Information Loss Is All You Need: TC^0 Pseudorandomness from Worst-Case Lattices

Chen Qian

State Key Laboratory of Cryptography and Digital Economy Security,
Shandong University, Qingdao 266237, China

Key Laboratory of Cryptologic Technology and Information Security of Ministry of Education,
Shandong University, Qingdao 266237, China

School of Cyber Science and Technology,
Shandong University, Qingdao 266237, China

Abstract

Pseudorandomness plays a central role in algorithms, complexity theory, and cryptography. We ask whether deterministic information loss alone can yield post-quantum pseudorandomness with TC^0 evaluation. Learning with errors (LWE) hides linear measurements with random noise. Learning with rounding (LWR) replaces noise with deterministic rounding, but a standard polynomial-modulus hardness reduction requires parameters tailored to a predetermined sample bound. Decisional Diffie–Hellman offers a noiseless algebraic route to pseudorandomness, but its standard instantiations are vulnerable to quantum algorithms. We show that deterministic rounding achieves this under quantum worst-case lattice hardness, with fixed polynomial parameters and no predetermined sample bound.

We introduce the *learning with algebraic rounding* (LAR) assumption, which applies an explicit integer rounding map to blocks of secret-dependent linear measurements and encodes the result in a lattice quotient. For one fixed family with polynomially bounded modulus and dimensions, its deterministic labels on exactly uniform, secret-independent public matrices are computationally indistinguishable from independent uniform labels for every polynomial number of classical random examples. Our classical reduction from Gaussian LWE, combined with the quantum worst-case reduction to LWE, bases LAR on GapSVP and SIVP over general lattices at approximation factor $\widetilde{O}(q)$, where q is the modulus.

To prove this, we develop a *source-independent affine transport framework* that may be of independent interest. Using public affine maps sampled independently of the hidden errors, it transports the Gaussian error law close to uniform on the rounding map's zero-label fiber while preserving exactly uniform public matrices. For every fixed inverse-polynomial accuracy, we give a polynomial-time sampler for the same public parameter family.

This yields the first post-quantum weak pseudorandom function with joint key-and-input evaluation in TC^0, secure against uniform quantum polynomial-time distinguishers given any polynomial number of classical uniform random examples. We also establish the hardness of learning polynomial-size two-hidden-layer ReLU networks from noiseless uniform random examples, even for quantum polynomial-time improper learners.

Contents

Main text

1 Introduction 1
2 Preliminaries 13
3 Learning with Algebraic Rounding 18
4 From Uniform-Zero-Fiber LWE to Decisional LAR 21
5 From Gaussian LWE to Uniform-Zero-Fiber LWE 22
6 Zero-Fiber Geometry and Affine Matrix Realization 25
7 Finite-Dimensional Probability Calibration 30
8 Chameleon-Lobster Family with Safe Local Sampling 31
9 Calibrating Public Randomness for Affine Fiber Transport 48
10 Applications 59
References 62

Appendices

A Foundations and the Main Reduction 66
B Incidence Geometry and Affine Transport Proofs 75
C Local Kernels and Bounded Sampling 78
D Shared Baseline and Joint Realization 88
E Fourier Targets and Protected Reference Tools 94
F Positive Proposal and Response Calibration 99
G Complete Fourier Tails and Public-Core Moments 116
H Charge Conditioning and Bounded Positive Sampling 125
I Proofs for the Applications 132

1 Introduction

Pseudorandomness plays a central role in algorithms, complexity theory, and cryptography, providing efficiently computable objects that appear random to computationally bounded observers. We ask the following natural question:

> **Can deterministic information loss alone yield post-quantum pseudorandomness with TC^0 evaluation?**

Pseudorandom functions. Weak pseudorandom functions (weak PRFs) give a concrete formulation of this goal. A weak PRF is a keyed family of deterministic functions. Under one uniformly random key, its evaluations at independently sampled uniform inputs are computationally indistinguishable from evaluations of a truly random function at the same inputs.

Pseudorandom functions were introduced by Goldreich, Goldwasser, and Micali [GGM86]. Their strong notion requires computational indistinguishability even under adaptively chosen oracle queries. The original construction builds PRFs from pseudorandom generators and hence from one-way functions. Naor and Reingold [NR04] gave more efficient algebraic constructions based on factoring and the decisional Diffie–Hellman assumption, with evaluation in TC^0. However, Shor [Sho94] gave quantum polynomial-time algorithms for the underlying factoring and discrete-logarithm problems.

Banerjee, Peikert, and Rosen [BPR12] initiated the study of direct, highly parallelizable lattice-based PRFs by introducing learning with rounding (LWR) as a deterministic counterpart of learning with errors (LWE). Their constructions include low-depth PRFs with TC^0 evaluation under suitable parameters. Subsequent work developed key-homomorphic PRFs from LWE [BLMR13, BP14] and constrained key-homomorphic PRFs from standard lattice assumptions [BV15]. Further work addressed efficiency, reduction tightness, and modulus size [JKP18, Mon18a, Kim20].

A central issue for deterministic rounding is to establish hardness with a polynomial modulus and no a priori bound on the number of samples. Alwen, Krenn, Pietrzak, and Wichs [AKPW13, Theorem 4.1] proved LWE-based hardness of LWR with polynomial modulus, but their reduction requires parameters that depend explicitly on the sample bound. Newton and Richelson [NR23] established limitations for non-aborting pointwise reductions that preserve the source dimension and modulus in a specified parameter regime. These limitations apply to the stated reduction model and do not establish a general impossibility of polynomial-modulus LWR hardness.

Another line of work studies PRFs from learning parity with noise (LPN). Ding, Jain, and Komargodski [DJK25] constructed NC^1 weak PRFs from standard LPN and stronger variants under structured LPN assumptions. The same authors [DJK26] subsequently gave a generic depth-preserving transformation from weak to strong PRFs and obtained NC^1 strong PRFs from standard LWE with a polynomial modulus-to-noise ratio, as well as from standard LPN and computational Diffie–Hellman.

Computational hardness of learning. Weak PRFs with low-complexity evaluation also give hardness of noise-free learning from uniform random examples. An efficient learner that predicts a fresh label with a nonnegligible advantage over random guessing would distinguish pseudorandom labels from independent uniform ones. In Valiant's probably approximately correct (PAC) framework [Val84], a concept class can have polynomial sample complexity and efficiently computable functions, yet admit no polynomial-time learner under suitable computational assumptions. Such hardness is especially significant for *improper learning*, where the learner may output any efficiently

evaluable hypothesis, including one outside the target class. Kearns and Valiant [KV94] established the cryptographic approach to this representation-independent hardness.

The input distribution is central to this separation. Distribution-free hardness rules out a learner that succeeds for every input distribution; it need not imply hardness from uniformly random examples. A reduction establishing the latter must preserve the prescribed input law. Kharitonov [Kha95] established cryptographic hardness under uniform inputs. For distribution-free learning, Klivans and Sherstov [KS09] based the hardness of learning intersections of half-spaces on worst-case lattice problems. Low-depth circuits make the role of computational resources concrete: Linial, Mansour, and Nisan [LMN93] gave a Fourier-based quasipolynomial-time uniform-distribution learning algorithm for constant-depth AND/OR/NOT circuits (AC^0). For constant-depth threshold circuits (TC^0), Arunachalam, Grilo, and Sundaram [AGS21] ruled out polynomial-time quantum learning under the uniform distribution, even with quantum membership queries, assuming quantum hardness of ring learning with errors.

Shallow neural networks can be hard to learn even without label noise. Daniely and Vardi [DV21] obtained hardness for shallow rectified linear unit (ReLU) networks under Gaussian inputs from local pseudorandom generators. Chen, Gollakota, Klivans, and Meka [CGKM22] proved superpolynomial statistical-query lower bounds for noise-free two-hidden-layer ReLU learning under Gaussian inputs. They also established cryptographic hardness results based on learning with rounding. These are distinct types of lower bounds: statistical-query hardness alone does not rule out general efficient learners.

We seek this separation under exactly uniform inputs, with deterministic TC^0 evaluation and a worst-case lattice foundation. We consider uniform quantum polynomial-time learners receiving classical noiseless random examples, with possibly improper hypotheses given by classical descriptions and admitting efficient deterministic classical evaluation as specified in Section 10.3. Our neural-network application uses a uniform finite grid. The same fixed family of polynomial parameters must support every polynomial sample budget for both pseudorandomness and this learning consequence.

1.1 Our Results

Under quantum worst-case hardness of general-lattice problems, we construct an explicit deterministic block-rounding family whose labels on exactly uniform, secret-independent public matrices are indistinguishable from independent uniform labels.

Learning with algebraic rounding (LAR) mixes a block of linear measurements using an integer matrix $\mathbf{T} \in \mathbb{Z}^{d\times d}$, divides by an odd prime q, and rounds down coordinatewise. The label is its coset modulo $\mathbf{T}\mathbb{Z}^d$, invariant under changes of integer representatives. For the rounding matrices of Definition 2.3, Smith coordinates encode these cosets in $\mathbb{F}_q^\ell$, where $\mathbb{F}_q = \mathbb{Z}/q\mathbb{Z}$ and $\ell = d - r \geq 1$ for modular rank r. These coordinates also give translations prescribing label changes (Lemma 3.5), as used below. With the quotient coordinate matrix $\mathbf{P}$ of Definition 3.2, the label map $f_{\mathbf{T}}$ of Definition 3.4, and one uniformly random secret $\mathbf{s} \in \mathbb{F}_q^N$, the examples are

$$\mathbf{M}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right), \qquad i = 1, \ldots, t,$$
$$\mathbf{y}_i = f_{\mathbf{T}}(\mathbf{M}_i\mathbf{s}) = \mathbf{P}\left\lfloor \frac{\mathbf{T}[\mathbf{M}_i\mathbf{s}]_q}{q} \right\rfloor \bmod q \in \mathbb{F}_q^\ell.$$

Here the matrix-secret product is over $\mathbb{F}_q$, $[\cdot]_q$ chooses representatives in $\{0, \ldots, q-1\}$. The subsequent matrix product is over the integers, and the floor acts coordinatewise. Labels use no fresh randomness. The public matrices are mutually independent and secret-independent; ideal labels

are independent and uniform in $\mathbb{F}_q^\ell$. Exact public-input uniformity gives a random-example learning problem on the full matrix space.

A fixed family for every polynomial sample budget. Our hardness result uses one explicit, polynomial-time computable family with $N = \lambda$, an odd prime q, and $d \geq q$, all polynomially bounded in λ, and

$$\mathbf{T} = \left[\mathbf{e}_1 - \mathbf{e}_2 \mid \cdots \mid \mathbf{e}_{d-1} - \mathbf{e}_d \mid \sum_{j=1}^{q} \mathbf{e}_j \right], \qquad \mathbf{P} = \mathbf{1}^\intercal, \qquad \ell = 1,$$

where the $\mathbf{e}_j$ are standard basis vectors and $\mathbf{1}$ is the all-ones vector. Thus each label sums the rounded coordinates modulo q. The modulus, dimensions, and rounding map are fixed before the adversary and its polynomial sample budget are chosen.

Theorem 1.1 (Main reduction, informal)**.** For this fixed public family, any uniform QPT distinguisher that separates real and ideal LAR distributions on polynomially many classical random examples with nonnegligible advantage yields uniform quantum polynomial-time algorithms for worst-case GapSVP_γ and SIVP_γ on general lattices of dimension N, with $\gamma = \widetilde{O}(q)$.

Quantum worst-case hardness therefore gives pseudorandom labels for every polynomial sample budget under the same fixed public family. The reduction passes through Gaussian LWE: the worst-case lattice step is quantum, while the subsequent sample transformations are classical and remain valid when composed with a QPT distinguisher.

From Gaussian noise to deterministic information loss. The zero fiber consists of inputs with label zero. Translation equivariance combines a uniform fiber residual and an independent uniform label into a uniform ambient vector, reducing uniform-zero-fiber LWE to LAR. Our *source-independent affine transport* reduces Gaussian LWE to uniform-zero-fiber LWE by making Gaussian errors nearly uniform on this fiber using public randomness.

Applying this public affine map to both matrix and noisy response preserves exact matrix uniformity and independence from the transported error through invertibility modulo q. The map observes neither hidden errors nor LWE data. Its accuracy is chosen after fixing the adversary and polynomial sample budget, without changing the public family. Transport is used only in the security reduction; labels remain deterministic.

Calibrating public randomness. Local calibration prescribes color probabilities at every protected input through one positive public law of safe maps. Balanced expected counts still leave endpoints correlated through shared source coordinates.

Global calibration reweights public choices to cancel bounded-support Fourier tests; a full-frequency tail bound controls remaining joint error. The solve temporarily omits the integrality congruence and may yield negative weights. Moment bounds justify bounded nonnegative clipping; restoring the congruence yields efficient, source-independent sampling.

Consequences for pseudorandomness and noiseless learning. Under the same worst-case hardness condition, restricting the secret to nonzero keys gives a field-valued weak PRF on $\mathbb{F}_q^{d \times N}$ with joint key-and-input TC^0 evaluation. Security holds against uniform QPT distinguishers receiving any polynomial number of classical uniform random examples. The normalized labels also

have polynomial-size two-hidden-layer ReLU representations on a uniform finite grid, with rational weights and biases of polynomially bounded magnitude and bit length.

The same assumption rules out uniform QPT PAC learning of these networks in squared loss from classical noiseless random examples, including improper learning with classically described hypotheses whose rational predictions admit uniform deterministic polynomial-time classical evaluation. The separation is computational: exhaustive search identifies the secret from $O(N)$ examples at fixed confidence.

Formal verification. The main theorems and lemmas listed in the Prove2Me project[1] have all been verified in Lean 4. These include the fixed-order Gaussian transport of Theorem 9.9 and the reduction from worst-case lattice problems to decisional LAR in Corollary 10.1.

1.2 Technical Overview

Gaussian LWE hides a linear function of the secret using random errors; a LAR label is deterministic. Simply rounding an LWE response can leave its label dependent on the hidden error. We instead transport Gaussian errors to a law that rounding absorbs exactly, while preserving exactly uniform public inputs. The transport uses randomness independent of the samples in the security reduction; label evaluation remains deterministic.

The resulting reduction chain is

$$\{\mathsf{GapSVP}, \mathsf{SIVP}\} \xrightarrow{\text{[Reg05]; Section I.1}} \text{Gaussian } \mathsf{LWE} \xrightarrow{\text{Section 5}} \text{uniform-zero-fiber } \mathsf{LWE}$$
$$\xrightarrow{\text{Section 4}} \mathsf{dLAR} \xrightarrow{\text{Sections 10.2 and 10.3}} \text{weak PRFs and noiseless learning hardness.}$$

The first arrow supplies the quantum worst-case lattice foundation; the last gives consequences for classical random-example access. For the middle arrows, we identify the required error law, construct affine maps that reach its support except on a Gaussian tail event, and choose a positive public law of maps giving joint uniformity up to controlled error.

I. Finding the error law that rounding can absorb. For the label map $f_{\mathbf{T}}$ recalled in Section 1.1, the zero fiber $\mathcal{F}_{\mathbf{T},0}$ consists of inputs with label zero. Its role follows from *translation equivariance*: the matrix $\mathbf{J}$ in Lemma 3.5 converts any desired label change $\mathbf{y} \in \mathbb{F}_q^\ell$ into an input translation. Adding $\mathbf{Jy}$ adds precisely $\mathbf{y}$ to the label. Thus every input has a unique description by a zero-fiber residual and a label. The decomposition of Lemma 4.1 is the bijection

$$(\mathbf{e}, \mathbf{y}) \longmapsto \mathbf{e} + \mathbf{Jy}, \qquad \mathcal{F}_{\mathbf{T},0} \times \mathbb{F}_q^\ell \xrightarrow{\sim} \mathbb{F}_q^d.$$

A uniform residual and an independent uniform label reconstruct an exactly uniform input whose rounded label is the chosen label. Thus fiber membership removes the error from the label, while uniformity supplies the required input law. The residual need not be coordinatewise small: zero-fiber membership concerns integer rounding and quotient encoding, not the modular kernel of $\mathbf{T}$. Coordinatewise error bounds alone do not solve this joint distributional problem.

Theorem 4.2 samples one uniform auxiliary vector $\mathbf{a}_0 \in \mathbb{F}_q^N$ shared across blocks and independent uniform labels, all independently of the input table. On real uniform-zero-fiber LWE input with hidden secret $\mathbf{s}$, its rank-one update makes the matrix-secret product equal the zero-fiber residual plus the chosen label direction when $\mathbf{a}_0^\intercal \mathbf{s} = 1$. With exact auxiliary draws, this gives real LAR

[1] `https://prove2.me/missions/Information_Loss_Is_All_You_Need%3A_TC%5E0_Pseudorandomness_from_Worst-Case_Lattices`

labels on exactly uniform matrices; otherwise the update gives the ideal product law. For each nonzero secret, alignment has probability $1/q$, explaining the advantage loss. The shared vector selects one branch for the whole table. The theorem also accounts for the zero-secret event and bounded sampling. It remains to reach the zero-fiber error law from Gaussian samples.

II. Moving errors that the reduction cannot see. An LWE reduction receives a public matrix $\mathbf{A}$ and a noisy response $\mathbf{b}$, but cannot extract the hidden Gaussian error $\mathbf{e}$. An error-dependent transport would thus be unusable. We instead sample an integer matrix $\mathbf{Q}$ and offset β independently of the secret and the entire sample table. Their affine action can be applied directly to the observed block:

$$(\mathbf{A}, \mathbf{b}) \longmapsto (\mathbf{QA}, \mathbf{Qb} + \beta), \qquad \mathbf{e} \longmapsto \mathbf{Qe} + \beta \pmod{q}.$$

Source independence makes this operation available to the reduction. If every sampled $\mathbf{Q}$ is invertible modulo q, the transformed public matrix is exactly uniform even conditional on the affine map and error. It is therefore independent of the transformed error. An ideal uniform response also remains uniform. It suffices to make the transformed error close to $\mathsf{U}(\mathcal{F}_{\mathbf{T},0})$, averaging over the independently drawn map and source; an individual fixed map need not achieve this closeness. Lemma 5.2 and Theorem 5.4 then give the reduction, with statistical loss at most the number of blocks times the single-block transport error.

Why incidence geometry is useful. We first design integer vertex *loads* $\mathbf{w}$, then recover the affine output through $\mathbf{T}^{-1}$. Their sum is the *charge*, their residues are endpoint *colors*, and their integer quotients on division by q are *carries*. For the matrix in Section 1.1, two distinct conditions govern a load's preimage:

$$\sum_{j=1}^{d} \mathbf{w}_j \equiv 0 \pmod{q} \quad \text{(integral preimage)},$$
$$\sum_{j=1}^{d} \left\lfloor \frac{\mathbf{w}_j}{q} \right\rfloor \equiv 0 \pmod{q} \quad \text{(zero label)}.$$

The first says $\mathbf{w} \in \mathbf{T}\mathbb{Z}^d$, so $\mathbf{T}^{-1}\mathbf{w}$ is integral (Lemma 6.4). Once it holds, the second ensures that this preimage modulo q belongs to the zero fiber (Lemma 6.6). Individual carries may be nonzero; only their sum must vanish modulo q.

The charge condition places the colors in the zero-sum hyperplane H_0. Taking their canonical representatives, applying $\mathbf{T}^{-1}$, and reducing modulo q gives a bijection to the zero fiber (Lemma 6.5). We must therefore make the colors jointly uniform while enforcing compatible carries. For the affine output $\mathbf{x}'$ and its colors $\bar{\mathbf{w}}$, the bound of Lemma 6.6 reads

$$\Delta\big(\mathcal{L}(\mathbf{x}')\,, \mathsf{U}(\mathcal{F}_{\mathbf{T},0})\big) \le \Delta\big(\mathcal{L}(\bar{\mathbf{w}})\,, \mathsf{U}(H_0)\big) + \Pr[\text{incompatible carry}],$$

where $\mathcal{L}(\cdot)$ denotes a probability law and $\Delta(\cdot,\cdot)$ statistical distance.

A directed edge has an incidence column with entries $+1$ and -1 at its endpoints. Multiplying by the edge's source coordinate transfers load without changing the charge. A tree's columns form an integer basis of all zero-sum loads; adjoining a column of charge q gives a basis of $\mathbf{T}\mathbb{Z}^d$. Applying $\mathbf{T}^{-1}$ produces an integer matrix of determinant ± 1. Source-coordinate changes invertible modulo q preserve this modular invertibility (Corollary 6.8). The remaining task is to choose a tree and local maps with safe carries and adjustable colors.

Safety is analyzed on a *protected source*, obtained by conditioning each underlying Gaussian on a prescribed bounded interval. Protection includes the independent public Gaussian used below to close the body path into a cycle. The sampler never tests hidden inputs. Coupling the full and protected experiments fails only when a Gaussian leaves its interval. For our parameter choice, the total probability is exponentially small (Section 8.8).

A geometry with safe carries and adjustable colors. The *chameleon lobster* in Figure 1 supplies public parameters that adjust colors while ensuring compatible carries throughout this protected experiment. We call these parameters the public *controls*.

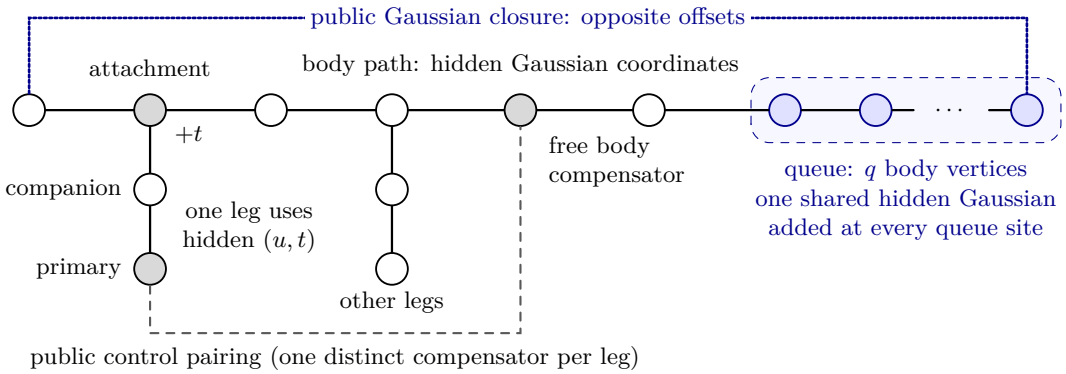


Figure 1: The chameleon lobster and its public controls. Solid segments are tree edges. The dotted closure supplies opposite public offsets; the dashed link pairs a primary with a separate body vertex whose distribution compensates for changes in the companion. Neither adds a hidden-source edge. The queue enclosure represents one shared source coordinate, with sites drawn together for clarity.

The *body* is a path of hidden Gaussian edge values, with vertex increments given by neighboring differences. An independent public Gaussian adds opposite endpoint offsets, closing these into a cycle law through β. All increments then have the same marginal law, but neighboring differences share a Gaussian and all increments sum to zero. Public initial loads, called *centers*, shift the body colors.

The *legs* provide more targeted color control. A leg runs from an outer *primary* through a *companion* to an *attachment* on the body. A public affine map acts on two hidden Gaussian coordinates (u, t). For a public shift ζ, the primary and companion loads sum to $q - 1 + \zeta - t$; the attachment receives $+t$ from the very same coordinate. This shared input makes the leg's total contribution independent of t. Safe affine maps also make the two off-body residue representatives have the same integer sum as their loads, so their carries cancel even when either load wraps. Body centers leave enough room for protected increments that body carries vanish. Proposition 8.17 proves the resulting carry compatibility.

The *queue* supplies the direction missing from tree incidence. It adds one hidden Gaussian scalar at all q queue sites, giving a column of total charge q and hence full rank. This shared scalar creates a global dependence. A congruence on public centers and leg shifts ensures that the initial loads also have charge divisible by q, making the affine offset integral. This is the *public charge condition*.

A leg's control changes both its primary and companion colors. To isolate the primary response, each leg is assigned a distinct free body vertex, its *compensator*, whose center law can offset the companion's change in expected color counts. These vertices lie outside the attachments and queue.

The initial assignment of attachment and queue sites, which also determines the compensators, is the public *placement*. Two subsequent permutations have separate purposes. A uniform body permutation distributes the cycle increments among body roles, making selected positions unlikely to be neighbors. An independent final permutation relabels all endpoints, so symmetric histogram and Fourier information applies to arbitrary endpoint positions. Neither permutation makes the endpoint colors independent.

Marking a color without observing the hidden input. At a fixed protected input, the leg's two canonical color representatives must sum to $q - 1 + \zeta - t$. The diffuse target law is uniform on such pairs with both colors in a trimmed interior (Definition 8.5); choosing the primary fixes its companion. This law misses boundary primary colors. A *marked color* is a public target, chosen independently of hidden inputs, that lets us add probability at any prescribed primary.

An affine map can create this bias without testing the hidden input. Fix its constant term to the marked color and its coefficient on t to zero. The primary then equals the marked color whenever $u = 0$, regardless of t. Mixing these *marked maps* with a diffuse branch gives a controllable contribution whose source-averaged mass is the marking probability times $\Pr[u = 0]$ under the protected Gaussian law. Because activation depends only on the first coordinate, it preserves the second coordinate's law at the attachment.

The difficulty is to retain only this intended bias. Marked maps also affect inputs off the line $u = 0$, and deleting unsafe maps biases the diffuse branch. Repairing the output after seeing the input would violate source independence. We must repair *one public distribution of maps* so that every protected input has its prescribed color law simultaneously.

Local calibration gives a positive law. For this repair, begin with uniform unrestricted affine maps. Their evaluations at distinct protected inputs are pairwise independent. A centered function of one evaluation therefore changes that input's color marginal and averages to zero at every other input. Summing such functions gives an exact reference response to the desired marginal adjustments.

We then exclude zero first-coordinate coefficients and unsafe maps. The excluded probability is small even conditional on an allowed evaluation, so deletion only slightly perturbs this response. A finite linear solve reweights the remaining, *regular* maps to restore the marginals and cancel the marked branch's off-line bias. A separate one-sided pointwise bound proves positivity on the regular support (Section C.5 and Corollary 7.2). Under its feasibility and marked-color conditions, Theorem 8.9 gives an exact public map law whose output at every protected input is uniform on the interior pairs above when $u \neq 0$, and, for our parameter choice, an equal mixture of that law and the pair with the marked primary and the companion fixed by the same integer sum when $u = 0$. Every supported map is safe at every protected input, with cancelling off-body carries, and has the nonzero first-coordinate coefficient required for modular invertibility. Bounded sampling approximates the complete flag–map law, including activation probabilities, while preserving this support.

One calibration principle, used twice. Both calibrations solve for weights on public choices. Local constraints specify color probabilities at each protected input; global constraints specify Fourier coefficients averaged over the protected source. The weight itself uses only public data.

The interface in Section 7 separates coefficients, weights, and responses. Synthesis Ψ maps coefficients z to weights relative to a reference law; response $\mathcal{M}$ records unnormalized constraints, including mass as the weight's expectation. Reference evaluation $\mathcal{E}$ converts coefficients into ideal

responses; $\mathcal{A}$ accounts for the actual construction:

$$\mathcal{M}\Psi = \mathcal{E}\mathcal{A}.$$

Solving $\mathcal{A}z = \tau$ for target coefficients τ gives response $\mathcal{E}\tau$. Only $\mathcal{A}$ must be invertible; $\mathcal{E}$ need not be. Both applications prove $\|\mathcal{A} - I\| \le \kappa < 1$, with I the identity, using the supremum norm locally and a weighted coefficient norm globally. Neumann inversion gives inverse norm at most $1/(1-\kappa)$ and a controlled coefficient correction (Theorem 7.1). The construction-specific work is to build this stable response and realize a probability law: positive directly in the local case, initially signed in the joint case.

III. From balanced colors to a uniform joint law. The local law gives safe, adjustable color marginals; joint uniformity still requires choosing the marked-color and body-center laws. Uniform marginals alone do not suffice: coordinates that all equal the same uniform field element have uniform marginals and a highly nonuniform joint law. Here body increments share Gaussians, each companion shares t with its attachment, and all queue sites share one scalar.

Charge conditioning would also couple the public controls. To keep an independent reference for these controls, we first omit charge and target the uniform law on the entire color space $\mathbb{F}_q^d$, in which coordinates are independently uniform. This is an auxiliary target: the candidate integer loads need not yet have integral affine preimages under $\mathbf{T}$. Conditioning the uniform target on zero color sum gives exactly $\mathsf{U}(H_0)$. We will impose public charge only after controlling the ambient error, restoring integral affine data.

A positive proposal that balances first moments. We first choose a positive public reference law, the *proposal*, that balances expected color counts. Marking can supply missing primary mass, but also moves the companion. To compensate using body centers, their effects must undergo the same random-shift averaging. We place the companion's reference color in a safe interval and add an independent public offset with the marginal law of one protected cycle difference. This offset copies a law, never an actual hidden increment. Companion and body-center responses then share the required averaging.

For the chosen public family and sufficiently large λ, these laws balance the histogram in expectation over the proposal controls, protected cycle, leg inputs, and shared queue scalar, at every fixed placement (Theorem 9.3). Independent uniform final relabeling gives uniform coordinate marginals. The proposal's control groups are independent at this placement; their endpoints remain dependent through the hidden inputs.

Pairing makes the first response invertible. Reweight a leg's marked color and its distinct compensator. Common shift averaging cancels the marked companion contribution. The diffuse target uses only the mark's side of the midpoint, mixing fairly at equality. Its change is therefore one scalar times a fixed profile difference. Dividing by the source-averaged activation probability leaves the primary adjustment plus this rank-one residual. Lemma 9.4 bounds this perturbation and supplies a stable inverse for the normalized *first response*. The attachment has zero first response, but its shared hidden coordinate remains in joint tests.

Hoeffding-type corrections give a triangular response. Joint uniformity requires every nonconstant product of Fourier characters to have zero expectation. Its *support* is the set of endpoints tested. We use a Hoeffding-type grading of public reweightings: degree counts how

many control groups interact in each term. This orders the correction of pairwise, triple, and higher biases.

A degree-r correction averages terms selecting r distinct control groups, each a primary with its compensator or an ordinary free body center. Every term is centered in each selected group: averaging that group under the proposal gives zero with the others fixed. Different groups affect disjoint endpoint sets. A test on fewer than r endpoints misses at least one selected group, whose averaging kills that term, even with the protected cycle fixed. This holds for every selection and hence their average, so higher-degree corrections preserve lower-degree constraints.

A small perturbation makes the response invertible. Fix a positive integer calibration order R and write $L = \lambda + 2$. The square coefficient response $\mathbf{A}_R$ of Section 9.5 realizes $\mathcal{A}$ through degree R. Triangularity orders the equations but does not bound their diagonal inverses. We compare the response with the identity after normalizing its leading terms.

Three effects remain. Unselected controls supply background correlations after first-moment balancing; selected groups can emit more endpoint factors than their leading response. Local profile bounds and control-pool normalization make these effects small, with only boundedly many pattern types at fixed order.

The third source is cycle dependence. Compare distinct body increments with independent copies of their protected marginal. Actual increments share Gaussians only at neighboring positions; the body permutation makes these overlaps rare. Theorem E.4 and Lemma F.12 bound this defect in a weighted L^2 coefficient norm matching distinct-group averaging. Each complete leg retains its shared hidden input, and complete expressions use one common queue scalar, averaged before inversion. Together the estimates give $\|\mathbf{A}_R - \mathbf{I}\| = O_R(L^{-3})$ (Section F.2.6). For each fixed order and sufficiently large λ, the common Neumann bound therefore gives a bounded inverse and a small coefficient correction.

Second application: exact signed joint calibration. We solve the public coefficient equation $\mathbf{A}_R\mathbf{z} = \mathbf{e}_0$, where $\mathbf{e}_0$ is the constant-one target and $\mathbf{z}$ is the coefficient vector to be chosen. At each degree, subtract the response of the coefficients already determined and solve the current diagonal block. The proposal has zero degree-one part, so corrections start at degree two (Section F.2.5), retaining all mixtures of paired and body controls. Theorem 9.5 gives a public weight W_R with expectation one under the fixed-placement proposal. After independent uniform endpoint relabeling, its *signed, uncharged, protected-source* endpoint law has every nontrivial Fourier coefficient of support at most R exactly zero. This is the global counterpart of the local marginal repair, but the weight may be negative.

At fixed order this polynomial-size system averages the known protected-source law in advance; its weight reads only public controls and placement.

Finite cancellation needs a complete Fourier tail bound. Finite cancellation leaves too many frequencies for individual bounds to suffice. We control their total energy using the *complete* Fourier norm $\mathfrak{F}_\rho$ of Section 9.1. Its square sums squared coefficients over all frequencies, weighted by support size. The analysis radius ρ penalizes large supports.

Set $\rho = L^{1/4}$, as in Section 9.9. For the public correction coefficients $\mathbf{c}$ and the endpoint law P_{W_R} weighted by W_R, Theorems G.5 and 9.5 and Proposition 9.6 give

$$\|\mathbf{c}\|_{\mathrm{simp},\rho} = O_R(L^{-3}), \qquad \mathfrak{F}_\rho(\mathsf{P}_{W_R}) = O_R(1).$$

By Lemma E.2, cancellation through R bounds the square root of the remaining Fourier energy by $O_R(\rho^{-(R+1)})$. Parseval and Cauchy–Schwarz give the same half-L^1 distance to uniformity on

$\mathbb{F}_q^d$ (Lemma E.1), hence decay $L^{-(R+1)/4}$ before charge conditioning. The radius affects only the analysis.

Fix placement, protected cycle, body permutation, and shared queue. The unweighted proposal factors into independent singletons and complete leg triples, so their character polynomials multiply. An all-degree product bound combines local coefficient bounds with exponential moments of the normalized imbalance of the conditional mean histogram. Changing one latent Gaussian affects only two cycle differences, controlling fluctuations; the queue adds uniformly small drift and is averaged after multiplication. For each summand in the correction average, local operators handle its bounded set of selected groups, including compensators; the product bound controls the remaining endpoints.

High moments turn the signed solution into a sampler. Mass one and Fourier cancellation do not ensure positivity, and pointwise closeness to one is unavailable. For each fixed calibration order R and each fixed positive integer moment order p, Proposition 9.6 gives, for sufficiently large L depending on these orders,

$$\|W_R - 1\|_{2p} \le C_{R,p} L^{-3}.$$

Here $C_{R,p}$ depends only on the two fixed orders. The norm uses the uncharged public proposal, uniformly over fixed placements. In the moment expansion, centering kills terms with a control index used only once. The large independent control pools bound the remaining overlaps, transferring coefficient control to high moments even for paired controls (Lemma G.6 and Corollary G.7).

Clip W_R to $[0, 2]$. Changed weights satisfy $|W_R - 1| \ge 1$, so the L^1 loss is bounded by the corresponding even moment. All fixed moment orders apply to the same weight at fixed calibration order, making clipping loss negligible uniformly over placements (Corollary 9.8). The protected endpoint channel contracts L^1, giving a nonnegative, unnormalized ambient measure negligibly close to the signed law.

Charge and the passage to the zero fiber. Charge is public: cycle increments sum to zero, each leg's second input cancels at its attachment, and the queue contributes a multiple of q. Conditioning on the public charge congruence therefore selects H_0 and restores integral affine offsets without inspecting the source. On this hyperplane, frequencies differing by a multiple of the all-ones vector define the same character.

We control the restricted law through the ambient approximation. The ambient uniform law assigns the hyperplane mass exactly $1/q$, so restricting and normalizing the nonnegative measure amplifies ambient error by only $O(q)$. For sufficiently large fixed calibration order, the ambient error is $o(1/q)$. Theorem 9.7 then gives clipped charge mass at least $1/(4q)$ and conditional error $O_R(qL^{-(R+1)/4})$ plus a negligible term. Polynomially bounded q preserves the negligibility of clipping.

Rejection with the clipped weight and public charge test samples this law. Retaining placement through retries preserves its marginal. Acceptance of order $1/q$ and a deterministic retry cap give worst-case polynomial time with controlled failure (Lemma H.1). Numerical approximations and fallbacks preserve integrality and modular invertibility on every execution (Section H.5). Conditional local maps, independent closure, and permutations complete the source-independent transport.

One public family for every polynomial sample budget. Use the public family and Gaussian source specified below. For each sufficiently large fixed calibration order R, an implementation

with numerical budget ϵ_{num} satisfies the bound of Theorem 9.9, for all sufficiently large λ, with C_R depending only on the fixed calibration order:

$$\Delta(\mathcal{L}(\mathbf{Qe} + \beta \bmod q), \mathsf{U}(\mathcal{F}_{\mathbf{T},0})) \le C_R q L^{-(R+1)/4} + \mathrm{negl}_R(L) + \epsilon_{\mathrm{num}}.$$

The terms are the charge-amplified Fourier remainder, the negligible sum of charge-amplified clipping loss and source-protection error, and numerical error. At fixed calibration order, the negligible bound is uniform over placements. One source coupling under the positive public law covers both endpoint disagreement and unsafe carries.

For the family below, $q = \Theta(L^{64})$. Fix an integer accuracy order $c \ge 1$ and choose a fixed calibration order $R \ge 4c + 267$. The Fourier term is then $O_R(L^{-c-3})$. With $\epsilon_{\mathrm{num}} \le \frac{1}{4} L^{-(c+2)}$ the total error is at most λ^{-c} eventually. The calibration order determines W_R and remains fixed as λ grows; the numerical budget may shrink. Both choices leave the public family and source unchanged.

For security, first fix this public family and Gaussian source. Given a uniform QPT distinguisher with classical sample access and nonnegligible advantage, fix its polynomial block-count bound t and an inverse-polynomial advantage lower bound on an infinite set of security parameters. By Corollary 5.5, the statistical loss is at most $qt\lambda^{-c} + q^{-N}$. Choose one fixed accuracy order making this loss smaller than half that advantage for all sufficiently large parameters in that set. The corresponding uniform PPT sampler's running-time polynomial may depend on the order. Thus the same LAR family supports every polynomial sample budget; this asserts neither uniform polynomial complexity in inverse accuracy nor negligible error for one fixed-order sampler. The quantum worst-case lattice foundation then yields the TC^0 weak PRF and noise-free finite-grid learning hardness under classical random-example access in Section 1.3.

These guarantees hold for the explicit family in Section 9.9, which takes q to be the least prime at least $2^{16} L^{64}$ and sets

$$N = \lambda, \qquad \sigma = L, \qquad \ell = 1, \qquad d = qL^{96} + 2L^{128}.$$

Thus $d = \Theta(L^{160})$, with the remaining choices verified in Section H.4. These conservative exponents reflect coarse bounds chosen to simplify the proof. We believe that a more refined analysis could reduce the modulus and block dimension.

1.3 Applications and Relation to Prior Work

Uniform inputs and nonstandard rounding. The Nearby Learning with Lattice Rounding (NLWLR) problem of Montgomery [Mon18b] provides fully deterministic rounding with polynomial modulus and no predetermined polynomial bound on the number of samples. Its real sampler first chooses a uniform lattice-coefficient vector and a small perturbation, then chooses a matrix constrained so that its product with the secret equals that lattice vector plus the perturbation. Deterministic lattice inversion recovers the coefficient vector, whose reduction modulo two is the label [Mon18b, full version, Section 3.1 and Definition 4.2]. The ideal sampler instead uses uniform matrices and independent uniform binary-vector labels. Thus the real-input interface uses a secret-dependent nearby-lattice sampler; it is not an experiment that draws real public matrices independently and exactly uniformly.

Our contribution combines an explicit deterministic quotient-rounding map, exactly uniform secret-independent public matrices, and fixed polynomial parameters supporting every polynomial sample budget. Blockwise rounding records a lattice coset, and the source-independent affine transport preserves the public-matrix law exactly. These structural properties supply the uniform random-example interface used in the weak-PRF and finite-grid learning consequences.

Compatibility with lower bounds for ordinary LWR. Newton and Richelson [NR23, Definition 5 and Theorem 2] analyze non-aborting pointwise reductions from LWE to ordinary scalar LWR: a fixed deterministic map acts on each sample separately, with the same dimension and source modulus. For prime polynomial source modulus and their specified range of rounding moduli, such a reduction with polynomial security loss yields an efficient attack on the source LWE problem. Our reduction changes the target to quotient-valued block rounding, uses a source-independent randomized affine transport, and shares randomness across the sample table in the zero-fiber step. It therefore lies outside that reduction model. We do not establish hardness of ordinary scalar LWR with polynomial modulus and unrestricted polynomial sample counts.

Weak pseudorandom functions over the field. Rounding-based weak PRFs with TC^0 evaluation appear in [BPR12, AKPW13]. Under the worst-case hardness condition above, the direct family in Section 10.2, with a uniformly sampled nonzero key, is a weak PRF on $\mathbb{F}_q^{d\times N}$ with range $\mathbb{F}_q$. Joint key-and-input evaluation is in TC^0, with polynomial-time constructible circuits, and each fixed nonzero key has exactly uniform output marginals under uniform inputs. For this same fixed public family, security holds against every uniform QPT adversary receiving classical random examples. Section 10.2 states the security and circuit guarantees; evaluation does not invoke the transport sampler.

Ding, Jain, and Komargodski [DJK26] obtain weak and strong PRFs in NC^1 from standard LWE with polynomial modulus-to-noise ratio. Our direct quotient-rounding construction is a weak PRF with TC^0 evaluation and security for random examples.

Noise-free learning. A learner that predicts fresh field labels with an inverse-polynomial improvement over success probability $1/q$ gives a decisional LAR distinguisher. We also give a concrete representation of the normalized labels by two-hidden-layer ReLU networks on a uniform finite grid, with polynomial size and polynomially bounded rational weights and biases of logarithmic bit length in the input dimension. A uniform QPT learner that, for every target, attains squared loss at most $1/100$ with probability at least $2/3$ yields the worst-case lattice algorithm above. These guarantees use independent classical noiseless random examples and permit improper hypotheses given by classical descriptions with deterministic classical finite-precision evaluation, as in Section 10.3.

The learning-to-distinguishing argument follows the rounding-based approach of Chen, Gollakota, Klivans, and Meka [CGKM22]. Section 10.3 states the precise guarantee, which does not cover arbitrary adaptive label queries.

A complementary route to improper learning hardness uses average-case assumptions about refuting random constraint satisfaction problems [DLSS14]. Within lattice-based learning hardness, Li, Zadik, and Zampetakis [LZZ25] study one-hidden-layer ReLU networks under standard Gaussian inputs with polynomially small additive Gaussian label noise. Their setting differs from our uniform finite-grid inputs and noiseless labels.

Organization. Sections 2 and 3 give the preliminaries and the formal rounding definitions. Sections 4 and 5 state the generic reductions and explain their distributional mechanisms. Section 6 develops the integer affine geometry. Section 7 gives the common calibration principle; Sections 8 and 9 apply it within the safe affine construction and joint endpoint analysis. Section 10 states the lattice, pseudorandomness, and learning consequences. Each section points to the appendix containing its complete proofs and technical verification.

2 Preliminaries

Throughout, residues modulo an integer $q \geq 2$ are represented by $\{0, \ldots, q-1\}$. For vectors and matrices, this convention is applied entrywise. We use $\|\cdot\|_2$ for the Euclidean norm, and $\mathbf{e}_v$ for the standard unit vector at coordinate v in the ambient space under discussion. Sample tables are written as indexed families in braces, retaining repeated entries and the order specified by their indices.

2.1 Lattices

We use the standard lattice terminology of [Pei15, Section 2.1].

Definition 2.1 (Lattice)**.** Let $m \geq 1$ be an integer. A *lattice* $\Lambda \subseteq \mathbb{R}^m$ is a discrete additive subgroup of $\mathbb{R}^m$.

Equivalently, for some integer $0 \leq n \leq m$, there are vectors $\mathbf{b}_1, \ldots, \mathbf{b}_n \in \mathbb{R}^m$ that are linearly independent over $\mathbb{R}$ and generate Λ by integer linear combinations. Write $\mathbf{B} = [\mathbf{b}_1 \mid \cdots \mid \mathbf{b}_n] \in \mathbb{R}^{m \times n}$. For integer coefficients $z_1, \ldots, z_n$, write $\mathbf{z} = (z_1, \ldots, z_n)^\intercal$. Then

$$\Lambda = \mathcal{L}(\mathbf{B}) := \{\mathbf{Bz} : \mathbf{z} \in \mathbb{Z}^n\} = \left\{ \sum_{i=1}^{n} z_i \mathbf{b}_i : z_1, \ldots, z_n \in \mathbb{Z} \right\}.$$

The matrix $\mathbf{B}$ is a *basis matrix* of Λ, and its columns form a *basis*: each point of Λ has a unique coefficient vector in $\mathbb{Z}^n$. The *rank* of Λ is $n = \dim(\mathsf{Span}_{\mathbb{R}}(\Lambda))$, the dimension of its real linear span, and m is its *ambient dimension*. The lattice is *full rank* if $n = m$.

Definition 2.2 (Zero-sum integer lattice)**.** For an integer $m \geq 1$, let $\mathbf{1} \in \mathbb{Z}^m$ be the column vector whose entries are all 1. Define

$$A_{m-1} := \{\mathbf{z} \in \mathbb{Z}^m : \mathbf{1}^\intercal \mathbf{z} = 0\}.$$

2.2 Rounding Matrices

We use the integer Smith normal form and its description of quotient groups as in [Sta16, Section 2]. A square integer matrix is *unimodular* if its determinant is ± 1, equivalently if it has an integer inverse.

Definition 2.3 (Rounding Matrix)**.** Let q be an odd prime and let $d \geq 1$ be an integer. Write $\mathbb{F}_q := \mathbb{Z}/q\mathbb{Z}$. Fix an integer $0 \leq r \leq d$ and set $\ell := d - r$. An integer matrix $\mathbf{T} \in \mathbb{Z}^{d \times d}$ is a *rounding matrix of modular rank r* if there exist unimodular matrices $\mathbf{U}, \mathbf{V} \in \mathbb{Z}^{d \times d}$ such that

$$\mathbf{UTV} = \mathrm{diag}(\mathbf{I}_r, q\mathbf{I}_\ell).$$

The Smith form connects modular rank to the integer quotient in which rounding outputs are encoded.

Lemma 2.4 (Modular rank and image lattice)**.** A rounding matrix as in Definition 2.3 satisfies

$$\operatorname{rank}_{\mathbb{F}_q}(\mathbf{T} \bmod q) = r, \qquad \dim_{\mathbb{F}_q} \ker(\mathbf{T} \bmod q) = \ell,$$
$$|\det \mathbf{T}| = q^\ell.$$

Its image lattice obeys

$$q\mathbb{Z}^d \subseteq \mathcal{L}(\mathbf{T}) \subseteq \mathbb{Z}^d, \qquad \mathbb{Z}^d/\mathcal{L}(\mathbf{T}) \cong \mathbb{F}_q^\ell.$$

For the fixed Smith matrix $\mathbf{U}$, the isomorphism takes the last ℓ coordinates of $\mathbf{Uz}$ modulo q. In particular the quotient has q^ℓ elements and is nontrivial exactly when $\ell \geq 1$.

The proof, including both endpoint ranks, is in Appendix A.2.

2.3 Probability Laws and Conditioning

For a random variable X taking values in a space Ω, write $\mathcal{L}(X)$ for its probability law: $\mathcal{L}(X)(A) = \Pr[X \in A]$ for every measurable set $A \subseteq \Omega$. On a finite or countable space all subsets are measurable; on a Euclidean space we use the Borel sets. For a law P on a finite or countable space, write $\mathsf{P}(\omega) := \mathsf{P}(\{\omega\})$ for its point probability. Thus $\mathsf{P}(A) = \sum_{\omega \in A} \mathsf{P}(\omega)$. The point-mass law δ_ω assigns probability one to ω. For laws P and Q, possibly on different spaces, their product $\mathsf{P} \otimes \mathsf{Q}$ is the joint law of independent variables $X_1 \sim \mathsf{P}$ and $X_2 \sim \mathsf{Q}$.

Definition 2.5 (Convolution of probability laws)**.** Let Ω be a finite or countable abelian group written additively, or $\mathbb{R}^m$ for an integer $m \geq 1$. For probability laws P, Q on Ω, their *convolution* is

$$\mathsf{P} * \mathsf{Q} := \mathcal{L}(X_1 + X_2)\,,$$

where $X_1 \sim \mathsf{P}$ and $X_2 \sim \mathsf{Q}$ are independent. In the finite or countable case, the point probabilities satisfy

$$(\mathsf{P} * \mathsf{Q})(\omega) = \sum_{\omega' \in \Omega} \mathsf{P}(\omega')\mathsf{Q}(\omega - \omega'), \qquad \omega \in \Omega.$$

Addition and subtraction are the group operations in Ω. In particular, on $\mathbb{Z}/q\mathbb{Z}$ for an integer $q \geq 2$, they are taken modulo q.

For the remainder of this subsection, all state spaces are finite or countable. For laws P and Q on the same space Ω, their *statistical distance* (SD), or total variation distance, is

$$\Delta(\mathsf{P}, \mathsf{Q}) := \frac{1}{2} \sum_{\omega \in \Omega} |\mathsf{P}(\omega) - \mathsf{Q}(\omega)|.$$

A randomized transformation assigns to each input point a probability law on a fixed output space. Applying the same assignment to inputs with laws P and Q produces output laws whose SD is at most $\Delta(\mathsf{P}, \mathsf{Q})$. A deterministic function is the special case in which each assigned law is a point mass.

Definition 2.6 (Coupling)**.** Let P and Q be probability laws on the same finite or countable

space Ω. A *coupling* of these laws is a joint law of a pair (X_1, X_2) in Ω^2 with marginal laws P and Q, respectively. The coordinates need not be independent.

Under any such coupling [LP17, Proposition 4.7],

$$\Delta(\mathsf{P}, \mathsf{Q}) \le \Pr[X_1 \neq X_2].$$

For laws $\mathsf{P}_1, \mathsf{Q}_1$ on one space and $\mathsf{P}_2, \mathsf{Q}_2$ on another, the product bound is [LP17, Exercise 4.4]

$$\Delta(\mathsf{P}_1 \otimes \mathsf{P}_2, \mathsf{Q}_1 \otimes \mathsf{Q}_2) \le \Delta(\mathsf{P}_1, \mathsf{Q}_1) + \Delta(\mathsf{P}_2, \mathsf{Q}_2)\,.$$

Iterating this inequality bounds the SD between products of any finite number of laws by the sum of the SDs of corresponding factors. These bounds also apply on countable spaces; the verification is in Section A.3.

For $A \subseteq \Omega$ with $\mathsf{P}(A) > 0$, the conditional law $\mathsf{P}(\cdot \mid A)$ is a law on the same space Ω, defined by

$$\mathsf{P}(A' \mid A) := \frac{\mathsf{P}(A' \cap A)}{\mathsf{P}(A)}, \qquad A' \subseteq \Omega.$$

It assigns zero probability to $A^c := \Omega \setminus A$. The cost of conditioning is exactly the discarded probability mass:

$$\Delta(\mathsf{P}, \mathsf{P}(\cdot \mid A)) = \mathsf{P}(A^c).$$

We call $\mathsf{P}(\cdot \mid A)$ the *truncation* of P to A; it includes normalization by $\mathsf{P}(A)$. If independent variables $X_1 \sim \mathsf{P}$ and $X_2 \sim \mathsf{Q}$ are jointly conditioned on $\{X_1 \in A_1,\, X_2 \in A_2\}$, where $\mathsf{P}(A_1) > 0$ and $\mathsf{Q}(A_2) > 0$, their conditional joint law is $\mathsf{P}(\cdot \mid A_1) \otimes \mathsf{Q}(\cdot \mid A_2)$. Thus they remain independent under this conditional law. The same statement holds for any finite number of independent variables, each restricted to a set of positive probability in its own coordinate. These conditioning identities are proved in Section A.3. They let us analyze protected sources by truncation without making the sampler inspect a source realization.

2.4 Finite-Space Densities

Let Ω be a nonempty finite set, equipped with the reference law $\mathsf{U}(\Omega)$. Write $\mathbb{R}^\Omega$ for the vector space of real functions on Ω, and 1 for the constant function with value one. Unless another law is explicitly specified, probabilities and expectations in this subsection refer to this law. In particular, for a real function $f : \Omega \to \mathbb{R}$,

$$\mathbb{E}f := \frac{1}{|\Omega|} \sum_{\omega \in \Omega} f(\omega).$$

For $A \subseteq \Omega$, superscript c denotes its complement in Ω, and $\mathbf{1}_A(\omega)$ is one when $\omega \in A$ and zero otherwise.

A function $\mathsf{f} : \Omega \to \mathbb{R}_{\ge 0}$ with $\mathbb{E}\mathsf{f} = 1$ is a *probability density relative to the uniform law.* It defines a probability law whose point mass at ω is $\mathsf{f}(\omega)/|\Omega|$. More generally, a real function f defines a signed measure with point mass $f(\omega)/|\Omega|$. Its total mass is $\mathbb{E}f$, which need not be one or nonnegative. For real functions f, g on Ω, define the inner product and supremum norm by

$$\langle f, g \rangle := \mathbb{E}[fg], \qquad \|f\|_\infty := \max_{\omega \in \Omega} |f(\omega)|.$$

For a linear operator between subspaces of real functions on Ω, its $\infty \to \infty$ norm is the supremum of the output supremum norm over input functions of supremum norm at most one.

2.5 Canonical Representatives and Carries

Use the prime field of Definition 2.3. The following notation makes the canonical residue convention explicit.

Definition 2.7 (Canonical lift)**.** For an integer $d \geq 1$ and $\mathbf{x} \in \mathbb{F}_q^d$, its *canonical integer lift* $[\mathbf{x}]_q$ is the unique vector in $\mathbb{Z}^d$ whose reduction modulo q equals $\mathbf{x}$ and whose entries lie in $\{0, \ldots, q-1\}$.

The same notation applies to scalars and vectors of any dimension. For an integer k, its canonical residue is $[k \bmod q]_q$ and its *carry* is $\lfloor k/q \rfloor$. They satisfy

$$k = [k \bmod q]_q + q\lfloor k/q \rfloor.$$

The carry is the integer quotient in this decomposition, including when $k < 0$. These conventions apply coordinatewise to integer vectors. We call an integer in $\{0, \ldots, q-1\}$ a *color*.

2.6 Error Distributions and Learning with Errors

For Gaussian errors, we use the discrete Gaussian convention of [GPV08]; see also [Pei15, Section 2.3]. Fix an integer $m \geq 1$. For a width parameter $\sigma > 0$ and a center $\mathbf{c} \in \mathbb{R}^m$, define

$$\rho_{\sigma,\mathbf{c}}(\mathbf{x}) := \exp\left(-\pi \frac{\|\mathbf{x} - \mathbf{c}\|_2^2}{\sigma^2}\right), \qquad \mathbf{x} \in \mathbb{R}^m.$$

For a nonempty countable set $S \subseteq \mathbb{R}^m$, write

$$\rho_{\sigma,\mathbf{c}}(S) := \sum_{\mathbf{y} \in S} \rho_{\sigma,\mathbf{c}}(\mathbf{y}).$$

Definition 2.8 (Discrete Gaussian distribution)**.** If $0 < \rho_{\sigma,\mathbf{c}}(S) < \infty$, the *discrete Gaussian distribution* $\mathsf{D}_{S,\sigma,\mathbf{c}}$ assigns to each $\mathbf{x} \in S$ probability

$$\mathsf{D}_{S,\sigma,\mathbf{c}}(\mathbf{x}) = \frac{\rho_{\sigma,\mathbf{c}}(\mathbf{x})}{\rho_{\sigma,\mathbf{c}}(S)}.$$

It assigns zero probability outside S.

The strict positivity of the normalizing sum follows from $S \neq \varnothing$. We omit the center when it is zero, writing $\mathsf{D}_{S,\sigma} := \mathsf{D}_{S,\sigma,\mathbf{0}}$ and $\rho_\sigma := \rho_{\sigma,\mathbf{0}}$. Every lattice and every translate of a lattice has finite Gaussian mass for all $\sigma > 0$ and all centers, so these sets are valid choices of S. For the integer lattice, the centered distribution factors coordinatewise:

$$\mathsf{D}_{\mathbb{Z}^m,\sigma} = \left(\mathsf{D}_{\mathbb{Z},\sigma}\right)^{\otimes m}.$$

The continuous Gaussian with Lebesgue density $\sigma^{-m}\rho_{\sigma,\mathbf{c}}(\mathbf{x})$ on $\mathbb{R}^m$ has mean $\mathbf{c}$ and covariance $\sigma^2\mathbf{I}_m/(2\pi)$. In one dimension its standard deviation is $\sigma/\sqrt{2\pi}$.

The next inequality bounds the exponential moment of each real linear functional of a discrete Gaussian whose weight is centered at zero [MP12, full version, Lemma 2.8]. It uses the Gaussian-width convention above and requires no lower bound on the width beyond positivity.

Lemma 2.9 (Exponential moments of a centered discrete Gaussian)**.** Let $\Lambda \subseteq \mathbb{R}^m$ be a lattice, let $\sigma > 0$, and sample $\mathbf{x} \hookleftarrow \mathsf{D}_{\Lambda,\sigma}$. For $\mathbf{u}_{\mathrm{test}} \in \mathsf{Span}_{\mathbb{R}}(\Lambda)$ and $t \in \mathbb{R}$,

$$\mathbb{E}\Big[e^{t\langle \mathbf{u}_{\mathrm{test}}, \mathbf{x}\rangle}\Big] \le e^{\sigma^2 t^2 \|\mathbf{u}_{\mathrm{test}}\|_2^2/(4\pi)}.$$

Learning with errors. The LWE problem was introduced by Regev [Reg05]. We use a shared-secret block formulation: one secret is sampled for the entire collection of blocks. When each error block consists of independent draws from the same scalar error law, this is the usual LWE experiment [Pei15, Section 4.2.1] with scalar samples grouped into blocks.

Definition 2.10 (Shared-secret block LWE)**.** Let q be an odd prime, let $N, d \ge 1$ be integers, and let χ be a distribution on $\mathbb{F}_q^d$. For an integer block count $t \ge 1$, draw one secret $\mathbf{s} \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right)$ and, independently for $i = 1, \dots, t$, draw

$$\mathbf{A}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right), \qquad \mathbf{e}_i \hookleftarrow \chi, \qquad \mathbf{u}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^d\right).$$

All these choices are mutually independent and independent of the secret. The real and ideal laws of indexed families of t blocks are

$$\mathsf{Real}_{\chi,t}^{\mathsf{LWE}} := \mathcal{L}\Big(\{(\mathbf{A}_i, \mathbf{A}_i\mathbf{s} + \mathbf{e}_i)\}_{i=1}^t\Big),$$
$$\mathsf{Ideal}_t^{\mathsf{LWE}} := \mathcal{L}\Big(\{(\mathbf{A}_i, \mathbf{u}_i)\}_{i=1}^t\Big),$$

where responses are computed in $\mathbb{F}_q^d$.

A *search* LWE algorithm receives a draw from the real experiment and succeeds exactly when its output is the sampled vector $\mathbf{s}$. *Decisional* LWE asks to distinguish the real and ideal laws. For a randomized algorithm $\mathcal{D}$ returning a bit, its advantage is

$$\Big|\Pr[\mathcal{D}(\mathsf{Real}_{\chi,t}^{\mathsf{LWE}}) = 1] - \Pr[\mathcal{D}(\mathsf{Ideal}_t^{\mathsf{LWE}}) = 1]\Big|.$$

Definition 2.11 (Gaussian LWE)**.** For $\sigma > 0$, *Gaussian* LWE is the specialization of Definition 2.10 to

$$\chi = \mathsf{D}_{\mathbb{Z}^d,\sigma} \bmod q = \left(\mathsf{D}_{\mathbb{Z},\sigma} \bmod q\right)^{\otimes d}.$$

The notation mod$q$ on a law means its image under coordinatewise reduction from $\mathbb{Z}^d$ to $\mathbb{F}_q^d$ (or from $\mathbb{Z}$ to $\mathbb{F}_q$ in the scalar factor). Equivalently, for every $\mathbf{e} \in \mathbb{F}_q^d$,

$$\chi(\mathbf{e}) = \frac{\rho_\sigma([\mathbf{e}]_q + q\mathbb{Z}^d)}{\rho_\sigma(\mathbb{Z}^d)}.$$

Here t blocks contain td scalar LWE samples with independent errors. We use this discrete Gaussian convention throughout, rather than rounding a continuous Gaussian to the nearest integer; see [Pei15, Section 4.2.1].

We also consider errors drawn uniformly from the inputs on which a public function is zero.

Definition 2.12 (Uniform-zero-fiber LWE)**.** In the setting of Definition 2.10, fix an integer $\ell \geq 1$ and a public function $g\colon \mathbb{F}_q^d \to \mathbb{F}_q^\ell$, whose *zero fiber*

$$\mathcal{F}_{g,0} := g^{-1}(\mathbf{0}) = \left\{\mathbf{e} \in \mathbb{F}_q^d : g(\mathbf{e}) = \mathbf{0}\right\}$$

is nonempty. *Uniform-zero-fiber* LWE associated with g is the specialization of Definition 2.10 to $\chi = \mathsf{U}(\mathcal{F}_{g,0})$.

Uniformity is over the entire fiber, which need not be a linear subspace. Thus every fiber point has probability $|\mathcal{F}_{g,0}|^{-1}$.

3 Learning with Algebraic Rounding

Definition 3.1 (Algebraic Rounding)**.** Let q be an odd prime, let $d \geq 1$ be an integer, and let $\mathbf{T} \in \mathbb{Z}^{d\times d}$ be a rounding matrix of modular rank r, as in Definition 2.3, with output dimension ℓ. The associated *algebraic rounding function* is

$$\begin{aligned}\mathsf{AR}_{\mathbf{T}}\colon \mathbb{Z}^d &\longrightarrow \mathbb{Z}^d,\\ \mathbf{x} &\longmapsto \lfloor \mathbf{T}\mathbf{x}/q \rfloor.\end{aligned}$$

The product and division inside the floor are computed over $\mathbb{Q}$.

An integer rounded output determines a coset in $\mathbb{Z}^d/\mathcal{L}(\mathbf{T})$ that depends only on the input modulo q. Smith coordinates encode this coset as a label in $\mathbb{F}_q^\ell$. The corresponding rounding directions translate labels bijectively, giving fibers of equal size. We use these coordinates and directions to define the label function.

Fix public Smith matrices $\mathbf{U}, \mathbf{V}$ as in Definition 2.3. The quotient isomorphism of Lemma 2.4 gives the following coordinate maps.

Definition 3.2 (Quotient coordinates)**.** Define the *coordinate embedding* $\mathbf{E}$ and the *quotient coordinate matrix* $\mathbf{P}$ by

$$\mathbf{E} := \begin{pmatrix}\mathbf{0}_{r\times\ell}\\ \mathbf{I}_\ell\end{pmatrix} \in \mathbb{Z}^{d\times\ell}, \qquad \mathbf{P} := \mathbf{E}^\intercal\mathbf{U} \in \mathbb{Z}^{\ell\times d}.$$

For $\mathbf{z} \in \mathbb{Z}^d$, $\mathbf{P}\mathbf{z}$ extracts the last ℓ coordinates of $\mathbf{U}\mathbf{z}$. Reducing them modulo q labels its quotient coset. The matching columns of $\mathbf{V}$ define the rounding directions.

Definition 3.3 (Rounding directions)**.** Define the *integer rounding direction matrix* $\mathbf{J}_{\mathrm{int}} := \mathbf{V}\mathbf{E} \in \mathbb{Z}^{d\times\ell}$ and its *modular reduction* $\mathbf{J} := \mathbf{J}_{\mathrm{int}} \bmod q \in \mathbb{F}_q^{d\times\ell}$.

These directions shift the argument of the floor by integer vectors. Its coordinatewise fractional parts therefore remain unchanged, while the rounded output shifts by the same integer vector. The quotient encoding turns this shift into the prescribed label change (Lemma 3.5).

For an input in $\mathbb{F}_q^d$, first use the canonical integer lift of Definition 2.7, then apply the integer rounding map and reduce its quotient coordinates modulo q.

Definition 3.4 (Encoded label function)**.** The *encoded label function* associated with $\mathbf{T}$ and $\mathbf{P}$ is defined by

$$f_{\mathbf{T}} \colon \mathbb{F}_q^d \longrightarrow \mathbb{F}_q^\ell,$$
$$\mathbf{x} \longmapsto \mathbf{P}\mathsf{AR}_{\mathbf{T}}([\mathbf{x}]_q) \bmod q.$$

The notation $f_{\mathbf{T}}$ suppresses the fixed choice of $\mathbf{P}$. When $\ell = 0$, the label space $\mathbb{F}_q^0$ consists of the single empty vector; the definitions and properties below include this case.

Lemma 3.5 (Basic properties of LAR)**.** Under the definitions above, the following properties hold.

Well-defined quotient encoding. For all $\mathbf{x}, \mathbf{z} \in \mathbb{Z}^d$,

$$\mathsf{AR}_{\mathbf{T}}(\mathbf{x} + q\mathbf{z}) = \mathsf{AR}_{\mathbf{T}}(\mathbf{x}) + \mathbf{T}\mathbf{z}.$$

The map $\mathbf{z} \mapsto \mathbf{P}\mathbf{z} \bmod q$ induces an isomorphism $\mathbb{Z}^d/\mathcal{L}(\mathbf{T}) \cong \mathbb{F}_q^\ell$. Consequently, for $\mathbf{x} \in \mathbb{F}_q^d$, $f_{\mathbf{T}}(\mathbf{x})$ encodes the quotient output $\mathsf{AR}_{\mathbf{T}}([\mathbf{x}]_q) + \mathcal{L}(\mathbf{T})$ and is unchanged if another integer lift is used.

Translation equivariance. The matrix $\mathbf{J}$ has full column rank over $\mathbb{F}_q$, and its columns form a basis of $\ker(\mathbf{T} \bmod q)$. For all $\mathbf{x} \in \mathbb{F}_q^d$ and $\mathbf{a} \in \mathbb{F}_q^\ell$,

$$f_{\mathbf{T}}(\mathbf{x} + \mathbf{J}\mathbf{a}) = f_{\mathbf{T}}(\mathbf{x}) + \mathbf{a}.$$

All operations in this identity are over $\mathbb{F}_q$.

Balanced fibers. For every label $\mathbf{y} \in \mathbb{F}_q^\ell$, its *fiber*, the set of inputs mapped to that label, satisfies

$$\left|\left\{\mathbf{x} \in \mathbb{F}_q^d : f_{\mathbf{T}}(\mathbf{x}) = \mathbf{y}\right\}\right| = q^r.$$

The complete proof is in Appendix A.5. The Smith coordinates make integral changes of lift disappear in the quotient. Their rounding directions then translate each label once along every kernel coset, which gives balanced fibers.

If $\ell = 1$ and the chosen Smith decomposition satisfies $\mathbf{P} = \mathbf{1}^\mathsf{T}$, where $\mathbf{1} \in \mathbb{Z}^d$ is the all-ones vector, then

$$f_{\mathbf{T}}(\mathbf{x}) = \sum_{j=1}^{d} \left\lfloor \frac{(\mathbf{T}[\mathbf{x}]_q)_j}{q} \right\rfloor \bmod q \qquad (\mathbf{x} \in \mathbb{F}_q^d).$$

Thus in this case the label is a single field element obtained by summing the integer rounded coordinates.

3.1 Decisional LAR assumption

Fix a family of public parameters $\mathsf{pp} = \mathsf{pp}(\lambda)$ indexed by an integer security parameter $\lambda \geq 1$. A fixed deterministic algorithm computes $\mathsf{pp}(\lambda)$ from the unary input 1^λ in time polynomial in λ.

These parameters specify the prime q, the integer dimensions N, d, r, ℓ, the rounding matrix $\mathbf{T}$, and the chosen Smith matrices $\mathbf{U}, \mathbf{V}$, hence also $\mathbf{P}$ and $f_{\mathbf{T}}$. Require $N \geq 1$ and $\ell \geq 1$. The dimensions and the binary encoding length of pp are polynomially bounded in λ. The parameter family is fixed before the sample-count function and the adversary are chosen.

A nonnegative function of λ is *negligible* if, for every constant $c > 0$, it is at most λ^{-c} for all sufficiently large λ.

We first define the real and ideal LAR distributions.

Definition 3.6 (Real and ideal LAR distributions)**.** For an integer sample count $t \geq 1$, draw one secret $\mathbf{s} \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right)$. Independently of that secret and independently for each $i = 1, \ldots, t$, draw

$$\mathbf{M}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d \times N}\right), \qquad \mathbf{u}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{\ell}\right)$$

with the matrices and uniform labels mutually independent. The *real* LAR *distribution* and the *ideal* LAR *distribution* are the following laws of indexed families of t samples:

$$\mathsf{Real}_t^{\mathsf{LAR}} := \mathcal{L}\Big(\{(\mathbf{M}_i, f_{\mathbf{T}}(\mathbf{M}_i\mathbf{s}))\}_{i=1}^t\Big),$$
$$\mathsf{Ideal}_t^{\mathsf{LAR}} := \mathcal{L}\Big(\{(\mathbf{M}_i, \mathbf{u}_i)\}_{i=1}^t\Big).$$

The products $\mathbf{M}_i\mathbf{s}$ are computed over $\mathbb{F}_q$ before applying $f_{\mathbf{T}}$. Conditional on the single sampled secret, the real samples are independent; their unconditional joint law averages over that same secret. In the ideal law, the matrices and labels are all independent.

Definition 3.7 (Decisional LAR assumption)**.** For a uniform QPT distinguisher $\mathcal{D}$ on classical inputs, define its decisional advantage by

$$\mathsf{Adv}_{\mathcal{D}}^{\mathsf{dLAR}}(\lambda, t) := \big|\Pr[\mathcal{D}(1^\lambda, \mathsf{pp}, \mathsf{Real}_t^{\mathsf{LAR}}) = 1] - \Pr[\mathcal{D}(1^\lambda, \mathsf{pp}, \mathsf{Ideal}_t^{\mathsf{LAR}}) = 1]\big|.$$

The parameter family satisfies the *decisional* LAR *assumption* if this advantage is negligible in λ for every such $\mathcal{D}$ and every positive integer-valued sample-count function $t = t(\lambda)$ bounded by a polynomial in λ.

This is the decisional LAR assumption against QPT adversaries. Restricting the same definition to PPT adversaries gives its classical version.

Lemma 3.8 (Information-theoretic identifiability)**.** Fix public parameters as in Definition 3.6, with $\ell \geq 1$, and let $t \geq 1$ be an integer. Let $\mathbf{M}_1, \ldots, \mathbf{M}_t$ be independent and uniform on $\mathbb{F}_q^{d \times N}$. For every fixed $\mathbf{s} \in \mathbb{F}_q^N$, the examples $\{(\mathbf{M}_i, f_{\mathbf{T}}(\mathbf{M}_i\mathbf{s}))\}_{i=1}^t$ uniquely determine $\mathbf{s}$ with probability at least $1 - (q^N - 1)q^{-\ell t}$.

The proof and recovery procedure are in Section A.6.

4 From Uniform-Zero-Fiber LWE to Decisional LAR

To establish the decisional LAR assumption from uniform-zero-fiber LWE hardness, we use a decomposition of each input into its encoded label and a residual in the zero fiber. Fix a public parameter family as in Section 3.1. For each security parameter, write

$$\mathcal{F}_{\mathbf{T},0} := f_{\mathbf{T}}^{-1}(\mathbf{0}) = \Big\{\mathbf{e} \in \mathbb{F}_q^d : f_{\mathbf{T}}(\mathbf{e}) = \mathbf{0}\Big\}.$$

By the balanced-fiber property in Lemma 3.5, $|\mathcal{F}_{\mathbf{T},0}| = q^r = q^{d-\ell}$, so each point has probability q^{-r} under $\mathsf{U}(\mathcal{F}_{\mathbf{T},0})$. In particular, this set is nonempty. Set $\chi := \mathsf{U}(\mathcal{F}_{\mathbf{T},0})$, the error law in Definition 2.12 for $g = f_{\mathbf{T}}$. Membership in this zero fiber means that the encoded label is zero; it does not require $(\mathbf{T} \bmod q)\mathbf{e} = \mathbf{0}$.

Define the *zero-fiber decomposition map* by

$$\begin{aligned}\Phi_{\mathbf{T}} \colon \mathcal{F}_{\mathbf{T},0} \times \mathbb{F}_q^\ell &\longrightarrow \mathbb{F}_q^d,\\ (\mathbf{e}, \mathbf{y}) &\longmapsto \mathbf{e} + \mathbf{J}\mathbf{y}.\end{aligned}$$

All operations are over $\mathbb{F}_q$.

Lemma 4.1 (Zero-fiber decomposition)**.** The zero-fiber decomposition map $\Phi_{\mathbf{T}}$ defined above is a bijection. For every $\mathbf{x} \in \mathbb{F}_q^d$, its inverse is given by

$$\Phi_{\mathbf{T}}^{-1}(\mathbf{x}) = (\mathbf{x} - \mathbf{J} f_{\mathbf{T}}(\mathbf{x}), f_{\mathbf{T}}(\mathbf{x})).$$

Proof. For every $\mathbf{e} \in \mathcal{F}_{\mathbf{T},0}$ and $\mathbf{y} \in \mathbb{F}_q^\ell$, translation equivariance in Lemma 3.5 gives

$$f_{\mathbf{T}}(\Phi_{\mathbf{T}}(\mathbf{e}, \mathbf{y})) = f_{\mathbf{T}}(\mathbf{e} + \mathbf{J}\mathbf{y}) = f_{\mathbf{T}}(\mathbf{e}) + \mathbf{y} = \mathbf{y}.$$

Thus the image uniquely determines $\mathbf{y}$; subtracting $\mathbf{J}\mathbf{y}$ then recovers $\mathbf{e}$. Hence $\Phi_{\mathbf{T}}$ is injective.

Conversely, fix any $\mathbf{x} \in \mathbb{F}_q^d$ and set $\mathbf{y} := f_{\mathbf{T}}(\mathbf{x})$ and $\mathbf{e} := \mathbf{x} - \mathbf{J}\mathbf{y}$. Applying the translation equivariance of Lemma 3.5 with offset $-\mathbf{y}$ yields

$$f_{\mathbf{T}}(\mathbf{e}) = f_{\mathbf{T}}(\mathbf{x} - \mathbf{J}\mathbf{y}) = f_{\mathbf{T}}(\mathbf{x}) - \mathbf{y} = \mathbf{0}.$$

Therefore $\mathbf{e} \in \mathcal{F}_{\mathbf{T},0}$, and $\Phi_{\mathbf{T}}(\mathbf{e}, \mathbf{y}) = \mathbf{x}$. This proves surjectivity and the stated inverse formula. □

By Lemma 4.1, if $\mathbf{x} \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^d\right)$, then $\Phi_{\mathbf{T}}^{-1}(\mathbf{x})$ is uniform over $\mathcal{F}_{\mathbf{T},0} \times \mathbb{F}_q^\ell$. Consequently, its residual $\mathbf{x} - \mathbf{J} f_{\mathbf{T}}(\mathbf{x})$ and label $f_{\mathbf{T}}(\mathbf{x})$ are independent and uniformly distributed over $\mathcal{F}_{\mathbf{T},0}$ and $\mathbb{F}_q^\ell$, respectively. Conversely, independent draws $\mathbf{e} \hookleftarrow \mathsf{U}(\mathcal{F}_{\mathbf{T},0})$ and $\mathbf{y} \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^\ell\right)$ give a uniform vector $\mathbf{e} + \mathbf{J}\mathbf{y}$ in $\mathbb{F}_q^d$ whose label is $\mathbf{y}$.

For any classical or quantum algorithm $\mathcal{B}$ with classical input and a classical output bit, define its advantage on these block LWE experiments by

$$\begin{aligned}\mathsf{Adv}_{\mathcal{B}}^{\mathsf{LWE}}(\lambda, t) := \big|\, &\Pr[\mathcal{B}(1^\lambda, \mathsf{pp}, \mathsf{Real}_{\chi,t}^{\mathsf{LWE}}) = 1]\\ &- \Pr[\mathcal{B}(1^\lambda, \mathsf{pp}, \mathsf{Ideal}_{t}^{\mathsf{LWE}}) = 1]\big|.\end{aligned}$$

Theorem 4.2 (From uniform-zero-fiber LWE to decisional LAR)**.** Fix an efficiently computable public parameter family as in Section 3.1, and set $\chi = \mathsf{U}(\mathcal{F}_{\mathbf{T},0})$. For every uniform QPT distinguisher $\mathcal{D}$ on classical sample tables and every positive integer-valued sample-count function $t = t(\lambda)$ bounded by a polynomial in λ, there is a uniform QPT distinguisher $\mathcal{B}$ using t input blocks and one call to $\mathcal{D}$ such that

$$\mathsf{Adv}^{\mathsf{dLAR}}_{\mathcal{D}}(\lambda, t) \le q\,\mathsf{Adv}^{\mathsf{LWE}}_{\mathcal{B}}(\lambda, t) + q^{-N}.$$

Proof sketch. First consider exact uniform auxiliary choices. The reduction samples one uniform vector $\mathbf{a}_0 \in \mathbb{F}_q^N$ shared across blocks and independent uniform labels $\mathbf{y}_i \in \mathbb{F}_q^\ell$, all independently of the input. Given input blocks $(\mathbf{A}_i, \mathbf{b}_i)$, it outputs $(\mathbf{M}_i, \mathbf{y}_i)$ with the rank-one update

$$\mathbf{M}_i := -\mathbf{A}_i + (\mathbf{b}_i + \mathbf{J}\mathbf{y}_i)\mathbf{a}_0^\intercal.$$

All arithmetic is over $\mathbb{F}_q$. For real input, $\mathbf{b}_i = \mathbf{A}_i\mathbf{s} + \mathbf{e}_i$. Set $\mathbf{x}_i := \mathbf{e}_i + \mathbf{J}\mathbf{y}_i$. By Lemma 4.1, these vectors are independent and uniform over $\mathbb{F}_q^d$, independent of the input matrices, secret, and $\mathbf{a}_0$, with $f_{\mathbf{T}}(\mathbf{x}_i) = \mathbf{y}_i$. Substitution gives the identity

$$\mathbf{M}_i = -\mathbf{A}_i(\mathbf{I}_N - \mathbf{s}\mathbf{a}_0^\intercal) + \mathbf{x}_i\mathbf{a}_0^\intercal.$$

Hence, whenever $\mathbf{a}_0^\intercal\mathbf{s} = 1$,

$$\mathbf{M}_i\mathbf{s} = \mathbf{x}_i, \qquad f_{\mathbf{T}}(\mathbf{M}_i\mathbf{s}) = \mathbf{y}_i.$$

For each fixed secret and auxiliary vector satisfying this alignment condition, the update is a surjective linear map of the independent uniform input matrix and vector, so the output matrices are independent and uniform. Averaging over aligned choices yields the real LAR table conditioned on a nonzero secret. Otherwise, $\det(\mathbf{I}_N - \mathbf{s}\mathbf{a}_0^\intercal) = 1 - \mathbf{a}_0^\intercal\mathbf{s} \neq 0$. Conditional on the vectors $\mathbf{x}_i$, the update is then an affine bijection in each input matrix, so the output matrices and labels have exactly the ideal product law.

For every nonzero secret, the alignment event $\mathbf{a}_0^\intercal\mathbf{s} = 1$ has probability $1/q$. Thus the transformed real LWE table mixes the nonzero-secret real LAR law with the ideal law, with real-branch weight $(1 - q^{-N})/q$. Ideal LWE input maps exactly to the ideal LAR law. Comparing these mixtures gives the factor q in the advantage bound; the zero-secret event accounts for the additive q^{-N} allowance.

Capped rejection sampling implements the auxiliary choices in worst-case polynomial time. Its error fits within the unused portion of the zero-secret allowance, retaining the stated bound. The reduction prepares a classical table and calls $\mathcal{D}$ once, so it remains uniform QPT. The full game-based proof, including the exact mixture identities and sampling-error budget, is deferred to Appendix A.7. □

5 From Gaussian LWE to Uniform-Zero-Fiber LWE

A sufficiently accurate source-independent affine transport reduces Gaussian LWE to uniform-zero-fiber LWE. We state the transport property and prove the generic reduction under that hypothesis. Section 9 constructs the transport for the public rounding family used in this paper.

5.1 T-adapted affine fiber transport

Retain the public rounding matrix $\mathbf{T}$, its quotient encoding, and the zero fiber $\mathcal{F}_{\mathbf{T},0}$ from Section 4, with $\ell \ge 1$. Let χ be a source distribution on $\mathbb{F}_q^d$.

Definition 5.1 (**T**-adapted affine fiber transport)**.** Fix an accuracy $\delta = \delta(\lambda) \in [0,1]$. A δ-*accurate* **T**-*adapted affine fiber transport* (**T**-AFT) for χ is a pair $\mathsf{AFT}_{\mathbf{T}} = (\mathsf{Sample}, \mathsf{Eval})$ of uniform classical PPT algorithms with access to the fixed public modulus, dimensions, and quotient encoding. The evaluator Eval is deterministic. Their input/output interfaces are:

- $\mathsf{Sample}(1^\lambda, \mathbf{T}, \chi, \delta) \to (\mathbf{Q}, \beta)$: uses only public descriptions and fresh randomness to return $\mathbf{Q} \in \mathbb{Z}^{d\times d}$ and $\beta \in \mathbb{Z}^d$. Its code and worst-case running-time polynomial in λ may depend on the designated accuracy.
- $\mathsf{Eval}(\mathbf{Q}, \beta, \mathbf{x}) \to \mathbf{x}'$: takes the integer matrix and offset and an input $\mathbf{x} \in \mathbb{F}_q^d$, and returns $\mathbf{x}' \in \mathbb{F}_q^d$.

The algorithms satisfy:

Affine evaluation. For all $\mathbf{Q} \in \mathbb{Z}^{d\times d}$, $\beta \in \mathbb{Z}^d$, and $\mathbf{x} \in \mathbb{F}_q^d$,

$$\mathsf{Eval}(\mathbf{Q}, \beta, \mathbf{x}) = \mathbf{Q}\mathbf{x} + \beta \bmod q \in \mathbb{F}_q^d.$$

Invertibility. Every output of Sample satisfies $\mathbf{Q} \bmod q \in \mathrm{GL}_d(\mathbb{F}_q)$.

Fiber transport. For independently sampled

$$(\mathbf{Q}, \beta) \hookleftarrow \mathsf{Sample}(1^\lambda, \mathbf{T}, \chi, \delta), \qquad \mathbf{e} \hookleftarrow \chi,$$

$$\Delta(\mathcal{L}(\mathsf{Eval}(\mathbf{Q}, \beta, \mathbf{e})), \mathsf{U}(\mathcal{F}_{\mathbf{T},0})) \le \delta.$$

The transport is *exact* if $\delta = 0$.

The accuracy is designated before choosing the algorithm; the notation does not require uniform complexity across all accuracies. Exact transport obeys the same worst-case running-time bound. The sampler may equivalently return the matrix and offset modulo q. We retain integer representatives for the construction, while the following LWE experiments operate over $\mathbb{F}_q$.

Lemma 5.2 (Basic properties of **T**-AFT)**.** Let $\mathsf{AFT}_{\mathbf{T}} = (\mathsf{Sample}, \mathsf{Eval})$ be a δ-accurate **T**-AFT for χ. Sample

$$\mathbf{A} \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right), \qquad \mathbf{e} \hookleftarrow \chi,$$
$$(\mathbf{Q}, \beta) \hookleftarrow \mathsf{Sample}(1^\lambda, \mathbf{T}, \chi, \delta)$$

independently, and set $\mathbf{A}' := \mathbf{Q}\mathbf{A}$ and $\mathbf{e}' := \mathsf{Eval}(\mathbf{Q}, \beta, \mathbf{e})$. Then $\mathbf{A}'$ is exactly uniform over $\mathbb{F}_q^{d\times N}$ and independent of $\mathbf{e}'$, and

$$\Delta\left(\mathcal{L}\left(\mathbf{e}'\right), \mathsf{U}\left(\mathcal{F}_{\mathbf{T},0}\right)\right) \le \delta.$$

If $\mathbf{u} \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^d\right)$ is independent of the joint choice of $\mathbf{A}$ and $(\mathbf{Q}, \beta)$, then

$$\mathcal{L}(\mathbf{Q}\mathbf{A}, \mathsf{Eval}(\mathbf{Q}, \beta, \mathbf{u})) = \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right) \times \mathsf{U}\left(\mathbb{F}_q^d\right).$$

Proof. Condition on $(\mathbf{Q}, \beta, \mathbf{e})$. Multiplication by the fixed invertible matrix $\mathbf{Q}$ is a bijection on

$\mathbb{F}_q^{d\times N}$. Thus the conditional law of $\mathbf{A}'$ is uniform and does not depend on the conditioning variables. Averaging shows that $\mathbf{A}'$ is uniform and independent of $(\mathbf{Q}, \beta, \mathbf{e})$, hence also of $\mathbf{e}'$. The error bound is the transport guarantee of Definition 5.1.

For the last property, condition on $(\mathbf{Q}, \beta)$. The two maps $\mathbf{A} \mapsto \mathbf{QA}$ and $\mathbf{u} \mapsto \mathbf{Qu} + \beta$ are bijections. They act separately on independent uniform inputs, so their outputs have the stated product law, independently of the conditioned pair. Averaging proves the claim even when $\mathbf{Q}$ and β are correlated. □

5.2 Gaussian errors and the reduction

Write χ_{G} for the error law of Definition 2.11. A *Gaussian* $\mathbf{T}$-*AFT* is a $\mathbf{T}$-AFT for this source.

Definition 5.3 (Fixed-order Gaussian $\mathbf{T}$-AFT)**.** Fix an efficiently computable public parameter family as in Section 3.1, augmented by a positive Gaussian width $\sigma = \sigma(\lambda)$ included in pp. This fixed family admits *fixed-order Gaussian* $\mathbf{T}$-*AFTs* if, for every fixed integer $c \geq 1$, there is a uniform classical PPT sampler Sample_c satisfying the algebraic requirements of Definition 5.1 at every security parameter, with

$$\Delta(\mathcal{L}(\mathsf{Eval}(\mathbf{Q}, \beta, \mathbf{e}))\,, \mathsf{U}\,(\mathcal{F}_{\mathbf{T},0})) \leq \lambda^{-c}$$

for all $\lambda \geq \lambda_0(c)$, where $(\mathbf{Q}, \beta) \hookleftarrow \mathsf{Sample}_c(1^\lambda, \mathsf{pp})$ and $\mathbf{e} \hookleftarrow \chi_{\mathrm{G}}$ independently. The running-time polynomial and threshold may depend on c. We also call this property *arbitrary fixed inverse-polynomial accuracy.*

Choosing the accuracy order. After fixing a hypothetical uniform QPT distinguisher and polynomial bounds on its block count and inverse advantage, choose one fixed c so that the accumulated transport error is smaller than that advantage. This chooses a sampler Sample_c, not a new public parameter family or Gaussian source. No single sampler taking the order as input is required. The reduction runs at every security parameter without computing $\lambda_0(c)$.

Theorem 5.4 (Gaussian LWE to uniform-zero-fiber LWE)**.** Fix an efficiently computable public parameter family as above. Suppose the family admits fixed-order Gaussian $\mathbf{T}$-AFTs as in Definition 5.3. For every fixed positive integer c, every uniform QPT distinguisher $\mathcal{D}$ on classical sample tables, and every polynomially bounded block count $t = t(\lambda)$, using Sample_c yields a uniform QPT reduction $\mathcal{B}_c$ with

$$\mathsf{Adv}^{\mathsf{UZF\text{-}LWE}}_{\mathcal{D}}(\lambda, t) \leq \mathsf{Adv}^{\mathsf{Gauss\text{-}LWE}}_{\mathcal{B}_c}(\lambda, t) + t\lambda^{-c}$$

for all sufficiently large λ. The reduction uses t input blocks and one call to $\mathcal{D}$.

Since the family admits such transports for every fixed accuracy order, Gaussian decisional LWE hardness against uniform QPT distinguishers on classical inputs for every polynomial block count implies the same hardness for uniform-zero-fiber LWE with this fixed public family.

Proof sketch. Fix c and its sampler Sample_c. For each input block, sample a fresh transport pair independently of the entire input experiment and apply the transformation

$$(\mathbf{A}_i, \mathbf{b}_i) \longmapsto (\mathbf{Q}_i\mathbf{A}_i, \mathbf{Q}_i\mathbf{b}_i + \beta_i)$$

over $\mathbb{F}_q$. On real Gaussian input this preserves the shared secret and replaces each error by its affine transport. Conditional on the shared secret, Lemma 5.2 makes the entire transformed matrix table exactly uniform and independent of the transported errors. The error law in each block, averaged over both the Gaussian source and its independently sampled transport pair, is within λ^{-c} of $\mathsf{U}(\mathcal{F}_{\mathbf{T},0})$. Independence across blocks gives total statistical distance at most $t\lambda^{-c}$ from the real uniform-zero-fiber table, also after averaging over the shared secret.

On ideal input, modular invertibility preserves exactly the product law of independent uniform matrices and responses. The reduction prepares the transformed classical table and invokes $\mathcal{D}$ once. Statistical distance bounds the resulting acceptance gap even for quantum postprocessing, giving the stated inequality. For fixed c, this procedure is uniform QPT at every security parameter; the distance bound holds for all sufficiently large λ as in Definition 5.3. The full game-based proof is deferred to Appendix A.8. □

Composition with the rounding reduction. The two reductions preserve the sample count and the fixed public family. Their quantitative losses combine as follows.

Corollary 5.5 (Gaussian LWE to decisional LAR)**.** Under the hypotheses of Theorem 5.4, for every fixed positive integer c, every polynomially bounded t, and every uniform QPT distinguisher $\mathcal{D}$ on classical sample tables, there is a uniform QPT reduction $\mathcal{B}_c$ such that, for all sufficiently large λ,

$$\mathsf{Adv}^{\mathsf{dLAR}}_{\mathcal{D}}(\lambda, t) \le q\,\mathsf{Adv}^{\mathsf{Gauss\text{-}LWE}}_{\mathcal{B}_c}(\lambda, t) + qt\lambda^{-c} + q^{-N}.$$

The reduction uses t source blocks and invokes $\mathcal{D}$ once. If q is polynomially bounded, q^{-N} is negligible, and Gaussian decisional LWE is hard against uniform QPT distinguishers on classical inputs for every polynomial block count, then decisional LAR is hard against such distinguishers for every polynomial sample count for the same family.

Proof. Apply Theorem 4.2 to $\mathcal{D}$, then apply Theorem 5.4 to the resulting uniform-zero-fiber distinguisher. Substitution yields the advantage bound of Corollary 5.5. Each reduction preserves the number of blocks and makes one call to its downstream distinguisher. Their composition is uniform QPT, with classical preprocessing and one invocation of $\mathcal{D}$ on a classical table.

If the LAR advantage is at least $1/\mathsf{p}(\lambda)$ infinitely often, choose fixed a, b, c_{adv} bounding q, t, p by the corresponding powers of λ. Taking $c > a + b + c_{\mathrm{adv}} + 2$ makes the statistical term at most $1/(4\mathsf{p}(\lambda))$ eventually, while negligibility gives the same bound for q^{-N}. The Gaussian advantage is then at least $1/(2q\mathsf{p}(\lambda))$ infinitely often, contradicting Gaussian LWE hardness. The parameter family is fixed first, followed by the hypothetical distinguisher and its polynomial bounds, then by one fixed accuracy order and its sampler. The reduction runs at every security parameter and need not compute the accuracy threshold. □

For the concrete family, Theorem 9.9 supplies the transport hypothesis. The following sections develop that construction. Both reductions remain uniform classical PPT when their downstream distinguisher is classical PPT.

6 Zero-Fiber Geometry and Affine Matrix Realization

We express affine transport as an integer incidence map and reduce its distributional guarantee to endpoint uniformity and carry safety. A fixed rounding matrix identifies the relevant image lattice;

a tree incidence basis then supplies integral affine maps with invertible linear parts modulo q.

6.1 Graph Interpretation

All vectors are columns, and a matrix acts from its input columns to its output rows. We represent affine transport in the image lattice as integer loads on the vertex set $V := \{1, \dots, d\}$. Rows of an integer matrix correspond to vertices, and each nonzero column represents a signed, integer-weighted hyperedge: positive entries are tail incidences, negative entries are head incidences, and their absolute values are the incidence weights. These columns need not have zero sum, and either the tail or head set may be empty. Appendix B records the full graph conventions.

Definition 6.1 (Incidence representation of affine transport)**.** Fix the rounding matrix $\mathbf{T}$ and a transport pair $(\mathbf{Q}, \beta)$ satisfying the exact algebraic requirements of Definition 5.3. Its *incidence representation* is the pair

$$\mathbf{D_T} := \mathbf{TQ} \in \mathbb{Z}^{d\times d}, \qquad \mathbf{z}_0 := \mathbf{T}\beta \in \mathbb{Z}^d.$$

For an integer representative $\mathbf{x} \in \mathbb{Z}^d$ of the source input, the associated *vertex load* is

$$\mathbf{w} := \mathbf{D_T x} + \mathbf{z}_0 = \mathbf{T}(\mathbf{Qx} + \beta).$$

Each source coordinate gives the coefficient of one hyperedge, and $\mathbf{z}_0$ is the initial vertex load. For the construction, we choose a tree incidence basis, mix the source coordinates before applying it, and append a charge-injection column. The image-lattice constraints ensure that solving these incidence equations gives integral affine data. For the fixed rounding matrix below, Corollary 6.8 reduces these constraints to total charge and the mixer determinant.

Ordinary directed edges supply charge-neutral updates: their incidence columns have $+1$ at the tail and -1 at the head, so each column sums to zero. Let $\mathbf{1} \in \mathbb{Z}^d$ denote the all-ones vector, and write A_{d-1} for the zero-sum integer lattice of Definition 2.2 in dimension d. The following standard fact lets us parametrize every such update using a tree.

Lemma 6.2 (Tree incidence basis)**.** Let $\mathscr{T}$ be a tree on d vertices with an arbitrary orientation. It has $d-1$ edges, and the columns of its signed incidence matrix $\mathbf{D}_{\mathscr{T}}$ form an integral basis of A_{d-1}. In particular,

$$\mathbf{D}_{\mathscr{T}}\mathbb{Z}^{d-1} = A_{d-1}, \qquad \operatorname{rank}_{\mathbb{R}}(\mathbf{D}_{\mathscr{T}}) = d-1.$$

The leaf-elimination proof is given in Appendix B.1.

6.2 Zero-Fiber Geometry

For a load $\mathbf{w} \in \mathcal{L}(\mathbf{T})$, define its canonical carry by

$$\mathbf{k} := \left\lfloor \frac{\mathbf{w}}{q} \right\rfloor \in \mathbb{Z}^d.$$

The load's membership in the image lattice guarantees an integral preimage. Reduction modulo q retains the load's remainders as endpoint colors and discards its integer quotients, which form the carry. The label of the preimage modulo q is $\mathbf{Pk} \bmod q$; the zero fiber requires this label to vanish. The criterion in Lemma 6.6 below permits nonzero carries, which is essential when off-body

loads wrap around. Canonical loads in the box $[0,q)^d$ have zero carry and supply the endpoint parametrization.

We now specialize to a rounding matrix built from a path incidence matrix and a column that injects charge at q vertices. This rounding matrix remains fixed throughout the subsequent tree-incidence transport construction.

Definition 6.3 (Charge distribution and simultaneous injection)**.** Let q be an odd prime and let $d \geq q$ be an integer. Write $\mathbf{e}_1, \ldots, \mathbf{e}_d$ for the standard basis of $\mathbb{Z}^d$. Define

$$\mathbf{D}_d := [\mathbf{e}_1 - \mathbf{e}_2 \mid \cdots \mid \mathbf{e}_{d-1} - \mathbf{e}_d],$$
$$\mathbf{v}_0 := \sum_{j=1}^{q} \mathbf{e}_j, \qquad \mathbf{T} := [\mathbf{D}_d \mid \mathbf{v}_0] \in \mathbb{Z}^{d\times d}.$$

We use the scalar quotient encoding

$$\mathbf{P} := \mathbf{1}^\intercal.$$

The path columns redistribute charge, while the last column injects one unit at each of the first q vertices:

$$\mathbf{1}^\intercal \mathbf{D}_d = \mathbf{0}^\intercal, \qquad \mathbf{1}^\intercal \mathbf{v}_0 = q.$$

The image lattice of $\mathbf{T}$ consists precisely of the integer vectors whose coordinate sum is divisible by q. Equivalently, it is the preimage of the zero-sum hyperplane in $\mathbb{F}_q^d$ under reduction modulo q.

Lemma 6.4 (Image lattice and Smith normal form)**.** Let $\mathbf{T}$ be the matrix from Definition 6.3. Its image lattice is

$$\mathcal{L}(\mathbf{T}) = \left\{ \mathbf{z} \in \mathbb{Z}^d : \sum_{j=1}^{d} z_j \equiv 0 \pmod{q} \right\}.$$

This lattice has index q in $\mathbb{Z}^d$, and the Smith normal form of $\mathbf{T}$ is $\operatorname{diag}(\mathbf{I}_{d-1}, q)$. Consequently, $\mathbf{T}$ is a rounding matrix with $r = d-1$ and $\ell = 1$ in the sense of Definition 2.3.

The proof is in Appendix B.2. The coordinate-sum map identifies the quotient and its scalar encoding.

For this rounding matrix, define the zero-sum hyperplane

$$H_0 := \left\{ \bar{\mathbf{w}} \in \mathbb{F}_q^d : \mathbf{1}^\intercal \bar{\mathbf{w}} = 0 \right\}.$$

By Lemma 6.4, the image lattice $\mathcal{L}(\mathbf{T}) = \mathbf{T}\mathbb{Z}^d$ is precisely the integer preimage of H_0 under reduction modulo q. In particular, canonical lifting identifies H_0 with $\mathcal{L}(\mathbf{T}) \cap [0,q)^d$. The preceding observation therefore defines the *canonical endpoint pullback*

$$c_{\mathbf{T}} \colon H_0 \longrightarrow \mathcal{F}_{\mathbf{T},0},$$
$$\bar{\mathbf{w}} \longmapsto \mathbf{T}^{-1}[\bar{\mathbf{w}}]_q \bmod q.$$

Here $\mathbf{T}^{-1}$ is taken over $\mathbb{Q}$ before reduction modulo q; its product with $[\bar{\mathbf{w}}]_q \in \mathcal{L}(\mathbf{T})$ is integral.

Lemma 6.5 (Canonical parametrization of the zero fiber)**.** The map $c_{\mathbf{T}}$ is a bijection from H_0 onto $\mathcal{F}_{\mathbf{T},0}$, with inverse

$$c_{\mathbf{T}}^{-1}(\mathbf{x}) = \mathbf{T}[\mathbf{x}]_q \bmod q, \qquad \mathbf{x} \in \mathcal{F}_{\mathbf{T},0}.$$

The proof in Appendix B.3 verifies the two inverse maps. Canonical representatives identify the finite hyperplane with the zero fiber.

The bijection in Lemma 6.5 transfers uniformity and statistical distance, but need not preserve affine structure. The following lemma bounds the distance of an affine output to uniform on $\mathcal{F}_{\mathbf{T},0}$ by its endpoint distance to uniform and the probability that its carry lies outside the image lattice.

Lemma 6.6 (Endpoint and carry bound for affine transport)**.** Let $\mathbf{T}$ be the matrix from Definition 6.3, and let $\mathbf{x} \hookleftarrow \chi_{\mathrm{G}}$. Sample $(\mathbf{Q}, \beta) \in \mathbb{Z}^{d\times d} \times \mathbb{Z}^d$ independently of $\mathbf{x}$, with $\mathbf{Q} \bmod q \in \mathrm{GL}_d(\mathbb{F}_q)$. Use the incidence load of Definition 6.1 at the canonical source lift $[\mathbf{x}]_q$, and its carry as defined in Section 6.2. Define the output and endpoint by

$$\begin{aligned} \mathbf{x}' &:= (\mathbf{Q}[\mathbf{x}]_q + \beta) \bmod q, \\ \bar{\mathbf{w}} &:= \mathbf{w} \bmod q = \mathbf{T}\mathbf{x}' \bmod q \in H_0. \end{aligned}$$

Then

$$\Delta\big(\mathcal{L}(\mathbf{x}')\,, \mathsf{U}\,(\mathcal{F}_{\mathbf{T},0})\big) \le \Delta\big(\mathcal{L}(\bar{\mathbf{w}})\,, \mathsf{U}\,(H_0)\big) + \Pr[\mathbf{k} \notin \mathcal{L}(\mathbf{T})].$$

The distributions and probability are over both the transport data and the Gaussian source. The exact pointwise criterion is

$$\mathbf{x}' = c_{\mathbf{T}}(\bar{\mathbf{w}}) \quad \Longleftrightarrow \quad \mathbf{k} \in \mathcal{L}(\mathbf{T}) \quad \Longleftrightarrow \quad \mathbf{1}^{\mathsf{T}}\mathbf{k} \equiv 0 \pmod q.$$

The proof is in Appendix B.4. The carry measures exactly the discrepancy between an affine output and the canonical pullback of its endpoint. Coupling these two outputs reduces the distributional question to the displayed endpoint and carry terms.

6.3 General Tree-Incidence Transport

For the rest of this paper, fix the rounding matrix $\mathbf{T}$ of Definition 6.3. Let $\mathscr{T}$ be any oriented spanning tree on V. Its signed incidence matrix is an integral basis for charge-neutral motion by Lemma 6.2. The following extension determines the effect of adjoining a column of nonzero total charge, independently of the shape of the tree.

Lemma 6.7 (Determinant of an extended tree incidence matrix)**.** For any oriented tree $\mathscr{T}$ on d vertices with signed incidence matrix $\mathbf{D}_{\mathscr{T}} \in \mathbb{Z}^{d\times(d-1)}$ and any integer column $\mathbf{v} \in \mathbb{Z}^d$,

$$|\det[\mathbf{D}_{\mathscr{T}} \mid \mathbf{v}]| = |\mathbf{1}^{\mathsf{T}}\mathbf{v}|\,.$$

The proof is in Appendix B.5. Appending one charge column measures the index of the resulting lattice in $\mathbb{Z}^d$. This gives the determinant formula independently of the tree shape.

Corollary 6.8 (General tree-incidence transport)**.** Let $\mathscr{T}$ be any oriented spanning tree, let $\mathbf{R}_{\mathrm{edge}} \in \mathbb{Z}^{(d-1)\times(d-1)}$, and let $\mathbf{v} \in \mathbb{Z}^d$ satisfy $\mathbf{1}^\intercal \mathbf{v} = q$. Define the extended incidence basis and source-coordinate mixer by

$$\overline{\mathbf{D}}_{\mathscr{T}} := [\mathbf{D}_{\mathscr{T}} \mid \mathbf{v}], \qquad \overline{\mathbf{R}} := \mathrm{diag}(\mathbf{R}_{\mathrm{edge}}, 1),$$
$$\mathbf{D}_{\mathbf{T}} := \overline{\mathbf{D}}_{\mathscr{T}} \overline{\mathbf{R}} \in \mathbb{Z}^{d\times d}.$$

Let the integer affine baseline and public closure offset together form $\mathbf{z}_0 \in \mathbb{Z}^d$ with $\mathbf{1}^\intercal \mathbf{z}_0 \equiv 0 \pmod q$, and set

$$\mathbf{Q} := \mathbf{T}^{-1}\mathbf{D}_{\mathbf{T}}, \qquad \beta := \mathbf{T}^{-1}\mathbf{z}_0.$$

Then $|\det \overline{\mathbf{D}}_{\mathscr{T}}| = q$, $\mathbf{D}_{\mathbf{T}}\mathbb{Z}^d \subseteq \mathcal{L}(\mathbf{T})$, and $\mathbf{z}_0 \in \mathcal{L}(\mathbf{T})$. Consequently $\mathbf{Q} \in \mathbb{Z}^{d\times d}$ and $\beta \in \mathbb{Z}^d$, with

$$\det \mathbf{D}_{\mathbf{T}} = \pm q \det \mathbf{R}_{\mathrm{edge}},$$
$$\det \mathbf{Q} = \frac{\det \overline{\mathbf{D}}_{\mathscr{T}} \det \overline{\mathbf{R}}}{\det \mathbf{T}} = \pm \det \mathbf{R}_{\mathrm{edge}}.$$

Every integer input $\mathbf{x} \in \mathbb{Z}^d$ gives

$$\mathbf{w} := \overline{\mathbf{D}}_{\mathscr{T}} \overline{\mathbf{R}} \mathbf{x} + \mathbf{z}_0 \in \mathcal{L}(\mathbf{T}), \qquad \mathbf{w} \bmod q \in H_0.$$

If $\det \mathbf{R}_{\mathrm{edge}} \not\equiv 0 \pmod q$, then $\mathbf{Q} \bmod q \in \mathrm{GL}_d(\mathbb{F}_q)$.

The tree columns have total charge zero, while the extra column has total charge q. The extended incidence basis therefore has determinant of absolute value q, matching that of $\mathbf{T}$. Pulling back through $\mathbf{T}^{-1}$ cancels this common factor and leaves the mixer determinant, up to sign. Appendix B.6 verifies integrality and this determinant identity.

The extended incidence basis is singular modulo q. The invertibility conclusion concerns the integral quotient $\mathbf{Q} = \mathbf{T}^{-1}\mathbf{D}_{\mathbf{T}}$: compute its determinant over the integers before reducing modulo q.

6.4 Probabilistic Requirements for Fiber Transport

Corollary 6.8 supplies the algebraic part of fiber transport. Fix a public parameter family and, for each fixed positive integer c, an admissible sampler $\mathsf{Sample}_c(1^\lambda, \mathsf{pp})$ satisfying the exact algebraic, source-independence, and computational requirements of Definition 5.3. With the endpoint and carry of Lemma 6.6, it remains to prove nonnegative error bounds

$$\Delta(\mathcal{L}(\bar{\mathbf{w}}), \mathsf{U}(H_0)) \le \varepsilon_{\mathrm{end}},$$
$$\Pr[\mathbf{k} \notin \mathcal{L}(\mathbf{T})] \le \varepsilon_{\mathrm{carry}},$$
$$\varepsilon_{\mathrm{end}} + \varepsilon_{\mathrm{carry}} \le \lambda^{-c}.$$

These estimates must hold for every $\lambda \ge \lambda_0(c)$, with laws averaged jointly over the public sampler and independent Gaussian source. The error terms may depend on the fixed accuracy order and security parameter. The endpoint-and-carry lemma then gives the fiber-transport accuracy required by Definition 5.3. The next section gives a common finite-dimensional probability calibration principle. The affine construction and joint endpoint analysis in Sections 8 and 9 apply it twice while establishing these endpoint and carry estimates for the fixed rounding matrix and source.

7 Finite-Dimensional Probability Calibration

We choose weights on public choices so that their outputs satisfy specified probability constraints. In the local application (Section 8.5.4), the choices are affine maps and the constraints prescribe color probabilities at every protected input. In the joint application (Section 9.6), the choices are public controls and the constraints prescribe Fourier coefficients after averaging the protected source. In both cases the unknown weight uses only public data. We compare the actual response to a solvable reference response; a small perturbation then permits stable inversion.

Neumann inversion. Let V be a finite-dimensional real normed space, with identity I, and let $\mathcal{A} : V \to V$ be linear. If its defect $\mathcal{K} := I - \mathcal{A}$ satisfies $\|\mathcal{K}\|_{V\to V} \le \kappa < 1$, the Neumann series converges in operator norm and gives

$$\mathcal{A}^{-1} = \sum_{k\ge 0} \mathcal{K}^k, \qquad \|\mathcal{A}^{-1}\|_{V\to V} \le \frac{1}{1-\kappa}.$$

Once an application constructs such a response and bounds its perturbation, stable solvability follows.

Finding such a linear response. Let $(\Omega, \mathfrak{\mu})$ be the reference probability space of public choices. A linear synthesis map $\Psi : V \to L^1(\mathfrak{\mu})$ turns coefficients into weights. A linear map $\mathcal{M}$ records their constraints in a real vector space Y, and a linear map $\mathcal{E} : V \to Y$ specifies the reference response. We seek a weight whose response is the prescribed $\mathcal{E}\tau$.

Factoring the actual response through $\mathcal{E}$ reduces calibration to a coefficient equation. The theorem bounds its correction by the residual of any reference coefficient z_0.

Theorem 7.1 (Stable finite-dimensional calibration)**.** Suppose an explicit linear operator $\mathcal{A} : V \to V$ satisfies

$$\mathcal{M}\Psi = \mathcal{E}\mathcal{A}, \qquad \|\mathcal{A} - I\|_{V\to V} \le \kappa < 1.$$

For every specified $\tau \in V$, there is a unique solution z of $\mathcal{A}z = \tau$, and $w = \Psi z$ satisfies $\mathcal{M}(w) = \mathcal{E}\tau$. For every reference $z_0 \in V$,

$$\|z - z_0\|_V \le \frac{\|\tau - \mathcal{A}z_0\|_V}{1-\kappa}.$$

If the constraints include total mass and the target mass is one, then $\mathbb{E}_{\mathfrak{\mu}} w = 1$.

Uniqueness concerns the full coefficient equation; neither $\mathcal{M}$ nor $\mathcal{E}$ need be invertible. Efficient finite computation additionally requires dimension, precision, and inverse-gap bounds. Proofs are in Section A.9.

Probability interpretation. Weighted expectations $\mathcal{M}(w)_j = \mathbb{E}_{\mathfrak{\mu}}[wh_j]$ give linear constraints when the products are integrable; $h_0 = 1$ records total mass. For local calibration these tests record colors at a fixed protected input. For joint calibration they are conditional expectations of the real and imaginary parts of endpoint characters, averaging the protected source and remaining channel randomness given the public choices. This averaging defines the response; evaluating the synthesized weight requires no hidden input.

Nonnegative weights. Mass one gives a signed law. A pointwise lower bound can establish nonnegativity directly; high-moment bounds instead support clipping to an approximate probability law.

Corollary 7.2 (One-sided pointwise positivity)**.** Under the hypotheses of Theorem 7.1, suppose V is a function space with the supremum norm. The solution coefficient is then itself a function, written $f = z = \mathcal{A}^{-1}\tau$; the public weight is Ψf. We have

$$\inf f \geq \inf \tau - \frac{\kappa \|\tau\|_\infty}{1-\kappa}.$$

In particular, if $\tau = 1 - h$, $\sup h \leq a$ and $\|h\|_\infty \leq b$, then

$$f \geq 1 - a - \frac{\kappa(1+b)}{1-\kappa}.$$

For a support set $S \subseteq \Omega$, masked synthesis means $\Psi f = \mathbf{1}_S f$. This weight vanishes outside S. If $f \geq 0$ on S and the masked weight has calibrated expectation one under μ, it is a probability density relative to μ.

Corollary 7.3 (High moments and bounded clipping)**.** For a real integrable weight W, put $V_{\mathrm{clip}} = \mathrm{clip}_{[0,2]}(W) = \min\{2, \max\{0, W\}\}$ and $\delta_{\mathrm{clip}} = \mathbb{E}_\mu |W - V_{\mathrm{clip}}|$. For every integer $p \geq 1$,

$$0 \leq V_{\mathrm{clip}} \leq 2, \qquad |W - V_{\mathrm{clip}}| \leq |W - 1|^{2p}.$$

Consequently,

$$\delta_{\mathrm{clip}} \leq \mathbb{E}_\mu |W - 1|^{2p}.$$

A small coefficient correction need not keep the public weight pointwise close to one. To control clipping, we instead bound a high moment of the weight correction under the reference law. Suppose $W = \Psi z$ and $\Psi z_0 = 1$. For an integer $p \geq 1$, assume

$$\|\Psi v\|_{L^{2p}(\mu)} \leq C_{2p} \|v\|_V$$

on a subspace V_{corr} containing $z - z_0$. Since $W - 1 = \Psi(z - z_0)$, this bound converts the coefficient estimate in Theorem 7.1 into a moment estimate for the public weight. Applying Corollary 7.3 gives

$$\delta_{\mathrm{clip}} \leq \left(\frac{C_{2p} \|\tau - \mathcal{A} z_0\|_V}{1-\kappa} \right)^{2p}.$$

This bounds the expected change in the public weight caused by clipping, before conditioning and normalization.

8 Chameleon-Lobster Family with Safe Local Sampling

The construction starts from the canonical arrangement in Figure 2. The horizontal path is the body, each leg adds two off-body vertices, and the highlighted queue is a set of body vertices that receive one shared scalar. Solid edges form a tree; the dashed enclosure denotes the additional queue hyperedge.

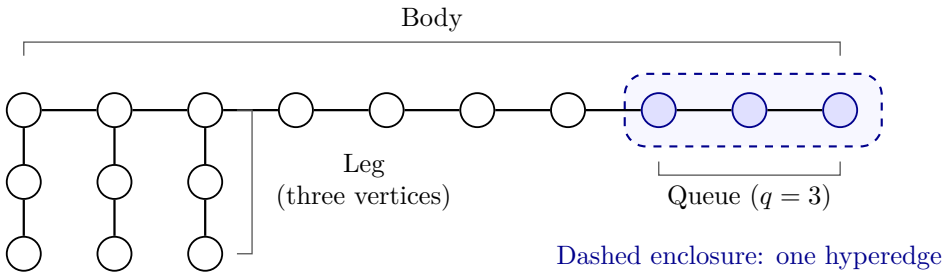


Figure 2: The canonical lobster-tree construction with $q = 3$. Solid edges form the underlying tree: the body is the horizontal path, and each leg adds two vertices to a body vertex. The three highlighted vertices at the tail form the queue; the four remaining body vertices are available for later calibration. The dashed enclosure represents the additional hyperedge supported on the queue.

We specialize the tree-incidence construction of Section 6 to this graph and safe local sampling laws. Given an admissible public-core sampler, Lemma 8.20 supplies integral, source-independent affine data, modular invertibility, and a Gaussian tail bound on carry failure. The joint calibration in Section 9 chooses the core law to make endpoints nearly uniform, completing the fixed-order Gaussian $\mathbf{T}$-AFT guarantee of Definition 5.3. The local and joint calibrations both use Theorem 7.1; the construction here provides their support and source-independence guarantees.

Construction Roadmap and Invariants

Recall the interface of Section 6.4. For $\mathbf{x} \hookleftarrow \mathsf{D}_{\mathbb{Z}^d,\sigma}$, with $\mathbf{e} = \mathbf{x} \bmod q$, sample public affine data $(\mathbf{Q}, \beta)$ independently of the source. We require integrality, invertibility of $\mathbf{Q}$ modulo q, and carry safety on protected inputs. With the endpoint and carry of Lemma 6.6, that lemma gives

$$\Delta(\mathcal{L}(\mathbf{Qe} + \beta \bmod q), \mathsf{U}(\mathcal{F}_{\mathbf{T},0})) \le \Delta(\mathcal{L}(\bar{\mathbf{w}}), \mathsf{U}(H_0)) + \Pr[\mathbf{k} \notin \mathcal{L}(\mathbf{T})].$$

This section bounds the carry term and prepares the endpoint term for Section 9. An arbitrary admissible core law need not make endpoints nearly uniform.

Closing the body path gives all body increments the same marginal law while keeping their sum zero. Each leg adds control over a primary color through a safe public affine map; the dependence of the off-body pair's total load on its second hidden input is canceled by the attachment's increment. Tree increments alone cannot change the total load, so the queue supplies the missing direction by adding one shared scalar to q vertices. The baseline's public charge congruence ensures that the complete offset has an integer preimage under $\mathbf{T}$. The component laws are defined in Definitions 8.2, 8.10 and 8.13 and Section 8.6.

The joint placement fixes disjoint attachments and queue vertices. An independent body permutation moves the cycle, and an independent final permutation relabels the complete load. These operations retain the cycle, leg, and queue dependencies; the endpoint analysis must therefore control their joint law. The affine validity and carry arguments appear in Propositions 8.17 and 8.19, with bounded implementation in Lemma 8.20.

8.1 Geometry and Placement

Let q be an odd prime. Choose the body size $n_{\mathcal{B}}$ and the number g of labeled legs so that

$$d = n_{\mathcal{B}} + 2g, \qquad n_{\mathcal{B}} > q + 2g.$$

The graph has the following components, illustrated in Figure 2.

- The *body* is a path on an ordered vertex set $\mathcal{B}$ of size $n_{\mathcal{B}}$.
- For $1 \le j \le g$, labeled *leg* j is a two-edge path with vertex order $(v_{\mathrm{L}j}, v_{\mathrm{M}j}, v_{\mathrm{R}j})$. The *primary* vertex $v_{\mathrm{L}j}$ and *companion* vertex $v_{\mathrm{M}j}$ lie outside the body; the *attachment* $v_{\mathrm{R}j}$ lies in $\mathcal{B}$. The legs have disjoint off-body vertices and distinct attachments.
- The *queue* is a set $S_{\mathrm{Q}} \subset \mathcal{B}$ of q vertices, disjoint from the attachments. It supports one additional hyperedge that adds the same scalar to every queue vertex.

The body path and the two edges of each leg form a spanning tree with $d-1$ edges. In the canonical arrangement, the attachments are the first g body vertices in labeled order, and the queue consists of the last q body vertices. The strict inequality in Section 8.1 leaves additional body vertices for the calibration controls of Section 9.2.

The following three public permutations place these components and relabel the resulting load.

1. *Jointly place the legs and queue.* A uniform body permutation $\mathbf{P}_{\mathrm{place}}$ places the queue set and labeled attachments on disjoint vertices. The resulting *joint placement* ξ consists of the queue set and the ordered tuple of labeled attachments, retaining only these images of the permutation.
2. *Independently place the body cycle.* The path will be closed by an independent public Gaussian offset at its two endpoints, as defined in Definition 8.2. A uniform body permutation $\mathbf{P}_{\mathrm{B}}$, independent of the placement, initial-load sampling, and Gaussian inputs, moves the path and its closure together. It leaves the attachments and queue fixed.
3. *Relabel the complete load.* An independent uniform permutation $\mathbf{P}_{\mathrm{glob}}$ acts on the entire vertex load, including its initial configuration and all increments.

The calibrated sampler postpones the body and final permutations until the other public sampling steps are complete. Rejection sampling keeps ξ fixed; the complete sampling order and independence conditions are in Section 9.2. The final permutation preserves the image lattice and carry condition because both depend on coordinate sums. Conditional on the endpoint histogram, it makes all arrangements of that histogram equally likely.

8.2 The Affine Realization Matrix

The affine family uses one shared baseline and three increment blocks. The public *core* specifies the body centers, leg shifts, and marks, where each mark is a target primary color and a mixing weight. Given this core, each leg samples a flag and a complete affine map. The same map supplies the baseline intercept and the two mixer coefficients; these quantities must retain their joint law. The deterministic baseline assembly is defined in Definition 8.1, and Section 8.7 specifies the admissible joint sampling laws. The independent public Gaussian closure contributes to the offset.

To evaluate the affine map, first mix the source coordinates and use the incidence columns to place the resulting increments on vertices. Add the sampled baseline and closure offset in these vertex coordinates, relabel the complete load, and apply $\mathbf{T}^{-1}$ to return to the rounding input coordinates. The formulas below separate the linear part from the offset, retaining their shared public randomness.

The placement ξ of Section 8.1 determines the embedding $\mathbf{E}_j : \mathbb{Z}^3 \to \mathbb{Z}^d$ in the leg vertex order fixed above. Write $\mathbf{E}_{\mathrm{B}} : \mathbb{Z}^{n_{\mathcal{B}}} \to \mathbb{Z}^d$ for zero extension from the body. Order the hidden source blocks as

$$\mathbf{x} = \begin{pmatrix} \mathbf{x}_{\mathrm{B}} \\ \mathbf{x}_{\mathrm{L},1} \\ \vdots \\ \mathbf{x}_{\mathrm{L},g} \\ x_{\mathrm{Q}} \end{pmatrix} \hookleftarrow \mathsf{D}_{\mathbb{Z}^d,\sigma}, \qquad \begin{array}{c} \mathbf{x}_{\mathrm{B}} \in \mathbb{Z}^{n_{\mathcal{B}}-1}, \quad \mathbf{x}_{\mathrm{L},j} \in \mathbb{Z}^2, \\ x_{\mathrm{Q}} \in \mathbb{Z}, \qquad d = (n_{\mathcal{B}} - 1) + 2g + 1. \end{array}$$

The fixed local directions, public leg mixers, and queue column are

$$\mathbf{D}_{\mathrm{B}} = [\mathbf{e}_1 - \mathbf{e}_2 \mid \cdots \mid \mathbf{e}_{n_{\mathcal{B}}-1} - \mathbf{e}_{n_{\mathcal{B}}}],$$
$$\mathbf{D}_{\mathrm{L}} = \begin{pmatrix} 1 & 0 \\ -1 & -1 \\ 0 & 1 \end{pmatrix}, \qquad \mathbf{R}_{\mathrm{L},\omega_j} = \begin{pmatrix} \widetilde{c}_j & \widetilde{C}_j \\ 0 & 1 \end{pmatrix}, \quad \mathbf{D}_{\mathrm{Q}} = \sum_{v \in S_{\mathrm{Q}}} \mathbf{e}_v.$$

The mixer entries are the canonical integer lifts of the nonzero scale and shear in Definition 8.7, sampled by the safe law of Section 8.5. The independent body permutation acts only on the path and closure; the closure Gaussian of Definition 8.2 is independent of the source and core. Before the independent final permutation $\mathbf{P}_{\mathrm{glob}}$, define

$$\begin{aligned} \mathbf{D}_\xi(\mathbf{P}_{\mathrm{B}}) &:= [\mathbf{E}_{\mathrm{B}}\mathbf{P}_{\mathrm{B}}\mathbf{D}_{\mathrm{B}} \mid \mathbf{E}_1\mathbf{D}_{\mathrm{L}} \mid \cdots \mid \mathbf{E}_g\mathbf{D}_{\mathrm{L}} \mid \mathbf{D}_{\mathrm{Q}}], \\ \mathbf{R}_{\mathrm{src}} &:= \mathrm{diag}(\mathbf{I}_{n_{\mathcal{B}}-1}, \mathbf{R}_{\mathrm{L},\omega_1}, \ldots, \mathbf{R}_{\mathrm{L},\omega_g}, 1), \\ \mathbf{z}_0^{\mathrm{pre}} &:= \boldsymbol{\Delta}_{\mathrm{BS}}^{\mathrm{pre}} + \mathbf{E}_{\mathrm{B}}\mathbf{P}_{\mathrm{B}}\mathbf{a}_{\mathrm{B}}. \end{aligned}$$

Here $\mathbf{a}_{\mathrm{B}}$ is the two-endpoint closure offset defined in Definition 8.2, and $\boldsymbol{\Delta}_{\mathrm{BS}}^{\mathrm{pre}}$ is the single shared baseline of Definition 8.1, which counts each attachment center once. The affine realization matrix and offset are

$$\mathbf{Q} = \mathbf{T}^{-1}\mathbf{P}_{\mathrm{glob}}\mathbf{D}_\xi(\mathbf{P}_{\mathrm{B}})\mathbf{R}_{\mathrm{src}}, \quad \beta = \mathbf{T}^{-1}\mathbf{P}_{\mathrm{glob}}\mathbf{z}_0^{\mathrm{pre}}.$$

All inverse products with $\mathbf{T}$ are computed over $\mathbb{Q}$ before reduction modulo q. For an admissible core and safe leg maps, Proposition 8.19 gives integral affine data: the public charge condition makes the offset integral, and the nonzero leg scales give modular invertibility.

Public calibration. The Gaussian-convolution viewpoint of [MP12, DPS23] motivates choosing public randomness to shape the output distribution. We use Section 7 twice. Locally, restricting affine maps to safe support biases their color marginals, and marked maps add unwanted bias away from the protected line. The calibration in Section 8.5.4 reweights one public distribution of complete maps to give every protected input its prescribed marginal simultaneously. Its output is a positive map law sampled without observing the source.

The joint calibration in Section 9 addresses the correlations that remain between endpoints. At each fixed placement, it reweights public controls in the protected experiment before imposing charge. After independent uniform endpoint relabeling, the resulting signed law has no nonzero Fourier coefficients up to the chosen support order. A complete Fourier bound controls the remaining distribution error. Moment bounds control the error from clipping the public weight to a bounded nonnegative function; public charge conditioning and normalization then restore an admissible law. Bounded rejection sampling keeps placement fixed and uses only public data. Coupling the protected inputs back to the original Gaussian source gives the AFT bound.

8.3 Safe Baseline Support

Use vertex labels before the final relabeling. Denote the placed increments by

$$\begin{aligned}\boldsymbol{\Delta}_{\mathrm{B}} &:= \mathbf{E}_{\mathrm{B}}\mathbf{P}_{\mathrm{B}}\boldsymbol{\Delta}_{\mathrm{B}}^{\mathrm{loc}},\\ \boldsymbol{\Delta}_{\mathrm{L}} &:= \sum_{j=1}^{g}\mathbf{E}_j\mathbf{D}_{\mathrm{L}}\mathbf{R}_{\mathrm{L},\omega_j}\mathbf{x}_{\mathrm{L},j},\\ \boldsymbol{\Delta}_{\mathrm{Q}} &:= \mathbf{D}_{\mathrm{Q}}x_{\mathrm{Q}}.\end{aligned}$$

Here $\boldsymbol{\Delta}_{\mathrm{B}}^{\mathrm{loc}}$ includes the public closure, as defined in Definition 8.2. A superscript loc denotes a local increment; subscripts $\mathrm{B}, \mathrm{L}, \mathrm{Q}, \mathrm{BS}$ refer to the body, leg, queue, and shared baseline.

Fix positive integer cutoffs K_{B}, K_{L}, and K_{Q}, and an integer margin K_{BS} with

$$K_{\mathrm{B}} + K_{\mathrm{L}} + K_{\mathrm{Q}} \le K_{\mathrm{BS}}, \qquad 2K_{\mathrm{BS}} \le q - 1.$$

The body baseline lies in

$$\mathcal{I}_{\mathrm{safe}} := \{K_{\mathrm{BS}}, \dots, q - 1 - K_{\mathrm{BS}}\}.$$

The centers, shifts, and affine maps determine the baseline below; Section 8.7 specifies their joint sampling law and charge constraint. All labels precede the final relabeling, so we suppress pre on $\boldsymbol{\Delta}_{\mathrm{BS}}^{\mathrm{pre}}$.

Definition 8.1 (Shared baseline vector)**.** Given public body centers b, leg shifts ζ, and one public affine map per leg with canonical intercept a_j as in Definition 8.7, define $\boldsymbol{\Delta}_{\mathrm{BS}}$ by

$$\begin{aligned}(\boldsymbol{\Delta}_{\mathrm{BS}})_v &:= b_v && (v \in \mathcal{B}),\\ (\boldsymbol{\Delta}_{\mathrm{BS}})_{v_{\mathrm{L}j}} &:= a_j && (1 \le j \le g),\\ (\boldsymbol{\Delta}_{\mathrm{BS}})_{v_{\mathrm{M}j}} &:= q - 1 + \zeta_j - a_j && (1 \le j \le g).\end{aligned}$$

Each attachment retains its body entry; the leg supplies no second baseline there. The same sampled map ω_j determines both the mixer $\mathbf{R}_{\mathrm{L},\omega_j}$ and the intercept a_j, preserving their joint law.

Section 8.7 specifies the joint sampling law and public charge condition. On the protected event of Definition 8.16, the margin and safe maps give the carry guarantee of Proposition 8.17.

8.4 Cycle Gaussian Body Injection

The public closure gives the two body endpoints the same Gaussian difference marginal as the internal vertices. It produces a cycle of correlated, charge-neutral increments, illustrated in Figure 3.

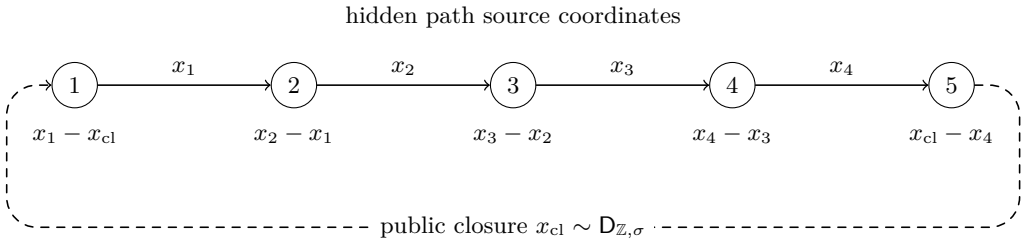


Figure 3: Cycle Gaussian Body Injection for $n_{\mathcal{B}} = 5$. The solid arrows are the four hidden path coordinates. The dashed closure denotes the public offset $x_{\mathrm{cl}}(\mathbf{e}_5 - \mathbf{e}_1)$; it is not a transport-matrix column. Its endpoint contributions produce the exact cyclic difference law. The body placement moves this entire increment together. Baseline, legs, and queue are omitted.

Definition 8.2 (Cycle Gaussian body injection)**.** For $\sigma > 0$, put $\mu_\sigma := \mathsf{D}_{\mathbb{Z},\sigma}$ and use the unit-weight path matrix $\mathbf{D}_{\mathrm{B}}$ from Section 8.2, without mixing. Draw the hidden source $\mathbf{x}_{\mathrm{B}} = (x_1, \ldots, x_{n_{\mathcal{B}}-1})^{\mathsf{T}} \hookleftarrow \mu_\sigma^{\otimes(n_{\mathcal{B}}-1)}$ and independently the public closure $x_{\mathrm{cl}} \hookleftarrow \mu_\sigma$. Set

$$\mathbf{a}_{\mathrm{B}} := x_{\mathrm{cl}}(\mathbf{e}_{n_{\mathcal{B}}} - \mathbf{e}_1) = -x_{\mathrm{cl}}\mathbf{D}_{\mathrm{B}}\mathbf{1}_{n_{\mathcal{B}}-1},$$
$$\boldsymbol{\Delta}_{\mathrm{B}}^{\mathrm{loc}} := \mathbf{D}_{\mathrm{B}}\mathbf{x}_{\mathrm{B}} + \mathbf{a}_{\mathrm{B}}.$$

With $x_{n_{\mathcal{B}}} := x_{\mathrm{cl}}$ and $x_0 := x_{n_{\mathcal{B}}}$, we have

$$(\boldsymbol{\Delta}_{\mathrm{B}}^{\mathrm{loc}})_i = x_i - x_{i-1} \quad (1 \le i \le n_{\mathcal{B}}), \qquad \mathbf{1}^{\mathsf{T}}\boldsymbol{\Delta}_{\mathrm{B}}^{\mathrm{loc}} = 0 \quad \text{over } \mathbb{Z}.$$

Both $\mathbf{a}_{\mathrm{B}}$ and $\boldsymbol{\Delta}_{\mathrm{B}}^{\mathrm{loc}}$ lie in $A_{n_{\mathcal{B}}-1}$ by Lemma 6.2. The placement of Section 8.2 moves the entire cycle, including both closure contributions.

The correlated cycle law averages over the public closure Gaussian; fixing its realized value generally changes the endpoint marginals. Appendix C.1 verifies this law and its coordinate identities. Appendix C.1.1 implements the public scalar with error ν_{G} in worst-case polynomial time, preserving every integer-lattice identity. The protected-edge reference and joint error bounds are in Section 8.8.

8.5 Leg Injection in Tree-Edge Coordinates

Use the leg vertex roles from Section 8.1, suppressing the leg index. The public mixer shapes the primary color while the second hidden input transfers bounded load between the companion and attachment. One public law on affine maps must cancel off-body carries at every protected input without observing either hidden input. For a protected input with nonzero first coordinate, it makes the off-body pair uniform over the allowed interior pairs. When this coordinate vanishes, it permits a prescribed point-mass component while retaining the second input's distribution and hence the attachment marginal. Endpoint assembly uses the joint law of the off-body pair and attachment increment.

8.5.1 Color Pairs and Safe Affine Maps

Use the prime field and canonical lifts of Definitions 2.3 and 2.7.

Definition 8.3 (Color slice)**.** For an integer offset ξ with $|\xi| < q$, define the *color slice* by

$$I_\xi := \{b \in \{0, \dots, q-1\} : q - 1 + \xi - b \in \{0, \dots, q-1\}\}.$$

This is the set of allowed primary colors for integer pair sum $q - 1 + \xi$. The companion of $b \in I_\xi$ is uniquely determined as $q - 1 + \xi - b$.

Definition 8.4 (Interior colors and trimmed slices)**.** For an integer margin $h \geq 0$ with $2h \leq q - 1$, define the *interior colors* and *boundary colors* by

$$C_h := \{h, \dots, q - 1 - h\}, \qquad E_{\partial,h} := \{0, \dots, q-1\} \setminus C_h.$$

For an integer offset ξ with $|\xi| + 2h < q$, define the *trimmed slice* by

$$J_{\xi,h} := \{b \in C_h : q - 1 + \xi - b \in C_h\}.$$

The trimmed slice keeps both colors at distance at least h from 0 and $q - 1$.

Both slices are nonempty under their stated parameter conditions. Appendix C.2 gives their interval forms and cardinalities. When $h = 0$, the boundary set is empty and $J_{\xi,h} = I_\xi$.

Definition 8.5 (Uniform law on trimmed pairs)**.** For ξ and h as in Definition 8.4, define

$$\mathsf{L}^{\mathrm{trunc}}_{\xi,h} := \mathsf{U}\left(\{(b, q - 1 + \xi - b)\}_{b \in J_{\xi,h}}\right).$$

Choosing the primary uniformly from $J_{\xi,h}$ fixes its companion and gives each allowed pair probability $1/|J_{\xi,h}|$.

Definition 8.6 (Protected inputs)**.** For a public cutoff $K_{\mathrm{L}} \in \mathbb{Z}_{>0}$, an integer input (u, t) is *protected* if $|u| < K_{\mathrm{L}}$ and $|t| < K_{\mathrm{L}}$. The protected set and its size are

$$\mathcal{E} := \{-(K_{\mathrm{L}} - 1), \dots, K_{\mathrm{L}} - 1\}^2, \qquad r_{\mathrm{L}} := |\mathcal{E}| = (2K_{\mathrm{L}} - 1)^2.$$

The *protected line* consists of the inputs with $u = 0$, namely $\{0\} \times \{-(K_{\mathrm{L}} - 1), \dots, K_{\mathrm{L}} - 1\}$. Let μ_{L} denote $\mathsf{D}_{\mathbb{Z},\sigma}$ conditioned on $|u| < K_{\mathrm{L}}$.

In affine evaluations, integer inputs are reduced modulo q, while offsets and prescribed pair sums remain integers.

Definition 8.7 (Affine evaluations and public leg maps)**.** Let $\omega = (a_0, c, C) \in \mathbb{F}_q^3$, with intercept a_0, scale c, and shear C. At a protected input, define the *affine evaluation* by

$$\ell_{u,t}(\omega) := [a_0 + cu + Ct]_q \in \{0, \dots, q-1\}.$$

The canonical integer coefficients are

$$a := [a_0]_q, \qquad \widetilde{c} := [c]_q, \qquad \widetilde{C} := [C]_q.$$

A *public leg map* is such a triple with $c \neq 0$.

The evaluation is the primary output color. Since the scale is nonzero, fixing the second field input makes the public map a bijection from the first field input onto $\mathbb{F}_q$.

Definition 8.8 (Safe affine map)**.** Fix a public integer offset ζ satisfying $|\zeta| + K_{\mathrm{L}} - 1 < q$. A public leg map ω is *safe* for ζ and K_{L} if

$$\ell_{u,t}(\omega) \in I_{\zeta - t} \quad \text{for every } (u,t) \in \mathcal{E}.$$

At second input t, the required integer pair sum is $q - 1 + \zeta - t$. The offset bound makes this slice nonempty at every protected input. The distinction between safety and trimming is as follows.

- Safety requires the same map to produce a valid primary and companion color at every protected input simultaneously. Boundary colors are allowed.
- To require both colors to lie in C_h, impose the stronger bound

  $$|\zeta| + K_{\mathrm{L}} - 1 + 2h < q$$

  and replace $I_{\zeta - t}$ in the safety condition by $J_{\zeta - t, h}$.

Appendix C.2 verifies these consequences of the offset bounds and the bijectivity of public leg maps.

8.5.2 Tree-Edge Injection and Carry Safety

Use $\mathbf{D}_{\mathrm{L}}$ from Section 8.2 in the vertex order $(v_{\mathrm{L}}, v_{\mathrm{M}}, v_{\mathrm{R}})$. Fix an integer shift ζ and an integer bound $B_{\mathrm{L}} \in \mathbb{Z}_{\geq 0}$ with $|\zeta| \leq B_{\mathrm{L}}$. For a public leg map ω and the hidden input $\mathbf{x}_{\mathrm{L}} = (u,t)^{\intercal}$, write $\mathbf{R}_{\mathrm{L},\omega}$ for its block in Section 8.2, suppressing the leg index, and set $\mathbf{r}_{\mathrm{L},\omega} := \mathbf{R}_{\mathrm{L},\omega}\mathbf{x}_{\mathrm{L}}$. The first edge redistributes load within the off-body pair, so its coefficients can mix colors without changing the pair sum. The second edge transfers exactly t from the middle vertex to the attachment; leaving this coordinate unmixed bounds the attachment perturbation by K_{L} on protected inputs, even when the first mixed coordinate is large. The mixer is invertible modulo q because $\det \mathbf{R}_{\mathrm{L},\omega} = \widetilde{c} \not\equiv 0 \pmod{q}$. Written in these tree-edge coordinates, the local increment is

$$\boldsymbol{\Delta}_{\mathrm{L}}^{\mathrm{loc}} := \mathbf{D}_{\mathrm{L}}\mathbf{r}_{\mathrm{L},\omega} = \mathbf{D}_{\mathrm{L}}\mathbf{R}_{\mathrm{L},\omega}\mathbf{x}_{\mathrm{L}} = \begin{pmatrix} \widetilde{c}u + \widetilde{C}t \\ -\widetilde{c}u - (\widetilde{C}+1)t \\ t \end{pmatrix}.$$

The vertex-load matrix is therefore the derived product $\mathbf{M}_{\mathrm{L},\omega} := \mathbf{D}_{\mathrm{L}}\mathbf{R}_{\mathrm{L},\omega}$. Restricting the shared baseline of Definition 8.1 and adding this increment gives the complete local load:

$$\boldsymbol{\Delta}_{\mathrm{BS,L}}^{\mathrm{loc}} := (\boldsymbol{\Delta}_{\mathrm{BS}})_{(v_{\mathrm{L}}, v_{\mathrm{M}}, v_{\mathrm{R}})}, \qquad \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \boldsymbol{\Delta}_{\mathrm{BS,L}}^{\mathrm{loc}} + \boldsymbol{\Delta}_{\mathrm{L}}^{\mathrm{loc}}.$$

The intercept $a = [a_0]_q$ and mixer come from the same public map. The attachment center z_{R} is its body entry in $\mathbf{\Delta}_{\mathrm{BS}}$. Thus $\mathbf{\Delta}^{\mathrm{loc}}_{\mathrm{BS,L}}$ is a restriction of the shared baseline, and the leg adds only $\mathbf{\Delta}^{\mathrm{loc}}_{\mathrm{L}}$, with $\mathbf{W}_{\mathrm{L}} = \mathbf{0}$ and $\mathbf{a}_{\mathrm{L}} = \mathbf{0}$.

The off-body pair loses exactly the load transferred to the attachment, so the complete triple has source-independent integer charge. For a safe map, the prescribed pair sum fixes the companion residue from the primary residue, forcing the two off-body carries to cancel. Appendix C.3 gives the charge identities and the precise deterministic statement in Lemma C.1.

At the attachment, only the bounded second input is added to the baseline center. If the body, leg, and queue cutoffs fit within the shared baseline margin, the assembled attachment stays in $[0, q)$ and has zero carry. Lemma C.2 in Appendix C.4 states the exact margin and contribution bounds and verifies this claim.

8.5.3 Target Distribution and Parameters

Choosing the prescribed primary residue. Choose the mixing weight and primary color so that

$$\alpha \in \mathbb{Q} \cap [0, 1/2], \qquad a_* \in \{0, \dots, q-1\}.$$

The required pair sum fixes the companion. Write

$$b_* := q - 1 + \zeta - a_*, \qquad b_* - t \in C_h \quad \text{for every } |t| < K_{\mathrm{L}}.$$

For each fixed t, the prescribed pair is therefore $(a_*, b_* - t)$. The last condition keeps this companion in the interior for every protected second input. Thus only the prescribed primary color can contribute boundary mass. These are the *marked-color admissibility conditions* used below.

Target pair laws. Fix any protected input (u, t) and consider the randomness of the public map. Its primary residue is $a_0 + cu + Ct$ modulo q; the fixed pair sum then determines the companion residue. We want the following two cases.

When $u \neq 0$. The residue pair should be uniform over the trimmed pairs: those with the required sum and both coordinates in C_h.

When $u = 0$. The residue pair should be a mixture of that uniform law and a point mass of weight α at one prescribed pair $(a_*, b_* - t)$ for each fixed t. The primary is chosen in advance; its companion varies with t to preserve the pair sum. The baseline calibration controls this extra mass through the prescribed primary and mixing weight.

One public sampling law must satisfy both cases at every protected input without observing the hidden source. The exact pair law is in Theorem 8.9. Figure 4 illustrates the resulting profiles after averaging over the protected Gaussian inputs. The marked contribution occurs only when the first input vanishes. Its total weight after averaging is

$$\gamma := \alpha \mu_{\mathrm{L}}(0).$$

Feasibility conditions. Use the protected-input count r_{L} and trimming margin h from Section 8.5.1. We impose three sufficient conditions.

(i) *Distinct protected inputs.* Require

$$2(K_{\mathrm{L}} - 1) < q.$$

Then different protected inputs remain different after reduction modulo q.

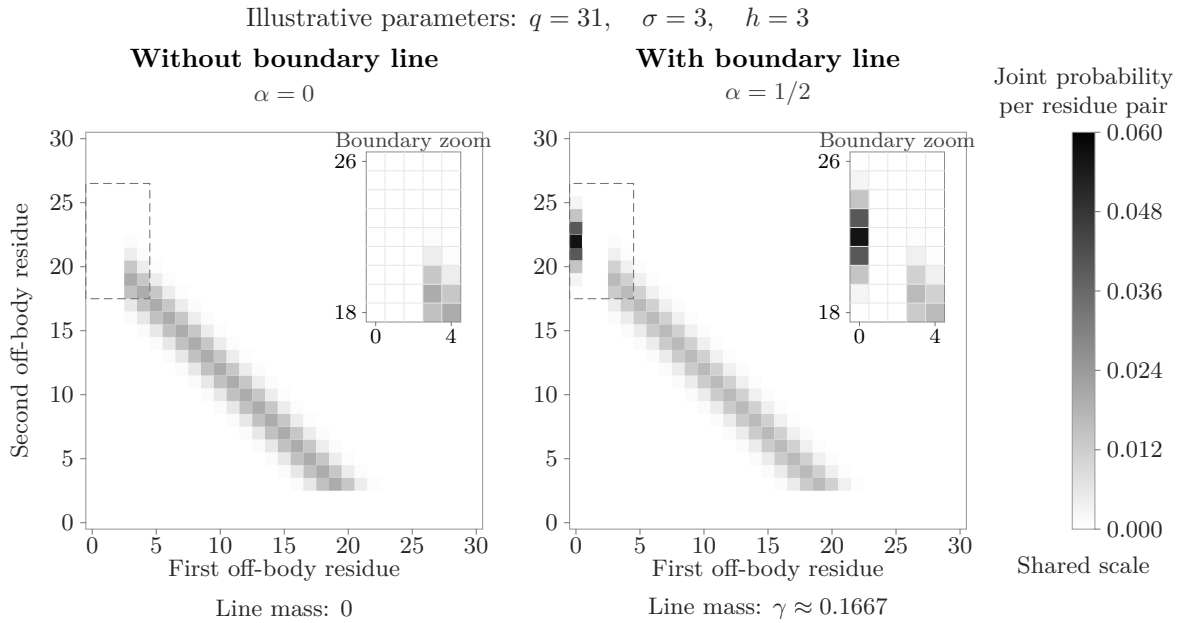


Figure 4: Normalized off-body pair profiles, with the attachment increment marginalized, for illustrative parameters $q = 31$, $\sigma = 3$, $K_{\mathrm{L}} = 5$, $h = 3$, and $\zeta = -8$. These profiles illustrate the target laws; the parameters are not required to satisfy the sampler's sufficient feasibility conditions. Left: $\alpha = 0$. Right: $\alpha = 1/2$ and $a_* = 0$ give the boundary line $(0, 22 - t)$ total mass γ, with $|t| < K_{\mathrm{L}}$; the unmarked component has weight $1 - \gamma$.

(ii) *Room for trimmed pairs.* Require

$$B_{\mathrm{L}} + K_{\mathrm{L}} - 1 + 2h < q.$$

Then every protected input has at least one pair with both residues in the interior.

(iii) *A small enough correction.* Filtering unsafe affine maps generally destroys the exact color marginals. The correction reweights the surviving safe affine maps; it changes neither their coefficients nor the hidden source. The following bounds ensure a strictly positive solution to the simultaneous marginal correction:

$$\begin{aligned}
\delta_{\mathrm{aff}} &:= \frac{r_{\mathrm{L}}(B_{\mathrm{L}} + K_{\mathrm{L}} - 1 + 2h) + 1}{q}, \\
\kappa_{\mathrm{aff}} &:= (2r_{\mathrm{L}} + 1)\delta_{\mathrm{aff}} \le \tfrac{1}{4}, \\
A_{\mathrm{loc}} &:= \frac{r_{\mathrm{L}}}{1 - \delta_{\mathrm{aff}}} \left(\delta_{\mathrm{aff}} + \frac{\kappa_{\mathrm{aff}}}{1 - \kappa_{\mathrm{aff}}} \right) \le \tfrac{1}{3}.
\end{aligned}$$

Here δ_{aff} bounds the deleted fraction of maps, including after fixing an allowed evaluation color, and the lost fraction of marked scales. The quantity κ_{aff} bounds the resulting perturbation of the marginal-response operator in the function supremum norm. Finally, A_{loc} bounds the possible density decrease from cancelling the marked branch's off-line bias after inversion. Appendix C.5 proves these bounds.

We refer to these three requirements as the *feasibility conditions*.

8.5.4 First Application: Safe Affine-Map Calibration

Filtering uniform affine maps to safe supports biases their color marginals; the marked branch adds an off-line bias. We compensate by reweighting one public law on complete maps, enforcing every protected input's target marginal without observing the hidden input. The sampler first chooses a regular or marked branch.

The regular branch keeps only maps whose primary and companion colors are interior at every protected input simultaneously. The marked branch fixes the intercept to a_* and the shear to zero, with nonzero scale. It therefore produces the prescribed primary exactly on the protected line $u = 0$; off that line, it is filtered to keep both colors interior. Marked-color admissibility keeps the companion interior on the line as well.

Choose the marked branch with probability α. Within that branch, sample uniformly from its safe support; on the regular branch, reweight the surviving maps to remove both the filtering bias and the marked branch's off-line bias. The same public law applies simultaneously at all protected inputs and is sampled independently of the hidden source. Write P_{fc} for the complete flag and map law and Π_{iso} for its map marginal. Appendix C.5 specifies its finite model, the supports S, S_{line}, and the branch densities $\mathsf{Q}_{\mathrm{reg}}, \mathsf{Q}_*$ relative to the uniform law on coefficient triples.

Theorem 8.9 (Constructive safe affine-map sampler)**.** Fix public core data satisfying the feasibility and marked-color admissibility conditions of Section 8.5.3. The source-independent law P_{fc} has the following guarantees.

(i) For every protected input (u, t), sampling from P_{fc} gives the residue-pair law

$$\mathcal{L}([X \bmod q]_q, [Y \bmod q]_q) = \begin{cases} (1-\alpha)\mathsf{L}^{\mathrm{trunc}}_{\zeta-t,h} + \alpha\delta_{(a_*,b_*-t)}, & u = 0, \\ \mathsf{L}^{\mathrm{trunc}}_{\zeta-t,h}, & u \neq 0. \end{cases} \tag{1}$$

(ii) Every supported flag and map pair has $c \neq 0$ and is safe at every protected input, with cancelling off-body carries. At each such input, an off-body residue lies in $E_{\partial,h}$ if and only if $B_{\mathrm{mark}} = 1$, $a_* \in E_{\partial,h}$, and $u = 0$. The assembled attachment has zero carry under the additional hypotheses of Lemma C.2.

(iii) The supports S and S_{line} are nonempty, and the calibration matrix $\mathbf{G}_{\mathrm{L}}$ of Section C.5.1 is nonsingular. The branch densities $\mathsf{Q}_{\mathrm{reg}}$ and Q_* are rational, normalized probability densities relative to $\mathsf{U}(\Omega)$. The regular density is at least $1/3$ on S and zero outside it; the marked density is uniform on S_{line}.

(iv) The exact branch tables are computable in time polynomial in the numerical values of q and r_{L} and the bit lengths of the public rational inputs. For every $\varepsilon_{\mathrm{c}} \in \mathbb{Q} \cap (0, 1)$, a worst-case bounded fair-bit sampler approximates P_{fc} within statistical distance ε_{c}, preserves each branch support, and runs in time polynomial in the same parameters and $\log(1/\varepsilon_{\mathrm{c}})$.

These are marginal laws for the same joint public distribution; outputs at distinct protected inputs are not asserted to be independent.

Proof sketch. Pairwise independence of unrestricted affine evaluations gives an exact reference response for additive color corrections. Deleting unsafe maps perturbs this response by at most

$\kappa_{\mathrm{aff}} \le 1/4$. A finite linear solve restores the marginals and cancels the marked branch's off-line bias. The regular target subtracts the marked branch's color imbalances, so their positive part directly threatens nonnegativity. Before excluding zero scale or unsafe colors, each off-line color has exactly one scale preimage; deleting only a small fraction of scales bounds this excess. The inverse-response bound controls the further change caused by filtering. Together, $A_{\mathrm{loc}} \le 1/3$ and $\kappa_{\mathrm{aff}} \le 1/4$ keep the regular density at least $1/3$ on its support. Appendix C.5 gives the complete calibration proof, and Section C.5.1 proves rational table computation and bounded sampling of the complete flag–map law.

8.5.5 Bounded Sampling and Leg Injection

At each fixed public core, Algorithms 1 and 2 in Appendix C.5.1 compute exact rational branch tables and use bounded fair bits to sample within the chosen support. Rounding and fallbacks preserve branch support, carry safety, and the theorem's pointwise boundary-activation criterion. Their effect on activation probabilities is covered by the complete flag–map sampling error. For the concrete public family, $q \le n_{\mathcal{B}} \le d$, so the modulus contributes only polynomial cost in d.

Definition 8.10 (Conditional line-mark leg injection)**.** Under the conditions of Section 8.5.3, the leg's *conditional public-data kernel* is P_{fc} for $(B_{\mathrm{mark}}, \omega)$, given the core data of the shared baseline sampler in Section 8.7. The public draw is independent of $\mathbf{x}_{\mathrm{L}} \sim \mathsf{D}_{\mathbb{Z}^2,\sigma}$. Given these public data, the leg contributes the increment $\boldsymbol{\Delta}_{\mathrm{L}}^{\mathrm{loc}}$ of Section 8.5.2; its complete local load is $\boldsymbol{\Delta}_{\mathrm{BS,L}}^{\mathrm{loc}} + \boldsymbol{\Delta}_{\mathrm{L}}^{\mathrm{loc}}$.

8.5.6 Joint Law for Endpoint Assembly

For analysis, condition both independent Gaussian inputs of Definition 8.10 to be protected. Their marginals become $\mathsf{\mu}_{\mathrm{L}}$, while the independent public kernel remains unchanged. Keep the two off-body residues together with the integer attachment increment to retain the dependence needed for charge cancellation. Section 8.8 bounds the full-source comparison and the bounded sampler's error.

Definition 8.11 (Protected reference experiment)**.** Fix public core data satisfying the feasibility and marked-color admissibility conditions of Section 8.5.3. Draw u, t independently from $\mathsf{\mu}_{\mathrm{L}}$, independently of $(B_{\mathrm{mark}}, \omega) \sim \mathsf{P}_{\mathrm{fc}}$. Define the activation event by

$$\mathsf{A}_{\mathrm{L}} := \{B_{\mathrm{mark}} = 1,\ u = 0\}.$$

Independence of the flag and source inputs gives $\Pr[\mathsf{A}_{\mathrm{L}}] = \gamma$, with the weight defined in Section 8.5.3.

Definition 8.12 (Joint reference law)**.** For the protected reference experiment, define the diffuse

and marked component laws and their mixture by

$$\mathsf{Q}_{\mathrm{L}} := \sum_{|t|<K_{\mathrm{L}}} \mu_{\mathrm{L}}(t)\mathsf{L}^{\mathrm{trunc}}_{\zeta-t,h} \otimes \delta_t,$$
$$\mathsf{Q}_{\mathrm{L},*} := \sum_{|t|<K_{\mathrm{L}}} \mu_{\mathrm{L}}(t)\delta_{(a_*,b_*-t,t)},$$
$$\mathsf{Q}_{\zeta,a_*,\alpha,h} := (1-\gamma)\mathsf{Q}_{\mathrm{L}} + \gamma\mathsf{Q}_{\mathrm{L},*}.$$

Averaging (1) over the two independent protected inputs gives

$$\mathcal{L}([X \bmod q]_q, [Y \bmod q]_q, t) = \mathsf{Q}_{\zeta,a_*,\alpha,h}.$$

The third coordinate is the integer attachment increment, before adding its baseline center. Marked-color admissibility ensures that the prescribed companion in the marked component is already a canonical residue for every protected second input. The two components have the following meaning.

- The marked component describes the activated contribution. The diffuse component describes the complement after averaging both flag branches. In particular, it includes marked-branch outcomes with $u \neq 0$.
- Each component has attachment-increment marginal μ_{L}, so their mixture does as well. Activation depends only on B_{mark} and u and therefore leaves the second input unrestricted.
- The companion's $-t$ and the attachment's $+t$ always use the same input, including on activation. Keeping this input in the joint law preserves their cancellation.

8.6 Queue Injection

The column $\mathbf{D}_{\mathrm{Q}}$ from Section 8.2 adds one scalar to every queue vertex. The queue set is marginally uniform and is independent of the source and body permutation; its dependence on the attachments is fixed by Section 8.1. For a public cutoff $K_{\mathrm{Q}} \in \mathbb{Z}_{>0}$, the analytical reference draws $x_{\mathrm{Q}} \hookleftarrow \mu_{\mathrm{Q}} := \mathsf{D}_{\mathbb{Z},\sigma} \mid \{|x_{\mathrm{Q}}| < K_{\mathrm{Q}}\}$ and sets $\boldsymbol{\Delta}^{\mathrm{ref}}_{\mathrm{Q}} := x_{\mathrm{Q}}\mathbf{D}_{\mathrm{Q}}$. All queue vertices share this bounded increment, whose total charge is $qx_{\mathrm{Q}} \equiv 0 \pmod q$. The actual injection $\boldsymbol{\Delta}_{\mathrm{Q}}$ of Section 8.3 uses the untruncated source without padding. Its truncation error is exactly the single-scalar tail probability in Section 8.8, independently of the queue size.

Conditioning on a placement-dependent core can bias the queue's posterior law, so the joint averaging in Section 8.9 retains that dependence.

8.7 Shared Baseline Sampler: Admissible Joint Distributions

The shared sampler first chooses body centers, leg shifts, and marks. It then draws a complete flag–map pair for each leg and assembles one baseline vector. Assume the geometry and safety margins of Sections 8.1 and 8.3 and the leg feasibility conditions of Section 8.5.3.

8.7.1 Admissible Core Laws

Use the safe-center interval $\mathcal{I}_{\mathrm{safe}}$ from Section 8.3.

Definition 8.13 (Admissible joint baseline law)**.** A public law D_{BS} on body centers b, leg shifts ζ, and marks m is *admissible* if it uses only permitted public parameters, is independent of the hidden Gaussian source and secret-dependent information, and has support satisfying

$$\begin{aligned} b &= (b_v)_{v\in\mathcal{B}} \in \mathcal{I}_{\mathrm{safe}}^{n_{\mathcal{B}}}, \\ \zeta &= (\zeta_1,\dots,\zeta_g) \in ([-B_{\mathrm{L}}, B_{\mathrm{L}}]\cap\mathbb{Z})^g, \\ \sum_{v\in\mathcal{B}} b_v + \sum_{j=1}^{g} \zeta_j &\equiv g \pmod{q}. \end{aligned}$$

For each leg, m records a target and mixing weight satisfying the marked-color admissibility conditions of Section 8.5.3. The common public truncation margin and all leg parameters satisfy the feasibility conditions of Section 8.5.3.

The primary target may be a boundary color. Only its moving companion is required to stay in the interior at every protected second input.

8.7.2 Dependence and Sampling Requirements

Admissibility permits the following dependencies.

- The centers, shifts, and marks may be arbitrarily correlated. In particular, the shifts may be randomized conditionally on the centers.
- The core may depend on preceding public template or local placement data. In that case, D_{BS} denotes the prescribed conditional law given those data. This sampling retains the joint disjoint placement law of Section 8.1. Conditioning on the resulting core may nevertheless bias the posterior law of this initial joint placement, even though the sampling procedure preserves its marginal.
- The public cycle-closing Gaussian and the final common relabeling remain independent of the baseline choices.

We separately require nonempty admissible support and a bounded, efficient public sampler preserving that support, as specified in Appendix D.1.1. Admissibility alone does not provide such a sampler.

The direct-Fourier construction in Section 9 uses a more restricted subclass:

- The joint placement ξ records only the labeled attachments and queue set. The core and rejection tests do not read the separate body-cycle permutation $\mathbf{P}_{\mathrm{B}}$ or the closure Gaussian. These are sampled independently after core acceptance, as specified in Section 9.2.
- The sampler first constructs an uncharged auxiliary proposal, which is not yet admissible, then imposes positive public charge for each fixed placement. All rejection trials retain the same placement, so acceptance does not reweight its marginal law.

These additional restrictions preserve the independence needed for the cycle density estimates. They are not consequences of admissibility.

8.7.3 Joint Public Sampling

Definition 8.14 (Complete baseline data)**.** Draw $(b, \zeta, m) \hookleftarrow \mathsf{D}_{\mathrm{BS}}$. Given this core, draw $(B_{\mathrm{mark}\,j}, \omega_j)$ conditionally independently for $1 \le j \le g$ from the corresponding kernels P_{fc} of Definition 8.10, using fresh public randomness independent of the hidden source. Define P_{BS} as the joint law of

$$(b, \zeta, m, ((B_{\mathrm{mark}\,j}, \omega_j))_{j=1}^{g}),$$

conditional on the same preceding public data when applicable.

The complete law retains every flag and map with its dependence on the core and uses the full two-branch kernel of Theorem 8.9; the regular branch alone generally has nonuniform off-line marginals.

8.7.4 Baseline Vector and Carry Safety

Assemble the complete data of Definition 8.14 using Definition 8.1, retaining the same maps for the baseline intercepts and leg mixers.

Proposition 8.15 (Joint baseline realization)**.** Every admissible core sample and its safe leg maps yield $\mathbf{\Delta}_{\mathrm{BS}} \in \mathcal{L}(\mathbf{T})$, with body entries in $\mathcal{I}_{\mathrm{safe}}$ and off-body entries in $\{0, \dots, q-1\}$. On the joint protected event, the body carries vanish and the off-body carries cancel pairwise. The sampling is source-independent.

Appendix D.1 verifies this assembly. The support and sampling-error guarantees are in Section 8.8.

8.8 Joint Protection and Approximation

Assume the geometry, safe maps, and baseline support of Section 8.1 and Definitions 8.10 and 8.13. The public core and conditional maps are independent of the hidden source; the closure Gaussian is independent of both. The actual sampler uses untruncated Gaussian inputs. We introduce protected inputs only for the analysis.

The protected reference in Section 8.8.3 clips the latent body edge scalars including the public closure, both inputs of every leg, and the single shared queue scalar. It retains the cycle, same-source leg triples, and shared queue correlations. This is the reference used for public calibration in Section 9. For an admissible positive core law, a successful source coupling also certifies carry safety.

The source comparisons below also apply to a positive, uncharged core law with the same local safety support. Integer-lattice validity and the zero-fiber conclusion require the admissible charge constraint.

8.8.1 Protected Sources and Tail Bounds

Definition 8.16 (Joint protected event)**.** Set

$$H_{\mathrm{e}} := \left\lfloor \frac{K_{\mathrm{B}} - 1}{2} \right\rfloor.$$

The event $\mathsf{E}_{\mathrm{safe}}$ requires all of the following:

(i) Every body edge scalar, including the public closure, has magnitude at most H_{e}.

(ii) Both inputs of every leg have magnitude less than K_{L}.

(iii) The queue scalar has magnitude less than K_{Q}.

The latent-scalar bounds protect every cyclic difference and retain the cycle's dependence and zero sum. Each leg contributes two scalar tails; the shared queue contributes one, with no factor q.

The scalar budgets θ_{B}, θ_{L}, and θ_{Q} are defined and bounded in Lemma D.1 in Appendix D.2. Their sum, with one body term, two leg terms per leg, and one queue term, bounds the failure of $\mathsf{E}_{\mathrm{safe}}$. The proof controls latent scalars before taking differences, so no independence of endpoint increments is needed.

8.8.2 Baseline Support and Carry Safety

Under both the ideal baseline-data law P_{BS} and its support-preserving implementation $\widehat{\mathsf{P}}_{\mathrm{BS}}$,

$$\Pr\Big[b \notin \mathcal{I}_{\mathrm{safe}}^{n_{\mathcal{B}}} \text{ or } \boldsymbol{\Delta}_{\mathrm{BS}} \notin \mathcal{L}(\mathbf{T}) \cap \{0,\dots,q-1\}^{d}\Big] = 0.$$

Thus sampling approximation adds no baseline support failure, including on fallback branches. The joint protected event then gives the following carry guarantee.

Proposition 8.17 (Carry safety of the matrix realization)**.** Assume an admissible core, safe leg maps, and $K_{\mathrm{B}} + K_{\mathrm{L}} + K_{\mathrm{Q}} \le K_{\mathrm{BS}}$. For the natural-source load in Section 8.2, $\mathsf{E}_{\mathrm{safe}}$ implies

$$\mathbf{1}^{\intercal}\mathbf{k} = 0, \qquad \mathbf{x}' = c_{\mathbf{T}}(\bar{\mathbf{w}}) \in \mathcal{F}_{\mathbf{T},0}.$$

For any positive source-independent sampling law with this support,

$$\Delta\big(\mathcal{L}(\mathbf{x}')\,,\mathsf{U}\,(\mathcal{F}_{\mathbf{T},0})\big) \le \Delta\big(\mathcal{L}(\bar{\mathbf{w}})\,,\mathsf{U}\,(H_0)\big) + \Pr[{\mathsf{E}_{\mathrm{safe}}}^{c}].$$

The baseline margin keeps body loads canonical, and each off-body pair has cancelling carries. The common relabeling preserves their total. Appendix D.3 proves the endpoint reduction and its invariance under the choice of source lift.

8.8.3 Protected References and Coupling

The fully protected reference uses the leg and queue laws of Definition 8.11 and Section 8.6 and protects the latent body scalars, including the public closure, as follows.

Definition 8.18 (Protected body cycle)**.** Define the protected edge law and its one-vertex residue marginal by

$$\mathsf{\mu}_{\mathrm{e}} := \mathsf{D}_{\mathbb{Z},\sigma} \mid \{|x| \le H_{\mathrm{e}}\}, \qquad \mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} := (\mathsf{\mu}_{\mathrm{e}} * \mathsf{\mu}_{\mathrm{e}}) \bmod q.$$

The *protected body cycle* consists of the cyclic differences of $n_{\mathcal{B}}$ independent samples from $\mathsf{\mu}_{\mathrm{e}}$.

These differences sum to zero and satisfy the body cutoff pointwise. Clipping the latent scalars gives a different law from conditioning the full cycle on safe differences.

Appendix D.2 bounds the whole-cycle coupling budget $\varepsilon_{\mathrm{B,clip}}$ by θ_{B} and proves the component comparisons. The following joint bound is the interface used by the endpoint analysis.

Joint coupling. For an exact positive, source-independent public-core law and exact conditional public sampling, the coupling in Appendix D.2 gives

$$\Pr[\text{coupling failure}] \le \varepsilon_{\mathrm{B,clip}} + 2g\theta_{\mathrm{L}} + \theta_{\mathrm{Q}}.$$

Coupling latent sources before assembly and sharing the complete public data preserves the body cycle, same-source leg triples, and shared queue scalar, with all prescribed public dependencies. On success the complete endpoint vectors agree. For an admissible core, success also implies $\mathsf{E}_{\mathrm{safe}}$ and hence carry safety by Proposition 8.17. After charge conditioning and positive repair, Theorem 9.9 therefore pays this source cost once.

8.8.4 Public Sampling Errors

Couple the implemented core first, then each complete conditional flag–map pair. Uniform conditional error bounds give total error at most $\varepsilon_{\mathrm{BS,joint}} + g\varepsilon_{\mathrm{c}}$ in the complete baseline data; these include the Bernoulli flags and are counted once. The formal statement and coupling proof are in Lemma D.2 in Appendix D.2.

8.9 Gadget Assembly and Realization

Use the source order of Section 8.2, with $\mathbf{x}_{\mathrm{L},j} = (u_j, t_j)^{\intercal}$, and the joint baseline sampler of Section 8.7. We establish affine validity and bounded implementation under these sampling assumptions. Endpoint approximation uses the protected comparison of Section 8.8.3.

8.9.1 Integer Affine Validity

The body path and legs form one spanning tree. The queue column supplies the remaining direction, so the incidence construction applies to the matrix in Section 8.2.

Proposition 8.19 (Integer affine realization)**.** For an admissible public core and its complete safe leg maps, the affine data of Section 8.2 satisfy

$$\mathbf{Q} \in \mathbb{Z}^{d\times d}, \qquad \beta \in \mathbb{Z}^{d}, \qquad \det \mathbf{Q} = \pm \prod_{j=1}^{g} \widetilde{c}_j \not\equiv 0 \pmod{q}.$$

These conclusions hold on every support-preserving fallback.

Proof sketch. The tree columns span the zero-sum integer lattice, and the queue column has total charge q. The general incidence result therefore makes the inverse product integral. The admissible baseline makes the offset integral, and the triangular leg mixers contribute exactly their nonzero scales to the determinant. Appendix D.4 gives the complete verification using Corollary 6.8.

Sampling assumptions and numerical budgets. Assume the geometry, safety margins, and local feasibility conditions above, with rational $\sigma^2 \ge 2$ and nonempty admissible core support. The public template and core samplers must be source-independent, run in worst-case polynomial time in the dimension and public input and precision lengths, and preserve admissibility and polynomial

output length on every branch. Their conditional errors are uniform over the permitted preceding public data and histories. Appendix D.5 specifies the complete contract and separate core, map, closure, and permutation budgets.

The carry-failure budget is

$$\varepsilon_{\mathrm{carry}} := \theta_{\mathrm{B}} + \nu_{\mathrm{G}} + 2g\theta_{\mathrm{L}} + \theta_{\mathrm{Q}}.$$

It bounds failure of the protected event in the actual experiment; support-preserving core, map, and permutation draws add no carry error. The final AFT analysis uses the single source coupling of Section 8.8.3, which covers endpoint disagreement and unsafe carries together, so it does not add this carry budget again.

8.9.2 Bounded Realization

All probabilities below average over the public sampler randomness and the independent Gaussian source.

Lemma 8.20 (Gadget assembly and realization)**.** Under the assumptions of Section D.5, there is a bounded public sampler of affine data $(\mathbf{Q}, \beta)$ with the following guarantees.

(i) It runs in worst-case polynomial time in d, the binary public-input lengths, and the requested precision lengths. Every output is integral and source-independent, with

$$\mathbf{Q} \bmod q \in \mathrm{GL}_d(\mathbb{F}_q),$$

and determinant as in Proposition 8.19.

(ii) For $\mathbf{x}' := \mathbf{Q}\mathbf{e} + \beta \bmod q$, where $\mathbf{e} = \mathbf{x} \bmod q$,

$$\Pr[\mathbf{x}' \notin \mathcal{F}_{\mathbf{T},0}] \le \varepsilon_{\mathrm{carry}}.$$

(iii) Each conditional per-leg flag–map draw is within statistical distance ε_{c} of P_{fc}, uniformly over admissible public histories. Every such output, including rounding and fallback, preserves its branch support and is safe with nonzero scale.

Proof sketch. Each bounded sampler preserves its exact support, so finite-bit errors cannot invalidate the affine data or local safety. Couple its closure to an ideal Gaussian and apply the protected-event carry guarantee; the remaining sampling errors are charged to the public-data coupling. Exact rational assembly then gives the stated matrix and running time. Appendix D.5 supplies the full implementation and error accounting.

The public calibration in Section 9 controls the fully protected joint law before positive repair, source coupling, and numerical implementation.

9 Calibrating Public Randomness for Affine Fiber Transport

Section 8 supplies the affine construction and bounds carry failure. We must still choose a positive, source-independent law for the public randomness so that the joint endpoint law, averaged over both the public choices and the source, is close to uniform on H_0. Balancing individual endpoints

does not suffice: the protected body cycle, same-source leg triples, and shared queue scalar create dependencies between endpoints.

We begin with a positive proposal that balances the expected endpoint histogram, temporarily leaving the public charge constraint aside. At each fixed placement, the proposal gives independent groups of public controls that we can reweight to correct correlations. The response bounds below place the calibration operator near the identity, so Theorem 7.1 yields a stable finite correction. After independent uniform final relabeling, this correction cancels all nonzero Fourier coefficients with support size at most the chosen fixed calibration order, but it may assign negative weights. We bound the total Fourier energy of the remaining modes to control the endpoint error, and use high moments to clip the weight to a bounded nonnegative function with negligible loss. Conditioning on public charge then returns the endpoints to H_0, amplifying the error by only a factor of order q. Rejection sampling keeps the placement fixed, and a single source coupling gives endpoint agreement and carry safety, completing the transport guarantee of Theorem 9.9.

9.1 The Joint Distribution Target

Once carries are safe, the endpoint-to-fiber correspondence of Section 6.2 reduces the transport problem to uniformity on H_0. Imposing charge now would couple the public controls, so we first work on the full ambient space $\mathbb{F}_q^d$ under the uncharged proposal. The linear calibration may produce a signed endpoint law. We bound its ambient error before restoring nonnegativity and charge in Section 9.8. The parameters remain symbolic as in Section 8.1 until we choose their sizes in Section 9.9.

For a mass-one real signed endpoint law P on $\mathbb{F}_q^d$ and frequency $\mathbf{a} \in \mathbb{F}_q^d$, we define the positive-sign characters and Fourier transform by

$$\chi_{\mathbf{a}}(\mathbf{Y}) = e^{2\pi \mathrm{i}\langle \mathbf{a}, \mathbf{Y}\rangle/q}, \qquad \widehat{\mathsf{P}}(\mathbf{a}) = \mathbb{E}_{\mathsf{P}}\chi_{\mathbf{a}}(\mathbf{Y}).$$

Independent uniform final relabeling makes the uncharged protected endpoint law exchangeable. To control correlations at every support size, we collect the Fourier coefficients in one polynomial. Choosing a nonconstant character from exactly r endpoint factors records a frequency with support size r. The product form will let us bound all degrees together using the local endpoint blocks.

Definition 9.1 (Weighted character-polynomial norm)**.** For a mass-one real signed law P on $\mathbb{F}_q^d$, define its character polynomial in the formal variables $\mathbf{z} = (\mathsf{z}_a)_{a\in\mathbb{F}_q^\times}$ by

$$\mathsf{P}_{\mathrm{char}}(\mathbf{z}) := \mathbb{E}_{\mathsf{P}} \prod_{v=1}^{d} \left(1 + \sum_{a\in\mathbb{F}_q^\times} \mathsf{z}_a \chi_a(\mathbf{Y}_v)\right),$$

using the same character convention for scalar arguments. Its ordered symmetric coefficient tensors T_r are specified by

$$\mathsf{P}_{\mathrm{char}}(\mathbf{z}) = \sum_{r=0}^{d} \sum_{a_1,\dots,a_r\in\mathbb{F}_q^\times} (T_r)_{a_1,\dots,a_r} \prod_{j=1}^{r} \mathsf{z}_{a_j}.$$

The tensor norm is Euclidean on ordered frequency tuples; repeated and distinct entries use this same convention, with no additional factorial in a tensor. In particular, $T_0 = 1$. For an

analysis radius $\rho > 0$, define

$$\mathfrak{F}_\rho(\mathsf{P})^2 := \sum_{r=0}^{d} \frac{r!\rho^{2r}}{d^r} \|T_r\|_2^2.$$

For distinct $a_1, a_2 \in \mathbb{F}_q^\times$, a coefficient of one on $\mathsf{z}_{a_1}^2$ is stored once, as $(T_2)_{a_1,a_1} = 1$. A coefficient of one on $\mathsf{z}_{a_1}\mathsf{z}_{a_2}$ is split equally between the two ordered entries: $(T_2)_{a_1,a_2} = (T_2)_{a_2,a_1} = 1/2$. These contribute 1 and 1/2, respectively, to $\|T_2\|_2^2$; the degree-two norm weight applies the same outer factor 2! to either contribution.

Exchangeability ensures that this polynomial retains the full Fourier information: all choices of the support vertices give the same marginal coefficient tensor. The polynomial sums over $\binom{d}{r}$ such choices, whereas Fourier energy sums their squared magnitudes. The normalization $r!/d^r$ compensates for this counting difference with at most a factor e^r loss, uniformly over all degrees, as proved in Lemma E.2. Thus the norm at radius ρ controls the total Fourier energy with support weight $(\rho/\sqrt{e})^{2r}$. When $\rho > \sqrt{e}$, higher support sizes receive larger weights. After low-support cancellation, a bounded norm therefore forces the entire remaining tail to be small. The constant term contributes exactly one, which explains the subtraction in the following criterion.

Lemma 9.2 (Fourier criterion for the joint distribution target)**.** Let P be an exchangeable real signed law of mass one on $\mathbb{F}_q^d$. Let R be an integer with $0 \le R < d$, and let $\rho > \sqrt{e}$, with $s_{\mathrm{F}} = \rho/\sqrt{e}$. If

$$\widehat{\mathsf{P}}(\mathbf{a}) = 0 \qquad \text{whenever} \qquad 1 \le |\operatorname{supp} \mathbf{a}| \le R,$$

then

$$\begin{aligned} \tfrac{1}{2}\|\mathsf{P} - \mathsf{U}\left(\mathbb{F}_q^d\right)\|_1 &\le \tfrac{1}{2}\left(\sum_{\mathbf{a}\neq 0} |\widehat{\mathsf{P}}(\mathbf{a})|^2\right)^{1/2} \\ &\le \frac{\left(\mathfrak{F}_\rho(\mathsf{P})^2 - 1\right)^{1/2}}{2s_{\mathrm{F}}^{R+1}}. \end{aligned}$$

For a probability law, the left-hand side is its statistical distance from the ambient uniform law.

Proof sketch. Exchangeability relates the character-polynomial norm to the complete weighted Fourier energy. After cancellation, every surviving frequency has support size at least $R+1$, so its weight is at least $s_{\mathrm{F}}^{2(R+1)}$. Parseval and Cauchy–Schwarz then give the distributional bound. The full proof, using Lemmas E.1 and E.2, appears in Appendix E.1.2.

9.2 Protected Cycle Reference and Independent Public Controls

To correct the joint endpoint law, we reweight public controls and measure their effect through the protected channel of Section 8.8.3. The root marginal $\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}}$ determines the averaged profiles used below. The channel itself retains the zero-sum cycle, the same-source leg triples of Definition 8.12, and one shared queue scalar.

Placements and pools. We work at a fixed placement ξ in the geometry of Section 8.1, with labeled attachments $R_j = v_{\mathrm{R}j}$. The attachment set and the remaining free body vertices are

$$A_{\mathrm{att}} = \{R_j\}_{j=1}^{g}, \qquad F_{\mathrm{free}} = \mathcal{B} \setminus (A_{\mathrm{att}} \cup S_{\mathrm{Q}}),$$
$$n_{\mathrm{F}} = n_{\mathcal{B}} - g - q.$$

Each leg supplies one primary control and will need its own compensator to offset the companion response. We order the free vertices by their public body labels and assign the first g, in that order, to the labeled legs as compensators. The remaining free vertices form the ordinary-body pool, so

$$m_{\mathrm{P}} = g, \qquad m_{\mathrm{B}} = n_{\mathrm{F}} - g = n_{\mathcal{B}} - q - 2g.$$

Disjointness couples the queue and attachment tuple. Their joint law comes from a uniform initial permutation, which gives equal numbers of preimages to every permitted queue set and ordered attachment tuple. Calibration at a positive integer order R requires selections of up to that many distinct controls. Such selections are available whenever both pool sizes are at least R. Lemma E.3 in Appendix E.2 gives the precise placement law, pool conditions, and bounded implementation.

Public sampling contract. To keep calibration independent of the source, we restrict the information used by the weight and the charge test. The public core θ contains the centers, primary colors, anchors, and shifts prescribed below. The weight uses only the primary colors, their assigned compensator centers, the ordinary-body centers, and ξ; the charge test may also use all body centers and leg shifts. Neither uses hidden source coordinates, actual roots, the shared queue scalar, closure randomness, body or final permutations, or completed flag–affine-map draws.

Normalization and every rejection trial keep ξ fixed, preserving its marginal. Once a core is accepted, we sample complete flag–affine-map pairs conditional on that core, and independently draw the closure Gaussian, body-cycle permutation, and final global permutation. These latter draws and the hidden sources are independent of ξ. The body and final permutations are separate from the initial permutation that chooses the joint placement.

At each fixed placement, the uncharged proposal has independent primary and body controls. Pairing each primary with a distinct compensator therefore gives the independent groups used to form canonical coefficients. Signed correction and charge conditioning can subsequently correlate these controls.

We estimate the finite response matrices using Theorem E.4. Finite support bounds the likelihood along each selected path, while centering eliminates configurations with an isolated selected site. Under the independent body permutation, the remaining path configurations are rare. This gives the fixed-order bound proved in Appendix E.3. The later estimate of the complete Fourier energy still uses the entire protected cycle.

9.3 An Explicit Positive Uncharged Proposal

Our first goal is to balance the expected endpoint histogram, averaging over the protected sources. Since a primary affects both its own endpoint and its companion, we choose the primary and free-center laws together at each fixed disjoint placement. Throughout this construction, convolutions are cyclic on $\mathbb{F}_q$, and support assertions use canonical integer lifts. We impose the structural relations

$$K_{\mathrm{B}} = K_{\mathrm{L}} = K_{\mathrm{Q}}, \qquad K_{\mathrm{BS}} = 3K_{\mathrm{L}}.$$

The common cutoff and all other parameter sizes remain free until Section 9.9. We use the reference laws and activation factor

$$\mathsf{u}_{\mathrm{safe}} = \mathsf{U}\left(\mathcal{I}_{\mathrm{safe}}\right), \qquad \mathsf{u} = \mathsf{U}\left(\mathbb{F}_q\right), \qquad \gamma_{\mathrm{cal}} = \tfrac{1}{2}\boldsymbol{\mu}_{\mathrm{L}}(0).$$

The queue and leg clipped scalar laws therefore coincide.

Convolution profiles. To express the averaged endpoint profiles, let $\mathbf{C}_{\mathrm{B}}$ and $\mathbf{C}_{\mathrm{L}}$ denote convolution with $\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}}$ and $\mathsf{\mu}_{\mathrm{L}}$, respectively, acting on mass profiles. The two body profiles we need are

$$\mathsf{b}_i = \mathbf{C}_{\mathrm{B}}\mathbf{C}_{\mathrm{L}}^{i}\mathsf{u}_{\mathrm{safe}} \quad (i = 0, 1).$$

Body inputs are extended by zero outside their safe support. Section F.1.1 gives the matrix-entry conventions.

Inward anchors and the independent blur. The companion must remain interior for every protected second input. We obtain its anchor by reflecting the primary color and shifting the result inward. We then add an independent public blur with law $\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}}$ and choose the leg shift to preserve the required pair sum, with mixing weight $\alpha = 1/2$. The blur is a fresh draw from this law and uses no actual root or closure value. We impose the anchor and general leg-feasibility conditions specified in Appendix F.1.1, which also gives the exact anchor kernel, blur assignments, and diffuse profiles. We write $\widetilde{\mathsf{A}}$ for the anchor kernel followed by convolution with $\mathsf{\mu}_{\mathrm{L}}$, as defined there.

Balancing the primary and companion together. The companion contribution and the free-center contribution pass through the same body convolution. We can therefore adjust the free-center law to cancel the companion mean, then choose the primary law to balance the remaining profiles. The resulting target $\mathbf{h}_{\mathrm{base}}$ is obtained from $d\mathsf{u}$ by subtracting the body and diffuse off-body means after this cancellation. Appendix F.1.1 gives its full profile expansion. The compensation factor is

$$\varepsilon_{\mathrm{anc}} = \frac{g\gamma_{\mathrm{cal}}}{n_{\mathrm{F}}}.$$

The two proposal laws are specified by

$$\begin{aligned} g\gamma_{\mathrm{cal}}\mathsf{\pi}_0 &= \mathbf{h}_{\mathrm{base}}, \\ \mathsf{p}_0 &= (1+\varepsilon_{\mathrm{anc}})\mathsf{u}_{\mathrm{safe}} - \varepsilon_{\mathrm{anc}}\widetilde{\mathsf{A}}\mathsf{\pi}_0. \end{aligned} \tag{2}$$

The proposal samples free centers independently from p_0, all other body centers independently from $\mathsf{u}_{\mathrm{safe}}$, and all leg primaries independently from $\mathsf{\pi}_0$. It then samples the anchors and public blurs using the kernel specified in Appendix F.1.1.

Theorem 9.3 (Positive proposal and exact averaged mean)**.** For the public family of Section 9.9 and all sufficiently large L, uniformly over disjoint placements, the proposal laws above are reflection-symmetric probability laws and

$$\begin{gathered} \mathsf{\pi}_0 \geq (2q)^{-1}, \\ \tfrac{1}{2}\mathsf{u}_{\mathrm{safe}} \leq \mathsf{p}_0 \leq (1+\varepsilon_{\mathrm{anc}})\mathsf{u}_{\mathrm{safe}} \leq 4/q \quad \text{on } \mathcal{I}_{\mathrm{safe}}, \end{gathered}$$

The center law vanishes off $\mathcal{I}_{\mathrm{safe}}$. In the auxiliary protected experiment, the expected histogram is exactly $d\mathsf{u}$.

Proof sketch. The inward anchors keep every sampled core on the safe local support. The numerator defining the primary law has mass $g\gamma_{\mathrm{cal}}$, and the profile losses are small enough to leave the stated positive margin. The free-center adjustment cancels the companion mean, so the primary equation balances the remaining histogram. Lemma F.1 gives the general loss conditions and quantitative bounds; its proof also verifies this public family.

This proposal has not yet imposed public charge. Once a positive charge-correct core law has been obtained, conditional completion by Definition 8.10 gives affine validity through Proposition 8.15.

9.4 The First Response

To correct higher correlations, we need to understand how reweighting one primary changes the expected endpoint histogram. Its companion also responds, so we pair the centered primary score with a score on its distinct free compensator. Their common body convolution allows the compensator score to cancel the marked companion response. For the fixed proposal, let $\mathbf{D}_\pi = \mathrm{diag}(\pi_0)$. For a centered primary density score f, write $\mathsf{v}_{\mathrm{in}} = \mathbf{D}_\pi \mathsf{f}$ for its mass-zero profile.

The resulting paired lift is

$$(J_{\mathrm{in}}\mathsf{f})(c,a) = \mathsf{f}(c) - \gamma_{\mathrm{cal}} \frac{(\widetilde{\mathsf{A}}\mathbf{D}_\pi \mathsf{f})(a)}{\mathsf{p}_0(a)}.$$

The primary and compensator inputs have independent laws π_0 and p_0, respectively. The ratio is used only on the safe support.

After companion cancellation, the remaining perturbation comes from the diffuse target's dependence on the marked color. The reflected anchor cancels that color from the pair sum, leaving only the choice of side of the midpoint; at the midpoint the two choices are averaged. Thus there are only two diffuse profiles, as specified in Appendix F.1.1. Using these profiles and the anchor sign, let $\mathbf{e}$ be the mass-coordinate column of $\mathsf{E} = (1-\gamma_{\mathrm{cal}})(\mathsf{v}_{+1} - \mathsf{v}_{-1})$, and let χ be the column with entries $\chi_{\mathrm{sign}}(c)$. For the mass-zero adjustment v_{in}, the scalar $\chi^{\mathsf{T}}\mathsf{v}_{\mathrm{in}}$ is the change in mass above the midpoint minus the change below it. Multiplying this scalar by the fixed profile $\mathbf{e}$ gives the diffuse residual. After division by the source-averaged activation probability, the response is

$$\mathbf{T}_1 = \mathbf{I} + {\gamma_{\mathrm{cal}}}^{-1}\mathbf{e}\chi^{\mathsf{T}}.$$

Lemma 9.4 (Exact and stable first response)**.** For the public family of Section 9.9 and all sufficiently large L, uniformly over disjoint placements, the paired lift of every centered primary score is centered under its independent proposal inputs. Its complete first histogram response, averaging the protected roots and the leg inputs, is $\gamma_{\mathrm{cal}}\mathbf{T}_1\mathsf{v}_{\mathrm{in}}$. The operator $\mathbf{T}_1$ preserves mass-zero profiles. In the norm

$$\|\mathsf{v}_{\mathrm{in}}\|^2_{\pi_0^{-1}} = \sum_c \mathsf{v}_{\mathrm{in}}(c)^2/\pi_0(c),$$

it satisfies

$$\|\mathbf{T}_1 - \mathrm{Id}\| = O(L^{-5}), \qquad \|\mathbf{T}_1^{-1}\| \le 2.$$

Proof sketch. The compensator score cancels the companion response, leaving the primary point-mass contribution and the diffuse sign term. Dividing by the activation factor makes the former the identity. The sign term has rank one and is small in the primary-reference norm, which gives the inverse bound. The attachment has zero first response, although its higher joint responses remain in the leg triple. Lemma F.2 gives the general bound and exact inverse; Section F.1.2 proves the claims.

The joint calibration applies this inverse separately in each primary slot, using the graded coefficient norm of Section 9.5. Lemma F.4 relates this norm control to the point evaluations needed for numerical implementation.

9.5 Hoeffding-Type Coordinates and Response Matrices

To correct joint correlations, we reweight groups of independent public controls at a fixed placement. Fix a positive integer calibration order R, a labeled disjoint placement, and its positive uncharged proposal ν_0 from Theorem 9.3. Each primary is paired with its assigned compensator, while ordinary-body centers form a separate pool; a compensator consumes no ordinary-body slot. As in a Hoeffding decomposition, the degree counts selected groups and each kernel is centered in every selected slot. The paired lift J_{in} of Section 9.4 adds the compensator score to a primary correction, and U averages over distinct indices in the two pools.

The coefficient space $\mathcal{V}_{\le R}$ records these corrections as a scalar part and centered, separately symmetric density kernels on the independent primary and body controls. Definition F.3 in Appendix F.2.1 gives the full definition and the coordinate map coord, whose constant-one vector is $\mathbf{e}_0$. The bound $m_R \le (2q+R)^R$ makes the system polynomial in size at every fixed order. All estimates below use the parameters of Section 9.9 and hold uniformly over placements at fixed R.

Three maps and the calibration equation. Calibration must connect a coefficient vector to both a public weight and its endpoint response. The physical row $\mathbf{\Psi}_R(\theta)$ maps input coefficients to a public weight: it sends $\mathbf{e}_0$ to one and sends coord(f) of positive degree to $(\mathsf{U}J_{\mathrm{in}}\mathsf{f})(\theta)$. On the output side, leading evaluation $\mathbf{E}_R$ converts coefficients to ordered Fourier tensors by Lemma F.5. The full response $\mathbf{A}_R$ acts on the input space and records the protected channel's effect before leading evaluation. For every coefficient vector $\mathbf{z}$, Section F.2.5 gives

$$W_R = \mathbf{\Psi}_R(\theta)\mathbf{z}, \qquad [\mathsf{P}_{\mathrm{char},W_R}]_{\le R} = \mathbf{E}_R\mathbf{A}_R\mathbf{z}.$$

Here $\mathsf{P}_{\mathrm{char},W_R}$ is the unnormalized character polynomial of the protected endpoint experiment weighted by W_R relative to the uncharged proposal ν_0. The response averages the hidden source conditional on public controls, without normalizing the weighted law, and treats the Fourier output as a real vector space. Since leading evaluation may be rectangular and noninjective, the Fourier equations need not determine the input coefficients. Calibration instead solves the full square system.

The background coefficient separates the contribution of unselected controls. Products in the coefficient algebra use disjoint slots; physical multiplication also includes overlapping control indices. Appendix F.2.1 defines the product precisely. Let B_{bg} be the unit-constant background coefficient constructed in Section F.2.5, put $\mathbf{b} = \mathrm{coord}(\mathsf{B}_{\mathrm{bg}})$, and let $\mathbf{M}_{\mathrm{bg}}$ represent multiplication by it. Then

$$\mathbf{A}_R = \mathbf{M}_{\mathrm{bg}}\widetilde{\mathbf{A}}_R = \mathbf{M}_{\mathrm{bg}}(\mathbf{I} + \mathbf{U}_R + \mathbf{Z}_R).$$

After removing the background, $\mathbf{U}_R$ strictly raises degree, while $\mathbf{Z}_R$ records protected-root effects and can mix types without lowering degree. Both annihilate $\mathbf{e}_0$ after the raw root constant column has been absorbed into the background. Each complete expression uses one common queue scalar; we average these expressions before normalization and inversion.

Centering explains the triangular structure. Consider one summand in the distinct-index average, with protected roots fixed. If a selected group emits no output, integrating its centered score makes that summand vanish. Every surviving summand therefore has output degree at least its input degree. Averaging over selections and multiplying by the unit-constant background preserve this grading. The normalized and full responses are consequently block lower triangular, with common diagonal blocks $\mathbf{A}_{rr} = \mathbf{I} + (\mathbf{Z}_R)_{r,r}$. These blocks retain the degree-preserving protected-root effects, as verified in Section F.2.5.

A norm adapted to the controls. The graded weighted Hilbert norm $\|\cdot\|_{\mathrm{simp},\rho}$ of Appendix F.2.1 controls both response inversion and public weight moments. Its population weights match averaging over distinct control groups, while its body-radius weight matches the Fourier estimate. The induced operator norm measures the response perturbation, and canonical centering transfers coefficient control to physical moment bounds. Clipping will use these moments, since coefficient control alone gives no pointwise closeness of the weight to one.

9.6 Second Application: Joint Endpoint Calibration

We enforce low-order cancellation by choosing coefficients whose full response is the constant-one vector. We apply Theorem 7.1 on $\mathcal{V}_{\le R}$, with the graded norm above, synthesis $\mathbf{\Psi}_R$, the unnormalized truncated character response, and target evaluation $\mathbf{E}_R$. The reference coefficient operator is the identity and the actual operator is $\mathbf{A}_R$; both the target and the reference coefficient vector are $\mathbf{e}_0$.

Theorem 9.5 (Finite-order joint endpoint calibration)**.** For the family and radius of Section 9.9 and every fixed positive integer R, the following holds for all sufficiently large L, with a threshold depending only on the fixed order and uniformly over placements. Each labeled disjoint placement admits the unique full coefficient solution

$$\mathbf{A}_R\mathbf{z} = \mathbf{e}_0, \qquad W_R(\theta) = \mathbf{\Psi}_R(\theta)\mathbf{z}, \qquad \mathbf{z} = \mathbf{e}_0 + \mathbf{c}.$$

Its scalar coordinate is one, the correction has scalar coordinate zero, and

$$\|\mathbf{c}\|_{\mathrm{simp},\rho} = O_R(L^{-3}), \qquad \mathbb{E}_{\nu_0} W_R = 1.$$

For the signed, uncharged protected-cycle endpoint law weighted by W_R,

$$[\mathsf{P}_{\mathrm{char},W_R}]_{\le R} = \mathbf{E}_R\mathbf{A}_R\mathbf{z} = 1.$$

After independent uniform final relabeling, every nonzero Fourier coefficient of support at most R vanishes.

The vector $\mathbf{c}$ represents the canonical kernel f_R. The calibrated weight may be signed; bounded clipping and public conditioning recover a positive sampling law in Section 9.8.

Response estimates. Stability requires a quantitative bound on the full response. The background, raised-response, and protected-root estimates of Proposition F.9 and Lemmas F.10 and F.12 give

$$\|\mathbf{b} - \mathbf{e}_0\|_{\mathrm{simp},\rho} + \|\mathbf{U}_R\|_{\mathrm{simp},\rho} + \|\mathbf{Z}_R\|_{\mathrm{simp},\rho} \le C_R L^{-3}.$$

Combining these estimates after the common-queue assembly of Section F.2.5, submultiplicativity gives $\|\mathbf{A}_R - \mathbf{I}\|_{\mathrm{simp},\rho} = O_R(L^{-3})$. The full proof is in Section F.2.6.

Finite block solution. Lower triangularity lets us solve one degree at a time: subtract the response of previously determined coefficients, then invert the current diagonal block. The full perturbation bound makes each diagonal block stably invertible, so the recurrence remains controlled. Appendix F.2.2 gives the explicit recurrence and certified finite computation. The proposal already balances degree one before leading evaluation, so $(\mathbf{b})_1 = (\mathbf{c})_1 = 0$ by Sections F.2.5 and F.2.6. For

$R \geq 2$, the first possible correction therefore lies in the degree-two block, where primary–primary, primary–body, and body–body kernels interact.

9.7 Fourier Tails and Weight Moments

The finite correction does not by itself bound the total energy of the remaining Fourier modes, and its weight may be negative. We therefore bound both that energy and the loss from clipping the weight. Both estimates concern the same uncharged weight at a fixed placement. The Fourier estimate averages the protected endpoint experiment, retaining the cycle, same-source leg inputs, shared queue scalar, and independent final relabeling. The moment estimate averages over the independent public controls under the uncharged proposal and will bound the clipping loss.

Proposition 9.6 (Complete Fourier tails and weight moments)**.** Use the public family of Section 9.9 and the analysis radius $\rho = L^{1/4}$, independent of the calibration and moment orders. For every fixed pair of positive integers R, p, there are positive constants C_R and $C_{R,p}$ such that the following bounds hold for all sufficiently large L, depending only on these fixed orders, uniformly over disjoint placements. Let W_R and P_{W_R} be the uncharged weight and protected endpoint law of Theorem 9.5, including independent uniform final relabeling, and write $\widehat{\mathsf{P}}$ for this law's Fourier transform. Then

$$\begin{aligned} \mathfrak{F}_\rho(\mathsf{P}_{W_R}) &\leq C_R, \\ \sum_{\mathbf{a}\neq 0} |\widehat{\mathsf{P}}(\mathbf{a})|^2 &\leq C_R L^{-(R+1)/2}, \\ \tfrac{1}{2}\|\mathsf{P}_{W_R} - \mathsf{U}\left(\mathbb{F}_q^d\right)\|_1 &\leq C_R L^{-(R+1)/4}, \\ \|W_R - 1\|_{2p} &\leq C_{R,p} L^{-3}. \end{aligned}$$

The moment norm uses the uncharged public proposal. For fixed R, the weight is the same for all moment orders.

Proof sketch. After conditioning on the protected cycle and shared queue scalar, we bound the product of the local character polynomials across all degrees. Exact low-support cancellation then converts this bound into control of the remaining Fourier energy. For the weight moments, canonical centering uses the large pools of independent public controls to turn coefficient bounds into a bound on every fixed moment. The full proof, including the local estimates, is in Sections G and G.5.

9.8 Public Charge and Bounded Positive Sampling

We recover a probability law by clipping the signed weight and then imposing public charge. For a fixed placement ξ, the charge event is

$$\mathcal{E}_\xi = \left\{ \sum_{v\in\mathcal{B}} b_v + \sum_{j=1}^{g} \zeta_j \equiv g \pmod q \right\}.$$

This test reads only the permitted core centers and leg shifts. For every choice of the remaining randomness, it is equivalent to the protected endpoint lying in H_0, as verified in Section H.1. With

$V_R = \mathrm{clip}_{[0,2]}(W_R)$ from Corollary 7.3, the desired public-core law is

$$d\nu_{V,\xi}(\theta) = \frac{V_R(\theta)\mathbf{1}_{\mathcal{E}_\xi}(\theta)}{\mathbb{E}_{\nu_{0,\xi}} V_R \mathbf{1}_{\mathcal{E}_\xi}}\, d\nu_{0,\xi}(\theta).$$

We normalize separately at each placement and keep that placement through every rejection trial. This prevents differences in acceptance rates from biasing the placement marginal. Passing the accepted core through the protected endpoint channel, including final relabeling, gives the probability law $\mathsf{P}_{V,\xi}$. In the estimates below, we suppress placement subscripts and write ν_0 for the uncharged proposal and $\mathcal{E}$ for its charge event.

Theorem 9.7 (Bounded clipping and public conditioning)**.** Fix a disjoint placement and the calibrated real weight $W = W_R$ of Theorem 9.5, with $\mathbb{E}_{\nu_0} W = 1$. Let P be its protected endpoint image, including independent uniform final relabeling, and put $V = \mathrm{clip}_{[0,2]}(W)$. Define

$$\delta_{\mathrm{amb}} = \tfrac{1}{2}\|\mathsf{P} - \mathsf{U}\left(\mathbb{F}_q^d\right)\|_1, \qquad \delta_{\mathrm{clip}} = \mathbb{E}_{\nu_0}|W - V|,$$
$$\epsilon_{\mathrm{ch}} = 2\delta_{\mathrm{amb}} + \delta_{\mathrm{clip}}, \qquad b_{\mathrm{ch}} = \mathbb{E}_{\nu_0} V \mathbf{1}_{\mathcal{E}}.$$

If $\epsilon_{\mathrm{ch}} \le 1/(2q)$, then

$$b_{\mathrm{ch}} \ge 1/(4q), \qquad \Delta(\mathsf{P}_{V,\xi}, \mathsf{U}(H_0)) \le 2q\epsilon_{\mathrm{ch}}.$$

Proof sketch. Passing through the endpoint channel cannot increase the weight's L^1 clipping loss. The ambient uniform target assigns mass $1/q$ to the charge hyperplane, so restricting to that hyperplane and renormalizing amplifies the error by a factor proportional to q. The full proof is in Section H.2.

Corollary 9.8 (Quantitative clipping loss)**.** Use the public family of Section 9.9. For every fixed pair of positive integers R, p, there is a positive constant $C_{R,p}$ such that, for all sufficiently large L, the clipping loss of Theorem 9.7 satisfies

$$\delta_{\mathrm{clip}} \le C_{R,p} L^{-6p},$$

uniformly over disjoint placements. For each fixed calibration order R, both δ_{clip} and $q\delta_{\mathrm{clip}}$ are negligible in L.

Proof sketch. Clipping changes the weight only where its deviation from one is at least one, so any fixed even moment bounds the clipping loss. Every such moment estimate applies to the same calibrated weight; increasing the moment order strengthens the analysis without changing the sampler. The full proof is in Section H.2.

Realizing the bounded positive core. To sample this law, draw a fresh core from the uncharged proposal at the retained placement and accept with probability

$$\Pr[\text{accept} \mid \theta] = \mathbf{1}_{\mathcal{E}_\xi}(\theta)\frac{V_R(\theta)}{2}.$$

The clipped weight lies in $[0, 2]$, giving a fixed rejection envelope. The high-moment bound controls the statistical cost of clipping, while finite-atom estimates control the numerical error in evaluating the acceptance probability.

The ideal acceptance rate is at least $1/(8q)$. For a requested timeout bound $0 < \epsilon_{\mathrm{cap}} < 1$, a cap of $\lceil 8q \log(1/\epsilon_{\mathrm{cap}}) \rceil$ trials gives failure probability at most ϵ_{cap}. The same bounded envelope lets us account for proposal and weight errors before charge normalization, which amplifies their numerical contribution by only $O(q)$. The geometric timeout estimate and normalization bounds are made precise in Lemma H.1, Appendix H.3, with the corresponding retry caps and numerical budgets. After acceptance, we complete the remaining public draws according to Section 9.2.

9.9 Choosing the Public Parameters

The estimates above require large control pools and small response defects. We choose one polynomial-size public family that meets these requirements for every fixed accuracy order. Its parameters are verified in Section H.4, and the resulting transport guarantee is stated in Theorem 9.9.

One public choice.

$$L = \lambda + 2, \quad t = L^8, \quad N = \lambda, \quad \sigma = L,$$
$$q = \text{the least prime at least } 2^{16} t^8,$$
$$n_{\mathcal{B}} = q t^{12}, \quad g = t^{16}, \quad d = n_{\mathcal{B}} + 2g, \quad \ell = 1,$$
$$K_{\mathrm{B}} = K_{\mathrm{L}} = K_{\mathrm{Q}} = L^2, \quad K_{\mathrm{BS}} = 3L^2, \quad a_{\mathrm{anc}} = t^6, \quad B_{\mathrm{L}} = 3a_{\mathrm{anc}}, \quad h = 1.$$

In particular, $q = \Theta(L^{64})$ and $d = \Theta(L^{160})$. For the Fourier analysis, we take $\rho = L^{1/4}$. This ensures $s_{\mathrm{F}} > 1$ for sufficiently large L, as needed to control the uncanceled modes.

The complete Fourier estimates remain valid up to radius 2ρ. These radii serve only the analysis and are independent of the calibration and moment orders. All public sizes are polynomial in λ. The family, including q, d, N, σ, and the rounding map, is independent of accuracy order, sample budget, and adversary. Outside the certified regime, the finite-prefix fallback of Section H.5 preserves affine validity; the stated accuracy holds for all sufficiently large λ.

9.10 Bounded Implementation and Fixed-Order Accuracy

With the public family fixed, we obtain the requested fixed positive accuracy order c by choosing a sufficiently large calibration order and a sufficiently small numerical budget:

$$R \geq 4c + 267, \qquad \epsilon_{\mathrm{num}} \leq \tfrac{1}{4} L^{-(c+2)}.$$

Write $\mathrm{negl}_R(L)$ for the nonnegative negligible bound on charge-amplified clipping and source-protection errors established in Section H.6, uniformly over placements. The calibration order changes the weight but stays fixed as λ grows. Increasing the moment order only strengthens the clipping estimate for that same weight. The numerical budget, in turn, determines implementation precision and may shrink with λ. Thus every fixed accuracy order has its own polynomial-time sampler; the degree of its running-time polynomial may depend on that order, while the public family remains the same.

The sampler draws one disjoint placement and constructs the finite calibration tables for it. It then draws fresh proposal cores at that placement until the clipped public-charge rule of Section 9.8

accepts or the retry cap is reached. After acceptance, it completes the conditional leg maps, independently samples the closure and the two subsequent permutations, and assembles the affine data. If certification fails or trials time out, the prescribed fallbacks still return integral affine data with modular invertibility. Every branch uses public data alone, preserving source independence.

Implementation sketch. Table construction and weight evaluation use rational intervals with certified error bounds. Deterministic work limits and retry caps make the worst-case running time polynomial for each fixed calibration order. The precision is chosen so that charge amplification and all subsequent public draws fit the numerical budget. Algorithm 3 in Appendix H.5 gives the complete procedure, precision allocation, and support-preserving fallbacks.

Theorem 9.9 (Constructive fixed-order Gaussian AFT)**.** For the family in Section 9.9, each fixed positive integer c admits a uniform PPT sampler Sample_c satisfying Definition 5.3. Every execution, including fallback, returns integral source-independent affine data with $\mathbf{Q} \bmod q \in \mathrm{GL}_d(\mathbb{F}_q)$. For the fixed calibration order R chosen above and all sufficiently large λ, its accuracy is bounded by

$$\begin{aligned}&\Delta(\mathcal{L}(\mathbf{Q}\mathbf{e} + \beta \bmod q)\,, \mathsf{U}\,(\mathcal{F}_{\mathbf{T},0}))\\ &\quad \le C_R q L^{-(R+1)/4}\\ &\qquad + \mathrm{negl}_R(L) + \epsilon_{\mathrm{num}}, \qquad \mathbf{e} \hookleftarrow \mathsf{D}_{\mathbb{Z}^d,\sigma}.\end{aligned}$$

The budget ϵ_{num} accounts for errors in drawing the joint placement, numerically sampling the charged core, capping rejection, completing the flag–affine-map pairs, sampling the closure, and drawing the two subsequent independent permutations. The charged-core contribution includes the amplification from normalization.

Proof sketch. The Fourier and clipping bounds first give a positive protected endpoint law close to uniform after charge conditioning. We then apply one source coupling to the underlying positive, source-independent public-core law. On a match, the protected and actual endpoints agree and carries are safe, so the endpoint-to-fiber correspondence gives the affine transport bound with a single source-protection cost. The numerical implementation preserves valid affine data on every branch, and its errors fit the stated budget. The full proof, including the accuracy-order calculation, is in Section H.6.

10 Applications

We give three consequences for the fixed parameter family of Section 9.9, with the rounding matrix and quotient encoding of Definition 6.3 and the transport of Theorem 9.9. All dimensions and moduli are polynomially bounded in λ. The transport is used in the security reduction; the functions evaluated in the applications are deterministic rounding functions.

10.1 Basing LAR on Worst-Case Lattice Problems

The direction of the computational reductions is

$$\left\{\mathsf{GapSVP}_\gamma, \mathsf{SIVP}_\gamma\right\} \le_{\mathrm{quantum}} \text{Gaussian } \mathsf{LWE} \le_{\mathrm{classical}} \text{uniform-zero-fiber } \mathsf{LWE} \le_{\mathrm{classical}} \mathsf{dLAR}.$$

Thus worst-case lattice hardness implies decisional LAR hardness. The classical arrows describe the sample transformations: with a QPT downstream distinguisher, their compositions are QPT on classical sample tables, as in Section 3.1. Here GapSVP_γ is the decision shortest-vector problem with approximation gap γ, and SIVP_γ asks for N independent lattice vectors of length at most γ times the last successive minimum.

The worst-case reduction uses continuous Gaussian errors, whereas our affine transport takes a discrete Gaussian source. Gaussian randomized rounding connects these interfaces while preserving the uniform experiment up to negligible implementation error, as verified in Section I.1.

For this public family, the approximation factor is

$$\gamma = \widetilde{O}\Big(\frac{Nq}{\sigma}\Big) = \widetilde{O}(q).$$

Corollary 10.1 (Worst-case foundation for decisional LAR)**.** For the fixed public family and the approximation factor above, a uniform QPT decisional LAR distinguisher on polynomially many classical random examples with nonnegligible advantage yields uniform quantum polynomial-time algorithms for worst-case GapSVP_γ and SIVP_γ in dimension N.

If the distinguishing advantage is guaranteed only on infinitely many security parameters, the resulting algorithms succeed on those parameters. Consequently, quantum worst-case hardness that excludes such solvers succeeding on infinitely many dimensions implies decisional LAR hardness against uniform QPT distinguishers for every polynomial classical sample count. The proof in Section I.1 chooses the transport accuracy after the adversary's sample and advantage bounds, then composes with the Gaussian rounding interface. This order of choices supports every polynomial sample count within one fixed public family.

10.2 Direct Field-Valued Weak PRFs in Constant-Depth Threshold Circuits

Deterministic rounding as a source of weak pseudorandom functions, including TC^0 evaluation, has a substantial history [BPR12, AKPW13]. Here we obtain a direct field-valued construction whose fixed polynomial parameters support every polynomial random-example budget. Its input law is exactly uniform on the full field-matrix space, and its evaluation is the explicit quotient-rounding map; the affine transport is needed only to prove security.

Let $\mathcal{K} := \mathbb{F}_q^N \setminus \{\mathbf{0}\}$. Sample a key $\mathbf{s} \hookleftarrow \mathsf{U}(\mathcal{K})$ and define

$$F_{\mathbf{s}} \colon \mathbb{F}_q^{d\times N} \longrightarrow \mathbb{F}_q, \qquad F_{\mathbf{s}}(\mathbf{M}) := f_{\mathbf{T}}(\mathbf{Ms}).$$

Excluding the zero key makes each single output exactly uniform for every fixed key under uniform inputs. Security concerns the joint law of examples sharing that key. This marginal property differs from the key-uniformity condition of Ding, Jain, and Komargodski [DJK26, Definition 15], which fixes an input and randomizes the key.

For implementation, use a bounded polynomial-time key sampler supported on $\mathcal{K}$ whose law is within $2^{-\lambda}$ in SD of $\mathsf{U}(\mathcal{K})$. Its sampling error and the effect of repeated inputs are negligible for every polynomial sample count. The joint key-and-input evaluation circuits are polynomial-time constructible. Security is against uniform QPT adversaries receiving any polynomial number of independent classical random examples from one fixed key, in the access model of Section 3.1.

Proposition 10.2 (A field-valued weak PRF in TC^0)**.** Under the worst-case hardness condition of Corollary 10.1, the bounded implementation above has post-quantum weak pseudorandomness against uniform QPT random-example distinguishers receiving polynomially many classical examples. Its joint key-and-input evaluation is in TC^0, and its single-output marginals are exactly uniform for every fixed key under uniform field-matrix inputs.

The evaluation circuit uses a fixed sequence of arithmetic stages: inner products, coordinatewise floors, and the quotient sum. Their polynomially bounded ranges permit parallel threshold comparisons. Section I.2 verifies key sampling, circuit construction, and security against a uniformly random function.

10.3 Noise-Free Learning on a Uniform Finite Grid

The weak PRF of Proposition 10.2 supplies a family of noise-free learning targets: each label is a deterministic function of a uniform input and a fixed unknown key. A learner that predicts fresh field labels with an inverse-polynomial advantage over the uniform-guess baseline $1/q$ yields a decisional LAR distinguisher, also when training is QPT. Thus the same worst-case lattice hardness condition rules out such efficient prediction for these polynomial-size TC^0 functions under uniform inputs in $\mathbb{F}_q^{d\times N}$. The learner may be improper, provided it outputs a classical hypothesis with efficient deterministic classical evaluation, as specified below. This is the standard learning-to-distinguishing argument used for rounding-based targets; see [CGKM22, Theorem 4.1] and the verification in Section I.3.

For the ReLU representation, set $D := dN$ and $\mathcal{X} := \{0,\dots,q-1\}^D$, identify each input matrix with its canonical entries in $\mathcal{X}$, and define

$$g_{\mathbf{s}}(\mathbf{M}) := \frac{[F_{\mathbf{s}}(\mathbf{M})]_q}{q}.$$

Grid coordinates are encoded as binary integers and labels as exact rational numerator and denominator pairs. Network depth counts ReLU hidden layers, excluding the final affine scalar output. For our public family, $q \le d \le D = \mathrm{poly}(\lambda)$, so time polynomial in the learning dimension is also polynomial in the security parameter.

Lemma 10.3 (Two-hidden-layer representation on a finite grid)**.** For every $\mathbf{s} \in \mathcal{K}$, the function $g_{\mathbf{s}}$ on $\mathcal{X}$ is representable by a two-hidden-layer ReLU network of size polynomial in D, with rational weights and biases of polynomially bounded magnitude and bit length.

The first layer implements the coordinatewise floors, whose sum gives an integer score. The second implements reduction modulo q and normalization on the score's bounded integer range. The construction is polynomial-time in D given the key and yields a $[0,1]$-valued extension to all real inputs. Section I.4 gives the complete construction, including the $O(Dq^2)$ unit bound and $O(\log D)$ parameter bit length.

Learning model. A learner is a uniform QPT algorithm whose running time is polynomial in D on every input transcript. It receives independent classical noiseless examples under $\mathsf{U}(\mathcal{X})$ and outputs a classical description of a possibly improper hypothesis $\widehat{g}$. A uniform deterministic classical evaluator must, for every output hypothesis and every grid point, approximate its prediction clipped

to $[0,1]$ by a rational number to additive error $1/1000$ in time polynomial in D. No quantum state is part of the hypothesis or supplied to its evaluator. For a target $g_{\mathbf{s}}$, squared loss means

$$\mathbb{E}_{\mathbf{M}\leftarrow\mathsf{U}(\mathcal{X})}\big(\widehat{g}(\mathbf{M}) - g_{\mathbf{s}}(\mathbf{M})\big)^2,$$

where the input is fresh and independent of training. Success probability is over the training examples and the learner's internal randomness and measurement outcomes.

Corollary 10.4 (Learning hardness on a uniform finite grid)**.** Under the preceding learning model, a learner attaining squared loss at most $1/100$ with success probability at least $2/3$ for every target of Lemma 10.3 yields a uniform QPT decisional LAR distinguisher on classical samples with advantage at least $1/2$ for all sufficiently large λ.

Independent validation separates the learner's small squared loss from the constant variance of an independent uniform label. Section I.3 gives the quantitative argument, including the evaluation error and the zero-secret event.

Together with Corollary 10.1, the corollary rules out such a QPT learner under the stated quantum worst-case lattice hardness condition. In particular, it rules out learning all polynomial-size two-hidden-layer ReLU networks with polynomially bounded rational weights and biases of polynomial bit length under the same classical sampling and hypothesis-evaluation conventions. In contrast, Lemma 3.8 gives exact identification from $O(N)$ classical examples at fixed confidence by computationally inefficient exhaustive search. The hardness concerns efficient learning in the stated classical-example model; it does not cover coherent quantum examples or quantum membership queries.

AI Disclosure

We used ChatGPT in preparing this paper. It assisted with developing and checking the algorithm as well as drafting parts of the text and proofs, reviewing the manuscript for correctness, consistency and presentation, and LaTeX typesetting. The author checked the manuscript's statements and proofs and takes full responsibility for its content.

The Prove2Me project[2] verifies four theorems, nine lemmas, and three corollaries from this paper in Lean 4. Gaussian LWE reduces to uniform-zero-fiber LWE (Theorem 5.4) and then to decisional LAR (Theorem 4.2), with the composed bound in Corollary 5.5. Lean also verifies stable finite-dimensional calibration (Theorem 7.1), fixed-order Gaussian affine fiber transport for the explicit family (Theorem 9.9), and the worst-case lattice reduction (Corollary 10.1).

The algebraic results cover modular rank and image lattices (Lemma 2.4), quotient encoding, equivariance, and balanced fibers (Lemma 3.5), and zero-fiber decomposition (Lemma 4.1). Under affine transport the public matrix stays exactly uniform and independent of the transported error (Lemma 5.2). For the explicit matrix, the proofs give its image lattice and Smith form (Lemma 6.4), the zero-fiber parametrization by zero-sum endpoints (Lemma 6.5), and the endpoint and carry bound on transport error (Lemma 6.6). The integral tree transport (Corollary 6.8) uses the tree incidence basis (Lemma 6.2) and the determinant formula for a one-column extension (Lemma 6.7).

AI generated all project-specific code for the author's earlier Lean formalization of Theorem 9.9 in the counted binary-operation model. The author checked its main statement against the paper, compiled the code, and verified that the proof uses only Lean's standard logical axioms.

[2] `https://prove2.me/missions/Information_Loss_Is_All_You_Need%3A_TC%5E0_Pseudorandomness_from_Worst-Case_Lattices`

References

[AGS21] Srinivasan Arunachalam, Alex Bredariol Grilo, and Aarthi Sundaram. Quantum hardness of learning shallow classical circuits. *SIAM Journal on Computing*, 50(3):972–1013, 2021. 2

[AKPW13] Joël Alwen, Stephan Krenn, Krzysztof Pietrzak, and Daniel Wichs. Learning with rounding, revisited - new reduction, properties and applications. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 57–74, Santa Barbara, CA, USA, August 18–22, 2013. Springer Berlin Heidelberg, Germany. 1, 12, 60

[Bar68] Erwin H. Bareiss. Sylvester's identity and multistep integer-preserving Gaussian elimination. *Mathematics of Computation*, 22(103):565–578, 1968. 85

[BBLM05] Stéphane Boucheron, Olivier Bousquet, Gábor Lugosi, and Pascal Massart. Moment inequalities for functions of independent random variables. *The Annals of Probability*, 33(2):514–560, 2005. 117

[BLMR13] Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic PRFs and their applications. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 410–428, Santa Barbara, CA, USA, August 18–22, 2013. Springer Berlin Heidelberg, Germany. 1

[BP14] Abhishek Banerjee and Chris Peikert. New and improved key-homomorphic pseudorandom functions. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 353–370, Santa Barbara, CA, USA, August 17–21, 2014. Springer Berlin Heidelberg, Germany. 1

[BPR12] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 719–737, Cambridge, UK, April 15–19, 2012. Springer Berlin Heidelberg, Germany. 1, 12, 60

[BV15] Zvika Brakerski and Vinod Vaikuntanathan. Constrained key-homomorphic PRFs from standard lattice assumptions - or: How to secretly embed a circuit in your PRF. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015: 12th Theory of Cryptography Conference, Part II*, volume 9015 of *Lecture Notes in Computer Science*, pages 1–30, Warsaw, Poland, March 23–25, 2015. Springer Berlin Heidelberg, Germany. 1

[CGKM22] Sitan Chen, Aravind Gollakota, Adam Klivans, and Raghu Meka. Hardness of noise-free learning for two-hidden-layer neural networks. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 10709–10724. Curran Associates, Inc., 2022. 2, 12, 61, 134

[Dev86] Luc Devroye. *Non-Uniform Random Variate Generation*. Springer, New York, 1986. 72

[DJK25] Youlong Ding, Aayush Jain, and Ilan Komargodski. A new approach for LPN-based pseudorandom functions: Low-depth and key-homomorphic. In Michal Koucký and Nikhil Bansal, editors, *57th Annual ACM Symposium on Theory of Computing*, pages 1898–1909, Prague, Czechia, June 23–27, 2025. ACM Press. 1

[DJK26] Youlong Ding, Aayush Jain, and Ilan Komargodski. Pseudorandom functions in NC^1 from LWE/LPN/CDH (or: How to build PRFs in NC^1, generically), 2026. arXiv:2608.25213. 1, 12, 60

[DLSS14] Amit Daniely, Nati Linial, and Shai Shalev-Shwartz. From average case complexity to improper learning complexity. In David B. Shmoys, editor, *46th Annual ACM Symposium on Theory of Computing*, pages 441–448, New York, NY, USA, May 31 – June 3, 2014. ACM Press. 12

[DPS23] Julien Devevey, Alain Passelègue, and Damien Stehlé. G+G: A Fiat-Shamir lattice signature based on convolved gaussians. In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology – ASIACRYPT 2023, Part VII*, volume 14444 of *Lecture Notes in Computer Science*, pages 37–64, Guangzhou, China, December 4–8, 2023. Springer, Singapore, Singapore. 34

[DV21] Amit Daniely and Gal Vardi. From local pseudorandom generators to hardness of learning. In Mikhail Belkin and Samory Kpotufe, editors, *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pages 1358–1394. PMLR, 2021. 2

[GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, October 1986. 1

[GLPN93] Giorgio Gallo, Giustino Longo, Stefano Pallottino, and Sang Nguyen. Directed hypergraphs and applications. *Discrete Applied Mathematics*, 42(2–3):177–201, 1993. 75

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th Annual ACM Symposium on Theory of Computing*, pages 197–206, Victoria, BC, Canada, May 17–20, 2008. ACM Press. 16, 79

[HAB02] William Hesse, Eric Allender, and David A. Mix Barrington. Uniform constant-depth threshold circuits for division and iterated multiplication. *Journal of Computer and System Sciences*, 65(4):695–716, 2002. 133

[Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. 134

[JKP18] Tibor Jager, Rafael Kurek, and Jiaxin Pan. Simple and more efficient PRFs with tight security from LWE and matrix-DDH. In Thomas Peyrin and Steven Galbraith, editors, *Advances in Cryptology – ASIACRYPT 2018, Part III*, volume 11274 of *Lecture Notes in Computer Science*, pages 490–518, Brisbane, Queensland, Australia, December 2–6, 2018. Springer, Cham, Switzerland. 1

[Kha95] Michael Kharitonov. Cryptographic lower bounds for learnability of Boolean functions on the uniform distribution. *Journal of Computer and System Sciences*, 50(3):600–610, 1995. 2

[Kim20] Sam Kim. Key-homomorphic pseudorandom functions from LWE with small modulus. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020, Part II*, volume 12106 of *Lecture Notes in Computer Science*, pages 576–607, Zagreb, Croatia, May 10–14, 2020. Springer, Cham, Switzerland. 1

[KS09] Adam R. Klivans and Alexander A. Sherstov. Cryptographic hardness for learning intersections of halfspaces. *Journal of Computer and System Sciences*, 75(1):2–12, 2009. 2

[KV94] Michael Kearns and Leslie G. Valiant. Cryptographic limitations on learning Boolean formulae and finite automata. *Journal of the ACM*, 41(1):67–95, 1994. 2

[LMN93] Nathan Linial, Yishay Mansour, and Noam Nisan. Constant depth circuits, Fourier transform, and learnability. *Journal of the ACM*, 40(3):607–620, 1993. 2

[LP17] David A. Levin and Yuval Peres. *Markov Chains and Mixing Times*. American Mathematical Society, 2 edition, 2017. With contributions by Elizabeth L. Wilmer. 15, 67, 73

[LZZ25] Shuchen Li, Ilias Zadik, and Manolis Zampetakis. On the hardness of learning one hidden layer neural networks. In Gautam Kamath and Po-Ling Loh, editors, *Proceedings of The 36th International Conference on Algorithmic Learning Theory*, volume 272 of *Proceedings of Machine Learning Research*, pages 700–701. PMLR, 2025. 12

[McD89] Colin McDiarmid. On the method of bounded differences. In J. Siemons, editor, *Surveys in Combinatorics, 1989*, volume 141 of *London Mathematical Society Lecture Note Series*, pages 148–188. Cambridge University Press, 1989. 117

[Mon18a] Hart Montgomery. More efficient lattice PRFs from keyed pseudorandom synthesizers. In Debrup Chakraborty and Tetsu Iwata, editors, *Progress in Cryptology - INDOCRYPT 2018: 19th International Conference in Cryptology in India*, volume 11356 of *Lecture Notes in Computer Science*, pages 190–211, New Delhi, India, December 9–12, 2018. Springer, Cham, Switzerland. 1

[Mon18b] Hart Montgomery. A nonstandard variant of learning with rounding with polynomial modulus and unbounded samples. In Tanja Lange and Rainer Steinwandt, editors, *Post-Quantum Cryptography - 9th International Conference, PQCrypto 2018*, pages 312–330, Fort Lauderdale, Florida, United States, April 9–11, 2018. Springer, Cham, Switzerland. 11

[MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 700–718, Cambridge, UK, April 15–19, 2012. Springer Berlin Heidelberg, Germany. 16, 34

[NR04] Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudorandom functions. *Journal of the ACM*, 51(2):231–262, March 2004. 1

[NR23] Parker Newton and Silas Richelson. A lower bound for proving hardness of learning with rounding with polynomial modulus. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part V*, volume 14085 of *Lecture*

Notes in Computer Science, pages 805–835, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland. 1, 12

[Pei10] Chris Peikert. An efficient and parallel Gaussian sampler for lattices. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 80–97, Santa Barbara, CA, USA, August 15–19, 2010. Springer Berlin Heidelberg, Germany. 132

[Pei15] Chris Peikert. A decade of lattice cryptography. Cryptology ePrint Archive, Report 2015/939, 2015. 13, 16, 17

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 84–93, Baltimore, MA, USA, May 22–24, 2005. ACM Press. 4, 17, 132

[Sho94] Peter W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *35th Annual Symposium on Foundations of Computer Science*, pages 124–134, Santa Fe, NM, USA, November 20–22, 1994. IEEE Computer Society Press. 1

[Sho04] Victor Shoup. Sequences of games: a tool for taming complexity in security proofs. Cryptology ePrint Archive, Report 2004/332, 2004. 69, 72, 133

[Sta16] Richard P. Stanley. Smith normal form in combinatorics, 2016. arXiv:1602.00166. 13

[Val84] Leslie G. Valiant. A theory of the learnable. In *16th Annual ACM Symposium on Theory of Computing*, pages 436–445, Washington, DC, USA, 1984. ACM Press. 1

[Zas91] Thomas Zaslavsky. Orientation of signed graphs. *European Journal of Combinatorics*, 12(4):361–375, 1991. 75

A Foundations and the Main Reduction

We complete the main reduction, then verify the lattice, probability, rounding, and calibration facts used in the main text.

A.1 The main reduction

Proof of Theorem 1.1. Fix the public family of Section 9.9 and the rounding matrix of Definition 6.3 before choosing an accuracy order, sample count, or distinguisher. Given a uniform QPT distinguisher with nonnegligible advantage on polynomially many classical examples, choose a fixed accuracy order as in Corollary 5.5. Theorem 9.9 supplies its source-independent uniform PPT sampler. It preserves integral affine validity and modular invertibility at every security parameter, with the approximation guarantee holding eventually. The composition corollary therefore gives a uniform QPT Gaussian LWE distinguisher with nonnegligible advantage. This sample-table composition preserves the block count, uses td scalar samples, and invokes the original distinguisher once after classical preprocessing.

Applying Corollary 10.1 gives the claimed uniform QPT algorithms for worst-case GapSVP_γ and SIVP_γ, with $\gamma = \widetilde{O}(q)$. Its proof in Section I.1 supplies the Gaussian interface and the quantum subroutine implementation. If the distinguishing advantage is guaranteed only infinitely often,

the lattice guarantees hold on the corresponding dimensions. The running-time polynomial and accuracy threshold may depend on the fixed order and adversary bounds; the public family and source law remain unchanged. □

A.2 Smith coordinates and the image lattice

Proof of Lemma 2.4. Fix a rounding matrix $\mathbf{T}$ together with matrices $\mathbf{U}, \mathbf{V}$ as in Definition 2.3. The reductions of $\mathbf{U}$ and $\mathbf{V}$ modulo q are invertible over $\mathbb{F}_q$. Hence the same identity modulo q transforms $\mathbf{T}$ mod q into $\operatorname{diag}(\mathbf{I}_r, \mathbf{0}_{\ell\times\ell})$. Consequently,

$$\operatorname{rank}_{\mathbb{F}_q}(\mathbf{T} \bmod q) = r, \qquad \dim_{\mathbb{F}_q} \ker(\mathbf{T} \bmod q) = \ell.$$

Taking determinants gives $|\det \mathbf{T}| = q^\ell > 0$, so $\mathbf{T}$ has full rank d over $\mathbb{Q}$.

Since $\mathbf{V}\mathbb{Z}^d = \mathbb{Z}^d$, the lattice generated by the columns of $\mathbf{T}$ satisfies

$$\mathbf{U}\mathcal{L}(\mathbf{T}) = \mathbb{Z}^r \times q\mathbb{Z}^\ell.$$

The unimodular map $\mathbf{U}$ preserves both $\mathbb{Z}^d$ and $q\mathbb{Z}^d$, so $q\mathbb{Z}^d \subseteq \mathcal{L}(\mathbf{T}) \subseteq \mathbb{Z}^d$. Its quotient is

$$\mathbb{Z}^d/\mathcal{L}(\mathbf{T}) \cong \mathbb{Z}^d/(\mathbb{Z}^r \times q\mathbb{Z}^\ell) \cong \mathbb{F}_q^\ell.$$

For a fixed choice of $\mathbf{U}$, this isomorphism sends the coset of $\mathbf{z} \in \mathbb{Z}^d$ to the last ℓ coordinates of $\mathbf{U}\mathbf{z}$, reduced modulo q. Thus the quotient has q^ℓ cosets.

If $r = 0$, then $\mathbf{T} \bmod q = \mathbf{0}$, $\mathcal{L}(\mathbf{T}) = q\mathbb{Z}^d$, and the quotient is $\mathbb{F}_q^d$. If $r = d$, then $\mathbf{T}$ is unimodular, $\mathcal{L}(\mathbf{T}) = \mathbb{Z}^d$, and the quotient has one element. For $0 < r < d$, the modular image and the quotient both have positive dimension strictly less than d. The condition $\ell \geq 1$ is exactly the requirement that the quotient contain more than one coset.

□

A.3 Conditioning and statistical distance

Proof of the conditioning identities in Section 2.3. On A, conditioning rescales each probability by $1/\mathsf{P}(A)$, and off the event it sets that probability to zero. Thus

$$\begin{aligned} 2\Delta(\mathsf{P}, \mathsf{P}(\cdot \mid A)) &= \sum_{\omega\in A} \mathsf{P}(\omega)\left(\frac{1}{\mathsf{P}(A)} - 1\right) + \mathsf{P}(A^c) \\ &= 2\mathsf{P}(A^c). \end{aligned}$$

For independent variables restricted separately to coordinate events, both the joint point probability and the probability of the restriction factor. Dividing the former by the latter proves the product conditional law and hence preservation of independence. □

The contraction, coupling, and product bounds in Section 2.3 are standard finite-state inequalities [LP17, Chapter 4]. Their proofs extend to countable spaces by absolute summability.

A.4 Gaussian exponential moments

Proof of Lemma 2.9. Completing the square expresses the expectation as the right-hand exponential times

$$\frac{\rho_{\sigma,\mathbf{0}}(\Lambda - \sigma^2 t\mathbf{u}_{\mathrm{test}}/(2\pi))}{\rho_{\sigma,\mathbf{0}}(\Lambda)}.$$

For rank zero, both sides of the claimed inequality equal one. For positive rank, regard Λ as a full-rank lattice in its Euclidean span. Poisson summation within $\mathsf{Span}_{\mathbb{R}}(\Lambda)$ expresses the translated Gaussian mass as a sum of positive dual Gaussian coefficients multiplied by cosines. Each cosine is at most one, and all are one at zero shift, so the ratio is at most one. □

A.5 Quotient labels and translation equivariance

Proof of Lemma 3.5. **Well-defined quotient encoding.** For integer vectors $\mathbf{x}, \mathbf{z}$, the identity

$$\lfloor \mathbf{T}\mathbf{x}/q + \mathbf{T}\mathbf{z} \rfloor = \lfloor \mathbf{T}\mathbf{x}/q \rfloor + \mathbf{T}\mathbf{z}$$

proves the stated translation formula. By Lemma 2.4, the quotient projection has kernel $\mathcal{L}(\mathbf{T})$ and is surjective. The change $\mathbf{T}\mathbf{z}$ lies in this kernel, proving independence of the integer lift.

Translation equivariance. Since $\mathbf{V}$ is unimodular, its reduction modulo q is invertible, so $\mathbf{J}$ has full column rank. Moreover,

$$\mathbf{T}\mathbf{J}_{\mathrm{int}} = q\mathbf{U}^{-1}\mathbf{E}.$$

Thus $(\mathbf{T} \bmod q)\mathbf{J} = \mathbf{0}$ over $\mathbb{F}_q$. Since the modular kernel has dimension ℓ by Section 2.2, the columns of $\mathbf{J}$ form a basis of that kernel. Applying coordinatewise floor to integer lifts therefore gives

$$\mathsf{AR}_{\mathbf{T}}([\mathbf{x}]_q + \mathbf{J}_{\mathrm{int}}[\mathbf{a}]_q) = \mathsf{AR}_{\mathbf{T}}([\mathbf{x}]_q) + \mathbf{U}^{-1}\mathbf{E}[\mathbf{a}]_q.$$

By the first property, reducing the shifted input to its canonical lift does not change the encoded label. Applying $\mathbf{P}$ modulo q proves the translation-equivariance claim of Lemma 3.5.

Balanced fibers. The image of $\mathbf{J}$ is an ℓ-dimensional subspace of $\mathbb{F}_q^d$. For each fixed $\mathbf{x} \in \mathbb{F}_q^d$, the map $\mathbf{a} \mapsto \mathbf{x} + \mathbf{J}\mathbf{a}$ is a bijection from $\mathbb{F}_q^\ell$ onto the coset $\mathbf{x} + \mathrm{im}(\mathbf{J})$. By translation equivariance in Lemma 3.5, its label is $f_{\mathbf{T}}(\mathbf{x}) + \mathbf{a}$, so each label occurs exactly once on that coset. There are $q^{d-\ell} = q^r$ such cosets, so each label has exactly that many preimages. □

A.6 Information-theoretic identifiability

Proof of Lemma 3.8. Let $\mathbf{M}$ be uniform on $\mathbb{F}_q^{d\times N}$, and fix distinct secrets $\mathbf{s}, \mathbf{s}' \in \mathbb{F}_q^N$. Choose $\mathbf{a} \in \mathbb{F}_q^N$ with $\mathbf{a}^\intercal(\mathbf{s} - \mathbf{s}') = 1$; such a vector exists because $\mathbf{s} - \mathbf{s}' \neq \mathbf{0}$. For any fixed $\mathbf{u} \in \mathbb{F}_q^\ell$, the matrix $\mathbf{M}' := \mathbf{M} + (\mathbf{J}\mathbf{u})\mathbf{a}^\intercal$ is also uniform. Translation equivariance from Lemma 3.5 gives

$$\begin{aligned} f_{\mathbf{T}}(\mathbf{M}'\mathbf{s}) - f_{\mathbf{T}}(\mathbf{M}'\mathbf{s}') &= f_{\mathbf{T}}(\mathbf{M}\mathbf{s}) + \mathbf{u}(\mathbf{a}^\intercal\mathbf{s}) \\ &\quad - f_{\mathbf{T}}(\mathbf{M}\mathbf{s}') - \mathbf{u}(\mathbf{a}^\intercal\mathbf{s}') \\ &= f_{\mathbf{T}}(\mathbf{M}\mathbf{s}) - f_{\mathbf{T}}(\mathbf{M}\mathbf{s}') + \mathbf{u}. \end{aligned}$$

The difference law is translation-invariant on the finite label space, hence uniform. The two labels therefore coincide with probability $q^{-\ell}$.

Fix the true secret and one incorrect candidate. By this uniform difference law and independence of the matrices, the candidate matches all labels with probability $q^{-\ell t}$. A union bound over the $q^N - 1$ incorrect candidates proves the stated recovery bound.

Exhaustive search recovers the secret whenever it is the unique candidate explaining every label. For $\ell = 1$, take $t = N + k$ with an integer $k \geq 0$. The recovery failure probability is at most

$$(1 - q^{-N})q^{-k} < q^{-k}.$$

For fixed output dimension and confidence, $O(N)$ examples therefore suffice for exact recovery, which implies information-theoretic probably approximately correct (PAC) learnability under uniform inputs.

The consistency test accepts exactly when some secret explains every label. The sampled common secret always explains a real table, so the test accepts $\mathsf{Real}_t^{\mathsf{LAR}}$ with probability one. In the ideal experiment, condition on any public matrix table and fix a candidate secret. The independent uniform labels match its outputs with probability $q^{-\ell t}$. A union bound over all q^N candidates bounds the acceptance probability on $\mathsf{Ideal}_t^{\mathsf{LAR}}$ by $q^{N-\ell t}$. Thus the distinguishing advantage is at least $1-q^{N-\ell t}$ when $\ell t > N$. Both procedures enumerate q^N secrets and generally take exponential time, so neither contradicts Definition 3.7 or the later conditional hardness results for efficient learning. □

A.7 From uniform-zero-fiber LWE to decisional LAR

Proof of Theorem 4.2. Fix $\mathcal{D}$ and t. We use the following sequence of games, using the standard indistinguishability and failure-event transitions described in [Sho04, Section 1.1]. Each game constructs the indexed family $\{(\mathbf{M}_i, \mathbf{y}_i)\}_{i=1}^t$, retaining repeated entries, gives it to $\mathcal{D}$ together with $(1^\lambda, \mathsf{pp})$, and outputs its decision bit. The displays below specify how this family is generated. The *view* of $\mathcal{D}$ means this input family; its public parameters are fixed throughout the games. We suppress these common public inputs in the acceptance probabilities below. Auxiliary random choices are kept internal to the game. Every sampling instruction uses fresh independent randomness; variables sampled before a loop are shared across its iterations. Vector and matrix arithmetic is over $\mathbb{F}_q$. The games specify probability experiments with exact uniform draws; a worst-case polynomial-time implementation is constructed at the end of the proof. All views are classical, so their distributional identities and mixture formulas are independent of the implementation of $\mathcal{D}$. The contraction bound in Section 2.3 applies throughout. The reduction generates the complete table classically, invokes $\mathcal{D}$ once, and returns its measured output bit without inspecting or rewinding its quantum state.

Game $\mathbf{G}_0$: This is the real LAR experiment:

$$\left\{ \begin{array}{rl} & \mathbf{s} \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right) \\ \{(\mathbf{M}_i, \mathbf{y}_i)\}_{i=1}^t \ : & \textbf{for } i = 1, \ldots, t \textbf{ do} \\ & \quad \mathbf{M}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right);\ \mathbf{y}_i := f_{\mathbf{T}}(\mathbf{M}_i\mathbf{s}) \end{array} \right\}$$

Difference. This is the starting game, with no modification to the real experiment: the view of $\mathcal{D}$ has law $\mathsf{Real}_t^{\mathsf{LAR}}$. Thus

$$\left|\Pr[\mathbf{G}_0 = 1] - \Pr[\mathcal{D}(\mathsf{Real}_t^{\mathsf{LAR}}) = 1]\right| = 0.$$

Game $\mathbf{G}_1$: We modify $\mathbf{G}_0$ by defining the event $\mathsf{E}_0 := \{\mathbf{s} = \mathbf{0}\}$. On this event, replace the labels by independent uniform vectors in $\mathbb{F}_q^\ell$:

$$\left\{ \begin{array}{rl} & \mathbf{s} \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right);\ \mathsf{E}_0 := \{\mathbf{s} = \mathbf{0}\} \\ & \textbf{for } i = 1, \ldots, t \textbf{ do} \\ \{(\mathbf{M}_i, \mathbf{y}_i)\}_{i=1}^t \ : & \quad \mathbf{M}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right) \\ & \quad \textbf{if } \neg\mathsf{E}_0 \textbf{ then } \mathbf{y}_i := f_{\mathbf{T}}(\mathbf{M}_i\mathbf{s}) \\ & \quad \textbf{else } \mathbf{y}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^\ell\right) \end{array} \right\}$$

Difference. Couple the two games to use the same classical secret and matrices. Their input views agree whenever E_0 does not occur. Let $\mathsf{Real}_t^{\mathsf{LAR},\neq 0}$ denote the law of the output family in the real LAR experiment conditioned on its sampled secret being nonzero. The conditional view in $\mathbf{G}_1$ has this law on $\overline{\mathsf{E}_0}$ and has law $\mathsf{Ideal}_t^{\mathsf{LAR}}$ on E_0. Its unconditional law is therefore

$$(1-q^{-N})\mathsf{Real}_t^{\mathsf{LAR},\neq 0} + q^{-N}\mathsf{Ideal}_t^{\mathsf{LAR}}.$$

On E_0, the labels in $\mathbf{G}_0$ are all zero because $f_{\mathbf{T}}(\mathbf{0}) = \mathbf{0}$; in $\mathbf{G}_1$ they are independent and uniform. The SD between the conditional laws of the two views on this event is exactly $1-q^{-\ell t}$. Since the secret is uniform, the difference in acceptance probabilities satisfies

$$\big|\Pr[\mathbf{G}_0 = 1] - \Pr[\mathbf{G}_1 = 1]\big| \le q^{-N}(1-q^{-\ell t}). \tag{3}$$

Game $\mathbf{G}_2$: Independently sample $\mathbf{a}_0 \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right)$ and define

$$\mathsf{E} := \{\mathbf{a}_0^{\intercal}\mathbf{s} \neq 1\}.$$

Modify $\mathbf{G}_1$ by using independent uniform labels whenever E occurs, and real labels otherwise. Since $\mathbf{a}_0^{\intercal}\mathbf{0} = 0 \neq 1$, we have $\mathsf{E}_0 \subseteq \mathsf{E}$. Thus the zero-secret case still uses uniform labels:

$$\left\{\{(\mathbf{M}_i,\mathbf{y}_i)\}_{i=1}^t \;:\; \begin{array}{l} \mathbf{s} \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right);\ \mathbf{a}_0 \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right);\ \mathsf{E} := \{\mathbf{a}_0^{\intercal}\mathbf{s} \neq 1\} \\ \textbf{for } i = 1,\dots,t \textbf{ do} \\ \quad \mathbf{M}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right) \\ \quad \textbf{if } \neg\mathsf{E} \textbf{ then } \mathbf{y}_i := f_{\mathbf{T}}(\mathbf{M}_i\mathbf{s}) \\ \quad \textbf{else } \mathbf{y}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{\ell}\right) \end{array}\right\}$$

Difference. For each nonzero $\mathbf{s}$, exactly a $1/q$ fraction of the choices of $\mathbf{a}_0$ satisfy $\mathbf{a}_0^{\intercal}\mathbf{s} = 1$: the nonzero linear functional $\mathbf{a}_0 \mapsto \mathbf{a}_0^{\intercal}\mathbf{s}$ has exactly q^{N-1} preimages of each field element. Thus

$$\Pr[\overline{\mathsf{E}}] = \frac{\Pr[\overline{\mathsf{E}_0}]}{q} = \frac{1-q^{-N}}{q},$$

and $\mathbf{s}$ conditioned on $\overline{\mathsf{E}}$ is uniform over $\mathbb{F}_q^N \setminus \{\mathbf{0}\}$. The conditional view in $\mathbf{G}_2$ has law $\mathsf{Real}_t^{\mathsf{LAR},\neq 0}$ on $\overline{\mathsf{E}}$ and law $\mathsf{Ideal}_t^{\mathsf{LAR}}$ on E. Its unconditional law is

$$\frac{1-q^{-N}}{q}\mathsf{Real}_t^{\mathsf{LAR},\neq 0} + \left(1 - \frac{1-q^{-N}}{q}\right)\mathsf{Ideal}_t^{\mathsf{LAR}}.$$

Game $\mathbf{G}_3$: We replace the direct sampling and event-dependent label rule in $\mathbf{G}_2$ by a transformation of real uniform-zero-fiber LWE blocks:

$$\left\{\{(\mathbf{M}_i,\mathbf{y}_i)\}_{i=1}^t \;:\; \begin{array}{l} \mathbf{s} \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right);\ \mathbf{a}_0 \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right) \\ \textbf{for } i = 1,\dots,t \textbf{ do} \\ \quad \mathbf{A}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right);\ \mathbf{e}_i \hookleftarrow \mathsf{U}\left(\mathcal{F}_{\mathbf{T},0}\right);\ \mathbf{y}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{\ell}\right) \\ \quad \mathbf{b}_i := \mathbf{A}_i\mathbf{s} + \mathbf{e}_i;\ \mathbf{M}_i := -\mathbf{A}_i + (\mathbf{b}_i + \mathbf{J}\mathbf{y}_i)\mathbf{a}_0^{\intercal} \end{array}\right\}$$

Here the secret, input matrices, and errors are drawn by the real LWE experiment. A procedure given only the blocks $\{(\mathbf{A}_i,\mathbf{b}_i)\}_{i=1}^t$ can produce the displayed output matrices by sampling $\mathbf{a}_0$ and the labels and applying the last assignment. It does not need the secret or errors.

Difference. This modification preserves the output distribution. To prove this, define $\mathbf{x}_i := \mathbf{e}_i + \mathbf{J}\mathbf{y}_i$. By Lemma 4.1, these vectors are independent and uniform over $\mathbb{F}_q^d$, independent of all input matrices $\mathbf{A}_i$, the secret, and $\mathbf{a}_0$, with $\mathbf{y}_i = f_{\mathbf{T}}(\mathbf{x}_i)$. The matrices in the rank-one transformation above satisfy

$$\mathbf{M}_i = -\mathbf{A}_i(\mathbf{I}_N - \mathbf{s}\mathbf{a}_0^\intercal) + \mathbf{x}_i\mathbf{a}_0^\intercal.$$

Fix $\mathbf{s}$ and $\mathbf{a}_0$ such that E does not occur. Multiplying this identity by $\mathbf{s}$ gives

$$\mathbf{M}_i\mathbf{s} = \mathbf{x}_i, \qquad f_{\mathbf{T}}(\mathbf{M}_i\mathbf{s}) = \mathbf{y}_i.$$

The linear map from $(\mathbf{A}_i, \mathbf{x}_i)$ to $\mathbf{M}_i$ is surjective: any prescribed $\mathbf{M}_i$ is obtained by taking $\mathbf{A}_i = -\mathbf{M}_i$ and $\mathbf{x}_i = \mathbf{M}_i\mathbf{s}$. Each fiber of this surjective linear map is a translate of its kernel and hence has the same cardinality. The map therefore sends the uniform input pair to a uniform matrix. The blocks remain independent for these fixed choices of $\mathbf{s}$ and $\mathbf{a}_0$, exactly as in $\mathbf{G}_2$.

Now fix $\mathbf{s}$ and $\mathbf{a}_0$ such that E occurs. Then

$$\det(\mathbf{I}_N - \mathbf{s}\mathbf{a}_0^\intercal) = 1 - \mathbf{a}_0^\intercal\mathbf{s} \neq 0.$$

For these fixed values, conditional on all $\mathbf{x}_i$, the transformed-matrix expression above defines an affine bijection in each input matrix $\mathbf{A}_i$. Hence the output matrices are independent and uniform, independently of the vectors $\mathbf{x}_i$ and their labels. These labels are themselves independent and uniform. This conditional output law is $\mathsf{Ideal}_t^{\mathsf{LAR}}$ for every such fixed secret and auxiliary vector, so it remains the same after averaging over their conditional law. On $\overline{\mathsf{E}}$ the previous argument gives the real branch of $\mathbf{G}_2$. The secret and auxiliary vector have the same joint law in both games, proving equality of the output laws. Therefore,

$$\big|\Pr[\mathbf{G}_2 = 1] - \Pr[\mathbf{G}_3 = 1]\big| = 0.$$

Game $\mathbf{G}_4$: We replace the real uniform-zero-fiber LWE blocks in $\mathbf{G}_3$ by ideal LWE blocks, keeping the transformation unchanged:

$$\left\{ \{(\mathbf{M}_i, \mathbf{y}_i)\}_{i=1}^t \ : \begin{array}{l} \mathbf{a}_0 \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right) \\ \textbf{for } i = 1, \ldots, t \textbf{ do} \\ \quad \mathbf{A}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right);\ \mathbf{b}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^d\right);\ \mathbf{y}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^\ell\right) \\ \quad \mathbf{M}_i := -\mathbf{A}_i + (\mathbf{b}_i + \mathbf{J}\mathbf{y}_i)\mathbf{a}_0^\intercal \end{array} \right\}$$

Difference. The gap between $\mathbf{G}_3$ and $\mathbf{G}_4$ is exactly the advantage of the following procedure $\mathcal{B}^{\mathrm{id}}$, defined using exact uniform auxiliary choices. On input $\{(\mathbf{A}_i, \mathbf{b}_i)\}_{i=1}^t$, it samples $\mathbf{a}_0 \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right)$ and $\mathbf{y}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^\ell\right)$ independently of the input and one another, computes

$$\mathbf{M}_i := -\mathbf{A}_i + (\mathbf{b}_i + \mathbf{J}\mathbf{y}_i)\mathbf{a}_0^\intercal,$$

and gives the resulting indexed family of matrices and labels, together with $(1^\lambda, \mathsf{pp})$, to $\mathcal{D}$. It returns the distinguisher's output bit. It simulates $\mathbf{G}_3$ on real input and $\mathbf{G}_4$ on ideal input. This procedure performs polynomially many field operations and one call to $\mathcal{D}$. It does not test either bad event or reveal $\mathbf{a}_0$. Hence

$$\big|\Pr[\mathbf{G}_3 = 1] - \Pr[\mathbf{G}_4 = 1]\big| = \mathsf{Adv}_{\mathcal{B}^{\mathrm{id}}}^{\mathsf{LWE}}(\lambda, t). \tag{4}$$

Game $\mathbf{G}_5$: We replace the affine construction in $\mathbf{G}_4$ by direct sampling of independent uniform matrices and labels:

$$\left\{ \{(\mathbf{M}_i, \mathbf{y}_i)\}_{i=1}^{t} \ : \ \begin{array}{l} \textbf{for } i = 1, \ldots, t \textbf{ do} \\ \quad \mathbf{M}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d \times N}\right);\ \mathbf{y}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{\ell}\right) \end{array} \right\}$$

Difference. This modification preserves the output distribution. To see this, condition in $\mathbf{G}_4$ on $\mathbf{a}_0$, all $\mathbf{b}_i$, and all $\mathbf{y}_i$. The rank-one matrix transformation defined above is an affine bijection in each independent uniform $\mathbf{A}_i$. Thus the output matrices are independent and uniform, independently of the labels, as in $\mathbf{G}_5$. By its sampling rule, $\mathbf{G}_5$ is exactly the ideal LAR experiment. Consequently,

$$\left|\Pr[\mathbf{G}_4 = 1] - \Pr[\mathbf{G}_5 = 1]\right| = 0.$$

Combining the games. The initial and final games satisfy

$$\mathsf{Adv}_{\mathcal{D}}^{\mathsf{dLAR}}(\lambda, t) = \left|\Pr[\mathbf{G}_0 = 1] - \Pr[\mathbf{G}_5 = 1]\right|.$$

Since $\mathbf{G}_2$ and $\mathbf{G}_3$ have the same output law, as do $\mathbf{G}_4$ and $\mathbf{G}_5$, Equation (4) gives

$$\left|\Pr[\mathbf{G}_2 = 1] - \Pr[\mathbf{G}_5 = 1]\right| = \mathsf{Adv}_{\mathcal{B}^{\mathrm{id}}}^{\mathsf{LWE}}(\lambda, t).$$

The laws of the views in $\mathbf{G}_1$ and $\mathbf{G}_2$ are mixtures of the same two distributions, $\mathsf{Real}_t^{\mathsf{LAR},\neq 0}$ and $\mathsf{Ideal}_t^{\mathsf{LAR}}$, with respective real-branch weights $\Pr[\overline{\mathsf{E}_0}]$ and $\Pr[\overline{\mathsf{E}_0}]/q$. Subtracting the ideal acceptance probability and taking absolute values gives

$$\begin{aligned} \left|\Pr[\mathbf{G}_1 = 1] - \Pr[\mathbf{G}_5 = 1]\right| &= q\left|\Pr[\mathbf{G}_2 = 1] - \Pr[\mathbf{G}_5 = 1]\right| \\ &= q\,\mathsf{Adv}_{\mathcal{B}^{\mathrm{id}}}^{\mathsf{LWE}}(\lambda, t). \end{aligned}$$

The triangle inequality and (3) yield

$$\begin{aligned} \mathsf{Adv}_{\mathcal{D}}^{\mathsf{dLAR}}(\lambda, t) &\le \left|\Pr[\mathbf{G}_0 = 1] - \Pr[\mathbf{G}_1 = 1]\right| + \left|\Pr[\mathbf{G}_1 = 1] - \Pr[\mathbf{G}_5 = 1]\right| \\ &\le q\,\mathsf{Adv}_{\mathcal{B}^{\mathrm{id}}}^{\mathsf{LWE}}(\lambda, t) + q^{-N}(1 - q^{-\ell t}). \end{aligned} \tag{5}$$

Implement the $N + \ell t$ auxiliary field choices of $\mathcal{B}^{\mathrm{id}}$ by standard capped rejection sampling [Dev86, Section III.5.1], independently of the input blocks, with joint statistical error at most

$$\varepsilon_{\mathrm{field}} := \tfrac{1}{2} q^{-(N+\ell t+1)}.$$

Both the number of choices and $\log(1/\varepsilon_{\mathrm{field}})$ are polynomially bounded in λ, so the resulting $\mathcal{B}$ is uniform QPT with a worst-case polynomial running-time bound, using the same input blocks and distinguisher call. The standard failure-event bound [Sho04, Section 1.1], applied in each LWE experiment, changes the advantage by at most $2\varepsilon_{\mathrm{field}}$. Thus the additional loss $2q\varepsilon_{\mathrm{field}} = q^{-N-\ell t}$ is absorbed by the unused zero-secret allowance in (5), proving the bound in Theorem 4.2. □

A.8 From Gaussian LWE to uniform-zero-fiber LWE

Proof of Theorem 5.4. Fix a positive integer c, a distinguisher $\mathcal{D}$, and a block count t. Use the sampler Sample_c from Definition 5.3. The reduction below is defined at every security parameter; we prove its advantage bound for $\lambda \ge \lambda_0(c)$. We use the sequence-of-games method [Sho04, Section 1.1]. Each game gives $\{(\mathbf{A}'_i, \mathbf{b}'_i)\}_{i=1}^{t}$ to $\mathcal{D}$, together with $(1^\lambda, \mathsf{pp})$, and outputs its decision bit. We suppress

these common public inputs in the acceptance probabilities below. Every sampling instruction uses fresh independent randomness; variables sampled before a loop are shared across its iterations. The transport pairs remain internal to the games. All arithmetic is over $\mathbb{F}_q$. The sampler and all table transformations are classical. Equality of table laws is preserved by the quantum distinguisher, and the contraction bound in Section 2.3 controls the approximate transition.

Game $\mathbf{G}_0$: This is the real uniform-zero-fiber LWE experiment:

$$\left\{ \{(\mathbf{A}'_i, \mathbf{b}'_i)\}_{i=1}^t \ : \begin{array}{l} \mathbf{s} \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right) \\ \textbf{for } i = 1, \ldots, t \textbf{ do} \\ \quad \mathbf{A}'_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right);\ \mathbf{e}'_i \hookleftarrow \mathsf{U}\left(\mathcal{F}_{\mathbf{T},0}\right) \\ \quad \mathbf{b}'_i := \mathbf{A}'_i\mathbf{s} + \mathbf{e}'_i \end{array} \right\}$$

Difference. This is the starting game, whose sample distribution is exactly $\mathsf{Real}^{\mathsf{LWE}}_{\mathsf{U}(\mathcal{F}_{\mathbf{T},0}),t}$. Hence

$$\left|\Pr[\mathbf{G}_0 = 1] - \Pr[\mathcal{D}(\mathsf{Real}^{\mathsf{LWE}}_{\mathsf{U}(\mathcal{F}_{\mathbf{T},0}),t}) = 1]\right| = 0.$$

Game $\mathbf{G}_1$: We modify $\mathbf{G}_0$ by replacing each uniform zero-fiber error with an independently transported Gaussian error:

$$\left\{ \{(\mathbf{A}'_i, \mathbf{b}'_i)\}_{i=1}^t \ : \begin{array}{l} \mathbf{s} \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right) \\ \textbf{for } i = 1, \ldots, t \textbf{ do} \\ \quad \mathbf{A}'_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right);\ \mathbf{e}_i \hookleftarrow \chi_{\mathrm{G}} \\ \quad (\mathbf{Q}_i, \beta_i) \hookleftarrow \mathsf{Sample}_c(1^\lambda, \mathsf{pp}) \\ \quad \mathbf{e}'_i := \mathsf{Eval}(\mathbf{Q}_i, \beta_i, \mathbf{e}_i) \\ \quad \mathbf{b}'_i := \mathbf{A}'_i\mathbf{s} + \mathbf{e}'_i \end{array} \right\}$$

Difference. By Definition 5.3, the law of each $\mathbf{e}'_i$, averaged over both its Gaussian source error and its transport pair, is within λ^{-c} of $\mathsf{U}(\mathcal{F}_{\mathbf{T},0})$. The transported errors are independent across blocks, so the product bound for statistical distance [LP17, Exercise 4.4] gives

$$\Delta\Big(\mathcal{L}\Big(\{\mathbf{e}'_i\}_{i=1}^t\Big), \mathsf{U}(\mathcal{F}_{\mathbf{T},0})^{\otimes t}\Big) \le t\lambda^{-c}.$$

The entire public-matrix table and the shared secret are independent of these errors. Appending them and forming the responses cannot increase statistical distance, nor can quantum postprocessing distinguish the resulting classical tables with a larger gap. Therefore,

$$\left|\Pr[\mathbf{G}_0 = 1] - \Pr[\mathbf{G}_1 = 1]\right| \le t\lambda^{-c}.$$

Game $\mathbf{G}_2$: We replace the direct sampling of the public matrices in $\mathbf{G}_1$ by a transformation of real Gaussian LWE blocks:

$$\left\{ \{(\mathbf{A}'_i, \mathbf{b}'_i)\}_{i=1}^t \ : \begin{array}{l} \mathbf{s} \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^N\right) \\ \textbf{for } i = 1, \ldots, t \textbf{ do} \\ \quad \mathbf{A}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right);\ \mathbf{e}_i \hookleftarrow \chi_{\mathrm{G}} \\ \quad (\mathbf{Q}_i, \beta_i) \hookleftarrow \mathsf{Sample}_c(1^\lambda, \mathsf{pp}) \\ \quad \mathbf{b}_i := \mathbf{A}_i\mathbf{s} + \mathbf{e}_i \\ \quad \mathbf{A}'_i := \mathbf{Q}_i\mathbf{A}_i;\ \mathbf{b}'_i := \mathsf{Eval}(\mathbf{Q}_i, \beta_i, \mathbf{b}_i) \end{array} \right\}$$

Difference. The affine evaluation rule gives

$$\mathbf{b}'_i = \mathbf{A}'_i\mathbf{s} + \mathbf{e}'_i, \qquad \mathbf{e}'_i := \mathbf{Q}_i\mathbf{e}_i + \beta_i.$$

Fix the shared secret and apply Lemma 5.2 independently in each block. The entire transformed matrix table is exactly uniform and independent of the table of transported errors; these joint laws do not depend on the fixed secret. The errors have exactly the joint law used in $\mathbf{G}_1$, with the same shared secret preserved. Thus

$$\left|\Pr[\mathbf{G}_1 = 1] - \Pr[\mathbf{G}_2 = 1]\right| = 0.$$

Game $\mathbf{G}_3$: We replace the real Gaussian LWE blocks in $\mathbf{G}_2$ by ideal LWE blocks, keeping the affine transformation unchanged:

$$\left\{ \{(\mathbf{A}'_i, \mathbf{b}'_i)\}_{i=1}^t \ : \begin{array}{l} \textbf{for } i = 1, \dots, t \textbf{ do} \\ \quad \mathbf{A}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right);\ \mathbf{b}_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d}\right) \\ \quad (\mathbf{Q}_i, \beta_i) \hookleftarrow \mathsf{Sample}_c(1^\lambda, \mathsf{pp}) \\ \quad \mathbf{A}'_i := \mathbf{Q}_i\mathbf{A}_i;\ \mathbf{b}'_i := \mathsf{Eval}(\mathbf{Q}_i, \beta_i, \mathbf{b}_i) \end{array} \right\}$$

Difference. Define the Gaussian LWE distinguisher $\mathcal{B}_c$ as follows. On receiving its input blocks, it samples a fresh transport pair for each block using $\mathsf{Sample}_c(1^\lambda, \mathsf{pp})$, independently of the entire input experiment, and applies the matrix and response transformation in $\mathbf{G}_2$. It gives only the transformed classical table to $\mathcal{D}$ and returns its measured output bit, without inspecting or rewinding its quantum state. For fixed c, the sampler runs in worst-case polynomial time at every security parameter, so the classical preprocessing followed by the QPT call makes $\mathcal{B}_c$ uniform QPT. It uses t input blocks and one call to $\mathcal{D}$, without needing to compute $\lambda_0(c)$. It simulates $\mathbf{G}_2$ on real Gaussian input and $\mathbf{G}_3$ on ideal input. Consequently,

$$\left|\Pr[\mathbf{G}_2 = 1] - \Pr[\mathbf{G}_3 = 1]\right| = \mathsf{Adv}^{\mathsf{Gauss\text{-}LWE}}_{\mathcal{B}_c}(\lambda, t).$$

Game $\mathbf{G}_4$: We replace the affine construction in $\mathbf{G}_3$ by direct sampling of independent uniform matrices and responses:

$$\left\{ \{(\mathbf{A}'_i, \mathbf{b}'_i)\}_{i=1}^t \ : \begin{array}{l} \textbf{for } i = 1, \dots, t \textbf{ do} \\ \quad \mathbf{A}'_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d\times N}\right);\ \mathbf{b}'_i \hookleftarrow \mathsf{U}\left(\mathbb{F}_q^{d}\right) \end{array} \right\}$$

Difference. Condition in $\mathbf{G}_3$ on all transport pairs. By Lemma 5.2, each transformed pair consists of an independent uniform matrix and response, and these pairs remain independent across blocks. This product law does not depend on the conditioned values, so averaging gives exactly $\mathsf{Ideal}_t^{\mathsf{LWE}}$, the distribution in $\mathbf{G}_4$. Hence

$$\left|\Pr[\mathbf{G}_3 = 1] - \Pr[\mathbf{G}_4 = 1]\right| = 0.$$

Combining the games. The initial and final games are the real and ideal uniform-zero-fiber LWE experiments, respectively. The triangle inequality and the bounds above give

$$\begin{aligned} \mathsf{Adv}^{\mathsf{UZF\text{-}LWE}}_{\mathcal{D}}(\lambda, t) &= \left|\Pr[\mathbf{G}_0 = 1] - \Pr[\mathbf{G}_4 = 1]\right| \\ &\le \left|\Pr[\mathbf{G}_0 = 1] - \Pr[\mathbf{G}_1 = 1]\right| + \left|\Pr[\mathbf{G}_2 = 1] - \Pr[\mathbf{G}_3 = 1]\right| \\ &\le t\lambda^{-c} + \mathsf{Adv}^{\mathsf{Gauss\text{-}LWE}}_{\mathcal{B}_c}(\lambda, t). \end{aligned}$$

For the hardness implication, suppose that some uniform QPT distinguisher $\mathcal{D}$, polynomially bounded t, and positive polynomial p satisfy $\mathsf{Adv}_{\mathcal{D}}^{\mathsf{UZF\text{-}LWE}}(\lambda, t) \geq 1/\mathsf{p}(\lambda)$ for infinitely many λ. Choose fixed positive integers b, c_{adv} such that, for all sufficiently large λ,

$$t(\lambda) \leq \lambda^b, \qquad \mathsf{p}(\lambda) \leq \lambda^{c_{\mathrm{adv}}}.$$

Fix one integer $c > b + c_{\mathrm{adv}} + 2$. For all sufficiently large λ,

$$t(\lambda)\lambda^{-c} \leq \lambda^{b-c} \leq \frac{1}{2\lambda^{c_{\mathrm{adv}}}} \leq \frac{1}{2\mathsf{p}(\lambda)}.$$

Apply the quantitative bound just proved with the fixed sampler Sample_c. For infinitely many sufficiently large λ, it gives

$$\mathsf{Adv}_{\mathcal{B}_c}^{\mathsf{Gauss\text{-}LWE}}(\lambda, t) \geq \frac{1}{2\mathsf{p}(\lambda)}.$$

Since c is fixed, $\mathcal{B}_c$ is uniform QPT, contradicting Gaussian decisional LWE hardness. This order depends on the fixed distinguisher and its polynomial bounds; the public parameter family and Gaussian source remain unchanged. □

A.9 Finite-dimensional calibration

Proof of Theorem 7.1. The defect of Section 7 satisfies $\|\mathcal{K}^k\| \leq \kappa^k$, so its Neumann series converges in operator norm. Multiplying the partial sum through power k by $\mathcal{A} = I - \mathcal{K}$ on either side gives $I - \mathcal{K}^{k+1}$, which tends to the identity. Thus the series is the inverse, with norm at most $\sum_{k\geq 0} \kappa^k = (1-\kappa)^{-1}$. Take $z = \mathcal{A}^{-1}\tau$, use the response identity, and apply the inverse bound to $z - z_0 = \mathcal{A}^{-1}(\tau - \mathcal{A}z_0)$. The mass constraint gives the final claim. □

Proof of Corollary 7.2. Apply Theorem 7.1 with reference $z_0 = \tau$; its residual has norm at most $\kappa\|\tau\|_\infty$. For the second claim use $\inf \tau \geq 1 - a$ and $\|\tau\|_\infty \leq 1 + b$. □

Proof of Corollary 7.3. Inside $[0, 2]$ the clipping error is zero; outside it, $|W - 1| \geq 1$ and the clipping error is at most $|W - 1|$. Integration gives the stated loss bound. □

B Incidence Geometry and Affine Transport Proofs

This appendix fixes the incidence conventions and verifies the geometric and algebraic statements of Section 6. In the construction, $\mathbf{D}_{\mathscr{T}}$ maps tree-edge coefficients to vertex loads; the product $\overline{\mathbf{D}}_{\mathscr{T}}\overline{\mathbf{R}}$ includes the source-coordinate mixing and the separate charge-injection column of Corollary 6.8.

Definition B.1 (Hypergraph)**.** Let $d \geq 1$ and use the vertex set V fixed in Section 6.1. A finite *hypergraph* $G = (V, E)$ consists of the vertex set V, a finite set E of edge labels, and a nonempty support $\mathrm{supp}(e) \subseteq V$ for each $e \in E$. Distinct labels may have the same support. A vertex v is *incident* to e if $v \in \mathrm{supp}(e)$; its *degree* is the number of incident hyperedges. An *ordinary loopless graph* is a hypergraph with $|\mathrm{supp}(e)| = 2$ for every $e \in E$; parallel edges are allowed.

Orienting individual vertex–edge incidences is standard in bidirected graph theory [Zas91, Section 2]. For hyperedges, we use the head–tail representation and signed incidence matrix of Gallo et al. [GLPN93, Section 2], with the opposite global sign convention: $+1$ at tails and -1 at heads.

Definition B.2 (Oriented hypergraph and signed incidence matrix)**.** An *orientation* of a hypergraph G assigns to each $e \in E$ a partition $\operatorname{tail}(e) \dot\cup \operatorname{head}(e) = \operatorname{supp}(e)$ into disjoint tail and head sets. Either set may be empty. The hypergraph together with these choices is an *oriented hypergraph*. Its *signed incidence matrix* is

$$\mathbf{D}_G \in \{-1, 0, 1\}^{d \times |E|}, \qquad (\mathbf{D}_G)_{v,e} := \begin{cases} +1 & \text{if } v \in \operatorname{tail}(e), \\ -1 & \text{if } v \in \operatorname{head}(e), \\ 0 & \text{otherwise.} \end{cases}$$

An *ordinary directed graph* is the special case with $|\operatorname{tail}(e)| = |\operatorname{head}(e)| = 1$ for every edge.

An ordinary edge column is the tail basis vector minus the head basis vector; reversing orientation negates it. Mixing tree-edge coefficients changes the source-to-load matrix while retaining this incidence convention.

Definition B.3 (Integer-weighted incidence matrix)**.** Let G be an oriented hypergraph. An *integer-weighted incidence matrix* of G is a matrix $\mathbf{D}_{G,\mathrm{wt}} \in \mathbb{Z}^{d \times |E|}$ satisfying, for every vertex v and hyperedge e,

$$(\mathbf{D}_{G,\mathrm{wt}})_{v,e} \in \begin{cases} \mathbb{Z}_{>0} & \text{if } v \in \operatorname{tail}(e), \\ \mathbb{Z}_{<0} & \text{if } v \in \operatorname{head}(e), \\ \{0\} & \text{otherwise.} \end{cases}$$

The absolute value of a nonzero entry is its incidence weight; weights may differ within a hyperedge.

Definition B.4 (Paths, trees, and spanning trees)**.** Let G be an ordinary loopless graph.

Path. A *path* is an alternating sequence of vertices and edges in which each edge joins its two neighboring vertices and no vertex is repeated.

Cycle. A *cycle* is a closed alternating sequence of vertices and edges satisfying the same incidence condition, with at least two edges, no repeated edge, and no repeated vertex other than the coincident endpoints.

Connected graph. A graph is *connected* if every two vertices are joined by a path.

Tree. A graph is a *tree* if it is connected and contains no cycle.

Spanning tree. A *spanning tree* of G is a subgraph that is a tree and contains all of V.

For directed graphs these notions refer to the underlying graph obtained by forgetting edge orientations.

B.1 Proof of the tree incidence basis

Proof of Lemma 6.2. For $d = 1$, the empty incidence matrix generates the trivial lattice. For $d > 1$, a tree has a leaf, and deleting the leaf and its unique incident edge leaves a tree. Induction therefore

gives $d-1$ edges. Every incidence column has zero sum, so its integer span is contained in A_{d-1}. Conversely, let $\mathbf{z} \in A_{d-1}$. At a leaf, the prescribed vertex value determines the unique incident edge coefficient up to its orientation sign. Subtract that multiple of the incidence column and delete the leaf. The remaining vector still has zero sum, so induction determines all remaining edge coefficients uniquely. The final vertex equation is automatic. Thus every vector in A_{d-1} has a unique integral expansion in the tree columns, proving the basis and rank assertions. □

B.2 Image lattice and quotient coordinates

Proof of Lemma 6.4. Every path column has coordinate sum zero, and $\mathbf{v}_0$ has coordinate sum q. Hence every vector in $\mathcal{L}(\mathbf{T})$ has coordinate sum divisible by q. Conversely, let $\mathbf{z} \in \mathbb{Z}^d$ satisfy this divisibility condition and set $k := \mathbf{1}^\intercal \mathbf{z}/q \in \mathbb{Z}$. Then

$$\mathbf{z} - k\mathbf{v}_0 \in A_{d-1} = \mathbf{D}_d \mathbb{Z}^{d-1},$$

where the equality follows from Lemma 6.2. Thus $\mathbf{z} \in \mathcal{L}(\mathbf{T})$, proving the characterization.

The coordinate-sum homomorphism $\mathbf{z} \mapsto \mathbf{1}^\intercal \mathbf{z} \bmod q$ from $\mathbb{Z}^d$ onto $\mathbb{F}_q$ has kernel $\mathcal{L}(\mathbf{T})$. Its quotient is therefore cyclic of order q, which gives the index and the claimed Smith normal form. The induced quotient coordinate map is the one specified in Definition 6.3. The rounding-matrix parameters follow from Definition 2.3. □

B.3 Canonical endpoint parametrization

Proof of Lemma 6.5. For $\bar{\mathbf{w}} \in H_0$, the integer vector $\mathbf{T}^{-1}[\bar{\mathbf{w}}]_q$ is a lift of $c_{\mathbf{T}}(\bar{\mathbf{w}})$. Using lift independence from Lemma 3.5, we obtain

$$f_{\mathbf{T}}(c_{\mathbf{T}}(\bar{\mathbf{w}})) = \mathbf{P}\left\lfloor \frac{[\bar{\mathbf{w}}]_q}{q} \right\rfloor \bmod q = \mathbf{0}.$$

Thus the pullback indeed takes values in $\mathcal{F}_{\mathbf{T},0}$. Multiplying its defining congruence by the integer matrix $\mathbf{T}$ gives

$$\mathbf{T}[c_{\mathbf{T}}(\bar{\mathbf{w}})]_q \equiv [\bar{\mathbf{w}}]_q \pmod q.$$

Hence the proposed inverse recovers $\bar{\mathbf{w}}$, proving injectivity.

The hyperplane H_0 has q^{d-1} elements. By Lemma 6.4, $r = d-1$; the balanced fibers of Lemma 3.5 therefore give $|\mathcal{F}_{\mathbf{T},0}| = q^{d-1}$. The injection is consequently a bijection between these finite sets, and the preceding congruence gives the stated inverse on all of $\mathcal{F}_{\mathbf{T},0}$. □

B.4 Endpoint and carry criterion

Proof of Lemma 6.6. The endpoint $\bar{\mathbf{w}}$ lies in H_0 because $\mathbf{1}^\intercal \mathbf{T} \equiv \mathbf{0}^\intercal \pmod q$. Substituting the definitions of the affine output and canonical pullback gives the following difference in $\mathbb{F}_q^d$, with all products involving $\mathbf{T}^{-1}$ evaluated over $\mathbb{Q}$ before reduction modulo q:

$$\begin{aligned}
\mathbf{x}' - c_{\mathbf{T}}(\bar{\mathbf{w}}) &= (\mathbf{Q}[\mathbf{x}]_q + \beta - \mathbf{T}^{-1}[\bar{\mathbf{w}}]_q) \bmod q \\
&= \mathbf{T}^{-1}(\mathbf{T}(\mathbf{Q}[\mathbf{x}]_q + \beta) - [\bar{\mathbf{w}}]_q) \bmod q \\
&= \mathbf{T}^{-1}(\mathbf{w} - [\bar{\mathbf{w}}]_q) \bmod q \\
&= q\mathbf{T}^{-1}\mathbf{k} \bmod q
\end{aligned}$$

In particular, $q\mathbf{T}^{-1}\mathbf{k}$ is an integer vector. Its reduction modulo q vanishes exactly when $\mathbf{T}^{-1}\mathbf{k} \in \mathbb{Z}^d$, or equivalently $\mathbf{k} \in \mathcal{L}(\mathbf{T})$. Combining this with the lattice characterization of Lemma 6.4 proves the carry criterion of Lemma 6.6.

Couple $\mathbf{x}'$ and $c_{\mathbf{T}}(\bar{\mathbf{w}})$ using the same transport data and source input. The carry criterion just proved and the coupling inequality give

$$\Delta\big(\mathcal{L}(\mathbf{x}')\,,\mathcal{L}(c_{\mathbf{T}}(\bar{\mathbf{w}}))\big) \le \Pr[\mathbf{k} \notin \mathcal{L}(\mathbf{T})].$$

By Lemma 6.5, the map $c_{\mathbf{T}}$ is a bijection from H_0 onto $\mathcal{F}_{\mathbf{T},0}$, so

$$\Delta(\mathcal{L}(c_{\mathbf{T}}(\bar{\mathbf{w}}))\,,\mathsf{U}\,(\mathcal{F}_{\mathbf{T},0})) = \Delta(\mathcal{L}(\bar{\mathbf{w}})\,,\mathsf{U}\,(H_0))\,.$$

The claim follows from the triangle inequality. □

B.5 Extended incidence determinants

Proof of Lemma 6.7. The tree columns are an integer basis of A_{d-1} by Lemma 6.2. Summing coordinates identifies $\mathbb{Z}^d/A_{d-1}$ with $\mathbb{Z}$. If $\mathbf{1}^{\intercal}\mathbf{v} = 0$, the appended column lies in A_{d-1}, so the determinant vanishes. Otherwise, the image of the lattice generated by the appended matrix in this quotient is $(\mathbf{1}^{\intercal}\mathbf{v})\mathbb{Z}$. Hence its index in $\mathbb{Z}^d$ is $|\mathbf{1}^{\intercal}\mathbf{v}|$. This finite index equals the absolute determinant of the appended square matrix, proving the claim. □

B.6 Integral transport from an incidence basis

Proof of Corollary 6.8. Every column of $\mathbf{D}_{\mathscr{T}}\mathbf{R}_{\mathrm{edge}}$ has integer sum zero, and the appended column has sum q. By Lemma 6.4, these columns and the admissible initial load lie in $\mathcal{L}(\mathbf{T})$. Their $\mathbf{T}$-preimages are therefore integral, and every integer input yields a load in that lattice and an endpoint in H_0. Lemma 6.7 gives $|\det \overline{\mathbf{D}}_{\mathscr{T}}| = q$. Taking determinants in the defining product and dividing by $\det \mathbf{T}$, whose absolute value is q by Lemma 6.4, gives the determinant identities of Corollary 6.8. A nonzero determinant modulo the prime q gives the asserted modular invertibility. □

C Local Kernels and Bounded Sampling

This appendix verifies the body and leg laws of Section 8 and their bounded public samplers. It defines the color-fiber operators used for local invertibility and positivity, followed by the rational calibration system and its implementation.

C.1 The cycle body law

The coordinate identity in Definition 8.2 follows by multiplying the path block of Section 8.2 and inserting the two closure entries specified in that definition. Summing its coordinates telescopes to zero.

With an ideal draw of x_{cl}, the $n_{\mathcal{B}}-1$ independent body source coordinates and this one independent public Gaussian are $n_{\mathcal{B}}$ independent samples from $\mathsf{D}_{\mathbb{Z},\sigma}$. Their consecutive cyclic differences are exactly the body increment $\boldsymbol{\Delta}_{\mathrm{B}}^{\mathrm{loc}}$ of Definition 8.2, without convolution or smoothing loss. This joint-law identity remains valid after the prescribed body placement and conditional on public data sampled independently of the closure and hidden source. It retains all shared-coordinate dependencies. The identity is not conditioned on the realized closure or affine offset.

C.1.1 Bounded Sampling of the Public Cycle-Closing Gaussian

The body sampler draws only the public closure $x_{\mathrm{cl}} \hookleftarrow \mathsf{D}_{\mathbb{Z},\sigma}$ of Definition 8.2, retaining the hidden body coordinates. We implement the scalar Gaussian rejection rule of [GPV08, full version, Section 4.1] with bounded running time, accounting for truncation and finite precision. Assume $\sigma^2 \in \mathbb{Q}$ and $\sigma^2 \geq 2$, as in Lemma 8.20, and fix a rational $0 < \nu_{\mathrm{G}} < 1$.

Truncation and trial cap. Use the integer parameters

$$R_{\mathrm{G}} := \left\lceil \sqrt{\lceil \log_2(8/\nu_{\mathrm{G}}) \rceil} \right\rceil, \quad H_{\mathrm{G}} := R_{\mathrm{G}} \lceil \sigma \rceil,$$
$$M_{\mathrm{G}} := 2H_{\mathrm{G}} + 1, \qquad T_{\mathrm{G}} := \lceil 120 R_{\mathrm{G}} \lceil \log_2(4/\nu_{\mathrm{G}}) \rceil \rceil .$$

The ideal trial proposes an integer z_{G} uniformly from $[-H_{\mathrm{G}}, H_{\mathrm{G}}] \cap \mathbb{Z}$ and accepts it with probability $e^{-\pi z_{\mathrm{G}}^2/\sigma^2}$. The first accepted proposal has the Gaussian law conditioned on that interval. By Lemma 2.9 and Chernoff's bound,

$$\Pr_{\mathbf{x} \hookleftarrow \mathsf{D}_{\mathbb{Z},\sigma}} [|\mathbf{x}| > H_{\mathrm{G}}] \leq 2e^{-\pi(H_{\mathrm{G}}+1)^2/\sigma^2} \leq 2e^{-\pi R_{\mathrm{G}}^2} \leq \nu_{\mathrm{G}}/4.$$

This uses the integer threshold $H_{\mathrm{G}} + 1$; no smoothing assumption is required.

The interval $[-\sigma, \sigma]$ contains at least σ integers, all with acceptance weight at least $e^{-\pi}$. Since $\sigma \geq 1$ and $R_{\mathrm{G}} \geq 1$, $M_{\mathrm{G}} \leq 5R_{\mathrm{G}}\sigma$. Consequently every ideal trial succeeds with probability at least $e^{-\pi}/(5R_{\mathrm{G}}) > 1/(120R_{\mathrm{G}})$. Stop after at most T_{G} trials and return $x_{\mathrm{cl}} = 0$ if none succeeds. The failure probability is at most

$$\exp\bigl(-T_{\mathrm{G}}/(120R_{\mathrm{G}})\bigr) \leq \nu_{\mathrm{G}}/4.$$

The zero fallback preserves exact integer affine validity.

Finite precision. Choose the fixed bit counts

$$b_{\mathrm{G,prop}} := \lceil \log_2(8M_{\mathrm{G}}T_{\mathrm{G}}/\nu_{\mathrm{G}}) \rceil ,$$
$$b_{\mathrm{G,acc}} := \lceil \log_2(16T_{\mathrm{G}}/\nu_{\mathrm{G}}) \rceil .$$

Reduce a uniform $b_{\mathrm{G,prop}}$-bit integer modulo M_{G} and subtract H_{G} to obtain the proposal. Its statistical distance from uniform is at most $M_{\mathrm{G}}\, 2^{-b_{\mathrm{G,prop}}} \leq \nu_{\mathrm{G}}/(8T_{\mathrm{G}})$. Evaluate the Gaussian acceptance weight with absolute error at most $2^{-b_{\mathrm{G,acc}}}$, clamp the approximation to $[0, 1]$, and round down to a multiple of $2^{-b_{\mathrm{G,acc}}}$. A fresh uniform $b_{\mathrm{G,acc}}$-bit integer implements this coin exactly. The acceptance probability differs from its ideal value by at most $2^{1-b_{\mathrm{G,acc}}} \leq \nu_{\mathrm{G}}/(8T_{\mathrm{G}})$. Coupling all proposals and coins adds at most $\nu_{\mathrm{G}}/4$ over the capped execution. Together with truncation and the failure branch, this proves

$$\Delta(\mathcal{L}(x_{\mathrm{cl}}), \mathsf{D}_{\mathbb{Z},\sigma}) \leq \nu_{\mathrm{G}}.$$

All random bits are fresh and independent of the hidden source and all other sampling blocks.

Worst-case complexity. Exact integer and rational comparisons on the binary inputs determine $\lceil \sigma \rceil$ and all parameter and precision ceilings. No proposal-interval enumeration is needed. Each candidate and output has $O(\log \lceil \sigma \rceil + \log(R_{\mathrm{G}} + 1))$ bits. On this interval the absolute Gaussian exponent is at most $4\pi R_{\mathrm{G}}^2$, so standard rational approximations to π and the exponential evaluate the acceptance weight to the stated absolute precision in polynomial time. The cap T_{G} and both bit counts are polynomial in the binary input lengths and $\log(1/\nu_{\mathrm{G}})$. Thus every execution has polynomial running time and output length. The sampler never observes, conditions on, or rejects a hidden source coordinate.

C.2 Safe color slices

Under their respective parameter conditions, the slices have the following interval forms and positive sizes:

$$\begin{aligned} I_\xi &= \{\max(0,\xi),\dots,\min(q-1,q-1+\xi)\},\\ |I_\xi| &= q-|\xi|>0,\\ J_{\xi,h} &= \{h+\max(0,\xi),\dots,q-1-h+\min(0,\xi)\},\\ |J_{\xi,h}| &= q-2h-|\xi|>0. \end{aligned}$$

Proof of the slice identities following Definitions 8.3 and 8.4. For an ordered pair with sum $q-1+\xi$, requiring both coordinates to lie between zero and $q-1$ imposes

$$0\le b\le q-1, \qquad \xi\le b\le q-1+\xi.$$

Intersecting the intervals gives the safe slice of Definition 8.3; counting its integer points gives $q-|\xi|$. Requiring both coordinates instead to lie in C_h adds h to both lower bounds and subtracts it from both upper bounds. Their intersection gives the trimmed slice of Definition 8.4, with $q-2h-|\xi|$ integer points. The stated hypotheses make these counts positive.

For a protected second input, $|\zeta-t|\le|\zeta|+K_{\mathrm{L}}-1<q$. Thus every slice in Definition 8.8 is defined and nonempty. The stronger trimmed bound gives the same conclusion after deleting h colors at each endpoint. Membership ensures that the evaluated primary and its prescribed integer companion are both colors. Finally $c\neq 0$ makes the affine function of the first field input a bijection for each fixed second input. □

C.3 Off-body carry cancellation

For every integer input $(u,t)\in\mathbb{Z}^2$, the load formulas give the charge identities

$$\begin{aligned} X+Y &= q-1+\zeta-t,\\ X+Y+Z &= q-1+\zeta+z_{\mathrm{R}}. \end{aligned}$$

We restrict the sampling support to maps ω for which the pair sum satisfies

$$[X \bmod q]_q \le X+Y = q-1+\zeta-t < [X \bmod q]_q + q$$

simultaneously for every protected input (u,t) with $|u|,|t|<K_{\mathrm{L}}$.

Lemma C.1 (Off-body carry cancellation)**.** Fix a public map ω with $c\neq 0$. If the off-body carry condition above holds for every protected input, then for every such input,

$$\lfloor X/q\rfloor+\lfloor Y/q\rfloor=0.$$

Proof of Lemma C.1. Multiplying the local incidence and mixer blocks reproduces the increment of Definition 8.10. Adding the restriction of the shared baseline gives its complete local load. Summing the first two, or all three, coordinates gives the pair and triple charge identities above. By the stated off-body support condition, the integer $X+Y-[X \bmod q]_q$ lies in $\{0,\dots,q-1\}$ and is congruent to Y modulo q. It therefore equals $[Y \bmod q]_q$. Subtracting the two canonical residues from the pair sum and dividing by q gives the claimed cancellation. □

C.4 Attachment safety

Attachment safety is controlled by the shared baseline margin. We require

$$K_{\mathrm{BS}} \le z_{\mathrm{R}} \le q-1-K_{\mathrm{BS}}, \qquad K_{\mathrm{B}}+K_{\mathrm{L}}+K_{\mathrm{Q}} \le K_{\mathrm{BS}},$$
$$|(\boldsymbol{\Delta}_{\mathrm{B}})_{v_{\mathrm{R}}}| < K_{\mathrm{B}}, \qquad |(\boldsymbol{\Delta}_{\mathrm{Q}})_{v_{\mathrm{R}}}| < K_{\mathrm{Q}},$$

Lemma C.2 (Attachment safety)**.** For every protected leg input and any body and queue contributions satisfying the attachment margin and contribution bounds above,

$$0 \le Z + (\boldsymbol{\Delta}_{\mathrm{B}})_{v_{\mathrm{R}}} + (\boldsymbol{\Delta}_{\mathrm{Q}})_{v_{\mathrm{R}}} < q.$$

In particular, the assembled attachment has zero carry:

$$\left\lfloor \frac{Z + (\boldsymbol{\Delta}_{\mathrm{B}})_{v_{\mathrm{R}}} + (\boldsymbol{\Delta}_{\mathrm{Q}})_{v_{\mathrm{R}}}}{q} \right\rfloor = 0.$$

Proof of Lemma C.2. By the local-load construction of Definition 8.10, $Z = z_{\mathrm{R}} + t$. Since $|t| < K_{\mathrm{L}}$, the triangle inequality gives

$$|t + (\boldsymbol{\Delta}_{\mathrm{B}})_{v_{\mathrm{R}}} + (\boldsymbol{\Delta}_{\mathrm{Q}})_{v_{\mathrm{R}}}| < K_{\mathrm{L}} + K_{\mathrm{B}} + K_{\mathrm{Q}} \le K_{\mathrm{BS}}.$$

Combining this bound with $K_{\mathrm{BS}} \le z_{\mathrm{R}} \le q-1-K_{\mathrm{BS}}$ places the assembled attachment load in $[0, q)$, so its carry is zero. □

C.5 Finite calibration of the line-mark law

We first specify the finite coefficient model and the complete public law used in Theorem 8.9.

Definition C.3 (Finite affine-map model)**.** Equip $\Omega := \mathbb{F}_q^3$ with its uniform law. Enumerate the protected inputs as $\{(u_i, t_i)\}_{i=1}^{r_{\mathrm{L}}}$, and set

$$\ell_i := \ell_{u_i,t_i}, \qquad J_i := J_{\zeta - t_i, h}.$$

Let I_0 index the inputs with $u_i = 0$, and let $I_{\neq 0}$ index the remaining inputs.

Definition C.4 (Regular and marked supports)**.** Define the *regular support* and *marked support* by

$$S := \{\omega \in \Omega : c \neq 0,\ \ell_i(\omega) \in J_i \text{ for every } 1 \le i \le r_{\mathrm{L}}\},$$
$$S_{\mathrm{line}} := \left\{ \begin{array}{l} (a_*, c, 0) \in \Omega : c \neq 0, \\ \qquad\qquad \ell_i(a_*, c, 0) \in J_i \text{ for every } i \in I_{\neq 0} \end{array} \right\}.$$

Regular maps keep both colors interior. Marked maps do so off the line; on the line their primary is a_* and their companion is interior by marked-color admissibility. Every protected second input remains available, as required by the joint law in Section 8.5.6.

The complete public law. Fix public data satisfying Section 8.5.3. All densities below are relative to $\mathsf{U}(\Omega)$. Let $\mathsf{Q}_{\mathrm{reg}}$ be the regular-branch density constructed below, with exact rational coordinates in Section C.5.1. Define the uniform marked-branch density by

$$\mathsf{Q}_*(\omega) := \frac{q^3 \mathbf{1}_{S_{\mathrm{line}}}(\omega)}{|S_{\mathrm{line}}|}.$$

The flag B_{mark} equals 0 on the regular branch and 1 on the marked branch. The joint flag and map law is

$$\mathsf{P}_{\mathrm{fc}}(0,\omega) := \frac{(1-\alpha)\mathsf{Q}_{\mathrm{reg}}(\omega)}{q^3},$$
$$\mathsf{P}_{\mathrm{fc}}(1,\omega) := \frac{\alpha \mathsf{Q}_*(\omega)}{q^3}.$$

Its flag marginal is $\mathsf{Bernoulli}(\alpha)$; write Π_{iso} for its map marginal.

We construct the regular-branch density and prove the pair law in Theorem 8.9, using the finite model, supports, and marked density specified above. Fix public data satisfying Section 8.5.3; Section C.5.1 gives the rational coordinates. Expectations over coefficient triples use the unrestricted uniform law unless otherwise specified.

Affine two-point counts. Under the unrestricted uniform law on coefficient triples, for distinct indices i, i' and any colors b, b',

$$\Pr[\ell_i = b, \ell_{i'} = b'] = q^{-2}.$$

Indeed, the two equations over $\mathbb{F}_q$ are

$$a_0 + cu_i + Ct_i = b,$$
$$a_0 + cu_{i'} + Ct_{i'} = b'.$$

Subtracting them gives one nonzero linear equation in (c, C): distinct protected inputs remain distinct modulo q by the input-separation condition of Section 8.5. It has exactly q solutions, each determining a_0 uniquely. A single evaluation equation has q^2 solutions. If the scale c is also fixed, each of the q choices of shear again determines one intercept. Hence each evaluation is uniform and independent of the scale, and distinct evaluations are pairwise independent. These assertions apply only before excluding zero scale or imposing either safe support; the deletion bounds below account for all such restrictions.

Lemma C.5 (Safe supports and deletion bounds)**.** Under the feasibility and marked-color admissibility conditions of Section 8.5, both supports are nonempty and consist of safe maps. For $b \in J_i$,

$$\Pr(S^c), \quad \Pr(S^c \mid \ell_i = b) \le \delta_{\mathrm{aff}},$$
$$p_* := \Pr(S_{\mathrm{line}} \mid a_0 = a_*, C = 0) = \frac{|S_{\mathrm{line}}|}{q} \ge 1 - \delta_{\mathrm{aff}}.$$

Proof. Each input excludes at most $B_{\mathrm{L}} + K_{\mathrm{L}} - 1 + 2h$ colors. The unrestricted counts above and the zero-scale probability $1/q$ give the unconditional union bound. Conditional on an allowed color at input i, its own color constraint cannot fail; each other evaluation and the scale remain uniform. The same union bound therefore proves the conditional estimate. On the unfiltered marked fiber $a_0 = a_*, C = 0$, each off-line evaluation is uniform as c ranges over all of $\mathbb{F}_q$. A union bound

over its color exclusions and the event $c = 0$ proves the survival estimate; no joint independence on this fiber is needed. Since $\delta_{\mathrm{aff}} < 1$, both supports are nonempty. Regular maps and off-line marked evaluations use trimmed slices. On the marked line, the primary is a canonical color and the companion is interior by the marked-color admissibility condition of Section 8.5. Thus every supported map is safe. □

Definition C.6 (Additive color-function space)**.** Equip the real function space

$$\mathcal{H} := \mathbb{R}1 \oplus \bigoplus_{i=1}^{r_{\mathrm{L}}} \left\{ g_i \circ \ell_i : \begin{array}{c} \mathrm{supp}(g_i) \subseteq J_i, \\ \displaystyle\sum_{b \in J_i} g_i(b) = 0 \end{array} \right\}$$

with the function supremum norm on Ω.

Proof of Theorem 8.9. For a real function h on coefficient triples, define its color-fiber averages and centered color imbalances by

$$\begin{aligned} m_i^h(b) &:= q^{-2} \sum_{\substack{\omega \in \Omega \\ \ell_i(\omega) = b}} h(\omega), \\ (\mathcal{C}_i h)(\omega) &:= \mathbf{1}_{\{\ell_i(\omega) \in J_i\}} \left(m_i^h(\ell_i(\omega)) - \frac{1}{|J_i|} \sum_{b \in J_i} m_i^h(b) \right), \\ \mathcal{C}h &:= q^{-3} \sum_{\omega \in \Omega} h(\omega) + \sum_{i=1}^{r_{\mathrm{L}}} \mathcal{C}_i h. \end{aligned} \tag{6}$$

In the unrestricted uniform L^2 space, $\mathcal{C}_i$ is the orthogonal projection onto the corresponding centered summand of $\mathcal{H}$: first condition on the evaluation, then subtract the average over allowed colors and set the remaining colors to zero. The two-point counts make these summands mutually orthogonal and orthogonal to constants. Thus $\mathcal{C}$ is the orthogonal projection onto $\mathcal{H}$. It preserves mass and centered color moments and restricts to the identity there. For $f \in \mathcal{H}$, define the masked synthesis and deletion operator by

$$\Psi_{\mathrm{leg}} f := \mathbf{1}_S f, \qquad Df := \mathcal{C}(\mathbf{1}_{S^c} f).$$

The actual coefficient response on $\mathcal{H}$ is therefore

$$\mathcal{A}_{\mathrm{leg}} := \mathcal{C}\Psi_{\mathrm{leg}} = I - D.$$

For $h = \mathbf{1}_{S^c} f$, Lemma C.5 bounds $|\mathbb{E}h|$ and every $|m_i^h(b)|$ at an allowed color by $\delta_{\mathrm{aff}} \|f\|_\infty$. Bounding the mean and the two terms per centered component gives

$$\|D\|_{\infty \to \infty} \le (2r_{\mathrm{L}} + 1)\delta_{\mathrm{aff}} = \kappa_{\mathrm{aff}} \le \tfrac{1}{4}.$$

Define the branch-selection odds and marked off-line defect by

$$s_\alpha := \frac{\alpha}{1 - \alpha}, \qquad D_{\mathrm{mark}} := \sum_{i \in I_{\neq 0}} \mathcal{C}_i \mathsf{Q}_*.$$

The defect is the sum of projections of Q_* onto the off-line summands. Its mass and its color-difference moments at line inputs vanish; at each off-line input those moments equal those of Q_*. To cancel these moments, set $\tau_{\mathrm{leg}} := 1 - s_\alpha D_{\mathrm{mark}}$ and solve

$$\mathcal{A}_{\mathrm{leg}} f = \tau_{\mathrm{leg}}.$$

By Theorem 7.1, this equation has a unique solution $f \in \mathcal{H}$. Define the regular-branch density by

$$\mathsf{Q}_{\mathrm{reg}} := \Psi_{\mathrm{leg}} f = \mathbf{1}_S f.$$

For positivity, fix an off-line input. Before excluding $c = 0$ or filtering unsafe colors on the marked fiber, each color has exactly one scale preimage. Thus

$$0 \le m_i^{\mathsf{Q}_*}(b) = q \Pr_{\mathrm{mark}}[\ell_i = b] \le p_*^{-1} \le (1 - \delta_{\mathrm{aff}})^{-1}.$$

This mean vanishes outside J_i, and its average on that set is $q/|J_i| \in [1, p_*^{-1}]$; the upper bound also follows from $|S_{\mathrm{line}}| \le |J_i|$. Subtracting the average in (6) and summing gives

$$\sup D_{\mathrm{mark}} \le \frac{r_{\mathrm{L}} \delta_{\mathrm{aff}}}{1 - \delta_{\mathrm{aff}}}, \qquad \|D_{\mathrm{mark}}\|_\infty \le \frac{r_{\mathrm{L}}}{1 - \delta_{\mathrm{aff}}}.$$

Apply Corollary 7.2 to $\tau_{\mathrm{leg}} = 1 - s_\alpha D_{\mathrm{mark}}$. Using the respective one-sided and absolute bounds gives

$$f \ge 1 - \frac{\kappa_{\mathrm{aff}}}{1 - \kappa_{\mathrm{aff}}} - s_\alpha A_{\mathrm{loc}} \ge \tfrac{1}{3},$$

by $s_\alpha \le 1$ and the feasibility bounds in Section 8.5.3. Reconstruction preserves mass and $\mathbb{E} D_{\mathrm{mark}} = 0$, so $\mathbb{E}[\mathbf{1}_S f] = 1$. This normalizes the masked density; f itself need not have mean one. The marked density is normalized by definition. The equivalent rational coordinate system in Section C.5.1 gives rational solution and branch tables, together with the stated Gram nonsingularity. Its bounded implementation is proved in Section C.5.1.

The color-difference equations, normalization, and support give the regular color probabilities:

$$\Pr_{\mathrm{reg}}[\ell_i = b] = \begin{cases} |J_i|^{-1}, & i \in I_0, \\ |J_i|^{-1} - s_\alpha(\Pr_{\mathrm{mark}}[\ell_i = b] - |J_i|^{-1}), & i \in I_{\neq 0}, \end{cases}$$

for $b \in J_i$. Since $(1 - \alpha)s_\alpha = \alpha$, the off-line mixture is uniform on this set. On the line, the regular primary is uniform there, whereas the marked branch gives the integer loads $X = a_*$ and $Y = b_* - t$. For every supported map, safety determines the second residue uniquely from the first and the prescribed integer pair sum. This proves the exact pair law. All regular colors and all off-line marked colors are interior, as is every marked companion, proving the boundary-activation claim. Off-body carry cancellation follows from Lemma C.1, and attachment safety follows from Lemma C.2 under its stated margin and contribution bounds. Normalization of the branches gives the Bernoulli flag law. All table entries depend only on public data; sampling uses no rejection or correction conditioned on the hidden inputs. If $K_{\mathrm{L}} = 1$, then $I_{\neq 0} = \varnothing$ and $D_{\mathrm{mark}} = 0$, and the same argument applies. □

C.5.1 Sampling the Sharp Line-Mark Flag and Affine Map

We prove the table-computation and bounded-sampling claims of Theorem 8.9 for P_{fc}. All parameters are public and satisfy the feasibility and marked-color admissibility conditions of Section 8.5. Use the supports of Definition C.4, enumerating the marked support separately on the fiber $a_0 = a_*$, $C = 0$. Only the $u \neq 0$ inputs impose trimmed color constraints there. The companion safety test holds on every point of the marked line.

Public support enumeration. Enumerate the q^3 coefficient triples in $\Omega = \mathbb{F}_q^3$ and compute $\ell_i(\omega)$ at every protected input to construct both supports. This takes $O(q^3 r_{\mathrm{L}})$ field operations and membership tests. These are evaluations at publicly listed protected inputs, not at the hidden source. By Lemma C.5, both supports are nonempty, so one affine map in each may be stored as a certified representative.

The finite rational calibration system. For each protected input, choose the least allowed canonical color $b_i^{\mathrm{ref}} := \min J_i$ as reference, and define

$$\varphi_0(\omega) := 1, \qquad \varphi_{i,b}(\omega) := \mathbf{1}_{\{\ell_i(\omega)=b\}} - \mathbf{1}_{\left\{\ell_i(\omega)=b_i^{\mathrm{ref}}\right\}},$$

for $b \in J_i \setminus \left\{b_i^{\mathrm{ref}}\right\}$. Let $\mathcal{J}$ index the constant and these real-valued indicator differences. Their number is

$$|\mathcal{J}| = 1 + \sum_{i=1}^{r_{\mathrm{L}}} (|J_i| - 1) \le 1 + r_{\mathrm{L}}(q-1).$$

For $j, k \in \mathcal{J}$, define the rational matrix and right-hand side by

$$(\mathbf{G}_{\mathrm{L}})_{j,k} := q^{-3} \sum_{\omega \in S} \varphi_j(\omega)\varphi_k(\omega), \qquad (\mathbf{d}_{\mathrm{L}})_0 := 1,$$

$$(\mathbf{d}_{\mathrm{L}})_{i,b} := -\frac{s_\alpha \mathbf{1}_{\{u_i \neq 0\}}}{|S_{\mathrm{line}}|} \sum_{\omega \in S_{\mathrm{line}}} \varphi_{i,b}(\omega).$$

The coefficient vector $\mathbf{z}_{\mathrm{L}}$ is determined by

$$\mathbf{G}_{\mathrm{L}} \mathbf{z}_{\mathrm{L}} = \mathbf{d}_{\mathrm{L}}. \tag{7}$$

The constant and indicator differences form a basis of $\mathcal{H}$. Write

$$f = \sum_{j \in \mathcal{J}} (\mathbf{z}_{\mathrm{L}})_j \varphi_j.$$

In the unrestricted uniform L^2 inner product, the orthogonal projection property of Section C.5 gives

$$\langle \varphi_j, \mathcal{A}_{\mathrm{leg}} f \rangle = \langle \varphi_j, \mathbf{1}_S f \rangle = (\mathbf{G}_{\mathrm{L}} \mathbf{z}_{\mathrm{L}})_j.$$

The marked color-moment identities there give $\langle \varphi_j, \tau_{\mathrm{leg}} \rangle = (\mathbf{d}_{\mathrm{L}})_j$. Both sides of the function response equation lie in $\mathcal{H}$, so testing against its basis makes that equation equivalent to (7). The proved invertibility of the function response therefore makes $\mathbf{G}_{\mathrm{L}}$ nonsingular. The regular density is then evaluated using its masked-function definition in Section C.5. The contraction uses the function supremum norm, not coordinate ℓ_∞ or closeness of $\mathbf{G}_{\mathrm{L}}$ to the identity.

Direct assembly uses at most $O(q^3|\mathcal{J}|^2)$ arithmetic operations. Every Gram entry is an integer of absolute value at most q^3, divided by q^3. Every nonconstant right-hand-side entry is $-s_\alpha$ times an integer average over S_{line}, or zero. Here $s_\alpha = \alpha/(1-\alpha)$ has polynomial bit length in the rational input, and $1 \le |S_{\mathrm{line}}| \le q$. Thus all entries are rational with polynomial bit length.

Fraction-free elimination with row interchanges to select nonzero pivots computes the unique solution [Bar68, Section V]; the original leading principal minors need not be nonzero. Clearing denominators in the polynomial-size matrix and right-hand side gives integer entries of polynomial bit length. Minors of the augmented matrix, including after row permutations, control the intermediate integers. The determinant bound therefore gives polynomial bit lengths for these integers and every solution numerator and denominator. Polynomially many arithmetic operations and exact divisions give polynomial bit time, and the branch-density formulas of Sections C.5 and 8.5.4 have polynomial bit length at every map. Their nonnegativity and normalization were established in Section C.5; no numerical positivity test or clipping is used.

The normalized branch tables have probabilities $\mathsf{Q}_{\mathrm{reg}}(\omega)/q^3$ and $\mathsf{Q}_*(\omega)/q^3$ on their respective supports. Multiplying these by $1-\alpha$ and α gives a finite rational table for the complete flag–affine-map law, for any admissible rational $0 \le \alpha \le 1/2$. These probabilities need not be dyadic; bounded fair-bit sampling therefore approximates the exact rational table.

Fixed random bits and branch-preserving rounding. For $0 < \varepsilon_{\mathrm{c}} < 1$, choose an integer b_{s} such that

$$(1 + |S| + |S_{\mathrm{line}}|)2^{-b_{\mathrm{s}}} \le \varepsilon_{\mathrm{c}}.$$

For the flag draw, replace α by

$$\widehat{\alpha} = 2^{-b_{\mathrm{s}}}\lfloor 2^{b_{\mathrm{s}}}\alpha \rfloor.$$

Both branch tables remain those computed at the original rational α. Round each probability of each normalized branch table down to a multiple of $2^{-b_{\mathrm{s}}}$, and assign all remaining mass to the stored representative of that same branch support. Each rounded table remains normalized and nonnegative. Its SD from the exact table is at most its support size times $2^{-b_{\mathrm{s}}}$: the total mass rounded down is at most that amount, and only this mass is moved to the representative. The construction never clips a negative exact table entry.

Use b_{s} unbiased bits to draw the flag with probability $\widehat{\alpha}$, and draw the selected table with exactly b_{s} further unbiased bits, using its integer cumulative probabilities. Both bit blocks are drawn even for a deterministic branch. Coupling the flag and then the selected table bounds the complete flag–map error by the prescribed rounding budget of Section C.5.1, and hence by ε_{c}. Every output in either branch remains in its certified support, including all rounding-remainder outcomes. Nonzero scale, protected carry cancellation, and the boundary-activation identity therefore hold exactly for the bounded implementation. The exact marginal identities hold for the ideal rational law; their numerical perturbation is covered by this complete-pair SD allowance.

Sharp-strength implementation. For $\alpha = 1/2$, one unbiased bit samples the flag exactly. The complete-pair error is then at most $\frac{1}{2}(|S| + |S_{\mathrm{line}}|)2^{-b_{\mathrm{s}}}$, so the extra flag-rounding allowance can be omitted.

The bounds are uniform over admissible public histories. Enumeration, exact elimination, table construction, and sampling take time polynomial in the numerical values of q and r_{L}, the public rational input bit lengths, and $\log(1/\varepsilon_{\mathrm{c}})$. The enumeration depends polynomially on q, rather than $\log q$. In the assembly, $q \le n_{\mathcal{B}} \le d$ and $2(K_{\mathrm{L}} - 1) < q$ imply $r_{\mathrm{L}} = O(q^2)$; hence these bounds are polynomial in the overall public dimension and precision lengths. All computations and random bits are independent of the hidden source.

Computing the public leg parameters. Given the core, including the attachment center z_{R}, Algorithm 1 computes and samples one leg's public affine data. It calls Algorithm 2 for the exact branch laws and keeps all finite-bit outputs in their supports, preserving safety on protected inputs. If the specification instead gives a Gaussian width $\sigma > 0$ and a desired local error $\varepsilon_{\mathrm{leg}} \in \mathbb{Q} \cap (0, 1)$, a sufficient allocation is

$$\varepsilon_{\mathrm{c}} := \frac{\varepsilon_{\mathrm{leg}}}{2}, \qquad K_{\mathrm{L}} := \max\left\{1, \left\lceil \sigma\sqrt{\frac{\log(8/\varepsilon_{\mathrm{leg}})}{\pi}} \right\rceil\right\}.$$

Any integer at least this cutoff gives the same error bound, provided the feasibility tests of Algorithm 1 pass; a preassigned cutoff and sampling budget may instead be supplied directly. The shift bound B_{L} can be set to $|\zeta|$ for one fixed shift, or kept as a common bound for a family of legs. Within the two algorithms, $\mathsf{UniformBits}(b_{\mathrm{s}})$ is the integer encoded by b_{s} fresh independent fair bits, and $\mathsf{RationalSolve}$ solves a rational linear system exactly. Lexicographic order on affine maps is taken in their canonical coordinates.

Algorithm 1 Public leg parameters

Require: $q \in \mathbb{Z}_{\geq 3}$, $\mathsf{IsPrime}(q)$; $K_{\mathrm{L}} \in \mathbb{Z}_{>0}$, $h, B_{\mathrm{L}} \in \mathbb{Z}_{\geq 0}$, $\zeta, z_{\mathrm{R}} \in \mathbb{Z}$, $|\zeta| \leq B_{\mathrm{L}}$; $a_* \in \{0, \ldots, q-1\}$; $\alpha \in \mathbb{Q} \cap [0, 1/2]$, $\varepsilon_{\mathrm{c}} \in \mathbb{Q} \cap (0, 1)$.

$r_{\mathrm{L}} := (2K_{\mathrm{L}} - 1)^2$; $b_* := q - 1 + \zeta - a_*$

$\delta_{\mathrm{aff}} := [r_{\mathrm{L}}(B_{\mathrm{L}} + K_{\mathrm{L}} - 1 + 2h) + 1]/q$; $\kappa_{\mathrm{aff}} := (2r_{\mathrm{L}} + 1)\delta_{\mathrm{aff}}$

if $2(K_{\mathrm{L}} - 1) \geq q \ \vee \ B_{\mathrm{L}} + K_{\mathrm{L}} - 1 + 2h \geq q \ \vee \ \kappa_{\mathrm{aff}} > 1/4$ **then**

  return $\perp$

end if

$$A_{\mathrm{loc}} := \frac{r_{\mathrm{L}}}{1 - \delta_{\mathrm{aff}}} \left(\delta_{\mathrm{aff}} + \frac{\kappa_{\mathrm{aff}}}{1 - \kappa_{\mathrm{aff}}} \right)$$

if $A_{\mathrm{loc}} > 1/3 \ \vee \ b_* \notin [h + K_{\mathrm{L}} - 1, q - h - K_{\mathrm{L}}]$ **then**

  return $\perp$

end if

$(S, \mathsf{Q}_{\mathrm{reg}}, S_{\mathrm{line}}, \mathsf{Q}_*) := \mathsf{LegTables}(q, K_{\mathrm{L}}, h, \zeta, a_*, \alpha)$

$b_{\mathrm{s}} := \lceil \log_2((1 + |S| + |S_{\mathrm{line}}|)/\varepsilon_{\mathrm{c}}) \rceil$; $J_{\mathrm{samp}} := \mathsf{UniformBits}(b_{\mathrm{s}})$

$B_{\mathrm{mark}} := \mathbf{1}_{\{J_{\mathrm{samp}} < \lfloor 2^{b_{\mathrm{s}}} \alpha \rfloor\}}$

if $B_{\mathrm{mark}} = 0$ **then**

  $\omega := \mathsf{DyadicSample}(S, \mathsf{Q}_{\mathrm{reg}}/q^3, b_{\mathrm{s}})$

else

  $\omega := \mathsf{DyadicSample}(S_{\mathrm{line}}, \mathsf{Q}_*/q^3, b_{\mathrm{s}})$

end if

$(a_0, c, C) := \omega$; $(a, \widetilde{c}, \widetilde{C}) := ([a_0]_q, [c]_q, [C]_q)$

Form $\mathbf{R}_{\mathrm{L},\omega}$, $\mathbf{M}_{\mathrm{L},\omega}$, and $\mathbf{\Delta}^{\mathrm{loc}}_{\mathrm{BS,L}}$ using Sections 8.2 and 8.5, with attachment entry z_{R} in Section 8.7.

return $(B_{\mathrm{mark}}, \omega, \mathbf{R}_{\mathrm{L},\omega}, \mathbf{M}_{\mathrm{L},\omega}, \mathbf{\Delta}^{\mathrm{loc}}_{\mathrm{BS,L}})$

function $\mathsf{DyadicSample}(T, \mathsf{p}, b_{\mathrm{s}})$

  $Q_{\mathrm{samp}} := 2^{b_{\mathrm{s}}}$; $\omega_* := \min_{\mathrm{lex}} T$

  for all $\omega \in T$ **do**

    $n_{\mathrm{s}}(\omega) := \lfloor Q_{\mathrm{samp}} \mathsf{p}(\omega) \rfloor$

  end for

  $n_{\mathrm{s}}(\omega_*) := n_{\mathrm{s}}(\omega_*) + Q_{\mathrm{samp}} - \sum_{\omega \in T} n_{\mathrm{s}}(\omega)$

  $J_{\mathrm{samp}} := \mathsf{UniformBits}(b_{\mathrm{s}})$

  return $\min_{\mathrm{lex}} \left\{ \omega \in T : J_{\mathrm{samp}} < \sum_{\substack{\omega' \in T \\ \omega' \leq_{\mathrm{lex}} \omega}} n_{\mathrm{s}}(\omega') \right\}$

end function

Algorithm 2 Exact leg tables

function $\mathsf{LegTables}(q, K_{\mathrm{L}}, h, \zeta, a_*, \alpha)$
    Enumerate Ω and the protected inputs lexicographically; form both supports of Definition C.4.
    Choose the least allowed color as each reference color and form the indicator differences in Section C.5.1.
    Assemble $\mathbf{G}_{\mathrm{L}}$ and $\mathbf{d}_{\mathrm{L}}$ by the finite sums in Section C.5.1.
    $\mathbf{z}_{\mathrm{L}} := \mathsf{RationalSolve}(\mathbf{G}_{\mathrm{L}}, \mathbf{d}_{\mathrm{L}})$.
    Recover f in the chosen basis and evaluate the branch densities of Sections C.5 and 8.5.4.
    return $(S, \mathsf{Q}_{\mathrm{reg}}, S_{\mathrm{line}}, \mathsf{Q}_*)$
end function

The interval test on b_* is equivalent to the companion condition in Section 8.5, because the second protected input ranges from $-(K_{\mathrm{L}}-1)$ to $K_{\mathrm{L}}-1$. On success, Theorem 8.9 gives complete flag–map error at most ε_{c}, while preserving exact carry cancellation and the pointwise boundary-activation criterion on protected inputs. Together, the error-budget choice above and the comparison of Section 8.8 bound the local output SD from $\mathsf{Q}_{\zeta,a_*,\alpha,h}$ by $\varepsilon_{\mathrm{leg}}$. A return value of $\perp$ records only that one of the sufficient tests has failed. The algorithm performs one exact linear solve and has the polynomial bit complexity established above. All computations and random bits are source-independent. The baseline and mixer use the same sampled map, and assembly counts the shared attachment center once. Its zero-carry guarantee follows from Lemma C.2 under its stated margin and contribution bounds.

D Shared Baseline and Joint Realization

We verify the shared baseline, source comparisons, carry safety, and integer affine realization using the local samplers of Sections C.1.1 and C.5.1 and the placement sampler of Section E.2.

D.1 Shared baseline realization

Proof of Proposition 8.15. By admissibility and the off-body pair sums,

$$\begin{aligned}\mathbf{1}^{\mathsf{T}}\boldsymbol{\Delta}_{\mathrm{BS}} &= \sum_{v\in\mathcal{B}} b_v + \sum_{j=1}^{g}(q-1+\zeta_j)\\ &= \sum_{v\in\mathcal{B}} b_v + g(q-1) + \sum_{j=1}^{g}\zeta_j \equiv 0 \pmod q.\end{aligned}$$

Hence Lemma 6.4 gives lattice membership. Body support is part of admissibility; off-body support follows from Definition 8.8 at zero input. The carry statement is Proposition 8.17, whose proof below uses exactly this safe-center support and local carry lemmas. All draws use public data and fresh randomness without inspecting, conditioning on, or rejecting the hidden source. □

D.1.1 Baseline Data

Let $\widehat{\mathsf{D}}_{\mathrm{BS}}$ be the law returned by the public-core implementation assumed in Section D.5, with ideal target D_{BS} from Section 8.7. That contract requires uniform conditional error $\varepsilon_{\mathrm{BS,joint}}$, bounded running time, and exact admissibility on every execution. Marking probabilities are rational. Any

timeout fallback must run within the same bound, return an admissible tuple for the current public context, and charge its probability to $\varepsilon_{\mathrm{BS,joint}}$; nonempty support alone does not provide a fallback.

After the core draw, use fresh independent public randomness for each complete flag–affine-map pair $(B_{\mathrm{mark}\,j}, \omega_j)$, using Appendix C.5.1. Its conditional error is at most ε_{c} relative to the full joint law P_{fc} of Definition 8.10, uniformly over admissible public histories, including the sampled shift and mark. This error includes the finite-bit Bernoulli draw. Branch zero always returns an affine map in S, and branch one always returns an affine map in S_{line}; in particular every output satisfies Definition 8.8. The complete implemented law $\widehat{\mathsf{P}}_{\mathrm{BS}}$ is on the same tuple of core data and complete flag–map pairs as its ideal target P_{BS} in Definition 8.14. The resulting baseline is the deterministic image specified in Definition 8.1. The g conditional pair errors are separate from the core error and counted once each in Section 8.8. All steps are independent of hidden source and secret-dependent information; none inspects, conditions on, or rejects that information.

D.2 Protected sources and comparison bounds

For the source comparisons, fix only preceding public template and placement data permitted by Section 8.7.2. Share the exact positive, source-independent core law and all complete conditional flag–map draws, retaining their prescribed dependencies. The closure Gaussian and final relabeling remain independent of the core. The initial placement and conditional core retain their prescribed dependence, so conditioning on the core may bias that placement. In the direct-Fourier sampler, the separate body-cycle permutation $\mathbf{P}_{\mathrm{B}}$ is independent of the placement, core, and latent cycle, as required by Section 9.2. The cycle comparison averages over the closure.

Lemma D.1 (Source tail bounds)**.** For the ideal Gaussian inputs and public closure,

$$\begin{aligned}\Pr[\left\|\boldsymbol{\Delta}_{\mathrm{B}}^{\mathrm{loc}}\right\|_{\infty} \geq K_{\mathrm{B}}] &\leq \theta_{\mathrm{B}} := n_{\mathcal{B}} \Pr_{x\sim\mathsf{\mu}_{\sigma}}[|x| > H_{\mathrm{e}}] \\ &\leq 2n_{\mathcal{B}} \exp\left(-\frac{\pi\lceil K_{\mathrm{B}}/2\rceil^2}{\sigma^2}\right).\end{aligned}$$

The leg and queue scalar tails satisfy

$$\begin{aligned}\theta_{\mathrm{L}} &:= \Pr_{u\sim\mathsf{D}_{\mathbb{Z},\sigma}}[|u| \geq K_{\mathrm{L}}] \leq 2\exp\left(-\frac{\pi K_{\mathrm{L}}^2}{\sigma^2}\right), \\ \theta_{\mathrm{Q}} &:= \Pr_{x_{\mathrm{Q}}\sim\mathsf{D}_{\mathbb{Z},\sigma}}[|x_{\mathrm{Q}}| \geq K_{\mathrm{Q}}] \leq \min\left\{1,\, 2\exp\left(-\frac{\pi K_{\mathrm{Q}}^2}{\sigma^2}\right)\right\}.\end{aligned}$$

Consequently,

$$\Pr[\mathsf{E}_{\mathrm{safe}}{}^{c}] \leq \theta_{\mathrm{B}} + 2g\theta_{\mathrm{L}} + \theta_{\mathrm{Q}}.$$

Here θ_{B} is a union-bound budget, rather than the body's exact failure probability. All bounds hold under every permitted placement and final relabeling. If the implemented closure law has statistical error at most ν_{G}, add that amount to the body bound and the joint event bound.

Proof of the scalar tail bounds. The two-sided Chernoff estimate from Lemma 2.9, applied at an integer threshold, gives the leg and queue bounds in Lemma D.1; the queue estimate is also bounded by one. For the body, there are $n_{\mathcal{B}}$ independent Gaussian edge scalars, including the public closure. If every scalar has magnitude at most H_{e}, each cyclic difference has magnitude

at most $2H_{\mathrm{e}} < K_{\mathrm{B}}$. The union bound and $H_{\mathrm{e}} + 1 = \lceil K_{\mathrm{B}}/2 \rceil$ prove the body tail bound in Lemma D.1. Taking the union over these edge scalars, the two coordinates of every leg, and the single queue scalar proves the joint protected-event bound. These events concern source coordinates, so supported placements and relabelings preserve the bounds. A maximal coupling of the public closure to its ideal Gaussian adds at most ν_{G} to either body-containing bound.

The three source blocks admit the following comparisons.

Body. Coupling the edge scalars gives a coupling of the entire body cycle, with failure bound $\varepsilon_{\mathrm{B,clip}}$ satisfying

$$\varepsilon_{\mathrm{B,clip}} \le n_{\mathcal{B}} \Pr_{x \sim \mu_\sigma}[|x| > H_{\mathrm{e}}] = \theta_{\mathrm{B}}. \tag{8}$$

This comparison retains all correlations among the cyclic differences.

Legs. For each fixed admissible core, keep the complete public map and the same second input for the companion and attachment. The implemented local law satisfies

$$\begin{aligned}&\Delta(\mathcal{L}(([X \bmod q]_q, [Y \bmod q]_q, t)), \mathsf{Q}_{\zeta,a_*,\alpha,h}) \\ &\quad \le \varepsilon_{\mathrm{c}} + 1 - (1-\theta_{\mathrm{L}})^2 \le \varepsilon_{\mathrm{c}} + 2\theta_{\mathrm{L}}.\end{aligned}$$

The bound remains valid after translating and reducing the attachment coordinate, or taking a histogram. Map-sampling error adds no carry failure because each branch support is preserved.

Queue. For a fixed queue subset, and after averaging any common source-independent placement law,

$$\Delta\Big(\mathcal{L}(\mathbf{\Delta}_{\mathrm{Q}}), \mathcal{L}\Big(\mathbf{\Delta}_{\mathrm{Q}}^{\mathrm{ref}}\Big)\Big) = \theta_{\mathrm{Q}}.$$

This equality concerns the integer increment. Reduction modulo q and passage to a histogram can only decrease the distance.

Proof of the component comparisons. Independently couple each latent edge scalar, including the public closure, maximally to its protected law:

$$x \sim \mathsf{D}_{\mathbb{Z},\sigma}, \qquad z \sim \mu_{\mathrm{e}}.$$

The couplings fail with probability at most the sum of their tail probabilities. On a match every cyclic difference agrees, proving the stated whole-cycle comparison. The protected cycle retains its telescoping sum zero and the dependence among its vertex differences.

Legs and the shared queue. The two independent inputs of a leg both lie in their protected interval with probability $(1-\theta_{\mathrm{L}})^2$. Their joint conditional law is $\mu_{\mathrm{L}}^{\otimes 2}$, so conditioning has exact statistical distance at the input level

$$1 - (1-\theta_{\mathrm{L}})^2 \le 2\theta_{\mathrm{L}}.$$

Keep the same complete flag–map pair on a match, and use the same second input for the companion and attachment. The exact law of Theorem 8.9 and a coupling of the implemented pair at cost at most ε_{c} give the stated leg comparison. Data processing preserves it under translation, reduction, or taking a histogram. Exact map support preserves carry safety even if this coupling fails.

Conditioning the single queue scalar costs exactly θ_{Q}. For each queue subset, the increment determines its scalar by dividing its total integer charge by q. This recovery also works after averaging any common source-independent placement law. Data processing in both directions therefore gives the exact queue clipping distance stated above, retaining the shared scalar and the dependence between queue and attachments.

Joint comparison and public sampling. Under the exact positive public-core law fixed above, the union bound over the source couplings proves the joint comparison bound of Section 8.8.3. On a match, the entire integer load agrees. Assembly, modular reduction, and the shared final permutation therefore give agreement of the complete endpoint vectors, with all cycle, leg, and queue dependencies retained. A match also places the actual source in $\mathsf{E}_{\mathrm{safe}}$; for an admissible core, Proposition 8.17 then supplies carry safety. This single source cost covers both conclusions. Positivity is required for the probability coupling, and admissibility for the carry and zero-fiber conclusions. Clipping is an analytical comparison and never tests or rejects a hidden source value.

Lemma D.2 (Complete baseline-data sampling error)**.** For the same preceding public data, suppose the implemented core law has statistical error at most $\varepsilon_{\mathrm{BS,joint}}$, and each conditional complete flag–map law has error at most ε_{c}, uniformly over admissible preceding public template and placement data and public histories, respectively. Then

$$\Delta\Big(\widehat{\mathsf{P}}_{\mathrm{BS}}, \mathsf{P}_{\mathrm{BS}}\Big) \le \varepsilon_{\mathrm{BS,joint}} + g\varepsilon_{\mathrm{c}}.$$

Proof of the public sampling bound. Couple the implemented core to its target, conditional on the same preceding public data. On a match couple each complete conditional flag–map pair using its uniform error guarantee. The union bound proves the bound of Lemma D.2. The leg comparison uses these same map errors, charged once in assembly. Support preservation leaves the zero baseline failure probability exact, including every fallback.

D.3 Carry safety and source representatives

Proof of Proposition 8.17. Every body vertex satisfies

$$\Big|\big(\boldsymbol{\Delta}_{\mathrm{B}} + \boldsymbol{\Delta}_{\mathrm{L}} + \boldsymbol{\Delta}_{\mathrm{Q}}\big)_v\Big| \le K_{\mathrm{B}} + K_{\mathrm{L}} + K_{\mathrm{Q}} \le K_{\mathrm{BS}}.$$

Its baseline is in $\mathcal{I}_{\mathrm{safe}}$, so its total load stays in $[0, q)$ and its carry is zero. For each off-body pair, Lemma C.1 gives cancellation of the two carries, using the same complete safe map and its shared second source coordinate. Body and queue increments vanish off the body. Summing these contributions proves the integer equality for the total carry, which survives the final common permutation. By Lemma 6.4, the carry lies in $\mathcal{L}(\mathbf{T})$. For the natural and canonical source lifts, write

$$[\mathbf{x} \bmod q]_q = \mathbf{x} + q\mathbf{z}_{\mathrm{lift}}, \qquad \mathbf{z}_{\mathrm{lift}} \in \mathbb{Z}^d.$$

Their loads differ by $q\mathbf{D}_{\mathbf{T}}\mathbf{z}_{\mathrm{lift}}$, so their canonical carries differ by $\mathbf{D}_{\mathbf{T}}\mathbf{z}_{\mathrm{lift}} \in \mathcal{L}(\mathbf{T})$. Hence changing the lift preserves the carry-membership criterion of Lemma 6.6.

The zero-fiber assertion and statistical bound therefore follow from Lemma 6.6. □

D.4 Integer affine realization

The following table records the dimensions and public dependencies of the realization objects defined in Section 8.2.

Proof of Proposition 8.19. The first $d-1$ columns of $\mathbf{D}_{\xi}(\mathbf{P}_{\mathrm{B}})$ form a spanning-tree incidence matrix: the placed body path spans the body, and each leg joins two new vertices to its specified attachment.

Object	Domain and codomain	Public dependence and construction
$\mathbf{E}_{\mathrm{B}}$, $\mathbf{E}_j$	$\mathbb{Z}^{n_{\mathcal{B}}} \to \mathbb{Z}^d$, $\mathbb{Z}^3 \to \mathbb{Z}^d$	Body labels and ξ; explicit sparse embeddings.
$\mathbf{D}_\xi(\mathbf{P}_{\mathrm{B}})$	$\mathbb{Z}^d \to \mathbb{Z}^d$	Placement and body permutation; explicit incidence columns.
$\mathbf{R}_{\mathrm{src}}$	$\mathbb{Z}^d \to \mathbb{Z}^d$	Complete public leg maps; explicit block diagonal matrix.
$\mathbf{z}_0^{\mathrm{pre}}$	An element of $\mathbb{Z}^d$	Baseline, closure, and body permutation; explicit vector.
$\mathbf{Q}$, β	$\mathbb{Z}^d \to \mathbb{Z}^d$, an element of $\mathbb{Z}^d$	All preceding data and final permutation; explicitly formed for an admissible core.

Table 1: Integer realization objects. Matrix dimensions are codomain dimension by domain dimension. The probability-response matrices in Section 9 act on different, real coefficient spaces.

Hence the general objects of Corollary 6.8 specialize to

$$\begin{aligned}
\overline{\mathbf{D}}_{\mathscr{T}} &= [\mathbf{D}_{\mathscr{T}} \mid \mathbf{D}_{\mathrm{Q}}] = \mathbf{D}_\xi(\mathbf{P}_{\mathrm{B}}),\\
\mathbf{R}_{\mathrm{edge}} &= \operatorname{diag}(\mathbf{I}_{n_{\mathcal{B}}-1}, \mathbf{R}_{\mathrm{L},\omega_1}, \dots, \mathbf{R}_{\mathrm{L},\omega_g}),\\
\overline{\mathbf{R}} &= \operatorname{diag}(\mathbf{R}_{\mathrm{edge}}, 1) = \mathbf{R}_{\mathrm{src}},\\
\mathbf{D}_{\mathbf{T}}^{\mathrm{pre}} &= \mathbf{D}_\xi(\mathbf{P}_{\mathrm{B}})\mathbf{R}_{\mathrm{src}},\\
\mathbf{D}_{\mathbf{T}} &= \mathbf{P}_{\mathrm{glob}}\mathbf{D}_{\mathbf{T}}^{\mathrm{pre}}.
\end{aligned}$$

Multiplying the source blocks gives, respectively, the zero extension of $\mathbf{P}_{\mathrm{B}}\mathbf{D}_{\mathrm{B}}\mathbf{x}_{\mathrm{B}}$, the increments $\mathbf{E}_j\mathbf{D}_{\mathrm{L}}\mathbf{R}_{\mathrm{L},\omega_j}\mathbf{x}_{\mathrm{L},j}$ from the leg construction of Definition 8.10, and $\mathbf{D}_{\mathrm{Q}}x_{\mathrm{Q}}$. The closure term in the matrix assembly of Section 8.2 is exactly the zero extension of $\mathbf{P}_{\mathrm{B}}\mathbf{a}_{\mathrm{B}}$ from Definition 8.2. The assembly definition in Section 8.2 therefore gives

$$\mathbf{z}_0 = \mathbf{P}_{\mathrm{glob}}(\boldsymbol{\Delta}_{\mathrm{BS}}^{\mathrm{pre}} + \mathbf{E}_{\mathrm{B}}\mathbf{P}_{\mathrm{B}}\mathbf{a}_{\mathrm{B}}).$$

The assembled load adds the three increments to the shared baseline before the common relabeling.

The queue column has total charge q. The shared baseline satisfies the charge condition of Definition 8.13, and the closure offset has zero charge. These are precisely the lattice hypotheses of Corollary 6.8; the final permutation preserves them because it preserves coordinate sums. That corollary therefore gives integral matrix and offset with the dimensions in Table 1, and determinant $\pm\det\mathbf{R}_{\mathrm{edge}}$. The block diagonal form above and the triangular local mixers give the scale product in the determinant formula of Proposition 8.19. Each scale is nonzero modulo q; the final permutation contributes only a sign. All inverse products are evaluated over $\mathbb{Q}$, as required by the general incidence construction. □

D.5 Gadget Assembly and Realization

We prove Lemma 8.20 for the affine data and sampling assumptions of Section 8.9, using the bounded implementations in Sections C.1.1, C.5.1, D.1.1 and E.2.

Sampling assumptions and numerical budgets. Assume the geometry and safety margins of Sections 8.1 and 8.3, the feasibility and marked-color admissibility conditions of Section 8.5.3 for

every leg, and an admissible core law D_{BS} with nonempty support as in Definition 8.13. Require

$$q \text{ odd prime}, \qquad n_{\mathcal{B}} \geq q, \qquad \sigma^2 \in \mathbb{Q}, \qquad \sigma^2 \geq 2.$$

The public template and core samplers must satisfy two conditions.

(i) They are source-independent and run in worst-case polynomial time in d, the binary public-input lengths, and the requested precision lengths.

(ii) The core sampler has statistical error at most $\varepsilon_{\mathrm{BS,joint}}$ relative to D_{BS}, uniformly over admissible preceding public template and placement data. Every output, including any fallback, has polynomial encoding length and satisfies all admissibility constraints exactly.

Fix $\nu_{\mathrm{num}} \in \mathbb{Q} \cap (0,1)$, rational $\nu_{\mathrm{G}}, \nu_{\mathrm{perm}} \in (0,1)$, and $\varepsilon_{\mathrm{c}} \in \mathbb{Q} \cap (0,1)$, with

$$\nu_{\mathrm{G}} + \nu_{\mathrm{perm}} \leq \nu_{\mathrm{num}}.$$

Here ν_{G} bounds the total variation error of the public closure Gaussian sampler. The quantity ν_{perm} bounds the joint error of the two local permutations and the common relabeling.

The carry-failure budget is defined in Section 8.9; the core and map errors are charged separately from the source comparison.

Finite-sampling interface. Use the prescribed public template and admissible core baseline laws of Sections 8.1 and 8.7, retaining all their permitted dependence on preceding public data. The bounded samplers of Sections C.1.1 and E.2 return an integer closure within ν_{G} of $\mathsf{D}_{\mathbb{Z},\sigma}$ and valid permutations with joint error at most ν_{perm} from the three independent uniform laws. The closure's zero fallback is integral, and placement preserves the disjointness of Lemma E.3. These draws use fresh independent randomness, independent of the hidden source, with the independence prescribed in Sections 8.1 and 8.7.2. The core retains its prescribed conditional law given preceding template and placement data, and its tests never read the body-cycle or final permutation. Section D.1.1 supplies the complete conditional flag–map extension with per-leg error at most ε_{c}, uniformly over admissible public histories, and exact branch support. The core error remains $\varepsilon_{\mathrm{BS,joint}}$ under Section D.5.

Proof of Lemma 8.20. Exact affine validity. Every supported placement is valid, every returned leg map has nonzero scale, and every returned core is admissible. The closure sampler returns an integer, including on its zero fallback. Hence Proposition 8.19 gives integral affine data with the stated determinant on every execution. The public sampling contract above makes the complete matrix and offset source-independent; their assembly is exactly that of Section 8.2.

Protected carry. By Proposition 8.17, failure of the zero-fiber conclusion is contained in the complement of the protected event. Coupling the implemented public closure to its ideal law adds ν_{G} to the joint protected-event bound of Section 8.8, which is exactly the carry budget in Section 8.9. Baseline and flag–map sampling preserve safe support, and valid permutations preserve the protected bounds even for biased finite-bit draws. Their errors add no carry-failure term, proving the stated zero-fiber failure bound.

Bounded implementation and numerical errors. The cited local samplers have deterministic work and bit caps and polynomial output length. Their exact leg-table counting and linear algebra are polynomial in the numerical modulus, protected-input count, and public rational bit lengths; since $q \leq n_{\mathcal{B}} \leq d$ and the leg-feasibility conditions of Section 8.5 bound K_{L} by $O(q)$, this is polynomial in the overall dimension. The core contract and permutation sampler give the

same bounded-time guarantees for the remaining draws. Exact integer and rational linear algebra compute the realization of Section 8.2 in polynomial time and bit length. Every branch returns valid affine data without reading the hidden source.

Couple the closure and the three permutations to their ideal laws at cost at most $\nu_{\mathrm{G}} + \nu_{\mathrm{perm}}$. Conditional on matching placements and the same preceding public data, Lemma D.2 couples the core and complete flag–map pairs at cost at most $\varepsilon_{\mathrm{BS,joint}} + g\varepsilon_{\mathrm{c}}$. Thus the numerical error in the complete public data is at most $\nu_{\mathrm{num}} + \varepsilon_{\mathrm{BS,joint}} + g\varepsilon_{\mathrm{c}}$. This coupling retains the prescribed dependencies and does not require placements to remain uniform conditional on the core. Applying the same independent source and affine assembly cannot increase its error. Source truncation is accounted for separately by the protected comparison of Section 8.8.3. For every fixed accuracy order c, all these numerical tolerances can be chosen inverse-polynomially while retaining a uniform worst-case polynomial implementation. The public $q, d, N, \sigma, \mathbf{T}$ and the d hidden source coordinates remain fixed. □

The selected parameters and protected cutoffs. For the fixed family of Section 9.9, Section H.4 verifies feasibility and negligible source tails for sufficiently large λ. The calibrated sampler of Algorithm 3 meets the admissible-sampling hypotheses in its certified regime and preserves exact affine validity on the remaining finite prefix, without computing an asymptotic accuracy threshold. Its fiber bound in Theorem 9.9 uses the fully protected reference: one source-coupling cost covers both endpoint disagreement and unsafe carries, so the separate carry budget above is not added again.

E Fourier Targets and Protected Reference Tools

E.1 Direct Fourier Targets and Energy Identities

Use the characters and mass-one real signed laws of Section 9.1.

E.1.1 The ambient uniform target

Lemma E.1 (Ambient Fourier energy and distributional error)**.** Let P be a mass-one real signed law on $\mathbb{F}_q^d$, and let $\mathsf{f} = \mathsf{P}/\mathsf{U}\left(\mathbb{F}_q^d\right)$. Then

$$\|\mathsf{f} - 1\|^2_{L^2(\mathsf{U}(\mathbb{F}_q^d))} = \sum_{\mathbf{a}\neq 0} |\widehat{\mathsf{P}}(\mathbf{a})|^2,$$

$$\frac{1}{2}\|\mathsf{P} - \mathsf{U}\left(\mathbb{F}_q^d\right)\|_1 \leq \frac{1}{2}\left(\sum_{\mathbf{a}\neq 0} |\widehat{\mathsf{P}}(\mathbf{a})|^2\right)^{1/2}.$$

For a probability law, the second inequality bounds statistical distance.

Proof of Lemma E.1. The ambient characters form an orthonormal basis under $\mathsf{U}\left(\mathbb{F}_q^d\right)$. Mass one gives zero-frequency coefficient one. Parseval with this coefficient removed proves the energy identity, and Cauchy–Schwarz gives the distance bound. □

E.1.2 The exact support-energy identity

Use the character polynomial, ordered coefficient tensors, and norm of Definition 9.1. For $s_{\mathrm{F}} > 0$, define the weighted support energy by

$$\mathcal{E}_{s_{\mathrm{F}}}(\mathsf{P}) := \sum_{\mathbf{a}\neq 0} s_{\mathrm{F}}^{2|\operatorname{supp}\mathbf{a}|} |\widehat{\mathsf{P}}(\mathbf{a})|^2.$$

Lemma E.2 (Character polynomial and support energy)**.** Let P be an exchangeable real signed law of mass one on $\mathbb{F}_q^d$. For every $0 \le r \le d$,

$$\sum_{|\operatorname{supp}\mathbf{a}|=r} |\widehat{\mathsf{P}}(\mathbf{a})|^2 = \frac{r!}{(d)_r} \|T_r\|_2^2. \tag{9}$$

For every $\rho > 0$, with $s_{\mathrm{F}} = \rho/\sqrt{e}$,

$$\mathcal{E}_{s_{\mathrm{F}}}(\mathsf{P}) \le \mathfrak{F}_\rho(\mathsf{P})^2 - 1.$$

Proof of Lemma E.2. For a fixed ordered nonzero frequency tuple, exchangeability makes the corresponding marginal Fourier coefficient independent of the chosen ordered distinct vertices. The coefficient tensor is therefore $\binom{d}{r}$ times that marginal tensor. The support energy in Lemma E.2 is $\binom{d}{r}$ times the squared norm of the marginal tensor. Dividing by the square of the binomial coefficient proves the identity, including degrees zero and d.

For the all-degree comparison, the geometric mean of the largest r elements of $\{1,\dots,d\}$ is at least that of the entire set. Also

$$\log(d!) = \sum_{j=1}^{d} \log j \ge \int_1^d \log x \,\mathrm{d}x \ge d\log d - d.$$

Consequently $(d)_r \ge (d/e)^r$. Substitute this inequality into (9), multiply by s_{F}^{2r}, and sum the positive degrees. This bound holds at every degree. □

Proof of Lemma 9.2. Low-support cancellation removes all nonzero frequencies of support size at most R. Since $s_{\mathrm{F}} > 1$, Lemma E.2 gives

$$\begin{aligned}\sum_{\mathbf{a}\neq 0} |\widehat{\mathsf{P}}(\mathbf{a})|^2 &\le s_{\mathrm{F}}^{-2(R+1)} \mathcal{E}_{s_{\mathrm{F}}}(\mathsf{P}) \\ &\le s_{\mathrm{F}}^{-2(R+1)} (\mathfrak{F}_\rho(\mathsf{P})^2 - 1).\end{aligned}$$

Lemma E.1 supplies the first inequality in the claimed conclusion. Substituting this energy bound and taking square roots gives the second. For a probability law, half the ℓ_1 distance is statistical distance. □

Complex coefficient convention. Our complex Hilbert inner product is conjugate linear in the first argument and linear in the second. The character-polynomial expansion in Definition 9.1 uses a bilinear coefficient contraction, equivalently the Hilbert pairing of $\overline{T_r}$ with $\mathbf{z}^{\otimes r}$. Real channels obey $\widehat{\mathsf{P}}(-\mathbf{a}) = \overline{\widehat{\mathsf{P}}(\mathbf{a})}$. All subsequent inverses act on real probability or signed profile tables, so their physical weights are real.

E.2 Placements and Relabeling

Use the placement, free vertices, and control pools of Section 9.2.

Lemma E.3 (Disjoint placement and control pools)**.** Every sampled placement satisfies $A_{\mathrm{att}} \cap S_{\mathrm{Q}} = \varnothing$ and the pool counts of Section 9.2. Under the ideal joint placement law, the queue is a uniform q-subset, and the ordered attachments are conditionally uniform without replacement in its complement. The attachment tuple also has a uniform-without-replacement unconditional marginal. For every positive integer R, any selection requiring at most R distinct controls from either pool is available when $g \geq R$ and $n_{\mathcal{B}} - q - 2g \geq R$. The placement and compensators admit bounded implementations whose running time is polynomial in the public dimension and the requested precision length.

Proof of the placement law and pool claims. We verify Lemma E.3 for the joint permutation of Section 8.1. Fix the public order on the body. For rank decoding, write $S_{\mathrm{Q}}^{(0)}$ for the last q reference vertices and $A_{\mathrm{att}}^{(0)} = \left\{R_j^{(0)}\right\}_{j=1}^{g}$ for the first g, in labeled order. Their images are

$$S_{\mathrm{Q}} = \mathbf{P}_{\mathrm{place}}(S_{\mathrm{Q}}^{(0)}), \qquad v_{\mathrm{R}\,j} = \mathbf{P}_{\mathrm{place}}(R_j^{(0)}).$$

The reference sets are disjoint by the geometric requirements, and permutations preserve their disjointness.

For a queue set S of size q and an ordered distinct attachment tuple $(a_1, \ldots, a_g)$ in its complement, the exact placement probability is

$$\Pr\begin{bmatrix} S_{\mathrm{Q}} = S, \\ (R_j)_{j=1}^{g} = (a_j)_{j=1}^{g} \end{bmatrix} = \frac{1}{\binom{n_{\mathcal{B}}}{q}(n_{\mathcal{B}} - q)_g}.$$

Indeed, there are $q!(n_{\mathcal{B}} - q - g)!$ permutations realizing those images, out of $n_{\mathcal{B}}!$ permutations. Summing over the $(n_{\mathcal{B}} - q)_g$ permitted attachment tuples gives a uniform queue subset. Conversely, a fixed attachment tuple permits $\binom{n_{\mathcal{B}} - g}{q}$ queue sets, so its unconditional probability is $1/(n_{\mathcal{B}})_g$. The conditional attachment law is uniform in the queue complement; the queue and attachments are dependent because they are disjoint.

Removing the queue and attachments leaves the free vertices defined in Section 9.2. Assigning the first g of them to the labeled legs leaves the ordinary-body pool in Section 9.2. The geometry inequality in Section 8.1 provides distinct compensators and a nonempty ordinary-body pool. Each pool supports every selection of at most R distinct controls under the two size inequalities in Lemma E.3.

The independent body-cycle permutation moves the path columns and closure together, and the independent final permutation relabels the whole load. Both obey Section 9.2. Distributional averages use Lemma E.3 with the prescribed conditional core law; conditioning on the core need not preserve queue uniformity.

Bounded implementation. Approximate the three independent uniform permutation laws using independent fixed-length random bit strings. Choose

$$b_{\mathrm{perm}} := \lceil \log_2(3d!/\nu_{\mathrm{perm}}) \rceil .$$

For each of the two local permutations and the final common relabeling, reduce a uniform b_{perm}-bit integer modulo the number of permutations of the relevant vertex set, and decode this rank into a permutation. Each rank probability differs from uniform by less than $2^{-b_{\mathrm{perm}}}$, so each permutation has statistical error at most $\nu_{\mathrm{perm}}/3$ and their joint error is at most ν_{perm}. Factorials and permutation ranks have $O(d \log d)$ bits; ranking inversion and the fixed bit count are polynomial in d and $\log(1/\nu_{\mathrm{perm}})$. Every output is a valid permutation preserving disjointness of the reference sets. Scanning the free vertices in public order assigns all compensators in $O(n_{\mathcal{B}} + g)$ operations. All exact affine and carry properties are preserved. Core sampling and rejection retain the resulting placement; the body-cycle and final permutations may be drawn after acceptance. Rank decoding enumerates neither permutations nor placements.

E.3 Finite-Support Path Likelihoods and Protected Cycle Correlations

The protected latent edges of Definition 8.18 are independent and supported on $[-H_{\mathrm{e}}, H_{\mathrm{e}}] \cap \mathbb{Z}$. These two properties suffice for the fixed-order likelihood bounds below. Constants may depend on the fixed order. Likelihoods are taken over integer differences until reduction modulo q and are set to zero outside their reference support.

E.3.1 Fixed-order root kernels

Put $\nu_H = \mu_{\mathrm{e}} * \mu_{\mathrm{e}}$, the integer difference marginal by symmetry, and define

$$\eta_{\mathrm{R}} := \sqrt{\frac{4H_{\mathrm{e}} + 1}{n_{\mathcal{B}}}}.$$

Fix r different labeled body positions after an independent uniform body permutation. Let P_{R} be their protected integer joint law, define its likelihood $\mathsf{k}_{\mathrm{R}} = \mathsf{P}_{\mathrm{R}}/\nu_H^{\otimes r}$, and define

$$\mathsf{H}_r = \prod_{i=1}^{r} (\mathrm{Id} - \mathbb{E}_{\nu_H, i}) \mathsf{k}_{\mathrm{R}}.$$

Here the coordinate expectations integrate only the indicated coordinate under the product reference.

Theorem E.4 (Centered finite-support cycle density)**.** For each fixed positive integer r, assume $n_{\mathcal{B}} \geq 2r$. Then

$$\begin{aligned} \mathsf{H}_1 &= 0, \\ \|\mathsf{H}_r\|_{L^2(\nu_H^{\otimes r})} &\leq C_r (4H_{\mathrm{e}} + 1)^{r/2} n_{\mathcal{B}}^{-\lceil r/2 \rceil}, \\ &\leq C_r \eta_{\mathrm{R}}^r. \end{aligned}$$

The same bounds hold after coordinatewise reduction modulo q, relative to $(\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}})^{\otimes r}$. They hold conditionally on any data independent of the body permutation and all latent edge variables.

The proof combines a finite-support path bound with the probability that the selected sites have no isolated vertex.

E.3.2 A finite-support path bound

Lemma E.5 (Finite-support path likelihood)**.** For every positive integer k, let $\mathsf{p}_{\mathrm{path},H}$ be the law of k successive differences of $k+1$ independent samples from μ_{e}. Then

$$\left\| \frac{\mathsf{p}_{\mathrm{path},H}}{\nu_H^{\otimes k}} \right\|_{L^2(\nu_H^{\otimes k})} \le (4H_{\mathrm{e}}+1)^{k/2}.$$

Proof. Index a vector and a family of matrices by integer latent values $x, z \in [-H_{\mathrm{e}}, H_{\mathrm{e}}]$. Set

$$(\mathbf{u}_{\mathrm{path}})_x := \sqrt{\mu_{\mathrm{e}}(x)},$$
$$(\mathbf{K}_{\mathrm{path}}(y))_{x,z} := \sqrt{\mu_{\mathrm{e}}(x)\mu_{\mathrm{e}}(z)}\mathbf{1}_{\{z-x=y\}}.$$

The vector has Euclidean norm one, and $\|\mathbf{K}_{\mathrm{path}}(y)\|_{\mathrm{F}}^2 = \nu_H(y)$, where $\|\cdot\|_{\mathrm{F}}$ is the Frobenius norm. Expanding the matrix product, ordered by increasing i, counts each latent probability exactly once. Thus, for every $\mathbf{y} \in \mathbb{Z}^k$,

$$\begin{aligned}\mathsf{p}_{\mathrm{path},H}(\mathbf{y}) &= \mathbf{u}_{\mathrm{path}}^{\mathsf{T}} \left(\prod_{i=1}^{k} \mathbf{K}_{\mathrm{path}}(y_i) \right) \mathbf{u}_{\mathrm{path}} \\ &\le \prod_{i=1}^{k} \|\mathbf{K}_{\mathrm{path}}(y_i)\|_{\mathrm{F}} = \prod_{i=1}^{k} \nu_H(y_i)^{1/2}.\end{aligned}$$

On the product reference support, the squared path probability divided by the product marginal is therefore at most one. Each difference has at most $4H_{\mathrm{e}}+1$ possible values. Summing over all such tuples bounds the squared likelihood norm by $(4H_{\mathrm{e}}+1)^k$. □

E.3.3 Random placement and centering

Proof of Theorem E.4. Condition first on the selected positions in the underlying cycle. Their induced graph is a disjoint union of paths, since fewer than all cycle vertices were selected. A path of k vertices uses $k+1$ consecutive latent variables. Distinct components use disjoint latent variables, so the conditional likelihood is the tensor product of their path likelihoods. By Lemma E.5, its norm is at most $(4H_{\mathrm{e}}+1)^{r/2}$, uniformly over the selected positions.

If a selected vertex is isolated, its difference is independent of all other selected differences and has marginal ν_H. Centering in that slot makes the whole conditional likelihood zero. A forest on r vertices with no isolated vertices has at most $\lfloor r/2 \rfloor$ components and hence at least $\lceil r/2 \rceil$ edges. For any specified labeled path forest with h edges, expose one vertex per component freely and then expose each adjacent vertex. There are at most two possible cycle neighbors at each such exposure, among at least $n_{\mathcal{B}} - r$ remaining positions. Thus its embedding probability is at most $(4/n_{\mathcal{B}})^h$. There are only a number depending on r of possible labeled induced forests. A union bound gives probability at most $C_r n_{\mathcal{B}}^{-\lceil r/2 \rceil}$ of having no isolated vertex.

As an orthogonal projection, coordinate centering contracts product L^2. Minkowski's inequality multiplies the path likelihood bound by the probability of having no isolated vertex, without a square root. This gives the first bound and also $\mathsf{H}_1 = 0$. The second bound uses $\lceil r/2 \rceil \ge r/2$ and the root-tag definition.

For a deterministic coordinatewise map ϕ, the pushforward likelihood is the conditional expectation of the original likelihood given the mapped coordinates, under the product reference.

This conditional expectation contracts L^2 and commutes with coordinate centering. Applying it to reduction modulo q proves the residue bound. Conditioning on the independent data in the statement preserves the argument. □

An unchanged placement marginal does not suffice for the conditional clause. A core depending on the body permutation or an acceptance event reading the actual closure Gaussian need not satisfy its independence condition.

F Positive Proposal and Response Calibration

F.1 Positive Proposal and First-Response Proofs

We specify the proposal data and verify the positive proposal and first-response estimates of Section 9.

F.1.1 Anchor support and positive proposal laws

Inward anchors and the independent blur. Let $c_{\mathrm{mid}} = (q-1)/2$ and let $\chi_{\mathrm{sign}}(c)$ equal $-1, 0, 1$ below, at, and above c_{mid}. Draw j_{anc} uniformly from $\{a_{\mathrm{anc}}, \dots, 2a_{\mathrm{anc}} - 1\}$. Given c, set $\varsigma = \chi_{\mathrm{sign}}(c)$ off the midpoint, and choose ς fairly from $\{-1, 1\}$ at the midpoint. Define the Markov kernel A by

$$A_c = q - 1 - c + \varsigma j_{\mathrm{anc}}.$$

Writing $\mathbf{A}_{\mathrm{anc}}$ for this kernel in mass coordinates, its smoothed version is

$$\widetilde{\mathsf{A}} = \mathbf{C}_{\mathrm{L}} \mathbf{A}_{\mathrm{anc}}.$$

Draw $z_{\mathrm{pub}} \sim \mathsf{R}_{\mathrm{B}}^{\mathrm{prot}}$ independently, using its supported integer lift, and set

$$a_* = c, \qquad b_* = A_c + z_{\mathrm{pub}}, \qquad \zeta = \varsigma j_{\mathrm{anc}} + z_{\mathrm{pub}}, \qquad \alpha = \tfrac{1}{2}.$$

The blur copies a distribution, not any actual root or closure value. Require

$$q \geq 6a_{\mathrm{anc}} + 1, \quad a_{\mathrm{anc}} \geq 4K_{\mathrm{L}}, \quad a_{\mathrm{anc}} \geq h + 2K_{\mathrm{L}}, \quad B_{\mathrm{L}} \geq 2a_{\mathrm{anc}} + K_{\mathrm{L}}.$$

Here a_{anc} is a positive integer, and the general leg feasibility conditions of Section 8.5 remain in force.

A diffuse profile defined before the primary law. For each fixed $\varsigma \in \{-1, 1\}$, let v_ς be either endpoint marginal of

$$\mathbb{E}_{j_{\mathrm{anc}}, z_{\mathrm{pub}}, t} \mathsf{L}^{\mathrm{trunc}}_{\varsigma j_{\mathrm{anc}} + z_{\mathrm{pub}} - t, h}, \qquad t \sim \mathsf{\mu}_{\mathrm{L}}, \qquad \mathsf{v} = \tfrac{1}{2}(\mathsf{v}_{+1} + \mathsf{v}_{-1}).$$

These profiles use only the anchor and scalar laws. Write $\iota(c) = q - 1 - c$ for reflection.

The balancing target. The profile used in Section 9.3 is

$$\begin{aligned} \mathbf{h}_{\mathrm{base}} = d\mathsf{u} &- (n_{\mathrm{F}} + g\gamma_{\mathrm{cal}})\mathsf{b}_0 \\ &- (g+q)\mathsf{b}_1 - 2g(1-\gamma_{\mathrm{cal}})\mathsf{v}. \end{aligned}$$

Matrix coordinates. Mass profiles are columns; matrix rows index outputs and columns index inputs. The convolution and anchor operators have entries

$$(\mathbf{C}_{\mathrm{B}})_{a,c} = \mathsf{R}_{\mathrm{B}}^{\mathrm{prot}}(a-c), \qquad (\mathbf{C}_{\mathrm{L}})_{a,c} = \mathsf{\mu}_{\mathrm{L}}(a-c),$$
$$(\mathbf{A}_{\mathrm{anc}})_{a,c} = \Pr[A_c = a \mid c].$$

These matrices and $\widetilde{\mathsf{A}}$ act on $\mathbb{R}^q$ and are column-stochastic.

Before choosing parameter sizes, define the losses

$$\begin{aligned} D_{\mathrm{sl}} &:= 2h + 2a_{\mathrm{anc}} + 2K_{\mathrm{L}}, \\ \delta_{\mathrm{base}} &:= \frac{(n_{\mathcal{B}} + g\gamma_{\mathrm{cal}})2K_{\mathrm{BS}}}{g\gamma_{\mathrm{cal}}(q - 2K_{\mathrm{BS}})} + \frac{2(1-\gamma_{\mathrm{cal}})D_{\mathrm{sl}}}{\gamma_{\mathrm{cal}}(q - D_{\mathrm{sl}})}, \\ \delta_{\mathrm{ctr}} &:= \frac{\varepsilon_{\mathrm{anc}}(q - 2K_{\mathrm{BS}})}{a_{\mathrm{anc}}}, \\ \delta_{\mathrm{prim}} &:= \frac{4(n_{\mathcal{B}} + g\gamma_{\mathrm{cal}})K_{\mathrm{BS}}}{g\gamma_{\mathrm{cal}}q} + \frac{4(1-\gamma_{\mathrm{cal}})D_{\mathrm{sl}}}{\gamma_{\mathrm{cal}}q}. \end{aligned}$$

The denominators are positive under the anchor-feasibility conditions above.

Lemma F.1 (Symbolic positive-proposal bounds)**.** Assume the structural relations, anchor-feasibility conditions, and general leg-feasibility conditions of Section 9.3. If

$$\delta_{\mathrm{base}}, \delta_{\mathrm{ctr}} \le \tfrac{1}{2}, \qquad \varepsilon_{\mathrm{anc}} \le 1,$$

then the proposal has the symmetry, atom, support, and averaged-mean guarantees of Theorem 9.3, uniformly over disjoint placements. Every sampled core satisfies the leg-feasibility conditions. Moreover,

$$\begin{aligned} \|\mathsf{\pi}_0\|_\infty \le \frac{d}{g\gamma_{\mathrm{cal}}q}, \qquad \|\mathsf{\pi}_0 - \mathsf{u}\|_1 \le \delta_{\mathrm{prim}}, \\ \gamma_{\mathrm{cal}}\|\widetilde{\mathsf{A}}\mathsf{\pi}_0\|_\infty \le \gamma_{\mathrm{cal}}\left(\frac{4}{q} + \frac{\delta_{\mathrm{prim}}}{a_{\mathrm{anc}}}\right), \\ \|\mathsf{p}_0 - \mathsf{u}\|_{\mathsf{u}^{-1}} + \|\mathsf{u}_{\mathrm{safe}} - \mathsf{u}\|_{\mathsf{u}^{-1}} \\ \le (2+\varepsilon_{\mathrm{anc}})\sqrt{\frac{2K_{\mathrm{BS}}}{q - 2K_{\mathrm{BS}}}} + \varepsilon_{\mathrm{anc}}\sqrt{3 + \frac{q\delta_{\mathrm{prim}}}{a_{\mathrm{anc}}}}. \end{aligned}$$

Proof of Lemma F.1 and Theorem 9.3. Every anchor lies in $[a_{\mathrm{anc}}, q-1-a_{\mathrm{anc}}]$: below the midpoint its lower endpoint is at least $c_{\mathrm{mid}} - 2a_{\mathrm{anc}} + 2 \ge a_{\mathrm{anc}}$, and its upper endpoint is at most $q-1-a_{\mathrm{anc}}$; reflection handles the other side. At the midpoint, either sign gives an anchor in $[c_{\mathrm{mid}} - 2a_{\mathrm{anc}} + 1, c_{\mathrm{mid}} + 2a_{\mathrm{anc}} - 1]$, which lies in the same safe interval because $c_{\mathrm{mid}} \ge 3a_{\mathrm{anc}}$. Consequently $\widetilde{\mathsf{A}}$ is supported on $\mathcal{I}_{\mathrm{safe}}$, and $b_* - t \in C_h$ for every protected second input. Also $|\zeta| < 2a_{\mathrm{anc}} + K_{\mathrm{L}} \le B_{\mathrm{L}}$. The primary and moving companion have the prescribed ordinary pair sum, so the primary belongs to the untrimmed safe slice required by Definition 8.10. Thus every chosen core satisfies the leg-feasibility conditions of Section 8.5. The anchor and smoothed anchor kernels have height at most $1/a_{\mathrm{anc}}$.

The exact populations satisfy $n_{\mathrm{F}} + g + q = n_{\mathcal{B}}$. Consequently the body coefficients in the primary numerator sum to $n_{\mathrm{F}} + g\gamma_{\mathrm{cal}} + g + q = n_{\mathcal{B}} + g\gamma_{\mathrm{cal}}$. Adding the diffuse coefficient gives $d - g\gamma_{\mathrm{cal}}$,

so the primary has mass one. The center has mass one because the anchor is Markov. Exchange symmetry makes the two endpoint marginals of each diffuse slice equal. Reflection exchanges the two sign profiles and preserves the body profiles and v, hence the primary. The anchor intertwines reflection and symmetric scalar convolution preserves it, so the center is symmetric. The symmetric primary law and fair midpoint convention give equal sign probabilities, so the diffuse profile is the prescribed mixture v.

Convolution bounds each body height by $(q - 2K_{\mathrm{BS}})^{-1}$. Each diffuse slice deletes at most D_{sl} colors, so its height is at most $(q-D_{\mathrm{sl}})^{-1}$. Dividing the first equation of (2) by $g\gamma_{\mathrm{cal}}$ and subtracting the expression with every profile replaced by uniform gives $q\boldsymbol{\pi}_0 \geq 1 - \delta_{\mathrm{base}}$. Anchor height and support imply $(1 - \delta_{\mathrm{ctr}})\mathsf{u}_{\mathrm{safe}} \leq \mathsf{p}_0 \leq (1 + \varepsilon_{\mathrm{anc}})\mathsf{u}_{\mathrm{safe}}$. The assumed loss bounds prove the asserted positivity; its upper height bound also uses $q - 2K_{\mathrm{BS}} \geq q/2$ and $\varepsilon_{\mathrm{anc}} \leq 1$.

The total variation distance of the safe uniform from full uniform is $2K_{\mathrm{BS}}/q$; for each diffuse slice it is at most D_{sl}/q. Convolution contracts this distance. Inserting the corresponding ℓ_1 bounds into the first equation of (2) proves $\|\boldsymbol{\pi}_0 - \mathsf{u}\|_1 \leq \delta_{\mathrm{prim}}$. Positivity of the subtracted profiles gives height at most $d/(g\gamma_{\mathrm{cal}}q)$. At most $2a_{\mathrm{anc}} + 2$ primary colors can contribute to one anchor output, and each conditional probability is at most $1/a_{\mathrm{anc}}$. Thus $\|\widetilde{\mathsf{A}}\mathsf{u}\|_\infty \leq 4/q$. Adding $\delta_{\mathrm{prim}}/a_{\mathrm{anc}}$ yields

$$\|\widetilde{\mathsf{A}}\boldsymbol{\pi}_0\|_\infty \leq 4/q + \delta_{\mathrm{prim}}/a_{\mathrm{anc}},$$

Also, $\|\mathsf{u}_{\mathrm{safe}} - \mathsf{u}\|_{\mathsf{u}^{-1}}^2 = 2K_{\mathrm{BS}}/(q - 2K_{\mathrm{BS}})$. For any probability vector, its squared distance from uniform in the uniform-reference norm is at most q times its height minus one. Applying this to $\widetilde{\mathsf{A}}\boldsymbol{\pi}_0$ and using the second equation of (2) with the triangle inequality proves the remaining center bound.

The n_{F} free body vertices have marginal $\mathbf{C}_{\mathrm{B}}\mathsf{p}_0$; the g attachment-only and q queue-only vertices have marginal b_1. The latter equality uses the common widths and cutoffs of the clipped second leg input and the shared queue scalar. By the exact leg kernel, the two off-body coordinates per leg have total mean

$$\gamma_{\mathrm{cal}}\boldsymbol{\pi}_0 + \gamma_{\mathrm{cal}}\mathbf{C}_{\mathrm{B}}\widetilde{\mathsf{A}}\boldsymbol{\pi}_0 + 2(1 - \gamma_{\mathrm{cal}})\mathsf{v}.$$

Multiplying the second equation of (2) by $n_{\mathrm{F}}\mathbf{C}_{\mathrm{B}}$ gives the companion cancellation described in Section 9.3, since $n_{\mathrm{F}}\varepsilon_{\mathrm{anc}} = g\gamma_{\mathrm{cal}}$. Summing the remaining profiles gives the primary-law numerator, which proves the mean claim after averaging the protected source. This proves the symbolic lemma. The direct substitutions in Section H.4 verify its structural and feasibility conditions for the public family, and show that $\delta_{\mathrm{base}}, \delta_{\mathrm{ctr}}, \delta_{\mathrm{prim}}$ and $\varepsilon_{\mathrm{anc}}$ tend to zero. Its hypotheses therefore hold for all sufficiently large L, proving Theorem 9.3. □

F.1.2 The first response

Lemma F.2 (Symbolic first-response bounds)**.** Under the hypotheses of Lemma F.1, the paired lift of every centered primary score is centered and its complete first histogram response is $\gamma_{\mathrm{cal}}\mathbf{T}_1\mathsf{v}_{\mathrm{in}}$. The operator $\mathbf{T}_1$ preserves mass-zero profiles and, in the primary-reference norm of Lemma 9.4,

$$\|\mathbf{T}_1 - \mathrm{Id}\| = O(\gamma_{\mathrm{cal}}{}^{-1}\sqrt{a_{\mathrm{anc}}/q}).$$

For any such parameter family satisfying ${\gamma_{\mathrm{cal}}}^{-1}\sqrt{a_{\mathrm{anc}}/q} = o(1)$, eventually

$$\mathbf{T}_1^{-1} = \mathrm{Id} - \frac{\mathbf{e}\chi^{\mathsf{T}}}{\gamma_{\mathrm{cal}} + \chi^{\mathsf{T}}\mathbf{e}}, \qquad \|\mathbf{T}_1^{-1}\| \le 2.$$

Proof of Lemmas F.2 and 9.4. The primary score is centered by assumption. The compensator score is centered because $\widetilde{\mathsf{A}}$ is column-stochastic and v_{in} has mass zero. Independence of the two proposal inputs then gives a centered paired lift. The sign profile also has mass zero, so $\mathbf{T}_1$ preserves centering.

At fixed primary color the two off-body mean profiles sum to

$$\gamma_{\mathrm{cal}}\delta_c + \gamma_{\mathrm{cal}}\mathbf{C}_{\mathrm{B}}(\widetilde{\mathsf{A}})_{\cdot,c} + 2(1-\gamma_{\mathrm{cal}})\mathsf{v} + \mathsf{E}\chi_{\mathrm{sign}}(c).$$

The primary score gives the corresponding signed combination. The attachment marginal is independent of the primary color, so its first response is zero. The compensator contributes $-\gamma_{\mathrm{cal}}\mathbf{C}_{\mathrm{B}}\widetilde{\mathsf{A}}\mathsf{v}_{\mathrm{in}}$. The centered public score integrates all other coordinate responses to zero even with correlated roots, leaving exact companion cancellation.

By the anchor-feasibility conditions of Section 9.3, $D_{\mathrm{sl}} = O(a_{\mathrm{anc}})$ and $q - D_{\mathrm{sl}} \ge q/3$. Each sign slice differs from uniform in ℓ_1 by $O(a_{\mathrm{anc}}/q)$ and has height $O(q^{-1})$. Hence

$$\|\mathsf{E}\|_1 = O(a_{\mathrm{anc}}/q), \quad \|\mathsf{E}\|_\infty = O(q^{-1}), \quad \|\mathsf{E}\|_{\pi_0^{-1}} = O(\sqrt{a_{\mathrm{anc}}/q}).$$

The dual norm of the sign functional is at most one by Cauchy–Schwarz. The perturbation norm is therefore $O({\gamma_{\mathrm{cal}}}^{-1}\sqrt{a_{\mathrm{anc}}/q})$. Under the lemma's smallness hypothesis, the denominator correction has absolute value at most $\|\mathsf{E}\|_{\pi_0^{-1}} = o(\gamma_{\mathrm{cal}})$. Multiplying the two rank-one formulas proves the inverse formula of Lemma F.2; the perturbation bound gives inverse norm at most two eventually. For the public family, Section H.4 gives ${\gamma_{\mathrm{cal}}}^{-1}\sqrt{a_{\mathrm{anc}}/q} = O(L/t) = O(L^{-5})$. Together with Theorem 9.3, this proves Lemma 9.4. □

F.2 Construction and Bounds for the Graded Response Matrices

Use the parameters and radius of Section 9.9, including t, and fix a labeled disjoint placement from Theorem 9.3. We construct coefficients in $\mathcal{V}_{\le R}$ as product-reference density kernels, following the coordinate convention of Definition F.3. All expectations are uncharged. The independent-root comparison retains the protected marginal $\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}}$, the complete same-source leg channel, and one common queue scalar. The protected-cycle response retains the dependent cycle law. Use the sampling contract of Section 9.2 and the populations and pairings of Section 9.2. Constants may depend on fixed R, but not on placement.

F.2.1 Canonical coefficient algebra and its Hilbert norm

Definition F.3 (Canonical public-control coordinates)**.** Let

$$\begin{aligned} V_{\mathrm{P}} &= \{\mathsf{f} : \mathbb{F}_q \to \mathbb{R} : \mathbb{E}_{\pi_0}\mathsf{f} = 0\}, \\ V_{\mathrm{B}} &= \{\mathsf{f} : \mathcal{I}_{\mathrm{safe}} \to \mathbb{R} : \mathbb{E}_{\mathsf{p}_0}\mathsf{f} = 0\}, \\ \mathcal{V}_r &= \bigoplus_{a_{\mathrm{P}}+b_{\mathrm{B}}=r} \mathrm{Sym}^{a_{\mathrm{P}}}(V_{\mathrm{P}}) \otimes \mathrm{Sym}^{b_{\mathrm{B}}}(V_{\mathrm{B}}), \\ \mathcal{V}_{\le R} &= \mathbb{R} \oplus \bigoplus_{r=1}^{R} \mathcal{V}_r. \end{aligned}$$

Let $\mathrm{coord} : \mathcal{V}_{\le R} \to \mathbb{R}^{m_R}$ be the coordinate isomorphism constructed in Section F.2.1, and put $\mathbf{e}_0 = \mathrm{coord}(1)$.

Positive-degree coefficients are separately symmetric density kernels relative to $\pi_0^{\otimes a_{\mathrm{P}}}\mathsf{p}_0^{\otimes b_{\mathrm{B}}}$, centered in every input slot. The coordinate dimension satisfies $m_R \le (2q+R)^R$; Section F.2.1 supplies the basis and conversion bounds. All estimates use Section 9.9, at fixed R, uniformly over placements. Throughout, coord identifies coefficient kernels with their vectors and coefficient-space endomorphisms with their matrices. Vector and induced operator norms are transported through this same isomorphism.

The disjoint-slot product. The formal product $\diamond$ tensors disjoint slots, symmetrizes each type and truncates above degree R. It differs from physical multiplication, which includes overlaps.

A graded weighted Hilbert norm. For a coefficient vector $\mathbf{f} = \mathrm{coord}(\mathsf{f})$, put $\|\mathbf{f}\|_{\mathrm{simp},\rho} := \mathcal{C}_\rho(\mathsf{f})$, where

$$\mathcal{C}_\rho(\mathsf{f}) := \sum_{a_{\mathrm{P}}+b_{\mathrm{B}}\le R} \rho^{b_{\mathrm{B}}} {m_{\mathrm{P}}}^{-a_{\mathrm{P}}/2} {m_{\mathrm{B}}}^{-b_{\mathrm{B}}/2} \|\mathsf{f}_{a_{\mathrm{P}},b_{\mathrm{B}}}\|_{L^2(\pi_0^{\otimes a_{\mathrm{P}}}\otimes \mathsf{p}_0^{\otimes b_{\mathrm{B}}})}.$$

Add contributions of the same arity before taking their Hilbert norm. The population factors match distinct-index averaging; the body-radius factor matches the Fourier estimate. Lemmas F.4 and F.5 establish submultiplicativity and the retained-input fiber estimates used below. Operator norms below are induced by this norm. It controls physical moments but does not imply pointwise closeness of the weight to one.

Lemma F.4 (Finite-atom comparison)**.** For a kernel of type $(a_{\mathrm{P}}, b_{\mathrm{B}})$, possibly taking values in a Hilbert space, fixing input subsets $I_{\mathrm{P}}, I_{\mathrm{B}}$ gives

$$\sup_{\text{fixed } I_{\mathrm{P}},I_{\mathrm{B}}} \|\mathsf{f}_{a_{\mathrm{P}},b_{\mathrm{B}}}\|_{L^2(\text{remaining})} \le (2q)^{(|I_{\mathrm{P}}|+|I_{\mathrm{B}}|)/2} \|\mathsf{f}_{a_{\mathrm{P}},b_{\mathrm{B}}}\|_2.$$

Proof. In the squared product-reference norm, the summand for any fixed input tuple has probability at least $(2q)^{-|I_{\mathrm{P}}|-|I_{\mathrm{B}}|}$ by Theorem 9.3. Dropping the other nonnegative summands proves the fiber bound, also for a Hilbert-valued kernel. □

Coordinate construction. Coordinates are ordered by total degree and then by hybrid type, with columns indexing inputs and rows outputs. Choose reference colors $c_* \in \mathbb{F}_q$ and $a_* \in \mathcal{I}_{\mathrm{safe}}$.

For each nonreference color use the centered point-mass density bases

$$\varphi_c^{\mathrm{P}} = \frac{\delta_c}{\pi_0(c)} - \frac{\delta_{c_*}}{\pi_0(c_*)},$$
$$\varphi_a^{\mathrm{B}} = \frac{\delta_a}{\mathsf{p}_0(a)} - \frac{\delta_{a_*}}{\mathsf{p}_0(a_*)}.$$

For each nondecreasing tuple of basis indices, average the tensor product over all permutations within the primary slots and, separately, all permutations within the body slots, including repeated permutations. These averaged tensors, together with the scalar unit, form the chosen basis. This defines the coordinate isomorphism coord declared in Definition F.3.

In degree one, the coefficient of φ_c^{P} in f is $\pi_0(c)\mathsf{f}(c)$; centering recovers its value at the reference color. The body formula uses p_0. In several slots, evaluate at nonreference colors and multiply by their product reference probability. These are the ordered tensor coefficients. Multiply by the orbit size o_{sym} of distinct separate permutations to obtain the averaged symmetric coefficients; divide by it to recover the ordered expansion. Iterated centering recovers tuples containing reference colors. A basis tensor is specified by a multiset of at most R labels from the two basis types, which together have at most $2q$ labels. Order these labels and pad with a distinct symbol to obtain a tuple of length R. This injective encoding gives $m_R \le (2q+1)^R \le (2q+R)^R$ for $R \ge 1$, proving the stated dimension bound. The lower atom bounds in Theorem 9.3 make both basis transformations polynomially bounded at fixed order: each uses at most R reference probabilities or their inverses and permutation factors at most $R!$. The Fourier output of $\mathbf{E}_R$ stores the ordered symmetric coefficient tensors of Section 9.1, including degree zero; Lemma F.5 specifies the evaluation maps.

Multiplication entries. For $\mathbf{f} = \mathrm{coord}(\mathsf{f})$ and $\mathbf{g} = \mathrm{coord}(\mathsf{f}_{\mathrm{aux}})$, define

$$\mathbf{M}_R(\mathbf{f})\mathbf{g} = \mathrm{coord}(\mathsf{f} \diamond \mathsf{f}_{\mathrm{aux}}), \qquad \mathbf{M}_{\mathrm{bg}} = \mathbf{M}_R(\mathbf{b}).$$

Let $\varphi_{i_{\mathrm{bas}}}$ range over the chosen basis, including the unit. Its structure constants are defined by

$$\varphi_{i_{\mathrm{bas}}} \diamond \varphi_{j_{\mathrm{bas}}} = \sum_{k_{\mathrm{bas}}} s^{\mathrm{alg}}_{k_{\mathrm{bas}},i_{\mathrm{bas}},j_{\mathrm{bas}}} \varphi_{k_{\mathrm{bas}}}.$$

The resulting matrix entries are

$$(\mathbf{M}_R(\mathbf{f}))_{k_{\mathrm{bas}},j_{\mathrm{bas}}} = \sum_{i_{\mathrm{bas}}} s^{\mathrm{alg}}_{k_{\mathrm{bas}},i_{\mathrm{bas}},j_{\mathrm{bas}}} (\mathbf{f})_{i_{\mathrm{bas}}}.$$

Lemma F.5 (Algebra and leading evaluation)**.** The finite coefficient algebra is commutative and associative, and

$$\mathcal{C}_\rho(\mathsf{f} \diamond \mathsf{f}_{\mathrm{aux}}) \le \mathcal{C}_\rho(\mathsf{f})\mathcal{C}_\rho(\mathsf{f}_{\mathrm{aux}}).$$

Define its leading evaluation by first multiplying a density by its product reference measure, then applying

$$\mathcal{E}_{a_{\mathrm{P}},b_{\mathrm{B}}} = \mathrm{Sym}\left[(\gamma_{\mathrm{cal}}\mathbf{T}_1)^{\otimes a_{\mathrm{P}}} \otimes (\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}}*)^{\otimes b_{\mathrm{B}}}\right].$$

Here Sym averages over all permutations of the output slots. Evaluation gives symmetric tensors of mass-zero profiles in positive degree and is a unital algebra homomorphism, not necessarily

injective. Its Fourier transform is $\mathbf{E}_R$. In coordinates,

$$\mathbf{M}_R(\mathsf{f})\mathbf{M}_R(\mathsf{g}) = \mathbf{M}_R(\mathrm{coord}(\mathsf{f} \diamond \mathsf{f}_{\mathrm{aux}})),$$
$$\mathbf{M}_R(\mathbf{e}_0) = \mathbf{I}.$$

For a positive-degree coefficient define the finite formal operations

$$\exp_\diamond(\mathsf{f}) = \sum_{k=0}^{R} \frac{\mathsf{f}^k}{k!}, \qquad \log_\diamond(1+\mathsf{f}) = \sum_{k=1}^{R} \frac{(-1)^{k+1}}{k}\mathsf{f}^k.$$

All powers use $\diamond$. These operations are mutually inverse, and the logarithm takes products of unit-constant coefficients to sums. Multiplication by a unit-constant coefficient is unit block lower triangular. Its reciprocal and inverse bound are in Proposition F.6. These formal identities do not assert pointwise identities for physical weights:

$$\mathsf{U}J_{\mathrm{in}}(\mathsf{f} \diamond \mathsf{f}_{\mathrm{aux}}) \neq (\mathsf{U}J_{\mathrm{in}}\mathsf{f})(\mathsf{U}J_{\mathrm{in}}\mathsf{f}_{\mathrm{aux}}) \quad \text{in general,}$$

because the product on the right includes overlapping indices.

Proof. Before symmetrization, the product-reference Hilbert norm of a tensor product is the product of its two norms. The population and radius weights multiply because their exponents add. Separate symmetrization averages isometries and therefore contracts the Hilbert norm. Summing over arities proves submultiplicativity; truncation only drops nonnegative terms from the bound. Associativity follows either by averaging all permutations once or by viewing these tensors as symmetric polynomial coefficients. The density-to-profile multiplication also factors over disjoint slots. The slotwise linear maps commute with averaged symmetrization, proving multiplicativity on the full coefficient algebra. In coordinates, associativity gives the multiplication-matrix homomorphism and multiplication by the scalar unit gives the identity. Every product of more than R positive-degree coefficients vanishes. The logarithm and exponential identities thus follow from the corresponding formal power-series identities modulo degree $R+1$, with only finitely many terms per coefficient. Multiplication by a positive-degree coefficient strictly raises degree, giving the asserted triangular form. □

Equalities with character polynomials below implicitly Fourier-transform the evaluated profiles. Evaluation followed by this transform is $\mathbf{E}_R\,\mathrm{coord}$, with the coefficient convention of Definition F.3.

F.2.2 Finite computation in the coefficient algebra

Finite algebra implements the coefficient solve; stability uses Theorem 7.1. The background reciprocal is formed before solving the response system.

Proposition F.6 (Finite coefficient computation)**.** Use the algebra and graded Hilbert norm of Lemma F.5, with coordinates respecting degree. For any unit-constant coefficient B_{bg}, let $\mathbf{M}_{\mathrm{bg}}$ denote multiplication by it. The formal reciprocal and multiplication inverse are

$$\mathsf{B}_{\mathrm{bg}}{}^{-1} = \sum_{k=0}^{R}(1-\mathsf{B}_{\mathrm{bg}})^k, \qquad \mathbf{M}_{\mathrm{bg}}^{-1}\mathbf{e}_0 = \mathrm{coord}(\mathsf{B}_{\mathrm{bg}}{}^{-1}),$$

where every power uses $\diamond$. No smallness assumption is needed for these finite identities. If $\mathcal{C}_\rho(\mathsf{B}_{\mathrm{bg}} - 1) \le \varepsilon_{\mathrm{bg}} < 1$, then

$$\mathcal{C}_\rho({\mathsf{B}_{\mathrm{bg}}}^{-1} - 1) \le \frac{\varepsilon_{\mathrm{bg}}}{1 - \varepsilon_{\mathrm{bg}}}, \qquad \|\mathbf{M}_{\mathrm{bg}}^{-1}\|_{\mathrm{simp},\rho} \le \frac{1}{1 - \varepsilon_{\mathrm{bg}}}.$$

Suppose $\widetilde{\mathbf{A}}_R = \mathbf{I} + \mathbf{U}_R + \mathbf{Z}_R$, where $\mathbf{U}_R$ strictly raises degree, $\mathbf{Z}_R$ does not decrease degree, both annihilate $\mathbf{e}_0$, and

$$\|\mathbf{U}_R\|_{\mathrm{simp},\rho} + \|\mathbf{Z}_R\|_{\mathrm{simp},\rho} \le \eta_{\mathrm{resp}} < 1.$$

The scalar block is the identity, with no coupling to positive degrees. On positive degrees, solve in increasing degree: subtract the response of the already determined lower-degree coefficients, then solve with the current diagonal block. The same small-perturbation bound makes every diagonal block invertible. Under both displayed smallness hypotheses, the solution to $\mathbf{M}_{\mathrm{bg}}\widetilde{\mathbf{A}}_R\mathbf{z} = \mathbf{e}_0$ is unique and is computed using the background reciprocal followed by these degree-block solves. The correction bound and its scalar and degree-one structure for the constructed response are proved in Section F.2.6.

At fixed R, computation uses polynomially many arithmetic operations in the matrix dimension. Section H.5 certifies the data approximations, coordinate conditioning, precision and work bounds for the public family, preserving structural zero blocks and the scalar identity. Nonrational data are approximated rather than computed exactly in finite bits.

Proof. Multiplication by $\mathsf{B}_{\mathrm{bg}} - 1$ is nilpotent, giving the finite reciprocal. Submultiplicativity bounds $\|\mathbf{M}_{\mathrm{bg}} - \mathbf{I}\|$ by $\varepsilon_{\mathrm{bg}}$; applying Theorem 7.1 to $\mathbf{M}_{\mathrm{bg}}$ gives the stated inverse bounds. The reference-residual estimate with reference $\mathbf{e}_0$ also gives the background reciprocal bound. Homogeneous projections contract the norm, so the same response bound makes each diagonal block invertible. The response does not lower degree, so each block equation involves only the current and already solved degrees. Unit-constant background and invertible diagonal blocks suffice for the finite solve; smallness certifies invertibility and stability. Finite products and block elimination use polynomially many arithmetic operations in the matrix dimension. □

Finite block solution. The finite solve of Proposition F.6 starts with scalar correction zero and proceeds in increasing degree $1 \le r \le R$ by

$$(\mathbf{c})_r = -\mathbf{A}_{rr}^{-1}\left((\mathbf{b})_r + \sum_{k=1}^{r-1}(\mathbf{A}_R)_{r,k}(\mathbf{c})_k\right).$$

F.2.3 Concrete background and local transfer estimates

Lemma F.7 (Weighted Hilbert transfer)**.** A local linear operator consumes u primary and v body slots and emits m centered primary output profiles. Suppose its raw density operator has Hilbert norm at most B_2, with the same bound for Hilbert-valued retained inputs. After lifting the new output profiles by ${\gamma_{\mathrm{cal}}}^{-1}\mathbf{T}_1^{-1}$ and storing the resulting densities, its graded operator norm is at most

$$C_R\rho^{-v}{m_{\mathrm{B}}}^{v/2}{m_{\mathrm{P}}}^{(u-m)/2}{\gamma_{\mathrm{cal}}}^{-m}B_2.$$

The same assertion holds before separate symmetrization.

Proof. The output/input weight ratio in $\mathcal{C}_\rho$ gives the radius and pool factors. Each first inverse costs at most two by Lemma 9.4; activation inverses give the remaining factors. Symmetrization

contracts, and retained joint Hilbert arguments are covered by the assumed local operator bound. □

For the public family, the two new-slot factors are

$$\gamma_{\mathrm{cal}}{}^{-1} m_{\mathrm{P}}{}^{-1/2} = O(Lt^{-8}), \qquad \rho m_{\mathrm{B}}{}^{-1/2} = O(L^{1/4} t^{-10}).$$

Consuming a primary slot costs $O(t^8)$; consuming an ordinary slot costs at most $O(t^{10})$. These factors also apply to retained joint inputs.

Independent-root background certificates. For a probability profile μ, write C_μ for the slot-wise projection subtracting total mass times μ. A profile certificate is a specified canonical density coefficient with that leading evaluation. Primary lifts use the first inverse; safe body seeds remain body slots without inverting their convolution.

Fringe and translated safe seeds. Set $h_{\mathrm{safe}} = (q - 2K_{\mathrm{BS}})^{-1}$ and $\mathsf{F}_\partial = h_{\mathrm{safe}} - \mathsf{b}_0$, using the body profile of Section 9.3. The positive-baseline formula and the pointwise height bounds used in Theorem 9.3, specialized as in Section 9.9, give

$$0 \le \mathsf{F}_\partial \le \varepsilon_{\mathrm{anc}} \pi_0, \qquad \sum \mathsf{F}_\partial = O(K_{\mathrm{L}}/q).$$

In the baseline formula, the coefficient of F_∂ is $n_{\mathrm{F}} + g\gamma_{\mathrm{cal}}$. The remaining positive constant is $g\gamma_{\mathrm{cal}} h_{\mathrm{safe}}$; its safe-interval and diffuse-slice losses are each $O(L^{-1})$ relative to this constant by Section F.1.1. Dividing by the fringe coefficient gives the displayed bound for sufficiently large parameters. Since $\mathsf{b}_0 - \mathsf{u} = (\sum \mathsf{F}_\partial)\mathsf{u} - \mathsf{F}_\partial$,

$$\|\mathsf{b}_0 - \mathsf{u}\|_{\pi_0^{-1}} = O(L^{1/2} t^{-6}).$$

The squared Hilbert estimate uses $\sum \mathsf{F}_\partial{}^2/\pi_0 \le \varepsilon_{\mathrm{anc}} \sum \mathsf{F}_\partial$.

The safe projection $\mathsf{B}_{\mathrm{safe}}$ keeps safe colors and replaces every unsafe color by an independent $\mathsf{u}_{\mathrm{safe}}$. For an integer shift t, define cyclic translation of a mass profile μ by

$$(\mathsf{T}_t \mu)(c) := \mu(c - t).$$

For $|t| < 2K_{\mathrm{L}}$, define

$$\begin{aligned} \mathsf{d}_{\mathrm{sh}} &= \mathsf{T}_t \mathsf{u}_{\mathrm{safe}} - \mathsf{u}_{\mathrm{safe}}, & \mathsf{e}_{\mathrm{sh}} &= \mathbf{1}_{\mathcal{I}_{\mathrm{safe}}^c} \mathsf{T}_t \mathsf{u}_{\mathrm{safe}}, \\ \beta_{\mathrm{sh}} &= \sum \mathsf{e}_{\mathrm{sh}}, & \mathsf{k}_{\mathrm{sh}} &= \mathsf{d}_{\mathrm{sh}} - \mathsf{e}_{\mathrm{sh}} + \beta_{\mathrm{sh}} \mathsf{u}_{\mathrm{safe}}, \\ \mathsf{r}_{\mathrm{sh}} &= \mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{e}_{\mathrm{sh}} - \beta_{\mathrm{sh}} \mathsf{b}_0. \end{aligned}$$

These give the exact mass-zero identity

$$\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{T}_t \mathsf{u}_{\mathrm{safe}} - \mathsf{u} = \mathsf{b}_0 - \mathsf{u} + \mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{k}_{\mathrm{sh}} + \mathsf{r}_{\mathrm{sh}}.$$

Only $O(K_{\mathrm{L}})$ boundary colors change, so $\|\mathsf{k}_{\mathrm{sh}}\|_{\mathsf{p}_0^{-1}} = O(Lt^{-4})$. Moreover $\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{e}_{\mathrm{sh}} \le \mathsf{F}_\partial$; hence the residual has the Hilbert bound in the fringe estimate established above. The new-slot costs show that the certificate J_t defined by this recipe has norm $O(L^{3/2} t^{-14})$. For the public family, $a_{\mathrm{anc}} \ge 5K_{\mathrm{L}}$ for $L \ge 2$. Hence every translated anchor column used below stays on $\mathcal{I}_{\mathrm{safe}}$ whenever $|t| < 2K_{\mathrm{L}}$. Two further useful safe/profile estimates are

$$\|\mathsf{T}_{-t} \mathsf{A} \pi_0 - \mathsf{u}_{\mathrm{safe}}\|_{\mathsf{p}_0^{-1}} = O(Lt^{-1}), \qquad \|\pi_0 - \mathsf{u}\|_{\pi_0^{-1}} = O(L^{1/2} t^{-1}).$$

For the first, the anchor applied to uniform input differs only on $O(a_{\mathrm{anc}})$ endpoint or midpoint transition colors; the remaining change is at most $C\sqrt{q/a_{\mathrm{anc}}}\|\pi_0 - \mathsf{u}\|_1$ by Minkowski. Protected translation and safe trimming add $O(\sqrt{K_{\mathrm{L}}/q})$. For the second use $\mathsf{u}/\pi_0 \le 2$ and $\sum(\pi_0 - \mathsf{u})^2/\pi_0 \le C\|\pi_0 - \mathsf{u}\|_1$. The first is retained as a safe body seed; the second has a primary certificate of norm $O(L^{3/2}t^{-9})$.

The first estimate also holds after averaging the protected translate, so the balancing formula of Section 9.3 gives

$$\|\mathsf{p}_0 - \mathsf{u}_{\mathrm{safe}}\|_{\mathsf{p}_0^{-1}} = \varepsilon_{\mathrm{anc}}\|\widetilde{\mathsf{A}}\pi_0 - \mathsf{u}_{\mathrm{safe}}\|_{\mathsf{p}_0^{-1}} = O(t^{-5}).$$

Its body certificate costs $O(L^{1/2}t^{-15})$. Together with the fringe primary certificate this defines a certificate F of norm $O(L^{3/2}t^{-14})$ for $\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{p}_0 - \mathsf{u}$. The same profile has raw primary-reference norm $O(t^{-5})$: convolution from the center reference is bounded by conditional Jensen and $\sup(\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{p}_0)/\pi_0 \le 8$, and its fringe part is smaller.

Lemma F.8 (Diffuse and marked pair certificates)**.** At fixed protected leg scalar t, let $\mathsf{P}_{\mathrm{diff}\,t}$ be the actual diffuse pair, including its random anchor sign, length-a_{anc} interval, and independent public blur. Define

$$\mathsf{S}_{\mathrm{diff}\,t} = \mathsf{C}_{\mathsf{p}_0}^{\otimes 2}(\mathsf{B}_{\mathrm{safe}} \otimes \mathsf{B}_{\mathrm{safe}})\mathsf{P}_{\mathrm{diff}\,t},$$
$$\mathsf{E}_{\mathrm{diff}\,t} = \mathsf{C}_{\mathsf{u}}^{\otimes 2}\mathsf{P}_{\mathrm{diff}\,t} - (\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \otimes \mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} *)\mathsf{S}_{\mathrm{diff}\,t}.$$

Their profile norms in the respective two-slot references are $O(t)$ and $O(Lt^{-2})$. For $\mathsf{A}_{\mathrm{sh}} = \mathsf{T}_{-t}\mathsf{A}$, put

$$\mathsf{P}_{\mathrm{mark}\,t}(c, a) = \pi_0(c)(\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{A}_{\mathrm{sh}c})(a), \quad \mathsf{K}_{\mathrm{mark}\,t} = (\mathsf{C}_{\mathsf{u}} \otimes \mathsf{C}_{\mathsf{p}_0})[\pi_0(c)\mathsf{A}_{\mathrm{sh}c}(a)].$$

Then $\mathsf{K}_{\mathrm{mark}\,t}$ is supported safely in its second slot, has norm $O(t)$ in $\pi_0^{-1} \otimes \mathsf{p}_0^{-1}$, and

$$\mathsf{C}_{\mathsf{u}}^{\otimes 2}\mathsf{P}_{\mathrm{mark}\,t} = (I \otimes \mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} *)\mathsf{K}_{\mathrm{mark}\,t} + (\pi_0 - \mathsf{u}) \otimes (\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{p}_0 - \mathsf{u}).$$

Proof. Conditioning on sign and public blur puts the diffuse pair in a band of $O(qa_{\mathrm{anc}})$ cells, each of mass at most $O((qa_{\mathrm{anc}})^{-1})$. Its uniform-profile norm is therefore $O(\sqrt{q/a_{\mathrm{anc}}})$. Safe projection has column sums one and row sums at most two; its marginals have height $O(q^{-1})$. Projection and centering are bounded in this profile norm, and $\mathsf{p}_0 \ge \mathsf{u}_{\mathrm{safe}}/2$ on safe support. This proves the safe-seed bound.

Protected translations shift the length-a_{anc} sum-offset interval by $O(K_{\mathrm{L}})$, costing $O(K_{\mathrm{L}}/a_{\mathrm{anc}})$ in total variation. At a matched sum ξ, the slice $J_{\xi - z_1 - z_2, h} + z_1$ and $J_{\xi,h}$ have endpoints within $O(K_{\mathrm{L}})$ and lengths $\Theta(q)$. Their uniform laws, including denominator changes, cost $O(K_{\mathrm{L}}/q)$ to couple. Safe projection has the same cost by the marginal heights. Modulo reduction and convolution contract ℓ_1, while translation and convolution preserve the height bound. Thus the uncentered difference from the convolved projected pair has ℓ_1 norm $O(K_{\mathrm{L}}/a_{\mathrm{anc}} + K_{\mathrm{L}}/q)$ and maximum entry $O((qa_{\mathrm{anc}})^{-1})$. Interpolation, also after averaging sign and blur, gives

$$C\left(\frac{\sqrt{qK_{\mathrm{L}}}}{a_{\mathrm{anc}}} + \sqrt{\frac{K_{\mathrm{L}}}{a_{\mathrm{anc}}}}\right) = O(Lt^{-2}).$$

Since

$$\mathbf{C}_{\mathrm{B}}\mathsf{C}_{\mathsf{p}_0} = \mathsf{C}_{\mathbf{C}_{\mathrm{B}}\mathsf{p}_0}\mathbf{C}_{\mathrm{B}},$$

the convolved seed is centered at $\mathbf{C}_{\mathrm{B}}\mathsf{p}_0$. Expanding the two centering differences uses only bounded one-coordinate marginal norms, so recentering at u costs at most $C\|\mathsf{p}_0 - \mathsf{u}\|_{\mathsf{u}^{-1}} = O(Lt^{-4})$ by convolution contraction and Section F.1.1. Combining the errors and using $\pi_0 \geq \mathsf{u}/2$ proves the residual bound.

For the marked seed, the uncentered squared norm is at most Cq/a_{anc}, by the anchor-column height and unit column mass. Each centering is bounded in the indicated profile norm: $\mathsf{u}/\pi_0 \leq 2$ and the second reference is its own centering measure. This proves its bound. Convolution commutes with total mass, so applying it to the seed centers the second output at $\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{p}_0$. Changing this centering to u gives exactly the product in the marked mixed identity of Lemma F.8. □

The diffuse safe seed uses two body slots, its residual two primary slots, and the marked seed one slot of each type. Including all g legs and the marked activation, their weighted costs are respectively

$$O(Lt^{-3}), \qquad O(L^3t^{-2}), \qquad O(L^{1/2}t^{-1}).$$

All are $O(t^{-1/2})$. The marked product residual is smaller: the primary and free certificates above give population-weighted cost $O(L^2t^{-7})$. For either diffuse marginal v_t, use the exact decomposition

$$\mathsf{v}_t - \mathsf{u} = \mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * (\mathsf{B}_{\mathrm{safe}}\mathsf{v}_t - \mathsf{p}_0) + (\mathsf{v}_t - \mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{B}_{\mathrm{safe}}\mathsf{v}_t) + (\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{p}_0 - \mathsf{u}).$$

The first parenthesis is a mass-zero safe seed with bounded center norm, since its marginal height is $O(q^{-1})$. The second is a mass-zero primary residual. Marginalizing the translation/projection comparison in Lemma F.8 contracts its uniform-profile norm, giving $O(Lt^{-2})$ also in the primary reference. Lift these two terms by their respective slot maps and use F for the third. This specifies the diffuse marginal certificate, of norm $O(L^2t^{-10})$.

One scalar per whole triple. At fixed t the attachment has law $\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{T}_t\mathsf{u}_{\mathrm{safe}}$. Disjoint placement excludes queue injection at this attachment. The off-body pair and this attachment are conditionally independent in the independent-root comparison. Their triple is

$$\mathbb{E}_t\left[((1-\gamma_{\mathrm{cal}})\mathsf{P}_{\mathrm{diff}\,t} + \gamma_{\mathrm{cal}}\mathsf{P}_{\mathrm{mark}\,t}) \otimes (\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{T}_t\mathsf{u}_{\mathrm{safe}})\right].$$

The constant activation is justified by independence of the first leg input from t. Construct quadratic and cubic certificates by expanding this displayed tensor, substituting the specified marginal and pair certificates, multiplying at the same t, and only then averaging that scalar. Denote them by h_{20} and h_{30}. They do not depend on the queue scalar. The pair bounds and attachment certificate give

$$g\mathcal{C}_\rho(\mathsf{h}_{20}) = O(t^{-1/2}), \qquad g\mathcal{C}_\rho(\mathsf{h}_{30}) = O(t^{-14}).$$

Indeed, multiplying the attachment certificate by the marked and diffuse marginal certificates, including activation and summing legs, costs $O(L^2t^{-7} + L^{7/2}t^{-8})$. The marked primary gives the first term; the marked companion is smaller. The third output adds the translation certificate to a pair.

Define the canonical coefficient $\bar{\mathsf{J}}_1 = \mathbb{E}_{t\sim\mu_{\mathrm{L}}}\mathsf{J}_t$. The subscript on this averaged certificate counts the one scalar injection; it does not denote translation by the integer 1. Choose the linear leg certificate to be

$$\mathsf{l}_0 = -\frac{n_{\mathrm{F}}}{g}\mathsf{F} - \frac{q}{g}\bar{\mathsf{J}}_1.$$

Its evaluation is the true sum of the three leg marginal deviations: the averaged mean identity of Theorem 9.3 gives $n_{\mathrm{F}}\mathcal{E}\mathsf{F} + q\mathcal{E}\bar{\mathsf{J}}_1$ plus g times that sum equal to zero. Since $n_{\mathrm{F}}/g = O(t^4)$

and $q/g = O(t^{-8})$, this linear certificate has norm $O(L^{3/2}t^{-10})$. With these choices the complete fixed-queue linear coefficient is

$$\begin{aligned}\mathsf{M}_y &= n_{\mathrm{F}}\mathsf{F} + q\mathsf{J}_y + g\mathsf{I}_0 \\ &= q(\mathsf{J}_y - \bar{\mathsf{J}}_1).\end{aligned} \tag{10}$$

This is an equality before evaluation. In the fixed public family, both μ_{L} and μ_{Q} are the same $\mathsf{D}_{\mathbb{Z},\sigma}$ conditioned on $|t| < K_{\mathrm{L}} = K_{\mathrm{Q}}$. Consequently $\mathbb{E}_{y\sim\mu_{\mathrm{Q}}}\mathsf{J}_y = \bar{\mathsf{J}}_1$ as canonical coefficients and $\mathbb{E}_y\mathsf{M}_y = 0$. Its uniform norm is $O(L^{3/2}t^{-6})$.

Proposition F.9 (Complete background coefficients)**.** There are explicit unit-constant canonical coefficients Q_y such that

$$\mathcal{E}\mathsf{Q}_y = [\mathsf{F}^0{}_y]_{\le R}, \qquad \sup_y[\mathcal{C}_\rho(\mathsf{Q}_y - 1) + \mathcal{C}_\rho(\mathsf{Q}_y^{-1} - 1)] = O_R(L^{-3}).$$

Here $\mathsf{F}^0{}_y$ is the character polynomial of the independent-root proposal at the fixed queue value, with n_{F} free bodies, q queue-only bodies, and g complete leg triples.

Proof. The free population includes the g compensators, each with an independent control and comparison root before scoring. Its positive factor multiplies the leg-triple factor; pairing them for the response leaves the vertex counts unchanged. Define Q_y by the finite formal product

$$(1+\mathsf{F})^{n_{\mathrm{F}}}(1+\mathsf{J}_y)^q(1+\mathsf{I}_0+\mathsf{h}_{20}+\mathsf{h}_{30})^g.$$

Each factor has the specified leading evaluation, so Lemma F.5 gives the complete independent-root polynomial. Its finite formal logarithm H_y is the sum of the three role logarithms by Lemma F.5. Its degree-one coefficient is exactly M_y by (10). The complete nonconstant leg factor has norm $O(L^{3/2}t^{-10})$, dominated by its linear certificate. The free and queue factors have norm $O(L^{3/2}t^{-14})$. There are no cross-role products in the logarithm. All powers of order at least two are therefore bounded by the corresponding squared norm times a fixed-order constant. Thus

$$\begin{aligned}\mathcal{C}_\rho(\mathsf{H}_y) &\le C_R(L^{3/2}t^{-6} + L^3t^{-8} + t^{-1/2} + L^3t^{-4}) \\ &= O_R(t^{-1/2}).\end{aligned}$$

The four terms account respectively for the common-queue linear residual, higher free and queue powers, leg terms of degree at least two, and higher powers of the complete leg factor. The finite exponential identity of Lemma F.5 gives the same bound for $\mathsf{Q}_y - 1$. Since $t = L^8$, this is $O_R(L^{-4})$ and hence implies the stated bound. The background reciprocal estimate proved in Proposition F.6 gives it also for $\mathsf{Q}_y^{-1} - 1$, uniformly over the queue. □

F.2.4 Normalized response blocks and protected-cycle defects

For a canonical coefficient f, let $\mathsf{S}_y(\mathsf{f})$ be the protected-cycle endpoint character polynomial weighted by $\mathsf{U}J_{\mathrm{in}}\mathsf{f}$. This map is linear; its scalar input one gives the unweighted polynomial $\mathsf{F}^{\mathrm{cyc}}{}_y$. Use a superscript 0 on the signed polynomial for independent roots. Omitting the queue subscript denotes averaging the complete fixed-queue polynomial under μ_{Q}. All polynomial products, reciprocals, and identities in the remainder of this appendix are truncated above R. Unit-constant reciprocals use the finite formal formula of Proposition F.6. We construct output coefficients as centered signed color profiles to apply the first inverse, then Fourier-transform them.

One local degree budget. Use the local scale Q_{loc} from Lemma G.4. Positive paired coefficients and signed paired numerators of degree m have retained-Hilbert norm at most $C_R Q_{\mathrm{loc}}^m$. The same bound holds for each degree of a finite positive-factor reciprocal: degrees add in every product of its truncated geometric series. An ordinary positive body has bounded profile norm, so its finite reciprocal obeys the same bound because $Q_{\mathrm{loc}} \geq 1$. All local polynomials retain the complete flag–map pair and same-source leg scalar. Define the new-primary-slot cost and the normalized output factor by

$$\tau_{\mathrm{R}} := {\gamma_{\mathrm{cal}}}^{-1} {m_{\mathrm{P}}}^{-1/2}, \qquad \omega_{\mathrm{R}} := \tau_{\mathrm{R}} Q_{\mathrm{loc}}.$$

For the public family, $\tau_{\mathrm{R}} = O(Lt^{-8})$ and $\omega_{\mathrm{R}} = O(L^{1/2}t^{-6}) = o(1)$.

Lemma F.10 (Independent-root raised response)**.** An explicit operator $\mathsf{U}_{\mathrm{raise}}$ on canonical coefficients satisfies

$$\mathsf{S}_y^0(\mathsf{f})/{\mathsf{F}^0}_y = \mathcal{E}(I + \mathsf{U}_{\mathrm{raise}})\mathsf{f}, \qquad \|\mathsf{U}_{\mathrm{raise}}\|_{\mathcal{C}_\rho \to \mathcal{C}_\rho} = O_R(L^{-3}).$$

It strictly raises represented degree, creates no new body slot, and is independent of y.

Proof. Divide each selected group's complete signed polynomial by its positive polynomial, canceling unselected factors. A primary group contains its leg triple and distinct compensator; an ordinary group contains one body output. Independent comparison roots and the complete conditional flag–map law permit these local operators to act on arbitrary retained joint Hilbert inputs.

The degree-one primary response is exactly $\gamma_{\mathrm{cal}}\mathbf{T}_1$; represent it by the unchanged primary slot. A nonleading primary response emits $m \geq 2$ profiles. The weighted transfer rule of Lemma F.7 and the common degree budget bound its cost by

$$C_R {m_{\mathrm{P}}}^{1/2} \omega_{\mathrm{R}}^m.$$

For an ordinary group, retain the leading safe body slot without inverting its convolution. Each added reciprocal output costs at most ω_{R}. Because this factor is eventually at most one, every nonleading local response gains at most

$$C_R \max\{{m_{\mathrm{P}}}^{1/2} \omega_{\mathrm{R}}^2, \omega_{\mathrm{R}}\} = O_R(Lt^{-4}).$$

Tensor the local expansions with the exact distinct-index normalizers and symmetrize. Their leading responses are the identity; every nonleading fixed-order pattern contains a small factor just bounded and strictly raises represented degree, creating no body slot. Extra outputs only multiply by a nonincreasing factor. The number of degree patterns depends on the fixed order alone. This defines $\mathsf{U}_{\mathrm{raise}}$ and proves its stated bound. Selected groups avoid the queue, whose unselected factors canceled before expansion, so the operator is queue-independent. □

Lemma F.11 (Coarse transfer of centered root patterns)**.** Let the source and public-control reference laws be independent, and let the centered joint-root density satisfy $\|\mathsf{H}_r\|_2 \leq C_r \eta_{\mathrm{R}}^r$ under the product root reference. Let f be a canonical public-control kernel and partition the consumed arguments into disjoint local groups, using each complete control once. Local maps may be restricted to the centered root and control subspaces. In an expansion through degree R,

with $r \le R$ and the slot factors of Lemma F.7, each group hit by a root tag emits $m \ge 1$ profiles. Let N_{R} bound unselected choices per group and M_{R} bound a selected group's consumed-input weight. Assume an unselected single-output root map has retained-Hilbert norm at most δ_{R}; all other degree-m numerators and finite reciprocal coefficients have norm at most $C_R Q_{\mathrm{loc}}^m$. Assume unhit unselected factors cancel, while unhit selected normalized responses and the number of incidence and degree patterns are bounded in terms of the fixed order. Suppose

$$\omega_{\mathrm{R}} \le 1, \qquad \varepsilon_{\mathrm{R}} := \eta_{\mathrm{R}} \max\{1, N_{\mathrm{R}}\tau_{\mathrm{R}}\delta_{\mathrm{R}}, N_{\mathrm{R}}\omega_{\mathrm{R}}^2, M_{\mathrm{R}}\omega_{\mathrm{R}}\} \le 1.$$

The full response from nonempty root patterns has norm at most $C_R\varepsilon_{\mathrm{R}}$.

Proof. Source independence gives $\|\mathsf{H}_r \otimes \mathsf{f}\|_2 = \|\mathsf{H}_r\|_2\|\mathsf{f}\|_2$. Slotwise centering fixes this joint input. Apply the joint-root norm once and tensor the local maps on the disjoint groups; their retained-Hilbert bounds apply without factoring the root density. Reserve one scalar factor η_{R} for each hit group. There are no more hit groups than tags, and every remaining tag factor is at most one since $\eta_{\mathrm{R}} \le \varepsilon_{\mathrm{R}} \le 1$.

By Lemma F.7, an unselected hit group costs N_{R} times either $\tau_{\mathrm{R}}\delta_{\mathrm{R}}$ for one output or ω_{R}^m for at least two outputs. A selected hit group costs $M_{\mathrm{R}}\omega_{\mathrm{R}}^m$. After its reserved tag factor, each is bounded by the common budget $C_R\varepsilon_{\mathrm{R}}$; extra outputs multiply by $\omega_{\mathrm{R}} \le 1$. A multiply tagged group still consumes its complete control once.

Every nonempty pattern has a hit group. Multiplying these common budgets and the bounded unhit responses costs at most $C_R\varepsilon_{\mathrm{R}}$, since $\varepsilon_{\mathrm{R}} \le 1$. All population factors are already included; normalized selected-index averages and symmetrization contract. The remaining incidence and degree patterns depend only on the fixed order, so summing them proves the claim. □

Lemma F.12 (Protected-cycle defects on joint input kernels)**.** There are explicit canonical coefficients V_y and linear operators $\mathsf{B}_{\mathrm{root}\,y}$ satisfying

$$\begin{aligned} \mathcal{E}\mathsf{V}_y &= \mathsf{F}^{\mathrm{cyc}}{}_y/\mathsf{F}^0{}_y - 1, \\ \mathcal{E}\mathsf{B}_{\mathrm{root}\,y}\mathsf{f} &= \mathsf{S}_y(\mathsf{f})/\mathsf{F}^0{}_y - \mathcal{E}(I + \mathsf{U}_{\mathrm{raise}})\mathsf{f}, \\ \mathsf{B}_{\mathrm{root}\,y}1 &= \mathsf{V}_y. \end{aligned} \tag{11}$$

Uniformly over protected queue values,

$$\mathcal{C}_\rho(\mathsf{V}_y) = O_R(L^{-3}), \qquad \|\mathsf{B}_{\mathrm{root}\,y}\| = O_R(L^{-3}).$$

The operators do not lower represented degree, and $\mathsf{V}_y^{[0]} = \mathsf{V}_y^{[1]} = 0$ as canonical coefficients.

Proof. Fix a protected queue value y. The two experiments use the same public controls and local endpoint maps; only the joint law of the body roots changes. We isolate that change, construct its canonical coefficients, and then bound their norm. All polynomial operations below are truncated above R.

1. Dependent-root contribution. By Lemma F.10, the required weighted defect has leading evaluation

$$\mathcal{E}\mathsf{B}_{\mathrm{root}\,y}\mathsf{f} = \frac{\mathsf{S}_y(\mathsf{f}) - \mathsf{S}_y^0(\mathsf{f})}{\mathsf{F}^0{}_y}.$$

To expand this difference, write the identity in each root coordinate as the sum of its mean projection and its centering projection under $\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}}$. Expanding the joint root likelihood gives the constant term 1, which produces the independent-root experiment, and terms centered in nonempty sets of root coordinates. Call those coordinates *tagged*. A term with r tags has the joint density H_r of Theorem E.4.

A root enters the local channel only by translating its own body endpoint. If that endpoint contributes no output to a polynomial term, the term is independent of the root, and integration against its centered tag gives zero. Thus every surviving tag requires its own body output, and at most R tags survive. The cycle-density theorem and the public parameter choice give

$$\|\mathsf{H}_r\|_{L^2((\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}})^{\otimes r})} \le C_r \eta_{\mathrm{R}}^r, \qquad \eta_{\mathrm{R}} = O(Lt^{-10}).$$

2. Canonical coefficient construction. Work first with one summand in the distinct-index control average. A selected primary group consists of its complete leg triple and assigned compensator; a selected ordinary group consists of one body. Call a group *hit* if it contains a tagged root. For each nonempty root pattern, form the complete local endpoint polynomials, multiplying the leg and attachment factors at the same leg scalar before averaging. A paired group uses its complete control once, including when both its attachment and compensator roots are tagged.

Divide by the independent-root polynomial using its local finite reciprocals. Unhit unselected factors cancel. Unhit selected groups retain the normalized response of Lemma F.10, while hit groups act on the joint root and control kernel. Center the new output profiles, lift them by ${\gamma_{\mathrm{cal}}}^{-1}\mathbf{T}_1^{-1}$, and store the resulting primary densities; retain any unchanged safe body seeds as body slots. Average with the exact distinct-index normalizers and symmetrize the two input types. These finite linear operations define $\mathsf{B}_{\mathrm{root}\,y}$ on canonical coefficients. Their evaluation is the difference in Step 1, term by term. For scalar input 1, the same construction defines $\mathsf{V}_y := \mathsf{B}_{\mathrm{root}\,y}1$. Since the two scalar responses are ${\mathsf{F}^{\mathrm{cyc}}}_y$ and ${\mathsf{F}^{0}}_y$, this proves all three response identities.

3. Local and global norm bounds. The root density need not factor across groups. Instead, source independence gives the single joint-input bound

$$\|\mathsf{H}_r \otimes \mathsf{f}\|_2 = \|\mathsf{H}_r\|_2 \|\mathsf{f}\|_2.$$

We apply local maps to disjoint groups while retaining all other arguments in a Hilbert space. The retained-input bounds of Lemma G.4 therefore apply to this joint kernel. Integration over roots uses Cauchy–Schwarz under their probability reference, so no inverse root-atom probabilities are needed.

The case requiring a sharper local bound is an unselected hit group with just one output. Its tagged root forces that output to be a body endpoint. Before the root translation, its profile $\mathsf{\mu}$ is p_0 for a free body or compensator, or a fixed translate of $\mathsf{u}_{\mathrm{safe}}$ for a queue or attachment. Root centering replaces the translated profile by its difference from the root average. Translation fixes the uniform law and preserves its profile norm, while convolution contracts that norm. Consequently,

$$\sup_z \|\mathsf{T}_z\mathsf{\mu} - \mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * \mathsf{\mu}\|_{\mathsf{u}^{-1}} \le 2\|\mathsf{\mu} - \mathsf{u}\|_{\mathsf{u}^{-1}}.$$

Conversion to the primary reference costs at most $\sqrt{2}$. Thus Sections F.1.1 and H.4 provide the single-output bound

$$\delta_{\mathrm{R}} := 2\sqrt{2}(\|\mathsf{p}_0 - \mathsf{u}\|_{\mathsf{u}^{-1}} + \|\mathsf{u}_{\mathrm{safe}} - \mathsf{u}\|_{\mathsf{u}^{-1}}) = O(Lt^{-4}).$$

All other local numerators and finite reciprocal coefficients obey the degree bound $C_R Q_{\mathrm{loc}}^m$ established at the start of Section F.2.4. This includes ordinary-center scores: their fixed-root operator norm is $O(1)$ because $\mathsf{p}_0 \le 4/q$ and $\pi_0 \ge 1/(2q)$.

We can now apply Lemma F.11. Unselected hit groups have at most the body population many choices; selected indices are already averaged, so they incur only a consumed-input weight. Accordingly, take

$$N_{\mathrm{R}} = n_{\mathcal{B}} = O(t^{20}), \qquad M_{\mathrm{R}} = \max\{{m_{\mathrm{P}}}^{1/2}, {m_{\mathrm{B}}}^{1/2}/\rho\} = O(t^{10}).$$

The slot factors in Section F.2.4 give the following three costs, before the root-tag factor:

$$\begin{aligned} N_{\mathrm{R}}\tau_{\mathrm{R}}\delta_{\mathrm{R}} &= O(L^2 t^8) && \text{(unselected, one output)},\\ N_{\mathrm{R}}\omega_{\mathrm{R}}^2 &= O(Lt^8) && \text{(unselected, at least two outputs)},\\ M_{\mathrm{R}}\omega_{\mathrm{R}} &= O(L^{1/2}t^4) && \text{(selected)}. \end{aligned}$$

Extra outputs cost at most $\omega_{\mathrm{R}} = O(L^{1/2}t^{-6}) \le 1$ for sufficiently large parameters. Every hit group has a tag, so the transfer lemma supplies one factor η_{R} to offset its cost. Its common budget is therefore

$$\begin{aligned} \varepsilon_{\mathrm{R}} &= \eta_{\mathrm{R}} \max\{1, N_{\mathrm{R}}\tau_{\mathrm{R}}\delta_{\mathrm{R}}, N_{\mathrm{R}}\omega_{\mathrm{R}}^2, M_{\mathrm{R}}\omega_{\mathrm{R}}\}\\ &= O(L^3 t^{-2}) = O(L^{-13}), \end{aligned}$$

using $t = L^8$. Thus the budget is at most one for sufficiently large parameters. The lemma sums all nonempty root patterns, including patterns with several tags in one group, and yields

$$\|\mathsf{B}_{\mathrm{root}\,y}\| = O_R(L^{-13}).$$

Since $\mathcal{C}_\rho(1) = 1$, the same bound holds for V_y. This is stronger than the claimed $O_R(L^{-3})$. All local estimates are uniform in the protected queue value, and the remaining incidence and degree counts depend only on the fixed truncation order.

4. **Degree assertions.** For scalar input, Theorem E.4 gives $\mathsf{H}_1 = 0$, so a nonzero correction has at least two tags on distinct body endpoints. Both endpoints must contribute outputs by Step 1, even if they lie in the same local group. The coefficient lifts preserve these outputs and finite reciprocals only add degree. Hence the constant and linear canonical coefficients of V_y vanish. For a general input kernel, fix the roots and all other controls. If a selected group contributes no output, integrating its centered paired or ordinary score gives zero. Every surviving selected group therefore contributes at least one output. Since input degree counts selected groups, and subsequent lifts, reciprocals, and symmetrization do not lower degree, neither defect lowers represented degree.

□

F.2.5 True-cycle response and common-queue assembly

Assembly of the response. We combine the background coefficients of Proposition F.9 with the raised response and joint root operators of Lemmas F.10 and F.12. In the fixed coordinates, let $\mathbf{q}_y$ and $\mathbf{v}_y$ represent Q_y and V_y, respectively; let $\mathbf{B}_{\mathrm{root},y}$ and $\mathbf{U}_R$ represent the operators $\mathsf{B}_{\mathrm{root}\,y}$ and $\mathsf{U}_{\mathrm{raise}}$. Thus $\mathbf{U}_R\mathbf{e}_0 = 0$ and $\mathbf{B}_{\mathrm{root},y}\mathbf{e}_0 = \mathbf{v}_y$. Absorb this scalar response into the complete background and normalize once:

$$\begin{aligned}\mathbf{b} &= \mathbb{E}_y[\mathbf{M}_R(\mathbf{q}_y)(\mathbf{e}_0 + \mathbf{v}_y)],\\ \mathbf{C}_y &= \mathbf{B}_{\mathrm{root},y} - \mathbf{M}_R(\mathbf{v}_y)(\mathbf{I} + \mathbf{U}_R),\\ \mathbf{Z}_R &= \mathbf{M}_{\mathrm{bg}}^{-1}\mathbb{E}_y[\mathbf{M}_R(\mathbf{q}_y)\mathbf{C}_y].\end{aligned} \tag{12}$$

Here $\mathbf{M}_{\mathrm{bg}} = \mathbf{M}_R(\mathbf{b})$, as in Section F.2.5; the corresponding canonical coefficient is B_{bg}. Every product inside an expectation uses one queue value. Only after averaging the complete products do we invert the background. Its constant coordinate is one, so its multiplication inverse exists by Proposition F.6.

Convexity, submultiplicativity, and the background and root bounds give

$$\|\mathbf{b} - \mathbf{e}_0\|_{\mathrm{simp},\rho} + \sup_y \|\mathbf{C}_y\|_{\mathrm{simp},\rho} = O_R(L^{-3}).$$

The same proposition bounds $\mathbf{M}_{\mathrm{bg}}^{-1}$ by $1 + O_R(L^{-3})$. Consequently,

$$\|\mathbf{Z}_R\|_{\mathrm{simp},\rho} = O_R(L^{-3}), \qquad \mathbf{Z}_R\mathbf{e}_0 = 0,$$

since $\mathbf{C}_y\mathbf{e}_0 = 0$ before evaluation. For matrices these are induced graded operator norms.

For a fixed queue value, the root identities (11) and multiplicativity of leading evaluation give the exact physical response

$$[\mathsf{S}_y(\mathsf{f})]_{\le R} = \mathbf{E}_R\mathbf{M}_R(\mathbf{q}_y)(\mathbf{I} + \mathbf{U}_R + \mathbf{B}_{\mathrm{root},y})\mathbf{f}.$$

Substitute $\mathbf{C}_y$ from (12) and average. Since $\mathbf{U}_R$ is queue-independent and multiplication matrices respect the coefficient product, the result is

$$\begin{aligned}[\mathsf{S}(\mathsf{f})]_{\le R} &= \mathbf{E}_R\mathbf{M}_{\mathrm{bg}}(\mathbf{I} + \mathbf{U}_R + \mathbf{Z}_R)\mathbf{f}\\ &= \mathbf{E}_R\mathbf{A}_R\mathbf{f},\\ [\mathsf{F}^{\mathrm{cyc}}]_{\le R} &= \mathbf{E}_R\mathbf{b}.\end{aligned}$$

The scalar-input case gives the second identity. Both use the physical synthesis, averaged symmetrization, and exact distinct-index normalizers.

Grading and the canonical degree-one term. The raw root operator does not lower degree, and its scalar image $\mathbf{v}_y$ starts in degree two, by Lemma F.12. Hence $\mathbf{C}_y$ does not lower degree. Multiplication by the unit-constant background inverse and averaging preserve this property, so $\mathbf{Z}_R$ does not lower degree either. Together with the strictly raising $\mathbf{U}_R$, this makes both the normalized and full response block lower triangular, with actual diagonal blocks

$$\mathbf{I} + (\mathbf{Z}_R)_{r,r}.$$

These blocks act on the full hybrid degree space, including all body-to-primary conversions. The protected-root diagonal is retained.

For $R \ge 1$, the background product in Proposition F.9 has canonical degree-one coefficient M_y. The coefficient identity (10) and the equality of the clipped leg and queue laws give $\mathbb{E}_y\mathsf{M}_y = 0$ before evaluation. Since the root defect starts in degree two, the definition of $\mathbf{b}$ therefore gives $(\mathbf{b})_1 = 0$ as a canonical coefficient.

F.2.6 Application to the public correction

Proof of Theorem 9.5. By Section F.2.5, the full response is $\mathbf{A}_R = \mathbf{M}_{\mathrm{bg}}(\mathbf{I}+\mathbf{U}_R+\mathbf{Z}_R)$ and $\mathbf{A}_R\mathbf{e}_0 = \mathbf{b}$. Submultiplicativity in Lemma F.5 and the assembled bounds give, in the induced graded norm,

$$\begin{aligned}\|\mathbf{A}_R - \mathbf{I}\|_{\mathrm{simp},\rho} &\le \|\mathbf{b}-\mathbf{e}_0\|_{\mathrm{simp},\rho} \\ &\quad + (1+\|\mathbf{b}-\mathbf{e}_0\|_{\mathrm{simp},\rho})\,(\|\mathbf{U}_R\|_{\mathrm{simp},\rho} + \|\mathbf{Z}_R\|_{\mathrm{simp},\rho}) \\ &= O_R(L^{-3}).\end{aligned}$$

For fixed R and sufficiently large L, this is at most $1/2$. Apply Theorem 7.1 with synthesis $\mathbf{\Psi}_R$, evaluation $\mathbf{E}_R$, and target and reference $\mathbf{e}_0$. Its reference-residual bound gives the unique solution with

$$\|\mathbf{z}-\mathbf{e}_0\|_{\mathrm{simp},\rho} \le 2\|\mathbf{b}-\mathbf{e}_0\|_{\mathrm{simp},\rho} = O_R(L^{-3}).$$

The scalar row is the identity, so the correction has scalar zero. The degree-one equation then has zero right-hand side because $(\mathbf{b})_1 = 0$. Homogeneous projections contract the norm, so the same small-perturbation bound makes the degree-one diagonal block invertible and gives $(\mathbf{c})_1 = 0$. Paired lifting preserves canonical centering, so distinct-index averaging gives mean zero for the correction. Exact calibration yields

$$\mathbb{E}_{\nu_0} W_R = 1, \qquad [\mathsf{P}_{\mathrm{char},W_R}]_{\le R} = \mathbf{E}_R\mathbf{e}_0 = 1.$$

Independent uniform relabeling makes the law exchangeable. Each homogeneous character coefficient is then a positive combinatorial multiple of its support-frequency coefficient, giving the claimed cancellation. The kernels and inverses are averaged public quantities, so evaluating the weight reads only the controls in Section 9.2. □

G Complete Fourier Tails and Public-Core Moments

Use the public family and radius of Section 9.9. For a finite polynomial P of degree at most d, write $\|\mathsf{P}\|_{\mathcal{F},\rho}$ for the ordered-coefficient norm of Definition 9.1. Applied to a character polynomial, this is $\mathfrak{F}_\rho(\mathsf{P})$. Symmetrization is the orthogonal projection that averages tensor-slot permutations. Contractions and conjugation follow Section E.1.2. We apply the bounds below to the physical weight W_R and its canonical public-control coefficient f_R from Section 9.6. The real public input coefficients and scores give conjugate-symmetric character coefficients. Fix a disjoint placement, with the free population and control pools of Section 9.2. The protected latent cycle, its independent body permutation, and the single clipped queue scalar are retained under the uncharged experiment of Section 9.7.

G.1 A mean bound for the full protected cycle

Write z_v for the protected cyclic difference at body vertex v after the fixed body permutation. Expectations in the following lemma average the latent edge samples of Definition 8.18.

Lemma G.1 (Cycle mean exponential moment)**.** Fix a disjoint placement, a body permutation sampled independently of the public core and latent cycle, and the shared queue scalar. At body vertex v, let p_v be its center profile after averaging its leg input, if present, but before adding

its cycle root. These profiles are fixed independently of the latent edge samples. Suppose

$$\|\mathsf{p}_v - \mathsf{u}\|_{\mathsf{u}^{-1}} \le \delta_0, \qquad \rho^2\delta_0^2 = o(1),$$
$$F_{\mathrm{cyc}} = d^{-1/2}\sum_{v\in\mathcal{B}} \left(T_{z_v}\mathsf{p}_v - \mathbb{E}T_{z_v}\mathsf{p}_v\right).$$

For every fixed positive C and $\rho = L^{1/4}$,

$$\mathbb{E}\exp\Big(C\rho^2\|F_{\mathrm{cyc}}\|^2_{\mathsf{u}^{-1}}\Big) = O(1).$$

The bound is uniform in the fixed data. It also holds for the entire normalized conditional mean when its root-averaged drift has norm $O(\rho^{-1})$.

Proof. The latent edge variables are independent samples from $\boldsymbol{\mu}_{\mathrm{e}}$ by Definition 8.18. Replacing one by an independent copy changes two roots. Translation fixes u and is an isometry, so each affected summand changes in norm by at most $2\delta_0$. Consequently

$$\|F_{\mathrm{cyc}} - F'_{\mathrm{cyc}}\| \le c_{\mathrm{bd}} := 4\delta_0/\sqrt{d}, \qquad v_{\mathrm{bd}} := n_{\mathcal{B}}c_{\mathrm{bd}}^2 \le 16\delta_0^2.$$

All norms in this proof are in the profile Hilbert space. The case $v_{\mathrm{bd}} = 0$ is immediate. Otherwise apply the Efron–Stein inequality [BBLM05, Proposition 1] to each orthonormal coordinate and sum, obtaining $\mathbb{E}\|F_{\mathrm{cyc}}\|^2 \le v_{\mathrm{bd}}/2$. The scalar norm obeys the same coordinate difference bound, so its Doob martingale under successive edge exposure has conditional ranges of length at most c_{bd}. The one-sided bounded-differences inequality [McD89, Theorem 6.7] therefore gives

$$\Pr\Big(\|F_{\mathrm{cyc}}\| > \sqrt{v_{\mathrm{bd}}/2} + t_{\mathrm{H}}\Big) \le e^{-2{t_{\mathrm{H}}}^2/v_{\mathrm{bd}}} \quad (t_{\mathrm{H}} \ge 0).$$

Tail integration gives, whenever $0 \le \alpha_{\mathrm{mgf}}v_{\mathrm{bd}} < 1$,

$$\mathbb{E}e^{\alpha_{\mathrm{mgf}}\|F_{\mathrm{cyc}}\|^2} \le \frac{e^{\alpha_{\mathrm{mgf}}v_{\mathrm{bd}}}}{1-\alpha_{\mathrm{mgf}}v_{\mathrm{bd}}}.$$

Take $\alpha_{\mathrm{mgf}} = C\rho^2$. The smallness hypothesis gives an absolute bound. Add the drift and use $\|x_{\mathrm{aux}} + y_{\mathrm{aux}}\|^2 \le 2\|x_{\mathrm{aux}}\|^2 + 2\|y_{\mathrm{aux}}\|^2$ to prove the last statement. □

For the proposal in Theorem 9.3, the body profiles before adding cycle roots are

body role	population	p_v
free	n_{F}	p_0
attachment	g	$\boldsymbol{\mu}_{\mathrm{L}} * \mathsf{u}_{\mathrm{safe}}$
queue	q	$T_y\mathsf{u}_{\mathrm{safe}}$.

The attachment row averages its own leg's second input; the queue row fixes the common value y. Translation is an isometry and convolution is a contraction in the uniform-reference norm. Thus the center estimate established in Section H.4 bounds every row by the same δ_0, uniformly in the placement. The off-body coordinates use independent public blurs with the protected marginal and have no actual-root dependence. Thus the *entire* histogram mean has exactly the body's root fluctuation.

The root/queue-averaged histogram mean equals $d\mathsf{u}$ by Theorem 9.3. For fixed queue value its root-averaged excess is exactly

$$q\mathsf{R}_{\mathrm{B}}^{\mathrm{prot}} * (T_y\mathsf{u}_{\mathrm{safe}} - \boldsymbol{\mu}_{\mathrm{Q}} * \mathsf{u}_{\mathrm{safe}}).$$

Only the q queue vertices change when the queue average is removed; attachments lie outside that set. Both terms in parentheses are within δ_0 of u. Hence the normalized drift is at most

$$2q\delta_0/\sqrt{d}.$$

For the selected family, the center deviation is $\delta_0 = O(Lt^{-4})$. Consequently

$$\rho^2\delta_0^2 + \frac{\rho^2 q^2 \delta_0^2}{d} = O(L^{5/2}t^{-8}) = o(1).$$

This verifies both hypotheses of Lemma G.1 at every fixed multiple of the radius, uniformly in the common queue value.

G.2 Gaussian orthogonality and the all-degree product bound

Lemma G.2 (Complex Gaussian orthogonality)**.** Let P be a polynomial of degree at most d with the ordered symmetric coefficient convention of Definition 9.1, and let Z have independent circular complex Gaussian coordinates, with real and imaginary variances $1/2$. Then

$$\mathbb{E}|\mathsf{P}(\rho Z/\sqrt{d})|^2 = \|\mathsf{P}\|_{\mathcal{F},\rho}^2 .$$

The identity also holds for coefficients in a finite-dimensional Hilbert space, with its squared norm in place of absolute square. It requires neither an endpoint probability law nor independence among the endpoints represented by the polynomial.

Proof. For nonnegative integers a, b, one-variable Gaussian integration gives

$$\mathbb{E} Z_{j_{\mathrm{slot}}}^a \overline{Z_{j_{\mathrm{slot}}}}^b = \mathbf{1}_{\{a=b\}} a!.$$

Independence makes distinct monomials orthogonal. Within degree r, grouping ordered tuples with the same multiplicities contributes exactly $r!$ times the squared Euclidean norm of the ordered symmetric coefficient tensor. Rescaling by $\rho/\sqrt{d}$ gives the claimed norm at every degree. Sum over degrees and an orthonormal basis of the coefficient Hilbert space to obtain both assertions. □

For any homogeneous symmetric tensor of degree k, symmetrization contraction and $(k+r)! \le 2^{k+r} k! r!$ give

$$\left\| \langle T_k, \mathbf{t}^{\otimes k} \rangle \mathsf{P} \right\|_{\mathcal{F},\rho} \le \sqrt{k!} \left(\frac{\sqrt{2}\rho}{\sqrt{d}} \right)^k \|T_k\|_2 \, \|\mathsf{P}\|_{\mathcal{F},\sqrt{2}\rho} . \tag{13}$$

This follows by squaring and summing the degree bounds, so applies to every finite degree and to Hilbert-valued input coefficient maps.

Lemma G.3 (Finite products with one nonlinear radius enlargement)**.** Let P be a product of blocks of size at most three, with block polynomials $1 + \ell_v + \mathsf{h}_{v,2} + \mathsf{h}_{v,3}$, with at most g

nonsingleton blocks. Define

$$\eta = d^{-1}\sum_v \|\ell_v\|_2^2, \qquad M = d^{-1/2}\left\|\sum_v \ell_v\right\|_2,$$
$$A_2 = d^{-1}\sum_v \|\mathsf{h}_{v,2}\|_2, \quad A_3 = d^{-3/2}\sum_v \|\mathsf{h}_{v,3}\|_2.$$

There are absolute positive constants c, C such that

$$\rho^2\eta \le c, \qquad \rho^2 A_2 + \sqrt{g}\rho^3 A_3 \le c$$

imply $\|\mathsf{P}\|_{\mathcal{F},\rho}^2 \le Ce^{C\rho^2 M^2}$. For a mixture, uniform conditional activity bounds and a bounded expectation of this exponential suffice.

Proof. Apply Lemma G.2 to the full product. For the product of linear factors use $|1+z_{\mathrm{ev}}|^2 \le \exp(2\operatorname{Re} z_{\mathrm{ev}} + |z_{\mathrm{ev}}|^2)$. The resulting quadratic matrix has trace $\rho^2\eta$. Completing the square in the Gaussian integral gives a determinant $\det(I-\rho^2 S)^{-1}$ and a mean factor bounded by $\exp(2\rho^2 M^2)$, where $S = d^{-1}\sum_v \ell_v\ell_v^*$. If the trace is at most $1/2$, then $-\log\det(I-\rho^2 S) \le 2\rho^2\eta$, proving the linear claim.

Expand the product by selecting a quadratic and b cubic terms, in disjoint blocks, and put $j = a+b$. Omitting their linear factors changes the squared normalized mean to at most $2M^2+2j\eta$, by Cauchy–Schwarz. Thus their remaining linear product costs at most

$$\exp\Big(2c + 4\rho^2 M^2 + 4jc\Big)$$

at radius $\sqrt{2}\rho$, provided $c \le 1/4$. The factor e^{2c} is retained also when no nonlinear block is selected and $M = 0$. Multiply *all* selected nonlinear tensors first and apply (13) once. Summing their norms and dropping disjointness gives the upper bound

$$e^{4jc}2^{(2a+3b)/2}\frac{\sqrt{(2a+3b)!}}{a!b!}(\rho^2 A_2)^a(\rho^3 A_3)^b$$

times $\exp(2c+4\rho^2M^2)$. The factorial quotient is at most $4^j(3g)^{b/2}$: use $\sqrt{(2j)!}/j! \le 2^j$, $\binom{j}{a} \le 2^j$, and bound the remaining b factorial factors by $(3g)^b$ before taking the square root. For $c \le 1/4$, these terms are bounded by the product of two geometric majorants whose ratios are absolute constant multiples of $\rho^2 A_2$ and $\sqrt{g}\rho^3 A_3$, respectively. Choose c so both ratios are at most $1/2$. The two geometric sums are then bounded by an absolute constant. Square the norm bound and enlarge C to obtain the claim. If there are no nonsingleton blocks, the linear estimate applies directly. The single radius enlargement is independent of the number of blocks and the total degree. Jensen in the finite Hilbert direct sum proves the mixture statement. □

G.3 Local polynomial operators and the weighted Fourier bound

For the paired lift of Section 9.4, write

$$r_{\mathrm{cmp}} := \gamma_{\mathrm{cal}}\frac{\widetilde{\mathsf{A}}(\pi_0 f_{\mathrm{in}})}{\mathsf{p}_0}$$

for the compensator summand. Its denominator is used only on safe support. The two public inputs are independent under the proposal.

Lemma G.4 (Local polynomial operators and paired lifts)**.** Use the public family of Section 9.9 and its positive proposal from Theorem 9.3, for all sufficiently large L. Conditional on the complete protected root vector, the body permutation, the placement, and the shared queue, the unweighted proposal is a product of $n_{\mathrm{F}} + q$ singleton laws and g same-source triple laws. A triple contains the two off-body leg coordinates and its own attachment; the attachment has no queue shift. Put

$$Q_{\mathrm{loc}} := \sqrt{\frac{\gamma_{\mathrm{cal}} d}{g}}.$$

The positive triple, the positive four-coordinate factor obtained by adjoining its distinct compensator, and the corresponding signed paired group induced by a centered unit score in $L^2(\pi_0)$ have degree-r raw coefficient operator norm at most

$$C Q_{\mathrm{loc}}^r \qquad (1 \le r \le 4).$$

The triple has no degree-four term, and the signed paired factor has zero constant coefficient. An ordinary-center score has only a degree-one output, with operator norm $O(1)$. The paired lift has Hilbert operator norm $O(1)$ and evaluation-row norm $O(q)$ on its paired reference support. These raw output bounds hold both in the uniform-reference profile norm and in the primary-reference profile norm, uniformly in fixed attachment and compensator roots. All operator statements allow arbitrary retained Hilbert-valued inputs.

Proof. Each fixed-sign diffuse pair averages a uniform slice of length $q - O(a_{\mathrm{anc}} + K_{\mathrm{L}})$ over an anchor interval of length a_{anc}. For a specified pair, its pair sum fixes at most one anchor value after fixing the blur and leg input. Its point mass is therefore at most $C/(q a_{\mathrm{anc}})$, giving pair likelihood norm at most $C\sqrt{q/a_{\mathrm{anc}}}$. Its marginals have height $O(1/q)$. Convolution with the independent public blur preserves these bounds.

For a marked pair with the same leg scalar fixed, its law is a translate of a kernel with first marginal π_0 and conditional second-coordinate height $1/a_{\mathrm{anc}}$. Since $\sum \pi_0^2 \le \|\pi_0\|_\infty$, Cauchy–Schwarz bounds both the unweighted pair and its operator on unit primary scores by

$$q\sqrt{\frac{\|\pi_0\|_\infty}{a_{\mathrm{anc}}}}.$$

The analogous primary marginal bound is $\sqrt{q\|\pi_0\|_\infty}$. Multiply these marked terms by γ_{cal}. For the diffuse part the primary enters only through its anchor sign. The absolute score coefficient of either sign-conditioned law is at most $\|f_{\mathrm{in}}\|_1 \le \|f_{\mathrm{in}}\|_2$. This controls that part without division by a small sign probability.

The same-source triple retains the leg scalar in both the companion and attachment, as in Section 8.5.6. Conditional on that scalar, the attachment has a fresh center of height $O(1/q)$, independent of the anchor and public blur. Average the pair over the anchor and blur, tensor with the attachment law whose uniform-reference likelihood norm is $O(1)$, and average over the *same* scalar. This gives the pair bounds for the triple. The actual root is fixed and only translates the attachment. Centering output coordinates is orthogonal projection, so it can only decrease those norms. For the linear coefficient, the other marginals are bounded by the height estimates. The same primary height thus controls both the unweighted product and the score operators.

For the compensating term, the Markov-kernel Cauchy–Schwarz inequality gives

$$\|r_{\mathrm{cmp}}\|_{L^2(\mathfrak{p}_0)}^2 \le C\gamma_{\mathrm{cal}} \|f_{\mathrm{in}}\|_2^2.$$

Here use $\gamma_{\mathrm{cal}}\widetilde{\mathsf{A}}\pi_0 \le C/q$ and $\mathsf{p}_0 \ge C^{-1}/q$ on its support. The compensator singleton's linear deviation is $o(1)$ by Section H.4. There are exactly g such four-coordinate paired groups, each containing one triple and a distinct free singleton. The other free singletons number $n_{\mathrm{F}} - g = m_{\mathrm{B}}$; the q queue singletons are never controls. The signed compensator polynomial has zero constant coefficient and linear norm $O(\sqrt{\gamma_{\mathrm{cal}}})$. Since $\gamma_{\mathrm{cal}} \le 1$, multiplying it by an unweighted triple is bounded by the stated score-operator bounds. The signed-primary triple times the positive compensator obeys the same bounds, since the compensator has constant coefficient one and linear norm $o(1)$. This controls every degree through four.

The preceding bounds give linear coefficient cost $O(1 + \gamma_{\mathrm{cal}}\sqrt{q\|\pi_0\|_\infty})$ and higher-degree cost $O(\sqrt{q/a_{\mathrm{anc}}}(1 + \gamma_{\mathrm{cal}}\sqrt{q\|\pi_0\|_\infty}))$. The proposal height bound of Lemma F.1 gives

$$\gamma_{\mathrm{cal}}\sqrt{q\|\pi_0\|_\infty} \le Q_{\mathrm{loc}}.$$

For the public family, $Q_{\mathrm{loc}} \asymp L^{-1/2}t^2 \ge 1$ eventually, and $\sqrt{q/a_{\mathrm{anc}}} = O(Q_{\mathrm{loc}})$. Consequently the linear bounds above are $O(Q_{\mathrm{loc}})$, while all higher-degree bounds are $O(Q_{\mathrm{loc}}^2)$ and hence $O(Q_{\mathrm{loc}}^r)$ in degree $r \ge 2$. The positive compensator has constant coefficient one and linear deviation $o(1)$, so adjoining it preserves this degree envelope. The bound $\pi_0 \ge \mathsf{u}/2$ converts a degree-r uniform-reference profile bound to the primary reference at cost at most $2^{r/2}$.

The two summands of the paired lift are centered and depend on independent public inputs. Orthogonality and the compensator bound give

$$\|J_{\mathrm{in}} f_{\mathrm{in}}\|_2^2 \le (1 + C\gamma_{\mathrm{cal}})\|f_{\mathrm{in}}\|_2^2.$$

Every atom of the paired reference $\pi_0 \otimes \mathsf{p}_0$ on its support has mass at least $(4q^2)^{-1}$. This atom lower bound gives evaluation operator norm at most $2q$ at every complete pair. Combined with the lift bound, this proves the evaluation-row claim, including Hilbert-valued retained arguments. Tensor extension applies to arbitrary joint kernels. □

Theorem G.5 (Complete weighted Fourier bound)**.** Fix a positive integer R. For all sufficiently large L, let f be a canonical public-control coefficient with positive degrees at most R. Let $\mathcal{C}_\rho$ be its graded weighted L^2 norm with body weight $\rho^{b_{\mathrm{B}}}$, where $\rho = L^{1/4}$. On the protected cycle proposal put $W = 1 + \mathsf{U}J_{\mathrm{in}}\mathsf{f}$. Then the direct character polynomial of the signed output law satisfies

$$\|\mathsf{P}_W\|_{\mathcal{F},\rho} \le C_R(1 + \mathcal{C}_\rho(\mathsf{f})).$$

The result concerns the signed weight before charge and bounded clipping.

Proof. The character polynomial depends linearly on the weight. We therefore bound the unweighted proposal and then the contribution of each arity of f. For a correction term, the selected control groups affect only a bounded number of endpoints; the remaining endpoints are handled together by the all-degree product bound. Throughout the proof, the normalization scale is the full endpoint count d, even when some groups are removed.

The unweighted proposal. Fix the disjoint placement, the body permutation, the complete protected root vector, and the shared queue scalar. By Lemma G.4, the conditional character polynomial factors into $g = t^{16}$ triple polynomials and $n_{\mathrm{F}} + q = n_{\mathcal{B}} - g$ singleton polynomials. Each triple keeps the two leg endpoints together with their own attachment. All queue singletons use the same fixed translation.

We check the two inputs to Lemma G.3: small local activities and an exponential-moment bound for the normalized linear mean M. Each singleton has deviation $\delta_0 = O(Lt^{-4})$, while the triple coefficients are bounded by Lemma G.4. Using $gQ_{\mathrm{loc}}^2/d = \gamma_{\mathrm{cal}}$ in the activity definitions gives

$$\begin{aligned}\rho^2\eta &\le C\rho^2(\gamma_{\mathrm{cal}} + \delta_0^2) = O(L^{-1/2}),\\ \rho^2 A_2 &\le C\rho^2\gamma_{\mathrm{cal}} = O(L^{-1/2}),\\ \sqrt{g}\rho^3 A_3 &\le C(\rho^2\gamma_{\mathrm{cal}})^{3/2} = O(L^{-3/4}).\end{aligned}$$

These bounds are uniform in the conditioned data and still tend to zero at every fixed multiple of ρ. The cycle-mean estimate in Lemma G.1, with the fixed-queue drift verified immediately after that lemma, gives

$$\mathbb{E}\exp(C\rho^2 M^2) = O(1)$$

for every fixed positive C, now averaging the protected latent cycle while keeping the placement, body permutation, and shared queue scalar fixed. Thus Lemma G.3 bounds the conditional product in mean square over the cycle. Jensen's inequality allows averaging the cycle and then the body permutation and shared queue scalar, proving the claim for $W = 1$.

The selected control groups. For the local estimates, condition again on the root vector, body permutation, and shared queue scalar. Fix an arity $(a_{\mathrm{P}}, b_{\mathrm{B}})$ with $1 \le a_{\mathrm{P}} + b_{\mathrm{B}} \le R$, and one choice of distinct groups in the corresponding U-statistic. Call these groups tagged. A tagged primary group consists of its triple and its assigned compensator singleton; a tagged ordinary group consists of one ordinary-body singleton. These endpoint sets are disjoint by Section 9.2.

Canonical centering lets us apply the centered-score operators of Lemma G.4 in every input slot. A tagged primary group has zero constant coefficient and can emit degrees $1 \le r \le 4$. After the radius and endpoint normalization, its degree-r operator costs at most

$$C\left(\frac{\rho Q_{\mathrm{loc}}}{\sqrt{d}}\right)^r = C\left(\frac{\rho\sqrt{\gamma_{\mathrm{cal}}}}{\sqrt{g}}\right)^r = O({m_{\mathrm{P}}}^{-1/2}).$$

Here $m_{\mathrm{P}} = g$ and $\rho\sqrt{\gamma_{\mathrm{cal}}} = O(L^{-1/4})$; the absence of degree zero ensures at least one factor ${m_{\mathrm{P}}}^{-1/2}$. An ordinary group emits only degree one, with raw operator norm $O(1)$, so its normalized cost is $O(\rho/\sqrt{m_{\mathrm{B}}})$, using $m_{\mathrm{B}} \asymp d$ from Section H.4.

The Hilbert tensor product of these local operators acts on the joint kernel $\mathsf{f}_{a_{\mathrm{P}},b_{\mathrm{B}}}$. Its norm is bounded by the product of the local operator norms, whether or not the kernel factors across slots. The total tagged output degree is at most $4a_{\mathrm{P}} + b_{\mathrm{B}} \le 4R$. Consequently, summing its degree allocations, symmetrizing the output tensors, and including the degree factorials in the Fourier norm cost only a constant depending on R.

Combining tagged and untagged outputs. To recover the full correction polynomial, multiply the tagged output by the unweighted product on the remaining endpoints. Delete each tagged triple together with its compensator, and each tagged ordinary singleton. Leave every unselected triple and singleton as a separate block, so the remaining product still has blocks of size at most three, as required by Lemma G.3.

Deleting blocks decreases all three activities. Since at most R control groups were selected, it changes the normalized linear mean by at most $C_R Q_{\mathrm{loc}}/\sqrt{d}$. The radius-scaled square of this change is $O_R(\rho^2\gamma_{\mathrm{cal}}/g) = o(1)$, so the remaining mean has the same bounded exponential moment over the cycle. The product lemma therefore also controls the remaining product in mean square.

For each total tagged degree, combine all tagged coefficients into one tensor and apply (13) once. This uses the remaining product at radius $\sqrt{2}\rho$; the proof of Lemma G.3 uses at most 2ρ.

Both are fixed multiples covered by the activity and mean bounds above. Thus the radius does not grow with the number of untagged blocks.

Average over the complete roots, the shared queue scalar, and the body permutation using Jensen's inequality. In particular, multiply all queue factors before averaging their one common scalar. The resulting norm bound for this specified selection is

$$C_R {m_{\mathrm{P}}}^{-a_{\mathrm{P}}/2} \left(\frac{\rho}{\sqrt{m_{\mathrm{B}}}} \right)^{b_{\mathrm{B}}} \|\mathsf{f}_{a_{\mathrm{P}},b_{\mathrm{B}}}\|_2.$$

The bound is uniform in the fixed disjoint placement, so it also survives averaging the joint placement law, including the dependence between attachment and queue positions.

The physical U-statistic is a normalized average over these distinct-group selections. The triangle inequality preserves the displayed bound without an additional counting factor. Its right-hand side is precisely the contribution of this arity to $C_R \mathcal{C}_\rho(\mathsf{f})$, by Section F.2.1. Summing over arities and adding the unweighted bound proves the theorem. □

G.4 Canonical moment transfer from independent control pools

Canonical means that integrating any one kernel argument against its reference law gives zero with all other arguments fixed. The public control pools are large enough relative to their smallest atoms that an elementary finite-alphabet estimate suffices.

Lemma G.6 (Finite-alphabet canonical moments)**.** Consider two independent pools of positive sizes m_{P} and m_{B}, each consisting of independent identically distributed variables on a finite space. Let their smallest positive atom probabilities be α_{P} and α_{B}. Let f be a real canonical, separately symmetric kernel of fixed arities $(a_{\mathrm{P}}, b_{\mathrm{B}})$. Suppose the pool sizes are at least their respective arities and, for some fixed $c_{\mathrm{at}} > 0$,

$$m_{\mathrm{P}}\alpha_{\mathrm{P}} \geq c_{\mathrm{at}}, \qquad m_{\mathrm{B}}\alpha_{\mathrm{B}} \geq c_{\mathrm{at}}.$$

Conditions on a zero-arity pool may be omitted. For every fixed positive integer p, the distinct-index average normalized by $\binom{m_{\mathrm{P}}}{a_{\mathrm{P}}}\binom{m_{\mathrm{B}}}{b_{\mathrm{B}}}$ over unordered subsets satisfies

$$\|\mathsf{U}\mathsf{f}\|_{2p} \leq C_{a_{\mathrm{P}},b_{\mathrm{B}},p,c_{\mathrm{at}}} {m_{\mathrm{P}}}^{-a_{\mathrm{P}}/2} {m_{\mathrm{B}}}^{-b_{\mathrm{B}}/2} \|\mathsf{f}\|_2.$$

The kernel norm uses the product reference law. The constant is independent of alphabet and pool sizes and of the individual atom probabilities subject to the displayed lower bounds.

Proof. First consider a single column of n_{slot} independent samples X_j from a finite law with smallest positive atom α_{at}. For a centered function f, put $\mathcal{T}f = {n_{\mathrm{slot}}}^{-1} \sum_{j=1}^{n_{\mathrm{slot}}} f(X_j)$. In its $2p$th moment expansion, any sample index occurring once makes the expectation zero. A remaining term with $k_{\mathrm{eq}} \leq p$ distinct indices has absolute expectation at most $\alpha_{\mathrm{at}}^{-(p-k_{\mathrm{eq}})} \|f\|_2^{2p}$, since, on the reference support, $\|f\|_\infty \leq \alpha_{\mathrm{at}}^{-1/2} \|f\|_2$. Each pattern has at most ${n_{\mathrm{slot}}}^{k_{\mathrm{eq}}}$ index assignments, and the number of patterns depends only on p. Hence

$$\mathbb{E}|\mathcal{T}f|^{2p} \leq C_p {n_{\mathrm{slot}}}^{-p} \|f\|_2^{2p} \sum_{j=0}^{p-1} (n_{\mathrm{slot}}\alpha_{\mathrm{at}})^{-j}.$$

Thus the column-average operator from centered L^2 functions to L^{2p} has norm $O_{p,c_{\mathrm{at}}}({n_{\mathrm{slot}}}^{-1/2})$ whenever $n_{\mathrm{slot}}\alpha_{\mathrm{at}} \geq c_{\mathrm{at}}$.

To tensorize over independent columns, apply the first bound for each fixed sample array of the second, use Minkowski's inequality $\|\cdot\|_{L^{2p}(L^2)} \leq \|\cdot\|_{L^2(L^{2p})}$, then apply the second bound for each fixed reference input of the first. Iteration multiplies the bounds for any joint kernel, since averaging other columns preserves centering in the remaining slots.

For each positive arity, divide its pool uniformly into that many labeled disjoint blocks of equal size, leaving any remainder unused. Each block size lies between the pool size divided by twice its arity and the pool size divided by its arity. The atom condition therefore persists with a lower bound depending only on the fixed arities and c_{at}. Apply the tensorized bound to the rectangular average selecting one index per block. Its division-average is exactly the distinct-index average, since every ordered distinct tuple has the same inclusion probability. Jensen completes the proof. The zero-total-arity case is immediate. □

Corollary G.7 (Moment transfer to physical corrections)**.** For every fixed positive integer p and all sufficiently large L, the signed weight of Theorem G.5 satisfies

$$\|W - 1\|_{2p} \leq C_{R,p}\mathcal{C}_\rho(\mathsf{f}).$$

The estimate holds before charge and clipping, uniformly in the fixed disjoint placement.

Proof. Take L large enough that $\rho \geq 1$ and both control pools have size at least R. At each fixed placement, the uncharged proposal has $m_{\mathrm{P}} = g$ independent primary–compensator pairs with law $\pi_0 \otimes \mathsf{p}_0$ and an independent pool of m_{B} ordinary centers with law p_0, by Section 9.2. Their smallest positive atoms satisfy

$$\alpha_{\mathrm{P}} \geq (4q^2)^{-1}, \qquad \alpha_{\mathrm{B}} \geq (2q)^{-1}$$

by Theorem 9.3. The parameter bounds $q^2/m_{\mathrm{P}} = O(1)$ and $q/m_{\mathrm{B}} = O(t^{-12})$ from Section H.4 supply a common fixed positive c_{at} in Lemma G.6.

The paired lift preserves centering in each complete pair and has Hilbert operator norm $O(1)$ by Lemma G.4. Its tensor extension therefore sends any joint kernel of arities $(a_{\mathrm{P}}, b_{\mathrm{B}})$ to a canonical kernel on these pools with L^2 norm at most an order-dependent constant times $\|\mathsf{f}_{a_{\mathrm{P}},b_{\mathrm{B}}}\|_2$. Applying Lemma G.6 to that lifted kernel gives

$$\|\mathsf{U}J_{\mathrm{in}}\mathsf{f}_{a_{\mathrm{P}},b_{\mathrm{B}}}\|_{2p} \leq C_{R,p}{m_{\mathrm{P}}}^{-a_{\mathrm{P}}/2}{m_{\mathrm{B}}}^{-b_{\mathrm{B}}/2}\|\mathsf{f}_{a_{\mathrm{P}},b_{\mathrm{B}}}\|_2.$$

Sum over arities and use $\rho \geq 1$ to obtain the stated coefficient-norm bound, uniformly over placements. The moment order affects only the constant. □

G.5 Bounds for the calibrated weight and its full Fourier tail

Proof of Proposition 9.6. For the calibrated coefficient, $\mathcal{C}_\rho(\mathsf{f}_R) = \|\mathbf{c}\|_{\mathrm{simp},\rho} = O_R(L^{-3})$. Substitution into Theorem G.5 and Corollary G.7 gives the complete Fourier norm bound and every fixed moment bound. The analysis radius of Section 9.9 is independent of either order. The constants and the sufficiently large parameter threshold are uniform over placements.

By Theorem 9.5, the protected endpoint law has mass one and, after independent uniform final relabeling, is exchangeable with all nonzero support frequencies through R canceled. For sufficiently

large L, $s_{\mathrm{F}} = \rho/\sqrt{e} > 1$. Using the Fourier transform of this law and Lemma E.2 gives

$$\begin{aligned}\sum_{\mathbf{a}\neq 0} |\widehat{\mathsf{P}}(\mathbf{a})|^2 &= \sum_{|\operatorname{supp}\mathbf{a}|>R} |\widehat{\mathsf{P}}(\mathbf{a})|^2 \\ &\le s_{\mathrm{F}}^{-2(R+1)} \mathcal{E}_{s_{\mathrm{F}}}(\mathsf{P}_{W_R}) \\ &\le s_{\mathrm{F}}^{-2(R+1)} \left(\mathfrak{F}_\rho(\mathsf{P}_{W_R})^2 - 1\right) \\ &\le C_R L^{-(R+1)/2}.\end{aligned}$$

The sum includes every remaining frequency, up to full support. Applying Lemma E.1 and enlarging the order-dependent constant gives the stated ambient half-L^1 bound. □

H Charge Conditioning and Bounded Positive Sampling

We prove the charge and clipping guarantees of Section 9.8, verify the parameters, and establish the bounded implementation and accuracy of Section 9.10. Signed expectations and variation distance follow the conventions of Section 9.1.

H.1 The public charge test

Each body cycle has integer sum zero. A leg's two off-body residues have sum $q-1+\zeta-t$ modulo q, and its attachment adds t. The queue adds q times its scalar. Thus the public charge test of Section 9.8 depends only on core centers and leg shifts and is equivalent, pointwise in all remaining randomness, to the auxiliary endpoint lying in H_0.

H.2 Clipping and conditioning from ambient error

Proof of Theorem 9.7. Fix the placement. Let P_V be the unnormalized endpoint image of the clipped core measure $V\nu_0$. Pushing the signed and clipped core measures through the same protected endpoint channel, including final relabeling, contracts their ℓ_1 difference. Thus

$$\|\mathsf{P}_V - \mathsf{P}\|_1 \le \mathbb{E}|W - V| = \delta_{\mathrm{clip}}.$$

The triangle inequality gives

$$\|\mathsf{P}_V - \mathsf{U}\left(\mathbb{F}_q^d\right)\|_1 \le 2\delta_{\mathrm{amb}} + \delta_{\mathrm{clip}} = \epsilon_{\mathrm{ch}}.$$

The public charge test is pointwise equivalent to membership in H_0, by Section H.1. Hence $\mathsf{P}_V(H_0) = b_{\mathrm{ch}}$. The ambient uniform law assigns this hyperplane mass $1/q$, so

$$|b_{\mathrm{ch}} - 1/q| \le \epsilon_{\mathrm{ch}}, \qquad b_{\mathrm{ch}} \ge \frac{1}{q} - \epsilon_{\mathrm{ch}} \ge \frac{1}{2q}.$$

This is stronger than the asserted mass lower bound. Restriction cannot increase ℓ_1 error. Normalize the restricted clipped measure by b_{ch} and the restricted uniform measure by $1/q$. Adding and subtracting the latter restriction divided by b_{ch} gives

$$\begin{aligned}\Delta(\mathsf{P}_{V,\xi}, \mathsf{U}(H_0)) &\le \frac{\epsilon_{\mathrm{ch}} + |b_{\mathrm{ch}} - 1/q|}{2b_{\mathrm{ch}}} \\ &\le \frac{\epsilon_{\mathrm{ch}}}{b_{\mathrm{ch}}} \le 2q\epsilon_{\mathrm{ch}}.\end{aligned}$$

□

Proof of Corollary 9.8. Fix the calibration order. For every fixed positive integer p, Corollary 7.3 and Proposition 9.6 give

$$\delta_{\mathrm{clip}} \le \mathbb{E}|W-1|^{2p} \le C_{R,p} L^{-6p},$$

uniformly over the disjoint placements. The calibrated weight and its proposal do not depend on the moment order. Taking a sufficiently large fixed moment order for each desired inverse-polynomial bound proves that this same clipping loss is negligible, uniformly over placements. Since $q = O(L^{64})$, its charge-amplified loss is negligible as well. □

H.3 Bounded clipping for a fixed placement

Lemma H.1 (Bounded rejection and approximation budget)**.** Under the hypotheses of Theorem 9.7, the ideal rule samples the clipped positive core law upon acceptance, with acceptance rate at least $1/(8q)$. For $0 < \epsilon_{\mathrm{cap}} < 1$, a cap

$$T_R = \lceil 8q \log(1/\epsilon_{\mathrm{cap}}) \rceil$$

fails with probability at most ϵ_{cap}.

Let δ_{unn} bound the ℓ_1 error in a numerical nonnegative charged measure, whose weight also lies in $[0,2]$. If $\delta_{\mathrm{unn}} \le 1/(8q)$, its normalized core law is within $4q\delta_{\mathrm{unn}}$ of the ideal clipped law in statistical distance. Its acceptance rate with denominator two is at least $1/(16q)$, so replacing 8 by 16 in the cap gives the same failure bound. For a proposal at statistical distance at most δ_{prop} from ν_0 and a weight satisfying $\|\widetilde{W} - W\|_\infty \le \delta_W$, clipping and charge restriction permit the choice

$$\delta_{\mathrm{unn}} = 4\delta_{\mathrm{prop}} + \delta_W.$$

Proof of Lemma H.1. For the fixed placement, the ideal acceptance rate is $p_{\mathrm{acc}} = b_{\mathrm{ch}}/2 \ge 1/(8q)$. Independence of the trials gives

$$\Pr[\text{timeout}] = (1-p_{\mathrm{acc}})^{T_R} \le e^{-p_{\mathrm{acc}} T_R}.$$

Conditional on acceptance, the core has the clipped charged law, proving the ideal cap and output claims.

Suppose the rational proposal $\nu_{0,\mathrm{num}}$ is within δ_{prop} of ν_0 in statistical distance, and $\|\widetilde{W} - W\|_\infty \le \delta_W$. Clip the rational weight by $\mathrm{clip}_{[0,2]}$. Changing the proposal costs at most $4\delta_{\mathrm{prop}}$ in ℓ_1 because the ideal clipped weight lies in $[0,2]$; changing the weight costs at most δ_W because clipping is 1-Lipschitz. Charge restriction cannot increase either error, so the unnormalized charged measures have ℓ_1 discrepancy at most

$$\delta_{\mathrm{unn}} := 4\delta_{\mathrm{prop}} + \delta_W.$$

Normalize these nonnegative measures as in the proof of Theorem 9.7. Its ideal mass bound $1/(4q)$ gives normalized statistical distance at most $4q\delta_{\mathrm{unn}}$. If $\delta_{\mathrm{unn}} \le 1/(8q)$, the numerical mass is at least $1/(8q)$, and acceptance with denominator two has probability at least $1/(16q)$. The same geometric estimate proves the numerical cap. □

Retaining the initial placement through retries and timeout fallback preserves its marginal despite placement-dependent acceptance rates. Conditional completion follows Section 9.2.

H.4 Verification of the public parameter family

Geometry and local feasibility. Use $t = L^8$ from Section 9.9. The public sizes give

$$\frac{q+2g}{n_{\mathcal{B}}} = O(t^{-4}), \qquad m_{\mathrm{P}} = t^{16}, \qquad m_{\mathrm{B}} \asymp n_{\mathcal{B}} \asymp d \asymp t^{20}.$$

Thus disjoint placement and distinct compensators exist eventually. The finite-atom losses are absorbed by the pools: $q^2/m_{\mathrm{P}} = O(1)$ and $q/m_{\mathrm{B}} = O(t^{-12})$. These bounds verify the full paired-alphabet moment estimate in Corollary G.7. The three source cutoffs sum to the safe-center margin, while $L^2 \ll a_{\mathrm{anc}} = t^6 \ll q$ verifies the anchor support conditions.

For local affine-map feasibility, the protected input count satisfies $r_{\mathrm{L}} = (2L^2-1)^2 \sim 4L^4$. Substituting $B_{\mathrm{L}} = 3t^6$ and $q \geq 2^{16}t^8$ into the three loss quantities of Section 8.5.3 gives

$$\begin{aligned}
\delta_{\mathrm{aff}} &= O(L^4 t^{-2}) = O(L^{-12}),\\
\kappa_{\mathrm{aff}} &= O(L^8 t^{-2}) = O(L^{-8}),\\
A_{\mathrm{loc}} &= O(L^{12} t^{-2}) = O(L^{-4}).
\end{aligned}$$

All three losses tend to zero, so the local feasibility and positivity conditions hold for all sufficiently large L.

Proposal and response bounds. Substitution in the symbolic estimates of Section F.1.1 gives

$$\begin{gathered}
\gamma_{\mathrm{cal}} = \Theta(L^{-1}), \qquad \varepsilon_{\mathrm{anc}} = O(L^{-1}t^{-4}),\\
\delta_{\mathrm{base}} + \delta_{\mathrm{ctr}} + \delta_{\mathrm{prim}} = O(Lt^{-2}),\\
\|\pi_0\|_\infty \leq \frac{d}{g\gamma_{\mathrm{cal}} q} = O(Lt^{-4}),\\
\gamma_{\mathrm{cal}} \|\widetilde{\mathsf{A}}\pi_0\|_\infty = O(q^{-1}),\\
\|\mathsf{p}_0 - \mathsf{u}\|_{\mathsf{u}^{-1}} + \|\mathsf{u}_{\mathrm{safe}} - \mathsf{u}\|_{\mathsf{u}^{-1}} = O(Lt^{-4}).
\end{gathered}$$

These bounds imply proposal positivity. The first inverse is bounded because

$$\gamma_{\mathrm{cal}}{}^{-1}\sqrt{a_{\mathrm{anc}}/q} = O(L/t) = o(1).$$

The local estimates of Sections F.2 and G then give an $O_R(L^{-3})$ coefficient correction, the same decay for every fixed physical moment norm, and a bounded complete Fourier norm. Constants may depend on the fixed orders. With the common local scale of Lemma G.4, the three product activities at radius $\rho = L^{1/4}$ are $O(L^{-1/2})$, $O(L^{-1/2})$, and $O(L^{-3/4})$, respectively. They remain small at 2ρ, as required by Lemma G.3 and Theorem G.5.

Source protection and the finite prefix. The ratio of each cutoff to σ grows linearly in L. Combining the body coupling bound in (8) with the tail estimates in Lemma D.1 therefore gives

$$\theta_{\mathrm{B}} + \varepsilon_{\mathrm{B,clip}} + 2g\theta_{\mathrm{L}} + \theta_{\mathrm{Q}} \leq \mathrm{poly}(L) e^{-\Omega(L^2)}.$$

On any finite prefix excluded by analytic or numerical certification, use the fallback of Section H.5. It is integral and invertible, and carry-safe on protected inputs. No accuracy guarantee is asserted on that prefix.

H.5 Numerical verification and deterministic bounds

The sampler uses certified rational intervals and bounded finite-bit draws. For fixed R, the certificates below prove polynomial running time, exact safe support, and implementation error at most ϵ_{num}.

Algorithm 3 Bounded public-control matrix sampler

Require: Fixed accuracy order c and the public parameters of Section 9.9.
Ensure: Integral affine data independent of the hidden source.

1: Fix the order and budgets of Section 9.10, and the precision, table-size and work caps of Section H.5.
2: Draw one disjoint placement ξ, pair its labeled legs with distinct free compensators, and retain this placement.
3: Construct the proposal and finite response matrices; solve for the signed correction by the finite block computation of Proposition F.6, certifying the proposal, inverse gaps, and numerical bounds. If certification fails, return the deterministic carry-safe affine pair of Section H.5.
4: For at most $16q\lceil \log_2(5/\epsilon_{\mathrm{num}})\rceil$ trials, draw a fresh core at this placement and implement the threshold $\mathbf{1}_{\mathcal{E}_\xi}\,\mathrm{clip}_{[0,2]}(\widetilde{W})/2$ to the certified finite-bit precision. Retain the first accepted core.
5: If no core was accepted, use the charge-correct regular-only core of Section H.5.
6: Draw the complete conditional flag–affine-map pairs at the retained core, then independently draw the actual closure Gaussian, body permutation, and final global permutation with bounded samplers.
7: Assemble and return $(\mathbf{Q}, \beta)$ by Proposition 8.19.

Certified precomputation. Certification in Algorithm 3 checks Lemma E.3 and Theorem 9.3 before using population normalizers or calibration inverses. It uses the rational reference laws below, averages complete fixed-queue expressions before inversion as in Section F.2.5, and checks the inverse gaps of Proposition F.6. The following bounds certify precision, work caps, and fallbacks.

Finite coordinates, matrices, and arithmetic. For fixed R, the canonical public-control coordinates of Definition F.3, the multiplication structure constants, and the normalized response blocks are obtained from the same finite pattern contractions. Their total stored size satisfies

$$S_R \le (q n_{\mathcal{B}} K_{\mathrm{L}} a_{\mathrm{anc}})^{C_R}.$$

Each input or output pattern has $O(R)$ slots. Enumerate its selected vertices, their role intersections, and their relative positions on the cycle. There are $n_{\mathcal{B}}^{O(R)}$ possibilities. For each path cluster sum the finitely supported Gaussian masses on at most $O(R)$ latent edges. This requires $K_{\mathrm{L}}^{O(R)}$ operations. Anchors, colors, public blurs, and the shared queue scalar add factors at most $(qK_{\mathrm{L}}a_{\mathrm{anc}})^{O(R)}$. The distinct-index weights are ratios of falling factorials. Multiplication matrices, finite logarithms and reciprocals, and the block solve in Proposition F.6 use polynomially many operations in this stored size, retaining the queue value until its prescribed average. The computation enumerates only these fixed-order patterns.

Gaussian probabilities are computable reals, so the implementation uses certified rational approximations. Every nonzero protected edge mass has negative logarithm $O(L^2 + \log L)$. A fixed-order joint mass has the same bound with an order-dependent constant. Source contractions can therefore use joint masses directly; if likelihood ratios are used, $O_R(L^2)$ extra precision bits suffice to control these possibly exponentially small denominators. The primary-color and free-center proposal probabilities are at least $1/(2q)$ on their supports.

The diagonal inverse in Proposition F.6 is bounded in the graded coefficient norm. At each fixed order, the lower reference atoms and the finite-atom inequality of Lemma F.4 give, for every hybrid type,

$$\begin{aligned}\|\mathsf{f}_{a_{\mathrm{P}},b_{\mathrm{B}}}\|_\infty \le (2q)^{(a_{\mathrm{P}}+b_{\mathrm{B}})/2}\rho^{-b_{\mathrm{B}}} \\ \cdot {m_{\mathrm{P}}}^{a_{\mathrm{P}}/2}{m_{\mathrm{B}}}^{b_{\mathrm{B}}/2}\mathcal{C}_\rho(\mathsf{f}).\end{aligned}$$

Conversely, a supremum bound controls every product L^2 norm. The coordinate basis changes of Section F.2.1 and these explicit population factors are polynomial at fixed order, so the induced coordinate inverse norm is polynomially bounded. These factors affect numerical point evaluation; the physical moment and complete Fourier bounds remain as stated.

The finite-atom bound for a complete paired control, the integrated paired-lift estimate of Lemma G.4, and the normalized U-statistic average give polynomial uniform propagation of coefficient errors at fixed order. Evaluation of $\mathbf{\Psi}_R(\theta)\mathbf{z}$ enumerates distinct tuples of at most R controls and the finite anchor contractions in each paired lift. It takes $(dq)^{O(R)}$ arithmetic operations.

Rational approximation with a law-level budget. Use outward-rounded rational intervals for Gaussian masses and all finite operations. Approximate positive masses on their exact supports and normalize the resulting rational vectors; the primary and center margins in Theorem 9.3 permit this while retaining, for example, lower bounds $1/(4q)$ and one quarter of the safe uniform. Under these rational reference laws, canonicalize every table by the exact commuting projections $I - \mathbb{E}_{\text{one input slot}}$. The scalar component is retained separately. Use the same reference colors and ordered basis indices for the ideal and rational coordinate constructions. Rebuild the paired lift, response tables, and physical evaluation row using these same rational laws and normalized rational local kernels. In particular, let $\pi_{0,\mathrm{num}}, \mathsf{p}_{0,\mathrm{num}}$ be the two rational control laws, let $\mathbf{A}_{\mathrm{anc,num}}$ be the rational smoothed anchor kernel, normalized column by column on its exact safe support, and let γ_{num} approximate the activation. The rational paired lift is

$$(J_{\mathrm{in,num}}\mathsf{f})(c,a) = \mathsf{f}(c) - \gamma_{\mathrm{num}}\frac{(\mathbf{A}_{\mathrm{anc,num}}\,\mathrm{diag}(\pi_{0,\mathrm{num}})\mathsf{f})(a)}{\mathsf{p}_{0,\mathrm{num}}(a)}.$$

The denominator is positive on this support. Since the anchor columns sum exactly to one,

$$\mathbb{E}_{\pi_{0,\mathrm{num}}\otimes\mathsf{p}_{0,\mathrm{num}}}J_{\mathrm{in,num}}\mathsf{f} = (1-\gamma_{\mathrm{num}})\mathbb{E}_{\pi_{0,\mathrm{num}}}\mathsf{f} = 0$$

for every rationally centered input. Slotwise tensoring and distinct-index averaging therefore preserve exact centering of the physical rational correction. Merely projecting the coefficient tables while retaining an inconsistent lift would not suffice. The uniform weight error below includes changes to reference laws, basis conversions, projections, and lifts. Exact Fourier cancellation is required only for the ideal weight.

Propagate interval errors through the finite arithmetic circuit using the displayed operand and denominator bounds. For each degree-diagonal block, let κ_{diag} bound its induced coordinate ℓ_∞ inverse norm, and let δ_{diag} bound the maximum entrywise approximation error. Require

$$m_R\delta_{\mathrm{diag}} \le \frac{1}{2\kappa_{\mathrm{diag}}}.$$

The factor m_R converts entrywise error to operator error. Apply Theorem 7.1 after factoring out the ideal block: the approximate inverse has norm at most $2\kappa_{\mathrm{diag}}$. The inverse-difference identity bounds its error by $2\kappa_{\mathrm{diag}}^2 m_R\delta_{\mathrm{diag}}$. The polynomial coordinate dimension and inverse bound thus

certify an inverse tolerance for the fixed public family. The fixed-degree logarithms, multiplication matrices, finite reciprocals, and graded-inverse products have polynomial total degree; source ratios have at most $2^{O_R(L^2)}$ conditioning. These rules give an explicit upper bound

$$\kappa_{\mathrm{num},R} \le 2^{\mathrm{poly}_R(L)}$$

on the propagated error in any table entry or physical weight. Compute this bound from the circuit and positive margins. Precision $O(\log \kappa_{\mathrm{num},R} + \log(S_R/\delta_{\mathrm{tab}}))$ is polynomial in λ and $\log(1/\delta_{\mathrm{tab}})$ at fixed order. Rational Gaussian approximations can be produced by the convergent exponential power series with range reduction; its remainder bound certifies each endpoint interval.

Choose the proposal and weight errors to satisfy

$$\delta_{\mathrm{prop}} \le \frac{\epsilon_{\mathrm{num}}}{320q}, \qquad \delta_W \le \frac{\epsilon_{\mathrm{num}}}{80q}.$$

A product of at most $O(d)$ conditional elementary draws is within the sum of their uniform conditional errors, which determines the local table tolerances. Consequently $\delta_{\mathrm{unn}} \le \epsilon_{\mathrm{num}}/(40q)$. For all sufficiently large security parameters, uniformly over placements, Theorem 9.7 gives $b_{\mathrm{ch}} \ge 1/(4q)$. In this regime, Lemma H.1 bounds the normalized core error by $\epsilon_{\mathrm{num}}/10$ and certifies the numerical mass and acceptance bounds for $0 < \epsilon_{\mathrm{num}} < 1$.

Rational discrete sampling and acceptance are implemented by fixed-bit inverse-transform sampling. For a finite law of support size S_R, cumulative dyadic rounding with $O(\log(S_R/\epsilon))$ bits gives distance at most ϵ, while never selecting a forbidden support point. The acceptance threshold is $\mathbf{1}_{\mathcal{E}_\xi} \operatorname{clip}_{[0,2]}(\widetilde{W})/2$. Round it within $[0,1]$, keeping it exactly zero outside the public charge event and whenever the clipped value is zero. Budget these errors over the predetermined maximum number of draws and trials. In the same asymptotic regime, exact trials from the certified rational law satisfy the geometric timeout bound. Couple these capped trials with their finite-bit implementation: all rounding effects, including changes to acceptance or timeout, cost at most the sum of their conditional sampling errors. Reserve at most $\epsilon_{\mathrm{num}}/5$ for them. Use the integer cap

$$T_R = 16q \left\lceil \log_2(5/\epsilon_{\mathrm{num}}) \right\rceil .$$

Compute the ceiling by exact integer comparisons. Since $\log_2(5/\epsilon_{\mathrm{num}}) \ge \log(5/\epsilon_{\mathrm{num}})$, Lemma H.1 bounds timeout by $\epsilon_{\mathrm{num}}/5$ with $O(q \log(1/\epsilon_{\mathrm{num}}))$ trials. Reserve at most $\epsilon_{\mathrm{num}}/5$ for the subsequent complete flag–map draws, closure Gaussian, body permutation, and final relabeling, using Sections C.1.1, C.5.1 and E.2. Reserve a separate $\epsilon_{\mathrm{num}}/10$ for the initial joint placement permutation. This error is not charge-amplified: both ideal and approximate algorithms normalize separately for each retained placement, and total variation contracts under the common conditional output kernel. Unused budget covers the finite table and interval endpoints. This allocation includes charge-normalizer amplification.

Fallbacks and deterministic caps. If certification exceeds the prescribed polynomial work or bit caps, return the deterministic data

$$\begin{aligned} &\delta_{\mathrm{fb}} := [-dK_{\mathrm{BS}}]_q, \qquad (\mathbf{b}_{\mathrm{fb}})_v := K_{\mathrm{BS}} + \mathbf{1}_{\{v \le \delta_{\mathrm{fb}}\}}, \qquad 1 \le v \le d, \\ &\mathbf{Q} := \mathbf{I}, \qquad \beta := \mathbf{T}^{-1}\mathbf{b}_{\mathrm{fb}}. \end{aligned}$$

Since $0 \le \delta_{\mathrm{fb}} < q \le d$, the load sum $dK_{\mathrm{BS}} + \delta_{\mathrm{fb}}$ is divisible by q; Lemma 6.4 gives an integral offset. The linear part is invertible modulo q. Each row of $\mathbf{T}$ has absolute sum at most three, and

a protected source satisfies $\|\mathbf{e}\|_\infty \le K_{\mathrm{L}}$ by Definitions 6.3 and 8.16. As $K_{\mathrm{BS}} = 3K_{\mathrm{L}}$, the signed load obeys

$$0 \le (\mathbf{b}_{\mathrm{fb}} + \mathbf{T}\mathbf{e})_v \le 2K_{\mathrm{BS}} + 1 < q.$$

Its carry is zero; passing to the canonical source lift changes it by an element of $\mathcal{L}(\mathbf{T})$, preserving Lemma 6.6's criterion. Compute the last offset coordinate as $(dK_{\mathrm{BS}} + \delta_{\mathrm{fb}})/q$ and, for $1 \le j < d$, set

$$(\beta)_j = \sum_{v=1}^{j} \left((\mathbf{b}_{\mathrm{fb}})_v - (\beta)_d \mathbf{1}_{\{v \le q\}} \right).$$

The computation uses public data with polynomial work and bit length and preserves protected-input carry safety on the excluded prefix. The asymptotic bounds ensure eventual certification, using interval inequalities and rational inverse residuals.

Charge-correct timeout. For rejection timeout in the certified regime, use a charge-correct regular-only graph core: every body center is $K_{\mathrm{BS}} = 3K_{\mathrm{L}}$, every leg shift is 1, every marking weight is zero, and the primary target is $(q-1)/2$ with companion $(q+1)/2$. These companion targets remain interior after every protected second input. The charge is correct because q divides $n_{\mathcal{B}}$:

$$\sum_{v \in \mathcal{B}} b_v + \sum_j \zeta_j = 3n_{\mathcal{B}} K_{\mathrm{L}} + g \equiv g \pmod{q}.$$

The regular complete-map law preserves safe support and nonzero leg scales. Together with the verified charge constraint, Proposition 8.19 therefore gives integral affine data and modular invertibility on this fallback. Its error is charged only to timeout. Complete this core with independent closure and final permutations; certification failure instead returns the deterministic affine pair directly. Both fallbacks use public data only.

H.6 Completion of fixed-order affine fiber transport

Proof of Theorem 9.9. Proposition 9.6 bounds the ambient error by $O_R(L^{-(R+1)/4})$. The chosen calibration order and $q = O(L^{64})$ give $q\delta_{\mathrm{amb}} = o(1)$. By Corollary 9.8, $q\delta_{\mathrm{clip}} = o(1)$ as well, uniformly over placements. Thus $q\epsilon_{\mathrm{ch}} = q(2\delta_{\mathrm{amb}} + \delta_{\mathrm{clip}}) \le 1/2$ for all sufficiently large L. Theorem 9.7 then bounds the positive protected endpoint error by

$$O_R(qL^{-(R+1)/4}) + 2q\delta_{\mathrm{clip}}.$$

For each fixed calibration order, Corollary 9.8 makes the second term negligible, uniformly over placements. Normalization is separate at each disjoint placement; averaging adds no error.

Apply the joint coupling of Section D.2 to the positive source-independent core. On a match, the pointwise criterion of Lemma 6.6 identifies the affine output with the canonical pullback of the protected endpoint. By Lemma 6.5, the fiber error therefore incurs the single additional cost

$$\varepsilon_{\mathrm{B,clip}} + 2g\theta_{\mathrm{L}} + \theta_{\mathrm{Q}} \le \mathrm{poly}(L) e^{-\Omega(L^2)},$$

as verified in Section H.4. Absorb this source-protection error and the charge-amplified clipping error into the negligible function $\mathrm{negl}_R(L)$.

The numerical budget of Section H.5 includes charged-core normalization, finite-bit trials, timeout, placement, and completion. Its precision and work caps give polynomial worst-case time for each fixed order. Bertrand's bound gives $q = O(L^{64})$. Find the least prime deterministically by scanning the defining interval and testing each integer by trial division. The interval length and

divisor bounds are polynomial in λ, so computing the public modulus takes polynomial time. The calibration order in Section 9.10 satisfies $R \geq 4c + 267$, so

$$64 - \frac{R+1}{4} \leq -c - 3.$$

Thus the Fourier term is $O_R(L^{-c-3})$. Together with the negligible remainder and numerical budget, the total error is eventually at most λ^{-c}. □

I Proofs for the Applications

We prove the applications of Section 10 for the family in Section 9.9 and Definition 6.3.

I.1 The Gaussian interface and worst-case reduction

Proof of Corollary 10.1. For the fixed public family, $N = \lambda$, $\sigma = \lambda + 2$, and q is an odd prime polynomially bounded in N. Set $\alpha = \sigma/(\sqrt{2}q)$. Then $0 < \alpha < 1$ and $\alpha q = \sigma/\sqrt{2} > 2\sqrt{N}$ for $\lambda \geq 3$, as required by Regev [Reg05, full version, Theorem 3.1]. Moreover, t blocks contain td scalar samples, still polynomially many. It remains to match the Gaussian and algorithmic interfaces.

Regev's theorem uses continuous Gaussian error on the torus. Scale each response by q, apply Gaussian randomized rounding to $\mathbb{Z}$ with width $\sigma/\sqrt{2}$, and reduce modulo q. By [Pei10, full version, Theorem 3.1], the continuous and rounding squared widths, both $\sigma^2/2$, yield error within $\exp(-\Omega(\sigma^2))$ in SD of $\mathsf{D}_{\mathbb{Z},\sigma} \bmod q$. Integer-translation equivariance preserves the secret term and sends a uniform torus response to an exactly uniform residue. Truncation and polynomial-bit arithmetic give a bounded polynomial-time implementation with negligible additional error. Thus rounding supplies the source in Definition 2.11, with negligible loss over polynomially many samples even under QPT postprocessing, by the contraction bound in Section 2.3.

Fix a uniform QPT LAR distinguisher and polynomial bounds on its sample count and inverse advantage. Corollary 5.5 then chooses one fixed transport accuracy order, without changing the public family, and provides a discrete Gaussian distinguisher with nonnegligible advantage. Regev's secret-shift and decision-to-search reductions [Reg05, full version, Lemmas 4.1–4.2] apply to this error law; their probability estimates use fresh classical tables and independent QPT executions with measured output bits. Enumerate all positive table lengths within a computable polynomial bound. For each length, run the amplified coordinate-recovery acceptance-gap tests with fixed polynomial time bounds to obtain a candidate secret or report failure. Undo any secret shift and verify each candidate against λ fresh, independent scalar Gaussian LWE samples before accepting it: require every response minus its predicted inner product to have a centered integer lift of absolute value at most $q/4$. For the correct candidate, the Gaussian tail bound gives rejection probability at most $\lambda \exp(-\Omega(q^2/\sigma^2))$. For any incorrect candidate, conditioned on its construction, the fresh residuals are independent and uniform over $\mathbb{F}_q$, so its acceptance probability is at most $(1/2 + 1/q)^\lambda$. The length witnessing the distinguishing advantage occurs in the enumeration and yields the correct candidate with inverse-exponentially small failure. A union bound over polynomially many candidates preserves this failure bound for the verified search solver. Enumeration, amplification, and verification have polynomial overhead and preserve uniformity even when the sample-count function is not computable. Feeding the verified search solver rounded samples supplies the continuous-LWE solver.

For the coherent decoding call in [Reg05, full version, Lemmas 3.11 and 3.14], defer measurements in the amplified uniform QPT decoder, compute the decoder, copy its output, and uncompute; the promised output is unique. The coherent error is at most twice the square root of the

decoding failure probability. The Gaussian query state is independent of the stored Gaussian sample pool, so averaging its amplitude-weighted failure also over that pool gives negligible error. Fresh ancillary registers and a hybrid over polynomially many steps preserve negligible total error with polynomial overhead. The composition therefore gives uniform QPT algorithms for GapSVP and SIVP with approximation factor $\widetilde{O}(N/\alpha) = \widetilde{O}(Nq/\sigma)$, using only the distinguisher's guarantees on classical random examples. Accuracy and amplification choices follow the adversary's polynomial bounds; if its advantage holds only infinitely often, the algorithms succeed on the same subsequence of dimensions. □

I.2 Key sampling, exact marginals, security, and evaluation

Proof of Proposition 10.2. For every fixed nonzero key, a uniform $\mathbf{M}$ gives $\mathbf{Ms} \sim \mathsf{U}\left(\mathbb{F}_q^d\right)$. The balanced fibers of Lemma 3.5 therefore make each output exactly uniform in $\mathbb{F}_q$. For bounded-time key generation, sample a key index in $\left\{1, \ldots, q^N - 1\right\}$ by standard rejection sampling capped at λ attempts and use its base-q expansion as the key; on exhaustion return $\mathbf{e}_1$. The resulting law is supported on $\mathcal{K}$ and is within $\delta_{\mathrm{key}} := 2^{-\lambda}$ of $\mathsf{U}(\mathcal{K})$ in SD. Since the marginal property holds for every key in this support, it remains exact for the implemented sampler.

For a uniform QPT adversary, let $\mathrm{Adv}_{\mathrm{wPRF}}$ denote its random-example advantage against a uniformly random function. Conditioning the LAR secret to be nonzero costs q^{-N}, and implementing key sampling costs δ_{key}. Independent uniform labels couple to a random function unless inputs collide, at cost at most $\binom{t}{2} q^{-dN}$ [Sho04, Lemma 1]. Thus, for a suitable LAR distinguisher,

$$\mathrm{Adv}_{\mathrm{wPRF}} \leq \mathsf{Adv}_{\mathcal{D}}^{\mathsf{dLAR}}(\lambda, t) + q^{-N} + \delta_{\mathrm{key}} + \binom{t}{2} q^{-dN}.$$

The reduction invokes the adversary once on a classical table and is uniform QPT; the statistical bounds hold under quantum postprocessing by the contraction bound in Section 2.3. All three extra terms are negligible for every polynomial sample count, so Corollary 10.1 gives security.

Standard TC^0 circuits for multiplication, iterated addition, and division [HAB02] compute $\mathbf{Ms} \bmod q$. For the matrix in Definition 6.3, each local floor depends on at most three canonical residues and is computable by a polynomial-size constant-depth lookup circuit. Summation and modular reduction, adding a sufficiently large public multiple of q to handle negative values, then give the quotient label. All intermediate integers have $O(\log \lambda)$ bits, and the polynomially many operations occur in a constant number of parallel stages. Since the public parameters and lookup tables are polynomial-time computable, the resulting joint key-and-input TC^0 circuits are polynomial-time constructible. □

I.3 Prediction and regression distinguishers

Independent validation turns an inverse-polynomial field-label prediction advantage over $1/q$ into distinguishing advantage using polynomially many classical examples.

Proof of Corollary 10.4. Use the independent-label experiments of Definition 3.6, dividing canonical label representatives by q. Train the QPT learner on an initial sample block and measure its classical hypothesis. Let $\widetilde{g}$ be the evaluator's rational approximation to the clipped hypothesis, clipped once more to $[0, 1]$. Clipping does not increase squared loss, and the evaluator's $1/1000$ approximation error changes it by at most $1/500$. Thus, for every nonzero secret, successful training gives risk at most $1/100 + 1/500 = 3/250$. In the ideal experiment, each validation label is

independent of training and its prediction, even when inputs repeat, so the conditional risk is at least its variance $v_q = (q^2-1)/(12q^2) \geq 2/27$.

Reserve $m_{\mathrm{val}} := 10^4$ fresh samples and let $\widehat{\mathcal{R}}$ be their average squared loss for $\widetilde{g}$; accept exactly when $\widehat{\mathcal{R}} \leq 1/20$. Conditioned on the training transcript and measured hypothesis, and also on the secret in the real experiment, these losses are independent and lie in $[0,1]$. The threshold is more than $1/50$ from both risk bounds, so Hoeffding's inequality [Hoe63, Theorem 2] bounds each one-sided validation failure probability by $1/100$. Training succeeds with probability at least $2/3$ for every nonzero secret; the zero-secret event has probability q^{-N}. Consequently, the distinguishing advantage is at least

$$(1-q^{-N})\frac{2}{3}\frac{99}{100} - \frac{1}{100} > \frac{1}{2}$$

for all sufficiently large λ.

The learning model bounds the learner's time on all transcripts, including unrealizable ideal transcripts, and supplies a uniform deterministic classical evaluator for every output hypothesis. Its rational predictions and the validation arithmetic have polynomial bit length. QPT training followed by this classical validation therefore gives a uniform QPT distinguisher with time and total sample count polynomial in D, hence in λ. Corollary 10.1 then gives the hardness consequence for the same public family. □

I.4 The two-hidden-layer representation

Proof of Lemma 10.3. Fix a nonzero key. Integer-lift invariance from Lemma 3.5, with the quotient sum in Definition 6.3, gives

$$H_{\mathbf{s}}(\mathbf{M}) := \sum_{j=1}^{d} \left\lfloor \frac{(\mathbf{T}[\mathbf{M}]_q[\mathbf{s}]_q)_j}{q} \right\rfloor, \qquad g_{\mathbf{s}} = \frac{H_{\mathbf{s}} \bmod q}{q}.$$

Each numerator is an integer linear form in the input of magnitude at most $3N(q-1)^2$. Put $R := 4Nq^2$ and $\mathrm{ReLU}(z) := \max\{0,z\}$. For every integer $|z| \leq R$, the exact identity

$$\left\lfloor \frac{z}{q} \right\rfloor = -R + \sum_{u=-R+1}^{R} \big(\mathrm{ReLU}(z-qu+1) - \mathrm{ReLU}(z-qu)\big)$$

expresses each floor by one ReLU layer. Indeed, each difference of ReLUs equals $\mathbf{1}_{\{z \geq qu\}}$ on integer inputs. Counting the thresholds gives $R + \lfloor z/q \rfloor$, including negative inputs. There are $4R$ units per row and hence at most $4dR$ units in the first hidden layer. Its affine scalar output defines the extension of $H_{\mathbf{s}}$. On the grid this output is an integer and satisfies $|H_{\mathbf{s}}| \leq dR$. Define ψ by linear interpolation of $z \mapsto (z \bmod q)/q$ at all integers in this interval, and extend it constantly outside the interval. It is globally $[0,1]$-valued. To make its network representation explicit, let a_u be its slope on $(u, u+1)$ for integer $-dR \leq u < dR$, and put $a_{-dR-1} = a_{dR} = 0$. The standard piecewise-linear representation recalled in [CGKM22, Remark 2.2] is, in this case,

$$\psi(z) = \psi(-dR) + \sum_{u=-dR}^{dR} (a_u - a_{u-1})\,\mathrm{ReLU}(z-u).$$

The slopes have absolute value at most one and denominator q. Substituting the first layer's affine output for z inside these ReLUs gives a second hidden layer with at most $2dR+1$ units and an affine output. The intermediate affine map is absorbed into the second-layer preactivations, so the composition has exactly two hidden layers.

The total number of hidden units is at most $6dR + 1 = 24Dq^2 + 1 = O(D^3)$. First-layer input weights have magnitude at most $q-1$ and biases at most $qR+1$. The intermediate affine output has coefficients in $\{-1, 1\}$ and bias $-dR$; hence the second-layer weights have magnitude one and biases at most $2dR$. Output weights have magnitude at most two and denominator q. Since $q \leq d \leq D$, all parameters have polynomial magnitude and $O(\log D)$ bit length. Enumerating the units and their connections takes polynomial time in D. Finally, the global range of ψ ensures that the composite extension stays in $[0, 1]$ even off the grid. □