

Generic products of linear forms saturate the shifted partial derivative measure

Brandon Hudgeons*

August 2026

Abstract

Let f be a product of D generic linear forms in m variables over a field of characteristic 0, and for integers $k, \ell \geq 0$ let $\Gamma_{k,\ell}(f) = \dim S_\ell \cdot \partial^k f$ be its shifted partial derivative measure, the complexity measure behind the known lower bounds for homogeneous depth-four algebraic circuits. Two universal upper bounds hold for every homogeneous f of degree D : $\Gamma_{k,\ell}(f) \leq \min(N_k N_\ell, N_{D-k+\ell})$, where $N_e = \binom{e+m-1}{m-1}$. Our main theorem determines the measure of a generic product at every fixed (k, ℓ) once the degree is polynomially large, and in particular proves that it meets this cap.

For $m = 2$ we prove the exact law $\Gamma_{\kappa,\lambda}(f) = \min((\kappa+1)(\lambda+1), D - \kappa + \lambda + 1)$ for all $0 \leq \kappa \leq D$ and all $\lambda \geq 0$. For every fixed $m \geq 3$ we prove that for all (k, ℓ) and all $D \geq D_0(m, k, \ell) = \text{poly}(k, \ell)$, a generic product of D linear forms satisfies $\Gamma_{k,\ell}(f) = \min(N_k N_\ell, N_{D-k+\ell})$ — full saturation of the universal cap. For derivative spaces ($\ell = 0$) we prove exact equality $\dim \partial^k f = \min(N_k, N_{D-k})$ for all $k \leq D/m$ and $k \geq D - D/m$, and (once $D \geq 2m^2$) within a factor $(2/e)^{m-1}/(2e^2m)$ of the cap at *every* k , via a standalone combinatorial comparison lemma for capped compositions (a Pólya-urn coupling plus log-concavity). We also compute *exactly*, by a filtration calculus, the measure of products with disjoint-pair block structure, and show these are genuinely deficient in the row-dominated regime — for even $m \geq 6$ and $k = \ell \geq m^2$, by a factor at least $(k/32m)^{(m-4)/2}$ — witness choice, not analysis slack, is what previously kept this regime open.

The complexity-theoretic reading: against a single product gate of generic linear forms in any fixed number of variables, shifted partial derivatives certify nothing beyond a polynomial degree threshold. This is an exact, per-gate form of the saturation phenomenon underlying the rank-measure barriers of Efremenko–Landsberg–Schenck–Weyman, Efremenko–Garg–Oliveira–Wigderson, and Bhargav–Dutta–Saxena, here established with exact constants in the few-variable, high-degree regime relevant to algebraic hardness–randomness bootstrapping. Characteristic 0 is essential: over small finite fields the evaluation variant of the measure (Armand–Behera–Tavenas, 2026) reverses the polarity.

The commutative-algebra reading: we determine the Hilbert function of the ideal generated by $\partial^k f$ in each degree $k + \ell$ — for f a generic hyperplane multi-arrangement form — extending the study of apolar algebras of products of linear forms initiated by DiPasquale–Flores–Peterson. The proofs are elementary throughout (no recourse to Fröberg-type conjectures or Alexander–Hirschowitz): the main theorem reduces, by an exact “master reduction,” to the rank of an explicit Laurent-polynomial family, which is resolved by a zero-multiplicity bound for exponential polynomials, layered confluent and twist Vandermonde arguments, and a residual core lemma valid over any field. Every machine-checkable step of the derivation has been verified exactly (integer or modular arithmetic, multiple primes and seeds), including the full generic- m code path at $m = 4, \dots, 9$.

*Email: brandon@atsoft.at. This work was produced with heavy reliance on Anthropic’s Claude (Fable 5); see the disclosure in Section 1.5.

Contents

1	Introduction	3
1.1	The measure and the question	3
1.2	Main results	4
1.3	The complexity-theoretic reading	6
1.4	The commutative-algebra reading	7
1.5	Proof overview and organization	7
2	Preliminaries	9
2.1	Notation	9
2.2	The universal caps	10
2.3	Semicontinuity and the variety of products	10
2.4	Euler nesting	11
3	Two variables: the exact min-law	11
4	Disjoint-pair products: an exact calculus	12
4.1	Profile decomposition and chains	13
4.2	The exact formula	14
5	Derivative spaces: the case $\ell = 0$	15
5.1	The generating-function formula	15
5.2	A comparison lemma for capped compositions	16
6	The two regimes for disjoint-pair products	17
6.1	Ambient regime: saturation up to boundary mass	17
6.2	Row regime: block witnesses are genuinely deficient	19
7	Chain witnesses and the master reduction	20
7.1	Conventions for the chain	21
7.2	The master reduction	22
8	The end peel	23
9	The middle peel: coordinatization, windows, and the twist Vandermonde	25
9.1	Slice structure	25
9.2	Coordinatization: the telescoped global relation	26
9.3	Tensor factorization, split recovery, windows	27
9.4	The twist Vandermonde	28
10	The residual core and the assembly	28
10.1	The core lemma, uniform in m	28
10.2	The base case $m = 3$	29
10.3	Assembly	30
11	Related work	31

A	Worked instances: $m \leq 6$	32
A.1	$m = 3$	32
A.2	$m = 4$	33
A.3	$m = 5$	33
A.4	$m = 6$	34
A.5	Slack observed at the instances	34
B	Machine verification record	34
B.1	Protocol	34
B.2	What is checked, per result	35
B.3	A control	36

1 Introduction

1.1 The measure and the question

Fix a field $\mathbb{F}$ of characteristic 0 and the polynomial ring $S = \mathbb{F}[x_1, \dots, x_m]$, with graded pieces S_e of dimension $N_e = \binom{e+m-1}{m-1}$. For a homogeneous polynomial f of degree D and integers $k, \ell \geq 0$, the *shifted partial derivative space* is

$$S_\ell \cdot \partial^k f = \text{span}\{x^\alpha \partial^\beta f : |\alpha| = \ell, |\beta| = k\} \subseteq S_{D-k+\ell},$$

and the *shifted partial derivative measure* is its dimension,

$$\Gamma_{k,\ell}(f) = \dim S_\ell \cdot \partial^k f.$$

Introduced by Kayal [19] and developed by Gupta–Kamath–Kayal–Saptharishi [16], this measure and its variants drive essentially all known lower bounds for homogeneous depth-four algebraic circuits [16, 21, 22, 12]. Every application has the same shape: one proves an *upper* bound on $\Gamma_{k,\ell}$ of a single product gate — a polynomial of the form $Q_1 Q_2 \cdots Q_t$ — so that a circuit with few gates has small measure, and a *lower* bound on $\Gamma_{k,\ell}$ of the explicit hard polynomial.

The stakes of this template are the highest in algebraic complexity. By the depth-reduction theorems of Agrawal–Vinay, Koiran, and Tavenas [2, 20, 25], every homogeneous n -variate polynomial of degree $d = n^{O(1)}$ computable by an algebraic circuit of polynomial size is computable by a *homogeneous depth-four* circuit of size $n^{O(\sqrt{d})}$ with bottom fan-in $O(\sqrt{d})$; consequently an explicit polynomial family (in VNP) requiring homogeneous depth-four circuits of size $n^{\omega(\sqrt{d})}$ — *any* super-constant improvement in the exponent over the known $n^{\Omega(\sqrt{d})}$ bounds — would separate Valiant’s classes VP and VNP [26], the algebraic analogue of P vs. NP. Shifted partial derivatives is the measure family that produced those $n^{\Omega(\sqrt{d})}$ bounds, and whether it can be pushed past the threshold is exactly the question of how tightly the measure of product gates can be upper-bounded. This is the sense in which exact saturation results matter: each one closes, definitively, one route across the depth-four chasm. (To keep the claim sharp: our theorems concern single product-of-linear-forms gates in few variables — an exact instance of the saturation phenomenon, not a formal barrier for general depth-four lower bounds; Section 1.3 states the precise relationship to the barrier results.)

Two upper bounds hold for *every* homogeneous f of degree D , for counting reasons alone:

$$\Gamma_{k,\ell}(f) \leq \min\left(\underbrace{N_k N_\ell}_{\text{row count}}, \underbrace{N_{D-k+\ell}}_{\text{ambient dimension}}\right), \quad (1)$$

since the space is spanned by $N_k N_\ell$ vectors inside $S_{D-k+\ell}$. We call the two terms the *row cap* and the *ambient cap*, and say f *saturates* the measure at (k, ℓ) when equality holds in (1). A saturating polynomial is a certificate of blindness: no lower-bound argument based on $\Gamma_{k,\ell}$ can distinguish any polynomial from a circuit whose gates saturate, because the gates already carry the maximal possible measure.

This paper answers the question:

What is the exact shifted partial derivative measure of a single product of D generic linear forms in m variables?

The regime of interest is few variables and high degree: m fixed (or growing slowly), $D \rightarrow \infty$. This is precisely the regime of constant-variate algebraic circuits that controls the hardness–randomness bootstrapping phenomenon [1, 23, 15], where any measure-based approach would have to certify lower bounds against products of many linear forms in few variables; it is also the natural first family in the apolarity program of DiPasquale–Flores–Peterson [8], our point of departure on the commutative-algebra side.

The answer is: *at every fixed (k, ℓ) the generic product saturates once the degree clears a (provably necessary) polynomial threshold* — and we further compute the measure *exactly* for the structured sub-family of witnesses (products with disjoint-pair block structure) that provably does *not* saturate, with a two-sided sandwich locating those witnesses’ measure in the ambient-dominated range. (One question remains open: the exact generic value in the ambient-dominated range — shift budget below the overshoot threshold — beyond what the sandwich certifies; Remark 6.5.)

1.2 Main results

Throughout, “generic” means: for all choices of coefficients outside a proper Zariski-closed subset of the parameter space (here, of the D -fold product of the space of linear forms). Since $g \mapsto \Gamma_{k,\ell}(g)$ is a matrix rank in the coefficients of g , it is lower semicontinuous, and a single witness certifies the generic value from below; all our theorems are proved by constructing explicit witnesses inside the variety of products of linear forms.

The headline is the following description. (The names R / A of the regimes refer to which cap in (1) binds: *row-dominated* when $N_k N_\ell \leq N_{D-k+\ell}$, *ambient-dominated* otherwise.)

Theorem 1.1 (Saturation; informal roll-up). *Let $\text{char } \mathbb{F} = 0$ and let f be a product of D generic linear forms in m variables.*

- (a) **Two variables, exact, no thresholds.** *For $m = 2$, all D , and all $0 \leq \kappa \leq D$, $\lambda \geq 0$:*

$$\Gamma_{\kappa,\lambda}(f) = \min((\kappa + 1)(\lambda + 1), D - \kappa + \lambda + 1).$$

- (b) **Saturation at every fixed m .** *For every $m \geq 3$ and every (k, ℓ) there is $D_0(m, k, \ell) = \text{poly}(k, \ell)$ (for fixed m) such that for all $D \geq D_0$:*

$$\Gamma_{k,\ell}(f) = \min(N_k N_\ell, N_{D-k+\ell}).$$

- (c) **Derivative spaces, all k .** *For $\ell = 0$: $\dim \partial^k f = N_k$ exactly for $k \leq D/m$, and $\dim \partial^k f = N_{D-k}$ exactly for $k \geq D - D/m$; and for every $0 \leq k \leq D$, if $D \geq 2m^2$,*

$$\dim \partial^k f \geq \frac{(2/e)^{m-1}}{2e^2 m} \cdot \min(N_k, N_{D-k}).$$

- (d) **Ambient regime, quantified.** *For the explicit block witnesses of Theorem 1.2, the measure is sandwiched between the ambient mass of an explicit overshoot-budget event and the mass of the same event shrunk by a boundary shell of width $O(\sqrt{m(D-k+\ell)})$; in particular the witness meets the ambient cap exactly once the shift budget ℓ clears an explicit threshold, and — a matching impossibility — no such block witness ever exceeds the budget-event mass.*

The formal statements are Theorem 3.1 (part (a)), Theorem 7.1 (part (b)), Theorem 5.1 with Lemma 5.2 (part (c)), and Theorem 6.1 (part (d)).

Two further results describe the fine structure behind Theorem 1.1 and are, we believe, of independent interest.

Theorem 1.2 (Exact block calculus; informal). *Let $m = 2\mu$ and let $W = \prod_{j=1}^{\mu} f_j(u_j, v_j)$ be a product of generic binary forms in disjoint pairs of variables, each of degree $q = D/\mu$. Then $\Gamma_{k,\ell}(W)$ is computed exactly by an explicit filtration/increment formula (Theorem 4.3). Consequently:*

- (i) *in the ambient-dominated regime, W 's measure is located by an overshoot budget: profiles beyond the budget contribute nothing (a ceiling), profiles clearing the budget by a boundary shell contribute fully, and W meets the ambient cap exactly once ℓ clears an explicit threshold (Theorem 6.1);*
- (ii) *in the row-dominated regime, W is genuinely deficient: for even $m \geq 6$ and $k = \ell \geq m^2$ its measure falls below the row cap by a factor at least $(k/32m)^{(m-4)/2}$, while (for $k \leq D/2$, $D \geq 2m^2$) the window never exceeds $(\min(k, \ell) + 1)^{O(m)}$ (Theorem 6.6).*

Part (ii) is a structural finding, not a failure of analysis: products with disjoint-pair block structure can *never* certify row-cap saturation at $m \geq 3$, while (by Theorem 1.1(b)) generic products do saturate. The witnesses that close the row-dominated regime are *chains*,

$$g_1(x_1, x_2) g_2(x_2, x_3) \cdots g_{m-1}(x_{m-1}, x_m),$$

products of generic binary forms in *overlapping* pairs of variables: the shared variables destroy the multidegree grading that provably caps disjoint-pair products.

Theorem 1.3 (Chain saturation, uniform in m). *Let $\text{char } \mathbb{F} = 0$, $m \geq 3$, and fix (k, ℓ) . There are thresholds $H_1, \dots, H_{m-1}$ with $H_1 = (k + \ell + 1)(2k + \ell + 1) - 1$ and $H_j \leq 3(k + \ell + 1)^{\lceil (m-2)/2 \rceil} (k + 1) + O(k + \ell)$ for $2 \leq j \leq m - 1$ — exponential in m but polynomial in (k, ℓ) for every fixed m — such that for all $h_j \geq H_j$, all $h_1 \geq H_1$ outside a finite set (possibly depending on the other data), and generic binary forms $g_1(x_1, x_2), \dots, g_{m-1}(x_{m-1}, x_m)$ of degrees $h_1, \dots, h_{m-1}$:*

$$\Gamma_{k,\ell}(g_1 g_2 \cdots g_{m-1}) = N_k N_\ell = \binom{k + m - 1}{m - 1} \binom{\ell + m - 1}{m - 1},$$

the full row cap. Consequently there is a finite $D_0(m, k, \ell) \leq \sum_j H_j + E$ — where E is the size of the finite exceptional set for h_1 , so $D_0 = \sum_j H_j$ whenever that set is empty, as it was in every machine-checked configuration — such that a generic product of $D \geq D_0$ linear forms in m variables satisfies $\Gamma_{k,\ell} = \min(N_k N_\ell, N_{D-k+\ell})$.

This is Theorem 7.1; it is proved by one derivation uniform in m , organized around three reusable lemmas (Section 1.5).

1.3 The complexity-theoretic reading

Exact per-gate saturation. The known barriers for rank-type measures — Efremenko–Landsberg–Schenck–Weyman [11] (shifted partials cannot separate permanent from determinant), Efremenko–Garg–Oliveira–Wigderson [10] (rank methods in general), and Bhargav–Dutta–Saxena [6] (shifted-partial variants against constant-depth circuits) — show, each in its own setting, that these measures stall because random or generic gates nearly saturate them; for homogeneous depth four specifically, [6] places the obstruction at roughly $n^{O(\sqrt{d})}$ — precisely the depth-reduction threshold of Section 1.1 at which VP and VNP would separate. Those results operate in the many-variable regime and prove near-saturation up to unspecified or coarse factors. The present paper establishes the exact statement, with the cap met on the nose, for the single most structured gate family there is — one product of linear forms — in the few-variable high-degree regime. In particular:

Against a single product gate of D generic linear forms in m variables, in characteristic 0, the shifted partial derivative measure certifies nothing: for every fixed m , beyond a $\text{poly}(k, \ell)$ degree threshold its value equals the universal cap (1) exactly, at every (k, ℓ) , in every regime.

Any lower-bound strategy hoping to exploit shifted partials against depth-two ($\Sigma\Pi$) or product-of-low-variate structure in this regime must therefore be abandoned or must change the measure. The result quantifies exactly *how* blind the measure is: our block calculus (Theorem 1.2) even identifies a structured sub-family of products — disjoint-variable blocks — against which the measure retains a residual advantage of shape $(\min(k, \ell)/m)^{\Theta(m)}$ (in the ranges quantified by Theorem 6.6), and shows that advantage is destroyed by any overlap in the variable structure.

Where the regime matters. The few-variable, high-degree regime is not a curiosity: by the bootstrapping theorems of Agrawal–Ghosh–Saxena [1] and Kumar–Saptharishi–Tengse [23], explicit hitting sets for constant-variate circuits of barely non-trivial size would cascade to essentially complete derandomization of polynomial identity testing, and by Guo–Kumar–Saptharishi–Solomon [15] constant-variate hardness yields hitting sets. Measure-based approaches to such lower bounds would confront exactly the gates studied here. The saturation theorem closes that door definitively for plain shifted partials, at every fixed m .

Characteristic 0 is essential, not bookkeeping. Our proofs consume characteristic 0 in several load-bearing places: Euler-identity chain nesting, fat-point apolarity, zero-multiplicity bounds for exponential polynomials, and integer-node confluent Vandermondes. This is not an artifact. Over small finite fields, Armand–Behera–Tavenas [5] prove depth-five lower bounds using an *evaluation* variant of shifted partials, and the polarity there is reversed: over $\mathbb{F}_q$ with $q = O(1)$, a single product of affine forms of high rank vanishes to high order at almost all points, so its evaluated shifted partials have measure near *zero* — the opposite of saturation. The two results are complementary and disjoint in regime (their variable count grows polynomially in the degree; in the few-variable regime over $\mathbb{F}_q$ the entire evaluation space has constant dimension $\leq q^m$), and together they locate precisely the boundary at which product gates stop saturating: small characteristic plus the evaluation twist. See Section 11.

1.4 The commutative-algebra reading

For $f \in S_D$ with apolar ideal $f^\perp$, the measure $\Gamma_{k,\ell}(f)$ determines the Hilbert function of the ideal generated by the k -th derivative space: our theorems compute $\dim((\partial^k f) \cdot S_\ell)$ for generic products of linear forms, i.e. for generic *hyperplane multi-arrangement* forms. In the language of apolarity: writing $R_f = S/f^\perp$ for the apolar algebra, our Theorem 5.1 determines the Hilbert function of R_f itself (the case $\ell = 0$: $\dim(R_f)_k = \dim \partial^k f$), exactly on the ranges $k \leq D/m$ and $k \geq D - D/m$ and up to an explicit $2^{-\Theta(m)}$ factor in between; the general- ℓ theorems determine the Hilbert function of the ideal generated by the k -th derivative space.

DiPasquale–Flores–Peterson [8] initiated the systematic study of apolar algebras of products of linear forms, with Waring-rank bounds and detailed computations for special sextics; the generic Hilbert function questions answered here were, to our knowledge, open in every case $m \geq 3$ (for $m = 2$, the $\ell = 0$ Hilbert function of a binary form is classical; our contribution at $m = 2$ is the exact min-law for all shifts λ , Theorem 3.1). Two points deserve emphasis for readers arriving from this direction.

First, the proofs are elementary and self-contained. The natural first guesses — that row-cap saturation at $m \geq 3$ should follow from maximal-rank statements for generic subspaces (Fröberg’s conjecture territory [13, 4, 18, 7, 14]) or from fat-point interpolation (Alexander–Hirschowitz territory [3]) — both fail structurally: $\partial^k f$ is far from a generic subspace, and the fat-point mechanism that closes $m = 2$ provably cannot produce witnesses inside the variety of products (power sums leave the Chow variety at $m \geq 3$; Remark 3.2). Our route uses neither. The genericity actually consumed is bivariate-or-lower throughout: one coefficient-sequence lemma in a single variable, plus generic scalars.

Second, the block calculus (Theorem 4.3) gives, for products with disjoint-pair block structure, the *exact* Hilbert function of $S_\ell \cdot \partial^k$, in closed combinatorial form, for every (k, ℓ) — a complete answer for a natural family of non-generic arrangements, and the tool behind both the ambient saturation and the deficiency theorems.

1.5 Proof overview and organization

Two variables (Section 3). The exact min-law at $m = 2$ is proved in three steps: a power-sum witness $g^* = \sum_{r=1}^{\kappa+1} L_r^q$ whose derivative space is spanned by powers $L_r^{q-\kappa}$ (a Vandermonde argument); apolarity identifying the orthogonal complement of the shifted span with a degree slice of a fat-point ideal in $\mathbb{P}^1$; and the observation that fat-point ideals in $\mathbb{P}^1$ are *principal*, so the slice has exactly computable dimension. Semicontinuity plus the density of products among binary forms transfers the count to generic products.

Block calculus and the $\ell = 0$ theory (Sections 4–5). For disjoint-pair products the shifted-derivative space decomposes over pair-multidegree profiles; along each profile the per-pair contributions form nested chains (an Euler-identity argument), and a filtration identity computes the dimension of a sum of tensor products of chain members *exactly*. At $\ell = 0$ the formula collapses to a coefficient extraction $\dim \partial^k = [z^k] G(z)^\mu$ for an explicit palindromic polynomial G . Comparing this to $\min(N_k, N_{D-k})$ at all k is a self-contained combinatorial problem — comparing capped to uncapped composition counts — solved by Lemma 5.2: the ratio $f(k)/N_k$ is monotone (a Pólya-urn coupling pushes the uniform composition measure up a level *exactly*, and the cap event is a down-set), and the capped count is palindromic and log-concave (by Hoggar’s theorem [17]), which pins its midpoint value. We state and prove Lemma 5.2 for arbitrary caps, as a

standalone tool.

The two regimes for block gates (Section 6). The exact calculus splits cleanly. On the ambient side, a profile contributes at all only if its overshoot above the per-pair degree stays within the shift budget ℓ (a necessary condition), and contributes *fully* once the budget clears an additional $O(\sqrt{m(D-k+\ell)})$ boundary shell (sufficient): the gate’s ambient measure is thereby sandwiched by two explicit masses, and meets the ambient cap exactly beyond an explicit threshold. On the row side the same calculus is an exact *upper* bound that sits below the row cap by a factor of shape $(\min(k, \ell)/m)^{\Theta(m)}$ (precise ranges in Section 6): disjoint-pair witnesses cannot close the row-dominated regime, which forces the chain witnesses of the main theorem.

The chain theorem (Sections 7–10). The proof of Theorem 1.3 is one derivation, uniform in m , structured as three reusable lemmas plus explicit bookkeeping:

- **Lemma M (master reduction, Section 7).** For $f = g(x, y) \cdot G$ with G an *arbitrary* x -free co-factor, an explicit monomial bijection identifies the row space $\{x^\alpha \partial^\beta f\}$ with an explicit family of Laurent polynomials in a tag variable t that records the x -filtration depth. The identification is exact — no genericity, no thresholds. Applied at both ends of the chain it converts $\Gamma_{k, \ell}(g_1 \cdots g_{m-1})$ into the rank of a completely explicit two-tag family.
- **Lemma P (end peel, Section 8).** With the first factor specialized to a truncated multi-exponential $\sum_c \lambda_c e^{\omega_c w}$, a zero-multiplicity bound for exponential polynomials splits any kernel vector into single-frequency conditions; a Vandermonde in ω refines the kernel into graded classes; one scaling limit $h_1 \rightarrow \infty$ and a triangularization in the tag degree then reduce everything to the independence of a smaller “residue family.” The lemma is stated and proved with the entire rest of the chain as an abstract spectator: it therefore applies at every m with no induction.
- **Middle peel (Section 9).** All remaining factors are specialized to multi-geometric sequences. A telescoping identity shows the residue family’s columns are faithfully coordinatized by $m - 2$ integer coordinates, with index recovery in closed prefix and suffix forms; the exponential parts factor as pure tensors over the coordinates (prefix-product and suffix-ratio bases); a confluent Vandermonde on a product window plus an $(m - 2)$ -fold twist Vandermonde reduce independence to a per-group core identity.
- **Lemma S_m (residual core, Section 10).** The core identity — a falling-factorial polynomial identity in $m - 2$ formal variables — forces all coefficients to vanish, by one graded elimination in the unique load carrying the coordinate w followed by a separated graded-monic tensor argument. The lemma is valid over *any* field, for arbitrary index sets and constants; characteristic 0 is consumed only upstream.

The case $m = 3$, where the middle peel degenerates (there is no middle factor to specialize), is closed by the same Lemma-M/Lemma-P pattern with a one-variable coefficient-sequence lemma (Lemma 10.2) in place of the middle peel; it serves as the anchor of the general statement, and its proof is where the method was found.

Discovery path and worked instances. The uniform- m derivation was not found in one step: the cases $m = 3, 4, 5, 6$ were closed successively, each revealing one additional structural

ingredient — $m = 4$ forced the sequential (rather than symmetric) treatment of the two chain ends, $m = 5$ forced the multi-geometric specialization of *all* middle factors and revealed the merged frequency phenomenon, $m = 6$ revealed the telescoped global relation and the all-formal residual core that make the uniform derivation possible. Appendix A presents these instances with their exact (smaller) thresholds and records the two negative findings — the symmetric double limit provably loses rank, and evaluating the formal core parameters provably loses rank — that forced the architecture. Readers wanting a worked example before the general machinery may read Appendix A alongside Sections 7–10.

Machine verification. Every theorem in this paper is accompanied by exact numerical verification — integer or multi-prime modular arithmetic, not floating point — of every intermediate identity, at parameter ranges that exercise each joint of each proof, including the fully general code path of the uniform- m derivation at $m = 4, \dots, 9$ (over 5×10^5 monomial identities checked per configuration at the largest sizes). The derivations were additionally re-derived independently and re-verified from scratch at a second prime and seed. Appendix B records the protocol and the complete configuration tables. Verification scripts will be made available in a public repository.

Disclosure: how this work was produced. This research was not carried out by the author alone: it was conducted in heavy reliance on Claude (Fable 5), a large language model built by Anthropic, over an extended sequence of recorded research sessions. The model produced the candidate derivations, the proofs, the verification code, and the text of this manuscript; the author directed the program — the questions asked, the standards of rigor, and the decisions about what to pursue, what to verify, what to retract, and what to claim — and takes full responsibility for the correctness of everything stated here. Because both the derivations and their adversarial re-derivation passes were model-executed (in separate sessions, from the statements alone, with independent implementations at different primes and seeds), the trust chain deliberately does not rest on the model’s prose: it rests on the exact machine verification of every intermediate identity, as described above and catalogued in Appendix B, together with proofs we have aimed to keep short and elementary enough to be checked line by line by a human reader.

Organization. Section 2 fixes notation and records the caps and the semicontinuity principle. Sections 3–6 develop the exact theory: the $m = 2$ min-law, the block calculus, the $\ell = 0$ theory with the capped-composition lemma, and the two-regime analysis. Sections 7–10 prove the chain saturation theorem, uniformly in m . Section 11 discusses related work. Appendix A presents the worked instances $m \leq 6$; Appendix B the verification record.

2 Preliminaries

2.1 Notation

$\mathbb{F}$ is a field with $\text{char } \mathbb{F} = 0$; all results hold verbatim over any such field (ranks of matrices with entries in the prime field $\mathbb{Q}$ are field-stable), and we freely pass to an algebraic closure when convenient. $S = \mathbb{F}[x_1, \dots, x_m]$ with graded pieces S_e ;

$$N_e := \dim S_e = \binom{e + m - 1}{m - 1}.$$

For a multi-index $\beta \in \mathbb{N}^m$ we write $\partial^\beta = \partial_{x_1}^{\beta_1} \cdots \partial_{x_m}^{\beta_m}$ and $x^\alpha = x_1^{\alpha_1} \cdots x_m^{\alpha_m}$. For homogeneous $f \in S_D$:

$$\partial^k f := \text{span}\{\partial^\beta f : |\beta| = k\} \subseteq S_{D-k}, \quad \Gamma_{k,\ell}(f) := \dim S_\ell \cdot \partial^k f.$$

Falling factorials are written $x^{\underline{j}} := x(x-1) \cdots (x-j+1)$; note $x^{\underline{j}}$ is a degree- j polynomial in x with leading coefficient 1, and $n^{\underline{j}} = 0$ for integers $0 \leq n < j$.

Binomial-sum polynomials appear repeatedly: $B_a(z) := 1 + z + \cdots + z^a$.

2.2 The universal caps

For every homogeneous f of degree D and all (k, ℓ) :

$$\Gamma_{k,\ell}(f) \leq \min(N_k N_\ell, N_{D-k+\ell}), \quad \dim \partial^k f \leq \min(N_k, N_{D-k}). \quad (2)$$

The first inequality is the span/ambient count. In the second, the bound $\dim \partial^k f \leq N_{D-k}$ is the ambient count, and $\dim \partial^k f \leq N_k$ is the span count; the symmetric refinement $\dim \partial^k f = \text{rank}(\text{Cat}_k f) = \text{rank}(\text{Cat}_{D-k} f) = \dim \partial^{D-k} f$ (catalecticant symmetry) holds in characteristic 0.

For a product of linear forms $f = \prod_{i=1}^D \ell_i$ there is one further cap, $\dim \partial^k f \leq \binom{D}{k}$ (each $\partial^\beta f$ with $|\beta| = k$ is a combination of the *subset products* $\prod_{i \notin T} \ell_i$, $|T| = k$). In the few-variable regime this cap never binds: for $m-1 \leq k \log \frac{D}{k} / \log(k+1)$ — in particular throughout the regime $m = o(\log D)$, $1 \leq k \leq D/2$ — one has $N_k \leq (k+1)^{m-1} \leq \binom{D}{k}$, and the range $k > D/2$ mirrors by catalecticant symmetry. We therefore compare against (2) throughout.

Definition 2.1 (Saturation; regimes). $f \in S_D$ *saturates* the shifted partial derivative measure at (k, ℓ) if equality holds in the first inequality of (2). The pair (k, ℓ) (for given m, D) is *row-dominated* if $N_k N_\ell \leq N_{D-k+\ell}$ and *ambient-dominated* otherwise.

For fixed (k, ℓ) and $D \rightarrow \infty$ the row cap is constant while $N_{D-k+\ell} = \Theta_m(D^{m-1})$, so every fixed (k, ℓ) is eventually row-dominated; the ambient-dominated regime requires k growing with D .

2.3 Semicontinuity and the variety of products

$\Gamma_{k,\ell}(g)$ is the rank of a matrix whose entries are linear in the coefficients of g ; hence $g \mapsto \Gamma_{k,\ell}(g)$ is lower semicontinuous in the Zariski topology, and for any irreducible parameter variety X of polynomials, a single witness $g^* \in X$ with $\Gamma_{k,\ell}(g^*) \geq c$ certifies $\Gamma_{k,\ell}(g) \geq c$ for generic $g \in X$.

We use two parameter varieties: the full space S_D (irreducible), and the *Chow variety* of degree- D products of linear forms, i.e. the image of the multiplication map $(S_1)^{\times D} \rightarrow S_D$ — irreducible, as the image of an irreducible variety. Two facts about it are used repeatedly:

1. every binary form is a product of linear forms (over $\overline{\mathbb{F}}$), so for $m = 2$ the Chow variety is all of S_D and witnesses may be chosen freely;
2. for any m , a product of polynomials each of which is a product of linear forms is a point of the Chow variety; in particular products of binary forms in subsets of the variables (our block and chain witnesses) are Chow points, since their binary factors split over $\overline{\mathbb{F}}$. (When a factor is later specified as a *univariate* polynomial — as in the specializations of Sections 8–10 — we mean its binary homogenization, which is again a binary form.)

Consequently, to prove a statement about products of D generic *linear forms*, it suffices to exhibit one witness that is a product of binary forms (of any shape) with the desired measure.

2.4 Euler nesting

We record the one identity used to organize shifted-derivative spaces into chains. For homogeneous h of degree $e \geq 1$, Euler's identity $e h = \sum_i x_i \partial_i h$ gives $h \in S_1 \cdot \partial^1 h$; applying this to each $\partial^{\kappa-1}$ -derivative of f (homogeneous of degree $\deg f - \kappa + 1 \geq 1$),

$$S_{\lambda-1} \cdot \partial^{\kappa-1} f \subseteq S_\lambda \cdot \partial^\kappa f \quad (1 \leq \kappa \leq \deg f, \lambda \geq 1), \quad (3)$$

valid in characteristic 0. (The restriction $\kappa \leq \deg f$ matters: at $\kappa = \deg f + 1$ the left side is $S_{\lambda-1} \neq 0$ while the right side vanishes.) Thus for fixed $k - \ell$ the spaces $S_\ell \cdot \partial^k f$ increase along $(\kappa, \lambda) = (k - \ell + \lambda, \lambda)$: they form a chain. This elementary observation is the backbone of the block calculus of Section 4.

3 Two variables: the exact min-law

In two variables the answer is exact, for every pair of parameters and every degree, with no thresholds.

Theorem 3.1 (Exact min-law at $m = 2$). *Let $\text{char } \mathbb{F} = 0$ and $m = 2$. For generic $g \in S_q$ — in particular for a generic product of q distinct linear forms — and every $0 \leq \kappa \leq q$ and $\lambda \geq 0$:*

$$\Gamma_{\kappa, \lambda}(g) = \min((\kappa + 1)(\lambda + 1), q - \kappa + \lambda + 1).$$

Here $(\kappa + 1)(\lambda + 1) = N_\kappa N_\lambda$ is the row cap and $q - \kappa + \lambda + 1 = N_{q-\kappa+\lambda}$ the ambient cap, so the theorem says: generic binary forms saturate at every parameter pair. Since every binary form is a product of linear forms (Section 2.3), “generic product” and “generic form” certify each other.

Proof. The inequality $\leq$ is the universal cap (2). For $\geq$ it suffices, by semicontinuity, to exhibit a single witness in S_q ; forms with nonzero discriminant are dense in S_q and factor into distinct linear forms, so the generic-product statement follows as well.

Witness. Take

$$g^* = \sum_{r=1}^{\kappa+1} L_r^q, \quad L_r = x + \theta_r y,$$

with $\theta_1, \dots, \theta_{\kappa+1}$ distinct and generic.

Step 1: the derivative space. For $a + b = \kappa$,

$$\partial_x^a \partial_y^b L_r^q = q^\kappa \theta_r^b L_r^{q-\kappa},$$

so $\partial^\kappa g^* \subseteq \text{span}\{L_r^{q-\kappa} : 1 \leq r \leq \kappa + 1\}$ (the scalar $q^\kappa = q(q-1) \cdots (q-\kappa+1)$ is nonzero in characteristic 0 since $\kappa \leq q$), and the $(\kappa + 1) \times (\kappa + 1)$ coefficient matrix $[\theta_r^b]_{r,b}$ is an invertible Vandermonde matrix (the θ_r are distinct), so equality holds:

$$\partial^\kappa g^* = \text{span}\{L_r^{q-\kappa} : 1 \leq r \leq \kappa + 1\}.$$

Step 2: the unbalanced range. Suppose $\kappa > q/2$, so $\kappa + 1 \geq q - \kappa + 1 = \dim S_{q-\kappa}$. Any $q - \kappa + 1$ of the powers $L_r^{q-\kappa}$ are linearly independent (a Vandermonde in the θ_r again), hence $\partial^\kappa g^* = S_{q-\kappa}$, so $S_\lambda \cdot \partial^\kappa g^* = S_{q-\kappa+\lambda}$ and $\Gamma_{\kappa, \lambda}(g^*) = q - \kappa + \lambda + 1$, which is the min in this range

(the row cap $(\kappa + 1)(\lambda + 1)$ exceeds the ambient cap). So assume $\kappa \leq q/2$ from now on; then $\dim \partial^\kappa g^* = \kappa + 1$ by the same independence.

Step 3: fat points in $\mathbb{P}^1$. Equip S with the apolarity pairing $\langle x^i y^j, x^k y^l \rangle = \delta_{(i,j),(k,l)} i! j!$ — perfect on each graded piece, being diagonal with entries $i! j! \neq 0$ in characteristic 0, so orthogonal complements have complementary dimension — under which $\langle u \cdot v, h \rangle = \langle v, u(\partial)h \rangle$ and, for $\deg u = e$, $\langle L_r^e, u \rangle = e! u(1, \theta_r)$. For $h \in S_{q-\kappa+\lambda}$:

$$\begin{aligned} h \perp S_\lambda \cdot \partial^\kappa g^* &\iff \langle x^c y^{\lambda-c} L_r^{q-\kappa}, h \rangle = 0 \quad \forall 0 \leq c \leq \lambda, \forall r \\ &\iff \langle L_r^{q-\kappa}, \partial_x^c \partial_y^{\lambda-c} h \rangle = 0 \quad \forall c, r \\ &\iff \text{every order-}\lambda \text{ partial of } h \text{ vanishes at } (1, \theta_r), \forall r \\ &\iff h \text{ vanishes to order } \geq \lambda + 1 \text{ at each point } [1 : \theta_r] \in \mathbb{P}^1 \\ &\iff h \in (F)_{q-\kappa+\lambda}, \quad F := \left(\prod_{r=1}^{\kappa+1} (y - \theta_r x)^{\lambda+1} \right), \end{aligned}$$

where the fourth equivalence uses characteristic 0 (vanishing of all order- λ partials at a point is vanishing to order $\lambda + 1$ there).

Step 4: principality. F is a *principal* ideal generated in degree $(\kappa + 1)(\lambda + 1)$, so its degree- $(q - \kappa + \lambda)$ slice has dimension exactly $\max(0, q - \kappa + \lambda + 1 - (\kappa + 1)(\lambda + 1))$. Hence $\Gamma_{\kappa,\lambda}(g^*) = (q - \kappa + \lambda + 1) - \max(0, q - \kappa + \lambda + 1 - (\kappa + 1)(\lambda + 1)) = \min((\kappa + 1)(\lambda + 1), q - \kappa + \lambda + 1)$. $\square$

Remark 3.2 (What makes $m = 2$ special). The proof yields both regimes at once because a fat-point ideal in $\mathbb{P}^1$ is principal: $(\lambda + 1)$ -fold points at $\kappa + 1$ generic points of $\mathbb{P}^1$ impose independent conditions in every degree $\geq (\kappa + 1)(\lambda + 1) - 1$, trivially. Both special features of $m = 2$ fail at $m \geq 3$, in an instructive way. First, the witness $\sum_r L_r^q$ is a product of linear forms only when $m = 2$: at $m \geq 3$ the Chow variety is a proper subvariety of S_q and power sums of $\kappa + 1 \geq 3$ summands generically lie outside it (two summands are a binary form in disguise and always split), while semicontinuity certifies only the family containing the witness. Second, the natural substitute — products whose derivative spaces impose independent fat-point conditions — turns row-cap saturation at $m \geq 3$ into a Hilbert-function-of-fat-points problem (Alexander–Hirschowitz territory [3]: solved for double points, open in general) *with the added constraint that the witness lie on the Chow variety*. The chain witnesses of Sections 7–10 bypass this entirely: no fat-point statement, and no maximal-rank conjecture, is used anywhere at $m \geq 3$.

4 Disjoint-pair products: an exact calculus

We now compute the measure *exactly* for products with disjoint-pair block structure. Throughout this section let $m = 2\mu$ be even (odd m works identically with $\lfloor m/2 \rfloor$ pairs plus one singleton variable, whose per-degree graded pieces are one-dimensional; we suppress this case to lighten notation). Pair the variables as (u_j, v_j) , $1 \leq j \leq \mu$, write $S^{(j)}$ for $\mathbb{F}[u_j, v_j]$, and let

$$W = \prod_{j=1}^{\mu} f_j(u_j, v_j),$$

where each f_j is a generic product of $q := D/\mu$ distinct binary linear forms. W is a product of D linear forms, i.e. a point of the Chow variety.

4.1 Profile decomposition and chains

Since $\partial^\beta W = \prod_j \partial^{\beta_j} f_j$ and monomial shifts factor across the pairs, the space $S_\ell \cdot \partial^k W$ decomposes over *profiles*: tuples $\delta = (\delta_1, \dots, \delta_\mu)$ with $\sum_j \delta_j = D - k + \ell$ recording the total degree landing in each pair. Writing $\delta_j = q - o_j$ (so $\sum_j o_j = k - \ell$), the degree- δ component is

$$V_\delta = \sum_{\substack{\kappa, \lambda \geq 0 \\ \sum_j \kappa_j = k, \kappa_j - \lambda_j = o_j}} \bigotimes_{j=1}^\mu V_{\kappa_j, \lambda_j}^{(j)}, \quad V_{\kappa, \lambda}^{(j)} := S_\lambda^{(j)} \cdot \partial^\kappa f_j \subseteq S_{\delta_j}^{(j)}, \quad (4)$$

and distinct profiles contribute a direct sum: $\Gamma_{k, \ell}(W) = \sum_\delta \dim V_\delta$.

Two structural facts make (4) computable.

Lemma 4.1 (Chain nesting). *In characteristic 0, $V_{\kappa-1, \lambda-1}^{(j)} \subseteq V_{\kappa, \lambda}^{(j)}$ for all $1 \leq \kappa \leq q$ and $\lambda \geq 1$. Hence along a fixed profile, the summands of (4) for pair j form a chain indexed by $\nu_j := \kappa_j$ alone (with $\lambda_j = \nu_j - o_j$), starting at $\nu_j \geq o_j^+ := \max(o_j, 0)$, with dimensions*

$$v_j(\nu) = \min((\nu + 1)(\nu - o_j + 1), \delta_j + 1)$$

and increments

$$\text{inc}_j(\nu) := v_j(\nu) - v_j(\nu - 1) = \min(2\nu - o_j + 1, \delta_j + 1 - v_j(\nu - 1))$$

(with $v_j(o_j^+ - 1) := 0$, and $\text{inc}_j(\nu) = 0$ once the cap $\delta_j + 1$ is reached or $\nu > q$).

Proof. The nesting is (3) applied in the pair (u_j, v_j) . The dimension formula is Theorem 3.1 applied to the generic binary form f_j with $(\kappa, \lambda) = (\nu, \nu - o_j)$: row cap $(\nu + 1)(\nu - o_j + 1)$, ambient cap $q - \nu + (\nu - o_j) + 1 = \delta_j + 1$. The increment formula is immediate. $\square$

Lemma 4.2 (Filtration identity). *Let $\{V_\nu^{(j)}\}_{\nu \geq s_j}$, $1 \leq j \leq \mu$, be ascending chains of subspaces of vector spaces $U^{(j)}$ ($V_{\nu-1}^{(j)} \subseteq V_\nu^{(j)}$; set $V_{s_j-1}^{(j)} = 0$), and let $\mathcal{S} \subseteq \prod_j \{\nu \geq s_j\}$ be down-closed (if $\nu \in \mathcal{S}$ and $s_j \leq \nu'_j \leq \nu_j$ for all j then $\nu' \in \mathcal{S}$). Then*

$$\dim \sum_{\nu \in \mathcal{S}} \bigotimes_{j=1}^\mu V_{\nu_j}^{(j)} = \sum_{\nu \in \mathcal{S}} \prod_{j=1}^\mu (\dim V_{\nu_j}^{(j)} - \dim V_{\nu_j-1}^{(j)}).$$

Proof. Induction on μ . For $\mu = 1$, $\mathcal{S}$ is an interval $[s_1, N]$ and the sum of the chain is $V_N^{(1)}$, whose dimension telescopes.

For $\mu \geq 2$, we first record the two-chain principle: if $A_0 \subseteq A_1 \subseteq \dots$ ascends and $B_0 \supseteq B_1 \supseteq \dots$ descends, then

$$\dim \sum_i A_i \otimes B_i = \sum_i (\dim A_i - \dim A_{i-1}) \dim B_i.$$

Indeed, choose complements $A_i = A_{i-1} \oplus C_i$; then $\sum_i A_i \otimes B_i = \sum_i (\bigoplus_{j \leq i} C_j) \otimes B_i = \sum_j C_j \otimes \sum_{i \geq j} B_i = \sum_j C_j \otimes B_j$, using that B descends; and this last sum is direct because the C_j are independent in the first factor, giving $\dim = \sum_j \dim C_j \dim B_j$.

Now slice $\mathcal{S}$ on the first coordinate: $\mathcal{S} = \bigsqcup_{\nu_1} \{\nu_1\} \times \mathcal{S}_{\nu_1}$, where each slice $\mathcal{S}_{\nu_1} \subseteq \prod_{j \geq 2} \{\nu \geq s_j\}$ is down-closed and the slices *descend* in ν_1 ($\mathcal{S}_{\nu_1} \supseteq \mathcal{S}_{\nu_1+1}$), both because $\mathcal{S}$ is down-closed. Set $B_{\nu_1} := \sum_{\nu' \in \mathcal{S}_{\nu_1}} \bigotimes_{j \geq 2} V_{\nu'_j}^{(j)}$; these descend in ν_1 . The two-chain principle with $A_{\nu_1} = V_{\nu_1}^{(1)}$ gives

$$\dim \sum_{\nu \in \mathcal{S}} \bigotimes_j V_{\nu_j}^{(j)} = \sum_{\nu_1} (\dim V_{\nu_1}^{(1)} - \dim V_{\nu_1-1}^{(1)}) \dim B_{\nu_1},$$

and the inductive hypothesis evaluates each $\dim B_{\nu_1}$ as $\sum_{\nu' \in \mathcal{S}_{\nu_1}} \prod_{j \geq 2} (\dim V_{\nu'_j}^{(j)} - \dim V_{\nu'_j-1}^{(j)})$. Combining the two sums gives the claim. $\square$

4.2 The exact formula

Theorem 4.3 (Block calculus). *Let $\text{char } \mathbb{F} = 0$ and let $W = \prod_{j=1}^{\mu} f_j(u_j, v_j)$ with f_j generic binary forms of degree q (equivalently, generic products of q distinct binary linear forms). Then, with notation as above,*

$$\Gamma_{k,\ell}(W) = \sum_{\substack{o \in \mathbb{Z}^{\mu} \\ \sum_j o_j = k-\ell}} \sum_{\substack{\nu_j \geq o_j^+ \ \forall j \\ \sum_j \nu_j \leq k}} \prod_{j=1}^{\mu} \text{inc}_j(\nu_j), \quad (5)$$

where inc_j is as in Lemma 4.1 (depending on o_j and $\delta_j = q - o_j$; by convention $\text{inc}_j(\nu) = 0$ past the ambient cap or past $\nu > q$).

Proof. Fix a profile δ (equivalently o). We must show $\dim V_{\delta}$ equals the inner sum; summing over profiles then gives (5) (the outer sum is over profiles, written as their o -tuples; profiles with empty inner index set contribute 0).

By Lemma 4.1, the summands of (4) are $\bigotimes_j V_{\nu_j}^{(j)}$ (writing $V_{\nu}^{(j)}$ for $V_{\nu, \nu-o_j}^{(j)}$) over the down-closed index set $\mathcal{S} = \{\nu : \nu_j \geq o_j^+, \sum_j \nu_j \leq k\}$: down-closedness is immediate, and the constraint $\sum \nu_j \leq k$ comes from $\sum \kappa_j = k$ together with the nesting (a summand with $\sum \kappa_j < k$ arises from one with $\sum \kappa_j = k$ by nesting, and conversely including all $\sum \nu_j \leq k$ does not change the sum, again by nesting). Lemma 4.2 computes the dimension of the sum as $\sum_{\nu \in \mathcal{S}} \prod_j \text{inc}_j(\nu_j)$, which is the inner sum. $\square$

Remark 4.4 (Two-sided exactness; certified values). Formula (5) is an equality, not a pair of bounds: the $\geq$ direction is Lemma 4.1 plus Theorem 3.1 per pair (each chain member has exactly the claimed dimension), and the $\leq$ direction is Lemma 4.2 read as an upper bound (each component V_{δ} lies inside the corresponding sum of chain tensor products). By semicontinuity, every value of (5) is also a *lower* bound on $\Gamma_{k,\ell}$ of a generic (non-block) product of D linear forms. The formula reproduced the measured rank in every machine-checked configuration; see Appendix B.

Example 4.5. At $m = 4$ ($\mu = 2$), $D = 12$ ($q = 6$): formula (5) evaluates to 36, 80, 128, 212, 443 at $(k, \ell) = (2, 1), (2, 2), (3, 2), (3, 3), (6, 6)$ respectively, matching the measured ranks exactly. For comparison, the caps $\min(N_k N_{\ell}, N_{D-k+\ell})$ at these parameters are 40, 100, 200, 400, 455: the block witness is *exactly* computable and *genuinely* below the row cap at $(3, 3)$ — see Theorem 6.6 — while nearly saturating the ambient cap at $(6, 6)$ (443 of 455; see Theorem 6.1).

5 Derivative spaces: the case $\ell = 0$

At $\ell = 0$ the block calculus collapses to a coefficient extraction, and the comparison with the cap $\min(N_k, N_{D-k})$ becomes a self-contained combinatorial problem about *capped compositions*, which we solve — for arbitrary caps — in Lemma 5.2. The lemma may be of independent use.

5.1 The generating-function formula

At $\ell = 0$ each profile in (4) carries a single split ($\lambda_j = 0$, $\nu_j = \kappa_j = o_j$), and Theorem 4.3 reads

$$\dim \partial^k W = [z^k] G(z)^\mu, \quad G(z) = \sum_{\kappa=0}^q \min(\kappa + 1, q - \kappa + 1) z^\kappa = B_{\lceil q/2 \rceil}(z) \cdot B_{\lfloor q/2 \rfloor}(z), \quad (6)$$

where $B_a(z) = 1 + z + \dots + z^a$; the product form is a direct check ($[z^\kappa] B_{\lceil q/2 \rceil} B_{\lfloor q/2 \rfloor}$ counts lattice points and produces the tent function $\min(\kappa + 1, q - \kappa + 1)$). Thus $\dim \partial^k W$ is the number of weak compositions of k into $m = 2\mu$ parts with caps $\lceil q/2 \rceil$ (μ parts) and $\lfloor q/2 \rfloor$ (μ parts).

Theorem 5.1 (Derivative spaces of generic products). *Let $\text{char } \mathbb{F} = 0$ and let f be a generic product of D linear forms in m variables, $m \mid D$ (the general case follows by adjusting one block; we omit the routine bookkeeping). Then:*

- (a) $\dim \partial^k f = N_k$ for all $k \leq D/m$, and $\dim \partial^k f = N_{D-k}$ for all $k \geq D - D/m$;
- (b) for all $0 \leq k \leq D$: $\dim \partial^k f \geq [z^k] G(z)^\mu$ (explicit and computable), with G as in (6), $q = 2D/m$, and $\mu = \lfloor m/2 \rfloor$ (odd m : $\dim \partial^k f \geq [z^k] G(z)^\mu B_{D/m}(z)$ — μ pairs of degree q plus the singleton factor $x_m^{D/m}$, so that all m caps equal D/m);
- (c) for all $0 \leq k \leq D$:

$$\dim \partial^k f \geq C \cdot \min(N_k, N_{D-k}), \quad C := \frac{\prod_j (c_j + 1)}{(D + 1) N_{\lfloor D/2 \rfloor}},$$

with $c_1, \dots, c_m$ the caps of the witness in (b); and $C \geq \frac{(2/e)^{m-1}}{2e^2 m}$ once $D \geq 2m^2$.

Part (c) is a two-sided $2^{-\Theta(m)}$ comparison: together with the universal cap it pins $\dim \partial^k f$ at every single k , up to a factor depending only on m — in the few-variable regime $m = o(\log D)$, a factor $D^{-o(1)}$.

Proof of (a) and (b). (b) is (6) plus semicontinuity (Section 2.3: the block witness is a Chow point). For (a): if $k \leq D/m = q/2$ then every composition $\kappa \vdash k$ into m parts automatically satisfies $\kappa_j \leq q/2$, so the caps never bind and $[z^k] G^\mu = [z^k] (1 - z)^{-m} = N_k$, matching the span cap (2); the mirror range follows from catalecticant symmetry $\dim \partial^k f = \dim \partial^{D-k} f$. (The subset-product cap $\binom{D}{k}$ does not bind here: $k \leq D/m$ gives $D \geq mk \geq k + m - 1$, so $\binom{D}{k} \geq \binom{k+m-1}{k} = N_k$.) $\square$

The proof of (c) rests on the following standalone lemma.

5.2 A comparison lemma for capped compositions

For caps $c_1, \dots, c_m \in \mathbb{N}$ with $D = \sum_j c_j$, let

$$f(k) := [z^k] \prod_{j=1}^m B_{c_j}(z) = \#\{\kappa \in \mathbb{N}^m : \sum_j \kappa_j = k, \kappa_j \leq c_j \forall j\}$$

be the capped composition count, and $N_k = \binom{k+m-1}{m-1}$ the uncapped count.

Lemma 5.2 (Capped-composition comparison). *For any caps $c_1, \dots, c_m$:*

- (i) $h(k) := f(k)/N_k$ is nonincreasing on $[0, D]$;
- (ii) f is palindromic about $D/2$ and log-concave; hence $\max_k f(k) = f(\lfloor D/2 \rfloor) \geq \prod_j (c_j + 1)/(D + 1)$;
- (iii) for every $0 \leq k \leq D$:

$$f(k) \geq \frac{\prod_j (c_j + 1)}{(D + 1) N_{\lfloor D/2 \rfloor}} \cdot \min(N_k, N_{D-k}).$$

Proof. (i) $h(k)$ is the probability that a *uniformly random* weak composition of k into m parts lies under the caps (f counts the down-set, N_k the level). Consider the Pólya-urn kernel: from $x \in \mathbb{N}^m$ at level k , move to $x + e_j$ with probability $(x_j + 1)/(k + m)$. This is a valid kernel since $\sum_j (x_j + 1) = k + m$, and it pushes the uniform measure on level k to the uniform measure on level $k + 1$ *exactly*: the mass arriving at y on level $k + 1$ is

$$\sum_{j: y_j \geq 1} \frac{1}{N_k} \cdot \frac{y_j}{k + m} = \frac{k + 1}{(k + m) N_k} = \frac{1}{N_{k+1}},$$

using $N_{k+1}/N_k = (k + m)/(k + 1)$. The kernel never decreases a coordinate, and the cap event $\{x_j \leq c_j \forall j\}$ is a down-set, so its probability is nonincreasing along the coupling: $h(k + 1) \leq h(k)$. No structure of the caps is used.

(ii) Each B_c is palindromic, and the centers $c_j/2$ sum to $D/2$, so $\prod_j B_{c_j}$ is palindromic about $D/2$. Each B_c has positive log-concave coefficients (all ones), and products of polynomials with positive log-concave coefficient sequences (no internal zeros) have log-concave coefficients by Hoggar's theorem [17]; induction over the m factors gives f log-concave with no internal zeros on $[0, D]$. For a symmetric log-concave positive sequence, the ratios $r_k = f(k)/f(k - 1)$ are nonincreasing and satisfy $r_{D+1-k} = 1/r_k$, so $r_k \geq 1$ for $k \leq (D + 1)/2$: f is nondecreasing up to the middle, and its maximum sits at $\lfloor D/2 \rfloor$ ($= \lceil D/2 \rceil$ value by symmetry). Hence

$$f(\lfloor D/2 \rfloor) \geq \frac{1}{D + 1} \sum_{k=0}^D f(k) = \frac{\prod_j (c_j + 1)}{D + 1}.$$

(iii) For $k \leq \lfloor D/2 \rfloor$: $f(k) = h(k)N_k \geq h(\lfloor D/2 \rfloor)N_k$ by (i), and $h(\lfloor D/2 \rfloor) \geq \prod_j (c_j + 1)/((D + 1)N_{\lfloor D/2 \rfloor})$ by (ii); here $\min(N_k, N_{D-k}) = N_k$. For $k \geq \lceil D/2 \rceil$ apply the same at $D - k$ and use $f(k) = f(D - k)$. $\square$

Remark 5.3. The routes one might try first — Bonferroni-type inclusion–exclusion over cap violations, negative-association inequalities, or pointwise generating-function bounds — all fail here (each was tried and each loses more than a constant per coordinate in the balanced range). The coupling in (i) is what makes the lemma cap-structure-free, and appears to be new in this use.

Proof of Theorem 5.1(c). Apply Lemma 5.2 to the block witness’s caps: μ copies each of $\lceil q/2 \rceil$ and $\lfloor q/2 \rfloor$ (total D , $m = 2\mu$ parts); for odd m the witness of (b) has m caps all equal to D/m , so $\prod_j (c_j + 1) = (D/m + 1)^m \geq (D/m)^m$ and the chain below applies verbatim. Then $\prod_j (c_j + 1) \geq ((q+1)/2)^m \geq (D/m)^m$ and $N_{\lfloor D/2 \rfloor} \leq (e(D/2 + m)/(m-1))^{m-1}$, so the constant $C = \prod_j (c_j + 1) / ((D+1)N_{\lfloor D/2 \rfloor})$ satisfies

$$C \geq \left(\frac{2}{e}\right)^{m-1} \left(1 - \frac{1}{m}\right)^{m-1} \cdot \frac{1}{m} \cdot \frac{D}{D+1} \cdot \left(1 + \frac{2m}{D}\right)^{-(m-1)} \geq \frac{(2/e)^{m-1}}{2e^2 m} \quad \text{for } D \geq 2m^2,$$

using $(1 - 1/m)^{m-1} \geq e^{-1}$, $D/(D+1) \geq 1/2$, and $(1 + 2m/D)^{m-1} \leq e$. Combining with Theorem 5.1(b) and Lemma 5.2(iii) gives (c). The bound holds for all k , subsumes the middle range, and degrades gracefully to the exact ends: for $k \leq \lfloor q/2 \rfloor$ one has $h \equiv 1$, re-proving (a) from the block formula. $\square$

Remark 5.4 (The block value is not the truth). The witness bound (b) is what the *block* witness certifies; for the generic product itself, machine experiments consistently measure the full cap $\min(N_k, N_{D-k})$ even at $k = D/2$ (e.g. 84/84 at $m = 4$, $D = 12$, where the block witness gives 44). We do not need, and do not prove, more than (c) at general k ; at every fixed (k, ℓ) with $\ell \geq 0$ the row-dominated saturation theorem (Theorem 7.1) gives the exact value $N_k N_\ell$ for large D , which covers the $\ell = 0$ ranges a fortiori.

6 The two regimes for disjoint-pair products

The exact calculus of Theorem 4.3 splits cleanly along the two caps. On the ambient side, the block witness saturates up to an explicitly located boundary mass — and *exactly*, once the shift budget clears an explicit threshold. On the row side it is genuinely, quantifiably deficient. Throughout this section W , μ , q , profiles δ , and inc_j are as in Section 4, and we write

$$R := D - k + \ell, \quad T(\delta) := \sum_{j=1}^{\mu} (\delta_j - q)^+,$$

and let $\mathbb{P}$ denote the *profile measure*: the probability measure on $\{\delta \in \mathbb{N}^\mu : \sum_j \delta_j = R\}$ with $\mathbb{P}(\delta) = \prod_j (\delta_j + 1) / N_R$ (this is a probability measure because $\sum_\delta \prod_j (\delta_j + 1) = [z^R](1-z)^{-2\mu} = N_R$). The statistic T is the total *overshoot* of the profile above the per-pair degree q .

6.1 Ambient regime: saturation up to boundary mass

Theorem 6.1 (Ambient-side sandwich). *Let $\text{char } \mathbb{F} = 0$ and let W be the block witness. Then:*

- (i) (Upper bound; overshoot is necessary.)

$$\Gamma_{k,\ell}(W) \leq N_R \cdot \mathbb{P}[T(\delta) \leq \ell].$$

- (ii) (Lower bound; overshoot plus a boundary shell is sufficient.) With $\beta(\delta) := \sum_j (\lceil \sqrt{\delta_j + 1} \rceil + 1)$ and the deterministic bound $\beta(\delta) \leq \bar{\beta} := \sqrt{m(R+m)} + 2m$:

$$\Gamma_{k,\ell}(W) \geq N_R \cdot \mathbb{P}[T(\delta) \leq \ell - \beta(\delta)] \geq N_R \cdot \mathbb{P}[T(\delta) \leq \ell - \bar{\beta}].$$

- (iii) (Exact full saturation.) If $\ell \geq (R-q)^+ + \bar{\beta}$, then $\Gamma_{k,\ell}(W) = N_R$ exactly.

Thus the gate's ambient deficit is located *exactly*: it is the profile mass within a boundary shell of width $\beta(\delta) = O(\sqrt{mR})$ of the overshoot budget ℓ , plus the mass beyond the budget — which no product with this block structure can ever reach, by (i).

Proof. (i) Fix a profile δ and suppose the component V_δ of (4) is nonzero, i.e. some summand $\bigotimes_j V_{\kappa_j, \lambda_j}^{(j)}$ with $\sum \kappa_j = k$, $\sum \lambda_j = \ell$ is nonzero. Landing in pair-degree δ_j forces $\delta_j = q - \kappa_j + \lambda_j$, so $\lambda_j = \delta_j - q + \kappa_j \geq 0$ forces $\kappa_j \geq (q - \delta_j)^+$, whence

$$k = \sum_j \kappa_j \geq \sum_j (q - \delta_j)^+ = (k - \ell) + T(\delta)$$

(using $\sum_j (q - \delta_j) = \mu q - R = k - \ell$), i.e. $T(\delta) \leq \ell$. Profiles violating the budget contribute zero; those satisfying it contribute at most their ambient dimension $\prod_j (\delta_j + 1)$. Summing gives (i).

(ii) We exhibit, for each profile with $T(\delta) \leq \ell - \beta(\delta)$, a feasible choice of chain indices reaching the profile's full ambient dimension. Write $o_j = q - \delta_j$ and recall (Lemma 4.1) that along the profile the pair- j chain has value $v_j(\nu) = \min((\nu + 1)(\nu - o_j + 1), \delta_j + 1)$. Take

$$\nu_j = \min(o_j^+ + \lceil \sqrt{\delta_j + 1} \rceil, q).$$

Then $v_j(\nu_j) = \delta_j + 1$: if the min is the first argument, then $(\nu_j + 1)(\nu_j - o_j + 1) \geq (\lceil \sqrt{\delta_j + 1} \rceil + 1)\lceil \sqrt{\delta_j + 1} \rceil \geq \delta_j + 1$ (both factors are at least $\lceil \sqrt{\delta_j + 1} \rceil$, using $o_j^+ \geq o_j$); if the min is q , then $(q + 1)(q - o_j + 1) \geq \delta_j + 1$ since $q - o_j = \delta_j$. Feasibility: using $\sum_j o_j^+ = (k - \ell) + T(\delta)$,

$$\sum_j \nu_j \leq \sum_j o_j^+ + \sum_j \lceil \sqrt{\delta_j + 1} \rceil \leq (k - \ell) + T(\delta) + \beta(\delta) \leq (k - \ell) + \ell = k,$$

so the tuple ν lies in the index set of Theorem 4.3's inner sum, and the inner sum for this profile is at least $v(\nu) = \prod_j (\delta_j + 1)$ (the filtration sum contains the single tensor product $\bigotimes_j V_{\nu_j}^{(j)}$, which already has full profile dimension). Finally $\beta(\delta) \leq \sum_j (\sqrt{\delta_j + 1} + 2) \leq \sqrt{\mu(R + \mu)} + 2\mu \leq \bar{\beta}$ by Cauchy–Schwarz.

(iii) $T(\delta) \leq (R - q)^+$ deterministically: if no coordinate exceeds q then $T = 0$; otherwise $T(\delta) \leq \sum_j \delta_j - q = R - q$. So the condition in (ii) holds for *every* profile, every profile is fully reached, and $\Gamma_{k,\ell}(W) = \sum_\delta \prod_j (\delta_j + 1) = N_R$. $\square$

Proposition 6.2 (Explicit tail bound). *For any $a \geq 0$ and $R \geq 1$, with $s_0 = q + a/\mu$:*

$$\mathbb{P}[T(\delta) > a] \leq \mu \cdot 2m \left(\frac{(R - s_0 + m)^+}{R} \right)^{m-2} \leq m^2 \left(\frac{(R - q - a/\mu + m)^+}{R} \right)^{m-2}.$$

Proof. If $T > a$ then some coordinate has $(\delta_j - q)^+ > a/\mu$; union-bound over the μ exchangeable coordinates. The coordinate law is $\mathbb{P}[\delta_1 = s] = (s + 1) \binom{R-s+m-3}{m-3} / \binom{R+m-1}{m-1}$, so

$$\mathbb{P}[\delta_1 \geq s_0] \leq \frac{(R+1) \binom{R-s_0+m-2}{m-2}}{\binom{R+m-1}{m-1}} \leq (m-1) \frac{R+1}{R} \left(\frac{R-s_0+m}{R} \right)^{m-2} \leq 2m \left(\frac{R-s_0+m}{R} \right)^{m-2},$$

using the hockey-stick identity for the numerator, $\binom{R+m-1}{m-1} \geq R^{m-1}/(m-1)!$, and $\binom{A}{m-2} \leq A^{m-2}/(m-2)!$. $\square$

Corollary 6.3 (Deep-ambient saturation). $\Gamma_{k,\ell}(W) \geq N_R \left(1 - m^2 \left(\frac{(R-q-(\ell-\bar{\beta})/\mu+m)^+}{R}\right)^{m-2}\right)$ whenever $\ell \geq \bar{\beta}$. In particular $\Gamma_{k,\ell}(W)/N_R \rightarrow 1$ along any parameter family with m fixed in which $(\ell - \bar{\beta})/\mu - (R - q) \rightarrow \infty$ or, more generally, in which the right-hand side tends to 1; and by Theorem 6.1(iii) equality $\Gamma_{k,\ell}(W) = N_R$ holds outright once $\ell \geq R - q + \bar{\beta}$.

Example 6.4. At $m = 4$, $D = 12$, $k = \ell = 6$ ($\mu = 2$, $q = 6$, $R = 12$): formula (5) gives 443 of the ambient 455 (97.4%), and the deficit of 12 sits exactly on the boundary-shell profiles, as located by Theorem 6.1.

Remark 6.5 (The sandwich is about block witnesses, not about generic products). Theorem 6.1(i) is a genuine *ceiling* for disjoint-pair witnesses: a block product can never reach ambient profiles whose overshoot exceeds the shift budget ℓ . It does not bound generic products, whose measure (by semicontinuity) is at least every block value and — in machine experiments — consistently higher in exactly the deficient ranges (cf. Remark 5.4). Determining the exact ambient-regime measure of the fully generic product beyond the sandwich, i.e. inside the boundary shell and beyond the overshoot ceiling, remains open; the shell is $O(\sqrt{mR})$ wide in the budget variable, so the gap is a vanishing fraction of the ambient dimension whenever the budget event has any margin (Proposition 6.2).

6.2 Row regime: block witnesses are genuinely deficient

In the row-dominated regime the exact calculus turns into a *negative* structural statement: no product with disjoint-pair block structure comes close to the row cap once $m \geq 3$ (more precisely, once $\mu \geq 3$; already at $m = 4$ the deficiency is visible numerically — 212 against a row cap of 400 at $D = 12$, $k = \ell = 3$ — though the crude bound below is vacuous there).

Theorem 6.6 (Row-side deficiency of block products). *Let W be the block witness, $m = 2\mu$. For $k = \ell \geq 1$:*

$$\frac{N_k N_\ell}{\Gamma_{k,\ell}(W)} \geq \left(\prod_{i=\mu+1}^{2\mu-1} \frac{k+i}{i} \right)^2 \left(\frac{\mu}{2k+\mu} \right)^\mu \geq \left(\frac{k+\mu}{2\mu} \right)^{2\mu-2} \left(\frac{\mu}{2k+\mu} \right)^\mu.$$

In particular, for $m \geq 6$ and $k = \ell \geq m^2$,

$$\frac{N_k N_\ell}{\Gamma_{k,\ell}(W)} \geq \left(\frac{k}{32m} \right)^{(m-4)/2},$$

a deficiency factor of shape $(\min(k, \ell)/m)^{\Theta(m)}$. Conversely, in the row-dominated regime with $k \leq D/2$ and $D \geq 2m^2$ the deficiency never exceeds $2e^2 m (e/2)^{m-1} (\min(k, \ell) + 1)^{m-1}$: the window between the block value and the row cap is bounded by $(\min(k, \ell))^{O(m)}$.

Proof. Upper bound on $\Gamma_{k,k}(W)$. In (5), reindex each pair (o, ν) by $(\kappa, \lambda) := (\nu, \nu - o)$: the constraints become $\kappa_j, \lambda_j \geq 0$, $\sum \kappa_j \leq k$, $\sum \lambda_j = \sum \kappa_j - (k - \ell) \leq \ell$, and each increment satisfies $\text{inc}_j(\nu_j) \leq 2\nu_j - o_j + 1 = \kappa_j + \lambda_j + 1$. Hence, for $k = \ell$,

$$\Gamma_{k,k}(W) \leq \sum_{\substack{\kappa, \lambda \geq 0 \\ \sum \kappa_j \leq k, \sum \lambda_j \leq k}} \prod_{j=1}^{\mu} (\kappa_j + \lambda_j + 1) \leq \binom{k+\mu}{\mu}^2 \left(\frac{2k+\mu}{\mu} \right)^\mu,$$

bounding the number of index tuples by $\binom{k+\mu}{\mu}^2$ and each summand by AM–GM ($\prod(\kappa_j + \lambda_j + 1) \leq ((2k + \mu)/\mu)^\mu$).

Comparison. $N_k = \binom{k+2\mu-1}{2\mu-1}$, and

$$\frac{\binom{k+2\mu-1}{2\mu-1}}{\binom{k+\mu}{\mu}} = \prod_{i=\mu+1}^{2\mu-1} \frac{k+i}{i} \geq \left(\frac{k+\mu}{2\mu}\right)^{\mu-1},$$

giving the displayed bound. (The weakened form $((k+\mu)/2\mu)^{2\mu-2}$ is too lossy at $\mu = 3$; the “in particular” bound is derived from the product form.) For $\mu \geq 4$, write

$$\rho_\mu := \prod_{i=\mu+1}^{2\mu-1} \frac{2\mu}{i} = 2^{\mu-1} / \prod_{j=1}^{\mu-1} \left(1 + \frac{j}{\mu}\right) \geq \left(\frac{2}{\sqrt{e}}\right)^{\mu-1},$$

using $1+x \leq e^x$ and $\sum_{j<\mu} j/\mu = (\mu-1)/2$; thus $\prod_{i=\mu+1}^{2\mu-1} \frac{k+i}{i} \geq \left(\frac{k}{2\mu}\right)^{\mu-1} \rho_\mu$. For $k \geq m^2 = 4\mu^2$ we also have $(1 + \mu/2k)^\mu \leq e^{\mu^2/2k} \leq e^{1/8}$, so $\left(\frac{\mu}{2k+\mu}\right)^\mu \geq \left(\frac{\mu}{2k}\right)^\mu e^{-1/8}$. Combining,

$$\frac{N_k N_\ell}{\Gamma_{k,k}(W)} \geq \left(\frac{k}{2\mu}\right)^{2\mu-2} \left(\frac{4}{e}\right)^{\mu-1} \left(\frac{\mu}{2k}\right)^\mu e^{-1/8} = \left(\frac{k}{\mu}\right)^{\mu-2} (2e)^{-\mu} e^{7/8} \geq \left(\frac{k}{64\mu}\right)^{\mu-2},$$

the last step because $64^{\mu-2} \geq (2e)^\mu e^{-7/8}$ for $\mu \geq 4$ ($4096 > 365$ at $\mu = 4$; each increment of μ multiplies the left side by 64 and the right by $2e$), and $64\mu = 32m$. For $\mu = 3$ the product form gives directly $\left(\frac{(k+4)(k+5)}{20}\right)^2 \left(\frac{3}{2k+3}\right)^3 \geq \frac{k}{192} = \left(\frac{k}{32m}\right)^{\mu-2}$ for every $k \geq 1$: clearing denominators, this is $5184((k+4)(k+5))^2 \geq 400k(2k+3)^3$, which holds coefficient by coefficient after expansion ($5184((k+4)(k+5))^2 - 400k(2k+3)^3 = 1984k^4 + 78912k^3 + 605664k^2 + 1855440k + 2073600$).

The converse window bound. $\Gamma_{k,\ell}(W) \geq N_\ell$ (multiplication by any single nonzero derivative embeds S_ℓ), and $\Gamma_{k,\ell}(W) \geq \dim \partial^k W \geq \frac{(2/e)^{m-1}}{2e^{2m}} N_k$ for $k \leq D/2$ (by (6) and Lemma 5.2(iii) with the block caps — the content of the proof of Theorem 5.1(c)). Hence $N_k N_\ell / \Gamma_{k,\ell}(W) \leq \min(N_k, 2e^2 m (e/2)^{m-1} N_\ell) \leq 2e^2 m (e/2)^{m-1} (\min(k, \ell) + 1)^{m-1}$, using $N_e \leq (e+1)^{m-1}$. $\square$

Remark 6.7 (Why this matters). Theorem 6.6 shows the row-dominated regime cannot be closed by disjoint-pair witnesses: the deficiency is a structural property of the block decomposition (the multidegree grading of (4) is the obstruction), not slack in the analysis — the calculus is exact. Together with Remark 3.2 (the $m = 2$ fat-point mechanism cannot produce Chow witnesses at $m \geq 3$), this eliminates the two natural witness families and forces the overlapping-chain witnesses of the next sections, whose shared variables destroy the multidegree grading. This is the honest reason the row-dominated regime was the hard one: for generic *subspaces* in place of $\partial^k f$, the analogous maximal-growth statement at general (k, ℓ) is Fröberg-conjecture territory, open for $m \geq 4$ [13, 4, 18]; the chain-witness proof below sidesteps this entirely, and in particular does not imply (nor use) any new case of Fröberg’s conjecture — products of linear forms are very far from generic forms.

7 Chain witnesses and the master reduction

We now prove the main saturation theorem. The witnesses are *chains* of binary forms in overlapping pairs of variables:

$$f = g_1(x_1, x_2) g_2(x_2, x_3) \cdots g_{m-1}(x_{m-1}, x_m), \quad \deg g_i = h_i.$$

Each g_i splits into linear forms over $\overline{\mathbb{F}}$, so f is a product of $D = \sum_i h_i$ linear forms — a Chow point (Section 2.3).

Theorem 7.1 (Chain saturation, uniform in m). *Let $\text{char } \mathbb{F} = 0$, $m \geq 3$, and fix (k, ℓ) . There are thresholds $H_1, \dots, H_{m-1}$ with*

$$H_1 = (k + \ell + 1)(2k + \ell + 1) - 1, \quad H_j \leq 3(k + \ell + 1)^{\lceil (m-2)/2 \rceil} (k + 1) + O(k + \ell) \quad (2 \leq j \leq m-1),$$

such that for all $h_j \geq H_j$ ($2 \leq j \leq m-1$), all $h_1 \geq H_1$ outside a finite exceptional set (which may depend on k, ℓ and $h_2, \dots, h_{m-1}$), and generic binary forms $g_1, \dots, g_{m-1}$ of those degrees:

$$\Gamma_{k, \ell}(g_1 g_2 \cdots g_{m-1}) = N_k N_\ell,$$

the full row cap. Consequently there is a finite $D_0(m, k, \ell) \leq \sum_j H_j + E$, with E the size of the exceptional set for h_1 at $h_j = H_j$ ($j \geq 2$), such that a generic product of $D \geq D_0$ linear forms in m variables satisfies

$$\Gamma_{k, \ell}(f) = \min(N_k N_\ell, N_{D-k+\ell}).$$

The thresholds are exponential in m but polynomial in (k, ℓ) for every fixed m ; the finite exceptional set for h_1 (a minor nonvanishing in a scaling limit, Lemma 8.1 step (3)) is the one non-effective ingredient, and is empty in every machine-checked configuration — so $D_0 = \sum_j H_j$ in every checked case. The final “consequently” is proved by decomposing $D = \sum_j h_j$ with $h_j = H_j$ for $j \geq 2$ and the excess in h_1 : for $D \geq \sum_j H_j + E$ the admissible interval for h_1 has more than E points, so some choice avoids the exceptional set. Then semicontinuity on the Chow variety applies, plus the observation that $N_{D-k+\ell} = \Theta_m(D^{m-1})$ dominates the constant row cap for $D \geq O_m((k+m)(\ell+m))$, which is in turn dominated by D_0 .

The proof occupies this section through Section 10, and is one derivation for all m , organized as: an exact *master reduction* (this section); an *end peel* whose statement quantifies over an abstract spectator family (Section 8); a *middle peel* with explicit coordinatization (Section 9, for $m \geq 4$); and a *residual core lemma* valid over any field (Section 10), where the case $m = 3$ — for which the middle peel is empty — is completed by a one-variable core lemma and the general assembly is carried out.

7.1 Conventions for the chain

Rows of the measure matrix are indexed by pairs $(a, b) \in \mathbb{N}^m \times \mathbb{N}^m$ with $|a| = \ell$, $|b| = k$:

$$E(a, b) = x^a \partial^b f = x_1^{a_1} \cdots x_m^{a_m} \partial_{x_1}^{b_1} \cdots \partial_{x_m}^{b_m} f,$$

so $\Gamma_{k, \ell}(f) = \text{rank}_{\mathbb{F}}\{E(a, b)\}$, and the target $N_k N_\ell$ is exactly the number of rows: saturation means the rows are *linearly independent*.

We write $\phi_1(w) := g_1(1, w)$ and $\phi_{m-1}(w) := g_{m-1}(w, 1)$ for the dehomogenizations at the two ends, and $\sigma_1 := a_1 - b_1 \in [-k, \ell]$ for the (first) *class* parameter. Middle factors are expanded as

$$g_2 = \sum_d c_d^{(0)} x_2^{h_2-d} x_3^d, \quad g_{j+2} = \sum_e c_e^{(j)} x_{j+2}^{h_{j+2}-e} x_{j+3}^e \quad (1 \leq j \leq m-4), \quad \phi_{m-1} = \sum_\mu e_\mu w^\mu.$$

7.2 The master reduction

The first move is an *exact* linear identification that eliminates the end variable x_1 in favor of a Laurent “tag” variable t recording the filtration depth in g_1 ’s coefficients. Its key feature: the rest of the polynomial is never inspected.

Lemma 7.2 (Master reduction at one end; arbitrary co-factor). *Let $g(x, y) = \sum_{\nu} \lambda_{\nu} x^{h-\nu} y^{\nu}$ be a binary form of degree h and let G be any polynomial in y and any further variables ζ , with x -free coefficients. Put $f = g \cdot G$ and $\phi(w) := g(1, w)$. For every row $x^{a_x} y^{a_2} \zeta^A \partial_x^{b_x} \partial_y^{b_2} \partial_{\zeta}^B f$, the monomial substitution $T : x^{\rho} \mapsto t^{h+\ell-\rho}$ (a bijective relabeling on the x -degree range $[0, h + \ell]$, which contains every row) satisfies*

$$T\left(x^{a_x} y^{a_2} \zeta^A \partial_x^{b_x} \partial_y^{b_2} \partial_{\zeta}^B f\right) = t^{\ell} \cdot t^{b_x - a_x} y^{a_2} \zeta^A (h - t\partial_t)^{\underline{b_x}} \partial_y^{b_2} \partial_{\zeta}^B [\phi(ty) \cdot G]. \quad (7)$$

Hence the two families (rows on the left, right-hand sides on the right) have equal rank — exactly, for every g , with no genericity assumption and no threshold on h .

Here $(h - t\partial_t)^{\underline{b_x}}$ denotes the operator $(h - t\partial_t)(h - t\partial_t - 1) \cdots (h - t\partial_t - b_x + 1)$, which acts on a monomial of t -degree ν as the scalar $(h - \nu)^{\underline{b_x}}$.

Proof. Differentiating the expansion of g ,

$$\partial_x^{b_x} \partial_y^j g = \sum_{\nu} \lambda_{\nu} (h - \nu)^{\underline{b_x}} \nu^j x^{h-\nu-b_x} y^{\nu-j},$$

and the Leibniz expansion of $\partial_y^{b_2}$ across $g \cdot G$ collapses against the substituted variable:

$$\sum_j \binom{b_2}{j} t^j \phi^{(j)}(ty) \partial_y^{b_2-j} G = \partial_y^{b_2} [\phi(ty) G].$$

In the t -form, the t -degree of the ν -slice is exactly ν : it tags *which coefficient of g contributed*, i.e. the x -filtration depth. (This is why the tag works where a y -grading would not: ∂_y lowers the y -degree but not the t -degree.) The weight $(h - \nu)^{\underline{b_x}}$ is therefore realized by the operator $(h - t\partial_t)^{\underline{b_x}}$. The x -exponent of the ν -slice is $\rho = a_x + h - \nu - b_x$, and $(h - \nu)^{\underline{b_x}} = 0$ whenever $\nu > h - b_x$, so $\rho \geq 0$ on the support: T is a bijective relabeling of an ambient monomial space containing every row, i.e. the restriction of a linear isomorphism — exactly rank-preserving. The co-factor G was used only through x -freeness. $\square$

Proposition 7.3 (Double master reduction). *For the chain $f = g_1 \cdots g_{m-1}$, write $y := x_2$, $z_j := x_{j+2}$ ($1 \leq j \leq m-3$), let t tag the x_1 -depth and u the x_m -depth. Then*

$$\Gamma_{k,\ell}(f) = \text{rank}_{\mathbb{F}} \{ R_{a,b} : |a| = \ell, |b| = k \},$$

where

$$R_{a,b} = t^{b_1 - a_1} u^{b_m - a_m} y^{a_2} \prod_{j=1}^{m-3} z_j^{a_{j+2}} (h_1 - t\partial_t)^{\underline{b_1}} (h_{m-1} - u\partial_u)^{\underline{b_m}} \partial_y^{b_2} \prod_{j=1}^{m-3} \partial_{z_j}^{b_{j+2}} \left[\phi_1(ty) g_2(y, z_1) \cdots g_{m-2}(z_{m-4}, z_{m-3}) \phi_{m-1}(u z_{m-3}) \right] \quad (8)$$

(Laurent monomials in t, u are permitted; at $m = 3$ the list $z_1, \dots, z_{m-3}$ is empty and $\phi_{m-1} = \phi_2$ is applied to uy — recorded for completeness, though the $m = 3$ assembly of Section 10.2 uses only the single-end reduction).

Proof. Apply Lemma 7.2 at $x = x_1$ with co-factor $G = g_2 \cdots g_{m-1}$, then again at $x = x_m$ — with the orientation reversed, i.e. with $g(x, y) := g_{m-1}(y, x)$, whence $\phi(w) = g_{m-1}(w, 1) = \phi_{m-1}$ and the tag obeys $x_m^\rho \mapsto u^{h_{m-1}+\ell-\rho}$ — with co-factor $\phi_1(ty) g_2 \cdots g_{m-2}$. The two relabelings act on disjoint exponent axes; t enters only through $\phi_1(ty)$ and u only through $\phi_{m-1}(uz_{m-3})$; all middle derivatives preserve both slice depths, and the middle variables are inert under both reductions. Dropping the global offsets t^ℓ, u^ℓ (injective on Laurent polynomials) gives (8). $\square$

The problem is now completely explicit: prove that the $N_k N_\ell$ Laurent polynomials (8) are linearly independent for a suitable (then generic) choice of the g_i .

8 The end peel

The second move eliminates the tag t and the variable y , at the price of specializing ϕ_1 to a truncated multi-exponential. The lemma is stated with everything beyond (t, y) as an abstract *spectator*: an operator family and a co-factor about which nothing is assumed beyond freeness and commutation. This is what makes the statement usable at every m with no induction.

Fix (k, ℓ) and set $C := k + \ell + 1$ throughout.

Lemma 8.1 (End peel). *Let $\text{char } \mathbb{F} = 0$. Let ζ be any tuple of variables, $G(y, \zeta)$ any polynomial, and for each row index (a, b) let $S_{a,b}$ be an $\mathbb{F}$ -linear operator that is a product of: multiplications by Laurent monomials in ζ , powers of ∂_ζ , and Euler-type weights $(h' - \zeta_i \partial_{\zeta_i})^\varepsilon$ — in particular t -free, y -free, ω -free and h_1 -free, and commuting with ∂_y and with multiplication by y and t . Let ϕ_1 be the degree- h_1 truncation of $\sum_{c=1}^C \lambda_c e^{\omega_c w}$ with distinct generic ω_c and $\lambda_c \neq 0$, and consider the family*

$$R_{a,b} = t^{b_1-a_1} y^{a_2} (h_1 - t\partial_t)^{\underline{b_1}} \partial_y^{b_2} S_{a,b}[\phi_1(ty) \cdot G].$$

For a class $\sigma_1 = a_1 - b_1$, call the tuple of S -indices the block blk of a row (in the application: $(b_3, \dots, b_m, a_3, \dots, a_m)$), set $d_t := b_2 - a_2 = (k - \ell + \sigma_1) - \delta_{\text{blk}}$ where δ_{blk} is block-constant (in the application: $\sum_{j \geq 3} (b_j - a_j)$), and define the residue family, one element per block:

$$Y'_{\text{blk}} := y^{-d_t} \cdot S_{\text{blk}}[G],$$

where S_{blk} denotes $S_{a,b}$ for any row of the block (block-constant by hypothesis). If for every $\sigma_1 \in [-k, \ell]$ the family $\{Y'_{\text{blk}}\}_{\text{blk}}$ is linearly independent, then for all $h_1 \geq H_1 := C(2k + \ell + 1) - 1$ outside a finite exceptional set of h_1 -values (which may depend on the remaining data), the full family $\{R_{a,b}\}$ is linearly independent.

Proof. Suppose $\sum_{a,b} c_{a,b} R_{a,b} = 0$; we show every $c_{a,b} = 0$ via four steps.

(1) *Truncation congruence and the zero bound.* $R_{a,b}$ is linear in ϕ_1 , and for the truncated exponential one has the exact congruence

$$R_{a,b}[\text{trunc } e^{\omega \cdot}] \equiv e^{\omega ty} \cdot w_{a,b}^{(\omega)} \pmod{t^{h_1+1-\ell}}, \quad w_{a,b}^{(\omega)} = t^{b_1-a_1} y^{a_2} (\omega t + \partial_y)^{b_2} S_{a,b}[P_{b_1}^{h_1}(\omega ty) \cdot G],$$

where $P_{b_1}^{h_1}(x) = \sum_j \binom{b_1}{j} (h_1 - j)^{\underline{b_1-j}} (-x)^j$ is the Euler weight conjugated by the unit e^{ty} , applied before the ∂_y -conjugation — the operator order that makes the bracket exact. (Derivation: the generating identity $\sum_\nu (h_1 - \nu)^{\underline{b_1}} x^\nu / \nu! = e^x P_{b_1}^{h_1}(x)$ plus an index shift resums the per-slice weights; the truncation tails $\nu > h_1$ carry $t^{\nu+b_1-a_1}$, of order $\geq h_1 + 1 - \ell$.) The t -window of $w_{a,b}^{(\omega)}$ is contained in $[-\ell, 2k]$: t appears only in $t^{b_1-a_1}$, in $P_{b_1}^{h_1}(\omega ty)$ (degree $\leq b_1$), and in $(\omega t + \partial_y)^{b_2}$

(degree $\leq b_2$); S contributes no t ; the top degree is $2b_1 + b_2 - a_1 \leq 2k$ and the bottom is $b_1 - a_1 \geq -\ell$. (The window $[-\ell, 2k]$, rather than the naive $[-\ell, k]$, is attained: already $k = 1$, $\ell = 0$, $b = (1, 0, \dots)$ gives $w = t(h_1 - \omega ty)G$ with a t^2 term. This is what sets H_1 .)

Now over the function field $K = \mathbb{F}(y, \zeta)$ the C frequencies $\omega_c y$ are distinct (only $y \neq 0$ is used — the spectators are inert), the polynomial parts have t -degree $\leq 2k + \ell$ after clearing t^ℓ , and the putative relation $\sum_c e^{\omega_c ty} Q_c \equiv 0 \pmod{t^{h_1+1-\ell}}$, with $Q_c := \lambda_c \sum_{a,b} c_{a,b} w_{a,b}^{(\omega_c)}$, vanishes to order $\geq h_1 + 1$ at $t = 0$ after clearing. The classical zero-multiplicity bound for exponential polynomials — a nonzero $\sum_{c=1}^C e^{\eta_c t} q_c(t)$ with distinct η_c and $\deg q_c \leq q$ has a zero of order at most $C(q+1) - 1$ at $t = 0$; formally valid over any characteristic-0 field via the constant-coefficient ODE $\prod_c (\partial_t - \eta_c)^{q+1}$ with nonzero extreme coefficient — forces each $Q_c = 0$ exactly, once $h_1 + 1 > C(2k + \ell + 1) - 1$, i.e. $h_1 \geq H_1$.

(2) *ω -Vandermonde.* Since S is ω -free, $w_{a,b}^{(\omega)} = \omega^{a_1-b_1} \cdot w_{a,b}^{(1)}|_{t \mapsto \omega t}$, and $t \mapsto \omega t$ is a monomial bijection. Grading the coefficient vector by the class $\sigma_1 = a_1 - b_1$, the condition $Q_c = 0$ for the $C = k + \ell + 1$ distinct values ω_c is, per t -degree, a Vandermonde system in ω^{σ_1} over the $k + \ell + 1$ classes $\sigma_1 \in [-k, \ell]$; its invertibility places every class-graded piece $c^{[\sigma_1]}$ in the kernel of the $\omega = 1$ family.

(3) *Scaling limit, per class.* Fix a class. Clear the class-constant prefactor $t^{-\sigma_1}$ first, then substitute $t \mapsto h_1 t$ and divide the row of index (a, b) by $h_1^{b_1+b_2}$: every entry becomes a polynomial in $\varepsilon_1 := 1/h_1$ — here h_1 -freeness of S is consumed (S -weights are ε_1 -free constants) — and at $\varepsilon_1 = 0$ the family degenerates to

$$v_{a,b} = t^{b_2} y^{a_2} (1 - ty)^{b_1} \cdot S_{a,b}[G].$$

The prefactor order matters for $\sigma_1 < 0$; with it, the polynomiality is a direct check:

$$(h_1 - j)^{b_1-j} / h_1^{b_1-j} = \prod_i (1 - (j+i)\varepsilon_1), \quad (h_1 t + \partial_y)^{b_2} / h_1^{b_2} = (t + \varepsilon_1 \partial_y)^{b_2}.$$

If the limit family is independent, a maximal minor — a polynomial in ε_1 nonvanishing at 0 — certifies independence for all but finitely many h_1 : this is the proof's only limit and only exceptional set.

(4) *Triangularization in the tag degree.* Within a block, rows are indexed by b_1 alone (given the block and the class: $b_2 = k - b_1 - \sum_{j \geq 3} b_j$, $a_2 = \ell - a_1 - \sum_{j \geq 3} a_j$, $a_1 = b_1 + \sigma_1$), and $d_t = b_2 - a_2$ is block-constant. Expanding $(1 - ty)^{b_1}$ and using $a_2 = b_2 - d_t$,

$$v_{a,b} = \sum_{p=0}^{b_1} (-1)^p \binom{b_1}{p} (ty)^{b_2+p} \cdot Y'_{\text{blk}},$$

after clearing the global y -unit: the t -support of the row with index b_1 is $[b_2, b_2 + b_1]$, whose lowest point $p = b_2 = k - b_1 - \sum_{j \geq 3} b_j$ is distinct for distinct b_1 within the block, with unit coefficient — a unitriangular system for any subset of b_1 's. A vanishing combination of the v 's therefore gives, at each t -degree, a vanishing combination of the residue elements Y'_{blk} across the class's blocks; their assumed joint independence reduces this to one vanishing coefficient per block at each t -degree, and the unitriangular structure in b_1 then forces $c^{[\sigma_1]} = 0$ for every class, hence $c = 0$. $\square$

Remark 8.2 (What the lemma consumed). Characteristic 0 (the zero bound and the Vandermondes); the listed freeness and commutation properties of the spectator operators S ; and *nothing* about the co-factor G . In the application at general m the spectator monomials include

$u^{b_m - a_m}$ with possibly negative exponent — Laurent monomials — and no step of the proof uses nonnegativity of the exponents: the statement holds verbatim on the Laurent ambient space, where the rows live anyway.

Remark 8.3 (Reduction achieved). Applying Lemma 8.1 to the family (8) with

$$\begin{aligned}\zeta &= (z_1, \dots, z_{m-3}, u), & G &= g_2 \cdots g_{m-2} \phi_{m-1}(uz_{m-3}), \\ S_{a,b} &= u^{b_m - a_m} \prod_j z_j^{a_j + 2} \cdot (h_{m-1} - u\partial_u)^{b_m} \prod_j \partial_{z_j}^{b_j + 2},\end{aligned}$$

Theorem 7.1 reduces to: *for every class σ_1 , the residue family*

$$\begin{aligned}Y'_{\text{blk}} &= y^{-d_t} u^{b_m - a_m} \prod_{j=1}^{m-3} z_j^{a_j + 2} (h_{m-1} - u\partial_u)^{b_m} \prod_{j=1}^{m-3} \partial_{z_j}^{b_j + 2} \\ &\quad \left[g_2(y, z_1) \cdots g_{m-2}(z_{m-4}, z_{m-3}) \phi_{m-1}(uz_{m-3}) \right],\end{aligned}\tag{9}$$

*indexed by blocks $(b_3, \dots, b_m, a_3, \dots, a_m)$ with $d_t = (k - \ell + \sigma_1) - \sum_{j \geq 3} (b_j - a_j)$, is linearly independent, for one (hence a generic) choice of the factors $g_2, \dots, g_{m-2}, \phi_{m-1}$ of large enough degrees. At $m = 3$ we instead apply the peel to the *single-end* reduction (Lemma 7.2 at x_1 only): $\zeta = (z)$ with $z := x_3$, $G = g_2(y, z)$, $S_{a,b} = z^{a_3} \partial_z^{b_3}$, blocks (b_3, a_3) , and residue family $Y'_{\text{blk}} = y^{-d_t} z^{a_3} \partial_z^{b_3} g_2(y, z)$; the middle peel of the next section is not needed, and this base case is completed directly in Section 10.2. For $m \geq 4$ the next section eliminates the middle factors.*

9 The middle peel: coordinatization, windows, and the twist Vandermonde

Throughout this section $m \geq 4$. Fix a class σ_1 and recall the residue family (9), indexed by blocks $\text{blk} = (b_3, \dots, b_m, a_3, \dots, a_m)$. We must prove its linear independence for suitable middle factors. We specialize *all* $m - 2$ remaining coefficient sequences to *multi-geometric* form with $W := k + \ell + 1$ terms each and generic frequencies:

$$c_e^{(j)} = \sum_{i=1}^W \rho_i^{(j)} \nu_{j,i}^e \quad (j = 0, \dots, m-4), \quad e_\mu = \sum_{x=1}^W \rho_x^{(\phi)} \omega_x^\mu,\tag{10}$$

with all ρ 's nonzero and, per family, distinct nonzero frequencies. (Independence of the residue family is a Zariski-open condition on the coefficient tuples; the multi-geometric witnesses below certify nonemptiness, and Section 10.3 performs the final intersection of open conditions.) When indices from a single generic frequency tuple suffice we drop the second subscript and write $\nu_j := \nu_{j,i}$, $\omega := \omega_x$ for a typical term; all displayed identities are per-term.

9.1 Slice structure

Fix a block. Expanding the bracket of (9) over coefficient tuples

$$(e_0, \dots, e_{m-3}) := (d, e_1, \dots, e_{m-4}, \mu)$$

of the $m - 2$ sequences, the pre-derivative exponent of the middle variable z_j is

$$E_j = e_{j-1} + h_{j+2} - e_j \quad (1 \leq j \leq m-4), \quad E_{m-3} = e_{m-4} + \mu,$$

and after the derivatives and monomial multiplications the tuple contributes to the *single* column monomial with exponents

$$w := h_2 - d_t - d \text{ (of } y), \quad Z_j := E_j - b_{j+2} + a_{j+2} \text{ (of } z_j), \quad q := \mu + b_m - a_m \text{ (of } u), \quad (11)$$

with coefficient (exact, per monomial)

$$c_d^{(0)} \left(\prod_{j=1}^{m-4} c_{e_j}^{(j)} \right) e_\mu \cdot \prod_{j=1}^{m-3} E_j^{b_{j+2}} \cdot (h_{m-1} - \mu)^{b_m} \quad (12)$$

— $m-2$ falling-factorial *loads*, indexed by $b_3, \dots, b_m$. The map $(d, \vec{e}, \mu) \mapsto (w, \vec{Z}, q)$ is injective per block (index recovery below), so every monomial of Y'_{blk} receives exactly one tuple contribution.

9.2 Coordinatization: the telescoped global relation

Set

$$\delta_i := \sum_{j=i+2}^m (b_j - a_j) \quad (i = 1, \dots, m-2), \quad \tau := a_m - b_m (= -\delta_{m-2}), \quad A := h_2 - (k - \ell + \sigma_1).$$

Telescoping $\sum_j E_j = d + \mu + \sum_{j=3}^{m-2} h_j$ against (11) gives the *global relation* — block-independent:

$$\sum_{j=1}^{m-3} Z_j = \sum_{j=2}^{m-2} h_j - (k - \ell + \sigma_1) - w + q. \quad (13)$$

Fix any split point $r \in [1, m-3]$ (we take $r = \lceil (m-2)/2 \rceil$; the window exponents below are $(r, m-2-r)$). Columns are then faithfully coordinatized by the $m-2$ values $(w, q, Z_i \text{ (} i \neq r))$ — Z_r is recovered from (13) — and the sequence indices recover in closed prefix *and* suffix forms:

$$\text{prefix: } e_j = C_j + \delta_{j+1} - w - \sum_{i \leq j} Z_i, \quad C_j := A + \sum_{i=3}^{j+2} h_i \quad (0 \leq j \leq m-4); \quad (14)$$

$$\text{suffix: } e_j = \sum_{i > j} Z_i - q + \delta_{j+1} - D_j, \quad D_j := \sum_{i=j+3}^{m-2} h_i \quad (0 \leq j \leq m-4), \quad (15)$$

in particular $d = e_0 = A + \delta_1 - w$ and $\mu = q + \tau$. (The two forms are proved by upward resp. downward induction on j from (11); their mutual consistency is the global relation, in the form $C_j + D_j = \sum_{i=2}^{m-2} h_i - (k - \ell + \sigma_1)$.)

The load arguments become univariate in the kept coordinates:

$$E_j = Z_j + (\delta_j - \delta_{j+1}) \text{ for every } j, \quad h_{m-1} - \mu = (h_{m-1} - \tau) - q,$$

and expressing Z_r via (13) makes the b_{r+2} -load *affine in all coordinates*; it is the only load that carries w .

9.3 Tensor factorization, split recovery, windows

Substituting (10) and the prefix forms (14) for the left families $j \leq r-1$, the suffix forms (15) for the right families $j \geq r$:

$$\nu_j^{e_j} = [\nu_j^{C_j}] \cdot \nu_j^{\delta_{j+1}} \cdot \nu_j^{-w} \prod_{i \leq j} \nu_j^{-Z_i} \text{ (left), } \nu_j^{e_j} = [\nu_j^{-D_j}] \cdot \nu_j^{\delta_{j+1}} \cdot \nu_j^{-q} \prod_{i > j} \nu_j^{Z_i} \text{ (right), } \omega^\mu = \omega^\tau \cdot \omega^q,$$

so the per-coordinate exponential bases are *pure tensors* over the coordinates:

coordinate w :	base $(\nu_0 \nu_1 \cdots \nu_{r-1})^{-1}$ — prefix products, at most W^r values;
coordinate Z_i ($i \leq r-1$):	base $(\nu_i \cdots \nu_{r-1})^{-1}$, at most W^{r-i} values;
coordinate q :	base $\omega/(\nu_r \cdots \nu_{m-4})$ — suffix ratios, at most W^{m-2-r} values;
coordinate Z_i ($i \geq r+1$):	base $\nu_r \cdots \nu_{i-1}$, at most W^{i-r} values.

There is no cross-leakage: each family's exponential touches exactly the displayed coordinates, and the polynomial loads have degrees $\deg_w \leq b_{r+2}$, $\deg_q \leq b_{r+2} + b_m$, $\deg_{Z_i} \leq b_{i+2} + b_{r+2}$ — all at most k . Frequency recovery peels *from the middles outward*: the Z_{r-1} -base is ν_{r-1} solo, then the Z_{r-2} -base divided by ν_{r-1} yields ν_{r-2} , and so on until the w -base yields ν_0 ; symmetrically the Z_{r+1} -base is ν_r solo, $\dots$, and the q -base finally yields ω . Hence per-family distinctness of the frequencies alone makes the map (base vector) $\mapsto$ (frequency tuple) injective: pairwise distinctness of the prefix products or suffix ratios is *not* needed (collisions merely shorten the per-coordinate base lists, weakening the window requirements below). Edge cases: at $r=1$ there are no left Z 's and the w -base ν_0 is solo; at $r=m-3$ there are no right Z 's and the q -base still yields ω last; at $m=4$ there are no middle coordinates at all and only the w - and q -bases remain.

Windows. All constraints below quantify over the full twist ranges $\delta_i \in [-\ell, k]$, $\tau \in [-k, \ell]$, and the class list $\sigma_1 \in [-k, \ell]$ (the latter absorbed into the $O(k+\ell)$ slack via A); window *positions* for the interior coordinates are per-class (the constants C_j contain A), while the *thresholds* H_j are class-uniform — the entire argument runs within one class, so per-class positioning is all that is needed. Requiring each sequence index to range over a full product window of consecutive integers:

- w : $d = A + \delta_1 - w \in [0, h_2]$ leaves an interval of length $h_2 - k - \ell + 1 \geq W^r(k+1)$ iff $h_2 \geq H_2 := W^r(k+1) + k + \ell - 1$;
- q : $\mu = q + \tau \in [0, h_{m-1}]$ leaves length $h_{m-1} - k - \ell + 1 \geq W^{m-2-r}(k+1)$ iff $h_{m-1} \geq H_{m-1} := W^{m-2-r}(k+1) + k + \ell - 1$;
- left interior $1 \leq i \leq r-1$: $e_i \in [0, h_{i+2}]$ over the whole w -window and all upstream Z -windows leaves an available Z_i -interval of length $h_{i+2} + 1 - (\text{len}_w - 1) - \sum_{j < i} (\text{len}_{Z_j} - 1) - (k + \ell) \geq W^{r-i}(k+1)$; since $\sum_{j \leq i} W^{r-j} \leq 2W^r$, it suffices that $h_{i+2} \geq W^{r-i}(k+1) + 2W^r(k+1) + k + \ell$;
- right interior $r \leq j \leq m-4$: symmetrically against the q -window and downstream Z -windows, $h_{j+2} \geq W^{j+1-r}(k+1) + 2W^{m-2-r}(k+1) + k + \ell$ (the case $j=r$ is the eliminated coordinate's sequence: it contributes a threshold for h_{r+2} but no new window; the windows are positioned nested from the ends inward, and are nonempty exactly when the thresholds hold).

Hence $H_j \leq 3W^{\lceil (m-2)/2 \rceil} (k+1) + O(k+\ell)$ for all $2 \leq j \leq m-1$ at the split $r = \lceil (m-2)/2 \rceil$ — exponential in m , polynomial in (k, ℓ) for fixed m , which is all Theorem 7.1 claims. (Other splits give other exact constants; $r = 1$ reproduces the thresholds of the $m = 5$ instance in Appendix A, $r = 2$ those of $m = 6$.)

On the product window, every family element is a pure tensor from the Kronecker product of per-coordinate univariate families $\{\beta^x x^s\text{-type loads} : \beta \text{ a base value, } \deg \leq k\}$; each per-coordinate family is independent on a window of $\geq (\# \text{bases})(k+1)$ consecutive integers (a confluent Vandermonde at consecutive integer nodes: the order- $(\# \text{bases})(k+1)$ linear recurrence $\prod_{\beta} (\Sigma - \beta)^{k+1}$ with invertible extreme coefficients annihilates exactly this family), and distinct (frequency tuple, degree vector) pairs give distinct tensor elements. Hence, per frequency tuple, the vanishing of a combination of the residue family holds as a *polynomial identity in the $m-2$ formal coordinates* $(w, q, Z_i (i \neq r))$.

9.4 The twist Vandermonde

The per-tuple identity still carries the *twist factor* $\prod_{j=0}^{m-4} \nu_j^{\delta_{j+1}} \cdot \omega^\tau$, together with nonzero units $\rho^{(0)} \cdots \rho^{(\phi)} \cdot \nu_j^{C_j}$ (left) $\cdot \nu_j^{-D_j}$ (right). The twist exponents range over $\delta_i \in [-\ell, k]$ ($i = 1, \dots, m-3$) and $\tau \in [-k, \ell]$, each of size at most $k + \ell + 1 = W$; vanishing on the $(m-2)$ -dimensional grid of W distinct nonzero values per axis (tensor Vandermonde; Laurent shifts are absorbed by the nonzero nodes) forces every fine $(\delta_1, \dots, \delta_{m-3}, \tau)$ -group of coefficients to vanish separately. Within a fine group, blocks are faithfully indexed by $(b_3, \dots, b_m)$ alone: $a_m = b_m + \tau$ and $a_j = b_j - \delta_{j-2} + \delta_{j-1}$ ($j = 3, \dots, m-1$).

The residual obligation, per fine group, is exactly the hypothesis of Lemma 10.1 below, with all $m-2$ coordinates formal, at every m — no integer window parameters and no collapsing of formal coordinates is needed. The next section proves it and assembles the theorem.

10 The residual core and the assembly

10.1 The core lemma, uniform in m

Lemma 10.1 (Residual core). *Let $\mathbb{F}$ be any field, $p \geq 0$, and let $w, q, X_1, \dots, X_p$ be variables. Let $c_1, \dots, c_p, K, H \in \mathbb{F}$ be arbitrary constants and $B \subseteq \mathbb{N}^{p+2}$ any finite set of tuples $\beta = (\beta_1, \dots, \beta_p, \beta_w, \beta_q)$. If*

$$\sum_{\beta \in B} \gamma_\beta \prod_{i=1}^p (X_i + c_i)^{\beta_i} \cdot \left(K - w + q - \sum_i X_i \right)^{\beta_w} \cdot (H - q)^{\beta_q} \equiv 0 \quad \text{in } \mathbb{F}[w, q, X_1, \dots, X_p],$$

then all $\gamma_\beta = 0$.

In the application: $p = m-4$, the X_i are the kept coordinates Z_i ($i \neq r$), the β -slots are the block indices $(b_{i+2})_{i \neq r}$, $\beta_w = b_{r+2}$, $\beta_q = b_m$, $c_i = \delta_i - \delta_{i+1}$, $K = \sum_{j=2}^{m-2} h_j - (k - \ell + \sigma_1) + \delta_r - \delta_{r+1}$, $H = h_{m-1} - \tau$ — but none of that structure is needed.

Proof. Only the β_w -load involves w ; as a polynomial in w it has degree exactly β_w with leading coefficient $(-1)^{\beta_w}$ — a nonzero *constant*, independent of q and $\vec{X}$. Extract the top w -power: with $B_w := \max_{\beta \in B} \beta_w$, the coefficient of w^{B_w} in the identity reads

$$\sum_{\beta: \beta_w = B_w} \gamma_\beta (-1)^{B_w} \prod_i (X_i + c_i)^{\beta_i} \cdot (H - q)^{\beta_q} \equiv 0 \quad \text{in } \mathbb{F}[q, \vec{X}].$$

The $p + 1$ remaining factors are univariate polynomials in *separate* variables with degrees $(\beta_1, \dots, \beta_p, \beta_q)$ and unit leading coefficients $(1, \dots, 1, (-1)^{\beta_q})$, so the products over distinct tuples $(\beta_1, \dots, \beta_p, \beta_q)$ form a subset of a graded tensor basis of $\mathbb{F}[\vec{X}, q]$ — linearly independent — hence all γ_β with $\beta_w = B_w$ vanish. Remove them and induct downward on β_w ; the elimination is contamination-free because the removed blocks are gone entirely. $\square$

Note what the proof did *not* use: no characteristic hypothesis (every leading coefficient is ± 1), no structure of B , no distinctness or nonvanishing of the constants. Characteristic 0 is consumed only upstream (Lemma 8.1 and the Vandermondes of Section 9).

10.2 The base case $m = 3$

At $m = 3$ there are no middle factors, and the residue family (after the single-end reduction; Remark 8.3) is

$$Y'_{(b_3, a_3)} = y^{-d_t} z^{a_3} \partial_z^{b_3} g_2(y, z), \quad d_t = (k - \ell + \sigma_1) - (b_3 - a_3),$$

one element per block (b_3, a_3) , within a class σ_1 . Writing $g_2 = \sum_d c_d y^{h_2-d} z^d$,

$$Y'_{(b_3, a_3)} = \sum_d c_d d^{b_3} y^{h_2-d-d_t} z^{d-b_3+a_3},$$

and since $(h_2 - d - d_t) + (d - b_3 + a_3) = h_2 - (k - \ell + \sigma_1)$ is class-constant, all elements of the class live on a single anti-diagonal, coordinatized by the z -exponent $e := d - b_3 + a_3$. In coordinate e the block (b_3, a_3) contributes the vector $((e + \delta)^{b_3} c_{e+\delta})_e$ with $\delta := b_3 - a_3$, and distinct blocks give distinct pairs (δ, b_3) . Independence is thus an instance of the following one-variable lemma.

Lemma 10.2 (Shifted sequences with falling-factorial loads). *Let $\text{char } \mathbb{F} = 0$. Let Ω be a finite set of distinct pairs (δ, n) with $|\delta| \leq \Delta$ and $0 \leq n \leq M$. Let $c_e = \sum_{x=1}^{W_0} \lambda_x \mu_x^e$ be a multi-geometric sequence with distinct nonzero μ_x , all $\lambda_x \neq 0$, and $W_0 \geq 2\Delta + 1$. Then the vectors $((e + \delta)^n c_{e+\delta})_{e \in \mathcal{W}}$, $(\delta, n) \in \Omega$, over any integer window $\mathcal{W}$ of length at least $W_0(M + 1)$, are linearly independent.*

Proof. The falling-factorial Vandermonde identity $(e + \delta)^n = \sum_i \binom{n}{i} \delta^{n-i} e^i$ expands each vector over the basis $u_{x,i} = (e^i \mu_x^e)_e$; the $u_{x,i}$ ($1 \leq x \leq W_0$, $0 \leq i \leq M$) are independent on a window of length $W_0(M + 1)$ (confluent Vandermonde at consecutive integer nodes, i.e. the recurrence $\prod_x (\Sigma - \mu_x)^{M+1}$ with invertible extreme coefficients). Vanishing of each $u_{x,i}$ -coefficient says that the polynomial $G_i(\mu) = \sum_\delta \mu^{\delta+\Delta} \sum_n \alpha_{\delta,n} \binom{n}{i} \delta^{n-i}$ (degree $\leq 2\Delta$) vanishes at the $W_0 \geq 2\Delta + 1$ points μ_x : so $G_i \equiv 0$ for every i . For each fixed δ , the matrix $[\binom{n}{i} \delta^{n-i}]_{n,i}$ is lower-triangular with unit diagonal, so all $\alpha_{\delta,n} = 0$. $\square$

Applying Lemma 10.2 with $\Delta = k + \ell$, $M = k$, $W_0 = 2\Delta + 1$: for a multi-geometric g_2 (which certifies a generic g_2 ; the window fits inside $[0, h_2]$ once $h_2 \geq W_0(M + 1) + 2\Delta - 1 = \text{poly}(k, \ell)$), every class's residue family is independent. Together with Lemmas 7.2 and 8.1 this proves Theorem 7.1 for $m = 3$, via the assembly below. (Boundary detail: coordinates e near the ends of $[0, h_2]$ where $d^{b_3} = 0$ vanish coherently with the formula, since $d^{b_3} = 0$ for $0 \leq d < b_3$; the window is placed away from the ends.)

10.3 Assembly

Proof of Theorem 7.1. Fix $m \geq 4$ and (k, ℓ) (for $m = 3$ replace the middle peel by Lemma 10.2 as just described). Chain the lemmas:

(1) *Exact reduction.* Applied at both ends of the chain, Lemma 7.2 (Proposition 7.3) identifies $\Gamma_{k,\ell}(g_1 \cdots g_{m-1})$ with the rank of the family (8), exactly, for every choice of factors.

(2) *Residue independence for witness middles.* Specialize the middle factors and ϕ_{m-1} multi-geometrically (10) with degrees meeting the window thresholds of Section 9.3. For such witnesses, Sections 9.1–9.4 prove: every σ_1 -class residue family (9) is linearly independent. Indeed, a vanishing combination, expanded over the slice structure, restricts on each product window to a Kronecker product of independent per-coordinate families (confluent Vandermonde), giving a polynomial identity per frequency tuple; the twist Vandermonde separates the fine $(\vec{\delta}, \tau)$ -groups; and within a fine group the identity is precisely the hypothesis of Lemma 10.1, which kills all coefficients. The finitely many conditions consumed — per-family distinct nonzero frequencies, nonzero ρ 's, window fit — are Zariski-open and nonempty.

(3) *End peel.* With every class's residue family independent, Lemma 8.1 (with the spectator data of Remark 8.3) gives independence of the full family $\{R_{a,b}\}$ for the truncated multi-exponential ϕ_1 with $h_1 \geq H_1 = (k + \ell + 1)(2k + \ell + 1) - 1$ outside a finite exceptional set. Hence $\Gamma_{k,\ell} = N_k N_\ell$ for the witness chain.

(4) *One tuple satisfies everything; genericity.* Each condition above is Zariski-open in the irreducible parameter space of coefficient tuples of $(g_1, \dots, g_{m-1})$ at the given degrees, and each is nonempty (the witnesses certify this), so their finite intersection is nonempty and dense: independence holds for one, hence for generic, chains of those degrees — and $\Gamma_{k,\ell}$ is lower-semicontinuous, so the generic chain of any degrees $h_j \geq H_j$ achieves the row cap.

(5) *Transfer to generic products of linear forms.* The witness factors have generic leading coefficients (one more open condition), so each g_i homogenizes to a binary form of degree h_i that splits into linear forms over $\overline{\mathbb{F}}$; the chain is then a product of $D = \sum h_i$ linear forms — a point of the irreducible Chow variety — and semicontinuity on it (Section 2.3) gives $\Gamma_{k,\ell}(f) \geq N_k N_\ell$ for the generic product of $D \geq D_0 = \sum_j H_j$ linear forms. Finally $N_{D-k+\ell} = \binom{D-k+\ell+m-1}{m-1} = \Theta_m(D^{m-1}) \geq N_k N_\ell$ for $D \geq O_m((k+m)(\ell+m))$, which D_0 dominates; so the row cap is the min, and the universal cap (2) closes the equality. $\square$

Remark 10.3 (Quantifier order). The order matters and is: fix the class-window positions and prove residue independence for the fixed multi-geometric middle tuple (per class); then Lemma 8.1 for the multi-exponential end; then intersect the finitely many open conditions. The per-class window *positions* depend on σ_1 ; only the *thresholds* are uniform, which is all that is quantified in the theorem.

Remark 10.4 (Two forced-order facts). Two natural simplifications provably fail, and the architecture is shaped by their failure. (i) Treating the two chain ends by a *symmetric* double scaling limit $h_1, h_{m-1} \rightarrow \infty$ loses rank: in the double limit the middle factor decouples and Pascal relations $x^b(1-x)^c = x^{b+1}(1-x)^c + x^b(1-x)^{c+1}$ create kernel vectors; the measured limit rank is deficient at every interior class already at $m = 4$. Hence the sequential order — peel one end by Lemma 8.1, keep the other end at finite degree inside the spectator family — is forced. (ii) In the middle peel, evaluating the formal coordinates at a generic point before the final elimination loses rank (measured deficient at every interior (δ, τ) -group at $m = 5$): the residual core must be run with all coordinates *formal*, which is exactly what Lemma 10.1 provides. Appendix A records the measurements.

Remark 10.5 (Effectivity). All thresholds are explicit except the finite exceptional set for h_1 (the nonvanishing at $\varepsilon_1 = 0$ of one maximal minor, Lemma 8.1 step (3)). In every machine-checked configuration the exceptional set is empty — at $m = 3$, saturation was measured down to h_1 as small as k — so the large- h_1 restriction appears to be proof slack, not truth. Making step (3) effective would make the entire theorem effective.

11 Related work

Shifted partials and depth-four lower bounds. The measure $\Gamma_{k,\ell}$ was introduced by Kayal [19] and used by Gupta–Kamath–Kayal–Saptharishi [16] to prove $n^{\Omega(\sqrt{d})}$ lower bounds for homogeneous depth-four circuits with bounded bottom fan-in; subsequent refinements pushed variants of the measure to $n^{\Omega(\sqrt{d})}$ for general homogeneous depth four (Kayal–Limaye–Saha–Srinivasan [21], Kumar–Saraf [22]) and, with bounded bottom fan-in, for iterated matrix multiplication (Fournier–Limaye–Malod–Srinivasan [12]); the bounds have remained at $n^{\Omega(\sqrt{d})}$ since 2014. The upper-bound step in all these arguments over-approximates the gate measure by the universal caps; our results show that for product-of-linear-forms gates in few variables the over-approximation is, beyond polynomial thresholds, no approximation at all.

Barriers for rank methods. Efremenko–Landsberg–Schenck–Weyman [11] proved that shifted partials cannot separate the permanent from the determinant beyond small polynomial factors; Efremenko–Garg–Oliveira–Wigderson [10] formalized barriers for general rank methods; Bhargav–Dutta–Saxena [6] extended the barrier analysis to more measures and models. These results concern the many-variable regime and prove that random/generic gates *nearly* saturate. The present paper proves *exact* saturation, with explicit constants and thresholds, for the most structured gate family, in the complementary few-variable high-degree regime — and localizes precisely which witness structures fail to saturate (Theorem 6.6) and why.

Evaluation variants over small fields. Armand–Behera–Tavenas [5] prove $2^{\Omega(\sqrt{d})}$ lower bounds for depth-five circuits with bounded top product fan-in over small finite fields, using *evaluations* of shifted partials on a well-chosen point set. The contrast with our results is instructive, and confirms that our characteristic-0 hypothesis is substantive rather than technical. Over $\mathbb{F}_q$ with $q = O(1)$, a product of many affine forms of high rank vanishes to order $> k$ everywhere off an exponentially small set — each factor vanishes at a random point with probability about $1/q$ — so the order- k shifted partials of such a gate, *evaluated* on the surviving points, are all zero: the gate’s evaluated measure is near zero, the opposite polarity to saturation. In characteristic 0 (or large characteristic) this multiplicity mechanism has no purchase, and our theorems show the classical (coefficient-space) measure is fully saturated. The two regimes are also disjoint quantitatively: in [5] the variable count grows polynomially in the degree, while in the few-variable regime over $\mathbb{F}_q$ every evaluation space has constant dimension $\leq q^m$. Together the results locate the exact boundary of the saturation phenomenon: small characteristic *plus* the evaluation twist restores measure-confinement against product gates; either ingredient alone does not.

Partial derivatives of products. Shalmon–Shpilka [24] prove an $\exp(\Omega(r^{1/6}))$ lower bound on the (unshifted) partial-derivative dimension of a product of r linearly independent *quadratics*

in many variables — a largeness result for nonlinear factors, motivated by the classical fact that a product of n independent linear forms has derivative dimension $\binom{n}{k}$ at order k . Our $\ell = 0$ results (Theorem 5.1) are complementary: few variables, linear factors, exact and near-exact values including the shifted ranges. Edwards [9] develops a coefficient-matrix formalism (“SPDP rank and codimension”) for shifted-partial measures; the framework is orthogonal to our concern (no dimension computations for concrete gate families are made there).

Apolarity of products of linear forms. DiPasquale–Flores–Peterson [8] study the apolar algebra of a product of linear forms — Waring rank bounds under genericity hypotheses via star configurations (Geramita–Harbourne–Migliore), and computational classification of certain sextic multi-arrangements — and connect it to the geometry of hyperplane arrangements. The generic Hilbert-function questions answered here (for $\partial^k f$ and its shifts) do not appear there or, to our knowledge, elsewhere; [8] is the closest prior work and our point of departure. On the commutative-algebra side of our proofs, we emphasize again the *negative* relationship to maximal-rank conjectures: Fröberg’s conjecture [13] (settled for $m \leq 2$ by Fröberg and for $m = 3$ by Anick [4], for $\ell = 1$ — i.e. in the first nontrivial degree — by Hochster–Laksov [18], and in various special ranges, recently including the second nontrivial degree by Boij–Dannetun–Lundqvist [7]; see also Fröberg–Löfwall [14]) concerns *generic* forms; derivative spaces of products of linear forms are highly non-generic subspaces, and our saturation theorems neither use nor imply any case of it. Similarly, the fat-point interpolation problems of Alexander–Hirschowitz type [3] are used only at $m = 2$, where they are trivial (principal ideals in $\mathbb{P}^1$).

Provenance of the questions. The block-gate $D^{O(m)}$ -window phenomenon quantified in Theorem 6.6, and the question of whether it reflects the truth about generic products, arose in an analysis of what shifted-partial-type measures can certify against constant-variate product gates in the bootstrapping regime [1, 23, 15]; the answer proved here — nothing, beyond polynomial thresholds — was conjectured from exact machine computations of the ranks at small parameters, which also guided every stage of the proofs.

A Worked instances: $m \leq 6$

The uniform derivation of Sections 7–10 was found by closing the cases $m = 3, 4, 5, 6$ in succession; each case contributed one structural ingredient, and each was proved with smaller (sometimes different) thresholds than the uniform statement instantiates. We record the instance statements, their exact thresholds, how each proof route relates to the uniform one, and the negative findings that forced the architecture. Full derivations of the instances are subsumed by the uniform proof (with the split parameter r chosen per instance); we indicate only what differs.

A.1 $m = 3$

Theorem A.1. $\text{char } \mathbb{F} = 0$. For every (k, ℓ) there are $H_1 = (k + \ell + 1)(2k + \ell + 1) - 1$ and $H_2 = (2k + 2\ell + 1)(k + 1) + 2(k + \ell) - 1$ such that for all $h_1 \geq H_1$ outside a finite exceptional set, $h_2 \geq H_2$, and generic binary $g_1(x, y)$, $g_2(y, z)$: $\Gamma_{k, \ell}(g_1 g_2) = N_k N_\ell$. Hence there is a finite $D_0(k, \ell) \leq H_1 + H_2 + E$ (E the size of the exceptional set at $h_2 = H_2$; $D_0 = H_1 + H_2$ when it is empty, as in every machine-checked configuration) such that a generic product of $D \geq D_0$ linear forms in 3 variables saturates $\min(N_k N_\ell, N_{D-k-\ell})$.

Route: single-end master reduction; Lemma 8.1; residue family resolved by Lemma 10.2 (Section 10.2). This is the anchor case of the general theorem and the origin of the method: the master reduction was found here as an exact replacement for an ε -degeneration bookkeeping, and the multi-exponential assembly of Lemma 8.1 was found as the fix for the failure of the naive global limit. Two measured degeneracy facts from this case carry structural information: exponential or geometric coefficient sequences for g_1 are *deficient* witnesses for the full family (measured 82/100 resp. 91/100 at $(3, 3)$ against 100/100 generic) — genericity of g_1 is load-bearing for the cross-class assembly only, since within a class the g_1 coefficients are removable scalings; and truncating ϕ_1 in degree shrinks the engaged tag-range, which is exactly what the assembly consumes.

A.2 $m = 4$

Theorem A.2. $\text{char } \mathbb{F} = 0$. For every (k, ℓ) there are $H_1, H_2, H_3 = \text{poly}(k, \ell)$ (H_1 as above) such that for all $h_1 \geq H_1$ (outside a finite exceptional set), $h_2 \geq H_2$, $h_3 \geq H_3$, and generic binary $g_1(x_1, x_2)$, $g_2(x_2, x_3)$, $g_3(x_3, x_4)$: $\Gamma_{k, \ell}(g_1 g_2 g_3) = N_k N_\ell = \binom{k+3}{3} \binom{\ell+3}{3}$; hence generic products of $D \geq \text{poly}(k, \ell)$ linear forms in 4 variables saturate.

New at this case: the double master reduction (Proposition 7.3) and the *sequential* assembly. The originally proved route kept the u -end at finite degree via a bespoke independence lemma (“Lemma Y”: multi-geometric witnesses for both g_2 and ϕ_3 , a two-variable confluent Vandermonde, and a twist Vandermonde onto (δ, τ) -groups with graded-monic loads); in the uniform framework this is the case $p = 0$ of Lemma 10.1 (no kept middle coordinates: only the w - and q -bases remain).

Forced order (measured). The symmetric double scaling limit $h_1, h_3 \rightarrow \infty$ provably loses rank: in the double limit g_2 factors out and the Pascal relation $x^{b_3}(1-x)^{b_4} = x^{b_3+1}(1-x)^{b_4} + x^{b_3}(1-x)^{b_4+1}$ produces kernel vectors; the limit family is deficient at every interior class (e.g. 30/50 at $(3, 3)$, central class). The $m = 3$ theorem therefore cannot be cited as a black box at the last level; the finite-degree treatment of the second end is forced. This is Remark 10.4(i).

A.3 $m = 5$

Theorem A.3. $\text{char } \mathbb{F} = 0$. For every (k, ℓ) there are $H_1 = (k + \ell + 1)(2k + \ell + 1) - 1$, $H_2 = (k + \ell + 1)(k + 1) + k + \ell - 1$, $H_3 = H_4 = (k + \ell + 1)^2(k + 1) + O(k + \ell)$ such that for all $h_1 \geq H_1$ (outside a finite exceptional set), $h_j \geq H_j$, and generic binary $g_1, \dots, g_4$ along the chain: $\Gamma_{k, \ell}(g_1 g_2 g_3 g_4) = \binom{k+4}{4} \binom{\ell+4}{4}$; hence generic products of $D \geq \text{poly}(k, \ell)$ linear forms in 5 variables saturate.

These thresholds are the uniform ones at split $r = 1$ (which is why H_2 is smaller here than at the balanced split). Two facts were discovered at this case. First, *all* middle sequences must be specialized multi-geometrically: with g_3 generic the q -side confluent Vandermonde is blocked (the g_3 -index depends on q), and the g_3/ϕ_4 frequencies *merge* in the private-end exponent as ratios ω_x/ν_y — the suffix-ratio bases of Section 9.3 in embryo. Second, the originally proved route collapsed one formal coordinate to a per-column integer (“ S -collapse”) and finished with a three-elimination core lemma; the $m = 6$ case later revealed the all-formal route, which retro-simplifies $m = 5$ (it is Lemma 10.1 at $p = 1$) — both routes are valid, and the instance thresholds agree. The measured counterpart of Remark 10.4(ii) is from this case: evaluating the two formal parameters of the residue at a point is deficient at every interior group, while the formal family is full-rank everywhere.

A.4 $m = 6$

Theorem A.4. $\text{char } \mathbb{F} = 0$. For every (k, ℓ) there are $H_1 = (k + \ell + 1)(2k + \ell + 1) - 1$, $H_2 = H_5 = (k + \ell + 1)^2(k + 1) + k + \ell - 1$, $H_3 = H_4 = (k + \ell + 1)^2(k + 1) + (k + \ell + 1)(k + 1) + O(k + \ell)$ such that for all $h_1 \geq H_1$ (outside a finite exceptional set), $h_j \geq H_j$, and generic binary $g_1, \dots, g_5$ along the chain: $\Gamma_{k, \ell}(g_1 \cdots g_5) = \binom{k+5}{5} \binom{\ell+5}{5}$; hence generic products of $D \geq \text{poly}(k, \ell)$ linear forms in 6 variables saturate.

These are the uniform thresholds at split $r = 2$ (the first balanced split; the “coupled” interior windows appear here). New at this case, and decisive for the uniform derivation: the telescoped global relation (13) was recognized (the middle-exponent sum is global, leaving $m - 2$ formal coordinates); the per-coordinate bases first exhibit a *product* merge ($\mu_i \nu_j$ in the w -exponent) alongside the ratio merge; and the residual core *simplified* relative to $m = 5$ — all coordinates formal, no integer windows — revealing the shape (Lemma 10.1) that generalizes. A verification by-product worth recording for anyone reproducing the computations: evaluating *both* tags (t, u) numerically before taking ranks is not rank-safe (it collapses pairs of rows related by the two end gradings; the deficiency equals the number of b -monomials at $\ell = 1$) — evaluating one tag only is the safe one-sided certificate.

A.5 Slack observed at the instances

Across all four instances and the uniform case, machine checks pass *below* the proof thresholds: saturation holds at h_1 as small as k (the exceptional set of Lemma 8.1(3) is empty in every tested configuration), with $W = 2$ frequency counts (the proof uses $k + \ell + 1$), and at degrees well under the window thresholds. None of the thresholds above should be read as sharp; they are what the uniform proof yields.

B Machine verification record

Every result in this paper was developed alongside exact machine verification, and every derivation was subsequently re-verified adversarially: each joint re-derived independently — in a separate session, from the statements alone, without access to the recorded derivation (see the disclosure in Section 1.5) — and diffed against the record, and an *independent* numerical battery implemented from the re-derivation, at a different prime and seed than the original. All ranks and identities are checked in exact arithmetic (integers, rationals, or $\mathbb{F}_p$ at primes ranging from 101 to $2^{31} - 1$); floating point enters only two peripheral comparisons (a total-positivity spot-check via rounded determinants, and the numerical evaluation of the transcendental closed form $(2/e)^{m-1}/(2e^2 m)$ when compared against the exact rational constant). The scripts are available at <https://github.com/bhudgeons/shifted-partials-saturation> and as ancillary files accompanying the arXiv version of this paper.

B.1 Protocol

For each derivation-level result the pass consists of: (i) independent re-derivation of every step, written before consulting the recorded derivation, then diffed; (ii) an independent implementation of every checkable identity, from the re-derivation, run at a second prime/seed; (iii) re-run of the original battery, with every reported figure reproduced exactly. Findings are classified as: quantitative repairs (a stated constant or window corrected), slack findings (a stated hypothesis

not actually consumed by the proof), or remark-level (phrasing). Across the five chain-theorem passes ($m = 3, 4, 5, 6$ and uniform) there was exactly one quantitative repair — the t -window of the end peel is $[-\ell, 2k]$, not $[-\ell, k]$, raising H_1 to $(k + \ell + 1)(2k + \ell + 1) - 1$, found at $m = 3$ and already incorporated in Lemma 8.1 — and no other; all later passes confirmed their derivations with only slack or remark-level findings (two of which are recorded in Section 9: per-family frequency distinctness suffices, and window positions are per-class).

B.2 What is checked, per result

- **Theorem 3.1 ($m = 2$ min-law).** Measured $\Gamma_{\kappa, \lambda}$ of dense forms, of products, and of the power-sum witness itself, against the min-law: all $q \leq 16$, all $\kappa, \lambda \leq 6$; exact agreement.
- **Theorem 4.3 (block calculus).** Formula vs. measured rank of the block gate: $m = 4$, $D = 12$, values 36/80/128/212/443 at $(2, 1), (2, 2), (3, 2), (3, 3), (6, 6)$ and $\ell = 0$ values 10/20/31/44 at $k = 2, 3, 4, 6$; exact agreement in every configuration, including total-positivity and min-law sub-checks per pair.
- **Theorem 6.1 and Proposition 6.2 (ambient sandwich).** An exact battery over full profile sweeps at $m = 4, 6$: overshoot *necessity* (every profile with $T > \ell$ has zero inner sum, and no profile exceeds its ambient dimension — zero violations); *sufficiency* (every profile with $T \leq \ell - \beta(\delta)$ is fully saturated); the exact-saturation threshold ($\ell \geq (R - q)^+ + \bar{\beta}$ forces $\Gamma_{k, \ell}(W) = N_R$, verified at every swept configuration); the tail bound of Proposition 6.2 dominating the exact tail, in rational arithmetic; and the deficit localization at $m = 4$, $D = 12$, $k = \ell = 6$: 443 of 455, with the deficit of 12 sitting exactly on the profiles $(0, 12), (12, 0)$. Measured mod- p rank agrees with the formula at every configuration, including $(6, 6)$.
- **Theorem 5.1 and Lemma 5.2.** The $G = B_{\lceil q/2 \rceil} B_{\lfloor q/2 \rfloor}$ identity ($q \leq 16$); the recorded midpoint ratios 44/84 ($m = 4$) and 141/462 ($m = 6$) reproduced exactly; the full block-cap sweep $m \leq 10$, $q \leq 14$: symmetry, log-concavity, monotonicity of h , boundary equality, midpoint bound, and the final bound (iii) at every k , in exact integer/rational arithmetic; 2000 random cap multisets for the cap-structure-free monotonicity; the Pólya-urn pushforward identity verified exactly in rational arithmetic at five (m, k) configurations; the exact rational constant verified to dominate the closed form $(2/e)^{m-1}/(2e^2m)$ for all block caps with $D \geq 2m^2$, $m \leq 10$, $q \leq 59$.
- **Chain theorem, instances $m = 3, 4, 5, 6$.** Per case: master reduction exactness (including off-cap configurations, e.g. $217 = 217$ at $m = 5$, $434 = 434 < 441$ at $m = 6$); truncation congruence and t -window attainment (exactly $[-\ell, 2k]$); scaling-limit and triangularization ranks per class; residue-family ranks per class (generic and multi-geometric witnesses, down to $W = 2$); the per-case core lemmas at every group of the full grids (e.g. 55/147/309 fine groups at $m = 5$, sizes $(2, 2)/(3, 3)/(4, 4)$; 131/471/1251 at $m = 6$); end-to-end saturation at the proof-witness configuration and thresholds (e.g. 36–441 rows at $m = 3$; 100/100, 200/200, 400/400 at $m = 4$; 225/225, 525/525 at $m = 5$; 126/126 at $m = 6$); and the two negative controls of Remark 10.4 (double-limit deficiency, e.g. 30/50; evaluated-parameter deficiency at every interior group).
- **Uniform derivation (Theorem 7.1).** The code path is uniform in m — the same functions build every object from $(m, k, \ell, \vec{h})$ with no per- m branches; $m = 4, 5, 6$ run as

regression against the closed instances, $m = 7, 8$ (and $m = 9$ in the adversarial battery) as the first uncovered cases. Checked as exact per-monomial dictionary identities, not just ranks: the double master reduction on every row ($m = 7$ at $(2, 1)$: all 196 rows; $m = 8$ at $(1, 1)$: all 64); the full coordinatization — global relation, prefix and suffix closed forms, injectivity, load arguments — on every monomial (1,850,492 monomials at $m = 7$ $(2, 1)$; 1,559,160 at $m = 8$ $(2, 1)$; 506,304 at $m = 9$ $(1, 1)$); Lemma 10.1 standalone at $m = 6..8$ for $k = 2, 3$ and $m = 9, 10$ for $k = 2$, every split point r , four constant regimes, in symbolic (polynomial-identity) form, plus $p = 101$ for the any-field claim and random sparse B beyond the application range; planted-collision recovery injectivity (243/729 frequency tuples with cross-family and product/ratio collisions); every fine group of the application grids full-rank (154/505 at $m = 7$, 869 at $m = 8$ $(2, 2)$); residue and end-peel families full-rank at every class; end-to-end saturation 49/49, 64/64, 196/196. The two joints genuinely new at $m \geq 7$ — multi-factor kept-coordinate bases, and nonempty suffix constants D_j — are exercised explicitly. Primes/seeds: 754974721/20260802 (original), 999999937/88 (independent).

B.3 A control

That the adversarial batteries genuinely bite is itself checked: in the uniform pass, the verifier’s first run of the symbolic Lemma 10.1 battery failed at $m = 6$ — traced to an off-by-one in the verifier’s own enumeration of kept coordinate slots, not to the lemma — and passed after the fix. We record this as evidence that the batteries are not vacuously green.

References

- [1] M. Agrawal, S. Ghosh, and N. Saxena. Bootstrapping variables in algebraic circuits. *Proc. Natl. Acad. Sci. USA*, 116(17):8107–8118, 2019.
- [2] M. Agrawal and V. Vinay. Arithmetic circuits: A chasm at depth four. In *Proc. 49th FOCS*, 67–75, 2008.
- [3] J. Alexander and A. Hirschowitz. Polynomial interpolation in several variables. *J. Algebraic Geom.*, 4(2):201–222, 1995.
- [4] D. J. Anick. Thin algebras of embedding dimension three. *J. Algebra*, 100(1):235–259, 1986.
- [5] J. Armand, A. R. Behera, and S. Tavenas. Lower bounds for depth-5 algebraic circuits with bounded fan-in of top product gates. ECCC Report TR26-104, 2026.
- [6] C. S. Bhargav, S. Dutta, and N. Saxena. Improved lower bound, and proof barrier, for constant depth algebraic circuits. In *Proc. 47th MFCS*, LIPIcs vol. 241, 18:1–18:16, 2022. Journal version: *ACM Trans. Comput. Theory*, 16(4):23:1–23:22, 2024.
- [7] M. Boij, E. Dannetun, and S. Lundqvist. Independence of generic forms and the Fröberg conjecture. arXiv:2605.03872, 2026.
- [8] M. DiPasquale, Z. Flores, and C. Peterson. On the apolar algebra of a product of linear forms. In *Proc. ISSAC 2020*, 130–137, ACM, 2020. arXiv:2002.04818.

- [9] D. J. Edwards. Shifted partial derivative polynomial rank and codimension. arXiv:2512.20729, 2025.
- [10] K. Efremenko, A. Garg, R. Oliveira, and A. Wigderson. Barriers for rank methods in arithmetic complexity. In *Proc. 9th ITCS, LIPIcs* vol. 94, 1:1–1:19, 2018.
- [11] K. Efremenko, J. M. Landsberg, H. Schenck, and J. Weyman. The method of shifted partial derivatives cannot separate the permanent from the determinant. *Math. Comp.*, 87(312):2037–2045, 2018. arXiv:1609.02103.
- [12] H. Fournier, N. Limaye, G. Malod, and S. Srinivasan. Lower bounds for depth-4 formulas computing iterated matrix multiplication. *SIAM J. Comput.*, 44(5):1173–1201, 2015.
- [13] R. Fröberg. An inequality for Hilbert series of graded algebras. *Math. Scand.*, 56:117–144, 1985.
- [14] R. Fröberg and C. Löfwall. Generic forms and algebras. arXiv:2504.13591, 2025 (v3 revised 2026).
- [15] Z. Guo, M. Kumar, R. Saptharishi, and N. Solomon. Derandomization from algebraic hardness. *SIAM J. Comput.*, 51(2):315–335, 2022. Preliminary version in FOCS 2019.
- [16] A. Gupta, P. Kamath, N. Kayal, and R. Saptharishi. Approaching the chasm at depth four. *J. ACM*, 61(6):33:1–33:16, 2014.
- [17] S. G. Hoggar. Chromatic polynomials and logarithmic concavity. *J. Combin. Theory Ser. B*, 16(3):248–254, 1974.
- [18] M. Hochster and D. Laksov. The linear syzygies of generic forms. *Comm. Algebra*, 15(1–2):227–239, 1987.
- [19] N. Kayal. An exponential lower bound for the sum of powers of bounded degree polynomials. ECCC Report TR12-081, 2012.
- [20] P. Koiran. Arithmetic circuits: The chasm at depth four gets wider. *Theoret. Comput. Sci.*, 448:56–65, 2012.
- [21] N. Kayal, N. Limaye, C. Saha, and S. Srinivasan. An exponential lower bound for homogeneous depth four arithmetic formulas. *SIAM J. Comput.*, 46(1):307–335, 2017. Preliminary version in FOCS 2014.
- [22] M. Kumar and S. Saraf. On the power of homogeneous depth 4 arithmetic circuits. *SIAM J. Comput.*, 46(1):336–387, 2017. Preliminary version in FOCS 2014.
- [23] M. Kumar, R. Saptharishi, and A. Tengse. Near-optimal bootstrapping of hitting sets for algebraic models. *Theory of Computing*, 19(12):1–30, 2023.
- [24] N. Shalmon and A. Shpilka. Partial derivative complexity of a product of linearly independent quadratics. ECCC Report TR26-065, 2026.
- [25] S. Tavenas. Improved bounds for reduction to depth 4 and depth 3. *Inform. and Comput.*, 240:2–11, 2015.
- [26] L. G. Valiant. Completeness classes in algebra. In *Proc. 11th STOC*, 249–261, 1979.