\pi}).$$

Each covering summand is divisible by $t^{\mathrm{wt}(H)}$ (since $\Delta(q;x;(tv_i)_{i\in T})$ is divisible by $t^{|T|}$), so $c_S\in\Bbbk[t]$. At $t=0$, only the coverings with $\mathrm{wt}(H)=|S|$, namely the partitions, contribute.

Proof. By (33) applied to the composite (extending scalars to $\Bbbk[t]$), $c_S(t)=t^{-|S|}\,\Delta(p\circ q;\,x;\,(tv_i)_{i\in S})$. The discrete covering formula (10), applied over $\Bbbk[t]$ with increments $(tv_i)_{i\in S}\in\Bbbk[t]^e$, expands this as $t^{-|S|}\sum_{H\in\mathrm{Cov}(S)}\Delta(p;\,y;\,(\Delta(q;\,x;\,(tv_i)_{i\in T}))_{T\in H})$. Each inner increment is divisible by $t^{|T|}$ by (33); write $\Delta(q;x;(tv)_T)=t^{|T|}\tilde q_T(t)$ with $\tilde q_T\in\Bbbk[t]^d$ and $\tilde q_T(0)=D(q;x;v_T)$. As a polynomial in $(w_1,\dots,w_r)\in(\Bbbk^d)^r$, the mixed difference $\Delta(p;y;w_1,\dots,w_r)$ vanishes on every hyperplane $w_j=0$, so each of its monomials contains at least one factor from each w_j; substituting $w_j=t^{|T_j|}\tilde q_{T_j}(t)$ therefore extracts $\prod_{T\in H}t^{|T|}=t^{\mathrm{wt}(H)}$. Hence $c_S(t)\in\Bbbk[t]$. At $t=0$, terms with $\mathrm{wt}(H)>|S|$ vanish; the remaining terms have $\mathrm{wt}(H)=|S|$, i.e. H is a partition (15). For a partition π with blocks $B_1,\dots,B_r$, the π-summand is $\Delta(p;\,y;\,(t^{|B_j|}\tilde q_{B_j}(t))_j)$; by the multilinear-part lemma (34), the multilinear monomials contribute $t^{\sum|B_j|}D(p;\,y;\,(\tilde q_{B_j}(t))_j)=t^{|S|}D(p;\,y;\,(\tilde q_{B_j}(t))_j)$, while all higher monomials contribute above order $|S|$. Dividing by $t^{|S|}$ and setting $t=0$ gives $D(p;\,y;\,(D(q;\,x;\,v_B))_{B\in\pi})$, since $\tilde q_B(0)=D(q;x;v_B)$ (33).

□

5 Fréchet Möbius calculus

The only analytic input needed to pass from polynomial maps to C^n maps between Banach spaces is that a C^n composite is computed by its Taylor polynomials up to order n; this is (36), proved in [Har25]. Everything else is the polynomial calculus of the previous section applied to jets.

The polynomial calculus of the previous section extends verbatim to polynomial maps between $\mathbb{k}$-modules, and in particular to continuous polynomial maps between Banach spaces: the proofs use only base change along $\mathbb{k} \to A_k^\nu$, which extends canonically to symmetric multilinear maps, and the branch maps ρ_β act on the ε-variables only, never on coordinates of the target.

Let $X_0, \dots, X_m$ be Banach spaces and let $f_r : X_{r-1} \to X_r$ be C^n near the relevant basepoints. Put $x_0 = x$ and $x_r = f_r(x_{r-1})$. We write $T^n(f_r; x_{r-1})$ for the order-n Taylor polynomial of f_r at x_{r-1}, regarded as a polynomial map on X_{r-1}. For a C^n map f, the mixed directional derivative $D^\alpha(f; x; v_\bullet) := D^{|\alpha|} f(x)[v_1^{\otimes \alpha_1}, \dots, v_k^{\otimes \alpha_k}]$ agrees with $D^\alpha(T^n(f; x); x; v_\bullet)$ for $|\alpha| \le n$; the iterated differential D^κ for $\kappa \in \mathcal{M}_+^m(k)$ is defined by the same recursion as (28), applied to the Taylor jets. Every κ with $\mathrm{Part}_m(\gamma, \kappa) \neq 0$ satisfies $\kappa \vdash \gamma$, so the sum in (37) is finite and only derivatives of order $\le |\gamma| \le n$ occur.

(36) Proposition (Taylor composition for Fréchet maps)**.** *For C^n maps between Banach spaces,*

$$T^n(f_m \circ \cdots \circ f_1; x) = \pi_{\le n}(T^n(f_m; x_{m-1}) \circ \cdots \circ T^n(f_1; x_0)),$$

where $\pi_{\le n}$ denotes truncation to total degree at most n.

Proof. Write $f_r = P_r + R_r$ where $P_r = T^n(f_r; x_{r-1})$ is the Taylor polynomial and R_r is the Peano residual with $R_r(h) = o(\|h\|^n)$. Substituting into the composite and expanding, every term containing at least one factor of R_r is $o(\|h\|^n)$, so the Taylor jet of the composite up to order n agrees with the truncated composite of the Taylor polynomials. Details are given in [Har25].

□

Thus the Taylor–Möbius duality and Faà di Bruno formulas of the previous section, which are identities for polynomial maps proved via the infinitesimal grid in A_k^ν, apply directly to the Taylor jets of C^n maps.

(37) Theorem (Fréchet iterated Faà di Bruno / recursive Constantine–Savits)**.** *Let $X_0, \dots, X_m$ be Banach spaces, let $f_r : X_{r-1} \to X_r$ be C^n near x_{r-1}, and put $x_0 = x$, $x_r = f_r(x_{r-1})$, $z = x_m$. Fix directions $v_1, \dots, v_k \in X_0$ and let $\gamma \in \mathbb{N}_0^k$ with $1 \le |\gamma| \le n$.*

1) Taylor composition:

$$(f_m \circ \cdots \circ f_1)(x + \textstyle\sum_i t_i v_i) = z + \displaystyle\sum_{\substack{0<|\gamma|\le n \\ \kappa \in \mathcal{M}_+^m(k)}} \frac{t^\gamma}{\gamma!} \mathrm{Part}_m(\gamma, \kappa)\, D^\kappa(f_1, \dots, f_m; x; v_\bullet) + o(|t|^n).$$

2) Faà di Bruno:

$$D^\gamma(f_m \circ \cdots \circ f_1; x; v_\bullet) = \sum_{\kappa \in \mathcal{M}_+^m(k)} \mathrm{Part}_m(\gamma, \kappa)\, D^\kappa(f_1, \dots, f_m; x; v_\bullet).$$

Here D^κ is the recursive iterated differential (28) and $\mathrm{Part}_m(\gamma, \kappa)$ is the partition grouping coefficient (30).

Proof. By Taylor composition (36), $F(x+h) = \pi_{\le n}(P_m \circ \cdots \circ P_1)(x+h) + o(\|h\|^n)$, where $P_r = T^n(f_r; x_{r-1})$ are polynomial maps between Banach spaces. The polynomial Faà di Bruno (29), which applies to polynomial maps between modules, gives 2) for the composite of the Taylor jets. For 1), restrict to $h = \sum_i t_i v_i$ with $\|h\| \le C|t|$, so $o(\|h\|^n) = o(|t|^n)$; expanding the truncated polynomial composite in t by binomial Taylor–Möbius duality (24) and inserting 2) gives the stated formula.

□

(38) Corollary (Multivariate Taylor formula, [Lan93])**.** *For a single C^n map $f : X \to Y$ between Banach spaces, $x \in X$, and $v_1, \dots, v_k \in X$:*

$$f(x+\textstyle\sum_i t_i v_i) = f(x) + \displaystyle\sum_{1\le|\alpha|\le n} \frac{t^\alpha}{\alpha!} D^\alpha(f;\, x;\, v_\bullet) + o(|t|^n).$$

(39) Corollary (Constantine–Savits, [CS96]). *For two C^n maps $f_1 : X_0 \to X_1$, $f_2 : X_1 \to X_2$ between Banach spaces, $x \in X_0$, $x_1 = f_1(x)$, and $\gamma \in \mathbb{N}_0^k$ with $|\gamma| \le n$:*

$$D(f_2 \circ f_1;\, x;\, v_\bullet^{\times\gamma}) = \sum_{\kappa\vdash\gamma} \frac{\gamma!}{\kappa!\, \prod_\alpha (\alpha!)^{\kappa(\alpha)}} D(f_2;\, x_1;\, (D(f_1;\, x;\, v_\bullet^{\times\alpha}))_\alpha^{\times\kappa(\alpha)}).$$

(40) Corollary (Univariate Faà di Bruno). *For C^n functions $f, g : \mathbb{R} \to \mathbb{R}$, the Constantine–Savits formula (39) with $k = 1$ and $\gamma = n$ gives:*

$$(f\circ g)^{(n)}(x) = \sum_{\substack{k_1+2k_2+\cdots+nk_n=n \\ k_j\ge 0}} \frac{n!}{k_1!\,k_2!\cdots k_n!} f^{(k_1+\cdots+k_n)}(g(x)) \prod_{j=1}^n \left(\frac{g^{(j)}(x)}{j!}\right)^{k_j}.$$

This is the classical Faà di Bruno formula in Bell polynomial form.

6 Applications

Each section exercises a different face of the covering calculus: a combinatorial self-count, a degree bound for polynomial maps between abelian groups, a composition law for Newton and Mahler coefficients, and a fast algorithm for composing discrete jets. The final section on product rules is a complement rather than an application: the same Möbius mechanism run for products instead of composites.

6.1 Self-enumeration

(41) Proposition (Covering counts). *Let $g : \mathbb{N}_0 \to \mathbb{Z}$, $g(x) = 2^x - 1$ and $f : \mathbb{N}_0 \to \mathbb{Z}$, $f(y) = 2^y$. Then*

$$|\mathrm{Cov}(k)| = \sum_{j=0}^k (-1)^{k-j} \binom{k}{j} 2^{2^j-1}.$$

Proof. Every increment equals 1: $\Delta^j g(0) = 1$ for $j \ge 1$ and $\Delta^p f(0) = 1$ for $p \ge 1$ at $y = g(0) = 0$. The covering formula (10) gives $\Delta^k(f \circ g)(0) = \sum_{H\in\mathrm{Cov}(k)} 1 = |\mathrm{Cov}(k)|$. The Newton expansion of the left side gives the stated identity.

□

The covering formula counts its own terms, the exact discrete parallel of the classical fact that e^{e^x-1} counts partitions (Bell numbers): exponential maps are the group-like elements where every increment is 1, and each formula enumerates its own index set.

	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$
$\lvert\mathrm{Cov}(k)\rvert$	1	5	109	32,297	2,147,321,017
$\lvert\mathrm{Part}(k)\rvert = B_k$	1	2	5	15	52

6.2 Degree bound for polynomial maps

A map $g : X \to Y$ between abelian groups is *polynomial of degree* $\le d$ if $\Delta^{d+1} g \equiv 0$ (see [MO34], [Lei02]).

(42) Corollary. *If $g : X \to Y$ is polynomial of degree $\le d$ and $f : Y \to Z$ is polynomial of degree $\le e$, then $f \circ g$ is polynomial of degree $\le ed$.*

Proof. In the covering formula (10), the H-term involves $\Delta^{|H|}f$ (vanishing for $|H| > e$) applied to directions $\Delta^{|T|}g$ (vanishing for $|T| > d$). A covering with $|H| \le e$ blocks of size $|T| \le d$ satisfies $k \le \sum_{T\in H} |T| \le ed$.

□

6.3 Newton–Mahler composition

For $g : \mathbb{N}_0 \to \mathbb{Z}$ with Newton coefficients $b_j = \Delta^j g(0)$ and any $f : \mathbb{Z} \to Y$, the binomial covering formula (18) with $\gamma = n$ and unit increments gives the Newton coefficients of $f \circ g$ in terms of the Newton coefficients of g and finite differences of f at $g(0)$:

$$\Delta^n(f \circ g)(0) = \sum_{\kappa} \mathrm{Cov}_2(n, \kappa)\, \Delta(f;\, g(0);\, (b_j)_j^{\times\kappa(j)}),$$

where the sum runs over all κ with $\mathrm{Cov}_2(n, \kappa) \neq 0$ (finitely many). Since Mahler's theorem identifies continuous maps $\mathbb{Z}_p \to \mathbb{Z}_p$ with convergent Newton series ($b_n \to 0$ p-adically), the same formula computes Mahler coefficients of composites of continuous p-adic maps.

6.4 Discrete jet composition

The fixed-basepoint property makes discrete jets composable without re-evaluation. A *discrete k-jet* of g at x in directions $u_\bullet$ is the table $(\Delta(g; x; u_T))_{T\subseteq[k]}$ of 2^k forward differences. Given this table, the discrete jet of $f \circ g$ is computed in three steps:

1. *Zeta (sum up):* recover the 2^k vertex values $g(x + \sum_{i\in S} u_i)$ from the differences by discrete Taylor (9).
2. *Evaluate:* apply f pointwise to get $(f \circ g)(x + \sum_{i\in S} u_i)$ for all $S \subseteq [k]$.
3. *Möbius (alternate):* extract the differences $\Delta(f \circ g; x; u_S)$ by the alternating sieve.

Steps 1 and 3 are the fast zeta and Möbius transforms on $\{0, 1\}^k$, each costing $O(k \cdot 2^k)$ operations. The total cost is $O(k \cdot 2^k)$, polynomial in the jet size 2^k. The covering formula is what this algorithm produces when expanded algebraically; the zeta/Möbius factorization is the efficient implementation.

6.5 Product rules

(43) Theorem (Discrete product rule)**.** *Let X be an abelian group, A an associative algebra, and $f_1, \dots, f_r : X \to A$. For $x \in X$ and $u_1, \dots, u_n \in X$:*

1) Boolean product rule:

$$\Delta(f_1 \cdots f_r; x; u_1, \dots, u_n) = \sum_{\substack{J_1,\dots,J_r\subseteq[n] \\ J_1\cup\cdots\cup J_r=[n]}} \Delta(f_1; x; u_{J_1}) \cdots \Delta(f_r; x; u_{J_r}).$$

2) Boolean Taylor product rule:

$$(f_1 \cdots f_r)(x + \textstyle\sum_{i=1}^n u_i) = \displaystyle\sum_{J_1,\dots,J_r\subseteq[n]} \Delta(f_1; x; u_{J_1}) \cdots \Delta(f_r; x; u_{J_r}).$$

The sum in 1) runs over all ordered coverings $(J_1, \dots, J_r)$ *of $[n]$: the J_a may overlap, and empty J_a are allowed. Part 2) is the product of Taylor expansions. The binomial versions follow by applying 1,2) to $S(\gamma)$.*

Proof. By discrete Taylor (9), $f_a(x + \sum_{i\in S} u_i) = \sum_{J_a\subseteq S} \Delta(f_a; x; u_{J_a})$. Substituting into the alternating sum for $f_1 \cdots f_r$ and expanding the product gives:

$$\Delta(f_1 \cdots f_r; x; u_1, \ldots, u_n) = \sum_{\substack{J_1,\ldots,J_r \subseteq [n] \\ S \supseteq J_1 \cup \cdots \cup J_r}} (-1)^{n-|S|}\, \Delta(f_1; x; u_{J_1}) \cdots \Delta(f_r; x; u_{J_r}).$$

Exchanging the order of summation, a fixed tuple $(J_1, \ldots, J_r) \in \mathcal{P}([n])^r$ appears in the S-summand exactly when $J_1 \cup \cdots \cup J_r \subseteq S$. Setting $M = [n] \setminus (J_1 \cup \cdots \cup J_r)$, the sign sum is $\sum_{R \subseteq M} (-1)^{|M|-|R|} = (1-1)^{|M|}$: this is 1 if $J_1 \cup \cdots \cup J_r = [n]$ and 0 otherwise.

□

(44) Theorem (Fréchet product rule)**.** *Let X, Y be Banach spaces, A a Banach algebra, and $f_1, \ldots, f_r : X \to A$ be C^n near x, with $n \geq k$. For $v_1, \ldots, v_k \in X$:*

1) Boolean product rule:

$$D(f_1 \cdots f_r; x; v_{[k]}) = \sum_{J_1 \sqcup \cdots \sqcup J_r = [k]} D(f_1; x; v_{J_1}) \cdots D(f_r; x; v_{J_r}).$$

2) Boolean Taylor product rule: *the square-free part of the order-n Taylor polynomial of $t \mapsto (f_1 \cdots f_r)(x + \sum_{i=1}^{k} t_i v_i)$ is*

$$\sum_{J_1 \sqcup \cdots \sqcup J_r \subseteq [k]} D(f_1; x; v_{J_1}) \cdots D(f_r; x; v_{J_r})\, t^{J_1 \sqcup \cdots \sqcup J_r}.$$

3) Binomial product rule: *For $\gamma \in \mathbb{N}_0^k$ with $|\gamma| \leq n$:*

$$D(f_1 \cdots f_r; x; v_\bullet^{\times\gamma}) = \sum_{\alpha_1 + \cdots + \alpha_r = \gamma} \frac{\gamma!}{\alpha_1! \cdots \alpha_r!}\, D(f_1; x; v_\bullet^{\times\alpha_1}) \cdots D(f_r; x; v_\bullet^{\times\alpha_r}).$$

4) Binomial Taylor product rule:

$$(f_1 \cdots f_r)(x + \textstyle\sum_i t_i v_i) = \displaystyle\sum_{\substack{\alpha_1,\ldots,\alpha_r \geq 0 \\ |\alpha_1 + \cdots + \alpha_r| \leq n}} \frac{t^{\alpha_1 + \cdots + \alpha_r}}{\alpha_1! \cdots \alpha_r!}\, D(f_1; x; v_\bullet^{\times\alpha_1}) \cdots D(f_r; x; v_\bullet^{\times\alpha_r}) + o(|t|^n).$$

The sum in 1) runs over all ordered partitions *$(J_1, \ldots, J_r)$ of $[k]$: the J_a are pairwise disjoint and empty J_a are allowed. This is the discrete product rule with coverings replaced by partitions.*

Proof. By Taylor composition (36), the coefficients of the product $f_1 \cdots f_r$ up to order n agree with those of the product of Taylor polynomials. For polynomials, the result follows from the discrete product rule applied in the Taylor algebra A_k: since $\varepsilon^{J_1} \cdots \varepsilon^{J_r} = 0$ unless the J_a are pairwise disjoint, the ordered coverings of the discrete rule reduce to ordered partitions.

□

7 Acknowledgements

We thank Duarte and Torres for their wonderfully dense paper [DT12], which inspired this work, and for sharing the source code accompanying their formula.

References

[BGN04] Wolfgang Bertram, Helge Glöckner, and Karl-Hermann Neeb. Differential calculus over general base fields and rings. *Expositiones Mathematicae*, 22:213–282, 2004.

[BJO+18] Kristine Bauer, Brenda Johnson, Christina Osborne, Emily Riehl, and Amelia Tebbe. Directional derivatives and higher order chain rules for abelian functor calculus. *Topology and its Applications*, 235:375–427, 2018.

[CS96] Gregory M. Constantine and Thomas H. Savits. A multivariate Faà di Bruno formula with applications. *Transactions of the American Mathematical Society*, 348(2):503–520, 1996. URL `https://www.ams.org/journals/tran/1996-348-02/S0002-9947-96-01501-2/`.

[DT12] P. Duarte and M. J. Torres. A discrete Faà di Bruno's formula. *Contributions to Discrete Mathematics*, 7(2):72–83, 2012. URL `https://doi.org/10.55016/ojs/cdm.v7i2.62119`.

[EL54] Samuel Eilenberg and Saunders Mac Lane. On the groups $H(\Pi, n)$, II: Methods of computation. *Annals of Mathematics*, 60:49–139, 1954.

[FL07] Michael S. Floater and Tom Lyche. Two chain rules for divided differences and Faaà di Bruno's formula. *Mathematics of Computation*, 76(258):867–877, 2007. URL `https://doi.org/10.1090/S0025-5718-06-01916-8`.

[FM14] Alessandra Frabetti and Dominique Manchon. Five interpretations of Faaà di Bruno's formula, 2014. URL `https://arxiv.org/abs/1402.5551`.

[Har06] Michael Hardy. Combinatorics of partial derivatives. *Electronic Journal of Combinatorics*, 13:Research Paper 1, 2006. URL `https://www.combinatorics.org/ojs/index.php/eljc/article/view/v13i1r1`.

[Har25] Heinrich Hartmann. Faà di Bruno is Taylor composition, 2025. URL `https://arxiv.org/abs/2606.26133`.

[JR79] S. A. Joni and G.-C. Rota. Coalgebras and bialgebras in combinatorics. *Studies in Applied Mathematics*, 61:93–139, 1979.

[Lan93] Serge Lang. *Real and Functional Analysis*, volume 142 of *Graduate Texts in Mathematics*. Springer, New York, 3rd edition, 1993. URL `https://link.springer.com/book/10.1007/978-1-4612-0897-6`.

[Lei02] Alexander Leibman. Polynomial mappings of groups. *Israel Journal of Mathematics*, 129:29–60, 2002.

[Lev06] Eliahu Levy. Why do partitions occur in Faa di Bruno's chain rule for higher derivatives?, 2006. arXiv preprint. URL `https://arxiv.org/abs/math/0602183`.

[MO34] Stanisław Mazur and Władysław Orlicz. Grundlegende Eigenschaften der polynomischen Operationen. *Studia Mathematica*, 5:50–68, 1934. URL `https://doi.org/10.4064/sm-5-1-50-68`.

[NR04] Pierpaolo Natalini and Paolo Emilio Ricci. An extension of the Bell polynomials. *Computers & Mathematics with Applications*, 47:719–725, 2004.

[Sta11] Richard P. Stanley. *Enumerative Combinatorics*, volume 1 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 2 edition, 2011.

[WW06] Weiping Wang and Tianming Wang. On the divided difference form of Faà di Bruno's formula. *Journal of Computational Mathematics*, 24:553–560, 2006.

[XW10] Aimin Xu and Chuanqing Wang. Some extensions of Faà di Bruno's formula with divided differences. *Computers & Mathematics with Applications*, 59:2047–2052, 2010.