hexagon:2611.00982v1

The Mean-Field Limit of High-Dimensional Euclidean Matching

Matteo D'Achille[1] Michael Goldman[2] Francesco Grotto[3]
Francesco Mattesini[4] Dario Trevisan[3]

[1]Institut Élie Cartan de Lorraine, Université de Lorraine, France
[2]Univ Paris Est Créteil, CNRS, Univ Gustave Eiffel, LAMA UMR 8050, F-94010 Créteil, France
[3]Dipartimento di Matematica, Università di Pisa, Italy
[4]School of Computation, Information and Technology, Technische Universität München, and Munich Center for Machine Learning, Germany

Abstract

We study the minimal cost of a bipartite matching between two independent samples of n uniform points in dimension d, with cost given by the p-th power of the distance, in the limit where first $n \to \infty$ and then $d \to \infty$. For an arbitrary sequence of norms, on the unit cube or on the flat torus, and for $1 \le p = o(d)$, the lim inf and the lim sup in n of the volume-normalized cost both equal $1 + (1 - 2\gamma_{\mathrm{E}})p/d + O((p/d)^2)$, where γ_{E} is Euler's constant. On the Euclidean flat torus and for $p = \alpha d$ with $\alpha > 0$ fixed, the lim inf and the lim sup in n of the normalized cost converge, as $d \to \infty$, to the limiting cost $\beta_{\mathrm{bip}}(1/\alpha)$ of the random bipartite assignment problem in pseudo-dimension $1/\alpha$; for $\alpha = 1$ this value is $\pi^2/6$. For arbitrary norms on the cube we obtain the same identification for α below an explicit threshold. The coefficient $1 - 2\gamma_{\mathrm{E}}$ is the right derivative at 0 of $\alpha \mapsto \beta_{\mathrm{bip}}(1/\alpha)$. The lower bounds rest on an averaging identity of Xu, a Gamma–Poisson representation of the targets and cavity dual prices. The upper bounds require turning a cavity allocation into an exact transport plan. For $p = o(d)$ this is done with a multiplicatively weighted Poisson race and a norm-adapted pairing of the residual mass. For fixed α we use a multiscale rescue scheme controlled by high Poisson moments, a periodic L^p flux estimate whose leading constant grows only polynomially in the dimension, and a single lift of the corrected plan to the atomic targets.

2020 Mathematics Subject CLassification Primary 60D05, 49Q22; Secondary 60C05, 60G55, 90C27

Keywords Random Euclidean bipartite matching, high-dimensional probability, random assignment problem, cavity equation, Poisson point processes, optimal transport

Contents

1 Introduction **3**
1.1 The problem 3
1.2 Main results 4
1.3 Background 4
1.4 Ideas of the proofs 5
1.5 Organization 6

2 Setting and main results **6**
2.1 Matching costs 6
2.2 The mean-field constant and the results we use 6

2.3 Main results 7
2.4 Conventions 8

3 Exact Poissonization **9**
3.1 Xu's identity 9
3.2 The Poissonized cost 10
3.3 Monotonicity in the sample size and de-Poissonization 11

4 The mean-field variational problem **12**
4.1 The Poisson race 13
4.2 The capped cavity law 15
4.3 Saturation and the trial limit 17
4.4 The logistic case 17

5 The vanishing-ratio regime **18**

6 Ordinary cubes and arbitrary norms **19**
6.1 Setting 20
6.2 Volume-coordinate race and the lower bound 21
6.3 Concentration of cells 23
6.4 Restricted allocations on the cube 24
6.5 Norm-net completion 26

7 Upper bounds on the cube **27**
7.1 Local allocations 27
7.2 Capped prices 28
7.3 The vanishing ratio 29
7.4 Proof of Theorem 2.3 29
7.5 Almost-sure costs 30

8 Cavity cells, exact cores and initial concentration **30**
8.1 Cavity cells 31
8.2 Exact cores and the initial exceptional mass 32
8.3 Concentration of the initial fields 33

9 Rescue maps and finite propagation **35**
9.1 The target rescue 36
9.2 The source rescue 36
9.3 The stages and their add-one envelopes 38
9.4 Finite propagation 38

10 Stage induction and choice of parameters **40**
10.1 Moments of the stage loads 40
10.2 Overloads and the Palm estimate 41
10.3 The two-phase induction 41
10.4 Choice of parameters 43

11 Finite-volume completion **44**
11.1 Terminal plans on the good event 45
11.2 The failure event 45

12 Density correction, flux and the atomic lift **46**
12.1 The reference plan 46
12.2 Labelled increments 46
12.3 A periodic flux estimate 47
12.4 Summation of the local currents 47
12.5 The nonlocal remainder and the static coupling 48
12.6 The typed lift 49

13 Proof of Theorem 2.5 **50**

A Proof of Xu's identity **51**
A.1 Convexity of the semidiscrete value 51
A.2 The leave-one-source-out allocation 52
A.3 The combinatorial bijection 53
A.4 Mixed moments and the Dirichlet law 54
A.5 Removing the genericity assumption 55
A.6 Proof of Corollary 3.2 56

B Proof of the periodic flux theorem **56**
B.1 Independent blocks 57
B.2 Heat kernel bounds 57
B.3 Time integration 58

Declaration on the use of artificial intelligence **59**

1 Introduction

1.1 The problem

Let $X_1, \dots, X_n$ and $Y_1, \dots, Y_n$ be independent uniform points in a d-dimensional domain, equipped with a distance dist, and let $p \geq 1$. The bipartite matching problem asks for

$$M_{n,d,p} = \min_{\sigma \in S_n} \sum_{i=1}^{n} \mathrm{dist}(X_i, Y_{\sigma(i)})^p. \tag{1.1}$$

For fixed $d \geq 3$ typical matching distances are of order $n^{-1/d}$, so that $\mathbb{E}M_{n,d,p}$ is of order $n^{1-p/d}$. For a norm N_d on $\mathbb{R}^d$ whose unit ball has volume v_d, the natural normalization is

$$(nv_d)^{p/d}\, \mathbb{E}\frac{M_{n,d,p}}{n}.$$

A ball of N_d-radius $(nv_d)^{-1/d}$ contains on average one sample point. In this normalization the cost of matching a point to the nearest point of the other sample has expectation of order one, uniformly in d and in the norm. This paper determines the limit of the normalized cost when first $n \to \infty$ at fixed d and then $d \to \infty$, in two regimes: $p = o(d)$, and $p = \alpha d$ with $\alpha > 0$ fixed.

In both regimes the answer is given by a mean-field model. Write $u = v_d N_d(x - y)^d$ for the volume coordinate of a displacement. Then the cost is $\mathrm{dist}^p = v_d^{-p/d} u^{p/d}$ and, around a fixed point, the Lebesgue measure of $\{u \leq t\}$ equals t. For a Poisson process of targets, the costs seen from a source therefore form a Poisson process on $(0, \infty)$ with intensity

$$q\, t^{q-1}\, \mathrm{d}t, \qquad q = \frac{d}{p}, \tag{1.2}$$

whatever the norm. This is the incident-edge process of the random-link (mean-field) bipartite assignment problem in pseudo-dimension q. That model ignores the geometric correlations between edge costs, which are present in every dimension. It suggests that the limiting normalized cost should be $\beta_{\mathrm{bip}}(q)$, the limiting cost per matched pair in that model, with $q \to \infty$ when $p = o(d)$ and $q = 1/\alpha$ when $p = \alpha d$. Our results show that the correlations do not affect the limit, so that the prediction is correct at leading order. In the first regime we identify the first-order correction, which is a statement about the behaviour of $q \mapsto \beta_{\mathrm{bip}}(q)$ as $q \to \infty$.

1.2 Main results

Write $\mathcal{C}(\alpha) = \beta_{\mathrm{bip}}(1/\alpha)$, so that $\mathcal{C}(\alpha)$ is the limiting cost of the random bipartite assignment problem in pseudo-dimension $1/\alpha$ (see Definition 2.1).

The regime $p = o(d)$. For an arbitrary sequence of norms N_d, on the cube with the distance $N_d(x - y)$ or on the flat torus with the induced periodic distance, and for $1 \le p_d = o(d)$, both the lim inf and the lim sup as $n \to \infty$ of the normalized cost equal

$$1 + (1 - 2\gamma_{\mathrm{E}})\frac{p_d}{d} + O\Big(\frac{p_d^2}{d^2}\Big)$$

(Theorems 2.3 and 2.4). The implied constant is numerical and does not depend on the sequence of norms. The coefficient $1 - 2\gamma_{\mathrm{E}}$ equals $\mathcal{C}'(0+)$: the vanishing-ratio regime is the tangent at the origin of the curve that governs the fixed-ratio regime.

The regime $p = \alpha d$. On the Euclidean flat torus, for every fixed $\alpha > 0$,

$$\lim_{d\to\infty} \liminf_{n\to\infty} (n\omega_d)^{\alpha}\, \mathbb{E}W_p^p(\mu_n, \nu_n) = \lim_{d\to\infty} \limsup_{n\to\infty} (n\omega_d)^{\alpha}\, \mathbb{E}W_p^p(\mu_n, \nu_n) = \mathcal{C}(\alpha)$$

(Theorem 2.5). Here μ_n, ν_n are the two empirical measures and ω_d is the volume of the Euclidean unit ball. For $\alpha = 1$ the limit is $\zeta(2) = \pi^2/6$. For arbitrary norms on the cube, the same identification holds when $\alpha < \alpha_{\mathrm{net}} \approx 0.0237$ (Theorem 2.3(i)).

The proofs also give the inequalities

$$1 + (1 - 2\gamma_{\mathrm{E}})r \le \mathcal{C}(r) \le \Gamma(1 + r)\Gamma(2 + r), \qquad r > 0$$

(Theorem 4.1 and Remark 4.5). The lower bound $1 + (1 - 2\gamma_{\mathrm{E}})p/d$ holds for the lim inf in n of the normalized cost at every fixed d, every norm and every $p \ge 1$ (Proposition 5.1).

1.3 Background

For $p = 1$, an arbitrary norm on the unit cube and $d \ge 3$, Talagrand [29] proved that, almost surely, the normalized matching cost is asymptotically at most $1 + O(\log d/d)$, and asked about its sharp high-dimensional behaviour. Our first result identifies the first-order term in the relative error, for all $1 \le p = o(d)$ and uniformly over norms.

The random-link model was introduced in statistical physics by Mézard and Parisi [23, 24, 25]. In [26] they compared it with the Euclidean problem, computing the corrections due to geometric correlations perturbatively and arguing that they vanish as $d \to \infty$. Numerical comparisons by Boutet de Monvel, Houdayer and Martin [6, 16] found the mean-field prediction accurate already in moderate dimension. Corrections in $1/d$ were studied further by Lucibello, Parisi and Sicuro [21]. Caracciolo, Lucibello, Parisi and Sicuro [9] proposed a field-theoretic description of the Euclidean bipartite problem in terms of density fluctuations, which captures the opposite, low-dimensional mechanism; see also [8]. D'Achille's thesis [10] compares the

Euclidean bipartite problem with the random assignment problem numerically and with the normalization used here.

On the mean-field side, Aldous [1] proved that the random assignment problem with exponential costs has limiting cost $\zeta(2)$, which is the case $q = 1$. Wästlund [31] introduced a finite-penalty relaxation and an associated exploration game, which yields the limit for every pseudo-dimension $q \geq 1$. Salez [27] proved uniqueness for the Mézard–Parisi cavity equation when $q > 1$, and Larsson [18] treated $0 < q < 1$. These results identify $\beta_{\mathrm{bip}}(q)$ but say nothing about geometric models.

For fixed $d \geq 3$ and every $p \geq 1$, Goldman and Trevisan [14] proved that the normalized expected cost on the Euclidean cube converges as $n \to \infty$; see also [15]. The dependence of the limit on d was not known. Concentration and regularity results for the power-cost matching problem are obtained in [11, 13].

1.4 Ideas of the proofs

Lower bounds. An identity of Xu [33] expresses the expected matching cost as the expected semidiscrete transport cost from the uniform measure to the targets weighted by an independent Dirichlet vector. Writing the Dirichlet weights as normalized exponential variables and Poissonizing the number of points turns the targets into a Poisson process with independent exponential marks (Section 3). Any choice of a price $h(E_i)$ for each target, as a function of its mark E_i, then gives a feasible dual pair, and its expected value can be computed exactly. For the logarithmic price $h(e) = r \log e$ this yields the lower bound $1 + (1 - 2\gamma_{\mathrm{E}})r$ for every norm (Proposition 5.1). For the prices given by the capped cavity equation it yields $\mathcal{C}(\alpha)$ in the limit (Section 4.2). Section 4 shows that the dual problem is the variational problem of a marked Poisson race. Its value is $\mathcal{C}(r)$, and as $r \to 0$ it reduces to a race of the ratios U_i/E_i, which explains the coefficient $1 - 2\gamma_{\mathrm{E}}$.

Upper bounds. The dual prices define an allocation of the source to the targets: the cell of target i is the set where i minimizes the reduced cost. In the mean-field limit the cell of a target with mark e has expected volume e. In high dimension it also concentrates, because the sharp Young inequality makes the overlap of two balls exponentially small compared with their volume. The main difficulty is that the realized cell volumes do not equal the realized marks. An exact transport plan must correct the discrepancy without paying the factor 2^p that a triangle inequality would cost on all of the mass.

For $p = o(d)$ (Sections 5–7) we use the ratio race, truncated to a fixed volume $L = d^4$, on the cube itself. The restricted allocation is centred exactly, for every Borel set, because the mean displacement kernel of the cells is even. The exponentially small residual mass is then paired directly along a hierarchy of nets adapted to the norm. This costs at most C^p times a power of the residual mass, which is negligible as long as $p = o(d)$. The bounds that survive the limit $n \to \infty$ do not depend on the norm. The same construction with capped cavity prices gives Theorem 2.3(i).

For $p = \alpha d$ (Sections 8–13) the cost dist^p is so sensitive to displacements that a residual of exponentially small mass may still be costly if it moves too far. We use bounded cavity prices with a cutoff $\lambda = (A_\alpha d)^\alpha$, a splitting of each cell into a core of volume $\min\{E_i, |C_i|\}$ and an exceptional remainder, and two rescue maps, one serving exceptional targets and one serving exceptional sources. They are applied on balls of geometrically increasing radii. The residual masses are controlled by Poisson moments of a fixed order $2s > 2\alpha$. The order must exceed 2α because a residual moved across the whole torus costs a factor V^α. The rescued plan has the correct target marginal and a source density that differs from the uniform one by $O(e^{-bd})$. This density is corrected by a periodic flux whose L^p norm is bounded with a leading constant $Cp\sqrt{d}$ (Theorem 12.1). The corrected continuous plan is then lifted once to

the atomic targets, and the reduced costs are estimated edge by edge, which avoids a loss of order 2^p on the whole matching.

1.5 Organization

Section 2 contains the definitions, the mean-field results we use and the statements of the main theorems. Section 3 recalls Xu's identity and proves the Poisson representation and a de-Poissonization lemma. Section 4 studies the mean-field variational problem and the capped cavity law. Sections 5–7 treat arbitrary norms: the regime $p = o(d)$ on the torus and on the cube, and the fixed-ratio regime on the cube for $\alpha < \alpha_{\mathrm{net}}$. Sections 8–13 prove Theorem 2.5. Appendix A contains a self-contained proof of Xu's identity, and Appendix B the proof of the periodic flux estimate.

2 Setting and main results

We define the normalized costs, recall the mean-field constant and the results from the literature that we use, state the main theorems, and fix the conventions used in the proofs.

2.1 Matching costs

Let N_d be a norm on $\mathbb{R}^d$, $K_d = \{x : N_d(x) \le 1\}$ its unit ball and $v_d = |K_d|$. Write $\|x\|_d = v_d^{1/d} N_d(x)$ for the volume-normalized norm, whose unit ball has volume one. On the unit torus $\mathbb{T}^d = \mathbb{R}^d/\mathbb{Z}^d$ the periodic distance is $d_{N_d,\mathbb{T}^d}(x,y) = \min_{z\in\mathbb{Z}^d} N_d(x-y+z)$. For $p \ge 1$ and probability measures μ, ν we write $W_p^p(\mu,\nu)$ for the optimal transport cost with cost dist^p; this is the integrated cost, not its p-th root.

Let $X_1, \dots, X_n, Y_1, \dots, Y_n$ be independent and uniform on $[0,1]^d$, and let $\mu_n = n^{-1}\sum_i \delta_{X_i}$, $\nu_n = n^{-1}\sum_i \delta_{Y_i}$. We consider the normalized costs on the cube and on the torus,

$$\mathcal{A}_{n,d,p,N_d} = (nv_d)^{p/d}\,\mathbb{E}W^p_{p,N_d;[0,1]^d}(\mu_n,\nu_n), \qquad \mathcal{A}^{\mathrm{per}}_{n,d,p,N_d} = (nv_d)^{p/d}\,\mathbb{E}W^p_{p,N_d;\mathbb{T}^d}(\mu_n,\nu_n), \quad (2.1)$$

where the first cost uses the distance $N_d(x-y)$ and the second the periodic distance $d_{N_d,\mathbb{T}^d}$. If $M^{N_d}_{n,p}$ denotes the minimal value of $\sum_i N_d(X_i - Y_{\sigma(i)})^p$, then $\mathcal{A}_{n,d,p,N_d} = \mathbb{E}M^{N_d}_{n,p}/(n^{1-p/d}v_d^{-p/d})$. Equivalently, $\mathcal{A}_{n,d,p,N_d}$ is n^{-1} times the expected minimal sum of $\|x-y\|_d^p$ for n uniform pairs in the cube $[0, n^{1/d}]^d$ of volume n. We also write

$$B^-_{d,p,N_d} = \liminf_{n\to\infty} \mathcal{A}^{\mathrm{per}}_{n,d,p,N_d}, \qquad B^+_{d,p,N_d} = \limsup_{n\to\infty} \mathcal{A}^{\mathrm{per}}_{n,d,p,N_d}.$$

Since $d_{N_d,\mathbb{T}^d}(x,y) \le N_d(x-y)$ for $x, y \in [0,1]^d$, we have $\mathcal{A}^{\mathrm{per}}_{n,d,p,N_d} \le \mathcal{A}_{n,d,p,N_d}$. For the Euclidean norm $v_d = \omega_d$.

2.2 The mean-field constant and the results we use

Fix $q > 0$. In the random bipartite assignment problem in pseudo-dimension q, the edges of the complete bipartite graph $K_{n,n}$ carry independent costs $\ell_e = n^{1/q}\ell'_e$, where the ℓ'_e are independent copies of a nonnegative random variable with $\mathbb{P}(\ell' \le t)/t^q \to 1$ as $t \downarrow 0$. Thus $n\mathbb{P}(\ell_e \le t) \to t^q$, and the costs of the edges at a vertex converge to a Poisson process with intensity $qt^{q-1}\,\mathrm{d}t$.

Definition 2.1. Let M_n be the minimal cost of a perfect matching of $K_{n,n}$ in this model. We let $\beta_{\mathrm{bip}}(q) = \lim_{n\to\infty} n^{-1}\mathbb{E}M_n$, the limiting cost per matched pair, and set

$$\mathcal{C}(\alpha) = \beta_{\mathrm{bip}}(1/\alpha), \qquad \alpha > 0.$$

The existence of this limit, its independence of the law of ℓ', and its description through the cavity equation are due to Aldous for $q = 1$ [1], to Wästlund for $q \geq 1$ [31], together with Salez for uniqueness when $q > 1$ [27], and to Larsson for $0 < q < 1$ [18]. We collect the precise inputs; in particular, $n^{-1}\mathbb{E}M_n \to \beta_{\mathrm{bip}}(q)$ for every $q > 0$.

For $\lambda > 0$ put $a = \lambda/2$. A *capped valuation* with cutoff λ is a random variable X with values in $[-a, a]$ such that $F(t) = \mathbb{P}(X \geq t)$ and $G = -\log F$ satisfy

$$G(t) = \mathbb{E}(t + X)_+^q, \qquad -a \leq t \leq a. \tag{2.2}$$

In particular $\mathbb{P}(X = a) = e^{-G(a)}$. For independent copies X, X' of a capped valuation put $E_{q,\lambda} = \frac{q}{q+1}\mathbb{E}(X + X')_+^{q+1}$. In the random-link model this is the limiting cost per pair of the edges used by an optimal λ-relaxed matching, in which a vertex may stay unmatched at price $\lambda/2$ [31, 18].

Theorem 2.2. *Let $q > 0$.*

(a) *If $q > 1$, there is a unique law of a real random variable X_q such that $F_q(t) = \mathbb{P}(X_q \geq t)$ satisfies $-\log F_q(t) = \mathbb{E}(t + X_q)_+^q$ for all $t \in \mathbb{R}$. It has moments of all orders, and $\beta_{\mathrm{bip}}(q) = -q \int_{\mathbb{R}} F_q \log F_q \, \mathrm{d}t$.*

(b) *For every $\lambda > 0$ there is a unique capped valuation $X = X_{q,\lambda}$ in the sense of (2.2).*

(c) *The limit in Definition 2.1 exists, is finite, and*

$$\limsup_{\lambda \to \infty} E_{q,\lambda} = \beta_{\mathrm{bip}}(q),$$

where $E_{q,\lambda}$ is computed with $X_{q,\lambda}$.

For $q \geq 1$, items (b) and (c) come from Wästlund's analysis of the relaxed (diluted) matching problem and of the associated exploration game [31, Sections 2–3]. He shows that the edge cost of the relaxed problem converges to $E_{q,\lambda}$ and that its limit as $\lambda \to \infty$ is the perfect-matching constant. Wästlund carries out the proofs on the complete graph K_n and states that the results also apply to $K_{n,n}$. Larsson [18, Section 3.5] gives the bipartite normalization explicitly, through the identities (3.5)–(3.7) there. For $q > 1$, uniqueness for the truncated and untruncated cavity equations, the tail bounds and the convergence of the entropy integrals are [27, Theorem 1 and Propositions 1–2]. Since $E_{q,\lambda} = -q \int F \log F$ (see (4.9)), the convergence of the entropy integrals together with (a) also gives (c) for $q > 1$. For $q = 1$ everything is explicit; see Section 4.4. For $0 < q < 1$, (b) and (c) are due to Larsson [18, Theorem 1.1 and Section 3.5]. Item (a) is used only in Theorem 4.1. The proofs of the main theorems use (b) and (c), and (c) only through Proposition 4.4.

2.3 Main results

Put

$$\chi = \frac{4}{3\sqrt{3}}, \qquad c_\chi = \log \frac{1}{\chi} = \log \frac{3\sqrt{3}}{4} \approx 0.2616, \qquad \alpha_{\mathrm{net}} = \frac{c_\chi/2}{\log 192 + c_\chi} \approx 0.0237.$$

The constant χ^d is the sharp Young bound for the normalized overlap of two balls of the same volume (Lemma 6.4). It governs the concentration of cell volumes in all our constructions.

Theorem 2.3. *Let $(N_d)_{d \geq 1}$ be an arbitrary sequence of norms.*

(i) *For every fixed $0 < \alpha < \alpha_{\mathrm{net}}$, with $p = \alpha d$,*

$$\lim_{d \to \infty} \liminf_{n \to \infty} \mathcal{A}_{n,d,p,N_d} = \lim_{d \to \infty} \limsup_{n \to \infty} \mathcal{A}_{n,d,p,N_d} = \mathcal{C}(\alpha).$$

(ii) *There are numerical constants $c, C > 0$ (any $c < c_\chi/2$ and $C \geq \log 192 + c_\chi$ will do, for instance $c = c_\chi/4$ and $C = 6$) such that, for every sequence $1 \leq p_d = o(d)$, with $r_d = p_d/d$, and all sufficiently large d,*

$$1+(1-2\gamma_{\mathrm{E}})r_d \leq \liminf_{n\to\infty} \mathcal{A}_{n,d,p_d,N_d} \leq \limsup_{n\to\infty} \mathcal{A}_{n,d,p_d,N_d} \leq \Gamma(1+r_d)\Gamma(2+r_d)+Ce^{-cd+Cp_d}. \tag{2.3}$$

In particular, in case (ii) both limits equal $1 + (1 - 2\gamma_{\mathrm{E}})r_d + O(r_d^2)$, with a numerical implied constant. In (ii), "sufficiently large d" means $d \geq d_0$ for a numerical d_0, together with $p_d \leq d/2 - 1$.

Theorem 2.4. *Let (N_d) be an arbitrary sequence of norms and $1 \leq p_d = o(d)$, $r_d = p_d/d$. With the numerical constants c, C of Theorem 2.3, for all sufficiently large d,*

$$1 + (1 - 2\gamma_{\mathrm{E}})r_d \leq B^-_{d,p_d,N_d} \leq B^+_{d,p_d,N_d} \leq \Gamma(1 + r_d)\Gamma(2 + r_d) + Ce^{-cd+Cp_d}. \tag{2.4}$$

In particular $B^{\pm}_{d,p_d,N_d} = 1 + (1 - 2\gamma_{\mathrm{E}})r_d + O(r_d^2)$, uniformly over the sequence of norms.

Theorem 2.5. *Let ω_d be the volume of the Euclidean unit ball and let W_p refer to the Euclidean distance on the torus $\mathbb{T}^d$. For every fixed $\alpha > 0$ and $p = \alpha d$,*

$$\lim_{d\to\infty} \liminf_{n\to\infty} (n\omega_d)^\alpha \, \mathbb{E}W_p^p(\mu_n, \nu_n) = \lim_{d\to\infty} \limsup_{n\to\infty} (n\omega_d)^\alpha \, \mathbb{E}W_p^p(\mu_n, \nu_n) = \mathcal{C}(\alpha). \tag{2.5}$$

In particular the common limit equals $\pi^2/6$ for $\alpha = 1$.

The quantity in (2.5) is $\mathcal{A}^{\mathrm{per}}_{n,d,\alpha d,|\cdot|}$. In the regime $p = \alpha d$ we write α for the ratio p/d, and in the regime $p = o(d)$ we write $r = p/d$.

In Theorems 2.3–2.5 the statements concern the lim inf and the lim sup as $n \to \infty$ at fixed d and fixed norm, followed by the limit $d \to \infty$. The constants c, C and the error terms that survive the limit $n \to \infty$ are independent of the norm; the rates in n may depend on it. The restriction $\alpha < \alpha_{\mathrm{net}}$ in Theorem 2.3(i) comes from the net completion of Section 6.

2.4 Conventions

Order of limits. Throughout, the sample size n, or the torus volume V in the Poissonized problem, is sent to infinity first, at fixed dimension, fixed norm and fixed values of all other parameters. Only afterwards do we let $d \to \infty$. Expressions such as "for all sufficiently large d" refer to this order.

Annealed and pathwise statements. Moment bounds, transport costs and L^p norms of random fields are annealed: they are expectations over the Poisson configuration, and an L^p norm of a random field u on a torus of volume V means the norm in $L^p(\mathbb{P} \otimes \mathrm{d}x/V)$. Marginal identities and capacity constraints of the constructed plans hold for every configuration in the event where they are asserted. We write $P_\alpha(d)$ for a quantity bounded by a polynomial in d whose coefficients depend only on the fixed parameters of the construction; it may change from line to line, but never depends on a stage index or on V.

Marked configurations and labels. On a torus $\Omega_T = (\mathbb{R}/T\mathbb{Z})^d$ of volume $V = T^d$ the target configuration is a finite marked point measure $\eta = \sum_i \delta_{(Y_i,E_i)}$ on $\Omega_T \times (0, \infty)$. These configurations form a standard Borel space, equipped with the evaluation σ-field. Under the Poisson law with intensity $\mathrm{d}y\, e^{-e}\, \mathrm{d}e$, almost surely the sites are distinct and finitely many. All maps below are defined on this full-measure set, with ties resolved by a fixed Borel rule, and arbitrarily on its complement. Each site carries its *physical label* (Y_i, E_i). When a site is added to the configuration, the existing labels are unchanged and a new one

is created. Comparisons between a configuration and the configuration with one added site are made on the union of the labels. A lexicographic Borel enumeration is used only to write sums. The objects we construct (cells, weights, rescue measures, densities) are built from finite minima of continuous costs, integrals of nonnegative Borel kernels and continuous truncations. They are therefore jointly measurable in the configuration and the space variable. Translation-covariant fields are defined from the configuration seen from the observation point, so they are stationary. For a functional Φ of η, we write $D_{z,e}\Phi = \Phi(\eta + \delta_{(z,e)}) - \Phi(\eta)$ for the effect of inserting a site at z with mark e. Integrals over the inserted mark are always taken with respect to $e^{-e}\,\mathrm{d}e$.

3 Exact Poissonization

The proofs of both main theorems work with a Poissonized form of the matching problem: the n sources are replaced by a multiple of Lebesgue measure on a torus of volume V, and the n targets by the sites of a unit Poisson process carrying independent exponential masses. By an identity of Xu [33], the Poissonized cost is an exact Poisson average of the fixed-n matching costs; this is the identity (3.7). Since the fixed-n costs are nonincreasing in n (Lemma 3.3), asymptotics of the Poissonized cost transfer to deterministic sample sizes (Lemma 3.4).

3.1 Xu's identity

Let (X, μ) be a probability space, let $m \geq 1$, and let $c_1, \dots, c_m \in L^1(\mu)$ be real-valued. For w in the simplex $\Delta_{m-1} = \{w \in [0,1]^m : \sum_j w_j = 1\}$ put

$$\mathcal{T}_c(\mu, w) := \inf\Big\{ \int_{\mathsf{X}} \sum_{j=1}^{m} c_j\, a_j \,\mathrm{d}\mu : \; a_j : \mathsf{X} \to [0,1],\ \sum_{j=1}^{m} a_j = 1 \ \mu\text{-a.e.},\ \int a_j \,\mathrm{d}\mu = w_j \Big\}, \quad (3.1)$$

the infimum being over measurable allocations (a_j). The constant allocation $a_j \equiv w_j$ is admissible, and every admissible allocation has cost bounded in absolute value by $\int \max_j |c_j| \,\mathrm{d}\mu$; hence $|\mathcal{T}_c(\mu, w)| \leq \int \max_j |c_j| \,\mathrm{d}\mu < \infty$. A convex combination of admissible allocations for w and w' is admissible for the same convex combination of w and w', so $\mathcal{T}_c(\mu, \cdot)$ is convex on Δ_{m-1}. Being finite and convex, it is continuous on the relative interior Δ°_{m-1} of the simplex, which carries the full mass of the Dirichlet law.

For finite measures μ, ν of equal total mass on measurable spaces X, Y and a measurable cost $c : \mathsf{X} \times \mathsf{Y} \to \mathbb{R}$, we write $\mathrm{OT}_c(\mu, \nu) = \inf_\pi \int c \,\mathrm{d}\pi$, the infimum being over all couplings π of μ and ν. When $c_j = c(\cdot, y_j)$ for points $y_1, \dots, y_m \in \mathsf{Y}$, with $c(\cdot, y_j) \in L^1(\mu)$ so that all the integrals below are defined, the allocation value is the semidiscrete Kantorovich cost:

$$\mathcal{T}_c(\mu, w) = \mathrm{OT}_c\Big(\mu, \sum_{j=1}^{m} w_j \delta_{y_j}\Big). \quad (3.2)$$

Indeed, if the y_j are distinct, a coupling π is concentrated on $\mathsf{X} \times \{y_1, \dots, y_m\}$, and the measures $\pi_j = \pi(\cdot \times \{y_j\})$ satisfy $\sum_j \pi_j = \mu$; hence $\pi_j = a_j \mu$ for an admissible allocation (a_j), and $\int c \,\mathrm{d}\pi = \int \sum_j c_j a_j \,\mathrm{d}\mu$. Conversely, every admissible allocation defines the coupling $\sum_j (a_j \mu) \otimes \delta_{y_j}$. If $y_j = y_k$ for some $j \neq k$, then $c_j = c_k$; replacing (a_j, a_k) by $a_j + a_k$, or splitting an allocation in the proportion $w_j : w_k$, shows that $\mathcal{T}_c(\mu, w)$ is unchanged when the two targets are merged into one of mass $w_j + w_k$ (if $w_j = w_k = 0$ both carry no mass and may be discarded).

Theorem 3.1. *Let (X, μ) be a probability space, $m \geq 1$ and $c_1, \dots, c_m \in L^1(\mu)$. Let $X_1, \dots, X_m$ be i.i.d. with law μ and, independently, let $W \sim \mathrm{Dir}(1, \dots, 1)$. Then*

$$\mathbb{E} \min_{\sigma \in S_m} \sum_{i=1}^{m} c_{\sigma(i)}(X_i) = m \, \mathbb{E} \, \mathcal{T}_c(\mu, W). \tag{3.3}$$

The functions c_j may take both signs and are arbitrary elements of $L^1(\mu)$, possibly discontinuous, equal to one another, or unrelated to any distance; μ may have atoms.

The identity is due to Xu [33, Theorem 2 and Section 2.1]. Since it enters every part of the argument, a self-contained proof is given in Appendix A. The form in which we use it is the following Gamma–Dirichlet reformulation with a common Poisson number of sources and targets.

Corollary 3.2. *Let (X, μ) be a probability space and $c : \mathsf{X} \times \mathsf{X} \to \mathbb{R}$ a measurable cost.*

(i) *Let $m \geq 1$ and let $y_1, \dots, y_m \in \mathsf{X}$ satisfy $c(\cdot, y_j) \in L^1(\mu)$. Let $X_1, \dots, X_m$ be i.i.d. with law μ, let $E_1, \dots, E_m$ be i.i.d.* $\mathrm{Exp}(1)$*, independent of the X_i, and put $S = \sum_{j=1}^{m} E_j$. Then*

$$\mathbb{E} \min_{\sigma \in S_m} \sum_{i=1}^{m} c(X_i, y_{\sigma(i)}) = \mathbb{E} \, \mathrm{OT}_c\Big(S\mu, \sum_{j=1}^{m} E_j \delta_{y_j}\Big). \tag{3.4}$$

(ii) *Assume that X is a compact metric space, μ is a Borel probability measure on X, and c is continuous. Let $V > 0$, let $X_1, X_2, \dots$ and $Y_1, Y_2, \dots$ be i.i.d. with law μ, $E_1, E_2, \dots$ i.i.d.* $\mathrm{Exp}(1)$ *and $N \sim \mathrm{Pois}(V)$, all independent, and put $S = \sum_{j=1}^{N} E_j$. Then*

$$\mathbb{E} \min_{\sigma \in S_N} \sum_{i=1}^{N} c(X_i, Y_{\sigma(i)}) = \mathbb{E} \, \mathrm{OT}_c\Big(S\mu, \sum_{j=1}^{N} E_j \delta_{Y_j}\Big), \tag{3.5}$$

both sides being finite, with the convention that both costs vanish when $N = 0$.

The proof is given in Appendix A.

3.2 The Poissonized cost

On $\Omega_T = (\mathbb{R}/T\mathbb{Z})^d$, of volume $V = T^d$, let $\{(Y_i, E_i)\}_{i \leq N_V}$ be the sites and marks of a Poisson process on $\Omega_T \times (0, \infty)$ with intensity $\mathrm{d}y \, e^{-e} \, \mathrm{d}e$. Thus $N_V \sim \mathrm{Pois}(V)$ and, given N_V, the sites are i.i.d. uniform and the marks i.i.d. $\mathrm{Exp}(1)$. Put $S = \sum_i E_i$ and $\theta = S/V$. For $\alpha > 0$ and $p = \alpha d$ we use the cost and the Poissonized cost

$$c_{d,\alpha}(x, y) = (\omega_d \, \mathrm{dist}(x, y)^d)^\alpha = \omega_d^\alpha \, \mathrm{dist}(x, y)^p, \qquad \mathcal{P}_{d,\alpha}(V) = \frac{1}{V} \mathbb{E} \, \mathrm{OT}_{c_{d,\alpha}}\Big(\theta \, \mathrm{d}x, \sum_i E_i \delta_{Y_i}\Big), \tag{3.6}$$

where dist is the periodized Euclidean distance on Ω_T. Let $b_m = \mathbb{E} W_p^p(\mu_m, \nu_m)$ be the expected matching cost of m uniform sources and m uniform targets on the unit torus, so that $m b_m = \mathbb{E} \min_{\sigma \in S_m} \sum_{i=1}^{m} \mathrm{dist}(X_i, Y_{\sigma(i)})^p$ by the Birkhoff–von Neumann theorem.

We apply Corollary 3.2(ii) with $\mathsf{X} = \Omega_T$, $\mu = \mathrm{d}x/V$ and $c = c_{d,\alpha}$, which is continuous on a compact space; note that $S\mu = \theta \, \mathrm{d}x$. The mechanism is as follows. Conditionally on $N_V = m \geq 1$ and on the sites, write $E_i = S W_i$; then $S \sim \mathrm{Gamma}(m, 1)$ is independent of $W \sim \mathrm{Dir}(1, \dots, 1)$, and $\mathbb{E}[S \mid N_V = m] = m$. By homogeneity in the common total mass,

$$\mathrm{OT}_{c_{d,\alpha}}\Big(\theta \, \mathrm{d}x, \sum_i E_i \delta_{Y_i}\Big) = S \, \mathrm{OT}_{c_{d,\alpha}}\Big(\frac{\mathrm{d}x}{V}, \sum_i W_i \delta_{Y_i}\Big),$$

and Theorem 3.1 (applied at fixed sites and then averaged over them, as in Corollary 3.2(ii)) identifies m times the conditional mean of the last transport cost with the expected cost of matching m independent uniform sources on Ω_T to the sites. Finally, the dilation $x \mapsto Tx$ from the unit torus onto Ω_T multiplies dist by T, hence the cost $c_{d,\alpha}$ by $\omega_d^\alpha T^p = \omega_d^\alpha V^\alpha$. Given $N_V = m$, the expected matching cost on Ω_T is therefore $\omega_d^\alpha V^\alpha m b_m$, and

$$\mathcal{P}_{d,\alpha}(V) = \omega_d^\alpha V^{\alpha-1} \mathbb{E}[N_V b_{N_V}]. \tag{3.7}$$

The derivation uses only the dilation and the homogeneity of the cost. Hence (3.7) holds with the same proof, and with v_d in place of ω_d, when dist is the periodized distance of any norm N_d and the cost is $(v_d \operatorname{dist}(x,y)^d)^\alpha$; it also holds on the cube $[0,T]^d$ in place of Ω_T, with b_m the corresponding matching cost on $[0,1]^d$.

3.3 Monotonicity in the sample size and de-Poissonization

The two lemmas of this subsection are known in the setting of uniform samples and power costs: monotonicity in the sample size is proved in [14, Section 6] by averaging over subsamples and using the convexity of the transport cost, and a general de-Poissonization result is [14, Proposition 6.1]. We include short proofs in the generality used here: arbitrary costs for the monotonicity, and a single Poisson number common to sources and targets, which is the form produced by Xu's identity.

Lemma 3.3. *Let μ and ν be probability measures on measurable spaces X and Y, and let $c : \mathsf{X} \times \mathsf{Y} \to [0,\infty)$ be measurable. Let $X_1, X_2, \dots$ be i.i.d. with law μ and, independently, $Y_1, Y_2, \dots$ i.i.d. with law ν, and put*

$$b_n = \mathbb{E}\,\frac{1}{n} \min_{\sigma \in S_n} \sum_{i=1}^{n} c(X_i, Y_{\sigma(i)}) \in [0,\infty].$$

Then $b_{n+1} \le b_n$ for every $n \ge 1$.

Proof. We compare the $(n+1)$-matching with an average of n-matchings obtained by deleting one source and one target. Fix a nonnegative $(n+1) \times (n+1)$ matrix $C = (c_{ij})$. For $a, b \in \{1, \dots, n+1\}$ let $C^{\widehat{a},\widehat{b}}$ be the matrix with row a and column b deleted, and let $P^{a,b}$ be the incidence matrix of an optimal assignment for $C^{\widehat{a},\widehat{b}}$, extended by zeros on row a and column b. In $\overline{P} = (n+1)^{-2} \sum_{a,b} P^{a,b}$ every row and every column sums to $n/(n+1)$, since a given row (or column) is matched exactly once whenever it survives the deletion. Hence $\frac{n+1}{n}\overline{P}$ is doubly stochastic. By the Birkhoff–von Neumann theorem it is a convex combination of permutation matrices, and since the cost is linear,

$$\min_{\sigma \in S_{n+1}} \sum_{i=1}^{n+1} c_{i\sigma(i)} \le \frac{n+1}{n} \cdot \frac{1}{(n+1)^2} \sum_{a,b=1}^{n+1} \min_{\tau} \sum_{i \ne a} c_{i\tau(i)}, \tag{3.8}$$

where τ ranges over bijections from $\{1, \dots, n+1\} \setminus \{a\}$ onto $\{1, \dots, n+1\} \setminus \{b\}$. Apply (3.8) to $c_{ij} = c(X_i, Y_j)$. For fixed (a,b), the remaining sources and targets are n i.i.d. sources and n independent i.i.d. targets, so each summand on the right has expectation $n b_n$. Taking expectations and dividing by $n+1$ gives $b_{n+1} \le b_n$. □

Lemma 3.4. *Let $\kappa > 0$, $\alpha > 0$ and $C_0 < \infty$, and let $(b_m)_{m \ge 1}$ be a nonincreasing sequence with $0 \le b_m \le C_0$ for all $m \ge 1$. For $V > 0$ let $N_V \sim \mathrm{Pois}(V)$ and put*

$$\mathcal{P}(V) = \kappa V^{\alpha-1} \mathbb{E}[N_V b_{N_V}],$$

where $N_V b_{N_V} = 0$ on $\{N_V = 0\}$. Then, for every $t \in (0,1)$,

$$\begin{aligned} \limsup_{n\to\infty} \kappa n^\alpha b_n &\le (1-t)^{-\alpha} \limsup_{V\to\infty} \mathcal{P}(V), \\ \liminf_{n\to\infty} \kappa n^\alpha b_n &\ge (1+t)^{-\alpha} \liminf_{V\to\infty} \mathcal{P}(V). \end{aligned} \tag{3.9}$$

Consequently $\limsup_n \kappa n^\alpha b_n \le \limsup_{V\to\infty} \mathcal{P}(V)$ *and* $\liminf_n \kappa n^\alpha b_n \ge \liminf_{V\to\infty} \mathcal{P}(V)$*, in* $[0,\infty]$.

Proof. This is a variant of [14, Proposition 6.1], in which the two samples have independent Poisson sizes; the proof is the same Poisson-tail sandwich. We use the Poisson tail bound: for every $s \in (0,1)$ there is $c_s > 0$ with $\mathbb{P}(|N_V - V| \ge sV) \le 2e^{-c_s V}$ for all $V > 0$; for instance $c_s = \min\{(1+s)\log(1+s) - s,\, s^2/2\}$. We also use $\mathbb{E}[N_V \mathbf{1}_{\{N_V \ge k+1\}}] = V\,\mathbb{P}(N_V \ge k)$ for $k \ge 0$.

For the upper bound fix n and take $V = (1-t)n$. Since $b_m \ge b_n$ for $m \le n$ and $b_m \ge 0$ for all m,

$$\mathbb{E}[N_V b_{N_V}] \ge b_n \mathbb{E}[N_V \mathbf{1}_{\{N_V \le n\}}] = b_n V\big(1 - \mathbb{P}(N_V \ge n)\big) \ge b_n V(1 - 2e^{-c_t V}),$$

because $n = V/(1-t) \ge (1+t)V$. Multiplying by $\kappa V^{\alpha-1}$ gives $\kappa n^\alpha b_n \le (1-t)^{-\alpha}(1 - 2e^{-c_t V})^{-1}\mathcal{P}(V)$ with $V = (1-t)n$, and the first inequality in (3.9) follows.

For the lower bound take $V = (1+t)n$. Since $b_m \le b_n$ for $m \ge n$ and $b_m \le C_0$ for all m,

$$\mathbb{E}[N_V b_{N_V}] \le b_n V + C_0\, n\, \mathbb{P}(N_V < n) \le b_n V + 2C_0 V e^{-c_s V}, \qquad s = \frac{t}{1+t},$$

because $n = (1-s)V$. Hence $\mathcal{P}(V) \le (1+t)^\alpha \kappa n^\alpha b_n + 2\kappa C_0 V^\alpha e^{-c_s V}$, and the second inequality in (3.9) follows.

Letting $t \downarrow 0$ gives the last assertion. □

We apply Lemma 3.4 with $\kappa = \omega_d^\alpha$ and the sequence b_m of (3.7). It is nonincreasing by Lemma 3.3 and satisfies $0 \le b_m \le (\sqrt{d}/2)^p$, the pth power of the diameter of the unit torus. Since $(n\omega_d)^\alpha b_n$ is the normalized cost (2.1) for the Euclidean norm, we obtain, at every fixed d,

$$\limsup_{n\to\infty} (n\omega_d)^\alpha b_n \le \limsup_{V\to\infty} \mathcal{P}_{d,\alpha}(V), \qquad \liminf_{n\to\infty} (n\omega_d)^\alpha b_n \ge \liminf_{V\to\infty} \mathcal{P}_{d,\alpha}(V). \tag{3.10}$$

The same holds for any norm and for the cube, with v_d in place of ω_d and the corresponding diameter in place of $\sqrt{d}/2$.

4 The mean-field variational problem

This section collects the mean-field objects on which both theorems rest. Theorem 4.1 represents the benchmark value $\mathcal{C}(r) = \beta_{\mathrm{bip}}(1/r)$ as an exact optimization over a marked Poisson race, and yields the explicit bracket (4.4) used in the vanishing-ratio regime. For a fixed ratio, the capped cavity law at a finite cutoff λ supplies the dual prices $h(E_i)$ of the marked targets and a trial value $C^{\mathrm{trial}}_{q,\lambda}$, which tends to $\beta_{\mathrm{bip}}(q)$ as $\lambda \to \infty$ (Proposition 4.4). Lemma 4.3 shows that the upper atom of the capped law, which carries the unmatched mass at the cutoff, has probability at most $e^{-c_q \lambda^q}$ for $\lambda \ge \lambda_0(q)$; this is why a cutoff polynomial in d suffices.

4.1 The Poisson race

Fix $0 < r < 1$ and let $\{(U_i, E_i)\}_{i\ge 1}$ be the atoms of a Poisson point process on $(0,\infty)^2$ with intensity $\mathrm{d}u\, e^{-e}\,\mathrm{d}e$. A selector ι chooses one atom (U_ι, E_ι), measurably and possibly with auxiliary randomization. We call it balanced if the selected mark has the size-biased exponential law

$$\mathbb{P}(E_\iota \in \mathrm{d}e) = e\, e^{-e}\,\mathrm{d}e, \qquad \text{that is,} \quad E_\iota \sim \mathrm{Gamma}(2,1). \tag{4.1}$$

For a measurable price $h : (0,\infty) \to \mathbb{R}$ for which $\mathbb{E}[E\,|h(E)|] < \infty$ and the first expectation below is finite, put

$$\Psi_r(h) = \mathbb{E}\inf_i\{U_i^r - h(E_i)\} + \mathbb{E}[E\,h(E)], \qquad E \sim \mathrm{Exp}(1). \tag{4.2}$$

The expression is unchanged if a constant is added to h. The integrability holds for bounded prices and for the logarithmic and cavity prices used below.

Theorem 4.1. *Let $0 < r < 1$ and $q = 1/r > 1$, and let X_q be the uncapped cavity valuation, the unique solution provided by Theorem 2.2(a), with $F_q(t) = \mathbb{P}(X_q \ge t)$. Then*

$$\mathcal{C}(r) = -q\int_{\mathbb{R}} F_q \log F_q\,\mathrm{d}t = \sup_h \Psi_r(h) = \inf_{\iota \text{ balanced}} \mathbb{E}U_\iota^r. \tag{4.3}$$

The supremum is attained by the cavity quantile price $h_q = G_q^{-1}$, where $G_q = -\log F_q$, and the infimum by the selector ι_q choosing the least reduced cost $U_i^r - h_q(E_i)$. Moreover,

$$1 + (1 - 2\gamma_{\mathrm{E}})r \le \mathcal{C}(r) \le \Gamma(1+r)\Gamma(2+r), \tag{4.4}$$

and, as $r \downarrow 0$,

$$\mathcal{C}(r) = 1 + (1 - 2\gamma_{\mathrm{E}})r + O(r^2), \qquad \mathcal{C}(0+) = 1, \quad \mathcal{C}'(0+) = 1 - 2\gamma_{\mathrm{E}}. \tag{4.5}$$

Here $\mathcal{C}'(0+)$ is the right derivative at 0 of $\mathcal{C}$ extended by $\mathcal{C}(0) := 1$. The logarithmic endpoint is the balanced optimization

$$\inf_{\iota \text{ balanced}} \mathbb{E}\log U_\iota = 1 - 2\gamma_{\mathrm{E}}, \tag{4.6}$$

where the infimum is over balanced selectors with integrable logarithmic cost; it is attained by the ratio race $\iota_0 = \arg\min_i U_i/E_i$.

Proof. By Theorem 2.2(a), $G_q(t) = \mathbb{E}(X_q + t)_+^q$ and $\mathcal{C}(r) = \beta_{\mathrm{bip}}(q) = -q\int F_q \log F_q$. Since $F_q > 0$ on $\mathbb{R}$, X_q is unbounded above, so $G_q(t) = \mathbb{E}(t + X_q)_+^q$ is continuous and strictly increasing from 0 at $-\infty$ to ∞, and $h_q = G_q^{-1} : (0,\infty) \to \mathbb{R}$ is well defined. The monotone inverse $h_q = G_q^{-1}$ pushes an Exp(1) mark to the law of X_q. In the Poisson race let

$$\phi_q = \inf_i\big(U_i^r - h_q(E_i)\big), \qquad \iota_q = \arg\min_i\big(U_i^r - h_q(E_i)\big).$$

The number of atoms with $U_i^r - h_q(E_i) \le t$ is Poisson with mean $\mathbb{E}(t + h_q(E))_+^q$, so the void probabilities give

$$\mathbb{P}(\phi_q > t) = \exp\big[-\mathbb{E}(t + h_q(E))_+^q\big] = e^{-G_q(t)} = F_q(t).$$

Thus ϕ_q has the law of X_q, and $\Psi_r(h_q) = \mathbb{E}X_q + \mathbb{E}[E\,h_q(E)]$. Both terms are finite, since X_q has finite moments by Theorem 2.2(a) and $\mathbb{E}[E|h_q(E)|] \le (\mathbb{E}E^2)^{1/2}(\mathbb{E}X_q^2)^{1/2}$.

Mecke's formula identifies the law of the selected mark. A candidate (u, e) wins precisely if no other atom has reduced cost smaller than $u^r - h_q(e)$. Therefore, using $\int_0^\infty F_q(u^r - s)\,\mathrm{d}u = \mathbb{E}(X_q + s)_+^q = G_q(s)$,

$$\mathbb{P}(E_{\iota_q} \in \mathrm{d}e) = e^{-e}\int_0^\infty F_q\big(u^r - h_q(e)\big)\,\mathrm{d}u\;\mathrm{d}e = e^{-e}G_q(h_q(e))\,\mathrm{d}e = e\,e^{-e}\,\mathrm{d}e.$$

Hence ι_q is balanced. Since $\phi_q = U_{\iota_q}^r - h_q(E_{\iota_q})$ and $\mathbb{E}h_q(E_{\iota_q}) = \mathbb{E}[E\,h_q(E)]$, we have

$$\Psi_r(h_q) = \mathbb{E}U_{\iota_q}^r. \tag{4.7}$$

We now compute this cost. For a candidate of physical cost $t = u^r$, which has intensity $qt^{q-1}\,\mathrm{d}t$, and an independent candidate price X_q' with the law of X_q, Campbell's formula gives

$$\mathbb{E}U_{\iota_q}^r = \int_0^\infty qt^q\,\mathbb{E}F_q(t - X_q')\,\mathrm{d}t = \frac{q}{q+1}\mathbb{E}(X_q + X_q')_+^{q+1} = -q\int_{\mathbb{R}} F_q(t)\log F_q(t)\,\mathrm{d}t = \mathcal{C}(r).$$

The third equality follows by Fubini from $G_q(t) = \mathbb{E}(X_q + t)_+^q$. On the other hand, for every balanced selector ι and every admissible h,

$$\Psi_r(h) \le \mathbb{E}\big[U_\iota^r - h(E_\iota)\big] + \mathbb{E}[E\,h(E)] = \mathbb{E}U_\iota^r,$$

because $\mathbb{E}h(E_\iota) = \mathbb{E}[E\,h(E)]$. Taking $\iota = \iota_q$ and using (4.7) proves (4.3) together with the attainment statements.

For the lower bound in (4.4) take $h(e) = r\log e$. The inequality $u^r \ge 1 + r\log u$ gives pointwise

$$\inf_i\{U_i^r - r\log E_i\} \ge 1 + r\log\Big(\inf_i \frac{U_i}{E_i}\Big).$$

The map $(u, e) \mapsto (u/e, e)$ sends the intensity to $e\,e^{-e}\,\mathrm{d}t\,\mathrm{d}e$, so the smallest ratio $T = \inf_i U_i/E_i$ is $\mathrm{Exp}(1)$ and $\mathbb{E}\log T = -\gamma_{\mathrm{E}}$. The other dual term is $\mathbb{E}[E\,r\log E] = r\Gamma'(2) = r(1 - \gamma_{\mathrm{E}})$. Consequently $\Psi_r(r\log e) \ge 1 + (1 - 2\gamma_{\mathrm{E}})r$.

For the upper bound take the ratio winner $\iota_0 = \arg\min_i U_i/E_i$. The same factorization of the intensity gives

$$T = U_{\iota_0}/E_{\iota_0} \sim \mathrm{Exp}(1), \qquad E_{\iota_0} \sim \mathrm{Gamma}(2, 1), \qquad T \perp E_{\iota_0}.$$

In particular ι_0 is balanced, and (4.3) gives $\mathcal{C}(r) \le \mathbb{E}(TE_{\iota_0})^r = \Gamma(1 + r)\Gamma(2 + r)$. Since $\Gamma(1+r)\Gamma(2+r) = (1+r)\Gamma(1+r)^2 = 1 + (1 - 2\gamma_{\mathrm{E}})r + O(r^2)$, the two bounds differ by $O(r^2)$, which proves (4.5). Finally, for every balanced ι with integrable logarithmic cost,

$$\mathbb{E}\log U_\iota = \mathbb{E}\log(U_\iota/E_\iota) + \mathbb{E}\log E_\iota \ge \mathbb{E}\log T + \psi(2) = 1 - 2\gamma_{\mathrm{E}},$$

with equality at ι_0. This proves (4.6). □

The expansion (4.5) concerns the mean-field value $\mathcal{C}(r)$ itself, and the endpoint (4.6) is attained by the explicit ratio race; the cavity laws X_q enter only through the identity (4.3) at each fixed q.

4.2 The capped cavity law

Fix $q > 0$ and a cutoff $\lambda > 0$, and put $a = \lambda/2$ and $r = 1/q$. By Theorem 2.2(b) there is a unique law of a random variable X with values in $[-a, a]$, the capped valuation, whose survival function $F(t) = \mathbb{P}(X \geq t)$ and $G = -\log F$ satisfy

$$F(t) = \exp[-G(t)], \qquad G(t) = \mathbb{E}(t + X)_+^q = q \int_{-a}^{a} (t + x)_+^{q-1} F(x)\, \mathrm{d}x, \qquad -a \leq t \leq a. \tag{4.8}$$

The function G is nondecreasing on $[-a, a]$, continuous by dominated convergence, and $G(-a) = 0$. It is strictly increasing, since $F(a) = e^{-G(a)} > 0$ means that X has an atom at a. Put

$$\bar{e}_\lambda = G(a), \qquad \text{so that} \quad \mathbb{P}(X = a) = e^{-\bar{e}_\lambda},$$

and define the price map $h_\lambda(e) = G^{-1}(e)$ for $0 \leq e < \bar{e}_\lambda$ and $h_\lambda(e) = a$ for $e \geq \bar{e}_\lambda$. For $E \sim \mathrm{Exp}(1)$ we have $\mathbb{P}(h_\lambda(E) \geq t) = \mathbb{P}(E \geq G(t)) = F(t)$ for $t \in (-a, a]$, so $h_\lambda(E)$ has the law of X, including the atom at a. Finally put

$$J = \int_{-a}^{a} F(t)G(t)\, \mathrm{d}t, \qquad E_{q,\lambda} = \frac{q}{q+1} \mathbb{E}(X + X')_+^{q+1}, \qquad C_{q,\lambda}^{\mathrm{trial}} = \mathbb{E}X + \mathbb{E}[E\, h_\lambda(E)],$$

where X' is an independent copy of X. When the dependence on the parameters matters we write $J = J_{q,\lambda}$.

Two elementary identities relate these quantities. First, writing $\frac{1}{q+1}(X + X')_+^{q+1} = \int_{-a}^{a} (s + X')_+^q \mathbf{1}_{\{s \leq X\}}\, \mathrm{d}s$, which holds because $-a \leq X \leq a$, and taking expectations, we get

$$E_{q,\lambda} = q \int_{-a}^{a} F(s)\, \mathbb{E}(s + X')_+^q\, \mathrm{d}s = q \int_{-a}^{a} F(s)G(s)\, \mathrm{d}s = qJ. \tag{4.9}$$

This is the area formula for the edge cost of the relaxed matching problem; cf. [31, Sections 2–3]. Second, h_λ is nondecreasing and continuous on $[0, \infty)$, with $h_\lambda(0) = -a$, and constant on $[\bar{e}_\lambda, \infty)$. Integrating by parts against $(e-1)e^{-e}\, \mathrm{d}e = \mathrm{d}(-e\, e^{-e})$ and substituting $e = G(t)$ on $[0, \bar{e}_\lambda]$,

$$\mathbb{E}[(E-1)h_\lambda(E)] = \int_{[0,\infty)} e\, e^{-e}\, \mathrm{d}h_\lambda(e) = \int_{-a}^{a} G(t)e^{-G(t)}\, \mathrm{d}t = J, \qquad \text{hence} \quad C_{q,\lambda}^{\mathrm{trial}} = 2\mathbb{E}X + J. \tag{4.10}$$

The capped law is the solution of a capped version of the Poisson race of Section 4. Let $\{(U_i, E_i)\}$ be a Poisson process on $(0, \infty)^2$ with intensity $\mathrm{d}u\, e^{-e}\, \mathrm{d}e$. For a measurable price $h : (0, \infty) \to (-\infty, a]$ with $\mathbb{E}[E|h(E)|] < \infty$ put

$$\Psi^\lambda(h) = \mathbb{E} \min\Big\{a, \inf_i (U_i^r - h(E_i))\Big\} + \mathbb{E}[E\, h(E)]. \tag{4.11}$$

Since $h \leq a$, the minimum lies in $[-a, a]$, so $\Psi^\lambda(h)$ is well defined and finite. The first term is the expected cost of a source which pays the reduced cost of the best target, or a if no target has reduced cost below a.

Proposition 4.2. *Let $q > 0$, $\lambda > 0$ and $r = 1/q$. Then:*

(i) $C_{q,\lambda}^{\mathrm{trial}} = \Psi^\lambda(h_\lambda) = \max\{\Psi^\lambda(h) : h \leq a \text{ admissible}\}$;

(ii) $C_{q,\lambda}^{\mathrm{trial}} = E_{q,\lambda} + \lambda e^{-\bar{e}_\lambda} = qJ + \lambda e^{-\bar{e}_\lambda}$;

(iii) $\lambda \mapsto C_{q,\lambda}^{\mathrm{trial}}$ *is nondecreasing on* $(0, \infty)$, *and* $0 \leq C_{q,\lambda}^{\mathrm{trial}} \leq \Gamma(1 + r)\Gamma(2 + r)$.

Proof. Write $\varphi_h = \min\{a, \inf_i(U_i^r - h(E_i))\}$. Since $h \le a$, only atoms with $U_i^r < 2a$ can have reduced cost below a, and there are finitely many of them; ties have probability zero.

The law of φ_{h_λ}. For $t \le a$, the number of atoms with $U_i^r - h_\lambda(E_i) < t$ is Poisson with mean $\iint \mathbf{1}_{\{u^r < t + h_\lambda(e)\}}\,\mathrm{d}u\, e^{-e}\,\mathrm{d}e = \mathbb{E}(t + h_\lambda(E))_+^q = \mathbb{E}(t+X)_+^q = G(t)$. Hence $\mathbb{P}(\varphi_{h_\lambda} \ge t) = e^{-G(t)} = F(t)$ for $-a \le t \le a$. For $t < -a$ the mean is zero, since $U_i^r - h_\lambda(E_i) > -a$, so $\varphi_{h_\lambda} \ge -a$; also $\varphi_{h_\lambda} \le a$, so φ_{h_λ} has the law of X. Therefore $\Psi^\lambda(h_\lambda) = \mathbb{E}X + \mathbb{E}[E h_\lambda(E)] = C_{q,\lambda}^{\mathrm{trial}}$.

A capped selector. Let ι be the atom minimizing $U_i^r - h_\lambda(E_i)$ if this minimum is smaller than a, and let $\iota = \varnothing$ ("unmatched") otherwise. By Mecke's formula, a candidate (u, e) is selected precisely if $u^r - h_\lambda(e) < a$ and no other atom has reduced cost at most $u^r - h_\lambda(e)$, so

$$\mathbb{P}(\iota \ne \varnothing,\ E_\iota \in \mathrm{d}e) = e^{-e}\,\mathrm{d}e \int_0^\infty F(u^r - h_\lambda(e))\mathbf{1}_{\{u^r - h_\lambda(e) < a\}}\,\mathrm{d}u.$$

With $u = t^q$ and $x = t - h_\lambda(e)$ the integral equals $q\int_{-h_\lambda(e)}^a (x + h_\lambda(e))^{q-1}F(x)\,\mathrm{d}x = G(h_\lambda(e))$ by (4.8), since $h_\lambda(e) \in [-a, a]$. As $G(h_\lambda(e)) = e \wedge \bar{e}_\lambda$,

$$\mathbb{P}(\iota \ne \varnothing,\ E_\iota \in \mathrm{d}e) = (e \wedge \bar{e}_\lambda)\, e^{-e}\,\mathrm{d}e, \qquad \mathbb{P}(\iota = \varnothing) = \mathbb{P}(\varphi_{h_\lambda} = a) = e^{-\bar{e}_\lambda}. \tag{4.12}$$

Weak duality. Let $h \le a$ be admissible. Pointwise, $\varphi_h \le (U_\iota^r - h(E_\iota))\mathbf{1}_{\{\iota \ne \varnothing\}} + a\mathbf{1}_{\{\iota = \varnothing\}}$. The variable $h(E_\iota)\mathbf{1}_{\{\iota \ne \varnothing\}}$ is integrable, since by (4.12) its absolute expectation is at most $\mathbb{E}[E|h(E)|]$. Taking expectations and using (4.12),

$$\Psi^\lambda(h) \le \mathbb{E}\big[U_\iota^r \mathbf{1}_{\{\iota \ne \varnothing\}}\big] + a\, e^{-\bar{e}_\lambda} + \int_0^\infty h(e)(e - e \wedge \bar{e}_\lambda)e^{-e}\,\mathrm{d}e.$$

Since $h \le a$ and $\int_0^\infty (e - \bar{e}_\lambda)_+ e^{-e}\,\mathrm{d}e = e^{-\bar{e}_\lambda}$, the right side is at most

$$R_\lambda = \mathbb{E}\big[U_\iota^r \mathbf{1}_{\{\iota \ne \varnothing\}}\big] + 2a\, e^{-\bar{e}_\lambda}. \tag{4.13}$$

For $h = h_\lambda$ both inequalities are equalities: the first by the definition of ι, the second because $h_\lambda = a$ on $[\bar{e}_\lambda, \infty)$. Hence $\Psi^\lambda(h) \le R_\lambda = \Psi^\lambda(h_\lambda)$ for every admissible $h \le a$, which proves (i).

The matched cost. By Mecke's formula and the substitution $t = u^r$, with $X' = h_\lambda(E')$ for an independent $E' \sim \mathrm{Exp}(1)$,

$$\mathbb{E}\big[U_\iota^r \mathbf{1}_{\{\iota \ne \varnothing\}}\big] = \mathbb{E}\int_0^\infty qt^q\, F(t - X')\mathbf{1}_{\{t < a + X'\}}\,\mathrm{d}t = \mathbb{E}\int_0^\infty qt^q\, \mathbf{1}_{\{t \le X + X'\}}\,\mathrm{d}t = E_{q,\lambda}, \tag{4.14}$$

where in the second equality X is independent of X' and we used that $t \le X + X'$ implies $t \le a + X'$; the event $\{t = a + X'\}$ has zero Lebesgue measure in t. Together with (4.13) and (4.9) this gives (ii).

Monotonicity and the bound. If $\lambda < \lambda'$, every price $h \le \lambda/2$ is admissible for λ', and $\Psi^{\lambda'}(h) \ge \Psi^\lambda(h)$ because $\min\{\lambda'/2, x\} \ge \min\{\lambda/2, x\}$. By (i), $C_{q,\lambda'}^{\mathrm{trial}} \ge C_{q,\lambda}^{\mathrm{trial}}$. For the upper bound, let $h \le a$ be admissible and let $\iota_0 = \arg\min_i U_i/E_i$ be the ratio winner. The map $(u, e) \mapsto (u/e, e)$ sends the intensity to $e\, e^{-e}\,\mathrm{d}t\,\mathrm{d}e$, so $T = U_{\iota_0}/E_{\iota_0} \sim \mathrm{Exp}(1)$ is independent of $E_{\iota_0} \sim \mathrm{Gamma}(2, 1)$, whatever $r > 0$. Then $\varphi_h \le U_{\iota_0}^r - h(E_{\iota_0})$, and $\mathbb{E}h(E_{\iota_0}) = \mathbb{E}[E\, h(E)]$, so

$$\Psi^\lambda(h) \le \mathbb{E}U_{\iota_0}^r = \mathbb{E}(TE_{\iota_0})^r = \Gamma(1 + r)\Gamma(2 + r).$$

Taking $h = h_\lambda$ gives the upper bound in (iii). Finally $C_{q,\lambda}^{\mathrm{trial}} \ge 0$ by (ii). □

By (ii) and (4.10), the two expressions of the trial value also give the identity $\mathbb{E}X = a\, e^{-\bar{e}_\lambda} + \frac{q-1}{2}J$. The quantity $C_{q,\lambda}^{\mathrm{trial}}$ is the limiting cost per pair of the λ-relaxed bipartite matching problem in pseudo-dimension q, in which each unmatched vertex pays $\lambda/2$; see [31] and [18, Section 3.5]. We do not use this interpretation.

4.3 Saturation and the trial limit

Lemma 4.3. *Let $q > 0$, $r = 1/q$, and put $c_q = \min\{e^{-1}2^{-q}, 2^{-2q-1}\}$ and $\lambda_0(q) = 12\,\Gamma(1+r)\Gamma(2+r)/q$. Then $\bar{e}_\lambda \ge c_q\lambda^q$ for every $\lambda \ge \lambda_0(q)$.*

Proof. If $F(0) \ge e^{-1}$, then $\bar{e}_\lambda = \mathbb{E}(a+X)_+^q \ge a^q F(0) \ge e^{-1}a^q = e^{-1}2^{-q}\lambda^q$.

Otherwise $G(0) > 1$, and the monotonicity of G gives $F(t) \le F(t)G(t)$ for $t \ge 0$. Hence $\mathbb{E}X_+ = \int_0^a F(t)\,\mathrm{d}t \le J$. By (4.10) and Proposition 4.2(iii), $\mathbb{E}X = (C_{q,\lambda}^{\mathrm{trial}} - J)/2 \ge -J/2$, so $\mathbb{E}X_- = \mathbb{E}X_+ - \mathbb{E}X \le \frac{3}{2}J$. By Proposition 4.2(ii)–(iii), $qJ \le C_{q,\lambda}^{\mathrm{trial}} \le \Gamma(1+r)\Gamma(2+r)$. By Markov's inequality,

$$\mathbb{P}(X < -a/2) \le \frac{2\mathbb{E}X_-}{a} \le \frac{3\Gamma(1+r)\Gamma(2+r)}{qa} \le \frac{1}{2}$$

as soon as $a \ge 6\Gamma(1+r)\Gamma(2+r)/q$, that is, $\lambda \ge \lambda_0(q)$. Then

$$\bar{e}_\lambda = \mathbb{E}(a+X)_+^q \ge (a/2)^q\,\mathbb{P}(X \ge -a/2) \ge 2^{-q-1}a^q = 2^{-2q-1}\lambda^q.$$

In both cases $\bar{e}_\lambda \ge c_q\lambda^q$. □

In particular, for the cutoff $\lambda = L^\alpha$ with $L = A_\alpha d$ used from Section 8 on, $\lambda^q = L$, and once $\lambda \ge \lambda_0(q)$ the quantities $e^{-\bar{e}_\lambda}$ and $\lambda e^{-\bar{e}_\lambda}$ are bounded by $(1+L+\lambda)e^{-c_q A_\alpha d}$, hence exponentially small in d.

Proposition 4.4. *For every $q > 0$,*

$$\lim_{\lambda\to\infty} C_{q,\lambda}^{\mathrm{trial}} = \lim_{\lambda\to\infty} E_{q,\lambda} = \beta_{\mathrm{bip}}(q). \tag{4.15}$$

Proof. By Proposition 4.2(iii), $C_{q,\lambda}^{\mathrm{trial}}$ is nondecreasing and bounded, so it has a finite limit ℓ. By Lemma 4.3, $\lambda e^{-\bar{e}_\lambda} \le \lambda e^{-c_q\lambda^q} \to 0$, so by Proposition 4.2(ii) also $E_{q,\lambda} \to \ell$. By Theorem 2.2(c), $\limsup_{\lambda\to\infty} E_{q,\lambda} = \beta_{\mathrm{bip}}(q)$; hence $\ell = \beta_{\mathrm{bip}}(q)$. □

Remark 4.5. Proposition 4.2(iii) and Proposition 4.4 give $\mathcal{C}(\alpha) \le \Gamma(1+\alpha)\Gamma(2+\alpha)$ for every $\alpha > 0$. The lower bound in (4.4) also extends to every $\alpha > 0$. Take $h_a(e) = \min\{a, \alpha\log e\}$ in (i), with $r = \alpha$. Since $h_a \le \alpha\log e$ and $u^\alpha \ge 1 + \alpha\log u$,

$$\min\Big\{a, \inf_i (U_i^\alpha - h_a(E_i))\Big\} \ge \min\{a, 1+\alpha\log T\}, \qquad T = \inf_i \frac{U_i}{E_i} \sim \mathrm{Exp}(1).$$

Monotone convergence (note that $1+\alpha\log T$ is integrable) and dominated convergence for $\mathbb{E}[E\,h_a(E)] \to \alpha(1-\gamma_{\mathrm{E}})$ give $\liminf_{a\to\infty} \Psi^{2a}(h_a) \ge 1 + (1-2\gamma_{\mathrm{E}})\alpha$. Hence $1 + (1-2\gamma_{\mathrm{E}})\alpha \le \mathcal{C}(\alpha) \le \Gamma(1+\alpha)\Gamma(2+\alpha)$ for all $\alpha > 0$, without using Theorem 2.2(a).

4.4 The logistic case

For $q = 1$ the capped law is explicit; it is a version, with a finite cutoff, of the classical logistic solution [1]. With $a = \lambda/2$ let $c_a \in (1,2)$ solve $c_a = 1 + e^{-ac_a}$ and put

$$F_a(t) = \frac{c_a}{1 + e^{c_a t}}, \quad -a \le t \le a, \qquad F_a(a) = e^{-ac_a}, \quad \bar{e}_\lambda = ac_a.$$

Indeed $F_a'(t) = -F_a(t)F_a(-t)$ and $F_a(-a) = 1$, which is the equation (4.8) for $q = 1$, where $G(t) = \int_{-t}^a F(x)\,\mathrm{d}x$. By Proposition 4.2(ii) the trial value is

$$C_{1,\lambda}^{\mathrm{trial}} = \int_{-a}^a F_a(t)(-\log F_a(t))\,\mathrm{d}t + 2aF_a(a), \qquad \lim_{\lambda\to\infty} C_{1,\lambda}^{\mathrm{trial}} = \frac{\pi^2}{6}.$$

The limit follows from $c_a \to 1$, dominated convergence (for $a \ge 1$ the integrand is bounded by $2e^{-|t|}(2|t|+1)$, and $2aF_a(a) \le 2ae^{-a} \to 0$), and the substitution $s = (1+e^t)^{-1}$, which gives $\int_{\mathbb{R}}(1+e^t)^{-1}\log(1+e^t)\,\mathrm{d}t = \int_0^1 \frac{-\log s}{1-s}\,\mathrm{d}s = \frac{\pi^2}{6}$. This agrees with Proposition 4.4 and the value $\beta_{\mathrm{bip}}(1) = \pi^2/6$ [1].

5 The vanishing-ratio regime

This section proves Theorem 2.4. The lower bound is a dual estimate on the torus which holds in every dimension, for every norm and every exponent (Proposition 5.1). The upper bound follows from the ordinary-cube Theorem 2.3 (ii), proved in Sections 6 and 7, because periodic distances are smaller than ordinary ones.

We use the normalized costs $\mathcal{A}^{\mathrm{per}}_{n,d,p,N_d}$ (torus) and $\mathcal{A}_{n,d,p,N_d}$ (cube) of (2.1), and the accumulation values $B^{\pm}_{d,p,N_d}$ of Section 2.

Proposition 5.1. *For every $d \geq 1$, every norm N_d on $\mathbb{R}^d$ and every $p \geq 1$, with $r = p/d$,*

$$B^{-}_{d,p,N_d} \geq 1 + (1 - 2\gamma_{\mathrm{E}})r.$$

The proof below uses the exponent only through the inequality $t^r \geq 1 + r \log t$, valid for all $r, t > 0$; hence the same bound holds for every exponent $p > 0$, that is, for every $r > 0$.

Proof. We dilate the unit torus to $\Omega_T = (\mathbb{R}/T\mathbb{Z})^d$ with $V = T^d = n$, and write $\|x - y\|_{d,T} = \min_{k \in \mathbb{Z}^d} \|x - y + Tk\|_d$ for the periodized normalized distance. Dilation multiplies p-costs by $T^p = n^r$, so

$$\mathcal{A}^{\mathrm{per}}_{n,d,p,N_d} = \frac{1}{n}\,\mathbb{E} \min_{\sigma \in S_n} \sum_{i=1}^{n} \|X_i - Y_{\sigma(i)}\|^p_{d,T},$$

where $X_1, \dots, X_n, Y_1, \dots, Y_n$ are independent and uniform on Ω_T. Attach to the targets independent marks $E_i \sim \mathrm{Exp}(1)$ and put $S = \sum_i E_i$, $W_i = E_i/S$. By Theorem 3.1 and (3.2), applied at fixed targets and then averaged over them (the integrand is measurable by the argument in the proof of Corollary 3.2(ii), since the torus is compact and the cost is continuous),

$$\mathcal{A}^{\mathrm{per}}_{n,d,p,N_d} = \mathbb{E}\,\mathrm{OT}_c\Big(\frac{\mathrm{d}x}{n}, \sum_{i=1}^{n} W_i \delta_{Y_i}\Big), \qquad c(x,y) = \|x - y\|^p_{d,T}.$$

We use the prices $h_i = r \log E_i$ and the potential $\phi(x) = \min_i\{c(x, Y_i) - r \log E_i\}$. Since $\phi(x) + h_i \leq c(x, Y_i)$ for all x and i, weak duality gives

$$\mathrm{OT}_c\Big(\frac{\mathrm{d}x}{n}, \sum_i W_i \delta_{Y_i}\Big) \geq \frac{1}{n}\int_{\Omega_T} \phi(x)\,\mathrm{d}x + r \sum_i W_i \log E_i.$$

Put $T_n(x) = \min_i \|x - Y_i\|^d_{d,T}/E_i$. Since $c(x, Y_i) = (\|x - Y_i\|^d_{d,T})^r$, applying $t^r \geq 1 + r \log t$ with $t = \|x - Y_i\|^d_{d,T}$ gives $c(x, Y_i) - r \log E_i \geq 1 + r \log(\|x - Y_i\|^d_{d,T}/E_i)$, hence $\phi(x) \geq 1 + r \log T_n(x)$. For the target term, $E_i = SW_i$ with $S \sim \mathrm{Gamma}(n, 1)$ independent of $W \sim \mathrm{Dir}(1, \dots, 1)$, and $W_1 \sim \mathrm{Beta}(1, n-1)$; therefore

$$\mathbb{E}\sum_i W_i \log E_i = \mathbb{E}\log S + n\mathbb{E}[W_1 \log W_1] = \psi(n) + \psi(2) - \psi(n+1) = 1 - \gamma_{\mathrm{E}} - \frac{1}{n}.$$

The law of $T_n(x)$ does not depend on x, so

$$\mathcal{A}^{\mathrm{per}}_{n,d,p,N_d} \geq 1 + r\,\mathbb{E}\log T_n(0) + r\Big(1 - \gamma_{\mathrm{E}} - \frac{1}{n}\Big). \tag{5.1}$$

It remains to show that $\liminf_n \mathbb{E}\log T_n(0) \geq -\gamma_{\mathrm{E}}$, which is $\mathbb{E}\log T$ for $T \sim \mathrm{Exp}(1)$. Set $H_n(t) = \int_{\Omega_T} \exp(-\|y\|^d_{d,T}/t)\,\mathrm{d}y$. By independence,

$$\mathbb{P}(T_n(0) > t) = \Big(1 - \mathbb{P}(E_1 > \|Y_1\|^d_{d,T}/t)\Big)^n = \Big(1 - \frac{H_n(t)}{n}\Big)^n.$$

Choose a Borel fundamental domain $D_T \subset \mathbb{R}^d$ of $T\mathbb{Z}^d$ on which $\|y\|_{d,T} = \|y\|_d$ (for instance, the set of y for which $k = 0$ is the first, in a fixed enumeration of $\mathbb{Z}^d$, minimizer of $k \mapsto \|y + Tk\|_d$). Since the unit ball of $\|\cdot\|_d$ has volume one, the volume coordinate $u = \|y\|_d^d$ pushes Lebesgue measure to $\mathrm{d}u$, and

$$H_n(t) = \int_{D_T} \mathrm{e}^{-\|y\|_d^d/t}\,\mathrm{d}y \le \int_{\mathbb{R}^d} \mathrm{e}^{-\|y\|_d^d/t}\,\mathrm{d}y = \int_0^\infty \mathrm{e}^{-u/t}\,\mathrm{d}u = t.$$

Conversely, let $\ell_T = T \min_{k\in\mathbb{Z}^d\setminus\{0\}} \|k\|_d$. If $\|y\|_d < \ell_T/2$, then $\|y + Tk\|_d \ge \ell_T - \|y\|_d > \|y\|_d$ for every $k \ne 0$. Hence the projection of $\{\|y\|_d < \ell_T/2\}$ to Ω_T is injective, $\|y\|_{d,T} = \|y\|_d$ on this set, and $H_n(t) \ge \int_{\{\|y\|_d<\ell_T/2\}} \mathrm{e}^{-\|y\|_d^d/t}\,\mathrm{d}y$. Since $\ell_T \to \infty$, monotone convergence gives $H_n(t) \to t$ for every $t > 0$. Hence $\mathbb{P}(T_n(0) > t) \to \mathrm{e}^{-t}$, that is, $T_n(0) \Rightarrow T$. Moreover $\mathbb{P}(T_n(0) \le t) \le H_n(t) \le t$, so $\mathbb{E}[(\log^- T_n(0))^2] \le \int_0^\infty 2s\,\mathrm{e}^{-s}\,\mathrm{d}s = 2$; the variables $\log^- T_n(0)$ are uniformly integrable and $\mathbb{E}\log^- T_n(0) \to \mathbb{E}\log^- T$. For the positive part, weak convergence and Fatou's lemma give $\liminf_n \mathbb{E}\log^+ T_n(0) \ge \mathbb{E}\log^+ T$. Thus $\liminf_n \mathbb{E}\log T_n(0) \ge \mathbb{E}\log T = -\gamma_{\mathrm{E}}$, and (5.1) yields the claim. □

Comparison of the torus and the cube. Map $[0,1)^d$ onto the unit torus by the quotient map, which sends uniform samples to uniform samples. For all $x, y \in [0,1]^d$,

$$d_{N_d,\mathbb{T}^d}(x,y) = \inf_{k\in\mathbb{Z}^d} N_d(x - y + k) \le N_d(x - y).$$

Hence, for every pair of samples, the optimal periodic matching cost does not exceed the optimal matching cost on the cube with the same norm. The normalization $(nv_d)^r$ is the same in both cases, so

$$\mathcal{A}^{\mathrm{per}}_{n,d,p,N_d} \le \mathcal{A}_{n,d,p,N_d} \qquad \text{for all } n, d, p, N_d. \tag{5.2}$$

Together with Proposition 5.1, this gives $\liminf_n \mathcal{A}_{n,d,p,N_d} \ge 1 + (1 - 2\gamma_{\mathrm{E}})r$ for every d, N_d and $p \ge 1$, which is the lower bound in Theorem 2.3 (ii).

Proof of Theorem 2.4. The lower bound is Proposition 5.1. By (5.2) and Theorem 2.3 (ii), there are numerical constants $c, C > 0$ such that, for all large d,

$$B^+_{d,p_d,N_d} \le \limsup_{n\to\infty} \mathcal{A}_{n,d,p_d,N_d} \le \Gamma(1 + r_d)\Gamma(2 + r_d) + C\exp(-cd + Cp_d).$$

Finally $\Gamma(1 + r)\Gamma(2 + r) = 1 + (1 - 2\gamma_{\mathrm{E}})r + O(r^2)$ as $r \downarrow 0$. Since $p_d = o(d)$ and $r_d \ge 1/d$, the term $\exp(-cd + Cp_d)$ is $o(r_d^2)$, and the expansion is uniform over norm sequences. □

6 Ordinary cubes and arbitrary norms

This section and the next prove Theorem 2.3. The lower bound in part (i) uses the capped mean-field prices as dual variables: in volume coordinates, the costs from a point to the nuclei of a unit Poisson process form the race of Section 4, whatever the norm (Proposition 6.2). The upper bounds are built in four steps. First, sources are allocated to the nuclei of a stationary marked Poisson process on all of $\mathbb{R}^d$ by a local race: capped additive prices at a fixed ratio, and, for the vanishing ratio, the multiplicatively weighted race $\|x - z\|_d^d/E_z$ restricted to a fixed range. Second, the volume of each cell concentrates around the mark of its nucleus; the variance bound $\chi^d L^3$ comes from the sharp Young inequality and holds uniformly over norm balls of a given volume. Third, the allocation is restricted to the cube and each cell is clipped to the mark of its nucleus. The resulting signed residual is exactly centred on every

Borel set, including sets meeting the boundary, because the mean displacement kernel of a reflection-covariant allocation is even. Fourth, the residual is paired directly at its original locations along a hierarchy of norm nets; since no mass is routed through intermediate centres, this costs a factor 24^p and no further factor 2^p per level. The residual mass is of order $\mathrm{e}^{-c_\chi d/2}$, which explains the threshold α_{net} in part (i).

6.1 Setting

Fix d, $p \geq 1$, $r = p/d$ and a norm N_d with unit ball K_d and $v_d = |K_d|$. The normalized norm $\|x\|_d = v_d^{1/d} N_d(x)$ has a unit ball of volume one. We write $B_\rho(x) = \{y : \|y - x\|_d \leq \rho\}$, so that $|B_\rho(x)| = \rho^d$; all radii below are measured in $\|\cdot\|_d$. The cost is

$$c_p(x, y) = \|x - y\|_d^p = \big(v_d N_d(x - y)^d\big)^r.$$

Let $Q_T = [0, T]^d$ and $V = T^d$. Dilating $[0, 1]^d$ by $T = n^{1/d}$ turns (2.1) into $\mathcal{A}_{n,d,p,N_d} = n^{-1}\mathbb{E} \min_\sigma \sum_i c_p(X_i, Y_{\sigma(i)})$ with independent uniform points of $Q_{n^{1/d}}$.

The shape of the norm enters only through the widths $w_i = \sup_{\|u\|_d \leq 1} |u_i|$ and $w_N = \max_i w_i$. If $x \in Q_T$ and $B_\rho(x) \not\subset Q_T$, then some coordinate x_i lies within $w_i\rho$ of $\{0, T\}$; moreover $Q_T + B_\rho \subset \prod_i [-w_i\rho, T + w_i\rho]$. Hence the inner and outer boundary layers satisfy

$$\big|\{x \in Q_T : B_\rho(x) \not\subset Q_T\}\big| + \big|(Q_T + B_\rho) \setminus Q_T\big| = O_{d,N_d,\rho}(T^{d-1}). \tag{6.1}$$

These constants depend on N_d, but they disappear when $T \to \infty$ at fixed d and N_d, which is the order of limits of Section 2.4.

Marked Poisson process and Palm measures. Let Π be a unit-intensity Poisson process on $\mathbb{R}^d$; each nucleus $z \in \Pi$ carries a mark $E_z \sim \mathrm{Exp}(1)$ and a tie mark U_z uniform on $[0, 1]$, all independent. Every allocation below assigns a point x to a nucleus minimizing a score, and equal scores are resolved in favour of the smaller tie mark. This matters for nonsmooth norms such as ℓ^∞, for which the set of points equidistant from two nuclei may have positive volume. Translations and the reflection $x \mapsto -x$ preserve the law of Π and carry the marks along, so all allocations below are translation- and reflection-covariant. For $z \in \mathbb{R}^d$ and $e > 0$ we write $\mathbb{P}^{z,e}$, $\mathbb{E}^{z,e}$ for the law of Π with an added nucleus at z of mark e (and an independent tie mark), and $\mathbb{E}^z = \int_0^\infty \mathbb{E}^{z,e}\, \mathrm{e}^{-e}\, \mathrm{d}e$. The Mecke formula [19] reads

$$\mathbb{E} \sum_{z \in \Pi} g(z, \Pi) = \int_{\mathbb{R}^d} \mathbb{E}^z g(z, \Pi)\, \mathrm{d}z \tag{6.2}$$

for measurable $g \geq 0$, where on the right Π contains the added nucleus. For a Borel set C of finite volume, a point z and $0 \leq v \leq |C|$, the radial clipping of C to volume v (around z) is $C \cap B_\rho(z)$, where ρ is the least radius with $|C \cap B_\rho(z)| \geq v$; since spheres of $\|\cdot\|_d$ are Lebesgue-null, its volume equals v.

Poissonization on the cube. The nuclei in Q_T are the targets. Let $N_V = \#(\Pi \cap Q_T)$, $S = \sum_{z \in \Pi \cap Q_T} E_z$ and $\theta = S/V$. Define

$$\mathcal{P}^{\mathrm{cube}}(V) = \frac{1}{V}\, \mathbb{E}\, \mathrm{OT}_{c_p}\Big(\theta \mathbf{1}_{Q_T}\, \mathrm{d}x, \sum_{z \in \Pi \cap Q_T} E_z \delta_z\Big), \tag{6.3}$$

with $\mathrm{OT} = 0$ when $S = 0$. Nuclei outside Q_T are used only to define auxiliary allocations; they are never targets.

Lemma 6.1. *For every d, $p \geq 1$ and norm N_d,*

$$\liminf_{n\to\infty} \mathcal{A}_{n,d,p,N_d} \geq \liminf_{V\to\infty} \mathcal{P}^{\mathrm{cube}}(V), \qquad \limsup_{n\to\infty} \mathcal{A}_{n,d,p,N_d} \leq \limsup_{V\to\infty} \mathcal{P}^{\mathrm{cube}}(V).$$

Proof. Let $b_m = \mathbb{E}W^p_{p,N_d;[0,1]^d}(\mu_m, \nu_m)$, so that $\mathcal{A}_{n,d,p,N_d} = (nv_d)^r b_n$. Conditionally on $N_V = m \geq 1$ the targets are independent and uniform in Q_T, and $E_z = SW_z$ with $W \sim \mathrm{Dir}(1,\dots,1)$ independent of S and $\mathbb{E}[S \mid N_V = m] = m$. Dilation by T^{-1} gives

$$\mathrm{OT}_{c_p}\Big(\theta \mathbf{1}_{Q_T}\,\mathrm{d}x, \sum_z E_z\delta_z\Big) = v_d^r T^p S\ \mathrm{OT}_{N_d^p}\Big(\mathbf{1}_{[0,1]^d}\,\mathrm{d}x, \sum_z W_z\delta_{z/T}\Big),$$

and Theorem 3.1 identifies the conditional expectation of the last transport cost given $N_V = m$ with b_m. As in the derivation of (3.7),

$$\mathcal{P}^{\mathrm{cube}}(V) = v_d^r\, V^{r-1}\, \mathbb{E}[N_V b_{N_V}].$$

The sequence (b_m) is nonincreasing by Lemma 3.3 and bounded by $\mathrm{diam}_{N_d}([0,1]^d)^p$. Lemma 3.4, applied with the exponent r and the constant v_d^r, gives the two inequalities. The averaging over the random targets is justified by Corollary 3.2(ii) with $\mathsf{X} = Q_T$ and $\mu = \mathbf{1}_{Q_T}\,\mathrm{d}x/V$, since Q_T is compact and c_p is continuous. □

Capped prices. Let $q = 1/r$ and $\lambda > 0$. Let $X_\lambda \in [-\lambda/2, \lambda/2]$ be the capped valuation of Theorem 2.2 (b) at pseudo-dimension q, with survival function $F_\lambda(t) = \mathbb{P}(X_\lambda \geq t)$ and $G_\lambda = -\log F_\lambda$. By (4.8),

$$G_\lambda(t) = \mathbb{E}(t + X_\lambda)_+^q, \qquad |t| \leq \lambda/2. \tag{6.4}$$

The function G_λ is continuous and increasing, with $G_\lambda(-\lambda/2) = 0$. Put

$$\bar{e}_\lambda = G_\lambda(\lambda/2), \qquad \rho_\lambda = \mathrm{e}^{-\bar{e}_\lambda} = \mathbb{P}(X_\lambda = \lambda/2), \qquad h_\lambda(e) = G_\lambda^{-1}(e \wedge \bar{e}_\lambda),$$

so that $h_\lambda(E)$ has the law of X_λ when $E \sim \mathrm{Exp}(1)$, and let $C^{\mathrm{trial}}_{q,\lambda} = \mathbb{E}X_\lambda + \mathbb{E}[E\, h_\lambda(E)]$ be the trial value of Section 4.2. By Proposition 4.4, $C^{\mathrm{trial}}_{q,\lambda} \to \beta_{\mathrm{bip}}(q) = \mathcal{C}(r)$ as $\lambda \to \infty$.

The capped allocation gives each nucleus the price $h_z = h_\lambda(E_z)$ and uses the potential

$$\phi_\lambda(x) = \min\Big\{\frac{\lambda}{2},\ \inf_{z\in\Pi}\big(c_p(x,z) - h_z\big)\Big\}.$$

The cell C_z of z is the set of x with $c_p(x,z) - h_z < \lambda/2$ at which z minimizes $c_p(x,\cdot) - h_\cdot$. Put

$$R = \lambda^{1/p}, \qquad L = R^d = \lambda^{1/r}.$$

Since $h_z \leq \lambda/2$, the condition $c_p(x,z) - h_z < \lambda/2$ forces $\|x - z\|_d < R$. Thus $C_z \subset B_R(z)$, only finitely many nuclei compete at each point, the cells are disjoint, $\phi_\lambda(x) = c_p(x,z) - h_z$ on C_z, and $\phi_\lambda = \lambda/2$ outside $\bigcup_z C_z$.

6.2 Volume-coordinate race and the lower bound

Proposition 6.2. *Let N_d be any norm.*

(i) *For all x and $t > 0$, $|\{y : c_p(x,y) < t\}| = t^{1/r}$. Consequently, for every x, the points $(\|x - z\|_d^d, E_z)$, $z \in \Pi$, form a Poisson process on $(0,\infty)^2$ with intensity $\mathrm{d}u\,\mathrm{e}^{-e}\,\mathrm{d}e$, which is the race of Section 4.*

(ii) *For every $\lambda > 0$ and every x, $\phi_\lambda(x)$ has the law of X_λ. In particular $\mathbb{P}(\phi_\lambda(x) = \lambda/2) = \rho_\lambda$.*

(iii) *For every $\lambda > 0$ and $e > 0$, $\mathbb{E}^{0,e}|C_0| = e \wedge \bar{e}_\lambda$.*

Proof. The set in (i) is the ball $B_{t^{1/p}}(x)$, of volume $t^{d/p} = t^{1/r}$; the second assertion follows from the mapping theorem for Poisson processes. For (ii), let $t \le \lambda/2$. The event $\{\phi_\lambda(x) \ge t\}$ is the event that no nucleus satisfies $c_p(x,z) < t + h_\lambda(E_z)$. By (i) and since $h_\lambda(E)$ has the law of X_λ, its Poisson void probability is

$$\mathbb{P}(\phi_\lambda(x) \ge t) = \exp\Big(-\int_0^\infty (t + h_\lambda(e))_+^q \, \mathrm{e}^{-e} \, \mathrm{d}e\Big) = \exp(-\mathbb{E}(t + X_\lambda)_+^q) = F_\lambda(t)$$

by (6.4). Both variables lie in $[-\lambda/2, \lambda/2]$, so their laws coincide; at $t = \lambda/2$ this gives the atom ρ_λ.

For (iii), let $h = h_\lambda(e)$. Under $\mathbb{P}^{0,e}$ the other nuclei form a copy of Π, and $x \in C_0$ exactly when $c_p(x,0) - h < \lambda/2$ and the infimum over the other nuclei exceeds $c_p(x,0) - h$; equality has probability zero because F_λ is continuous on $[-\lambda/2, \lambda/2)$. By (ii), with $F_\lambda(t) = 0$ for $t > \lambda/2$, and in the volume coordinate $u = \|x\|_d^d$,

$$\mathbb{E}^{0,e}|C_0| = \int_0^\infty F_\lambda(u^r - h) \, \mathrm{d}u = \mathbb{E} \int_0^\infty \mathbf{1}_{\{u^r \le X_\lambda + h\}} \, \mathrm{d}u = \mathbb{E}(h + X_\lambda)_+^q = G_\lambda(h).$$

This is the fixed-point relation $G_\lambda(t) = \int_0^\infty F_\lambda(u^r - t) \, \mathrm{d}u$, that is, (4.8) in volume coordinates. Since $G_\lambda(h_\lambda(e)) = e \wedge \bar{e}_\lambda$, (iii) follows. □

Proposition 6.3. *Let $d \ge 1$, let N_d be any norm and let $p \ge 1$, $r = p/d$. Then*

$$\liminf_{n \to \infty} \mathcal{A}_{n,d,p,N_d} \ge \mathcal{C}(r). \tag{6.5}$$

Proof. By Lemma 6.1 it suffices to bound $\liminf_V \mathcal{P}^{\mathrm{cube}}(V)$ from below. Fix $\lambda > 0$ and use the prices $h_z = h_\lambda(E_z) \in [-\lambda/2, \lambda/2]$ of the targets $z \in \Pi \cap Q_T$ together with the potential

$$\phi(x) = \min\Big\{\frac{\lambda}{2}, \min_{z \in \Pi \cap Q_T} \big(c_p(x,z) - h_z\big)\Big\}, \qquad x \in Q_T.$$

Then $\phi(x) + h_z \le c_p(x,z)$, and weak duality gives

$$\mathrm{OT}_{c_p}\Big(\theta \mathbf{1}_{Q_T} \, \mathrm{d}x, \sum_z E_z \delta_z\Big) \ge \theta \int_{Q_T} \phi(x) \, \mathrm{d}x + \sum_{z \in \Pi \cap Q_T} E_z h_z.$$

By (6.2) the target term has expectation $V\mathbb{E}[E\, h_\lambda(E)]$. Since $|\phi| \le \lambda/2$ and $\mathbb{E}|\theta - 1| \le (\mathrm{Var}\,\theta)^{1/2} = \sqrt{2/V}$, replacing θ by 1 in the source term changes its expectation by at most $\lambda\sqrt{2V}/2$. If $B_R(x) \subset Q_T$, every nucleus with $c_p(x,z) - h_z < \lambda/2$ lies in Q_T, so $\phi(x) = \phi_\lambda(x)$ and $\mathbb{E}\phi(x) = \mathbb{E}X_\lambda$ by Proposition 6.2 (ii). The remaining points of Q_T form the inner layer of (6.1), of volume $O_{d,N_d,\lambda}(T^{d-1})$, on which $|\phi| \le \lambda/2$. Therefore

$$\mathcal{P}^{\mathrm{cube}}(V) \ge \mathbb{E}X_\lambda + \mathbb{E}[E\, h_\lambda(E)] - O_{d,N_d,\lambda}(T^{-1} + V^{-1/2}),$$

and $\liminf_V \mathcal{P}^{\mathrm{cube}}(V) \ge C^{\mathrm{trial}}_{1/r,\lambda}$. Letting $\lambda \to \infty$ and using Proposition 4.4 gives (6.5). □

The limiting bound does not depend on N_d; the norm enters only through the boundary layer, which vanishes as $n \to \infty$ at fixed d.

6.3 Concentration of cells

The following estimates use only that the competing sets are translates of a set of given volume; they hold for arbitrarily eccentric norms.

Lemma 6.4. *For measurable sets $A, D \subset \mathbb{R}^d$ with $|A| = a$ and $|D| = b$ finite,*

$$\int_A \int_A |(x + D) \cap (y + D)| \, \mathrm{d}x \, \mathrm{d}y = \|\mathbf{1}_A * \mathbf{1}_D\|_{L^2}^2 \le \chi^d a^{3/2} b^{3/2}, \qquad \chi = \frac{4}{3\sqrt{3}} < 1. \tag{6.6}$$

Proof. Writing the overlap as $\int \mathbf{1}_D(z - x)\mathbf{1}_D(z - y) \, \mathrm{d}z$, Fubini gives the equality. The sharp Young inequality of Beckner and Brascamp–Lieb [4, 7] with exponents $(4/3, 4/3, 2)$ reads $\|f * g\|_2 \le A_{4/3}^{2d} \|f\|_{4/3} \|g\|_{4/3}$, where $A_s = s^{1/(2s)} (s')^{-1/(2s')}$ and $A_{4/3}^4 = 4/(3\sqrt{3}) = \chi$. We apply it to the two indicators, whose $L^{4/3}$ norms are $a^{3/4}$ and $b^{3/4}$. □

Proposition 6.5. *For every norm N_d, every $d \ge 1$, $\lambda > 0$ and $e > 0$, the cell of a nucleus planted at the origin in the capped allocation satisfies*

$$\mathrm{Var}^{0,e}(|C_0|) \le \varsigma_d := \chi^d L^3, \qquad L = \lambda^{1/r}. \tag{6.7}$$

Proof. Let $h = h_\lambda(e)$. Under $\mathbb{P}^{0,e}$ the cell C_0 lies in $B_R(0)$. Call a probe x eligible if $c_p(x, 0) - h < \lambda/2$; other probes have zero indicator. For eligible x let F_x be the set of marked points (w, e', u) that beat the planted nucleus at x. Then $x \in C_0$ exactly when no other nucleus lies in F_x. Since $h_\lambda(e') \le \lambda/2$, every $(w, e', u) \in \mathsf{F}_x$ satisfies $c_p(x, w) < \lambda$, that is, $w \in B_R(x)$. Let a_x, a_y be the intensities of $\mathsf{F}_x, \mathsf{F}_y$ and $I(x, y)$ that of $\mathsf{F}_x \cap \mathsf{F}_y$. The Poisson covariance is

$$\mathrm{e}^{-a_x - a_y}\big(\mathrm{e}^{I(x,y)} - 1\big) = \mathrm{e}^{-(a_x + a_y - I(x,y))}\big(1 - \mathrm{e}^{-I(x,y)}\big) \le I(x, y) \le |B_R(x) \cap B_R(y)|.$$

Integrating over $B_R(0) \times B_R(0)$ and applying Lemma 6.4 with $A = D = B_R(0)$, both of volume L, gives (6.7). □

For the vanishing ratio we use a multiplicatively weighted race with a fixed range. Fix $L \ge 1$ and $R = L^{1/d}$. For $x \in \mathbb{R}^d$ let $J_R(x)$ be the nucleus minimizing $\|x - z\|_d^d / E_z$ among the nuclei z with $\|x - z\|_d \le R$, and $J_R(x) = \varnothing$ if there is none. Let $C_z^R = \{x : J_R(x) = z\}$, and let J_∞, C_z^∞ be the same objects without the range restriction. Put

$$\Delta_L = (2 + \sqrt{L}) \, \mathrm{e}^{-\sqrt{L}}.$$

Lemma 6.6. *For every norm N_d, every $d \ge 1$ and $L \ge 1$, with $r = p/d$:*

$$\mathbb{P}(J_R(x) = \varnothing) = \mathrm{e}^{-L}, \qquad \sup_{e > 0} \mathrm{Var}^{0,e}(|C_0^R|) \le \chi^d L^3, \tag{6.8}$$

$$\int_0^\infty \big|\mathbb{E}^{0,e}|C_0^R| - e\big| \, \mathrm{e}^{-e} \, \mathrm{d}e \le 2\Delta_L, \qquad \mathbb{E}^0\big(|C_0^R| - E_0\big)_+ \le \chi^{d/2} L^{3/2} + 2\Delta_L, \tag{6.9}$$

$$\mathbb{E}\big[c_p(x, J_R(x)) \mathbf{1}_{\{J_R(x) \ne \varnothing\}}\big] \le \Gamma(1 + r)\Gamma(2 + r) + L^r \Delta_L, \tag{6.10}$$

where E_0 denotes the mark of the planted nucleus.

Proof. The number of nuclei in $B_R(x)$ is Poisson with mean L, which gives the first identity. Under $\mathbb{P}^{0,e}$ the cell C_0^R lies in $B_R(0)$, and a nucleus that beats the planted one at a probe x lies in $B_R(x)$. The covariance argument of Proposition 6.5 and Lemma 6.4 give the variance bound.

By Proposition 6.2 (i) and the factorization of the ratio race in the proof of Proposition 4.2(iii), which holds for every $r > 0$, the unrestricted winning ratio $T_* = \min_z \|x - z\|_d^d / E_z$

is Exp(1), the winning mark E_* is Gamma$(2,1)$ and independent of T_*, and $U_* = \|x - J_\infty(x)\|_d^d = T_* E_*$ satisfies $\mathbb{E}U_*^r = \Gamma(1+r)\Gamma(2+r)$. Moreover

$$\mathbb{P}(U_* > L) \le \mathbb{P}(T_* > \sqrt{L}) + \mathbb{P}(E_* > \sqrt{L}) = \mathrm{e}^{-\sqrt{L}} + (1+\sqrt{L})\mathrm{e}^{-\sqrt{L}} = \Delta_L.$$

If $U_* \le L$, the unrestricted winner lies within range, so $J_R(x) = J_\infty(x)$.

To compare the cells we use the mass-transport principle: for every measurable $g(z,x,\Pi) \ge 0$ that is invariant under joint translation of z, x and Π, one has $\mathbb{E}^0 \int g(0,x,\Pi)\,\mathrm{d}x = \mathbb{E}\sum_{z\in\Pi} g(z,0,\Pi)$, a consequence of (6.2) and stationarity. Take $g(z,x,\Pi) = |\mathbf{1}_{\{J_R(x)=z\}} - \mathbf{1}_{\{J_\infty(x)=z\}}|$. The left side dominates $\mathbb{E}^0||C_0^R| - |C_0^\infty||$. The right side vanishes unless $J_R(0) \ne J_\infty(0)$, in which case it is at most 2; hence it is at most $2\mathbb{P}(U_* > L) \le 2\Delta_L$. The one-point void probability gives $\mathbb{P}^{0,e}(x \in C_0^\infty) = \exp(-\|x\|_d^d/e)$, so $\mathbb{E}^{0,e}|C_0^\infty| = e$, and

$$\int_0^\infty \big|\mathbb{E}^{0,e}|C_0^R| - e\big|\,\mathrm{e}^{-e}\,\mathrm{d}e \le \mathbb{E}^0||C_0^R| - |C_0^\infty|| \le 2\Delta_L.$$

This is the first inequality in (6.9). For the second, write $(|C_0^R| - e)_+ \le \big||C_0^R| - \mathbb{E}^{0,e}|C_0^R|\big| + \big|\mathbb{E}^{0,e}|C_0^R| - e\big|$ and use the conditional Cauchy–Schwarz inequality with (6.8).

Finally, on $\{U_* \le L\}$ the finite-range cost equals U_*^r, and on its complement every accepted finite-range edge has cost at most L^r. This gives (6.10). □

Lemma 6.7. *Let $R > 0$ and consider either*

(i) *the capped allocation, with $R = \lambda^{1/p}$, or*

(ii) *the finite-range race J_R.*

For every nucleus z, including nuclei added under a Palm measure, the cell of z lies in $B_R(z)$, and its indicator is a measurable function of the marks of z and of the marked nuclei in $B_{2R}(z)$. Consequently the radial clipping of the cell to a volume determined by its own volume and the mark of z is determined by the same data, and it is translation- and reflection-covariant.

Proof. In case (i), if z wins at x then $c_p(x,z) - h_z < \lambda/2$, so $c_p(x,z) < \lambda$ and $\|x - z\|_d < R$. A nucleus w can beat z at such an x only if $c_p(x,w) - h_w < \lambda/2$, which again gives $\|x - w\|_d < R$ and therefore $\|w - z\|_d < 2R$. In case (ii) both the winner and every admissible competitor lie within R of x, which gives the same bounds. The cell determines its volume and hence the clipping radius. Norm balls are symmetric and the tie marks travel with the nuclei, so translations and the reflection act covariantly on all these objects. □

6.4 Restricted allocations on the cube

The next lemma removes the boundary bias of a stationary allocation restricted to Q_T, and controls the global factor θ by a single total-variation estimate.

Lemma 6.8. *Let $R > 0$ and $L = R^d$. Suppose a translation- and reflection-covariant rule assigns to each nucleus $z \in \Pi$ a Borel set $\widetilde{C}_z \subset B_R(z)$, measurable with respect to the marks of z and the marked nuclei in $B_{2R}(z)$, such that the sets $\widetilde{C}_z$ are pairwise disjoint and $|\widetilde{C}_z| \le E_z$. Define*

$$\sigma_1 = \mathbf{1}_{Q_T \setminus \bigcup_{z\in\Pi\cap Q_T} \widetilde{C}_z}\,\mathrm{d}x, \qquad \tau_1 = \sum_{z\in\Pi\cap Q_T} (E_z - |\widetilde{C}_z \cap Q_T|)\delta_z, \qquad \eta_1 = \sigma_1 - \tau_1.$$

Then, for every Borel set $A \subset Q_T$,

$$\mathbb{E}\eta_1(A) = 0, \qquad \operatorname{Var}\eta_1(A) \le 8^d(1+L)|A| \quad (d \ge 5). \tag{6.11}$$

Put $D_z = \widetilde{C}_z \cap Q_T$ for $z \in \Pi \cap Q_T$. For $\theta \geq 0$ set

$$\sigma_\theta = \Big(\theta \mathbf{1}_{Q_T} - \min\{\theta, 1\} \mathbf{1}_{\bigcup_z D_z}\Big)\,\mathrm{d}x, \qquad \tau_\theta = \sum_{z \in \Pi \cap Q_T} (E_z - \min\{\theta, 1\}|D_z|)\delta_z, \qquad \eta_\theta = \sigma_\theta - \tau_\theta;$$

for $\theta = 1$ these are the measures defined above. With $\theta = S/V$, the measures $\sigma_\theta, \tau_\theta$ are nonnegative and have equal total mass, and

$$\|\sigma_\theta - \sigma_1\|_{\mathrm{TV}} + \|\tau_\theta - \tau_1\|_{\mathrm{TV}} \leq V|\theta - 1|, \qquad \mathbb{E}\|\eta_\theta - \eta_1\|_{\mathrm{TV}} \leq \sqrt{2V}. \tag{6.12}$$

Proof. We first prove exact centring. Let $k(u) = \mathbb{P}^0(u \in \widetilde{C}_0)$, where $\mathbb{P}^0$ includes the planted mark. By (6.2) and stationarity, the mean measure of the retained pairs $\{(x, z) : x \in \widetilde{C}_z\}$ is $k(x - z)\,\mathrm{d}x\,\mathrm{d}z$. Reflection covariance gives $k(u) = k(-u)$, for every norm. Since the sets are disjoint, for $x \in Q_T$,

$$\frac{\mathbb{E}\sigma_1(\,\mathrm{d}x)}{\mathrm{d}x} = 1 - \int_{Q_T} k(x - z)\,\mathrm{d}z, \qquad \frac{\mathbb{E}\tau_1(\,\mathrm{d}x)}{\mathrm{d}x} = 1 - \int_{Q_T} k(y - x)\,\mathrm{d}y.$$

By evenness the two densities agree, so $\mathbb{E}\eta_1(A) = 0$ for every Borel $A \subset Q_T$, including sets meeting ∂Q_T. This is a statement about means; pathwise, $\eta_1(Q_T) = V - S$.

We now bound the variance. Write

$$f(x) = \mathbf{1}_{\{x \notin \bigcup_{z \in \Pi \cap Q_T} \widetilde{C}_z\}}, \quad \xi_z = E_z - |\widetilde{C}_z \cap Q_T|, \quad \mathcal{X}_A = \int_A f(x)\,\mathrm{d}x, \quad \mathcal{Y}_A = \sum_{z \in \Pi \cap A} \xi_z,$$

so that $\eta_1(A) = \mathcal{X}_A - \mathcal{Y}_A$. Under $\mathbb{E}^z$ and $\mathbb{E}^{z,w}$ the quantities f and ξ are recomputed in the configuration with the added nuclei. Since $0 \leq f \leq 1$ and $0 \leq \xi_z \leq E_z$, both $\mathcal{X}_A$ and $\mathcal{Y}_A$ are square integrable, and the Mecke formula of first and second order gives

$$\mathrm{Var}(\mathcal{X}_A) = \int_{A \times A} \mathrm{Cov}(f(x), f(y))\,\mathrm{d}x\,\mathrm{d}y, \tag{6.13}$$

$$\mathrm{Var}(\mathcal{Y}_A) = \int_A \mathbb{E}^z[\xi_z^2]\,\mathrm{d}z + \int_{A \times A} (\mathbb{E}^{z,w}[\xi_z \xi_w] - \mathbb{E}^z[\xi_z]\mathbb{E}^w[\xi_w])\,\mathrm{d}z\,\mathrm{d}w, \tag{6.14}$$

$$\mathrm{Cov}(\mathcal{X}_A, \mathcal{Y}_A) = \int_{A \times A} (\mathbb{E}^z[f(x)\xi_z] - \mathbb{E}f(x)\,\mathbb{E}^z[\xi_z])\,\mathrm{d}x\,\mathrm{d}z. \tag{6.15}$$

A nucleus whose set covers x lies in $B_R(x)$, and its set is determined by the nuclei in its own $2R$-ball; hence $f(x)$ is measurable with respect to the marked nuclei in $B_{3R}(x)$, while ξ_z depends only on the marks of z and the nuclei in $B_{2R}(z)$. If $\|z - w\|_d > 4R$, these neighbourhoods are disjoint, neither insertion affects the other variable, and the integrand in (6.14) vanishes. For $\|x - z\|_d > 5R$ the same holds in (6.15), and for $\|x - y\|_d > 6R$ the variables $f(x)$ and $f(y)$ are independent. At closer separations, $|\,\mathrm{Cov}(f(x), f(y))| \leq 1$, and the integrands in the last two displays are at most 2 in absolute value, since the added marks are independent Exp(1) variables; on the diagonal $\mathbb{E}^z\xi_z^2 \leq 2$. Since $|B_\rho| = \rho^d$ and $R^d = L$,

$$\mathrm{Var}(\mathcal{X}_A) \leq 6^d L|A|, \qquad \mathrm{Var}(\mathcal{Y}_A) \leq (2 + 2 \cdot 4^d L)|A|, \qquad |\,\mathrm{Cov}(\mathcal{X}_A, \mathcal{Y}_A)| \leq 2 \cdot 5^d L|A|.$$

Hence $\mathrm{Var}\,\eta_1(A) \leq (6^d L + 2 + 2 \cdot 4^d L + 4 \cdot 5^d L)|A| \leq 8^d(1 + L)|A|$ for $d \geq 5$. No regularity of ∂A is used: the restriction to Q_T only multiplies the integrands by deterministic indicators.

It remains to treat the factor θ. The sets D_z are disjoint subsets of Q_T with $|D_z| \leq E_z$, so $\sigma_\theta \geq (\theta - \min\{\theta, 1\})\mathbf{1}_{Q_T}\,\mathrm{d}x \geq 0$ and $\tau_\theta \geq 0$, and both measures have total mass $\theta V - \min\{\theta, 1\}\sum_z |D_z|$ because $\theta V = S$. If $\theta \geq 1$, then $\sigma_\theta - \sigma_1 = (\theta - 1)\mathbf{1}_{Q_T}\,\mathrm{d}x$ and $\tau_\theta = \tau_1$. If $\theta < 1$, then $\sigma_1 - \sigma_\theta = (1 - \theta)\mathbf{1}_{Q_T \setminus \bigcup_z D_z}\,\mathrm{d}x$ and $\tau_\theta - \tau_1 = (1 - \theta)\sum_z |D_z|\delta_z$, of total variation $(1 - \theta)(|Q_T \setminus \bigcup_z D_z| + \sum_z |D_z|) = (1 - \theta)V$. This proves the first bound in (6.12). Finally $\mathbb{E}S = V$ and $\mathrm{Var}\,S = 2V$, so $\mathbb{E}|\theta - 1| \leq \sqrt{2/V}$. □

6.5 Norm-net completion

The residual measures of Lemma 6.8 are transported by a hierarchy adapted to the norm. Coordinate dyadic cubes cannot be used with norm-independent constants, since the $\|\cdot\|_d$-diameter of a coordinate cube of unit volume is unbounded over norms; maximal nets of $\|\cdot\|_d$ have the required packing properties for every norm.

Lemma 6.9. *Let $d \ge 3$, let $\|\cdot\|_d$ be a norm whose unit ball has volume one, let $1 \le p < d/2$, $r = p/d$, and $s > 0$, $s_j = 2^j s$. For every T large enough, depending on d, the norm and s, the cube Q_T admits nested finite Borel partitions $\mathcal{Q}_0, \dots, \mathcal{Q}_{j_{\max}}$ such that*

(a) *every $Q \in \mathcal{Q}_j$ has $\|\cdot\|_d$-diameter at most $4s_j$;*

(b) *if $s_j \le T/(dw_N)$, then $\#\mathcal{Q}_j \le 3^d V/s_j^d$;*

(c) *the number of the remaining levels and the number of their cells are bounded in terms of d and the norm; $\mathcal{Q}_{j_{\max}} = \{Q_T\}$ and $s_{j_{\max}} = O_{d,N_d}(T)$.*

Let σ, τ be random nonnegative measures on Q_T with equal total mass, and $\eta = \sigma - \tau$. Suppose a random signed measure η^0 satisfies

$$\mathbb{E}\eta^0(A) = 0, \qquad \operatorname{Var}\eta^0(A) \le \Upsilon|A| \qquad \text{for every Borel } A \subset Q_T, \tag{6.16}$$

and

$$\mathbb{E}\sigma(Q_T) \le (m + o_T(1))V, \qquad \mathbb{E}\|\eta - \eta^0\|_{\mathrm{TV}} = O(\sqrt{V}),$$

where $\Upsilon \ge 1$ and $0 < m \le 1$ do not depend on T. Then, for $s = 3(\Upsilon/m^2)^{1/d}$,

$$\limsup_{T\to\infty} \frac{1}{V}\mathbb{E}\,\mathrm{OT}_{c_p}(\sigma,\tau) \le \Big(12^p + \frac{24^p}{2(1-2^{p-d/2})}\Big)\Upsilon^r m^{1-2r}. \tag{6.17}$$

In particular, if $p \le d/2 - 1$, the right side is at most $2 \cdot 24^p \Upsilon^r m^{1-2r}$. No constant on the right depends on the norm.

Proof. Let $Z_0 \subset Q_T$ be a maximal s-separated set, and inductively let $Z_j \subset Z_{j-1}$ be maximal s_j-separated within Z_{j-1}. By maximality every point of Q_T lies within s of Z_0, and every $z \in Z_{j-1}$ has a parent in Z_j at distance at most s_j. Let $\mathcal{Q}_0$ be the partition of Q_T into closest-point cells of Z_0, with a Borel tie-break, and let the cells of $\mathcal{Q}_j$ be the unions of the cells of $\mathcal{Q}_0$ sharing the same ancestor in Z_j. A point of a level-j cell lies within $s + \sum_{i=1}^{j} s_i < 2s_j$ of that ancestor, which proves (a).

The balls of radius $s_j/2$ centred at the points of Z_j are disjoint and lie in $Q_T + B_{s_j/2} \subset \prod_i [-w_i s_j/2, T + w_i s_j/2]$. Hence

$$\#Z_j\,(s_j/2)^d \le V \prod_i \Big(1 + \frac{w_i s_j}{T}\Big).$$

If $s_j \le T/(dw_N)$, the product is at most $(1+1/d)^d \le \mathrm{e}$, so $\#Z_j \le \mathrm{e}\,2^d V/s_j^d \le 3^d V/s_j^d$ because $d \ge 3$; this is (b). At the first level with $s_j > T/(dw_N)$, we have $\#Z_j \le \#Z_{j-1} \le 3^d(2dw_N)^d$, independently of T, and nestedness bounds all later counts by the same constant. Once s_j exceeds the $\|\cdot\|_d$-diameter of Q_T, which is $O_{d,N_d}(T)$, the set Z_j is a single point; this gives (c) with $O_{d,N_d}(1)$ remaining levels.

We pair the positive and negative parts of η at their original locations: first within the cells of $\mathcal{Q}_0$, then, at each level $j+1$, the mass left unpaired in the children of each cell of $\mathcal{Q}_{j+1}$. Mass is never moved through intermediate centres. The unpaired mass in a cell Q after its own stage is $|\eta(Q)|$, so the amount paired at level $j+1$ is at most $\frac{1}{2}\sum_{Q\in\mathcal{Q}_j} |\eta(Q)|$,

and by (a) each paired unit travels at most $4s_{j+1}$. With $M = \sigma(Q_T) = \tau(Q_T)$, this gives the standard multiscale bound over nested partitions (see [32, Proposition 1], and [5, 12] for earlier versions)

$$\mathrm{OT}_{c_p}(\sigma,\tau) \le (4s)^p M + \frac{1}{2}\sum_{j=0}^{j_{\max}-1} (4s_{j+1})^p \sum_{Q\in\mathcal{Q}_j} |\eta(Q)|. \tag{6.18}$$

For each partition, $\sum_Q |\eta(Q)| \le \sum_Q |\eta^0(Q)| + \|\eta - \eta^0\|_{\mathrm{TV}}$. Since $\sum_{j<j_{\max}} s_{j+1}^p \le 2s_{j_{\max}}^p = O_{d,N,p}(T^p)$, the total-variation errors contribute $O_{d,N,p}(T^p\sqrt{V})/V = O(V^{r-1/2})$ per unit volume, which vanishes because $r < 1/2$.

On the levels of (b), $\mathbb{E}|\eta^0(Q)| \le (\Upsilon|Q|)^{1/2}$ and the Cauchy–Schwarz inequality give

$$\mathbb{E}\sum_{Q\in\mathcal{Q}_j} |\eta^0(Q)| \le \bigl(\#\mathcal{Q}_j\,\Upsilon V\bigr)^{1/2} \le V\sqrt{\Upsilon}\,3^{d/2}s_j^{-d/2}.$$

The remaining levels have $O_{d,N_d}(1)$ cells and $s_{j+1} = O_{d,N_d}(T)$, so they contribute $O_{d,N,p}(T^p\sqrt{\Upsilon V})/V = O(V^{r-1/2})$ per unit volume. Summing the geometric series over $j \ge 0$,

$$\limsup_{T\to\infty} \frac{1}{V}\mathbb{E}\,\mathrm{OT}_{c_p}(\sigma,\tau) \le (4s)^p m + \frac{8^p}{2(1-2^{p-d/2})}\sqrt{\Upsilon}\,3^{d/2}s^{p-d/2}.$$

With $s = 3(\Upsilon/m^2)^{1/d}$ the first term equals $12^p\Upsilon^r m^{1-2r}$ and the second equals $24^p\Upsilon^r m^{1-2r}/(2(1-2^{p-d/2}))$, which proves (6.17). If $p \le d/2-1$, then $2^{p-d/2} \le 1/2$ and the bracket in (6.17) is at most $2\cdot 24^p$. □

7 Upper bounds on the cube

We now combine the allocations of Section 6 with the completion lemma. Proposition 7.1 reduces the upper bound to three quantities of a stationary allocation of $\mathbb{R}^d$: the mean cost at a point, the probability that a point is not allocated, and the mean excess of a cell over the mark of its nucleus. The nuclei outside Q_T only shape the cells.

7.1 Local allocations

An allocation rule assigns to each marked configuration and each $x \in \mathbb{R}^d$ a nucleus $J(x)$, or $J(x) = \varnothing$, measurably in the configuration and in x. Its cells are $C_z = \{x : J(x) = z\}$. Under $\mathbb{E}^0 = \int_0^\infty \mathbb{E}^{0,e}\,\mathrm{e}^{-e}\,\mathrm{d}e$ we write C_0 for the cell of the planted nucleus and E_0 for its mark.

Proposition 7.1. *Let $d \ge 5$, $1 \le p \le d/2-1$, $r = p/d$, $R > 0$, $L = R^d$, and let N_d be a norm. Let J be a translation- and reflection-covariant allocation rule such that $\|x - J(x)\|_d \le R$ whenever $J(x) \ne \varnothing$, and such that for every nucleus z, including nuclei added under a Palm measure, the cell C_z is determined by the marks of z and the marked nuclei in $B_{2R}(z)$. Put*

$$g(x) = c_p(x, J(x))\mathbf{1}_{\{J(x)\ne\varnothing\}}, \qquad \kappa = \mathbb{E}g(0), \qquad m = \mathbb{P}(J(0) = \varnothing) + \mathbb{E}^0\bigl(|C_0| - E_0\bigr)_+.$$

Then, with $\Upsilon = 8^d(1+L)$,

$$\limsup_{n\to\infty} \mathcal{A}_{n,d,p,N_d} \le \kappa + 2\cdot 24^p\Upsilon^r m^{1-2r}. \tag{7.1}$$

Proof. By Lemma 6.1 it suffices to bound $\limsup_V \mathcal{P}^{\mathrm{cube}}(V)$. We keep the part of each cell inside the cube up to the mark of its nucleus, bound its cost by g, and complete the residual with Lemma 6.9.

Let $\widetilde{C}_z$ be the radial clipping of C_z around z to volume $\min\{|C_z|, E_z\}$. Since $C_z \subset B_R(z)$ and the clipping radius is determined by C_z and E_z, the sets $\widetilde{C}_z$ are pairwise disjoint subsets of $B_R(z)$ with $|\widetilde{C}_z| \le E_z$, determined by the marks of z and the marked nuclei in $B_{2R}(z)$. Norm balls are symmetric, so they inherit the covariance of J. Thus Lemma 6.8 applies; let D_z, σ_θ, τ_θ, η_θ be as there, with $\theta = S/V$. The plan $\pi_{\mathrm{keep}} = \min\{\theta, 1\} \sum_{z\in\Pi\cap Q_T} \mathbf{1}_{D_z}(x)\,\mathrm{d}x\,\delta_z$ has marginals $\theta\mathbf{1}_{Q_T}\,\mathrm{d}x - \sigma_\theta$ and $\sum_z E_z\delta_z - \tau_\theta$. Gluing it with an optimal coupling of σ_θ and τ_θ gives

$$\mathrm{OT}_{c_p}\Big(\theta\mathbf{1}_{Q_T}\,\mathrm{d}x, \sum_{z\in\Pi\cap Q_T} E_z\delta_z\Big) \le \int c_p\,\mathrm{d}\pi_{\mathrm{keep}} + \mathrm{OT}_{c_p}(\sigma_\theta, \tau_\theta). \tag{7.2}$$

Every $x \in D_z$ satisfies $J(x) = z$, and the sets D_z are disjoint subsets of Q_T, so $\int c_p\,\mathrm{d}\pi_{\mathrm{keep}} \le \int_{Q_T} g(x)\,\mathrm{d}x$. By stationarity this has expectation κV.

We now bound the residual mass. A point of $Q_T \setminus \bigcup_z D_z$ either is not allocated, or is allocated to a nucleus outside Q_T, or lies in $C_z \setminus \widetilde{C}_z$ for some $z \in \Pi \cap Q_T$, a set of volume $(|C_z| - E_z)_+$. Hence

$$\sigma_1(Q_T) \le \big|\{x \in Q_T : J(x) = \varnothing\}\big| + \big|\{x \in Q_T : J(x) \notin Q_T \cup \{\varnothing\}\}\big| + \sum_{z\in\Pi\cap Q_T} \big(|C_z| - E_z\big)_+.$$

The first term has expectation $V\mathbb{P}(J(0) = \varnothing)$ by stationarity. Since $\|x - J(x)\|_d \le R$, the second set lies in the inner layer of (6.1). By (6.2) and stationarity the third term has expectation $V\mathbb{E}^0(|C_0| - E_0)_+$. By (6.12), $\sigma_\theta(Q_T) \le \sigma_1(Q_T) + V|\theta - 1|$, and $\mathbb{E}|\theta - 1| \le \sqrt{2/V}$. Therefore

$$\mathbb{E}\sigma_\theta(Q_T) \le (m + O_{d,N_d,R}(T^{-1} + V^{-1/2}))V.$$

Moreover $\sigma_\theta(Q_T) \le \theta V = S$, so $\mathbb{E}\sigma_\theta(Q_T) \le V$.

By Lemma 6.8, $\eta^0 = \eta_1$ satisfies (6.16) with $\Upsilon \ge 1$, and $\mathbb{E}\|\eta_\theta - \eta_1\|_{\mathrm{TV}} \le \sqrt{2V}$. For every $m' \in (0, 1]$ with $m' \ge \min\{m, 1\}$, the two mass bounds give $\mathbb{E}\sigma_\theta(Q_T) \le (m' + o_T(1))V$, and Lemma 6.9, with $p \le d/2 - 1$, gives

$$\limsup_{T\to\infty} \frac{1}{V}\mathbb{E}\,\mathrm{OT}_{c_p}(\sigma_\theta, \tau_\theta) \le 2 \cdot 24^p \Upsilon^r m'^{1-2r}.$$

Letting $m' \downarrow \min\{m, 1\}$ and using $1 - 2r > 0$, the right side becomes at most $2 \cdot 24^p\Upsilon^r m^{1-2r}$. Taking expectations in (7.2) and dividing by V gives (7.1). □

7.2 Capped prices

Proposition 7.2. *Let $d \ge 5$, $1 \le p \le d/2 - 1$, $r = p/d$, $q = 1/r$ and $\lambda > 0$. Let L, ρ_λ and $\varsigma_d = \chi^d L^3$ be as in Section 6.1 and Proposition 6.5, and $m = \rho_\lambda + \sqrt{\varsigma_d}$. For every norm N_d,*

$$\limsup_{n\to\infty} \mathcal{A}_{n,d,p,N_d} \le E_{q,\lambda} + 2 \cdot 24^p\big[8^d(1 + L)\big]^r m^{1-2r}. \tag{7.3}$$

Proof. Let $J(x)$ be the nucleus whose capped cell C_z contains x, and $J(x) = \varnothing$ if $\phi_\lambda(x) = \lambda/2$. As noted in Section 6.1, $\|x - J(x)\|_d < R = \lambda^{1/p}$, and $R^d = L$. By Lemma 6.7 (i), J satisfies the hypotheses of Proposition 7.1. We compute the three quantities.

By Proposition 6.2 (ii), $\mathbb{P}(J(0) = \varnothing) = \mathbb{P}(\phi_\lambda(0) = \lambda/2) = \rho_\lambda$.

By Proposition 6.2 (iii), $\mathbb{E}^{0,e}|C_0| = e \wedge \bar{e}_\lambda \le e$. Hence, by the Cauchy–Schwarz inequality and Proposition 6.5, $\mathbb{E}^{0,e}(|C_0| - e)_+ \le \mathbb{E}^{0,e}||C_0| - \mathbb{E}^{0,e}|C_0|| \le \sqrt{\varsigma_d}$ for every $e > 0$, and $\mathbb{E}^0(|C_0| - E_0)_+ \le \sqrt{\varsigma_d}$.

By Proposition 6.2 (i), the points $(U_z, E_z) = (\|z\|_d^d, E_z)$, $z \in \Pi$, form the Poisson process of Section 4.2, and $c_p(0, z) = U_z^r$. Thus $J(0)$ is the atom minimizing $U_z^r - h_\lambda(E_z)$ when this minimum is smaller than $\lambda/2$, and $J(0) = \varnothing$ otherwise; since ties have probability zero, this is the capped selector ι of the proof of Proposition 4.2. By (4.14), $\mathbb{E}g(0) = \mathbb{E}[U_\iota^r \mathbf{1}_{\{\iota\ne\varnothing\}}] = E_{q,\lambda}$.

Proposition 7.1, together with the monotonicity of $m \mapsto m^{1-2r}$, gives (7.3). □

By Proposition 4.2 (ii)–(iii) and Proposition 4.4, $E_{q,\lambda} \le C^{\mathrm{trial}}_{q,\lambda} \le \mathcal{C}(r)$ for every $\lambda > 0$.

7.3 The vanishing ratio

For the logarithmic endpoint we replace the capped prices by the multiplicative race of Lemma 6.6. A fixed physical range, rather than a finite penalty on additive prices, keeps the cells local.

Proposition 7.3. *Let $d \ge 5$, $1 \le p \le d/2 - 1$, $r = p/d$, $L = d^4$, and*

$$m_d = \mathrm{e}^{-L} + \chi^{d/2} L^{3/2} + 2\Delta_L, \qquad \Upsilon_d = 8^d(1 + L).$$

For every norm N_d,

$$\limsup_{n\to\infty} \mathcal{A}_{n,d,p,N_d} \le \Gamma(1 + r)\Gamma(2 + r) + L^r \Delta_L + 2 \cdot 24^p \Upsilon_d^r m_d^{1-2r}. \tag{7.4}$$

Proof. We apply Proposition 7.1 to the finite-range race J_R, $R = L^{1/d}$. By definition $\|x - J_R(x)\|_d \le R$, and Lemma 6.7 (ii) gives the locality of the cells. By (6.10), $\kappa \le \Gamma(1 + r)\Gamma(2 + r) + L^r\Delta_L$. By (6.8) and (6.9), $\mathbb{P}(J_R(0) = \varnothing) + \mathbb{E}^0(|C_0^R| - E_0)_+ \le m_d$. □

7.4 Proof of Theorem 2.3

Proof of Theorem 2.3. We first prove (i). Fix $0 < \alpha < \alpha_{\mathrm{net}}$, let $p = \alpha d$ and $q = 1/\alpha$. Since $\alpha < \alpha_{\mathrm{net}}$, we have $\alpha \log 192 < (c_\chi/2)(1 - 2\alpha)$, so we can choose $\mu < \mu_1 < c_\chi/2$ with $\alpha \log 192 < \mu(1 - 2\alpha)$. Let c_q be the constant of Lemma 4.3, choose $\bar{A} > \mu_1/c_q$, and put $L_d = \bar{A}d$, $\lambda_d = L_d^\alpha$. By Lemma 4.3, $\bar{e}_{\lambda_d} \ge c_q \lambda_d^q = c_q \bar{A} d$ for large d, so $\rho_{\lambda_d} \le \mathrm{e}^{-\mu_1 d}$. Since $\chi^{d/2} = \mathrm{e}^{-c_\chi d/2}$, we have $\sqrt{\varsigma_d} = \mathrm{e}^{-c_\chi d/2}(\bar{A}d)^{3/2}$, and therefore $m_d = \rho_{\lambda_d} + \sqrt{\varsigma_d} \le \mathrm{e}^{-\mu d}$ for all large d. If $d \ge 5$, $\alpha d \ge 1$ and $d \ge 2/(1 - 2\alpha)$, then $1 \le p \le d/2 - 1$, and Proposition 7.2 bounds the residual term in (7.3) by

$$2 \cdot 192^{\alpha d}(1 + \bar{A}d)^\alpha \mathrm{e}^{-\mu(1-2\alpha)d},$$

which tends to zero by the choice of μ. Also $E_{q,\lambda_d} \to \mathcal{C}(\alpha)$ by Proposition 4.4. Hence $\limsup_d \limsup_n \mathcal{A}_{n,d,p,N_d} \le \mathcal{C}(\alpha)$. Proposition 6.3 gives $\liminf_n \mathcal{A}_{n,d,p,N_d} \ge \mathcal{C}(\alpha)$ for every large d, and (i) follows.

We now prove (ii). The lower bound was proved in Section 5, from Proposition 5.1 and (5.2). For the upper bound, since $p_d = o(d)$, we have $p_d \le d/2 - 1$ for large d, and we apply Proposition 7.3 with $L = d^4$. For large d,

$$m_d = \mathrm{e}^{-d^4} + \mathrm{e}^{-c_\chi d/2} d^6 + 2(2 + d^2)\mathrm{e}^{-d^2} \le 2d^6 \mathrm{e}^{-c_\chi d/2}, \qquad \Upsilon_d^r = 8^p(1 + d^4)^r.$$

Since $r \le 1/2$, the residual term in (7.4) is at most

$$2 \cdot 24^p 8^p (1 + d^4)^{1/2}\, 2d^6\, \mathrm{e}^{-(c_\chi/2)(1-2r)d} \le 8d^8 \exp\bigl(-(c_\chi/2)d + (\log 192 + c_\chi)p\bigr),$$

and $L^r\Delta_L \le d^2(2 + d^2)\mathrm{e}^{-d^2}$ is superexponentially small. Hence, for any $c < c_\chi/2$ and $C \ge \log 192 + c_\chi$, the upper bound in (7.4) is at most $\Gamma(1 + r_d)\Gamma(2 + r_d) + C\mathrm{e}^{-cd+Cp_d}$ for all large d. The expansion follows from $\Gamma(1 + r)\Gamma(2 + r) = 1 + (1 - 2\gamma_{\mathrm{E}})r + O(r^2)$ and $\mathrm{e}^{-cd+Cp_d} = o(r_d^2)$, as in the proof of Theorem 2.4. □

7.5 Almost-sure costs

The paper [29] studies the matching problem on the ordinary cube $[0,1]^d$ for an arbitrary norm N on $\mathbb{R}^d$, normalizes the cost by $\varrho_N = |\{N \le 1\}|^{-1/d}$, and raises questions about the sharp high-dimensional behaviour of the normalized cost, for $p = 1$ and for higher powers of the distance. Write $M_{n,p}^N = \min_\sigma \sum_{i=1}^n N(X_i - Y_{\sigma(i)})^p$ for two independent uniform n-samples of $[0,1]^d$; then $M_{n,p}^N/(n^{1-r}\varrho_N^p) = (nv_d)^r W_{p,N;[0,1]^d}^p(\mu_n,\nu_n)$, so the normalization (2.1) is that of [29].

Corollary 7.4. *Let N_d be a norm on $\mathbb{R}^d$ and $1 \le p < d/2$. Let $\mathcal{A}_{d,N_d,p}^{\mathrm{as},-}$ and $\mathcal{A}_{d,N_d,p}^{\mathrm{as},+}$ be the almost-sure* lim inf *and* lim sup *of $M_{n,p}^{N_d}/(n^{1-r}\varrho_{N_d}^p)$ as $n \to \infty$. They are deterministic and equal to $\liminf_n \mathcal{A}_{n,d,p,N_d}$ and $\limsup_n \mathcal{A}_{n,d,p,N_d}$. Consequently, for every sequence of norms N_d (in particular for the Euclidean norm),*

$$\mathcal{A}_{d,N_d,1}^{\mathrm{as},\pm} = 1 + \frac{1-2\gamma_{\mathrm{E}}}{d} + O(d^{-2}), \qquad \lim_{d\to\infty} d\,|\mathcal{A}_{d,N_d,1}^{\mathrm{as},\pm} - 1| = 2\gamma_{\mathrm{E}} - 1,$$

and, for $1 \le p_d = o(d)$,

$$\mathcal{A}_{d,N_d,p_d}^{\mathrm{as},+} \le 1 + (1-2\gamma_{\mathrm{E}})\frac{p_d}{d} + O((p_d/d)^2).$$

The O-terms are uniform over norm sequences; the convergence in n is not. When the fixed-d limit defining Talagrand's constant $a_d(N_d)$ exists, these statements apply to $a_d(N_d)/\varrho_{N_d}$; the existence of this limit is not derived here.

Proof. Changing one of the $2n$ sample points changes $M_{n,p}^{N_d}$ by at most $D = \mathrm{diam}_{N_d}([0,1]^d)^p$. At fixed (d, N_d, p) the bounded differences inequality [22] gives, for every $\varepsilon > 0$,

$$\mathbb{P}\Big(|M_{n,p}^{N_d} - \mathbb{E}M_{n,p}^{N_d}| > \varepsilon n^{1-r}\varrho_{N_d}^p\Big) \le 2\exp(-c_{d,N_d,p,\varepsilon}\, n^{1-2r}).$$

This is summable in n because $r < 1/2$, and the Borel–Cantelli lemma identifies the almost-sure lim inf and lim sup with those of the expected normalized cost. The asymptotics then follow from Theorem 2.3 (ii). For $p = 1$, $r_d = 1/d$, and both bounds in Theorem 2.3 (ii) equal $1 + (1-2\gamma_{\mathrm{E}})/d + O(d^{-2})$; since $1 - 2\gamma_{\mathrm{E}} < 0$, this gives the limit of $d|\mathcal{A}_{d,N_d,1}^{\mathrm{as},\pm} - 1|$. □

8 Cavity cells, exact cores and initial concentration

Sections 8–10 construct the local part of the transport plan used for the upper bound in Theorem 2.5. In this section the capped dual potential partitions the torus into cavity cells. Each cell keeps as much of its volume as the mark of its site allows. This splits source and target into a core part, transported inside the cells, and small exceptional parts. The main result, Proposition 8.1, gives exponentially small means and Poisson-type concentration of the exceptional parts at every dyadic moment order. Concentration at an order $2s > 2\alpha$ is what later allows rare events to be paid for at the power cost $p = \alpha d$.

Throughout Sections 8–10, $\alpha > 0$ is fixed, $q = 1/\alpha$, and $L = A_\alpha d$, $\lambda = L^\alpha$, $a = \lambda/2$ and R with $\omega_d R^d = L$ are the cutoff scales. The capped valuation X, its survival function F, $G = -\log F$, the mark cutoff $\bar{e}_\lambda = G(a)$ and the price map $h(e)$ are those of Section 4.2; in particular F and G satisfy (4.8) and $h(E)$ has the law of X when $E \sim \mathrm{Exp}(1)$. Since $X \le a$, we have $\bar{e}_\lambda = \mathbb{E}(a + X)_+^q \le \lambda^q = L$. The marked Poisson process $\eta = \sum_i \delta_{(Y_i,E_i)}$ on $\Omega_T \times (0,\infty)$ has intensity $\mathbf{m}(dz,de) = dz\, e^{-e}de$, and $c_{d,\alpha}(x,y) = \omega_d^\alpha \,\mathrm{dist}(x,y)^p$. A radius r is called nonwrapping if $2r < T$; balls of nonwrapping radius are isometric to Euclidean balls. We assume throughout that $3R$ is nonwrapping. The conventions of Section 2.4 (order of limits, annealed moments, physical labels) are in force.

8.1 Cavity cells

Let $h_i = h(E_i)$. The capped dual potential and the cavity cells are

$$\phi(x) = \min\Big\{a,\ \min_i\big(c_{d,\alpha}(x,Y_i) - h_i\big)\Big\}, \quad C_i = \{x:\ c_{d,\alpha}(x,Y_i) - h_i = \phi(x) < a\}, \quad m_i = |C_i|. \tag{8.1}$$

All prices lie in $[-a,a]$. If $x \in C_i$, then $c_{d,\alpha}(x,Y_i) < 2a = \lambda$, that is $\omega_d \operatorname{dist}(x,Y_i)^d < L$; hence $C_i \subset B_R(Y_i)$ and $m_i \le L$. A site Y_j can win a point of $B_R(Y_i)$ against Y_i only if it lies in $B_{2R}(Y_i)$, so C_i is determined by the marked sites in $B_{2R}(Y_i)$. The cells are disjoint up to null sets: for $Y_i \neq Y_j$ the difference of the two reduced costs is almost everywhere differentiable with nonvanishing gradient, so its level sets are Lebesgue-null. The price of a site depends only on its own mark; adding a new site therefore leaves all old prices unchanged and can only shrink the old cells.

We compute the law of a cell by the void formula. For $t < a$, the set $\{y:\ c_{d,\alpha}(x,y) \le t + h(e)\}$ is a ball of volume $(t+h(e))_+^q$, of radius less than R. Since $h(E)$ has the law of X,

$$\mathbb{P}\Big(\min_j\big(c_{d,\alpha}(x,Y_j) - h_j\big) > t\Big) = \exp\Big(-\int_0^\infty (t+h(e))_+^q\, e^{-e} de\Big) = \exp\big(-\mathbb{E}(t+X)_+^q\big) = F(t). \tag{8.2}$$

By the Mecke formula [19], statements about a typical site with mark e refer to a site planted at the origin with mark e, added to η; we write $C^{(e)}$ for its cell, $z = h(e)$ for its price, and $m^{(e)} = |C^{(e)}|$. Conditional expectations given $E_i = e$ are understood in this Palm sense. With the radial volume $v = \omega_d |x|^d$, a point x lies in $C^{(e)}$ if and only if $v^\alpha - z < a$ and no other site has $c_{d,\alpha}(x,Y_j) - h_j \le v^\alpha - z$. By (8.2),

$$\mathbb{P}(x \in C^{(e)}) = \mathbf{1}_{\{v^\alpha - z < a\}} F(v^\alpha - z).$$

Integrating in the volume coordinate and substituting $v = (t+z)^q$ gives $\mathbb{E}m^{(e)} = q\int_{-a}^a (z+t)_+^{q-1} F(t)\,dt = G(z)$ by (4.8). Since $z = h(e) = G^{-1}(e \wedge \bar e_\lambda)$,

$$\mathbb{E}[m_i \mid E_i = e] = e \wedge \bar e_\lambda. \tag{8.3}$$

The variance of a cell is exponentially small relative to L^2, uniformly in the mark and in α:

$$\operatorname{Var}(m_i \mid E_i = e) \le \chi^d L^3, \qquad \chi = \frac{4}{3\sqrt{3}}, \qquad c_\chi = \log\frac{1}{\chi} > 0. \tag{8.4}$$

To prove (8.4), write $\operatorname{Var} m^{(e)} = \iint \operatorname{Cov}(\mathbf{1}_{\{x\in C^{(e)}\}}, \mathbf{1}_{\{y\in C^{(e)}\}})\,dx\,dy$. The indicator of x vanishes identically unless $v_x^\alpha - z < a$, which forces $x \in B_R := B_R(0)$. For two such probes x, y, let A_x and A_y be the $\mathbf{m}$-measures of the sets of marked sites that beat the planted site at x and at y, and let I_{xy} be the measure of their intersection. The void formula gives

$$\operatorname{Cov}(\mathbf{1}_{\{x\in C^{(e)}\}}, \mathbf{1}_{\{y\in C^{(e)}\}}) = e^{-A_x - A_y}\big(e^{I_{xy}} - 1\big) = e^{-A_x - A_y + I_{xy}}\big(1 - e^{-I_{xy}}\big) \le I_{xy}.$$

A site that beats the planted site at x lies in $B_R(x)$, and the marks have a probability density, so $I_{xy} \le |B_R(x) \cap B_R(y)| = (\mathbf{1}_{B_R} * \mathbf{1}_{B_R})(x-y)$. Since B_R is symmetric,

$$\operatorname{Var} m^{(e)} \le \iint \mathbf{1}_{B_R}(x)\mathbf{1}_{B_R}(y)\,(\mathbf{1}_{B_R} * \mathbf{1}_{B_R})(x-y)\,dx\,dy = \|\mathbf{1}_{B_R} * \mathbf{1}_{B_R}\|^2_{L^2(\mathbb{R}^d)}.$$

By the sharp Young inequality of Beckner and Brascamp–Lieb [4, 7] with exponents $(4/3, 4/3, 2)$, $\|f * g\|_2 \le A_{4/3}^{2d}\|f\|_{4/3}\|g\|_{4/3}$, where $A_p = (p^{1/p}/p'^{1/p'})^{1/2}$ is Beckner's constant; here $A_{4/3}^2 = (4/3)^{3/4}4^{-1/4} = 2\cdot 3^{-3/4}$. Hence $\|\mathbf{1}_{B_R} * \mathbf{1}_{B_R}\|_2^2 \le (A_{4/3}^4)^d |B_R|^3$ with $A_{4/3}^4 = 4/(3\sqrt{3}) = \chi$, and $|B_R| = L$. This proves (8.4).

8.2 Exact cores and the initial exceptional mass

Let $b > 0$ be the exponent fixed in (10.15) and put $\delta = e^{-2bd}$; this level is used from Section 9 on. With the convention $E_i/0 = +\infty$, the core weights, the core source measures, the core capacities and the two exceptional measures are

$$\begin{aligned} w_i &= \min\Big\{1, \frac{E_i}{m_i}\Big\}, \quad \mu_i^{\mathrm{c}} = w_i \mathbf{1}_{C_i}\, dx, \qquad \nu_i = w_i m_i = \min\{E_i, m_i\}, \\ \nu &= \sum_i \nu_i \delta_{Y_i}, \qquad \mu_0 = \Big(1 - \sum_i w_i \mathbf{1}_{C_i}\Big) dx, \quad \zeta_0 = \sum_i (E_i - m_i)_+ \delta_{Y_i}. \end{aligned} \tag{8.5}$$

The measure μ_0 is the exceptional source, absolutely continuous with density in $[0,1]$; ζ_0 is the exceptional target; and ν is the capacity of the core targets, each core site Y_i receiving the core source μ_i^{c} of mass $\nu_i \le L$. A site with $m_i = 0$ has no core, and its whole mark is exceptional. For every label $(E_i - m_i)_+ + \nu_i = E_i$ and $(m_i - E_i)_+ + \nu_i = m_i$. Moreover $\nu(\Omega_T) + \mu_0(\Omega_T) = V$.

We now show that the exceptional measures have exponentially small mean density: for all sufficiently large d,

$$M_0 := \frac{1}{V}\mathbb{E}\big[\mu_0(\Omega_T) + \zeta_0(\Omega_T)\big] \le P_\alpha(d)\, e^{-c_\chi d/2}, \qquad \frac{1}{V}\mathbb{E}\nu(\Omega_T) = 1 - \frac{1}{V}\mathbb{E}\mu_0(\Omega_T). \tag{8.6}$$

Here and below $P_\alpha(d)$ denotes a quantity bounded by a polynomial in d (for instance $L^{3/2} \le (A_\alpha d)^2$) whose coefficients depend only on the fixed parameters $\alpha, s, b, c_0, c_1, A_\alpha, k_*$ (the parameters s, c_0, c_1, k_* are fixed in Section 10). It may change from line to line, but never depends on the stage index k or on V.

The second identity in (8.6) is the expectation of $\nu(\Omega_T) + \mu_0(\Omega_T) = V$. For the first, the source not covered by the cores consists of $\{\phi = a\}$ and of the parts $(1 - w_i)\mathbf{1}_{C_i}\, dx$, of mass $m_i - \nu_i = (m_i - E_i)_+$. Hence

$$\mu_0(\Omega_T) + \zeta_0(\Omega_T) = |\{\phi = a\}| + \sum_i |m_i - E_i|.$$

By (8.2) at $t = a$, $\mathbb{P}(\phi(x) = a) = F(a) = e^{-\bar{e}_\lambda}$. By the Mecke formula, the sum has mean $V \int_0^\infty e^{-e}\, \mathbb{E}|m^{(e)} - e|\, de$. By (8.3) and (8.4),

$$\mathbb{E}|m^{(e)} - e| \le \big|\mathbb{E}m^{(e)} - e\big| + \big(\operatorname{Var} m^{(e)}\big)^{1/2} \le (e - \bar{e}_\lambda)_+ + \chi^{d/2} L^{3/2},$$

and $\int_0^\infty (e - \bar{e}_\lambda)_+ e^{-e} de = e^{-\bar{e}_\lambda}$. Altogether

$$M_0 \le 2e^{-\bar{e}_\lambda} + L^{3/2} e^{-c_\chi d/2}.$$

For the first term we use Lemma 4.3, valid for every $q > 0$: for d large, $\bar{e}_\lambda \ge c_q \lambda^q = c_q A_\alpha d$. We choose A_α so large that $c_q A_\alpha \ge c_\chi/2 + 1$. Then $2e^{-\bar{e}_\lambda} \le 2e^{-(c_\chi/2+1)d}$, which proves (8.6). The rate is $c_\chi/2$ rather than c_χ because the exceptional mass is a first absolute moment of the cell fluctuations.

The parameters are chosen in the following order: first α; then the moment parameter s and the exponents c_0, c_1, b in (10.15); then A_α with $c_q A_\alpha \ge c_\chi/2 + 1$, so that the saturated marks are dominated by the cell fluctuations in (8.6); then the preparatory length k_* by (10.17); then the dimension d large; and finally the volume V large. The polynomial factors in L and λ are absorbed at large fixed d.

8.3 Concentration of the initial fields

For a functional Φ of the marked configuration and $(z,e)\in\Omega_T\times(0,\infty)$, the add-one derivative is $D_{z,e}\Phi=\Phi(\eta+\delta_{(z,e)})-\Phi(\eta)$. The inserted site carries a new physical label and the old labels are kept, so labelled measures are compared on the union of the old and new labels. For a functional Φ and $t\ge 2$ put

$$S_2(\Phi)=\int |D_{z,e}\Phi|^2\,\mathbf{m}(dz,de),\qquad \mathcal{T}_t(\Phi)=\int |D_{z,e}\Phi|^t\,\mathbf{m}(dz,de).$$

For a signed measure σ, $|\sigma|$ is its total variation measure and $\|\sigma\|_{TV}=|\sigma|(\Omega_T)$; for labelled atomic measures this is $\sum_i|\sigma_i|$ over the union of labels.

Proposition 8.1. *Let $B=3L$ be the envelope of Lemma 8.2. Let $A\subset\Omega_T$ be a deterministic Borel set of volume $v=|A|\ge 1$, and let $\xi\in\{\mu_0,\nu,\zeta_0\}$. Then for every $s\in\{1,2,4,8,\dots\}$*

$$\mathbb{E}\Big|\frac{\xi(A)-\mathbb{E}\xi(A)}{v}\Big|^{2s}\le C_s(1+L)^sB^{2s}v^{-s},\tag{8.7}$$

$$\mathbb{E}\,\mathcal{T}_t(\xi(A)/v)\le C_tB^tv^{1-t},\qquad t\ge 2.\tag{8.8}$$

The constants C_s,C_t depend only on s and t.

The proof combines configurationwise add-one bounds with a recursion over dyadic moment orders. We first estimate the effect of an insertion. Let a new site be planted at (z,e), and let the old cell C_i lose the volume $\ell_i(z,e)=|C_i\setminus C_i'|$, where primes denote quantities after the insertion. The lost points belong to the new cell, which lies in $B_R(z)$ and has volume at most L. Hence

$$\ell_i(z,e)\le|C_i\cap B_R(z)|,\qquad \sum_i\ell_i(z,e)\le L,\qquad \int_{\Omega_T}\ell_i(z,e)\,dz\le Lm_i.\tag{8.9}$$

Lemma 8.2. *Let $B=3L$. For every configuration, every (z,e) and every Borel set A of volume v:*

$$\|D_{z,e}\mu_0\|_{TV}\le B,\qquad \|D_{z,e}\nu\|_{TV}\le B,\qquad \|D_{z,e}\zeta_0\|_{TV}\le B+e,\tag{8.10}$$

$$\begin{aligned}&\int dz\,|D_{z,e}\mu_0|(A)\le 3Lv,\\ &\int dz\,|D_{z,e}\nu|(A)\le Lv+L\sum_{Y_i\in A}m_i,\qquad \int dz\,|D_{z,e}\zeta_0|(A)\le ev+L\sum_{Y_i\in A}m_i.\end{aligned}\tag{8.11}$$

Proof. Fix an old label i with $m_i>0$; labels with $m_i=0$ do not change. Write $\ell_i=\ell_i(z,e)$ and $m_i'=m_i-\ell_i$, and let primes denote the quantities after the insertion. The mark E_i does not change. Since $m\mapsto\min\{E_i,m\}$ and $m\mapsto(E_i-m)_+$ are 1-Lipschitz,

$$|\nu_i'-\nu_i|\le\ell_i,\qquad \big|(E_i-m_i')_+-(E_i-m_i)_+\big|\le\ell_i.\tag{8.12}$$

The weight can only increase, and

$$0\le w_i'-w_i\le\frac{\ell_i}{m_i},\qquad (w_i'-w_i)\,m_i'\le\ell_i.\tag{8.13}$$

Indeed, if $m_i\le E_i$, then $w_i'=w_i=1$. If $m_i'\ge E_i$, then $w_i'-w_i=E_i\ell_i/(m_im_i')$, and $E_i\le m_i'$ gives both bounds. If $m_i'<E_i<m_i$, then $w_i'-w_i=(m_i-E_i)/m_i<\ell_i/m_i$, and $m_i'\le m_i$ gives the second bound. Since $C_i'\subset C_i$ and $|C_i\setminus C_i'|=\ell_i$,

$$\|\mu_i^{\mathrm{c}\prime}-\mu_i^{\mathrm{c}}\|_{TV}=(w_i'-w_i)\,m_i'+w_i\ell_i\le 2\ell_i.$$

The new site contributes a core atom $\nu_{\mathrm{new}} \le L$, a core source of the same mass, and an exceptional atom $(e - m_{\mathrm{new}})_+ \le e$. Summing over the old labels with $\sum_i \ell_i \le L$ gives $\|D_{z,e}\mu_0\|_{TV} \le 3L$, $\|D_{z,e}\nu\|_{TV} \le 2L$ and $\|D_{z,e}\zeta_0\|_{TV} \le L + e$, which is (8.10).

For (8.11), decompose $-D_{z,e}\mu_0$ into the new core source, the portions $w_i \mathbf{1}_{C_i \setminus C_i'}\,dx$ lost by the old cores, and the weight increments $(w_i' - w_i)\mathbf{1}_{C_i'}\,dx$. The first two are carried by the new cell, which lies in $B_R(z)$, so each has $\int dz\,|\cdot|(A) \le \iint \mathbf{1}_A(x)\mathbf{1}_{B_R(z)}(x)\,dx\,dz = Lv$. By (8.13) and (8.9), the third has $\int dz \sum_i (w_i' - w_i)|A \cap C_i| \le L \sum_i |A \cap C_i| \le Lv$, because the cells are disjoint. For the atomic measures only the old roots in A and the new atom at $z \in A$ contribute to $|D_{z,e}\xi|(A)$. Integrating (8.12) in z with (8.9) bounds the old core atoms and the old exceptional atoms by $L \sum_{Y_i \in A} m_i$ each, the new core atom by Lv, and the new exceptional atom by ev. □

The bounds (8.11) are uniform in e except for the term ev, which is the actual fresh exceptional atom. For a Borel set A and $t \ge 2$ we combine them in the weighted integrals

$$I_A^{(t)}(\xi) = \int dz \int_0^\infty e^{-e}\, b_\xi(e)^{t-1}\, |D_{z,e}\xi|(A)\, de, \qquad b_\xi(e) = \begin{cases} B, & \xi \in \{\mu_0, \nu\}, \\ B + e, & \xi = \zeta_0. \end{cases}$$

Since $m_i \le L$, $L \le B$ and $\int_0^\infty (B + e)\, e\, e^{-e} de = B + 2$, Lemma 8.2 gives, before any expectation,

$$I_A^{(2)}(\mu_0) \le CB^2 v, \qquad I_A^{(2)}(\nu) + I_A^{(2)}(\zeta_0) \le CB^2(v + LN_A), \qquad N_A = \#\{i :\ Y_i \in A\}. \tag{8.14}$$

Since $N_A \sim \mathrm{Pois}(v)$, every fixed moment of $v + LN_A$ is $O_m\big((1 + L)^m v^m\big)$ when $v \ge 1$.

The second tool is a moment recursion for Poisson functionals.

Lemma 8.3. *Let η be a Poisson process on a measurable space with σ-finite intensity $\mathbf{m}$, let Φ be a functional of η with $\mathbb{E}\Phi^2 < \infty$, write $D_w\Phi = \Phi(\eta + \delta_w) - \Phi(\eta)$, define S_2 and $\mathcal{T}_t$ as above with this intensity, and let $Z = \Phi - \mathbb{E}\Phi$. Then $\mathbb{E}Z^2 \le \mathbb{E}S_2(\Phi)$. If $m \ge 2$ is an integer and $\mathbb{E}|\Phi|^{2m} < \infty$, then*

$$\mathbb{E}|Z|^{2m} \le C_m \Big[\big(\mathbb{E}|Z|^m\big)^2 + \mathbb{E}S_2(\Phi)^m + \mathbb{E}\,\mathcal{T}_{2m}(\Phi) \Big], \tag{8.15}$$

where C_m depends only on m.

Proof. The first claim is the Poisson Poincaré inequality [19]. The second is the standard iteration of that inequality: apply it to $Z^m \in L^2$. Since $D_w(Z^m) = (Z + D_w\Phi)^m - Z^m$,

$$\mathbb{E}|Z|^{2m} - \big(\mathbb{E}Z^m\big)^2 \le \mathbb{E} \int \big|(Z + D_w\Phi)^m - Z^m\big|^2\, \mathbf{m}(dw).$$

The elementary inequality $|(x + y)^m - x^m|^2 \le C_m(|x|^{2m-2}|y|^2 + |y|^{2m})$, valid for real x, y of either sign, bounds the right side by $C_m[\mathbb{E}(|Z|^{2m-2}S_2(\Phi)) + \mathbb{E}\,\mathcal{T}_{2m}(\Phi)]$. Hölder's inequality with exponents $m/(m-1)$ and m gives

$$\mathbb{E}\big(|Z|^{2m-2}S_2(\Phi)\big) \le \big(\mathbb{E}|Z|^{2m}\big)^{(m-1)/m} \big(\mathbb{E}S_2(\Phi)^m\big)^{1/m}.$$

By Young's inequality, $C_m x^{(m-1)/m} y^{1/m} \le x/2 + C_m' y$. Since $\mathbb{E}|Z|^{2m} < \infty$, the term $\frac{1}{2}\mathbb{E}|Z|^{2m}$ can be absorbed into the left side, and $(\mathbb{E}Z^m)^2 \le (\mathbb{E}|Z|^m)^2$. This gives (8.15). □

For $s \in \{1, 2, 4, \dots\}$ one starts from the Poincaré inequality at order 2 and applies (8.15) successively with $m = 2, 4, \dots, s$; each application passes from order m to order $2m$ for the moments of Z, so that every intermediate dyadic order is used. In all applications

below, η is the marked Poisson process on $\Omega_T \times (0,\infty)$ and $|\Phi| \le C(V + N + \sum_i E_i)$, where $N = \eta(\Omega_T \times (0,\infty)) \sim \mathrm{Pois}(V)$. These variables have moments of all orders, so the integrability hypothesis holds on the finite torus. (Alternatively, one may apply the lemma with marks truncated at a level M, where all functionals are bounded by $C_M(V+N)$, and let $M \to \infty$ by Fatou's lemma, using that the derivative bounds below are uniform in M.) The constants depend on m only, and not on d, V or the stage index.

Proof of Proposition 8.1. Let $Z = (\xi(A) - \mathbb{E}\xi(A))/v$. Since $|D_{z,e}\xi(A)| \le |D_{z,e}\xi|(A) \le \|D_{z,e}\xi\|_{TV} \le b_\xi(e)$ by (8.10), we have $|D_{z,e}Z|^t \le v^{-t} b_\xi(e)^{t-1} |D_{z,e}\xi|(A)$, hence

$$S_2(Z) \le v^{-2} I_A^{(2)}(\xi), \qquad \mathcal{T}_t(Z) \le v^{-t} I_A^{(t)}(\xi).$$

By (8.14) and the moments of N_A, for every integer $m \ge 1$,

$$\mathbb{E} S_2(Z)^m \le C_m (1+L)^m B^{2m} v^{-m}. \tag{8.16}$$

For $\mathcal{T}_t$ we bound the expectation of $I_A^{(t)}$, which is smaller than the configurationwise bound (8.14) suggests. For μ_0, the first bound in (8.11) is uniform in e, so $I_A^{(t)}(\mu_0) \le 3B^{t-1}Lv = B^t v$. For ν we use the Mecke formula and (8.3):

$$\mathbb{E} \sum_{Y_i \in A} m_i = v\, \mathbb{E}_{E \sim \mathrm{Exp}(1)} (E \wedge \bar{e}_\lambda) \le v. \tag{8.17}$$

Together with (8.11), the old atoms contribute at most $B^{t-1}Lv \le B^t v$ to $\mathbb{E} I_A^{(t)}(\nu)$, and the new core atom at most $B^{t-1}Lv \le B^t v$. For ζ_0 the fresh mark is integrated against its density:

$$\mathbb{E} I_A^{(t)}(\zeta_0) \le L\, \mathbb{E} \sum_{Y_i \in A} m_i \int_0^\infty e^{-e} (B+e)^{t-1} de + v \int_0^\infty e^{-e}\, e\, (B+e)^{t-1} de \le C_t B^t v, \tag{8.18}$$

since $B \ge 1$ and t is fixed. Thus $\mathbb{E}\,\mathcal{T}_t(Z) \le C_t B^t v^{1-t}$ for all three measures, which is (8.8). This bound is annealed: its configurationwise input is (8.11), the Palm identity (8.17) removes the factor L carried by the old roots in (8.14), and the fresh mark is integrated against $e^{-e}de$.

We now prove (8.7) by induction over dyadic s. For $s = 1$, $\mathbb{E}Z^2 \le \mathbb{E}S_2(Z) \le C(1+L)B^2 v^{-1}$. If (8.7) holds for s, then Lemma 8.3 with $m = 2s$ (the functional satisfies $|\xi(A)| \le C(V + N + \sum_i E_i)$), (8.16) and (8.8) give

$$\mathbb{E}|Z|^{4s} \le C_s\big[(1+L)^{2s} B^{4s} v^{-2s} + (1+L)^{2s} B^{4s} v^{-2s} + B^{4s} v^{1-4s}\big] \le C_s (1+L)^{2s} B^{4s} v^{-2s},$$

because $v^{1-4s} \le v^{-2s}$ for $v \ge 1$. The constants are independent of d, V, A and α. □

9 Rescue maps and finite propagation

The core plan of Section 8 leaves an exceptional source μ_0 and an exceptional target ζ_0 of mean density at most $P_\alpha(d) e^{-c_\chi d/2}$. We transport them in stages $k = 1, 2, \dots$ on geometrically growing radii r_k, with two different partial plans. The target rescue $\mathcal{F}$ serves exceptional target atoms from Lebesgue measure, with source density at most a small level δ_k. The source rescue $\mathcal{H}$ sends exceptional source mass to the core sites, at most $\delta_k \nu_i$ to the site Y_i. This section proves the deterministic properties of these maps that the moment estimates of Section 10 require: monotonicity, stability in total variation under the insertion of a site, and finite range of dependence. Only the geometry of the normalized ball kernel enters; the exponent p plays no role.

Let $K_r = |B_r|^{-1} \mathbf{1}_{B_r}$ be the averaging kernel of a nonwrapping radius r.

9.1 The target rescue

For a finite positive atomic measure ζ on Ω_T, a radius r and a level $\delta_\star \in [0,1]$, let $Q = K_r * \zeta$ and

$$\mathcal{F}_{r,\delta_\star}(dx,dy) = K_r(x-y)\min\Big\{1, \frac{\delta_\star}{Q(x)}\Big\}\,dx\,\zeta(dy),$$

with the multiplier equal to 1 where $Q(x)=0$. Let $\mathcal{R}_{r,\delta_\star}\zeta$ be the part of ζ that is not served, that is, ζ minus the target marginal of $\mathcal{F}_{r,\delta_\star}$.

Lemma 9.1. *The plan $\mathcal{F}_{r,\delta_\star}$ moves mass over distances less than r and has source marginal $\min(Q,\delta_\star)\,dx \le \delta_\star dx$. The unserved target is*

$$\mathcal{R}_{r,\delta_\star}\zeta(dy) = \zeta(dy)\int K_r(x-y)\Big(1-\frac{\delta_\star}{Q(x)}\Big)_+ dx,$$

so that $0 \le \mathcal{R}_{r,\delta_\star}\zeta \le \zeta$ and

$$\mathcal{R}_{r,\delta_\star}\zeta(\Omega_T) = \int_{\Omega_T} (Q-\delta_\star)_+\,dx. \tag{9.1}$$

The map $\mathcal{R} = \mathcal{R}_{r,\delta_\star}$ is order preserving, and $\|\mathcal{R}\zeta - \mathcal{R}\zeta'\|_{TV} \le \|\zeta-\zeta'\|_{TV}$ for all finite positive atomic measures ζ,ζ'.

Proof. The formula for $\mathcal{R}\zeta$ follows from $\int K_r(x-y)\,dx = 1$ and $1-\min\{1,\delta_\star/Q\} = (1-\delta_\star/Q)_+$. By Fubini, $\mathcal{R}\zeta(\Omega_T) = \int Q(1-\delta_\star/Q)_+dx = \int (Q-\delta_\star)_+dx$. If $\zeta \le \zeta'$, then $Q \le Q'$, and the unserved multiplier $(1-\delta_\star/Q)_+$ is nondecreasing in Q, so $\mathcal{R}\zeta \le \mathcal{R}\zeta'$. Since $Q \mapsto (Q-\delta_\star)_+$ is 1-Lipschitz and nondecreasing,

$$\|\mathcal{R}\zeta' - \mathcal{R}\zeta\|_{TV} = \int [(Q'-\delta_\star)_+ - (Q-\delta_\star)_+]dx \le \int (Q'-Q)\,dx = \|\zeta'-\zeta\|_{TV}$$

for ordered inputs. For general ζ,ζ', compare both with their common part $\zeta\wedge\zeta'$ and add the two ordered bounds. □

All these statements are configurationwise; in particular they apply to the measures obtained after the insertion of a site.

9.2 The source rescue

The source rescue caps the mass sent to each core site. We use the cap map, defined for a finite positive measure π and $c \ge 0$ by $\mathscr{C}_c(\pi) = \pi\min\{1, c/\pi(\Omega_T)\}$ if $\pi(\Omega_T) > 0$ and $\mathscr{C}_c(0) = 0$.

Lemma 9.2. *For finite positive measures π,π' and $c,c' \ge 0$,*

$$\|\mathscr{C}_c(\pi) - \mathscr{C}_{c'}(\pi')\|_{TV} \le 2\|\pi-\pi'\|_{TV} + |c-c'|. \tag{9.2}$$

Proof. At fixed π the difference has total variation $|\min(c,\pi(\Omega_T)) - \min(c',\pi(\Omega_T))| \le |c-c'|$. At fixed c, let $m = \pi(\Omega_T)$, $m' = \pi'(\Omega_T)$, $w = \min(1,c/m)$ and $w' = \min(1,c/m')$. Adding and subtracting $w\pi'$ gives $\|w\pi - w'\pi'\|_{TV} \le \|\pi-\pi'\|_{TV} + m'|w-w'|$. The scalar bound $m'|w-w'| \le |m-m'|$ follows by checking the four positions of m,m' relative to c: it is trivial if $m,m' \le c$; if $m,m' \ge c$, then $m'|c/m - c/m'| = c|m-m'|/m \le |m-m'|$; if $m \le c \le m'$, then $m'(1-c/m') = m'-c \le m'-m$; and if $m' \le c \le m$, then $m'(1-c/m) \le m-c \le m-m'$. Since $|m-m'| \le \|\pi-\pi'\|_{TV}$, this proves (9.2); the case of a zero measure follows by continuity. □

Labels are physical labels: we take a common index set containing all sites of both configurations under comparison and give zero capacity to missing labels. Let $\nu = \sum_i \nu_i \delta_{Y_i}$ be a capacity measure, let $\mu = u\,dx$ be a source with $0 \le u \le 1$, and let $\tau \in (0,1]$. Put

$$t_\nu = K_r * \nu, \qquad P_i^\nu(x) = \frac{\nu_i K_r(x - Y_i)}{\max(t_\nu(x), \tau)}, \qquad \pi_i^\nu(dx) = u(x) P_i^\nu(x)\,dx.$$

Then $\sum_i P_i^\nu \le 1$, and $\sum_i P_i^\nu(x) = 1$ wherever $t_\nu(x) \ge \tau$. For $\delta_\star \in [0,1]$ the source rescue and its residual are

$$\mathcal{H}_i = \mathscr{C}_{\delta_\star \nu_i}(\pi_i^\nu) = u\, P_i^\nu \min\Big\{1, \frac{\delta_\star \nu_i}{\pi_i^\nu(\Omega_T)}\Big\} dx, \qquad \mathcal{S}_{r,\delta_\star}(\mu;\nu) = \mu - \sum_i \mathcal{H}_i. \tag{9.3}$$

The plan $\mathcal{H}_i$ sends the source mass $\mathcal{H}_i$ to the site Y_i over distances less than r. We have $\sum_i \mathcal{H}_i \le \mu$ and $\mathcal{H}_i(\Omega_T) \le \delta_\star \nu_i$. The residual mass splits into a normalization defect and a cap rejection:

$$\begin{aligned} \mathcal{S}_{r,\delta_\star}(\mu;\nu)(\Omega_T) &= \int u\Big(1 - \sum_i P_i^\nu\Big) dx + \sum_i \big(\pi_i^\nu(\Omega_T) - \delta_\star \nu_i\big)_+ \\ &\le |\{t_\nu < \tau\}| + \sum_i \big(\pi_i^\nu(\Omega_T) - \delta_\star \nu_i\big)_+. \end{aligned} \tag{9.4}$$

Lemma 9.3. *Let $\tau \in (0,1]$ and $\delta_\star \in [0,1]$. At fixed ν, the map $\mu \mapsto \mathcal{S}_{r,\delta_\star}(\mu;\nu)$ is order preserving on sources with density in $[0,1]$, and*

$$\|\mathcal{S}_{r,\delta_\star}(\mu;\nu) - \mathcal{S}_{r,\delta_\star}(\mu';\nu)\|_{TV} \le \|\mu - \mu'\|_{TV}. \tag{9.5}$$

At fixed μ,

$$\|\mathcal{S}_{r,\delta_\star}(\mu;\nu) - \mathcal{S}_{r,\delta_\star}(\mu;\nu')\|_{TV} \le \sum_i \|\mathcal{H}_i^\nu - \mathcal{H}_i^{\nu'}\|_{TV} \le \Big(\frac{4}{\tau} + \delta_\star\Big)\|\nu - \nu'\|_{TV}, \tag{9.6}$$

where $\mathcal{H}_i^\nu$ denotes the plan (9.3) *built with the capacities ν. In particular, for $\tau = 1/2$ the Lipschitz constant in ν is at most* 9.

Proof. Fix ν and let $\mu \le \mu'$. The provisional masses $\pi_i(\mu)(\Omega_T) = \int P_i^\nu\, d\mu$ increase, so the multipliers $w_i(\mu) = \min\{1, \delta_\star \nu_i / \pi_i(\mu)(\Omega_T)\}$ decrease. Since $\mathcal{S}(\mu;\nu) = \mu(1 - \sum_i w_i(\mu) P_i^\nu)$,

$$\mathcal{S}(\mu';\nu) - \mathcal{S}(\mu;\nu) = (\mu' - \mu)\Big(1 - \sum_i w_i(\mu') P_i^\nu\Big) + \mu \sum_i [w_i(\mu) - w_i(\mu')] P_i^\nu \ge 0.$$

The served mass $\sum_i \min\{\pi_i(\mu)(\Omega_T), \delta_\star \nu_i\}$ is nondecreasing in μ, so the residual mass increases by at most $(\mu' - \mu)(\Omega_T)$. Comparing general μ, μ' with $\mu \wedge \mu'$ gives (9.5).

For a change of capacities, the pointwise bound

$$\sum_i |P_i^\nu(x) - P_i^{\nu'}(x)| \le \tau^{-1} \sum_i |\nu_i - \nu_i'| K_r(x - Y_i) + \tau^{-1} |t_\nu(x) - t_{\nu'}(x)|$$

follows from $|P_i^\nu - P_i^{\nu'}| \le |\nu_i - \nu_i'| K_r(\cdot - Y_i)/\max(t_\nu, \tau) + \nu_i' K_r(\cdot - Y_i)\,|\max(t_\nu, \tau) - \max(t_{\nu'}, \tau)|/(\max(t_\nu, \tau)\max(t_{\nu'}, \tau))$ and $\sum_i \nu_i' K_r(\cdot - Y_i) = t_{\nu'} \le \max(t_{\nu'}, \tau)$. Integrating against $u\,dx \le dx$ and using $\int K_r = 1$ gives $\sum_i \|\pi_i^\nu - \pi_i^{\nu'}\|_{TV} \le 2\tau^{-1}\|\nu - \nu'\|_{TV}$. Lemma 9.2 with caps $\delta_\star \nu_i$ and $\delta_\star \nu_i'$ gives the second inequality in (9.6). The first holds because $\mathcal{S}(\mu;\nu) - \mathcal{S}(\mu;\nu') = \sum_i (\mathcal{H}_i^{\nu'} - \mathcal{H}_i^\nu)$. □

9.3 The stages and their add-one envelopes

Fix $\tau = 1/2$ and positive parameters c_0, c_1, chosen in (10.15). The stage radii, their volumes and the stage levels are

$$r_k = Re^{a_k}, \quad a_k = c_0 + (k-1)c_1, \qquad v_k = |B_{r_k}| = Le^{a_k d}, \qquad \delta_k = \delta 2^{-k-3}, \qquad k \geq 1. \tag{9.7}$$

Note that $\sum_{k\geq 1} \delta_k = \delta/8$. Starting from ζ_0 and $\mu_0 = u_0\, dx$ of (8.5), stage k applies the target rescue to ζ_{k-1} and the source rescue to μ_{k-1}, both at radius r_k and level δ_k, with the fixed capacities ν:

$$\mathcal{F}_k = \mathcal{F}_{r_k,\delta_k}(\zeta_{k-1}), \qquad \zeta_k = \mathcal{R}_{r_k,\delta_k}\zeta_{k-1}, \qquad \mathcal{H}_{i,k} = \mathcal{H}_i, \qquad \mu_k = u_k\, dx = \mathcal{S}_{r_k,\delta_k}(\mu_{k-1};\nu),$$

where $\mathcal{H}_i$ is defined by (9.3) with $(u,\nu,r,\delta_\star) = (u_{k-1},\nu,r_k,\delta_k)$. We write $\zeta_k = \sum_i \zeta_{k,i}\delta_{Y_i}$, $Q_k^\zeta = K_{r_k} * \zeta_{k-1}$ and $Q_k^\mu = K_{r_k} * \mu_{k-1}$ for the stage-k ball loads, $t_k = K_{r_k} * \nu$, and $P_{i,k}$, $\pi_{i,k}$ and $w_{i,k} = \min\{1, \delta_k\nu_i/\pi_{i,k}(\Omega_T)\}$ for the functions P_i^ν, the provisional measures and the cap multipliers of stage k, so that $\mathcal{H}_{i,k} = w_{i,k}\pi_{i,k}$. The cumulative served target density and the cumulative source rescue of the label i are

$$f_k = \sum_{j\leq k} \min(Q_j^\zeta, \delta_j) \leq \frac{\delta}{8}, \qquad \mathcal{H}_i^{[k]} = \sum_{j\leq k} \mathcal{H}_{i,j}, \qquad \mathcal{H}_i^{[k]}(\Omega_T) \leq \frac{\delta}{8}\,\nu_i.$$

Stage k is called admissible if $4\Lambda r_k < T$, with Λ defined in (9.10); then all balls used up to stage k, including those of Lemma 9.5, have nonwrapping radii.

Lemma 9.4. *For every $k \geq 0$, every configuration and every (z,e),*

$$\|D_{z,e}\zeta_k\|_{TV} \leq B + e, \qquad \|D_{z,e}\mu_k\|_{TV} \leq (1+9k)B, \qquad 0 \leq \zeta_k \leq \zeta_0, \qquad 0 \leq \mu_k \leq \mu_0 \leq dx. \tag{9.8}$$

Proof. The order relations follow from $\mathcal{R}\zeta \leq \zeta$ and $\mathcal{S}(\mu;\nu) \leq \mu$; in particular $u_k \leq 1$ at every stage. For $k = 0$ the envelopes are (8.10). The maps $\mathcal{R}$ and $\mathcal{S}$ are applied in the same way to the configuration with and without the inserted site. By Lemma 9.1, $\|D_{z,e}\zeta_k\|_{TV} \leq \|D_{z,e}\zeta_{k-1}\|_{TV}$. Writing μ^+, ν^+ for the measures after insertion, Lemma 9.3 gives

$$\begin{aligned}\|D_{z,e}\mu_k\|_{TV} &\leq \|\mathcal{S}(\mu_{k-1}^+;\nu^+) - \mathcal{S}(\mu_{k-1}^+;\nu)\|_{TV} + \|\mathcal{S}(\mu_{k-1}^+;\nu) - \mathcal{S}(\mu_{k-1};\nu)\|_{TV} \\ &\leq 9\|D_{z,e}\nu\|_{TV} + \|D_{z,e}\mu_{k-1}\|_{TV},\end{aligned}$$

and $\|D_{z,e}\nu\|_{TV} \leq B$. Induction on k concludes. □

The source envelope grows linearly in k. Consequently the late stages of Section 10 use the insertion shift $(1+9(k-1))B/v_k$ of the ball load Q_k^μ, whereas the preparatory stages compare μ_{k-1} with the initial field μ_0 and use the shift B/v_k.

9.4 Finite propagation

The density that the flux argument of Section 12 transports is built from the following auxiliary objects. For labels with $\nu_i > 0$ let

$$\mathcal{D}_{i,k} = (1-f_k)(\mu_i^{\mathrm{c}} + \mathcal{H}_i^{[k]}), \qquad d_{i,k} = \mathcal{D}_{i,k}(\Omega_T), \qquad \beta_{i,k} = \frac{\nu_i}{d_{i,k}},$$

and define the auxiliary source density

$$\rho_k = 1 + \sum_{i:\,\nu_i>0} (\beta_{i,k} - 1)\frac{d\mathcal{D}_{i,k}}{dx}. \tag{9.9}$$

Since $f_k \le \delta/8 < 1$ and $\mu_i^{\mathrm{c}}(\Omega_T) = \nu_i > 0$, the normalizers $d_{i,k}$ are positive. At $k = 0$ we have $f_0 = 0$ and $\mathcal{H}_i^{[0]} = 0$, hence $d_{i,0} = \nu_i$, $\beta_{i,0} = 1$ and $\rho_0 \equiv 1$.

An observable attached to a site Y_i is rooted-local within r if it depends only on the marked configuration in $B_r(Y_i)$. A field $\Phi(x)$ is point-local within r if $\Phi(x)$ depends only on the marked configuration in $B_r(x)$. These are deterministic statements about all configurations, including those with an inserted site. Let $\bar r_k = \sum_{j\le k} r_j$, with $\bar r_0 = 0$.

Lemma 9.5. *Let $k \ge 0$ and assume that the stages $1, \dots, k$ are admissible. Then:*

(i) *the residual target coefficient $\zeta_{k,i}$ of each label is rooted-local within $2R + 2\bar r_k$, and f_k is point-local within $2R + 2\bar r_k$;*

(ii) *u_k is point-local within $3R + 3\bar r_k$;*

(iii) *$d_{i,k}$ and $\beta_{i,k}$ are rooted-local within $3R + 3\bar r_k$;*

(iv) *ρ_k is point-local within $\ell_k = 3R + 4\bar r_k$.*

For $k \ge 1$, the ball loads $Q_k^{\zeta}(x)$ and $Q_k^{\mu}(x)$ depend only on the marked configuration in $B_{\Lambda r_k}(x)$, where

$$\Lambda = 4 + \frac{3}{e^{c_1} - 1}, \tag{9.10}$$

and $\ell_k \le 2\Lambda r_k$.

Proof. At $k = 0$, the quantities m_i, w_i, ν_i and $\zeta_{0,i} = (E_i - m_i)_+$ are rooted-local within $2R$ by Section 8, and $u_0(x) = 1 - \sum_i w_i \mathbf{1}_{C_i}(x)$ involves only roots within R of x, so it is point-local within $3R$. Assume (i) and (ii) at stage $k - 1$.

Target rescue. A site contributing to $Q_k^{\zeta}(x)$ lies within r_k of x and brings its rooted radius, so Q_k^{ζ} is point-local within $2R + 2\bar r_{k-1} + r_k$. The new coefficient $\zeta_{k,i} = \zeta_{k-1,i} \int_{B_{r_k}(Y_i)} K_{r_k}(x - Y_i)(1 - \delta_k/Q_k^{\zeta}(x))_+ dx$ integrates over $B_{r_k}(Y_i)$, which adds a further r_k; it is therefore rooted-local within $2R + 2\bar r_k$. The increment $\min(Q_k^{\zeta}, \delta_k)$ of f_k is point-local within $2R + 2\bar r_{k-1} + r_k$, which proves (i).

Source rescue. The field t_k is point-local within $2R + r_k$. The provisional mass

$$\pi_{i,k}(\Omega_T) = \int_{B_{r_k}(Y_i)} u_{k-1}(x)\, \frac{\nu_i K_{r_k}(x - Y_i)}{\max(t_k(x), \tau)}\, dx$$

and the cap multiplier of the label i are rooted-local within $\max\{3R + 3\bar r_{k-1} + r_k,\, 2R + 2r_k\}$. The new density $u_k(x) = u_{k-1}(x)\big(1 - \sum_i w_{i,k} P_{i,k}(x)\big)$ involves the roots within r_k of x, which adds a last r_k: u_k is point-local within $\max\{3R + 3\bar r_{k-1} + 2r_k,\, 2R + 3r_k\} \le 3R + 3\bar r_k$. The density of $\mathcal{H}_{i,j}$ at x is $u_{j-1}(x) P_{i,j}(x) w_{i,j}$, and these components keep the radii just established. This proves (ii).

Normalizers. The measures μ_i^{c} and $\mathcal{H}_i^{[k]}$ are supported in $B_{r_k}(Y_i)$, since $R \le r_j \le r_k$ for $j \le k$. Hence $d_{i,k}$ involves f_k on $B_{r_k}(Y_i)$, which is rooted-local within $2R + 2\bar r_k + r_k \le 3R + 3\bar r_k$, and the components of $\mathcal{H}_i^{[k]}$, which are rooted-local within $3R + 3\bar r_k$. The same holds for $\beta_{i,k}$, proving (iii). In (9.9), only roots within r_k of x contribute, each with rooted radius $3R + 3\bar r_k$, and $d\mathcal{D}_{i,k}/dx(x)$ has the same locality. Thus ρ_k is point-local within $3R + 3\bar r_k + r_k \le 3R + 4\bar r_k$, which is (iv).

Finally let $k \ge 1$. Since $r_j = r_k e^{-(k-j)c_1}$, we have $\bar r_{k-1} \le r_k/(e^{c_1} - 1)$, and $R \le r_k$. The load $Q_k^{\mu}(x)$ is point-local within $3R + 3\bar r_{k-1} + r_k \le (4 + 3/(e^{c_1} - 1)) r_k = \Lambda r_k$, and $Q_k^{\zeta}(x)$ within the smaller radius $2R + 2\bar r_{k-1} + r_k$. Moreover $\ell_k = 3R + 4\bar r_{k-1} + 4r_k \le (7 + 4/(e^{c_1} - 1)) r_k \le 2\Lambda r_k$. □

The finite-range flux estimate of Section 12 is applied to the auxiliary densities ρ_k; the density obtained after the completion of Section 11 is generally nonlocal.

10 Stage induction and choice of parameters

We now show that the mean residual mass after stage k,

$$M_k = m_F(k) + m_H(k), \qquad m_F(k) = \frac{1}{V}\mathbb{E}\zeta_k(\Omega_T), \qquad m_H(k) = \frac{1}{V}\mathbb{E}\mu_k(\Omega_T),$$

decays fast enough in k and d for the total cost of all stages to be exponentially small in d (Proposition 10.4). For $k = 0$ this is the quantity M_0 of (8.6). The induction has two phases. During a fixed number k_* of preparatory stages the residuals are dominated by the initial fields, whose moments are controlled by Proposition 8.1. At later stages we use moment bounds for the current residual fields; these pay the locality loss Λ^{sd} once, in the current stage. The parameters are fixed at the end of the section.

10.1 Moments of the stage loads

We first prove the moment bounds for the current residual fields. They rely on the envelopes (9.8) and the finite propagation of Lemma 9.5.

Lemma 10.1. *Let $s \in \{1, 2, 4, \dots\}$ and let $k \geq 1$ be an admissible stage with $v_k \geq 1$. Let Q_k denote either $Q_k^{\zeta}(x)$ or $Q_k^{\mu}(x)$ at a fixed point x. Then*

$$\mathbb{E}|Q_k - \mathbb{E}Q_k|^{2s} \leq C_s(k+1)^{2s}B^{2s}\frac{\Lambda^{sd}}{v_k^s}. \tag{10.1}$$

Let Ξ be a functional of the marked configuration with $|\Xi| \leq C(V + N + \sum_i E_i)$ and $|D_{z,e}\Xi| \leq C(1+k)B + e$ for all (z, e), for instance the total mass of ν, of μ_j or of ζ_j with $j \leq k$. Then, for $V \geq 1$,

$$\mathbb{E}|\Xi - \mathbb{E}\Xi|^{2s} \leq C_s(1+k)^{2s}B^{2s}(V^s + V). \tag{10.2}$$

The constants C_s depend only on s and on the constant C in the hypotheses.

Proof. By (9.8), $\|D_{z,e}\mu_{k-1}\|_{TV} \leq (1 + 9(k-1))B \leq (1+9k)B$ and $\|D_{z,e}\zeta_{k-1}\|_{TV} \leq B + e$, and $\|K_{r_k}\|_\infty = v_k^{-1}$. By Lemma 9.5, an insertion outside $B_{\Lambda r_k}(x)$ does not change Q_k. Hence

$$|D_{z,e}Q_k| \leq \big[(1+9k)B + e\big]v_k^{-1}\mathbf{1}_{B_{\Lambda r_k}(x)}(z).$$

The influence region has volume $\Lambda^d v_k$, so, deterministically,

$$S_2(Q_k) \leq C(k+1)^2B^2\Lambda^d/v_k, \qquad \mathcal{T}_{2m}(Q_k) \leq C_m(k+1)^{2m}B^{2m}\Lambda^d v_k^{1-2m}.$$

We apply Lemma 8.3 along the orders $2, 4, \dots, 2s$; the functional satisfies $|Q_k| \leq (V + \sum_i E_i)/v_k$. At order 2 the bound is $C(k+1)^2B^2\Lambda^d v_k^{-1}$. If (10.1) holds at order $2m$, the three terms of (8.15) at order $4m$ are at most $C_m(k+1)^{4m}B^{4m}\Lambda^{2md}v_k^{-2m}$: for the jump term this uses $\Lambda^d v_k^{1-4m} \leq \Lambda^{2md}v_k^{-2m}$, valid since $\Lambda, v_k \geq 1$. This proves (10.1).

For Ξ the influence region is all of Ω_T, so $S_2(\Xi) \leq CV(1+k)^2B^2$ and $\mathcal{T}_{2m}(\Xi) \leq C_mV(1+k)^{2m}B^{2m}$. The same recursion gives (10.2). At order 2 it gives $\mathbb{E}|\Xi - \mathbb{E}\Xi|^2 \leq CV(1+k)^2B^2$. If $\mathbb{E}|\Xi - \mathbb{E}\Xi|^{2m} \leq C_m(1+k)^{2m}B^{2m}(V^m + V)$, then at order $4m$ the three terms of (8.15) are at most $C_m(1+k)^{4m}B^{4m}(V^m+V)^2 \leq 4C_m(1+k)^{4m}B^{4m}V^{2m}$, $C_m(1+k)^{4m}B^{4m}V^{2m}$ and $C_m(1+k)^{4m}B^{4m}V$, using $(V^m+V)^2 \leq 4V^{2m}$ for $m \geq 1$ and $V \geq 1$. This gives the bound at order $4m$ with $V^{2m} + V$. □

The locality loss Λ^{sd} in (10.1) comes from the volume $\Lambda^d v_k$ of the influence region of the current stage only. The term V in (10.2) is dominated by V^s for $s \geq 1$; it records the contribution of the Poisson jumps.

10.2 Overloads and the Palm estimate

The residual masses are controlled through positive overloads. If a random variable Q has mean m, then for every $s \geq 1$

$$\mathbb{E}(Q-t)_+ \leq \frac{\mathbb{E}|Q-m|^{2s}}{(t-m)^{2s-1}} \quad (t>m), \qquad \mathbb{P}(Q \leq t) \leq \frac{\mathbb{E}|Q-m|^{2s}}{(m-t)^{2s}} \quad (t<m). \tag{10.3}$$

The first inequality holds because $0 < Q-t \leq |Q-m| \leq |Q-m|^{2s}/(t-m)^{2s-1}$ on $\{Q>t\}$; the second is Markov's inequality.

By (9.1) and stationarity, $m_F(k) = \mathbb{E}(Q_k^{\zeta}(0)-\delta_k)_+$. For the source rescue we use (9.4). Its first term is the normalization defect $\mathbb{P}(K_{r_k} * \nu(0) < \tau)$ per unit volume. The mean of $K_{r_k} * \nu(0)$ is $\mathbb{E}\nu(\Omega_T)/V$, which tends to 1 exponentially fast by (8.6), so the defect has a fixed positive mean gap. For the cap rejection, K_{r_k} is symmetric and $\max(t_k, \tau) \geq \tau$, so $\pi_{i,k}(\Omega_T) \leq \tau^{-1}\nu_i Q_k^{\mu}(Y_i)$. Hence, configurationwise,

$$\sum_i \big(\pi_{i,k}(\Omega_T) - \delta_k\nu_i\big)_+ \leq \tau^{-1}\sum_i \nu_i\big(Q_k^{\mu}(Y_i) - \tau\delta_k\big)_+. \tag{10.4}$$

Lemma 10.2. *Let $k \geq 1$ be admissible, and let μ' be a source field, defined for every configuration, such that $\mu_{k-1} \leq \mu'$ in every configuration and $\|D_{z,e}\mu'\|_{TV} \leq \Delta$ for all (z,e). Let $Q = K_{r_k} * \mu'$ and assume*

$$\mathbb{E}Q(0) \leq \tau\delta_k/4, \qquad \Delta/v_k \leq \tau\delta_k/4.$$

Then

$$\frac{1}{V}\mathbb{E}\sum_i \big(\pi_{i,k}(\Omega_T) - \delta_k\nu_i\big)_+ \leq \frac{1}{\tau}\,\mathbb{E}\Big(Q(0) + \frac{\Delta}{v_k} - \tau\delta_k\Big)_+ \leq C_s\,\delta_k^{1-2s}\,\mathbb{E}|Q(0)-\mathbb{E}Q(0)|^{2s}. \tag{10.5}$$

We apply the lemma in two ways. In the preparatory stages $\mu' = \mu_0$ and $\Delta = B$, which is admissible by (9.8) and (8.10). In the late stages $\mu' = \mu_{k-1}$ and $\Delta = \Delta_k$, where

$$\Delta_k = (1 + 9(k-1))B. \tag{10.6}$$

Proof. By (10.4) and $\mu_{k-1} \leq \mu'$, the sum on the left is bounded, before expectation, by $\tau^{-1}\sum_i \nu_i(K_{r_k} * \mu'(Y_i) - \tau\delta_k)_+$. The Mecke formula evaluates this sum in the configuration $\eta + \delta_{(y,e)}$ with a planted site:

$$\frac{1}{V}\mathbb{E}\sum_i \nu_i(K_{r_k} * \mu'(Y_i) - \tau\delta_k)_+ = \frac{1}{V}\int_{\Omega_T} dy \int_0^{\infty} e^{-e}de\,\mathbb{E}\Big[\nu_y^{\eta+\delta_{(y,e)}}\,\big(K_{r_k} * \mu'^{\,\eta+\delta_{(y,e)}}(y) - \tau\delta_k\big)_+\Big].$$

Here $\nu_y^{\eta+\delta_{(y,e)}}$ is the core capacity of the planted site; by (8.5) it equals $\min\{e,m\} \leq e$, where m is the volume of the planted cell. The planted field differs from the unplanted one by at most Δ in total variation, so its ball load at y is at most $K_{r_k} * \mu'^{\,\eta}(y) + \Delta/v_k$. The resulting expression no longer depends on the fresh mark, and $\int_0^{\infty} e\,e^{-e}de = 1$. Stationarity in y gives the first inequality in (10.5). By the two assumptions, $\tau\delta_k - \Delta/v_k - \mathbb{E}Q(0) \geq \tau\delta_k/2$, and the first inequality of (10.3) gives the second, with $C_s = \tau^{-1}(2/\tau)^{2s-1}$. □

10.3 The two-phase induction

A late-type estimate already at stage 1 would require the mean condition $c_0 > \log\Lambda \geq \log 4$ (see (10.8)), whereas the cost of the first stage requires $\alpha c_0 < c_\chi/2 \approx 0.13$ (see (10.16)).

These are incompatible once $\alpha \geq c_\chi/(2\log 4)$. This is why the first k_* stages are treated by domination by the initial fields, which costs no locality factor.

The induction uses the following hypotheses at stage k:

$$\begin{aligned} m_F(0) \leq \delta_k/2, \quad m_H(0) \leq \tau\delta_k/4, \quad B/v_k \leq \tau\delta_k/4 \qquad (1 \leq k \leq k_*), \\ m_F(k-1) \leq \delta_k/2, \quad m_H(k-1) \leq \tau\delta_k/4, \quad \Delta_k/v_k \leq \tau\delta_k/4 \qquad (k > k_*), \end{aligned} \tag{10.7}$$

together with $\mathbb{E}\nu(\Omega_T)/V \geq 3/4$. They follow from the exponent conditions

$$\frac{c_\chi}{2} - 2b > 0, \qquad c_0 - 2b > 0, \qquad sc_0 - 4sb > 0, \qquad s(c_0 + k_* c_1) - s\log\Lambda - 4sb > 0, \tag{10.8}$$

and the dimension conditions (10.18) below.

Proposition 10.3. *Fix $\alpha > 0$, $s \in \{1, 2, 4, \dots\}$ with $s > \alpha$, parameters $b, c_0, c_1 > 0$ and an integer $k_* \geq 1$ satisfying* (10.8)*, and let A_α be chosen as in Section 8. There is d_0 such that for every $d \geq d_0$ satisfying* (10.18)*, every torus and every admissible stage $k \geq 1$, the hypotheses* (10.7) *hold and*

$$M_k \leq P_\alpha(d)\,(k+1)^{2s}\,2^{(2s-1)k}\exp\Big\{-d\big[sa_k - (4s-2)b - s\log\Lambda\,\mathbf{1}_{\{k>k_*\}}\big]\Big\}. \tag{10.9}$$

Proof. We first derive a bound on M_k from the hypotheses at stage k, separately in the two phases, and then show that the hypotheses propagate.

Preparatory stages, $1 \leq k \leq k_*$. By (9.8), $\zeta_{k-1} \leq \zeta_0$, so the positive overload is ordered:

$$m_F(k) = \mathbb{E}\big(K_{r_k} * \zeta_{k-1}(0) - \delta_k\big)_+ \leq \mathbb{E}\big(K_{r_k} * \zeta_0(0) - \delta_k\big)_+.$$

The mean of $K_{r_k} * \zeta_0(0)$ is $m_F(0) \leq \delta_k/2$, so (10.3) and Proposition 8.1 with $A = B_{r_k}$ bound this by $C_s(1+L)^s B^{2s}\delta_k^{1-2s}v_k^{-s}$. The normalization defect satisfies, by (10.3) with the mean gap $1/4$ and Proposition 8.1,

$$\mathbb{P}\big(K_{r_k} * \nu(0) < \tau\big) \leq C_s(1+L)^s B^{2s} v_k^{-s}. \tag{10.10}$$

The cap rejection is bounded by Lemma 10.2 with $\mu' = \mu_0$, $\Delta = B$, and Proposition 8.1. By (9.4) and $\delta_k \leq 1$,

$$M_k \leq C_s \frac{(1+L)^s B^{2s}}{\delta_k^{2s-1} v_k^s}, \qquad 1 \leq k \leq k_*. \tag{10.11}$$

Only positive overloads are compared here: the domination $\zeta_{k-1} \leq \zeta_0$, $\mu_{k-1} \leq \mu_0$ transfers the moment bounds of the initial fields to the overloads of the current residuals.

Late stages, $k > k_*$. The normalization defect is still bounded by (10.10), and $(1+L)^s \leq \Lambda^{sd}$ for large d. The target overload follows from (9.1), (10.3) with $m_F(k-1) \leq \delta_k/2$, and (10.1) for the current residual. The cap rejection follows from Lemma 10.2 with $\mu' = \mu_{k-1}$, the shift (10.6), and (10.1). Hence

$$M_k \leq C_s \frac{(k+1)^{2s} B^{2s} \Lambda^{sd}}{\delta_k^{2s-1} v_k^s}, \qquad k > k_*. \tag{10.12}$$

Since $B = 3L$, $\delta_k^{1-2s} = 2^{(k+3)(2s-1)} e^{2(2s-1)bd}$ and $v_k^{-s} = L^{-s} e^{-sa_k d}$, the bounds (10.11) and (10.12) give (10.9), after absorbing the powers of L into $P_\alpha(d)$. The exponent $(4s-2)b$ comes from δ_k^{1-2s}; the envelope B^{2s} is polynomial in d. Each late estimate contains the factor Λ^{sd} exactly once, through (10.1) at the current stage.

Propagation of the hypotheses. By (8.6),

$$\frac{M_0}{\delta_{k_*}} \leq C_{k_*} P_\alpha(d) e^{-(c_\chi/2 - 2b)d} \longrightarrow 0,$$

so the mean hypotheses of the preparatory stages hold for large d, as does $\mathbb{E}\nu(\Omega_T)/V = 1 - m_H(0) \geq 3/4$. The insertion shifts satisfy

$$\frac{B}{v_k\delta_k} \leq \frac{\Delta_k}{v_k\delta_k} \leq P_\alpha(d)(k+1)2^k \exp\bigl[-(a_k - 2b)d\bigr] \leq P_\alpha(d)e^{-(c_0-2b)d}\,2(k+1)\bigl(2e^{-c_1 d}\bigr)^{k-1}. \tag{10.13}$$

By (10.18), $2e^{-c_1 d} < 1/2$, so the right side is bounded uniformly in k by $C\,P_\alpha(d)e^{-(c_0-2b)d}$, which tends to 0. For the late mean hypotheses it suffices that $M_{k-1}/\delta_k \leq \tau/4$ for $k > k_*$. Dividing (10.9) by δ_{k+1} gives

$$\frac{M_k}{\delta_{k+1}} \leq P_\alpha(d)(k+1)^{2s}2^{2sk} \exp\Bigl\{-d\bigl[sa_k - 4sb - s\log\Lambda\,\mathbf{1}_{\{k>k_*\}}\bigr]\Bigr\}. \tag{10.14}$$

For $k \leq k_*$ the exponent is at least $sc_0 - 4sb$. At $k = k_*+1$ it equals $s(c_0 + k_*c_1) - 4sb - s\log\Lambda$, and each further stage increases it by sc_1. Since $2^{2s}e^{-sc_1 d} < 1/2$ by (10.18),

$$\sup_{k>k_*}(k+1)^{2s}2^{2sk}e^{-sc_1(k-k_*-1)d} \leq 2^{2s(k_*+1)} \sup_{k>k_*}(k+1)^{2s}2^{-(k-k_*-1)} \leq C_{s,k_*}.$$

By (10.8), the ratio (10.14) therefore tends to 0 as $d \to \infty$, uniformly in k and V. Hence, for d large, the hypotheses (10.7) hold at every admissible stage, by induction on k. Only the sum M_k enters these hypotheses. □

10.4 Choice of parameters

We now fix the parameters for every $\alpha > 0$ by one formula, and verify (10.8) together with the cost margins. Let $\bar\alpha = \max\{1, \alpha\}$ and let $s \in \{2, 4, 8, \dots\}$ be the smallest dyadic integer with $s > \alpha$, so that the moment order is $2s \geq 4$. Put

$$c_0 = \frac{c_\chi}{4\bar\alpha}, \qquad c_1 = \frac{(s-\alpha)c_0}{4\bar\alpha}, \qquad b = \frac{(s-\alpha)c_0}{200s}, \qquad g = s\log\Lambda, \tag{10.15}$$

with $\Lambda = 4 + 3/(e^{c_1} - 1)$ as in (9.10). Since $0 < s - \alpha < s$, we have $b < c_0/200 \leq c_\chi/800$, $\alpha c_0 \leq c_\chi/4$, $\alpha c_1 \leq (s-\alpha)c_0/4$ (because $\alpha/\bar\alpha \leq 1$) and $4sb = 0.02(s-\alpha)c_0$. Hence the first-stage margin κ_1, the basic late margin κ_2 and the conditions (10.8) not involving k_* satisfy

$$\begin{aligned}
&\kappa_1 = \frac{c_\chi}{2} - \alpha c_0 \geq \frac{c_\chi}{4} > 0,\\
&\kappa_2 = (s-\alpha)c_0 - \alpha c_1 - (4s-2)b \geq (s-\alpha)c_0\Bigl(1 - \frac14 - \frac{2}{100}\Bigr) = 0.73(s-\alpha)c_0 > 0,\\
&c_0 - 2b > 0.99\,c_0, \qquad sc_0 - 4sb = \bigl(s - 0.02(s-\alpha)\bigr)c_0 > 0, \qquad \frac{c_\chi}{2} - 2b \geq 0.495\,c_\chi.
\end{aligned} \tag{10.16}$$

Next choose a fixed integer k_* such that

$$(s-\alpha)k_*c_1 > g + 1, \qquad s(c_0 + k_*c_1) > g + 4sb + 1. \tag{10.17}$$

It exists because $(s-\alpha)c_1 > 0$ and $sc_1 > 0$, and it is independent of d and V. The second condition is the last one in (10.8). The integer k_* may be very large when $s - \alpha$ is small, and the dimension threshold below depends on it; the argument therefore gives existence of the limits but no effective rate in d. After k_* we take d so large that

$$2e^{-c_1 d} < \frac12, \qquad 2^{2s}e^{-sc_1 d} < \frac12, \qquad 2^{4s}e^{-(s-\alpha)c_1 d} < \frac12, \qquad p = \alpha d \geq 2. \tag{10.18}$$

Proposition 10.4. *Let the parameters be chosen by* (10.15), (10.17) *and* (10.18), *and let* $\kappa_\alpha = \min\{\kappa_1, \kappa_2\} > 0$. *For every sufficiently large* d *and every admissible final stage* k',

$$\frac{1}{V}\mathbb{E}\,\mathrm{Cost}\Big(\sum_{k\le k'}\mathcal{F}_k + \sum_{k\le k'}\sum_i \mathcal{H}_{i,k}\Big) \le P_\alpha(d)e^{-\kappa_\alpha d}, \tag{10.19}$$

where $\mathcal{H}_{i,k}$ *is transported to the site* Y_i.

Proof. Each rescue plan moves mass directly from a source point to its target atom, so the left side of (10.19) is the sum of the stage costs. A stage-k rescue moves mass over distances less than r_k, at power cost at most $\omega_d^\alpha r_k^p = v_k^\alpha = \lambda e^{\alpha a_k d}$. It moves at most the residual mass left after stage $k-1$, so its expected cost per unit volume is at most $\lambda e^{\alpha a_k d} M_{k-1}$. For $k = 1$, (8.6) bounds this by $\lambda P_\alpha(d)e^{-\kappa_1 d}$. For $k \ge 2$, inserting (10.9) at $k-1$ and using $sa_{k-1} - \alpha a_k = (s-\alpha)c_0 + (s-\alpha)(k-2)c_1 - \alpha c_1$ gives

$$\text{stage-}k\text{ cost} \le P_\alpha(d)\,\lambda\, k^{2s} 2^{(2s-1)(k-1)} \exp\Big\{-d\big[\kappa_2 + (s-\alpha)(k-2)c_1 - g\,\mathbf{1}_{\{k\ge k_*+2\}}\big]\Big\}. \tag{10.20}$$

Stage k_*+1 is the first stage whose mean bound (10.9) contains Λ^{sd}, and stage k_*+2 is the first whose cost uses such a bound. At $k = k_*+2$ the exponent in (10.20) equals $\kappa_2 + (s-\alpha)k_*c_1 - g > \kappa_2 + 1$ by (10.17). Apart from this one switch, consecutive bounds in (10.20) differ by a factor at most $2^{2s}\cdot 2^{2s-1}e^{-(s-\alpha)c_1 d} < 1/2$ by (10.18), where 2^{2s} bounds $((k+1)/k)^{2s}$. Summing the two geometric series (stages $2,\dots,k_*+1$ and stages $k \ge k_*+2$) bounds the total by $\lambda P_\alpha(d)(e^{-\kappa_1 d} + e^{-\kappa_2 d} + e^{-(\kappa_2+1)d})$, and $\lambda = L^\alpha$ is polynomial in d. □

11 Finite-volume completion

Proposition 10.3 bounds the mean residual mass M_k at every locally admissible stage, but the local rescues never exhaust the exceptional mass. We stop them at a stage k_V of order $\log V$ and send what remains across the torus by explicit product plans. An edge of the torus costs at most $C_{d,\alpha}V^\alpha$, so the residual mass per unit volume has to decay faster than $V^{-\alpha}$; this is where the moment order $2s > 2\alpha$ enters. The same moment order controls the probability of the event on which this completion fails.

Fix

$$\gamma = \frac{1}{2}\Big(1 + \frac{\alpha}{s}\Big) \in \Big(\frac{\alpha}{s}, 1\Big).$$

For each sufficiently large fixed d, let k_V be the largest stage with $r_{k_V} \le RV^{\gamma/d}$. At fixed d we have $k_V = O_d(\log V)$, and $k_V > k_*$ once V is large. By (9.10) the dependence balls of the stages $k \le k_V$ have radius at most Λr_{k_V}, and $\Lambda r_{k_V}/T \le \Lambda R V^{-(1-\gamma)/d} \to 0$; hence all these stages are locally admissible for V large. The maximality of k_V gives $a_{k_V} \ge (\gamma/d)\log V - c_1$, while $a_{k_V} \le (\gamma/d)\log V$ and $c_0 > 0$ give $2^{(2s-1)k_V} \le 2^{2s-1}V^{\gamma(2s-1)\log 2/(c_1 d)}$. Inserting these bounds in (10.9) yields

$$M_{k_V} \le C_{d,\alpha,s}(1+\log V)^{2s} V^{-\beta_{d,s}}, \qquad \beta_{d,s} = \gamma\Big(s - \frac{(2s-1)\log 2}{c_1 d}\Big), \tag{11.1}$$

and $\beta_{d,s} > \alpha$ for all sufficiently large fixed d. Indeed $\gamma s - \alpha = (s-\alpha)/2$, so $\beta_{d,s} > \alpha$ is equivalent to $(s-\alpha)c_1 d > 2\gamma(2s-1)\log 2$, and this holds because $\gamma < 1$ and the third condition of (10.18) gives $(s-\alpha)c_1 d > (4s+1)\log 2$.

11.1 Terminal plans on the good event

Define the good event

$$\mathcal{G}_V = \Big\{\nu(\Omega_T) \geq \frac{V}{2},\ \zeta_{k_V}(\Omega_T) \leq \frac{\delta V}{4},\ \mu_{k_V}(\Omega_T) \leq \frac{\delta\,\nu(\Omega_T)}{4}\Big\}.$$

Since $\nu(\Omega_T) \geq V/2$ implies $\delta\nu(\Omega_T)/4 \geq \delta V/8$,

$$\mathcal{G}_V^c \subset \{\nu(\Omega_T) < V/2\} \cup \{\zeta_{k_V}(\Omega_T) > \delta V/4\} \cup \{\mu_{k_V}(\Omega_T) > \delta V/8\},$$

and the failure bound below is applied to these three events, each of which is a deviation of a single global mass. On $\mathcal{G}_V$ we complete the two rescues by the far plans

$$\mathcal{F}^{\mathrm{far}}(dx, dy) = V^{-1} dx\, \zeta_{k_V}(dy), \qquad \mathcal{H}_i^{\mathrm{far}}(dx) = \frac{\nu_i}{\nu(\Omega_T)}\, \mu_{k_V}(dx). \tag{11.2}$$

The first has target marginal ζ_{k_V} and constant source density $\zeta_{k_V}(\Omega_T)/V \leq \delta/4$, while the local target rescues have total source density at most $\sum_k \delta_k = \delta/8$. The second exhausts the exceptional source μ_{k_V} and assigns to the core label i the mass $\nu_i \mu_{k_V}(\Omega_T)/\nu(\Omega_T) \leq \delta\nu_i/4$, while the local source rescues assign at most $\delta\nu_i/8$ to it. Write f for the total source density of all target rescues, local and far, and $\mathcal{H}_i$ for the total source-rescue flow to the label i. Then, for every configuration in $\mathcal{G}_V$,

$$f \leq \frac{3\delta}{8} \leq \delta, \qquad \mathcal{H}_i(\Omega_T) \leq \frac{3\delta}{8}\, \nu_i \leq \delta\nu_i. \tag{11.3}$$

Each far edge costs at most $C_{d,\alpha} V^\alpha$, with $C_{d,\alpha} = (\omega_d(\sqrt{d}/2)^d)^\alpha$, and this bound is pathwise. By (11.1), the expected cost of the two far plans, divided by V, is at most $C_{d,\alpha} V^\alpha M_{k_V} \leq C_{d,\alpha,s}(1 + \log V)^{2s} V^{\alpha - \beta_{d,s}}$, which tends to zero as $V \to \infty$.

11.2 The failure event

The three masses $\nu(\Omega_T)$, $\zeta_{k_V}(\Omega_T)$ and $\mu_{k_V}(\Omega_T)$ are global mass functionals with add-one envelope $O((1 + k_V)B + e)$. Their means are separated from the thresholds defining $\mathcal{G}_V$ by gaps of order V (at fixed d): indeed $\mathbb{E}\nu(\Omega_T) \geq 3V/4$ by (8.6), and $\mathbb{E}[\zeta_{k_V}(\Omega_T) + \mu_{k_V}(\Omega_T)] = VM_{k_V} = o(\delta V)$ by (11.1), so that $VM_{k_V} \leq \delta V/16$ for V large and the gaps are at least $\delta V/16$. Markov's inequality at order $2s$ and (10.2) therefore give

$$\mathbb{P}(\mathcal{G}_V^c) \leq C_{d,\alpha,s}(1 + k_V)^{2s} V^{-s} \leq C_{d,\alpha,s}(1 + \log V)^{2s} V^{-s}. \tag{11.4}$$

The total-mass ratio $\theta = S/V$ is correlated with $\mathcal{G}_V$. The total mark S is a compound Poisson variable with $\mathbb{E}\theta = 1$, and Lemma 8.3 gives $\mathbb{E}|\theta - 1|^{2s} \leq C_s V^{-s}$: the add-one derivative of S at (z, e) is e, so $S_2(S) = 2V$ and $\mathcal{T}_t(S) = t!\, V$, whence $\mathbb{E}|S - V|^{2s} \leq C_s(V^s + V)$. By Hölder's inequality,

$$\mathbb{E}[\theta \mathbf{1}_{\mathcal{G}_V^c}] \leq \mathbb{P}(\mathcal{G}_V^c) + \|\theta - 1\|_{2s}\, \mathbb{P}(\mathcal{G}_V^c)^{1 - 1/(2s)} \leq C_{d,\alpha,s}(1 + \log V)^{2s} V^{-s}. \tag{11.5}$$

On $\mathcal{G}_V^c$ we use the product plan

$$\Pi_V^{\mathrm{prod}} = V^{-1} dx \otimes \sum_i E_i \delta_{Y_i},$$

which has source $\theta\, dx$ and target $\sum_i E_i \delta_{Y_i}$. Its expected cost on $\mathcal{G}_V^c$, divided by V, is at most $C_{d,\alpha} V^\alpha \mathbb{E}[\theta \mathbf{1}_{\mathcal{G}_V^c}]$, and this tends to zero as $V \to \infty$ by (11.5), because $s > \alpha$.

The far plans (11.2) are not local, and neither is the completed source density built from them in Section 12. We therefore build the transport current from the local auxiliary densities ρ_k of (9.9), and treat only the small difference between the completed density and ρ_{k_V} by a deterministic primitive on the torus.

12 Density correction, flux and the atomic lift

On $\mathcal{G}_V$ the rescues and the far plans form a reference plan that serves each target Y_i exactly, with a source density ρ uniformly close to 1 but different from θ. We move $\theta\,dx$ to $\rho\,dx$ by a current of small L^p norm, built stage by stage from a periodic flux estimate, turn this current into a static coupling Γ_V, and compose Γ_V once with the labels of the reference plan. The outcome is the reduced-cost estimate (12.12), which compares the cost of the resulting plan with the dual value Ψ_V. Throughout, $p = \alpha d \geq 2$, and the identities of this section hold pathwise on $\mathcal{G}_V$, while moments and costs are annealed, as in Section 2.4.

12.1 The reference plan

Work on $\mathcal{G}_V$. Let $\mathcal{F}_i$ be the completed target-rescue flow into Y_i, local and far; it serves the exceptional target demand $(E_i - m_i)_+$, and the source densities of all $\mathcal{F}_i$ add up to $f \leq \delta$. Let $\mathcal{H}_i$ be the completed source-rescue flow to the label i, so that $\sum_i \mathcal{H}_i = \mu_0$ and $\mathcal{H}_i(\Omega_T) \leq \delta\nu_i$ by (11.3). Recall that the core source measure $\mu_i^{\mathrm{c}} = w_i \mathbf{1}_{C_i} dx$ has mass ν_i. For $\nu_i > 0$ set

$$\mathcal{D}_i = (1-f)(\mu_i^{\mathrm{c}} + \mathcal{H}_i), \qquad d_i = \mathcal{D}_i(\Omega_T), \qquad \beta_i = \frac{\nu_i}{d_i}, \qquad \Pi_i^{\mathrm{ref}} = \mathcal{F}_i + \beta_i \mathcal{D}_i.$$

For $\nu_i = 0$ we have $\mu_i^{\mathrm{c}} = 0$ and $\mathcal{H}_i = 0$; we set $\Pi_i^{\mathrm{ref}} = \mathcal{F}_i$ and omit these labels from all sums over $\mathcal{D}_i$. Since $f \leq \delta$ and $\mathcal{H}_i(\Omega_T) \leq \delta\nu_i$,

$$(1-\delta)\nu_i \leq d_i \leq (1+\delta)\nu_i, \qquad |\beta_i - 1| \leq 2\delta.$$

By (8.5), each target receives exactly $(E_i - m_i)_+ + \beta_i d_i = (E_i - m_i)_+ + \min\{E_i, m_i\} = E_i$. Before the reweighting by β_i, the source density of the plan is

$$f + (1-f)\Big(\sum_i w_i \mathbf{1}_{C_i} + \frac{d\mu_0}{dx}\Big) = 1.$$

Hence the reference plan $\sum_i \Pi_i^{\mathrm{ref}}$ has source density

$$\rho = 1 + \sum_i (\beta_i - 1)\frac{d\mathcal{D}_i}{dx}, \qquad \|\rho - 1\|_\infty \leq 2\delta, \qquad \fint_{\Omega_T} \rho = \theta, \tag{12.1}$$

where the last identity expresses that the total source mass equals $\sum_i E_i = \theta V$. The target marginal of the reference plan is exactly $\sum_i E_i \delta_{Y_i}$, but its source is $\rho\,dx$ rather than $\theta\,dx$.

The completed density ρ depends on the far plans and is not a local field. For the flux construction we use instead the auxiliary densities ρ_k of (9.9). They are built from the cumulative local flows $\mathcal{F}^{[k]}$ and $\mathcal{H}_i^{[k]}$, with f_k the source density of $\mathcal{F}^{[k]}$, through $\mathcal{D}_{i,k}$, $d_{i,k}$ and $\beta_{i,k}$, exactly as ρ is built from the completed flows. In general ρ_k is not the source density of the unfinished plan at stage k. What matters is that $\|\rho_k - 1\|_\infty \leq 2\delta$, by the argument just given, and that, by Lemma 9.5, ρ_k is a translation-covariant field which is point-local within a radius ℓ_k, with $\ell_k \leq C_{c_1} r_k$ for $k \geq 1$, while $\rho_0 \equiv 1$.

12.2 Labelled increments

For $j < k$ let $\Delta\mathcal{F} = \mathcal{F}^{[k]} - \mathcal{F}^{[j]} \geq 0$ and $\Delta\mathcal{H}_i = \mathcal{H}_i^{[k]} - \mathcal{H}_i^{[j]} \geq 0$. Directly from the definitions,

$$\mathcal{D}_{i,k} - \mathcal{D}_{i,j} = (1 - f_k)\Delta\mathcal{H}_i - (f_k - f_j)(\mu_i^{\mathrm{c}} + \mathcal{H}_i^{[j]}). \tag{12.2}$$

Since $\sum_i (d\mu_i^{\mathrm{c}}/dx + d\mathcal{H}_i^{[j]}/dx) \leq 1$, summing total variations over the labels gives

$$\sum_i \|\mathcal{D}_{i,k} - \mathcal{D}_{i,j}\|_{TV} \leq \sum_i \Delta\mathcal{H}_i(\Omega_T) + \Delta\mathcal{F}(\Omega_T \times \Omega_T) \leq \mu_j(\Omega_T) + \zeta_j(\Omega_T),$$

the last inequality because the stages after j only serve the residual mass left at stage j. The normalizers satisfy $\beta_{i,k}d_{i,k} = \nu_i = \beta_{i,j}d_{i,j}$, hence

$$d_{i,k}|\beta_{i,k} - \beta_{i,j}| = \beta_{i,j}|d_{i,k} - d_{i,j}| \le C\|\mathcal{D}_{i,k} - \mathcal{D}_{i,j}\|_{TV}.$$

Writing

$$\rho_k - \rho_j = \sum_i (\beta_{i,k} - \beta_{i,j})\frac{d\mathcal{D}_{i,k}}{dx} + \sum_i (\beta_{i,j} - 1)\frac{d(\mathcal{D}_{i,k} - \mathcal{D}_{i,j})}{dx}$$

and using $|\beta_{i,j}-1| \le 2\delta$, we obtain $\int_{\Omega_T} |\rho_k-\rho_j| \le C\sum_i \|\mathcal{D}_{i,k}-\mathcal{D}_{i,j}\|_{TV}$. We take expectations, and for consecutive stages we interpolate with the pointwise bound $|\rho_k - \rho_{k-1}| \le 4\delta$:

$$\mathbb{E}\fint_{\Omega_T} |\rho_k - \rho_j| \le CM_j, \qquad \|\rho_k - \rho_{k-1}\|_{L^p(\mathbb{P}\otimes dx/V)} \le C\delta^{1-1/p}M_{k-1}^{1/p}. \tag{12.3}$$

The identity (12.2), applied with the far plans as increments, compares the completed density with ρ_{k_V}:

$$\mathbb{E}\Big[\mathbf{1}_{\mathcal{G}_V}\fint_{\Omega_T} |\rho - \rho_{k_V}|\Big] \le CM_{k_V}, \qquad \mathbf{1}_{\mathcal{G}_V}|\rho - \rho_{k_V}| \le 4\delta. \tag{12.4}$$

12.3 A periodic flux estimate

The local currents come from the following estimate, proved in Appendix B. Currents of the form $-\int_0^\infty \nabla P_t \cdot dt$, estimated through heat-kernel bounds and moment inequalities for sums of independent variables, are standard in the matching problem [3, 20]; the point of Theorem 12.1 is the explicit dependence $Cp\sqrt{d}$ of the leading constant on p and d. We write $P_t = e^{t\Delta}$ for the heat semigroup on Ω_T.

Theorem 12.1. *Let $d \ge 3$, $p \ge 2$, $\ell > 0$, and let $\Omega_T = (\mathbb{R}/T\mathbb{Z})^d$, $V = T^d$, with $T \ge 640\ell$. Let $u = (u(x))_{x\in\Omega_T}$ be a jointly measurable, translation-covariant field of the marked Poisson process such that $u(x)$ depends only on the marked points in $B_\ell(x)$. Put $v = u - \mathbb{E}u(0)$ and $A = \|v(0)\|_{L^p(\mathbb{P})} < \infty$. Then the periodic current*

$$\mathbf{j} = -\int_0^\infty \nabla P_t v \, dt$$

converges in $L^p(\mathbb{P}\otimes dx/V;\mathbb{R}^d)$, satisfies $\operatorname{div}\mathbf{j} = u - \fint_{\Omega_T} u$ in the sense of distributions, and

$$\|\mathbf{j}\|_{L^p(\mathbb{P}\otimes dx/V)} \le Cp\sqrt{d}\,\ell A + C_{d,p}T(\ell/T)^{d/2}A. \tag{12.5}$$

Here C is a numerical constant, independent of d, p, T, ℓ, and $C_{d,p}$ depends only on d and p.

The current removes the random spatial average $\fint_{\Omega_T} u$, which differs from the ensemble mean $\mathbb{E}u(0)$. Only the second term of (12.5), which vanishes as $T/\ell \to \infty$ at fixed d, has a constant depending on (d, p).

12.4 Summation of the local currents

Since $\rho_0 \equiv 1$, no current is needed at stage 0. We apply Theorem 12.1 with $p = \alpha d$, for $1 \le k \le k_V$, to $u = \rho_k - \rho_{k-1}$, with locality radius $\ell_k \le C_{c_1}r_k$ and, by stationarity and (12.3), $A \le C\delta^{1-1/p}M_{k-1}^{1/p}$. Since $\ell_k \le \ell_{k_V} \le C_{c_1}RV^{\gamma/d}$, the condition $T \ge 640\ell_k$ holds for V large at fixed d. Let $\mathbf{j}_{\mathrm{loc}}$ be the sum of these k_V currents, so that $\operatorname{div}\mathbf{j}_{\mathrm{loc}} = \rho_{k_V} - \fint_{\Omega_T}\rho_{k_V}$. Since $r_k = Re^{a_k}$, the leading terms of (12.5), divided by R, add up to at most

$$C_{c_1}p\sqrt{d}\,\delta^{1-1/p}\sum_{k=1}^{k_V} e^{a_k}M_{k-1}^{1/p}. \tag{12.6}$$

We show that the series in (12.6) is bounded by a constant depending only on α, uniformly in V. Bound each M_{k-1} by (10.9). For late stages, the ratio of the resulting consecutive terms is at most

$$\Big(\frac{k+1}{k}\Big)^{2s/p} 2^{(2s-1)/p} \exp\Big[-\Big(\frac{s}{\alpha}-1\Big)c_1\Big], \tag{12.7}$$

and hence at most $2^{(4s-1)/p}\exp[-(s/\alpha-1)c_1]$. Because $s>\alpha$, this is bounded by a constant smaller than one for all sufficiently large fixed d. At the passage from the preparatory to the late regime, the bound (10.9) acquires the factor Λ^{sd} once; after the pth root this is the factor $\Lambda^{s/\alpha}$, which does not depend on d. The finitely many terms with $k \le k_*+1$ are bounded by constants depending only on α, since k_* is fixed by (10.17) before d is chosen. Hence the series is bounded by C_α (here $P_\alpha(d)^{1/p} \le C_\alpha$ because $p=\alpha d$), and (12.6) is at most $P_\alpha(d)\delta^{1-1/p} \le P_\alpha(d)e^{-bd}$, since $\delta^{-1/p}=e^{2b/\alpha}$ is a constant.

The periodic remainders are negligible. For $k \le k_V$,

$$T(\ell_k/T)^{d/2} \le C_{d,\alpha,s} V^{1/d-(1-\gamma)/2},$$

which tends to zero when $d > 2/(1-\gamma)$. Multiplied by the corresponding $A \le C\delta$ and summed over the $k_V = O_d(\log V)$ stages, these terms vanish as $V\to\infty$ at fixed d. Their (d,p)-dependent constants therefore do not enter the leading estimate.

12.5 The nonlocal remainder and the static coupling

Put $e_V = \rho - \rho_{k_V}$ on $\mathcal{G}_V$ and $e_V = 0$ on $\mathcal{G}_V^c$. By (12.4) and interpolation, $\|e_V\|_p \le C\delta^{1-1/p} M_{k_V}^{1/p}$. Integrating coordinate by coordinate on the torus gives a periodic current $\mathbf{j}_{\mathrm{far}}$, a deterministic linear function of e_V, as follows. Put $g = e_V - \fint_{\Omega_T} e_V$, let $g_0 = g$ and let g_i be the average of g_{i-1} over the ith coordinate. Set $\mathbf{j}_{\mathrm{far}} = (F_1,\dots,F_d)$ with $F_i(x) = \int_0^{x_i}(g_{i-1}-g_i)\,dx_i'$, which is periodic because $g_{i-1}-g_i$ has zero mean in x_i. Then $\operatorname{div}\mathbf{j}_{\mathrm{far}} = g_0 - g_d = g$ since $g_d = \fint g = 0$, and $\mathbf{j}_{\mathrm{far}}$ satisfies

$$\operatorname{div}\mathbf{j}_{\mathrm{far}} = e_V - \fint_{\Omega_T} e_V, \qquad \|\mathbf{j}_{\mathrm{far}}\|_p \le 2dT\Big\|e_V - \fint_{\Omega_T} e_V\Big\|_p.$$

By (11.1), with $T = V^{1/d}$ and $p=\alpha d$,

$$\frac{\|\mathbf{j}_{\mathrm{far}}\|_p}{R} \le C_{d,\alpha,s} V^{(\alpha-\beta_{d,s})/(\alpha d)}(1+\log V)^{2s/p} \longrightarrow 0 \qquad (V\to\infty), \tag{12.8}$$

since $\beta_{d,s}>\alpha$. On $\mathcal{G}_V$, the current $\mathbf{j}_{\mathrm{loc}}+\mathbf{j}_{\mathrm{far}}$ has divergence $\rho - \fint\rho = \rho-\theta$ by (12.1). Let $\mathbf{j}_V = -(\mathbf{j}_{\mathrm{loc}}+\mathbf{j}_{\mathrm{far}})$, so that $\operatorname{div}\mathbf{j}_V = \theta-\rho$ on $\mathcal{G}_V$. Combining (12.6), the vanishing periodic remainders and (12.8), we obtain

$$U_d := \limsup_{V\to\infty} \frac{\|\mathbf{1}_{\mathcal{G}_V}\mathbf{j}_V\|_{L^p(\mathbb{P}\otimes dx/V)}}{R} \le P_\alpha(d)e^{-bd}. \tag{12.9}$$

The divergence identity holds pathwise on $\mathcal{G}_V$; the bound on U_d is annealed.

On $\mathcal{G}_V$ the densities θ and ρ take values in $[1-C\delta, 1+C\delta]$ and have the same total mass. The interpolation $\bar\rho_t = (1-t)\theta + t\rho$, $t\in[0,1]$, together with the time-independent current $\mathbf{j}_V$, solves $\partial_t\bar\rho_t + \operatorname{div}\mathbf{j}_V = (\rho-\theta)+(\theta-\rho) = 0$, and $\bar\rho_t \ge 1-C\delta$. The velocity $\mathbf{v}_t = \mathbf{j}_V/\bar\rho_t$ satisfies $\int_0^1\int|\mathbf{v}_t|^p\bar\rho_t\,dx\,dt<\infty$. For $p>1$, the characterization of absolutely continuous curves in the Wasserstein space [2, Theorem 8.3.1] (stated on $\mathbb{R}^d$; its proof applies verbatim on the flat torus), applied after normalizing the common mass, gives $W_p(\theta\,dx,\rho\,dx) \le \int_0^1\|\mathbf{v}_t\|_{L^p(\bar\rho_t)}dt$.

This is the Benamou–Brenier bound (see also [28, Chapter 5]), and Jensen's inequality in t gives

$$\mathrm{OT}_{\mathrm{dist}^p}(\theta\,dx, \rho\,dx) \le \int_0^1 \int_{\Omega_T} |\mathbf{j}_V|^p \bar{\rho}_t^{1-p}\,dx\,dt \le (1 - C\delta)^{1-p} \int_{\Omega_T} |\mathbf{j}_V|^p\,dx. \tag{12.10}$$

We take for Γ_V an optimal coupling for $\mathrm{OT}_{\mathrm{dist}^p}(\theta\,dx, \rho\,dx)$, chosen measurably in the configuration: on the compact torus the cost dist^p is continuous, so by [30, Corollary 5.22], applied after normalizing the common mass, there is a Borel map from pairs of marginals to optimal couplings, which we compose with the jointly Borel marginals $(\theta\,dx, \rho\,dx)$ on $\mathcal{G}_V$. By (12.10), the pth root of the annealed normalized cost of Γ_V, divided by R, is at most $(1 - C\delta)^{(1-p)/p} \|\mathbf{1}_{\mathcal{G}_V} \mathbf{j}_V\|_p / R$, and $(1 - C\delta)^{(1-p)/p} \to 1$ as $d \to \infty$.

12.6 The typed lift

On $\mathcal{G}_V$ the reference plan has two types of edges into each label i: core edges $\beta_i(1 - f)\mu_i^{\mathrm{c}}$, and exceptional edges $\mathcal{F}_i + \beta_i(1 - f)\mathcal{H}_i$. Write $w_{i,\star}(y)$, $\star \in \{\mathrm{core}, \mathrm{exc}\}$, for their source densities, so that $\sum_{i,\star} w_{i,\star} = \rho$. Labels are physical labels, enumerated as in Section 2.4, and all coefficients and densities are jointly Borel. Since $\rho \ge 1 - 2\delta$, the conditional weights

$$P_y(i, \star) = \frac{w_{i,\star}(y)}{\rho(y)}$$

are well defined. We compose them once with Γ_V:

$$\Pi_V(dx, dy, i, \star) = \Gamma_V(dx, dy)\, P_y(i, \star).$$

The x-marginal of Π_V is $\theta\,dx$ and its y-marginal is $\rho\,dy$, so its $(y, i, \star)$-marginal is the typed reference plan $w_{i,\star}(y)\,dy$, and its label marginal is $\sum_i E_i \delta_{Y_i}$. In particular the lift preserves the mass and the pth moment of the exceptional reference edges. On $\mathcal{G}_V^c$ we use Π_V^{prod}.

We compare the cost of Π_V with the dual value

$$\Psi_V = \theta \int_{\Omega_T} \phi\,dx + \sum_i E_i h_i.$$

Define the slack $\sigma_i(x) = c_{d,\alpha}(x, Y_i) - h_i - \phi(x)$, which is nonnegative by the definition of ϕ. Since the (x, i)-marginal of Π_V couples $\theta\,dx$ with $\sum_i E_i \delta_{Y_i}$,

$$\mathrm{Cost}(\Pi_V) - \Psi_V = \int \sigma_i(x)\,d\Pi_V(x, y, i, \star).$$

Recall that $c_{d,\alpha}(x, Y_i) = \lambda(\mathrm{dist}(x, Y_i)/R)^p$. The function $c_{d,\alpha}(\cdot, Y_i) - h_i$ is smaller than a only where $\mathrm{dist}(\cdot, Y_i) < R$, and there its gradient is at most $p\lambda/R$; hence $\mathrm{Lip}\,\phi \le p\lambda/R$. More precisely, if $\phi(y) < a$ is attained at j and $g_j = c_{d,\alpha}(\cdot, Y_j) - h_j$, then $\phi(x) - \phi(y) \le \min\{a, g_j(x)\} - g_j(y)$, and $\min\{a, g_j\}$ is a nondecreasing function of $\mathrm{dist}(\cdot, Y_j)$ with slope at most $p\lambda/R$; if $\phi(y) = a$ there is nothing to prove. For a lifted triple (x, y, i) let $\xi = \mathrm{dist}(y, Y_i)/R$ and $t = \mathrm{dist}(x, y)/R$. The triangle inequality gives

$$\sigma_i(x) \le \sigma_i(y) + \lambda\big[(\xi + t)^p - \xi^p + pt\big]. \tag{12.11}$$

On core edges $\sigma_i(y) = 0$ and $\xi \le 1$, so the right-hand side is at most $\lambda[(1 + t)^p - 1 + pt]$, since $p \ge 1$. On exceptional edges $\sigma_i(y) \le \lambda(\xi^p + 1)$, because $|h_i| \le a$ and $|\phi| \le a$, and the term $-\xi^p$ cancels; the right-hand side is at most $\lambda[(\xi + t)^p + 1 + pt]$.

Let $\mathcal{L}_V$ be the annealed, volume-normalized lifted measure on successful configurations: for Borel $g \geq 0$,

$$\int g \, d\mathcal{L}_V = \frac{1}{V}\mathbb{E}\Big[\mathbf{1}_{\mathcal{G}_V} \sum_{i,\star} \int g(x, y, i, \star)\, \Gamma_V(dx, dy) P_y(i, \star)\Big].$$

Its total mass is $\mathbb{E}[\theta \mathbf{1}_{\mathcal{G}_V}] \leq 1$. Define the deterministic quantities

$$U_V^p = \int t^p \, d\mathcal{L}_V, \qquad Y_V^p = \int_{\{\star=\mathrm{exc}\}} \xi^p \, d\mathcal{L}_V, \qquad Z_V = \mathcal{L}_V(\{\star = \mathrm{exc}\}).$$

We integrate (12.11) against $\mathcal{L}_V$. On the core submeasure, extended to mass one by a dummy mass at $t = 0$ (where $(1+t)^p - 1$ vanishes), Minkowski's inequality bounds the integral of $(1+t)^p - 1$ by $(1+U_V)^p - 1$. On the exceptional submeasure, Minkowski's inequality bounds the integral of $(\xi + t)^p$ by $(Y_V + U_V)^p$. The term pt is integrated once over $\mathcal{L}_V$ and contributes at most pU_V. This proves

$$\frac{1}{V}\mathbb{E}[\mathbf{1}_{\mathcal{G}_V}(\mathrm{Cost}(\Pi_V) - \Psi_V)] \leq \lambda\big[(1+U_V)^p - 1 + pU_V + (Y_V + U_V)^p + Z_V\big]. \tag{12.12}$$

Applying the triangle inequality to the full lifted measure, instead of separating the two types, would cost a factor 2^p on a mass of order one.

It remains to bound the right-hand side. By (12.10) and (12.9), $\limsup_{V\to\infty} U_V \leq (1 - C\delta)^{(1-p)/p} U_d$. By (8.6), $Z_d := \limsup_{V\to\infty} Z_V \leq P_\alpha(d) e^{-c_\chi d/2}$. The exceptional reference edges are the local rescue edges and the far edges; their cost is bounded by (10.19) and by the vanishing far cost of Section 11. Since the lift preserves the exceptional reference-edge moment,

$$Y_d^p := \limsup_{V\to\infty} Y_V^p \leq P_\alpha(d) e^{-\kappa_\alpha d}.$$

Thus $\limsup_{d\to\infty} Y_d \leq e^{-\kappa_\alpha/\alpha} < 1$, while $pU_d \to 0$ exponentially in d. The strict inequality $\limsup_d Y_d < 1$ keeps $(Y_d + U_d)^p$ exponentially small. Since $\lambda = (A_\alpha d)^\alpha$ grows only polynomially in d, each term on the right of (12.12) tends to zero after the limit $V \to \infty$ at fixed d and then $d \to \infty$. As $\sigma_i \geq 0$, this gives

$$\lim_{d\to\infty} \limsup_{V\to\infty} \frac{1}{V}\mathbb{E}[\mathbf{1}_{\mathcal{G}_V}(\mathrm{Cost}(\Pi_V) - \Psi_V)] = 0. \tag{12.13}$$

13 Proof of Theorem 2.5

We first prove the lower bound for the Poissonized cost $\mathcal{P}_{d,\alpha}(V)$ of Section 3; it uses only weak duality and the trial prices of Section 4.2. The upper bound combines (12.13) with the failure bound (11.5). Lemmas 3.3 and 3.4 then return to deterministic sample sizes. Throughout, $q = 1/\alpha$, $p = \alpha d$ and $\lambda = \lambda(d) = (A_\alpha d)^\alpha$.

Lower bound. Since $\phi(x) + h_i \leq c_{d,\alpha}(x, Y_i)$ for all x and i, every coupling of $\theta\, dx$ and $\sum_i E_i \delta_{Y_i}$ has cost at least Ψ_V. Hence, for every configuration,

$$\mathrm{OT}_{c_{d,\alpha}}\Big(\theta\, dx, \sum_i E_i \delta_{Y_i}\Big) \geq \Psi_V.$$

The dual value has mean

$$\lim_{V\to\infty} \frac{1}{V}\mathbb{E}\Psi_V = \mathbb{E}X + \mathbb{E}[Eh(E)] = C_{q,\lambda}^{\mathrm{trial}}, \qquad E \sim \mathrm{Exp}(1). \tag{13.1}$$

Indeed, ϕ is stationary and $\phi(0)$ has the capped law of X, including the atom at a; the Campbell–Mecke formula gives $V^{-1}\mathbb{E}\sum_i E_i h(E_i) = \mathbb{E}[Eh(E)]$; and since $|\phi| \le \lambda/2$, replacing θ by 1 in $\mathbb{E}[\theta \mathbin{⨍} \phi]$ costs at most $\frac{\lambda}{2}\mathbb{E}|\theta - 1| = O_\lambda(V^{-1/2})$. Therefore, for every d,

$$\liminf_{V\to\infty} \mathcal{P}_{d,\alpha}(V) \ge C^{\mathrm{trial}}_{q,\lambda(d)}. \tag{13.2}$$

Upper bound. For sufficiently large fixed d we use Π_V on $\mathcal{G}_V$ and Π_V^{prod} on $\mathcal{G}_V^c$. Then

$$0 \le \mathrm{OT}_{c_{d,\alpha}} - \Psi_V \le \mathbf{1}_{\mathcal{G}_V}\big(\mathrm{Cost}(\Pi_V) - \Psi_V\big) + \mathbf{1}_{\mathcal{G}_V^c}\big(\mathrm{Cost}(\Pi_V^{\mathrm{prod}}) - \Psi_V\big). \tag{13.3}$$

Since $\mathrm{Cost}(\Pi_V^{\mathrm{prod}}) \le C_{d,\alpha} V^\alpha S$ and $|\Psi_V| \le \lambda S$, the expectation of the second term, divided by V, is at most $(C_{d,\alpha}V^\alpha + \lambda)\mathbb{E}[\theta \mathbf{1}_{\mathcal{G}_V^c}]$, which tends to zero by (11.5) because $s > \alpha$. Taking expectations in (13.3), dividing by V and using (13.1), we get

$$\limsup_{V\to\infty} \mathcal{P}_{d,\alpha}(V) \le C^{\mathrm{trial}}_{q,\lambda(d)} + \limsup_{V\to\infty} \frac{1}{V}\mathbb{E}\big[\mathbf{1}_{\mathcal{G}_V}\big(\mathrm{Cost}(\Pi_V) - \Psi_V\big)\big].$$

By (12.13) the last term tends to zero as $d \to \infty$. Together with (13.2) and Proposition 4.4, which gives $C^{\mathrm{trial}}_{q,\lambda(d)} \to \beta_{\mathrm{bip}}(q)$, this proves

$$\lim_{d\to\infty} \liminf_{V\to\infty} \mathcal{P}_{d,\alpha}(V) = \lim_{d\to\infty} \limsup_{V\to\infty} \mathcal{P}_{d,\alpha}(V) = \beta_{\mathrm{bip}}(q) = \mathcal{C}(\alpha). \tag{13.4}$$

De-Poissonization. By (3.10), at every fixed d, the lim sup as $n \to \infty$ of $(n\omega_d)^\alpha \mathbb{E}W_p^p(\mu_n, \nu_n)$ is at most $\limsup_V \mathcal{P}_{d,\alpha}(V)$, and its lim inf is at least $\liminf_V \mathcal{P}_{d,\alpha}(V)$. By (13.4), both iterated limits equal $\mathcal{C}(\alpha)$, which proves Theorem 2.5.

A Proof of Xu's identity

Theorem 3.1 is due to Xu [33, Theorem 2 and Section 2.1]; we include a self-contained proof for completeness, including the combinatorial lemma of [33, Appendix A.1]. Under a genericity assumption on the costs, removing one source reduces the expected matching cost to $m\,\mathbb{E}\,\mathcal{T}_c(\mu, A)$ for a random mass vector A (Lemma A.2). A finite combinatorial bijection (Lemma A.3) shows that A has the law $\mathrm{Dir}(1, \dots, 1)$, whatever the costs. A small random perturbation then removes the genericity assumption. The appendix ends with the proof of Corollary 3.2.

Throughout, $(\mathsf{X}, \mathcal{A}, \mu)$ is a probability space, $m \ge 1$, and $c_1, \dots, c_m \in L^1(\mu)$ are real-valued. For w in the simplex $\Delta_{m-1} = \{w \in [0,\infty)^m : \sum_j w_j = 1\}$, an *allocation with masses* w is a family of measurable functions $a_j : \mathsf{X} \to [0,1]$ with $\sum_j a_j = 1$ μ-a.e. and $\int a_j \,\mathrm{d}\mu = w_j$ for every j. Thus $\mathcal{T}_c(\mu, w)$ is the infimum of $\int \sum_j c_j a_j \,\mathrm{d}\mu$ over all allocations with masses w. We write $\Delta^\circ_{m-1} = \{w \in \Delta_{m-1} : w_j > 0 \text{ for all } j\}$ for the relative interior of the simplex; it has full $\mathrm{Dir}(1, \dots, 1)$ measure.

A.1 Convexity of the semidiscrete value

Lemma A.1. *The function $w \mapsto \mathcal{T}_c(\mu, w)$ is convex on Δ_{m-1} and satisfies*

$$|\mathcal{T}_c(\mu, w)| \le \int \max_{1\le j\le m} |c_j| \,\mathrm{d}\mu < \infty \qquad (w \in \Delta_{m-1}). \tag{A.1}$$

It is continuous on Δ°_{m-1} and Borel measurable on Δ_{m-1}.

Proof. The constant allocation $a_j \equiv w_j$ has masses w, even when μ has atoms, so the infimum is taken over a nonempty family. The cost of every allocation lies between $-\int \max_j |c_j| \, \mathrm{d}\mu$ and $\int \max_j |c_j| \, \mathrm{d}\mu$, which gives (A.1).

If (a_j) and (a'_j) are allocations with masses w and w' and $t \in [0,1]$, then $(ta_j + (1-t)a'_j)$ is an allocation with masses $tw + (1-t)w'$, and its cost is the same convex combination of the two costs. Taking (a_j) and (a'_j) arbitrarily close to optimal gives

$$\mathcal{T}_c(\mu, tw + (1-t)w') \le t\mathcal{T}_c(\mu, w) + (1-t)\mathcal{T}_c(\mu, w').$$

A convex function that is finite on a relatively open convex subset of $\mathbb{R}^m$ is continuous there. The simplex is the disjoint union of the $2^m - 1$ sets $\{w \in \Delta_{m-1} : w_j > 0 \text{ exactly for } j \in Z\}$, $\emptyset \ne Z \subset \{1, \dots, m\}$, the relative interiors of its faces. Each of them is a Borel, relatively open convex set, and Δ°_{m-1} is the one with $Z = \{1, \dots, m\}$. By convexity and (A.1), $\mathcal{T}_c(\mu, \cdot)$ is continuous on each of them; hence it is continuous on Δ°_{m-1} and Borel on Δ_{m-1}. □

A.2 The leave-one-source-out allocation

In this subsection and the next we assume $m \ge 2$ and the genericity condition

$$\mu\{x : c_j(x) - c_l(x) = t\} = 0 \qquad \text{for all } j \ne l \text{ and } t \in \mathbb{R}. \tag{A.2}$$

Let $X_1, \dots, X_m$ be i.i.d. with law μ, and call the first $m-1$ sources the *base*. For $j \in \{1, \dots, m\}$ define the deletion cost

$$D_j = \min_{\pi : \{1,\dots,m-1\} \to \{1,\dots,m\}\setminus\{j\} \text{ bijective}} \sum_{i=1}^{m-1} c_{\pi(i)}(X_i). \tag{A.3}$$

An additional source at x has optimal assignment cost $\min_j\{c_j(x) + D_j\}$. Let

$$\iota(x) = \text{the smallest index } j \text{ minimizing } c_j(x) + D_j, \qquad A_j = \mu\{x : \iota(x) = j\}. \tag{A.4}$$

Conditional on the base, a tie between $j \ne l$ at x means $c_j(x) - c_l(x) = D_l - D_j$, so by (A.2) the minimizer is unique for μ-a.e. x and the tie-breaking convention is immaterial. Each D_j is a minimum of finitely many measurable functions of the base, so $(X_1, \dots, X_{m-1}, x) \mapsto \iota(x)$ is jointly measurable. By Fubini's theorem, $A = (A_1, \dots, A_m)$ is a measurable function of the base with values in Δ_{m-1}, and $\mathcal{T}_c(\mu, A)$ is a bounded random variable by Lemma A.1.

Lemma A.2. *For almost every base, the allocation $a_j = \mathbf{1}_{\{\iota = j\}}$ attains $\mathcal{T}_c(\mu, A)$, and*

$$\mathbb{E}_X \min_{\sigma \in S_m} \sum_{i=1}^{m} c_{\sigma(i)}(X_i) = m\,\mathbb{E}\,\mathcal{T}_c(\mu, A). \tag{A.5}$$

Proof. Fix the base and let (a_j) be any allocation with masses A. Since $\sum_j a_j = 1$ a.e., the minimizing property of ι gives

$$\begin{aligned}\int c_{\iota(x)}(x) \, \mathrm{d}\mu(x) &= \int \min_j (c_j(x) + D_j) \, \mathrm{d}\mu(x) - \sum_j D_j A_j \\ &\le \int \sum_j (c_j(x) + D_j) a_j(x) \, \mathrm{d}\mu(x) - \sum_j D_j A_j = \int \sum_j c_j(x) a_j(x) \, \mathrm{d}\mu(x).\end{aligned}$$

The indicator allocation has masses A, hence it attains $\mathcal{T}_c(\mu, A)$ and $\int c_{\iota(x)}(x) \, \mathrm{d}\mu(x) = \mathcal{T}_c(\mu, A)$.

The optimal full assignment is almost surely unique. Indeed, if $\sigma \neq \sigma'$, pick i_0 with $\sigma(i_0) \neq \sigma'(i_0)$; after conditioning on all sources except X_{i_0}, the difference of the two assignment costs equals $c_{\sigma(i_0)}(X_{i_0}) - c_{\sigma'(i_0)}(X_{i_0})$ plus a constant, and it vanishes with probability zero by (A.2). Let σ^* be the optimal assignment. Optimizing first over the target of X_m shows that the optimal cost is $c_{\iota(X_m)}(X_m) + D_{\iota(X_m)}$. It is therefore attained by sending X_m to $\iota(X_m)$ and the base optimally to the other labels, and by uniqueness $\sigma^*(m) = \iota(X_m)$ almost surely. Since X_m is independent of the base, Fubini's theorem gives

$$\mathbb{E}\big[c_{\sigma^*(m)}(X_m) \mid X_1, \dots, X_{m-1}\big] = \int c_{\iota(x)}(x)\, \mathrm{d}\mu(x) = \mathcal{T}_c(\mu, A).$$

All these quantities are integrable, as $|c_{\sigma^*(i)}(X_i)| \leq \max_j |c_j(X_i)|$. The identity $\sigma^*(m) = \iota(X_m)$ holds for any optimal assignment; uniqueness is used only to make σ^* a well-defined function of the sources. The sources are exchangeable and σ^* is a permutation-equivariant function of them, so the expected cost $\mathbb{E}\, c_{\sigma^*(i)}(X_i)$ contributed by source i does not depend on i. Summing the m contributions proves (A.5). □

A.3 The combinatorial bijection

The law of A does not depend on the costs. This rests on the following finite-dimensional statement, Xu's combinatorial lemma [33, Appendix A.1]. Fix $k \geq 0$, set $N = m - 1 + k$, and consider an $N \times m$ real matrix $C = (c_{ij})$, viewed as a cost on the edges of the complete bipartite graph between sources $\{1, \dots, N\}$ and labels $\{1, \dots, m\}$. We call C *cycle-generic* if no alternating sum of its entries along a simple cycle of this graph vanishes. For $I \subset \{1, \dots, N\}$ with $|I| = m - 1$, define $D_j(I)$ as in (A.3), using the rows indexed by I. Each $i \notin I$ chooses the minimizing label

$$\iota_I(i) = \operatorname*{argmin}_j \{c_{ij} + D_j(I)\}, \qquad K_j(I) = \#\{i \notin I : \iota_I(i) = j\},$$

with the smallest-index convention in case of ties. The label of $i \notin I$ depends only on row i and on the rows in I, not on the other sources outside I. The counts satisfy $\sum_j K_j(I) = k$.

Lemma A.3. *Let C be cycle-generic. Then every minimizer $\iota_I(i)$ is unique, and*

$$I \longmapsto (K_1(I), \dots, K_m(I)) \tag{A.6}$$

is a bijection from the $(m-1)$-element subsets of $\{1, \dots, N\}$ onto the weak compositions $n = (n_1, \dots, n_m) \in \mathbb{N}_0^m$ of k, that is, $\sum_j n_j = k$.

Proof. Fix a weak composition n of k and put $\eta = (m-1)/m \in (0, 1)$. Consider the finite transportation problem with one unit of supply at each of the N sources and demand $n_j + \eta$ at label j. It is balanced, because $\sum_j (n_j + \eta) = k + m - 1 = N$. By finite-dimensional linear programming duality there are an optimal plan $P = (P_{ij})$ and dual numbers (u_i, ξ_j) such that

$$u_i \leq c_{ij} + \xi_j \quad \text{for all } i, j, \qquad P_{ij} > 0 \implies u_i = c_{ij} + \xi_j. \tag{A.7}$$

We fix the gauge $\xi_m = 0$. The *support graph* has the edges (i, j) with $P_{ij} > 0$, and the *equality graph* has the edges with $u_i = c_{ij} + \xi_j$; by (A.7) the first is contained in the second.

We first show that the support graph is a spanning tree equal to the equality graph. Every source has an incident support edge, since its supply is one, and so does every label, since $n_j + \eta > 0$. Mass is conserved within each connected component of the support graph. If a component contains m' labels, its sources carry an integer amount of mass equal to $\sum_j n_j + m'\eta$, the sum running over the labels of the component. Hence $m'\eta = m'(m-1)/m$

is an integer. As $\gcd(m, m-1) = 1$, m' is a positive multiple of m, so $m' = m$ and the support graph is connected and spans all $N+m$ vertices. Along a cycle of the equality graph the alternating sum of the $c_{ij} = u_i - \xi_j$ telescopes to zero, which cycle-genericity forbids. So the equality graph is acyclic. It contains the connected spanning support graph, and an acyclic graph containing a spanning tree coincides with it. Hence both graphs are the same spanning tree, with $N+m-1$ edges.

Next we show that exactly $m-1$ sources are split between two labels. Call a source *split* if its degree is at least two, and let I be the set of split sources. The source degrees sum to $N+m-1$, so the total excess $\sum_i(\deg i - 1)$ equals $m-1$; in particular $|I| \le m-1$. Let ℓ_j be the number of unsplit sources attached to label j. The remaining demand $n_j + \eta - \ell_j$ at label j is nonnegative and differs from an integer by $\eta \in (0,1)$, so it is at least η. This remaining demand is supplied by the split sources, whose total mass is $|I|$. Hence $|I| \ge m\eta = m-1$. Thus $|I| = m-1$, and since the total excess is $m-1$, every split source has degree exactly two. Moreover $n_j - \ell_j \ge 0$ for every j, while $\sum_j \ell_j = N - |I| = k = \sum_j n_j$; therefore $\ell_j = n_j$ for every j.

We now show that $K(I) = n$. Deleting the k unsplit sources, which are leaves, leaves a tree in which every split source has degree two. Contracting each split source to an edge between its two labels produces a tree on the m labels. For a root label j, orient this tree away from j and match each split source to the label at the far end of its edge. This is a bijection from I onto the labels other than j that uses only edges of the equality graph. For any bijection π from I onto $\{1,\dots,m\} \setminus \{j\}$, the inequality in (A.7) gives $\sum_{i\in I} c_{i\pi(i)} \ge \sum_{i\in I} u_i - \sum_{l\ne j} \xi_l$, with equality for the bijection just constructed. Therefore

$$D_j(I) = \sum_{i\in I} u_i - \sum_{l\ne j} \xi_l, \qquad D_j(I) - D_m(I) = \xi_j - \xi_m = \xi_j. \tag{A.8}$$

Let $i \notin I$ be an unsplit source and j its label in the tree. By (A.8),

$$c_{il} + D_l(I) = c_{il} + \xi_l + D_m(I) \ge u_i + D_m(I) \qquad (1 \le l \le m),$$

with equality exactly when (i,l) is an edge of the equality graph. That graph is the tree, in which i has degree one, so equality holds only for $l = j$. Hence $\iota_I(i) = j$ is the unique minimizer, and $K_j(I) = \ell_j = n_j$ for every j.

We have constructed, for each weak composition n, a subset I_n with $K(I_n) = n$ and unique minimizers. Distinct n give distinct I_n. There are $\binom{k+m-1}{m-1} = \binom{N}{m-1}$ weak compositions and as many $(m-1)$-element subsets, so $n \mapsto I_n$ is onto. Every subset I is therefore some I_n; its minimizers are unique, and (A.6) is the inverse bijection. □

A.4 Mixed moments and the Dirichlet law

We return to the i.i.d. setting under (A.2). Fix $k \ge 0$, set $N = m-1+k$, and let $X_1, \dots, X_N$ be i.i.d. with law μ; the base is $X_1, \dots, X_{m-1}$ and the k additional sources $X_m, \dots, X_N$ are independent of it. Each additional source receives the label $\iota(X_i)$ given by the same map (A.4); the additional sources do not compete with one another. Conditional on the base, the labels are i.i.d. with law $(A_1, \dots, A_m)$, so the count vector $K = (K_1, \dots, K_m)$, $K_j = \#\{m \le i \le N : \iota(X_i) = j\}$, is multinomial:

$$\mathbb{P}(K = n \mid X_1, \dots, X_{m-1}) = \frac{k!}{n_1! \cdots n_m!} \prod_{j=1}^{m} A_j^{n_j}, \qquad \textstyle\sum_j n_j = k. \tag{A.9}$$

The matrix $C = (c_j(X_i))_{i\le N,\, j\le m}$ is cycle-generic almost surely. Indeed, there are finitely many simple cycles. For a fixed one, pick a source i_0 on it, with neighbouring labels $j \ne l$ on

the cycle (a simple cycle visits i_0 exactly once), and condition on all sources except X_{i_0}: the alternating sum becomes $c_j(X_{i_0}) - c_l(X_{i_0})$ plus a constant, which vanishes with probability zero by (A.2) and Fubini's theorem. With the notation of Lemma A.3 applied to C, we have $K = K(I_0)$ for $I_0 = \{1, \dots, m-1\}$ almost surely. The law of C is invariant under permutations of its rows, so $\mathbb{P}(K(I) = n)$ is the same for every $(m-1)$-element subset I. By Lemma A.3, almost surely exactly one subset I has $K(I) = n$. Averaging over I,

$$\mathbb{P}(K = n) = \binom{N}{m-1}^{-1} \sum_I \mathbb{P}(K(I) = n) = \binom{N}{m-1}^{-1} \mathbb{E}\,\#\{I : K(I) = n\} = \frac{(m-1)!\,k!}{(m+k-1)!}.$$

Taking expectations in (A.9) yields, for every $n \in \mathbb{N}_0^m$,

$$\mathbb{E} \prod_{j=1}^m A_j^{n_j} = \frac{(m-1)! \prod_{j=1}^m n_j!}{(m+k-1)!}, \qquad k = \sum_j n_j. \tag{A.10}$$

The right side is the mixed moment of the uniform law $\mathrm{Dir}(1, \dots, 1)$ on Δ_{m-1}, whose density in the coordinates $(w_1, \dots, w_{m-1})$ is $(m-1)!$:

$$(m-1)! \int_{\Delta_{m-1}} \prod_j w_j^{n_j}\, \mathrm{d}w_1 \cdots \mathrm{d}w_{m-1} = \frac{(m-1)! \prod_j n_j!}{(m+k-1)!}.$$

By the Stone–Weierstrass theorem, polynomials are uniformly dense in the continuous functions on the compact simplex, so mixed moments determine a probability law on Δ_{m-1}. Hence

$$(A_1, \dots, A_m) \overset{\text{law}}{=} \mathrm{Dir}(1, \dots, 1). \tag{A.11}$$

Since $\mathcal{T}_c(\mu, \cdot)$ is Borel on Δ_{m-1} (Lemma A.1), (A.11) gives $\mathbb{E}\,\mathcal{T}_c(\mu, A) = \mathbb{E}_W \mathcal{T}_c(\mu, W)$ for $W \sim \mathrm{Dir}(1, \dots, 1)$. Together with (A.5), this proves Theorem 3.1 when $m \geq 2$ and (A.2) holds.

A.5 Removing the genericity assumption

If $m = 1$, both sides of the identity in Theorem 3.1 equal $\int c_1 \,\mathrm{d}\mu$, since $\mathcal{T}_c(\mu, (1)) = \int c_1 \,\mathrm{d}\mu$. Let $m \geq 2$ and $c_j \in L^1(\mu)$ be arbitrary. We augment every source point x by a vector $z = (z_1, \dots, z_m)$ of i.i.d. uniform $[0,1]$ variables, and work on $\mathsf{X} \times [0,1]^m$ with $\widetilde{\mu} = \mu \otimes \mathrm{Unif}([0,1]^m)$ and the costs

$$\widetilde{c}_j^{\,\varepsilon}(x, z) = c_j(x) + \varepsilon z_j, \qquad \varepsilon > 0. \tag{A.12}$$

For $j \neq l$ and fixed x, the difference $\widetilde{c}_j^{\,\varepsilon} - \widetilde{c}_l^{\,\varepsilon}$ has an atomless law in z; hence $\widetilde{\mu}$ and $\widetilde{c}^{\,\varepsilon}$ satisfy (A.2), even when μ is purely atomic or $c_j = c_l$. The theorem therefore holds for the perturbed costs.

Let $(X_i, Z_i)_{i \leq m}$ be i.i.d. with law $\widetilde{\mu}$, and let M_0 and M_ε be the optimal assignment costs for c at the points X_i and for $\widetilde{c}^{\,\varepsilon}$ at the points (X_i, Z_i). Every assignment has its cost increased by a number in $[0, m\varepsilon]$, so

$$0 \leq M_\varepsilon - M_0 \leq m\varepsilon. \tag{A.13}$$

For the semidiscrete problem, fix $w \in \Delta_{m-1}$. If $(\widetilde{a}_j)$ is an allocation on $\mathsf{X} \times [0,1]^m$ with masses w, then $\bar{a}_j(x) = \int \widetilde{a}_j(x, z)\,\mathrm{d}z$ is, by Fubini's theorem, an allocation on X with masses w, and the unperturbed part of the cost of $(\widetilde{a}_j)$ equals $\int \sum_j c_j \bar{a}_j \,\mathrm{d}\mu \geq \mathcal{T}_c(\mu, w)$; the perturbation is nonnegative. Conversely, every allocation (a_j) on X lifts to $\widetilde{a}_j(x, z) = a_j(x)$, with an additional cost of at most ε. Taking infima gives the deterministic bound

$$0 \leq \mathcal{T}_{\widetilde{c}^{\,\varepsilon}}(\widetilde{\mu}, w) - \mathcal{T}_c(\mu, w) \leq \varepsilon \qquad (w \in \Delta_{m-1}). \tag{A.14}$$

By (A.13) and (A.14), replacing c by $\widetilde{c}^{\,\varepsilon}$ changes each side of the identity in Theorem 3.1 by at most $m\varepsilon$. Letting $\varepsilon \downarrow 0$ proves Theorem 3.1 for general costs $c_j \in L^1(\mu)$, without continuity, atomlessness or uniqueness of optimizers.

A.6 Proof of Corollary 3.2

The corollary records the normalization used in (3.7). For $y = (y_1, \dots, y_m)$ and $w \in \Delta_{m-1}$, write $\nu_{w,y} = \sum_j w_j \delta_{y_j}$ and $c_j^y = c(\cdot, y_j)$, and let G_m be the set of y with $\int \max_j |c(x, y_j)| \, \mathrm{d}\mu(x) < \infty$. By Tonelli's theorem G_m is measurable. In part (i) the hypothesis is $y \in G_m$; in part (ii) the cost is bounded, so $G_m = \mathsf{X}^m$.

We first identify the two semidiscrete values: for $y \in G_m$ and $w \in \Delta_{m-1}$,

$$\mathrm{OT}_c(\mu, \nu_{w,y}) = \mathcal{T}_{c^y}(\mu, w). \tag{A.15}$$

This is (3.2), applied with $c_j = c_j^y$; the integrals are finite since $y \in G_m$.

In part (ii) we also need that $(w, y) \mapsto \mathrm{OT}_c(\mu, \nu_{w,y})$ is measurable on $\Delta_{m-1}^\circ \times \mathsf{X}^m$. By (A.15) this map is $(w, y) \mapsto \mathcal{T}_{c^y}(\mu, w)$. For fixed y it is continuous on Δ_{m-1}° by Lemma A.1. For fixed w it is continuous in y: the admissible allocations do not depend on y, so $|\mathcal{T}_{c^y}(\mu, w) - \mathcal{T}_{c^{y'}}(\mu, w)| \le \max_j \sup_x |c(x, y_j) - c(x, y_j')|$, which tends to zero as $y' \to y$ by uniform continuity of c. A function on $\Delta_{m-1}^\circ \times \mathsf{X}^m$ that is continuous in y for fixed w and Borel in w for fixed y is measurable for $\mathcal{B}(\Delta_{m-1}^\circ) \otimes \mathcal{B}(\mathsf{X}^m)$, because X^m is a separable metric space (a Carathéodory function); this product σ-field is the Borel σ-field of $\Delta_{m-1}^\circ \times \mathsf{X}^m$. Since a $\mathrm{Dir}(1, \dots, 1)$ vector lies in Δ_{m-1}° almost surely, $\mathrm{OT}_c(\mu, \nu_{W,Y})$ is a random variable.

We now prove part (i) of the corollary. Change variables from the independent marks $E_1, \dots, E_m$ to (S, W), with $W_j = E_j / S$ and $W_m = 1 - \sum_{j<m} W_j$. The Jacobian is S^{m-1}, so the joint density is

$$e^{-S} S^{m-1} \, \mathrm{d}S \, \mathrm{d}w_1 \cdots \mathrm{d}w_{m-1} = \Big(\frac{e^{-S} S^{m-1}}{(m-1)!} \, \mathrm{d}S \Big) \big((m-1)! \, \mathrm{d}w_1 \cdots \mathrm{d}w_{m-1} \big).$$

Thus $S \sim \mathrm{Gamma}(m, 1)$ is independent of $W \sim \mathrm{Dir}(1, \dots, 1)$, and $\mathbb{E}S = m$. Since $S > 0$ a.s. and transport cost is homogeneous in the common total mass,

$$\mathrm{OT}_c\Big(S\mu, \sum_j E_j \delta_{y_j} \Big) = S \ \mathrm{OT}_c(\mu, \nu_{W,y}),$$

which is measurable by Lemma A.1 and (A.15). By the independence of S and W, (A.15) and Theorem 3.1 applied to $c_j = c_j^y \in L^1(\mu)$,

$$\mathbb{E}_E \ \mathrm{OT}_c\Big(S\mu, \sum_j E_j \delta_{y_j} \Big) = m \, \mathbb{E}_W \mathcal{T}_{c^y}(\mu, W) = \mathbb{E}_X \min_{\sigma \in S_m} \sum_{i=1}^m c(X_i, y_{\sigma(i)}).$$

Both sides are bounded in absolute value by $m \int \max_j |c(x, y_j)| \, \mathrm{d}\mu(x) < \infty$, which proves part (i).

For part (ii), condition on $N = m \ge 1$ and on the target locations $Y_1, \dots, Y_m$. Given $N = m$, the sources $X_1, \dots, X_m$ are i.i.d. with law μ and independent of the targets and of the marks, and $S = \sum_{j \le m} E_j$. Part (i) applies. By the measurability established above, we may average over the locations and over N, since the left side is bounded by $N \sup |c|$ and the right side by $S \sup |c|$, both of which are integrable. On $\{N = 0\}$ both sides vanish by convention.

B Proof of the periodic flux theorem

We prove Theorem 12.1. The current is the time integral of the gradient of the heat flow applied to v. For short times it is controlled by heat contraction. For intermediate times we need the decay of $\|P_t v - \fint v\|_p$, which comes from the independence of distant blocks of

the field, through Rosenthal's inequality, combined with Gaussian bounds. The long-time part is controlled by the spectral gap and vanishes as $T/\ell \to \infty$. The proof uses only the stationarity of u, its ℓ-locality (through the independence of blocks), the joint measurability of $(\omega, x) \mapsto u(x)$ (Section 2.4) and $v(0) \in L^p(\mathbb{P})$.

Throughout, C denotes numerical constants, independent of d, p, T, ℓ, and $C_{d,p}$ constants depending only on d and p. For random fields, $\|\cdot\|_p$ is the norm of $L^p(\mathbb{P} \otimes dx/V)$; for deterministic kernels on Ω_T, $\|\cdot\|_{L^r(\Omega_T)}$ is taken with respect to dx. By stationarity, $\|w\|_p = \|w(0)\|_{L^p(\mathbb{P})}$ for every translation-covariant field w. Since $p \geq 2$, $\|v(0)\|_{L^2(\mathbb{P})} \leq A$.

B.1 Independent blocks

Let $m = \lfloor T/\ell \rfloor \geq 640$, so that the side T/m lies in $[\ell, 2\ell]$, and partition Ω_T into m^d cubes of side T/m. On the cyclic index set $\mathbb{Z}/m\mathbb{Z}$, consider the square of the cycle, in which two indices are adjacent when their cyclic distance is at most two. Its maximal degree is four, so a greedy colouring uses at most five colours. Give each cube the product of the colours of its d coordinate indices; there are at most 5^d colour classes. Two distinct cubes of the same colour have indices at cyclic distance at least three in some coordinate, hence are at distance at least $2T/m \geq 2\ell$, and their ℓ-neighbourhoods are disjoint up to boundaries. The values of v on cubes of the same colour are therefore determined by the Poisson process on disjoint sets, and are independent.

Let K be a deterministic periodic kernel. By stationarity, $\|K * v\|_p = \|K * v(0)\|_{L^p(\mathbb{P})}$, and

$$K * v(0) = \sum_Q \xi_Q, \qquad \xi_Q = \int_Q K(-x) v(x)\, dx.$$

Each ξ_Q is centred and depends only on the Poisson points within distance ℓ of Q. Minkowski's inequality and stationarity give $\|\xi_Q\|_{L^r(\mathbb{P})} \leq A I_Q$ for $r \in \{2, p\}$, where $I_Q = \int_Q |K(-x)|\, dx$. Within a colour class c the ξ_Q are independent, and Rosenthal's inequality, whose constant for independent centred variables is at most Cp [17], gives

$$\Big\| \sum_{Q \in c} \xi_Q \Big\|_{L^p(\mathbb{P})} \leq CpA \Big[\Big(\sum_{Q \in c} I_Q^2 \Big)^{1/2} + \Big(\sum_{Q \in c} I_Q^p \Big)^{1/p} \Big].$$

We sum over the at most 5^d classes, using the Cauchy–Schwarz inequality for the first term and Hölder's inequality for the second:

$$\sum_c \Big(\sum_{Q \in c} I_Q^2 \Big)^{1/2} \leq 5^{d/2} \Big(\sum_Q I_Q^2 \Big)^{1/2}, \qquad \sum_c \Big(\sum_{Q \in c} I_Q^p \Big)^{1/p} \leq 5^{d(1-1/p)} \Big(\sum_Q I_Q^p \Big)^{1/p}.$$

Finally, Hölder's inequality on each cube, of volume at most $(2\ell)^d$, gives $I_Q^2 \leq (2\ell)^d \int_{-Q} |K|^2$ and $I_Q^p \leq (2\ell)^{d(p-1)} \int_{-Q} |K|^p$. Hence

$$\|K * v\|_p \leq CpA \Big[(10\ell)^{d/2} \|K\|_{L^2(\Omega_T)} + (10\ell)^{d(1-1/p)} \|K\|_{L^p(\Omega_T)} \Big]. \tag{B.1}$$

The factors $5^{d/2}$ and $5^{d(1-1/p)}$ from the colouring are the source of the bases 10ℓ.

B.2 Heat kernel bounds

Let $k_t^{T,1}$ be the one-dimensional periodized heat kernel. Its Fourier coefficients are nonnegative, so its maximum is attained at 0. Since

$$\sum_{j \in \mathbb{Z}} e^{-j^2 T^2/(4t)} \leq 1 + 2 \int_0^\infty e^{-u^2 T^2/(4t)}\, du = 1 + \frac{2\sqrt{\pi t}}{T},$$

we get

$$\|k_t^{T,1}\|_\infty \le (4\pi t)^{-1/2} + T^{-1} \le 2t^{-1/2}, \qquad 0 < t \le T^2.$$

The d-dimensional kernel k_t^T is the tensor product of these kernels and has $\|k_t^T\|_{L^1(\Omega_T)} = 1$. Interpolation gives, for $d \ge 3$, $p \ge 2$ and $0 < t \le T^2$,

$$\|k_t^T - V^{-1}\|_{L^p(\Omega_T)} \le 2\|k_t^T\|_{L^p(\Omega_T)} \le (8/\sqrt{t})^{d(1-1/p)},$$

and the same bound with $p = 2$.

Set $t_0 = (160\ell)^2$. For $T \ge 640\ell$ we have $2t_0 \le T^2/2$. We apply (B.1) with $K = k_t^T - V^{-1}$, for which $K * v = P_t v - \fint v$. Both terms then have the base $80\ell/\sqrt{t}$, which is at most $1/2$ for $t \ge t_0$; since $d(1-1/p) \ge d/2$, the second term is bounded by the first. As $(80\ell/\sqrt{t})^{d/2} \le (t_0/t)^{d/4}$, this gives

$$\Big\| P_t v - \fint v \Big\|_p \le CpA\,(t_0/t)^{d/4}, \qquad t_0 \le t \le T^2. \tag{B.2}$$

The gradient of the periodized heat kernel satisfies $\|\nabla k_t^T\|_{L^1(\Omega_T;\mathbb{R}^d)} \le \sqrt{d/(2t)}$, for the Euclidean norm of the gradient. Indeed, on $\mathbb{R}^d$ the heat kernel $g_t(x) = (4\pi t)^{-d/2} e^{-|x|^2/(4t)}$ satisfies $|\nabla g_t(x)| = |x| g_t(x)/(2t)$, and the Cauchy–Schwarz inequality with $\int |x|^2 g_t\,dx = 2dt$ yields $\int |\nabla g_t|\,dx \le \sqrt{d/(2t)}$; periodization does not increase the L^1 norm.

For long times we use the spectral gap. Let w be a mean-zero field. For $\tau \ge T^2$, write $P_\tau = P_{T^2} P_{\tau - T^2}$. The spectral gap $4\pi^2/T^2$ of $-\Delta$ on mean-zero functions gives $\|P_{\tau-T^2} w\|_2 \le e^{-4\pi^2(\tau-T^2)/T^2}\|w\|_2$, and $\|w\|_2 \le \|w\|_p$ since $p \ge 2$. To return from L^2 to L^p we use the smoothing of P_{T^2}: by scaling to the unit torus, P_{T^2} maps $L^2(dx/V)$ to L^∞ with a norm C_d depending only on d. Hence $\|P_\tau w\|_p \le C_d e^{4\pi^2} e^{-4\pi^2\tau/T^2}\|w\|_p$ for $\tau \ge T^2$, and by contraction the same bound holds for $0 \le \tau < T^2$ (with $C_d \ge 1$). We apply it to $w = P_{T^2} v - \fint v$, whose norm is at most $CpA(160\ell/T)^{d/2}$ by (B.2), and obtain

$$\Big\| P_t v - \fint v \Big\|_p \le C_{d,p} A\,(\ell/T)^{d/2} e^{-4\pi^2 t/T^2}, \qquad t \ge T^2. \tag{B.3}$$

The constant $C_{d,p}$ absorbs C_d, the factor $e^{8\pi^2}$ and $Cp\,160^{d/2}$.

B.3 Time integration

We bound $\|\nabla P_t v\|_p$ on three ranges of t.

On $(0, 2t_0)$, heat contraction and the gradient bound give $\|\nabla P_t v\|_p \le \|\nabla k_t^T\|_{L^1(\Omega_T)}\|v\|_p \le \sqrt{d/(2t)}\,A$, whose integral is $2\sqrt{dt_0}\,A$.

On $(2t_0, 2T^2)$, we write $\nabla P_t v = \nabla P_{t/2}(P_{t/2} v - \fint v)$ and use (B.2) at time $t/2 \in [t_0, T^2]$:

$$\int_{2t_0}^{2T^2} \|\nabla P_t v\|_p\,dt \le Cp\sqrt{d}\,A \int_{2t_0}^{\infty} t^{-1/2}(2t_0/t)^{d/4}\,dt = Cp\sqrt{d}\,A\,\frac{4\sqrt{2t_0}}{d-2} \le Cp\sqrt{dt_0}\,A,$$

where $d \ge 3$ is used for the convergence of the time integral.

On $(2T^2, \infty)$, the same splitting and (B.3) at time $t/2 \ge T^2$ give $\|\nabla P_t v\|_p \le \sqrt{d/t}\,C_{d,p} A(\ell/T)^{d/2} e^{-2\pi^2 t/T^2}$, whose integral is at most $C_{d,p} T(\ell/T)^{d/2} A$.

Since $\sqrt{t_0} = 160\ell$ and $p \ge 2$, the three bounds give the convergence of $\mathbf{j} = -\int_0^\infty \nabla P_t v\,dt$ in $L^p(\mathbb{P} \otimes dx/V; \mathbb{R}^d)$ and the estimate (12.5).

For the divergence, $\partial_t P_t v = \Delta P_t v = \operatorname{div} \nabla P_t v$ gives, for $0 < \tau < \tau'$,

$$\operatorname{div}\Big(- \int_\tau^{\tau'} \nabla P_t v\,dt \Big) = P_\tau v - P_{\tau'} v$$

in the sense of distributions. As $\tau \to 0$, $P_\tau v \to v$ in L^p; as $\tau' \to \infty$, $P_{\tau'} v \to \fint v$ by (B.3). Hence $\operatorname{div} \mathbf{j} = v - \fint v = u - \fint u$, which proves Theorem 12.1.

Remark B.1. The argument gives the leading coefficient more precisely. The first time range contributes $C\sqrt{d}\,\ell A$, and the middle range, whose time integral carries the factor $1/(d-2)$, contributes $Cp\,\ell A/\sqrt{d}$; hence $\|\mathbf{j}\|_p \le C(\sqrt{d}+p/\sqrt{d})\ell A + C_{d,p}T(\ell/T)^{d/2}A$.

Declaration on the use of artificial intelligence

This work was developed through an iterative, AI-assisted research process, principally using OpenAI's ChatGPT (Astra 6), with additional contributions from Anthropic's Claude (Opus 5.5) to a lesser extent. The authors formulated the research questions, provided mathematical ideas and preliminary arguments, and directed the development of the proofs. AI tools were used extensively to explore alternative proof strategies, formulate and refine mathematical arguments, check calculations and parameter estimates, identify potential gaps, consolidate independently developed approaches, and assist with literature research and LaTeX preparation. The manuscript underwent several rounds of adversarial mathematical review, including AI-assisted critiques that prompted substantial revisions and simplifications. Computational checks were also used to test selected identities and numerical estimates. These procedures do not constitute formal verification or independent certification of the results.

References

[1] D. J. Aldous, *The $\zeta(2)$ limit in the random assignment problem*, Random Structures Algorithms **18** (2001), 381–418.

[2] L. Ambrosio, N. Gigli and G. Savaré, *Gradient Flows in Metric Spaces and in the Space of Probability Measures*, 2nd ed., Lectures in Mathematics ETH Zürich, Birkhäuser, Basel, 2008.

[3] L. Ambrosio, F. Stra and D. Trevisan, *A PDE approach to a 2-dimensional matching problem*, Probab. Theory Related Fields **173** (2019), 433–477.

[4] W. Beckner, *Inequalities in Fourier analysis*, Ann. of Math. (2) **102** (1975), 159–182.

[5] E. Boissard and T. Le Gouic, *On the mean speed of convergence of empirical and occupation measures in Wasserstein distance*, Ann. Inst. Henri Poincaré Probab. Stat. **50** (2014), 539–563.

[6] J. H. Boutet de Monvel and O. C. Martin, *Mean field and corrections for the Euclidean minimum matching problem*, Phys. Rev. Lett. **79** (1997), 167–170.

[7] H. J. Brascamp and E. H. Lieb, *Best constants in Young's inequality, its converse, and its generalization to more than three functions*, Adv. Math. **20** (1976), 151–173.

[8] S. Caracciolo, M. D'Achille and G. Sicuro, *Random Euclidean matching problems in one dimension*, Phys. Rev. E **96** (2017), 042102.

[9] S. Caracciolo, C. Lucibello, G. Parisi and G. Sicuro, *Scaling hypothesis for the Euclidean bipartite matching problem*, Phys. Rev. E **90** (2014), 012118.

[10] M. P. D'Achille, *On two linear assignment problems: random assignment and Euclidean bipartite matching*, MSc thesis, Università degli Studi di Milano, 2016.

[11] M. D'Achille, F. Mattesini and D. Trevisan, *Concentration for random Euclidean combinatorial optimization*, preprint, arXiv:2602.21851, 2026.

[12] N. Fournier and A. Guillin, *On the rate of convergence in Wasserstein distance of the empirical measure*, Probab. Theory Related Fields **162** (2015), 707–738.

[13] M. Goldman, L. Koch and D. Trevisan, *Large-scale regularity theory for optimal p-transport maps and application to concentration in the matching problem*, preprint, arXiv:2611.25660, 2026.

[14] M. Goldman and D. Trevisan, *Convergence of asymptotic costs for random Euclidean matching problems*, Probab. Math. Phys. **2** (2021), 341–362.

[15] M. Goldman and D. Trevisan, *Optimal transport methods for combinatorial optimization over two random point sets*, Probab. Theory Related Fields **188** (2024), 1315–1384.

[16] J. Houdayer, J. H. Boutet de Monvel and O. C. Martin, *Comparing mean field and Euclidean matching problems*, Eur. Phys. J. B **6** (1998), 383–393.

[17] W. B. Johnson, G. Schechtman and J. Zinn, *Best constants in moment inequalities for linear combinations of independent and exchangeable random variables*, Ann. Probab. **13** (1985), 234–253.

[18] J. Larsson, *The minimum perfect matching in pseudo-dimension* $0 < q < 1$, Combin. Probab. Comput. **30** (2021), 374–397.

[19] G. Last and M. Penrose, *Lectures on the Poisson Process*, IMS Textbooks, Cambridge University Press, Cambridge, 2018.

[20] M. Ledoux, *On optimal matching of Gaussian samples*, Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI) **457** (2017), 226–264.

[21] C. Lucibello, G. Parisi and G. Sicuro, *One-loop diagrams in the random Euclidean matching problem*, Phys. Rev. E **95** (2017), 012302.

[22] C. McDiarmid, *On the method of bounded differences*, in: Surveys in Combinatorics 1989, London Math. Soc. Lecture Note Ser. **141**, Cambridge University Press, 1989, 148–188.

[23] M. Mézard and G. Parisi, *Replicas and optimization*, J. Physique Lett. **46** (1985), L771–L778.

[24] M. Mézard and G. Parisi, *Mean-field equations for the matching and the travelling salesman problems*, Europhys. Lett. **2** (1986), 913–918.

[25] M. Mézard and G. Parisi, *On the solution of the random link matching problems*, J. Physique **48** (1987), 1451–1459.

[26] M. Mézard and G. Parisi, *The Euclidean matching problem*, J. Physique **49** (1988), 2019–2025.

[27] J. Salez, *The Mézard–Parisi equation for matchings in pseudo-dimension* $d > 1$, Electron. Commun. Probab. **20** (2015), no. 13.

[28] F. Santambrogio, *Optimal Transport for Applied Mathematicians*, Progress in Nonlinear Differential Equations and their Applications **87**, Birkhäuser, Cham, 2015.

[29] M. Talagrand, *Matching random samples in many dimensions*, Ann. Appl. Probab. **2** (1992), 846–856.

[30] C. Villani, *Optimal Transport. Old and New*, Grundlehren der mathematischen Wissenschaften **338**, Springer, Berlin, 2009.

[31] J. Wästlund, *Replica symmetry of the minimum matching*, Ann. of Math. (2) **175** (2012), 1061–1091.

[32] J. Weed and F. Bach, *Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance*, Bernoulli **25** (2019), 2620–2648.

[33] Y. Xu, *Random matching between distinct distributions in two dimensions*, preprint, arXiv:2610.40509, 2026.