A deterministic FPTAS for weakly interacting continuous factor models

Abstract

We give a deterministic fully polynomial-time approximation scheme for partition functions of continuous factor models on the unit cube, with fixed bounds on factor arity, variable occurrence, and Lipschitz constants. For factors $1 \le \Psi_e \le e^q$, the algorithm applies when $3c\sqrt{r-1}\,q < 1$ with fixed slack, where $r \ge 2$ bounds the arity and c bounds occurrence. Under polynomial-time finite-precision value access, it has polynomial bit complexity in the input length and inverse relative accuracy. In particular, it gives polynomial time in the weak-interaction regime of the quasi-polynomial algorithm of Gamarnik and Smedira (Combin. Probab. Comput. 2026). The obstacle to a direct use of zero-free interpolation is numerical: logarithmic-order moments involve logarithmically many coordinates, on which a uniformly fine grid is still too large. We expand each factor in dyadic increments and truncate by total refinement level. Even for overlapping scopes, the common grid has size exponential only in that total level. Combined with Lipschitz decay, this evaluates a local moment of order s to absolute error e^{-as}, for any fixed $a > 0$, using $e^{O(s)}$ oracle calls, without mixed derivative assumptions. Stability estimates keep value precision and dyadic coordinate lengths to $O(\log(n/\varepsilon))$ bits, where n is the number of variables and ε the relative accuracy.

Note. This paper was generated entirely by AI, including MiMo, using an automated research pipeline developed by Chenghua Liu and Hanyu Li.

1 Introduction

Let V be a finite set of variables, and let E be a labelled collection of nonempty subsets of V. Different factor occurrences may have the same scope. For positive functions $\Psi_e : [0,1]^e \to \mathbb{R}_{>0}$, consider the partition function

$$Z = \int_{[0,1]^V} \prod_{e \in E} \Psi_e(x_e)\, dx = \mathbb{E} \prod_{e \in E} \Psi_e(X_e), \tag{1.1}$$

where the X_v are independent uniform random variables on $[0,1]$. We write $n = |V| \ge 1$ and $m = |E|$. The empty instance has $Z = 1$. The arity of a factor is its scope size; a variable's occurrence count is the number of factor occurrences containing it.

When every factor depends on only a few coordinates and each variable occurs in only a few factors, the interaction structure is sparse even though the integral is high-dimensional. Gamarnik and Smedira exploit this structure by proving correlation decay for near-constant factors with bounded variable occurrence [7, 8]. Their algorithm gives a deterministic $1 \pm O(1/n)$ approximation in $n^{O(\log n)}$ time, and they ask whether the exponent can be made independent of n [7, p. 2 and Thm. 1]. We obtain an FPTAS for bounded-arity Lipschitz factors throughout that regime, with an explicit finite-precision model for evaluating the factors.

A natural route is to replace the correlation-decay recursion by zero-free interpolation. Barvinok's multispin theorem gives a disk on which the partition function does not vanish [2, Thm. 1.1], and a logarithmic-order Taylor expansion then suffices. Bounded dependence also permits the coefficients

to be assembled from only polynomially many connected subproblems. For a fixed finite spin space, this leads to polynomial-time coefficient algorithms [4, 11]. Continuous variables leave another cost inside each subproblem, however: the local coefficients themselves are integrals.

Uniform discretization does not remove that cost. At accuracy $\varepsilon = 1/n$, the expansion order is $s = \Theta(\log n)$, and a local term can involve $\Theta(\log n)$ coordinates. Even if a grid has only n^a intervals per coordinate, using it on $b \log n$ coordinates costs $(n^a)^{b \log n} = \exp(ab(\log n)^2)$ cells for fixed $a, b > 0$. This is the same loss seen when a polynomial-size alphabet is substituted into finite-spin coefficient routines, whose local cost is exponential in the product of coefficient order and log-alphabet size [11, Rem. 3.2 and proof of Thm. 6.1] [4, Thm. 1.6]. Restricting to connected supports controls how many integrals must be computed, but not the cost of computing one of them.

We compute each local integral by preserving its factorization through the discretization. Write each factor as a sum of dyadic step-function increments, with the increment at level ν_i decaying as $2^{-\nu_i}$. Terms are retained according to the sum of their factor levels. An overlapping coordinate need only be refined to the largest level among the factors containing it: if those scopes are S_i, then

$$\sum_v \max_{i : v \in S_i} \nu_i \le \sum_i |S_i| \nu_i \le r \sum_i \nu_i. \tag{1.2}$$

Thus a term of total level J is integrated on at most 2^{rJ} cells. The decay of the increments makes $J = O(s)$ sufficient for error $e^{-O(s)}$, and the number of retained level vectors is also $e^{O(s)}$. This replaces a grid that simultaneously resolves every coordinate at the finest scale by a sum of grids whose refinement is shared among the factors.

The construction uses the total-level principle of sparse grids [6], but assigns levels to factors rather than individual coordinates. Only the individual factors need be Lipschitz; no mixed derivatives of the product are assumed. For moments containing powers of $\log \Psi_e$, their Lipschitz constants must also be accounted for: the product of these constants is exponential in the *total moment order*, not in the number of factors times that order. We show further that sampled-value errors and the conversion from moments to logarithmic coefficients each incur only an exponential loss in s. Consequently, the local routine and the full algorithm use $O(\log(n/\varepsilon))$ precision bits, rather than requiring an exact real-arithmetic oracle.

This local-computation requirement also appears in other approximation frameworks. Intersection-probability algorithms use exponential-time routines for small intersections [1, Sec. 2.2] and [9, Thm. 2], and Barvinok's Gaussian and exponential integration framework requires low-dimensional moments [3, Sec. 1.1]. Bencs and Regts obtain an FPTAS for a truncated independence polytope using explicit local integrals; their discussion of general graphon integrals identifies efficient coefficient computation as the unresolved step [5, Sec. 5]. The quadrature below supplies such a step for the Lipschitz factor model on the cube.

1.1 Main theorem and access model

Assumption 1.1 (Effective value access)**.** For a factor occurrence e, a rational point $x_e \in [0,1]^e$, and an integer $b \ge 1$, a deterministic oracle returns a binary rational y with

$$|y - \Psi_e(x_e)| \le 2^{-b}.$$

The oracle's running time is bounded uniformly by a polynomial in b, the query encoding length, and the factor-label encoding length. The input includes the incidence list, the factor labels or oracle handles, and a binary rational accuracy $0 < \varepsilon \le 1/2$. Their total encoding length, including that of ε, is denoted by L.

The structural and analytic parameters in each theorem are fixed, effectively computable constants. Certified rational bounds for the constants used by the algorithm are available. These parameters are not part of the varying input. No derivative values are supplied or used.

Theorem 1.2 (Fixed arity and bounded occurrence)**.** *Fix integers $r \geq 2$, $c \geq 1$, and effectively computable constants $K, q > 0$, $\sigma \in (0,1)$, such that*

$$3c\sqrt{r-1}\, q \leq 1 - \sigma. \tag{1.3}$$

Suppose every factor has arity at most r, every variable occurs in at most c factors, and

$$1 \leq \Psi_e \leq e^q, \qquad \mathrm{Lip}_2(\Psi_e) \leq K. \tag{1.4}$$

Under Assumption 1.1, there is a deterministic algorithm returning a positive rational $\widehat{Z}$ with

$$e^{-\varepsilon} \leq \frac{\widehat{Z}}{Z} \leq e^{\varepsilon}.$$

For a constant $C = C(r, c, K, q, \sigma)$, it uses at most $(n/\varepsilon)^C$ oracle calls, each at $O(\log(n/\varepsilon))$ bits of absolute precision and at dyadic points with $O(\log(n/\varepsilon))$ fractional bits. Its non-oracle bit complexity is $\mathrm{poly}(L)(n/\varepsilon)^C$. Including the oracle running time, the algorithm is polynomial in L, n, and $1/\varepsilon$.

Here Lip_2 denotes the Lipschitz constant for Euclidean distance. For a function continuous on the cube and differentiable on its relative interior, the promise $\|\nabla \Psi_e\|_2 \leq K$ implies $\mathrm{Lip}_2(\Psi_e) \leq K$: apply the mean-value theorem to interior line segments and then take limits at the boundary. Thus the theorem also applies under the differentiability assumptions of the source problem. Running it with logarithmic accuracy $\varepsilon/2$ gives the conventional $1 - \varepsilon \leq \widehat{Z}/Z \leq 1 + \varepsilon$ guarantee.

To state the original regime, the *primal graph* has vertex set V, with two distinct variables adjacent when they lie in a common scope. If its maximum degree is at most η, every scope has size at most $\eta + 1$.

Corollary 1.3 (The Gamarnik–Smedira regime)**.** *Fix positive integers Δ, η, constants $k > 0$, $\delta \in (0,1)$, and put*

$$\rho = (1-\delta)^{1/(2\Delta)} \left(1 + \frac{\log(1 + 1/\eta)}{2\Delta}\right). \tag{1.5}$$

Assume $\rho > 1$, every variable occurs in at most Δ factors, and the primal graph has maximum degree at most η. For factors satisfying

$$1 \leq \Psi_e \leq \rho, \qquad \mathrm{Lip}_2(\Psi_e) \leq k,$$

the conclusions of Theorem 1.2 hold with constants depending on Δ, η, k, δ. In particular, for $n \geq 2$, taking $\varepsilon = 1/n$ gives a deterministic $e^{\pm 1/n}$ approximation in polynomial time in the stated oracle model.

For $\varepsilon = 1/n$, Corollary 1.3 replaces the quasi-polynomial bound by polynomial time under the stated access model. The proof in Appendix A checks the range conversion and treats primal degree at most one directly. The parameters r, c, K, q and the slack σ are fixed in the FPTAS; the polynomial exponent may depend on them.

2 From zero-freeness to local moments

Taking the logarithm serves two purposes. A zero-free disk gives a geometrically convergent Taylor expansion, while independence makes its coefficients vanish on disconnected supports. Both facts are needed: the first bounds the order of the moments, and the second avoids the m^s enumeration of arbitrary factor tuples. We derive bounds for every subcollection of factors because coefficient extraction will use inclusion–exclusion over those subcollections.

Throughout Sections 2–4, assume the hypotheses of Theorem 1.2. Put

$$\ell_e = \log \Psi_e, \qquad 0 \le \ell_e \le q, \qquad \mathrm{Lip}_2(\ell_e) \le K.$$

For $B \subseteq E$, write $V(B) = \bigcup_{e \in B} e$ and define

$$Q_B(z) = \mathbb{E} \exp\left(z \sum_{e \in B} \ell_e(X_e) \right), \qquad Q(z) = Q_E(z). \tag{2.1}$$

These functions are entire: the exponential series converges uniformly on every compact set of z, since $\sum_{e \in B} \ell_e \le |B| q$. Also $Q(1) = Z$ and $Q_B(0) = 1$.

Lemma 2.1 (A common zero-free disk)**.** *Let*

$$\beta = \frac{1}{3c\sqrt{r-1}\, q}.$$

For every $B \subseteq E$, $Q_B(z) \neq 0$ whenever $|z| \le \beta$. One can fix an effectively chosen rational R with $1 < R < \beta$.

Proof. Each ℓ_e/q takes values in $[0,1]$, and hence is 1-Lipschitz for Hamming distance. It depends on at most r coordinates, each of which appears in at most c functions. Apply Barvinok's theorem [2, Thm. 1.1] with $\phi_e = \ell_e/q$ and $\lambda = zq$, on the product space indexed by $V(B)$. This gives the asserted disk, including for every subcollection B. Finally,

$$\beta \ge \frac{1}{1-\sigma} > 1 + \frac{\sigma}{2} > 1,$$

so a rational $R \in (1, 1+\sigma/2)$ can be selected using certified bounds on the fixed parameters. □

Let $L_B(z)$ be the analytic logarithm of $Q_B(z)$ on a neighbourhood of $|z| \le R$, normalized by $L_B(0) = 0$. Write

$$L_B(z) = \sum_{p \ge 1} d_p(B) z^p, \qquad c_p = d_p(E).$$

Since $Q(t) > 0$ for real $t \in [0,1]$, $L_E(1) = \log Z$ with the ordinary real logarithm.

Lemma 2.2 (Coefficient and truncation bounds)**.** *For $h = |B|$ and $p \ge 1$,*

$$|d_p(B)| \le 2Rhq\, R^{-p}. \tag{2.2}$$

Consequently, for every integer $s \ge 0$,

$$\left| \log Z - \sum_{p=1}^{s} c_p \right| \le \frac{2Rmq}{R-1} R^{-s}. \tag{2.3}$$

Moreover, for $|z| \le 1/2$,

$$|L_B(z)| \le \kappa h, \qquad |Q_B(z)| \ge e^{-\kappa h}, \qquad \kappa = \frac{2Rq}{2R-1}. \tag{2.4}$$

Proof. For $|z| \le R$,

$$\operatorname{Re} L_B(z) = \log |Q_B(z)| \le \log \mathbb{E} \exp\left(|z| \sum_{e \in B} \ell_e \right) \le Rhq.$$

Carathéodory's coefficient estimate for an analytic function with constant term zero and real part at most Rhq gives (2.2). This estimate also follows by applying the Fourier coefficient formula on $|z| = u < R$ to the nonnegative function $Rhq - \operatorname{Re} L_B(ue^{it})$, whose mean is Rhq, and letting $u \uparrow R$. Summing (2.2) over $p > s$ proves (2.3). On $|z| \le 1/2$, summing over $p \ge 1$ instead gives

$$|L_B(z)| \le 2Rhq \sum_{p \ge 1} (2R)^{-p} = \frac{2Rhq}{2R-1}.$$

The lower bound for $|Q_B| = \exp(\operatorname{Re} L_B)$ follows. □

If $m = 0$, return 1. Henceforth $m \ge 1$. Choose a rational upper bound $q_+ \ge q$, and let $s \ge 1$ be the least integer such that

$$\frac{2Rmq_+}{R-1} R^{-s} \le \frac{\varepsilon}{3}. \tag{2.5}$$

This is a rational comparison. Since $R > 1$ and $q_+ > 0$ are fixed,

$$s = \Theta\left(\log \frac{m+1}{\varepsilon} \right) = O\left(\log \frac{n+1}{\varepsilon} \right), \tag{2.6}$$

where the last inequality uses $m \le cn$.

The Taylor order is now logarithmic, but expanding all s-tuples of factors would still take quasi-polynomial time. To retain only the terms that can contribute to the logarithm, let H be the intersection graph on factor occurrences: distinct $e, f \in E$ are adjacent exactly when $e \cap f \ne \varnothing$. Its maximum degree is at most

$$D = r(c-1). \tag{2.7}$$

A disconnected set of factors depends on disjoint groups of independent variables. Its partition function factors, and its logarithm is additive. Inclusion–exclusion isolates precisely the monomials whose support is one specified set, making this cancellation explicit.

Consider the formal series

$$\mathcal{G}(\boldsymbol{w}) = \log \mathbb{E} \exp\left(\sum_{e \in E} w_e \ell_e \right) \in \mathbb{R}[[\boldsymbol{w}]],$$

so that $L_B(z) = \mathcal{G}(z\mathbf{1}_B)$. For $A \subseteq E$, set

$$\Theta_A(z) = \sum_{B \subseteq A} (-1)^{|A|-|B|} L_B(z). \tag{2.8}$$

For a multi-index $\boldsymbol{\alpha} \in \mathbb{N}_0^E$, write $\operatorname{supp} \boldsymbol{\alpha} = \{e : \alpha_e > 0\}$ and $|\boldsymbol{\alpha}| = \sum_e \alpha_e$.

Lemma 2.3 (Support extraction)**.** *For every $A \subseteq E$,*

$$\Theta_A(z) = \sum_{\operatorname{supp} \boldsymbol{\alpha} = A} [\boldsymbol{w}^{\boldsymbol{\alpha}}] \mathcal{G}(\boldsymbol{w})\, z^{|\boldsymbol{\alpha}|}. \tag{2.9}$$

Thus $L_E = \sum_{A \subseteq E} \Theta_A$, and $[z^p]\Theta_A = 0$ if $|A| > p$ or if $H[A]$ is disconnected. Also $\Theta_\varnothing = 0$.

Proof. A monomial with support $S \subseteq A$ receives in (2.8) the multiplier

$$\sum_{B: S \subseteq B \subseteq A} (-1)^{|A|-|B|} = (1-1)^{|A \setminus S|}.$$

This is one exactly when $S = A$, and zero otherwise. It proves (2.9), the inversion identity, and the degree bound. The empty-support term is zero because $\mathcal{G}(\mathbf{0}) = 0$.

If $A_1, \dots, A_t$ are the components of $H[A]$, their coordinate sets $V(A_i)$ are pairwise disjoint. Independence gives

$$\mathcal{G}(\boldsymbol{w}_A, \mathbf{0}) = \sum_{i=1}^{t} \log \mathbb{E} \exp\left(\sum_{e \in A_i} w_e \ell_e \right).$$

When $t \geq 2$, no monomial on the right has support exactly A. Hence $\Theta_A = 0$. □

It follows that, for $1 \leq p \leq s$,

$$c_p = \sum_{\substack{\varnothing \neq A \subseteq E \\ |A| \leq p,\ H[A] \text{ connected}}} \sum_{B \subseteq A} (-1)^{|A|-|B|} d_p(B). \tag{2.10}$$

There are at most $m(4 \max\{1, D\})^{a-1}$ connected vertex sets of size a. To see this and obtain an enumeration, root a set at its least vertex and use its canonical depth-first spanning tree, with neighbours ordered by their labels. There are at most 4^{a-1} rooted plane-tree shapes and at most $\max\{1, D\}^{a-1}$ choices of graph neighbours along the tree edges. Reject encodings with repeated vertices or whose tree is not the canonical tree of the resulting induced subgraph. This produces each connected set once, with polynomial work in a per candidate, after constructing the adjacency lists.

Put

$$A_0 = 8 \max\{1, D\}. \tag{2.11}$$

The number of pairs (A, B) with A nonempty and connected, $|A| \leq s$, and $B \subseteq A$, is at most

$$m \sum_{a=1}^{s} (4 \max\{1, D\})^{a-1} 2^a \leq m A_0^s.$$

In particular, the number of local coefficients in $\sum_{p=1}^{s} c_p$, counted with their multiplicities in (2.10), is bounded by

$$N_s \leq s m A_0^s. \tag{2.12}$$

We have reduced the global sum to $e^{O(s)}$ local computations per factor occurrence. Each is obtained from ordinary moments and an exact formal recurrence. For $|B| \leq s$, write

$$Q_B(z) = \sum_{j \geq 0} M_j(B) z^j, \qquad M_0(B) = 1.$$

The multinomial expansion gives, for $j \geq 1$,

$$M_j(B) = \sum_{\substack{\boldsymbol{\alpha} \in \mathbb{N}_0^B \\ |\boldsymbol{\alpha}| = j}} \frac{I_{\boldsymbol{\alpha}}}{\prod_{e \in B} \alpha_e!}, \qquad I_{\boldsymbol{\alpha}} = \int_{[0,1]^{V(\operatorname{supp} \boldsymbol{\alpha})}} \prod_{\alpha_e > 0} \ell_e(x_e)^{\alpha_e} \, dx. \tag{2.13}$$

The number of exponent vectors across $0 \leq j \leq s$ is

$$\binom{s + |B|}{|B|} \leq \binom{2s}{s} \leq 4^s. \tag{2.14}$$

The formal identity $Q'_B = L'_B Q_B$ yields the recurrence

$$d_j(B) = M_j(B) - \frac{1}{j}\sum_{i=1}^{j-1} i\, d_i(B) M_{j-i}(B), \qquad 1 \le j \le s. \tag{2.15}$$

The factors in $I_{\boldsymbol{\alpha}}$ can share coordinates, so this integral need not separate. Nevertheless, each factor still has arity at most r, and the exponents sum to at most s. The next section uses these two facts to evaluate it to error $e^{-O(s)}$ in time $e^{O(s)}$. The coefficient stability analysis will then verify that this accuracy suffices in (2.10).

3 Factor-level dyadic quadrature

A tensor grid for $I_{\boldsymbol{\alpha}}$ would refine as many as rs coordinates to the same accuracy. Instead, we first expand each local function in a telescoping sequence of step functions. A product of increments has small norm when the sum of its refinement levels is large, regardless of how those levels are distributed. This permits a truncation by total level, whose retained terms have small common grids even when their scopes overlap.

3.1 The total-level rule

Lemma 3.1 (Total-level quadrature)**.** *Let $t \ge 1$, let $S_1, \dots, S_t$ be nonempty scopes of size at most r, and put $\mathcal{U} = \bigcup_i S_i$. Suppose $u_i : [0,1]^{S_i} \to \mathbb{R}$ satisfy*

$$\|u_i\|_\infty \le A_i, \qquad \mathrm{Lip}_2(u_i) \le K_i.$$

Define

$$I = \int_{[0,1]^{\mathcal{U}}} \prod_{i=1}^{t} u_i(x_{S_i})\, dx, \quad B_i = \max\{1, A_i, 3\sqrt{r}\, K_i\}, \quad \mathcal{B} = \prod_{i=1}^{t} B_i, \quad C_0 = (1 - 2^{-1/2})^{-1}.$$

For $0 < \tau < 1$, choose an integer $J \ge 0$ satisfying

$$2^{J/2} \ge \frac{2\mathcal{B} C_0^t}{\tau}. \tag{3.1}$$

There is a finite step-function quadrature I_J with $|I - I_J| \le \tau/2$, using at most

$$2t\binom{J+t}{t} 2^{rJ} \le 2t\, 2^{(r+1)J+t} \tag{3.2}$$

function-value queries and $O_r(t\binom{J+t}{t}2^{rJ})$ arithmetic operations. Each query coordinate is dyadic with at most J fractional bits.

Proof. Partition each coordinate into half-open dyadic intervals, closing the last interval at 1. Let $\Pi_\ell u_i$ be the step function equal to the lower-corner sample on every cube of side $2^{-\ell}$. Set

$$\Delta_{i,0} = \Pi_0 u_i, \qquad \Delta_{i,\ell} = \Pi_\ell u_i - \Pi_{\ell-1} u_i \quad (\ell \ge 1).$$

Since

$$\|\Pi_\ell u_i - u_i\|_\infty \le \sqrt{r}\, K_i 2^{-\ell},$$

we have uniform convergence and

$$u_i = \sum_{\ell \geq 0} \Delta_{i,\ell}, \qquad \|\Delta_{i,\ell}\|_\infty \leq B_i 2^{-\ell}. \tag{3.3}$$

The product series is absolutely convergent in the uniform norm, because the sum of the norms of all its terms is at most $2^t \mathcal{B}$. It can therefore be integrated termwise. Define

$$I_J = \sum_{\substack{\boldsymbol{\nu} \in \mathbb{N}_0^t \\ |\boldsymbol{\nu}| \leq J}} \int_{[0,1]^{\mathcal{U}}} \prod_{i=1}^{t} \Delta_{i,\nu_i}. \tag{3.4}$$

For $|\boldsymbol{\nu}| > J$, $2^{-|\boldsymbol{\nu}|} \leq 2^{-J/2} 2^{-|\boldsymbol{\nu}|/2}$. Hence

$$|I - I_J| \leq \mathcal{B} 2^{-J/2} \prod_{i=1}^{t} \sum_{\ell \geq 0} 2^{-\ell/2} = \mathcal{B} C_0^t 2^{-J/2} \leq \tau/2. \tag{3.5}$$

For a retained vector $\boldsymbol{\nu}$, define

$$m_v = \max_{i : v \in S_i} \nu_i \qquad (v \in \mathcal{U}).$$

The product in (3.4) is constant on the dyadic grid having coordinate levels m_v, and

$$\sum_{v \in \mathcal{U}} m_v \leq \sum_{i=1}^{t} |S_i| \nu_i \leq rJ. \tag{3.6}$$

Thus it is integrated exactly by summing over at most 2^{rJ} cells, each of volume $2^{-\sum_v m_v}$. Its value on a cell requires at most $2t$ samples. There are $\binom{J+t}{t} \leq 2^{J+t}$ retained level vectors. Computing their coordinate levels, increments, cell products, and weighted sums gives the stated bounds. All lower-corner coordinates have at most J fractional bits. □

3.2 Finite-precision samples and primitive moments

The rule (3.4) contains signed increments. Its exact-sample error bound therefore does not by itself justify replacing the samples by approximate oracle values: cancellation in the rule could amplify those errors. The next estimate sums perturbations at the level of increments. It introduces only a factor exponential in the number of factors, which is acceptable at logarithmic moment order.

Lemma 3.2 (Sample perturbations)**.** *In Lemma 3.1, replace each sampled value of u_i by a number at absolute distance at most ζ_i, using the same stored value whenever a sample is reused. Let $\widetilde{I}_J$ be the resulting formula. If*

$$2(J+1)\zeta_i \leq B_i \qquad (1 \leq i \leq t),$$

then

$$|\widetilde{I}_J - I_J| \leq 2(J+1)3^{t-1}\mathcal{B} \sum_{i=1}^{t} \frac{\zeta_i}{B_i} \leq 2t(J+1)3^{t-1}\mathcal{B} \max_i \zeta_i. \tag{3.7}$$

Proof. The stored samples define perturbed step functions and increments. Their errors satisfy

$$\|\widetilde{\Delta}_{i,0} - \Delta_{i,0}\|_\infty \leq \zeta_i, \qquad \|\widetilde{\Delta}_{i,\ell} - \Delta_{i,\ell}\|_\infty \leq 2\zeta_i \quad (\ell \geq 1).$$

Therefore

$$\sum_{\ell=0}^{J} \|\widetilde{\Delta}_{i,\ell} - \Delta_{i,\ell}\|_\infty \le 2(J+1)\zeta_i,$$
$$\sum_{\ell=0}^{J} \|\Delta_{i,\ell}\|_\infty \le 2B_i,$$
$$\sum_{\ell=0}^{J} \|\widetilde{\Delta}_{i,\ell}\|_\infty \le 3B_i.$$

Apply the identity

$$\prod_{i=1}^{t} a_i - \prod_{i=1}^{t} b_i = \sum_{i=1}^{t} (a_i - b_i) \prod_{h<i} a_h \prod_{h>i} b_h$$

to each integrand. Take absolute values and enlarge the sum over $|\boldsymbol{\nu}| \le J$ to the box $\{0,\dots,J\}^t$. The resulting sums factor by i, and the three preceding bounds give the first inequality in (3.7). The second uses $B_i \ge 1$. □

For the primitive integral $I_{\boldsymbol{\alpha}}$ in (2.13), take $u_e = \ell_e^{\alpha_e}$ for $\alpha_e \ge 1$. Then

$$\|u_e\|_\infty \le q^{\alpha_e}, \qquad \mathrm{Lip}_2(u_e) \le \alpha_e K q^{\alpha_e - 1}. \tag{3.8}$$

Fix an integer

$$C_* \ge 2\max\{1, q, 3\sqrt{r}\,K\}.$$

For $\alpha \ge 1$, $\alpha \le 2^\alpha$ implies

$$B_e = \max\{1, q^{\alpha_e}, 3\sqrt{r}\,\alpha_e K q^{\alpha_e - 1}\} \le C_*^{\alpha_e}, \qquad \prod_{\alpha_e > 0} B_e \le C_*^s \quad \text{if } |\boldsymbol{\alpha}| \le s. \tag{3.9}$$

The total-order bound in (3.9) is essential. Bounding every B_e separately by C_*^s and multiplying over s factors would give $C_*^{s^2}$, require a quadratic total refinement level, and restore the quasi-polynomial cost. Charging the power α_e only to its own factor keeps the combined budget exponential in $\sum_e \alpha_e$.

Lemma 3.3 (Effective primitive moments)**.** *Let $1 \le |\boldsymbol{\alpha}| \le s$ and let $0 < \tau < 1$ be rational. Under Assumption 1.1, a deterministic procedure returns a rational $\widetilde{I}_{\boldsymbol{\alpha}}$ with*

$$|\widetilde{I}_{\boldsymbol{\alpha}} - I_{\boldsymbol{\alpha}}| \le \tau.$$

It uses parameters

$$J = O(s + \log(1/\tau)), \qquad P, b = O(s + \log(J+1) + \log(1/\tau)), \tag{3.10}$$

where query coordinates have at most J fractional bits, the oracle precision is b, and each local-function sample is rounded to P fractional bits. It makes $2^{O_r(J+s)}$ oracle calls. Its non-oracle bit complexity is $2^{O_r(J+s)}\,\mathrm{poly}(L + J + P + b + s)$.

Proof. Put $t = |\operatorname{supp}\boldsymbol{\alpha}| \le s$. Since $C_0 < 4$, a common rational choice of total level for all such exponent vectors is

$$J = \min\left\{ j \in \mathbb{N}_0 : 2^j \ge \left(\frac{2(4C_*)^s}{\tau}\right)^2 \right\}. \tag{3.11}$$

By (3.9), this satisfies (3.1), giving exact-sample error at most $\tau/2$. It also has the order asserted in (3.10).

The hypothesis (1.3) gives $q < 1/3$, so $e^q < 2$. Project each oracle answer for Ψ_e onto $[1, 2]$, which cannot increase its error. On this interval, log is 1-Lipschitz and lies in $[0, 1]$. Approximate the logarithm within 2^{-P}, project the approximation onto $[0, 1]$, raise it to the power α_e, and round once more to the nearest multiple of 2^{-P}. The resulting sample error is at most

$$\zeta_e \le \alpha_e(2^{-b} + 2^{-P}) + 2^{-P} \le s2^{-b} + (s + 1)2^{-P}. \tag{3.12}$$

All samples are stored for reuse.

Choose $b = P \ge 1$ so that the last expression is at most

$$\zeta = \min\left\{\frac{1}{2(J + 1)}, \frac{\tau}{4s(J + 1)3^{s-1}C_*^s}\right\}. \tag{3.13}$$

This is an effective rational comparison and gives (3.10). The first term verifies the hypothesis of Lemma 3.2; the second, together with (3.9), bounds the sample perturbation by $\tau/2$. Adding the two errors proves the approximation guarantee.

For clarity, the logarithms require no real-arithmetic oracle. For a rational $y \in [1, 2]$, put $v = (y - 1)/(y + 1) \in [0, 1/3]$. Then

$$\log y = 2\sum_{j\ge0} \frac{v^{2j+1}}{2j + 1}, \qquad 0 \le \log y - 2\sum_{j=0}^{N-1} \frac{v^{2j+1}}{2j + 1} \le \frac{2v^{2N+1}}{(2N + 1)(1 - v^2)}.$$

Taking $N = O(P + 1)$ makes this tail at most 2^{-P-1}; rounding the rational partial sum to the nearest P-bit dyadic gives the stated logarithm error. Exact summation, integer powering, and rounding take polynomial time in the precision, s, and the length of y. The oracle model bounds that last length by a polynomial in $L + b + J$.

Each sampled u_e requires one oracle call. The number of samples and cells is bounded by (3.2). Products of at most s P-bit dyadics and dyadic cell volumes have a common denominator 2^{Ps+rJ}. Thus accumulating the $2^{O_r(J+s)}$ signed terms uses numbers of bit length polynomial in $P + J + s$. This proves the bit-complexity bound. □

4 From local moments to the FPTAS

The primitive routine gives exponentially accurate moments at exponential cost in their order. To obtain an FPTAS, we must ensure that forming the logarithm and summing over connected supports do not demand substantially finer moments. Directly propagating errors through (2.15) is unnecessarily costly. Instead, perturb Q_B analytically on a fixed disk about zero. The lower bound on $|Q_B|$ from Lemma 2.2 controls the change in its logarithm, and Cauchy's estimate converts this to coefficient errors. We then allocate $\varepsilon/3$ each to Taylor truncation, coefficient computation, and final exponentiation.

Lemma 4.1 (Stability of local logarithmic coefficients)**.** *Fix an integer $F \ge 4e^\kappa$, with κ as in* (2.4)*. Let $|B| \le s$, and suppose*

$$\max_{1\le j\le s} |\widetilde{M}_j(B) - M_j(B)| \le \xi \le F^{-s}, \qquad \widetilde{M}_0(B) = 1.$$

Let $\widetilde{d}_j(B)$ be computed from these moments by the exact recurrence (2.15)*. Then*

$$|\widetilde{d}_j(B) - d_j(B)| \le F^s\xi \qquad (1 \le j \le s). \tag{4.1}$$

Proof. Define

$$E_B(z) = \sum_{j=1}^{s} (\widetilde{M}_j(B) - M_j(B)) z^j, \qquad \widehat{Q}_B(z) = Q_B(z) + E_B(z).$$

On $|z| \le 1/2$, $|E_B(z)| \le \xi$, while (2.4) gives $|Q_B(z)| \ge e^{-\kappa s}$. Therefore

$$w(z) = \frac{E_B(z)}{Q_B(z)}, \qquad |w(z)| \le e^{\kappa s}\xi \le 4^{-s} \le \frac{1}{4}.$$

In particular, $\widehat{Q}_B = Q_B(1+w)$ is zero-free on this disk, $w(0) = 0$, and the analytic logarithm $\log(1+w)$ has constant term zero. The convergent logarithm series gives

$$|\log(1+w)| \le \frac{|w|}{1-|w|} \le 2e^{\kappa s}\xi.$$

Cauchy's coefficient estimate on the circle of radius $1/2$ yields

$$|[z^j]\log \widehat{Q}_B - [z^j]\log Q_B| \le 2^{j+1} e^{\kappa s}\xi \le (4e^{\kappa})^s \xi \le F^s \xi.$$

The first s coefficients of $\widehat{Q}_B$ are precisely $1, \widetilde{M}_1(B), \dots, \widetilde{M}_s(B)$. Its logarithmic coefficients through degree s are therefore exactly the numbers returned by (2.15). This proves the claim. □

The factor F^s from this lemma and the count smA_0^s in (2.12) determine the required moment accuracy. All constants A_0, C_*, F are fixed integers selected from certified parameter bounds. With s chosen by (2.5), set

$$\tau_0 = \frac{\varepsilon}{3sm(4A_0F)^s}. \tag{4.2}$$

Evaluate every required primitive integral to absolute error at most τ_0. The coefficients $(\prod_e \alpha_e!)^{-1}$ in (2.13) have absolute value at most one; hence (2.14) gives moment error

$$\xi \le 4^s \tau_0 = \frac{\varepsilon}{3sm(A_0F)^s} \le F^{-s}.$$

Combining Lemma 4.1 with (2.12),

$$\sum_{p=1}^{s} |\widetilde{c}_p - c_p| \le smA_0^s F^s 4^s \tau_0 = \frac{\varepsilon}{3}. \tag{4.3}$$

Since $\log(1/\tau_0) = O(s)$, the choices in (3.11)–(3.13) give common global parameters

$$J, P, b = O(s) = O(\log((n+1)/\varepsilon)). \tag{4.4}$$

Every tolerance and every growing precision parameter is determined by rational arithmetic. The possibly longer original encoding of ε is charged to L.

It remains to bound the cost of manipulating the approximations. Controlling the number of arithmetic operations is insufficient unless the numerators and denominators also remain short. After a local-function value is rounded to a P-bit dyadic, perform all quadrature, multinomial, recurrence, and support-extraction arithmetic exactly over $\mathbb{Q}$, reducing fractions after each operation. Each primitive quadrature denominator divides 2^{Ps+rJ}. Since $\prod_e \alpha_e!$ divides $j!$, all computed moments of degree at most s have denominators dividing

$$D_0 = 2^{Ps+rJ} s!.$$

Induction in (2.15) shows that the denominator of $\widetilde{d}_j(B)$ divides $D_0^j j!$: in its i-th summand the additional factors are one moment denominator and a divisor of j, which are absorbed by $D_0^j j!$. All these denominators divide $D_0^s s!$, also across different sets B.

For numerator bounds, each sampled local function lies in $[0,1]$. A retained primitive quadrature contains $2^{O(s)}$ signed cell terms, each of absolute value at most one, so $|\widetilde{M}_j(B)| \le e^{O(s)}$. The recurrence then gives the coarse bound $|\widetilde{d}_j(B)| \le e^{O(s^2)}$, by induction on $j \le s$. The sums in (2.10) contain at most $me^{O(s)}$ terms. Since $\log D_0 = O(s^2)$, all numerators and denominators in the post-sampling computation have

$$O(s^3 + \log m) \tag{4.5}$$

bits. Exact rational arithmetic thus introduces only a polynomial factor. In particular, the constant moment remains exactly one.

The algorithm can now be specified entirely in rational arithmetic. The first three steps compute an approximation to $\log Z$; the last produces the rational output required by Theorem 1.2.

1. If $m = 0$, return 1. Otherwise construct H, choose s by (2.5), and choose τ_0, J, P, b as above.
2. Enumerate the nonempty connected sets $A \subseteq E$ of size at most s. For each $B \subseteq A$, compute the moments $\widetilde{M}_j(B)$, $1 \le j \le s$, from (2.13) using Lemma 3.3. Set $\widetilde{M}_0(B) = 1$; for $B = \varnothing$, all positive-degree moments are zero.
3. Use (2.15) and (2.10) to compute $\widetilde{c}_1, \dots, \widetilde{c}_s$ exactly, and put $\widetilde{L} = \sum_{p=1}^{s} \widetilde{c}_p$.
4. Return a positive rational $\widehat{Z}$ satisfying $|\log \widehat{Z} - \widetilde{L}| \le \varepsilon/3$, using the rational construction below.

For correctness, (2.3), (2.5), and (4.3) already give

$$|\widetilde{L} - \log Z| \le 2\varepsilon/3, \qquad |\widetilde{L}| \le mq + 1 = O(n).$$

Although Z can be exponential in n, its binary length is only polynomial. We implement the last step by first reducing the argument of the exponential to a fixed interval. Choose the least $h \ge 0$ with $|\widetilde{L}|/2^h \le 1/2$, and put $t = \widetilde{L}/2^h$. Then $2^h = O(n+1)$. For

$$P_N(t) = \sum_{j=0}^{N} \frac{t^j}{j!},$$

Taylor's theorem gives

$$|P_N(t) - e^t| \le \frac{2^{-N}}{(N+1)!}.$$

Choose N so that this rational bound is at most $a = \varepsilon/(12 \cdot 2^h)$. It suffices to take $N = O(h + \log(1/\varepsilon) + 1)$. Since $e^t > 1/2$ and $a \le 1/24$, $P_N(t) > 1/4$. The mean-value theorem for log on the interval between these two positive numbers gives

$$|\log P_N(t) - t| \le 4a.$$

Compute the positive rational $\widehat{Z} = P_N(t)^{2^h}$ by repeated exact squaring. It satisfies

$$|\log \widehat{Z} - \widetilde{L}| \le 2^h 4a = \varepsilon/3.$$

The rational partial sum has polynomial bit length in the length of t and N; its exact 2^h-th power increases this length by a factor of at most $2^h = O(n+1)$. Exponentiation therefore has polynomial bit complexity as well.

There are at most mA_0^s support–subset pairs and at most 4^s primitive exponent vectors per pair. Each primitive uses $2^{O_r(J+s)}$ samples and a polynomial factor of additional bit work. Together with (4.4) and (4.5), this gives

$$mA_0^s 4^s 2^{O_r(J+s)} \operatorname{poly}(L+s) = \operatorname{poly}(L)(n/\varepsilon)^{O_{r,c,K,q,\sigma}(1)}. \tag{4.6}$$

The same count without the bit-arithmetic factor bounds the oracle calls, and every call has precision and coordinate length as in (4.4). Finally,

$$|\log \widehat{Z} - \log Z| \le |\log \widehat{Z} - \widetilde{L}| + \left|\widetilde{L} - \sum_{p=1}^{s} c_p\right| + \left|\sum_{p=1}^{s} c_p - \log Z\right| \le \varepsilon.$$

This proves Theorem 1.2.

A The original weak-interaction regime

We first record the direct treatment needed when the primal graph has very small degree. No zero-free condition is required in this case.

Lemma A.1 (Primal degree at most one)**.** *Fix an integer $\Delta \ge 1$ and constants $k > 0$, $U > 1$. Suppose every variable occurs in at most Δ factors, the primal graph has maximum degree at most one, and $1 \le \Psi_e \le U$, $\mathrm{Lip}_2(\Psi_e) \le k$. Under Assumption 1.1, the partition function has a deterministic $e^{\pm\varepsilon}$ approximation with polynomial oracle-call and bit complexity. Precision and dyadic coordinate lengths are $O(\log(n/\varepsilon))$.*

Proof. Every primal component C has at most two vertices and contains at most $t_C \le 2\Delta$ factor occurrences. Put

$$F_C(x_C) = \prod_{e \subseteq C} \Psi_e(x_e), \qquad I_C = \int_{[0,1]^C} F_C(x_C)\, dx_C.$$

Then

$$Z = \prod_C I_C, \qquad I_C \ge 1, \qquad \mathrm{Lip}_2(F_C) \le 2\Delta k U^{2\Delta-1}.$$

The Lipschitz bound follows by telescoping the product and using the bound U on all remaining factors. Fix a positive rational upper bound $T \ge 2\Delta k U^{2\Delta-1}$. Apply the midpoint rule with uniform dyadic mesh

$$a \le \frac{\varepsilon}{4Tn}.$$

Pairing a point with the midpoint of its cell bounds the exact-sample error by $T\sqrt{2}\, a/2 \le \varepsilon/(4n)$.

Fix a rational $U_0 > U$, and project every oracle answer onto $[1, U_0]$. At any grid point the product error is at most

$$t_C U_0^{t_C - 1} 2^{-b} \le 2\Delta U_0^{2\Delta-1} 2^{-b}.$$

Choose $b = O(\log(n/\varepsilon))$ to make this at most $\varepsilon/(2n)$. The resulting rational midpoint estimate satisfies

$$\widehat{I}_C \ge 1, \qquad |\widehat{I}_C - I_C| \le \varepsilon/n.$$

Since log is 1-Lipschitz on $[1, \infty)$,

$$\left|\log \prod_C \widehat{I}_C - \log Z\right| \le \sum_C |\widehat{I}_C - I_C| \le \varepsilon.$$

Each grid has polynomially many points in dimension at most two. There are at most n components, and midpoint coordinates have $O(\log(n/\varepsilon))$ bits. Each product contains at most 2Δ oracle answers. Exact rational summation and the final product have polynomial bit length, because the sum of the bit lengths of all sampled rational numbers is polynomial under Assumption 1.1. Components without factors contribute exactly one. □

Proof of Corollary 1.3. For $\eta = 1$, apply Lemma A.1 with $U = \rho$. Suppose $\eta \geq 2$, and set $q = \log \rho > 0$. Equation (1.5) gives

$$q < \frac{\log(1 + 1/\eta)}{2\Delta}. \tag{A.1}$$

For $\eta = 2$, $\log(3/2) < 5/12$, so

$$\frac{11}{10}\, 3\Delta\sqrt{\eta}\, q < \frac{11}{10}\frac{3\sqrt{2}}{2} \log\frac{3}{2} < \frac{11\sqrt{2}}{16} < 1.$$

For $\eta \geq 3$, $\log(1 + 1/\eta) < 1/\eta$, and hence

$$\frac{11}{10}\, 3\Delta\sqrt{\eta}\, q < \frac{33}{20\sqrt{\eta}} \leq \frac{11\sqrt{3}}{20} < 1.$$

Thus $3\Delta\sqrt{\eta}\, q < 10/11$. Apply Theorem 1.2 with $r = \eta + 1$, $c = \Delta$, $K = k$, and $\sigma = 1/11$. □

Several conventions complete the comparison with the source formulation. If positive normalization constants are given as binary rationals, $\Phi_e = c_e \Psi_e$, then the original integral is $(\prod_e c_e)Z$; the product is restored exactly, with its input length included in L. Rational empty-scope factors are handled the same way before applying the theorem. If a source formulation gives a real neighbour bound $\eta_{\mathrm{src}} > 0$, put $d = \lfloor \eta_{\mathrm{src}} \rfloor$. For $d \geq 1$,

$$\log(1 + 1/\eta_{\mathrm{src}}) \leq \log(1 + 1/d),$$

so the corresponding range envelope is no larger than the one with integer parameter d. For $d \leq 1$, Lemma A.1 applies directly. If the neighbour parameter is literally zero, do not evaluate (1.5); the unary instance is instead handled with any separately supplied fixed range bound. Finally, $\rho = 1$ forces all normalized factors to be one, while $\rho < 1$ is inconsistent with their range promises. For $n = 1$, accuracy $\varepsilon = 1/2$ already gives an $e^{\pm 1/n}$ approximation.

The general theorem also separates the analytic range from the original envelope. For pairwise factors it allows

$$q \leq \frac{1 - \sigma}{3\Delta},$$

whereas (A.1) gives $q < 1/(2\Delta^2)$ when $\eta = \Delta$. This wider range comes from Barvinok's zero-free theorem, not a new zero-free estimate. The algorithm does not establish the full Dobrushin regime, nor does it give a sampling or mixing-time result.

B A deterministic oracle-information benchmark

Even on disjoint scopes, deterministic value access has a polynomial cost that depends on the factor arity. A finite set of queries can miss smooth bumps inside each factor's domain; their integrals accumulate across the factors in $\log Z$. The following form of the classical fooling-function argument [10] makes this dependence explicit. It concerns black-box values, so additional readable symbolic descriptions are outside its lower-bound model. A matching query upper bound comes from uniformly accurate step-function surrogates; computing their partition function is immediate on disjoint scopes but can remain costly when scopes overlap.

Proposition B.1 (Deterministic value-query scale)**.** *Fix an integer $r \geq 1$ and constants $K > 0$, $U > 1$. For $m \geq 1$ pairwise disjoint r-ary scopes with one factor on each scope, assume*

$$1 \leq \Psi_e \leq U, \qquad \mathrm{Lip}_2(\Psi_e) \leq K.$$

Every deterministic adaptive black-box algorithm guaranteeing $|\log(\widehat{Z}/Z)| \leq \varepsilon$ on every such instance, where $0 < \varepsilon \leq 1$, has worst-case query count

$$N_{\mathrm{qry}} \geq \left(\frac{c_0}{2U\varepsilon}\right)^r m^{r+1} - m \tag{B.1}$$

for a constant $c_0 = c_0(r, K, U) > 0$. This holds even with exact value queries and even for continuously differentiable factors.

Conversely, $M \geq 1$ arbitrary factors of arity at most r with the same range and Lipschitz bounds admit rational step-function surrogates whose partition function has logarithmic error at most ε, after $O_{r,K}(M^{r+1}\varepsilon^{-r})$ finite-precision value queries. For disjoint scopes this also gives a partition-function algorithm with the same query bound.

Proof. Define

$$p(t) = 16t^2(1-t)^2 \quad (0 \leq t \leq 1), \qquad \varphi(u) = \prod_{j=1}^{r} p(u_j),$$

and extend p by zero outside $[0, 1]$. Then

$$0 \leq p \leq 1, \quad p(0) = p(1) = p'(0) = p'(1) = 0, \quad |p'| \leq 8, \quad \int_0^1 p(t)\,dt = \frac{8}{15}.$$

It follows that

$$\|\nabla\varphi\|_2 \leq 8\sqrt{r}, \qquad J_r = \int_{[0,1]^r} \varphi(u)\,du = (8/15)^r.$$

Choose a positive rational

$$a_0 \leq \min\{U - 1, K/(8\sqrt{r}), 1\}, \qquad c_0 = a_0 J_r/6.$$

Run the algorithm on the all-one instance, answering every query by exactly one. If it does not terminate, the query lower bound is immediate. Otherwise let N_e be the number of queries to factor e. If $N_e \geq 1$, put

$$L_e = \left\lceil (2N_e)^{1/r} \right\rceil$$

and divide the cube into L_e^r congruent cells. Mark a cell when its interior contains a queried point. At most N_e cells are marked; since $L_e^r \geq 2N_e$, at least half remain unmarked. In each unmarked cell $C = z/L_e + [0, 1/L_e]^r$, place

$$g_{e,C}(x) = \frac{a_0}{L_e}\varphi(L_e x - z),$$

and set $g_e = 0$ on marked cells. These pieces define a C^1 function: each bump and its first derivatives vanish on every cell face. All queried points see value zero, including points on faces, and

$$0 \leq g_e \leq a_0 \leq U - 1, \qquad \|\nabla g_e\|_2 \leq 8a_0\sqrt{r} \leq K.$$

If $N_e = 0$, use $g_e = a_0\varphi$.

For $N_e \geq 1$, each active cell contributes $a_0 J_r L_e^{-(r+1)}$ to the integral. Since $L_e \leq 3N_e^{1/r}$,

$$\mu_e := \int g_e \geq \frac{a_0 J_r}{2L_e} \geq c_0 N_e^{-1/r}.$$

Together with the case $N_e = 0$, this gives

$$\mu_e \geq c_0 (N_e + 1)^{-1/r}. \tag{B.2}$$

The factors $\Psi_e = 1 + g_e$ satisfy all the promises and produce the same transcript as the all-one factors. Indeed, identical earlier answers force the same next query, and its answer is again one. The scopes are disjoint, so their alternative partition function is $Z_1 = \prod_e (1 + \mu_e)$, whereas $Z_0 = 1$ on the all-one instance. Since $\mu_e \leq U - 1$,

$$\log Z_1 \geq \frac{1}{U} \sum_e \mu_e \geq \frac{c_0}{U} \sum_e (N_e + 1)^{-1/r}.$$

Here we used $\log(1 + x) \geq x/(1 + x) \geq x/U$. A common output can be ε-accurate in logarithm for both instances only if $\log Z_1 \leq 2\varepsilon$. Jensen's inequality for $x \mapsto x^{-1/r}$, with $N_{\mathrm{qry}} = \sum_e N_e$, gives

$$\sum_e (N_e + 1)^{-1/r} \geq m^{1+1/r} (N_{\mathrm{qry}} + m)^{-1/r}.$$

Combining these inequalities proves (B.1).

The construction is compatible with finite-precision value access: all-one replies are also admissible finite-precision replies. The completed factors have finite rational piecewise-polynomial descriptions, using the grid sizes and marked-cell indices, so they can be evaluated at rational points effectively. These descriptions are hidden in this black-box argument. The proposition makes no lower-bound claim for an algorithm allowed to inspect such descriptions.

For the upper bound, use on each scope a dyadic grid of mesh

$$a \leq \frac{\varepsilon}{2K\sqrt{r}\, M}.$$

If the right-hand side exceeds one, take $a = 1$. Query the value at each lower corner to absolute error at most $\varepsilon/(2M)$, and replace the answer y by $\max\{1, y\}$. Define $\widetilde{\Psi}_e$ by these rational values on the corresponding cells. Projection onto $[1, \infty)$ does not increase the error from the true factor value. Thus

$$\widetilde{\Psi}_e \geq 1, \qquad \|\widetilde{\Psi}_e - \Psi_e\|_\infty \leq \varepsilon/M.$$

Since both functions are at least one,

$$|\log \widetilde{\Psi}_e(x_e) - \log \Psi_e(x_e)| \leq \varepsilon/M.$$

The two product integrands therefore differ by a factor in $[e^{-\varepsilon}, e^{\varepsilon}]$, as do their integrals. Each factor requires $O_{r,K}((M/\varepsilon)^r)$ grid values, for the stated total query count, and precision $O(\log(M/\varepsilon))$ suffices. For disjoint scopes, each surrogate integral is just its grid average; multiplying these averages computes the surrogate partition function without any further queries. □

In the setting of Corollary 1.3, take $r = \eta + 1$, $m = \lfloor n/r \rfloor$, $U = \rho > 1$, and $\varepsilon = 1/n$. Disjoint scopes satisfy every occurrence bound $\Delta \geq 1$, and (B.1) gives

$$N_{\mathrm{qry}} = \Omega(n^{2r+1}) = \Omega(n^{2\eta+3}).$$

Thus a parameter-dependent polynomial sampling cost is unavoidable in this black-box model. For overlapping scopes, however, evaluating the partition function of a uniformly accurate surrogate may still take exponential time. The query upper bound alone is not an FPTAS and does not show that the exponent in (4.6) is optimal. The algorithmic result depends on combining local integration with zero-freeness and connected-support extraction; it does not extend to arbitrary graphon representations without an effective regularity and access model.

References

[1] A. Barvinok. Computing the probability of intersection. *Combinatorics, Probability and Computing*, First View, 1–20, 2026. doi:10.1017/S0963548326100492.

[2] A. Barvinok. On the zeros of partition functions with multi-spin interactions. *Annales de l'Institut Henri Poincaré D*, 13(4):747–774, 2026. doi:10.4171/AIHPD/228.

[3] A. Barvinok. Computing Gaussian and exponential integrals in $\mathbb{R}^n$. arXiv:2606.23556v3, 2026. https://arxiv.org/abs/2606.23556v3.

[4] F. Bencs and G. Regts. Barvinok's interpolation method meets Weitz's correlation decay approach. arXiv:2507.03135v2, 2025. https://arxiv.org/abs/2507.03135v2.

[5] F. Bencs and G. Regts. Approximating the volume of a truncated relaxation of the independence polytope. *Discrete & Computational Geometry*, 76:508–525, 2026. doi:10.1007/s00454-026-00824-y.

[6] H.-J. Bungartz and M. Griebel. Sparse grids. *Acta Numerica*, 13:147–269, 2004. doi:10.1017/S0962492904000182.

[7] D. Gamarnik and D. Smedira. Integrating high-dimensional functions deterministically. arXiv:2402.08232v1, 2024. https://arxiv.org/abs/2402.08232v1.

[8] D. Gamarnik and D. Smedira. Computing volume of a polytope and integrals in high dimensions deterministically. *Combinatorics, Probability and Computing*, First View, 1–33, 2026. doi:10.1017/S0963548326100558.

[9] R. L. Mann and G. Waite. Approximate counting in local lemma regimes. arXiv:2512.10134v1, 2025. https://arxiv.org/abs/2512.10134v1.

[10] E. Novak. Some results on the complexity of numerical integration. In R. Cools and D. Nuyens, editors, *Monte Carlo and Quasi-Monte Carlo Methods*, Springer Proceedings in Mathematics & Statistics 163, pages 161–183. Springer, 2016. doi:10.1007/978-3-319-33507-0_6.

[11] V. Patel and G. Regts. Deterministic polynomial-time approximation algorithms for partition functions and graph polynomials. *SIAM Journal on Computing*, 46(6):1893–1919, 2017. doi:10.1137/16M1101003.