

Disjoint Optimizers and the Directed Landscape

Duncan Dauvergne

Lingfu Zhang

Author address:

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF TORONTO, ON M5S 1A1,
TORONTO, CANADA

Email address: `duncan.dauvergne@utoronto.ca`

DEPARTMENT OF MATHEMATICS, PRINCETON UNIVERSITY, PRINCETON,
NEW JERSEY 08544

Email address: `lingfuz@math.princeton.edu`

NOT FOR DISTRIBUTION

Contents

Chapter 1. Introduction	1
1.1. The KPZ universality class	1
1.2. The directed landscape	2
1.3. Two perspectives on last passage percolation	4
1.4. Brownian LPP and the extended Airy sheet	6
1.5. The full scaling limit of multi-point Brownian LPP	8
1.6. Consequences	10
1.7. More related work	11
1.8. Outline of the paper and a primer about the proofs	12
Acknowledgments	14
Chapter 2. Last passage percolation across lines	15
2.1. Last passage with multiple paths	16
2.2. Basic properties of disjoint optimizers and last passage values	17
2.3. Melons	21
2.4. Brownian melons and the parabolic Airy line ensemble	22
2.5. Melon geodesics and the Airy sheet	25
Chapter 3. Tightness	29
3.1. Tightness of prelimiting sheets	29
3.2. Tightness of melon optimizers	32
Chapter 4. Last passage percolation across the parabolic Airy line ensemble	37
4.1. Parabolic paths, length, and geodesics in $\mathcal{B}$	37
4.2. Disjoint optimizers in $\mathcal{B}$	43
Chapter 5. Limits of Melon Optimizers and the Extended Airy Sheet	47
5.1. Convergence of $\mathcal{S}^n$	47
5.2. Properties of the extended Airy sheet $\mathcal{S}$	49
5.3. Metric composition law	53
Chapter 6. The Scaling Limit of Multi-point Brownian LPP	55
6.1. Tightness of the prelimiting extended landscape	55
6.2. The explicit construction of $\mathcal{L}^*$	56
Chapter 7. Paths in the Extended Landscape	63
7.1. Path weights	63
7.2. Optimizers and transversal fluctuation	64
7.3. Monotonicity of optimizers	70
7.4. Sums of disjoint paths	71

Chapter 8. Disjointness of Optimizers	75
8.1. Convergence in the overlap topology	75
8.2. Two paths	77
8.3. Multiple paths	83
Chapter 9. Convergence of Optimizers	85
Appendix A. Appendix	89
A.1. Brownian melon estimates	89
A.2. Proof of Theorem 8.10	94
Bibliography	101

Abstract

We study maximal length collections of disjoint paths, or ‘disjoint optimizers’, in the directed landscape. We show that disjoint optimizers always exist, and that their lengths can be used to construct an extended directed landscape. The extended directed landscape can be built from an independent collection of extended Airy sheets, which we define from the parabolic Airy line ensemble. We show that the extended directed landscape and disjoint optimizers are scaling limits of the corresponding objects in Brownian last passage percolation (LPP). As two consequences of this work, we show that one direction of the Robinson-Schensted-Knuth bijection passes to the KPZ limit, and we find a criterion for geodesic disjointness in the directed landscape that uses only a single parabolic Airy line ensemble.

The proofs rely on a new notion of multi-point LPP across the parabolic Airy line ensemble, combinatorial properties of multi-point LPP, and probabilistic resampling ideas.

Received by the editor October 20, 2021, and, in revised form, October 21, 2021, May 13, 2022, and June 6, 2022.

Article electronically published on Month Day, Year.

DOI: <https://doi.org/1090/???>

2020 *Mathematics Subject Classification*. Primary 60K35; Secondary 82B23, 82B43.

Key words and phrases. directed landscape, KPZ universality, geodesics, disjoint optimizers, Airy sheet, last passage percolation, Airy line ensemble, Robinson-Schensted-Knuth correspondence.

©XXXX American Mathematical Society

NOT FOR DISTRIBUTION

CHAPTER 1

Introduction

1.1. The KPZ universality class

We start by considering a simple model of a random metric on the plane, see Figure 1.1 for an example. Take the lattice $\mathbb{Z}^2$, and assign i.i.d. positive edge weights to all edges. For vertices $p, q \in \mathbb{Z}^2$, let $d(p, q)$ denote graph distance in this randomly weighted graph. This model is an example of first passage percolation on $\mathbb{Z}^2$.

What does this metric look like as the points p and q become further and further away from each other? As long as the weight distribution is sufficiently nice, we expect the following picture. For concreteness, we take $p = (0, 0)$ and $q = (0, n)$. At leading order, $d(p, q)$ grows proportionally to n and the fluctuations of $d(p, q)$ are of order $O(n^{1/3})$. Moreover, geodesics from p to q are expected to only use edges in an $O(n^{2/3})$ -window around the vertical axis. In particular, $d(p, q)$ will




FIGURE 1.1. First passage percolation on $\mathbb{Z}^2$ with i.i.d. edge weights equal to 1 or 2 with equal probability. This model is expected to lie in the KPZ universality class. The red path is a geodesic and here the weighted graph is drawn using a spring electrical embedding (Mathematica). Figure generated by Bálint Virág.

only have nontrivial correlations with another distance $d((x_1, y_1), (x_2, y_2))$ when $x_1, x_2 = O(n^{2/3})$.

This $n : n^{1/3} : n^{2/3}$ scaling (usually referred to as a $1 : 2 : 3$ scaling) is also observed in other models with an underlying random planar geometry including random interface growth (e.g. the 1 + 1-dimensional Kardar-Parisi-Zhang (KPZ) equation, TASEP), random polymers, and last passage percolation. Collectively, these models are said to lie in the *KPZ universality class*. See Section 1.7 for background on this area.

1.2. The directed landscape

Any model satisfying the $1 : 2 : 3$ scaling described above is expected to converge to a scaling limit. Indeed, for first passage percolation $d(\cdot, \cdot)$ on $\mathbb{Z}^2$ defined using a sufficiently nice edge weight distribution we expect that there is a four-parameter continuum object $\mathcal{L}$ such that

$$(1.1) \quad d((xn^{2/3}, sn), (yn^{2/3}, tn)) = c_1(t - s)n - n^{1/3}\mathcal{L}(x, s; y, t) + o(n^{1/3})$$

for some constant c_1 .¹ The continuum limit $\mathcal{L}$ should be universal: up to possibly a linear rescaling of parameters, we should see the same limit regardless of our initial choice of weight distribution. Moreover, it should be the limit of all models in the KPZ universality class.

While the possibility of verifying (1.1) for general first passage percolation is pure conjecture at this point, over the past twenty-five years a handful of integrable models of *last passage percolation* have been discovered where establishing the existence of this scaling limit is possible. Indeed, recently Dauvergne, Ortmann, and Virág [15] constructed the limit object $\mathcal{L}$ and verified the analogue of (1.1) for one particular integrable model: Brownian last passage percolation. They called the identified limit the **directed landscape**. This result was extended to other integrable models in [18]. The papers [15, 18] build on many previous results in the field. We refer the reader to Section 1.7 for more about of this history, and continue describing the directed landscape $\mathcal{L}$.

The directed landscape $\mathcal{L}$ is a random continuous function from the parameter space

$$\mathbb{R}_\uparrow^4 = \{u = (p; q) = (x, s; y, t) \in \mathbb{R}^4 : s < t\}$$

to $\mathbb{R}$.

As with first passage percolation, the value $\mathcal{L}(p; q) = \mathcal{L}(x, s; y, t)$ is best thought of as a distance between two points p and q . Here we think of x, y as spatial coordinates and s, t as time coordinates. The two types of coordinates play very different roles, as is evident from the scaling in (1.1). The domain of $\mathcal{L}$ is not equal to all of $\mathbb{R}^4$ since the tilt of $c_1(t - s)n$ in (1.1) forces rescaled distances to tend to $-\infty$ if their time coordinates s, t are in the wrong order. In other words, in the limit we cannot move backwards or instantaneously in time.

Unlike an ordinary metric, $\mathcal{L}$ is not symmetric and may take negative values. It also satisfies the triangle inequality backwards (i.e. because of the sign convention in (1.1)):

$$(1.2) \quad \mathcal{L}(p; r) \geq \mathcal{L}(p; q) + \mathcal{L}(q; r) \quad \text{for all } (p; r), (p; q), (q; r) \in \mathbb{R}_\uparrow^4.$$

¹The negative sign in front of the limit $\mathcal{L}$ here is a convention in the field.

Just as in true metric spaces, we can define path lengths in $\mathcal{L}$, see [15, Section 12]. In the limiting setup, a **path** from (x, s) to (y, t) is a continuous function $\pi : [s, t] \rightarrow \mathbb{R}$ with $\pi(s) = x$ and $\pi(t) = y$. We can define the **length** of a path by

$$(1.3) \quad \|\pi\|_{\mathcal{L}} = \inf_{k \in \mathbb{N}} \inf_{s=t_0 < t_1 < \dots < t_k=t} \sum_{i=1}^k \mathcal{L}(\pi(t_{i-1}), t_{i-1}; \pi(t_i), t_i).$$

This is analogous to defining the length of a curve in Euclidean space by piecewise linear approximation. A path π is a **directed geodesic**, or geodesic for brevity, if $\|\pi\|_{\mathcal{L}}$ is maximal among all paths with the same start and endpoints. Geodesics maximize, rather than minimize, path length because the triangle inequality (1.2) is backwards. Equivalently, a geodesic is any path π with $\|\pi\|_{\mathcal{L}} = \mathcal{L}(\pi(s), s; \pi(t), t)$. Almost surely, directed geodesics exist between every pair of points $(x, s), (y, t)$ with $s < t$. Moreover, there is almost surely a unique geodesic between any fixed pair $(x, s), (y, t)$. In models which converge to the directed landscape, directed geodesics are limits of geodesics in the prelimiting metrics.

Before moving on to the main goal of this paper – understanding disjoint optimizers in the directed landscape – we mention a few relationships between the directed landscape and other well-known limit objects in order to better orient the reader. The first fluctuation limit theorem in the KPZ universality class is the Baik-Deift-Johansson theorem on the length of the longest increasing subsequence in a uniform permutation [3], see also [29] for the same result in a different model. From our present perspective, these theorems show that for any fixed point $u = (x, s; y, t) \in \mathbb{R}_{\uparrow}^4$, we have

$$\mathcal{L}(u) \stackrel{d}{=} (t-s)^{1/3} T - \frac{(x-y)^2}{t-s},$$

where T is a Tracy-Widom GUE random variable. In other words, these results identify the one-point distributions of $\mathcal{L}$.

Shortly afterwards, Prähofer and Spohn [43] found a richer one-parameter scaling limit while studying the polynuclear growth model. Again, from our perspective this theorem amounts to identifying the process

$$y \mapsto \mathcal{L}(0, 0; y, t) + y^2$$

as a stationary Airy₂ process. The work of Prähofer and Spohn also introduces a system of functions $\mathcal{A} = \{\mathcal{A}_i : \mathbb{R} \rightarrow \mathbb{R}, i \in \mathbb{N}\}$ with $\mathcal{A}_1$ being the Airy₂ process. This system was shown to consist of locally Brownian, ordered curves $\mathcal{A}_1 > \mathcal{A}_2 > \dots$ by Corwin and Hammond [8]. The collection $\mathcal{A}$ is called the **Airy line ensemble**, and it is the crucial integrable input needed in the construction of the directed landscape. In particular, the marginal $\mathcal{L}(\cdot, s; \cdot, t)$ is expressed in terms of an Airy line ensemble, see Chapter 2.5 for more details.

Finally, there are exact formulas for marginals of the form $\mathcal{L}(0, 0; y_i, t_i), i = 1, \dots, k$ from Johansson and Rahman [30] and Liu [33], as well as formulas for the Markov process

$$\mathfrak{h}_t(y) = \max_{x \in \mathbb{R}} f(x) + \mathcal{L}(x, 0; y, t)$$

for any upper semicontinuous function $f : \mathbb{R} \rightarrow \mathbb{R} \cup \{-\infty\}$ satisfying a certain growth condition. This Markov process is the KPZ fixed point, constructed by Matetski, Quastel, and Remenik [37].

1.3. Two perspectives on last passage percolation

The goal of this paper is to study and understand the structure of maximal length collections of disjoint paths in the directed landscape. This may initially seem like a rather peripheral object of study in a random metric. However, it turns out that optimal collections of disjoint paths plays a central role in random planar geometry. To properly motivate the study of these objects, we must first return to the prelimit and describe **Brownian last passage percolation** (henceforth Brownian LPP).

Let $f = \{f_i : i \in \mathbb{Z}\}$ be a sequence of continuous functions. For a nonincreasing cadlag function π from $[x, y]$ to the integer interval $\llbracket m, n \rrbracket$ with $\pi(y) = m$, henceforth a **path** from (x, n) to (y, m) , define the length of π with respect to f by

$$\|\pi\|_f = \sum_{i=m}^n f_i(z_i) - f_i(z_{i+1}).$$

Here $z_m = y$ and for $i > m$, z_i is the first time when π is less than i . We remark that we use the indexing convention where the path π is nonincreasing (rather than nondecreasing) to be consistent with [15]; this indexing convention is also convenient to work with in the limit transition to the Airy line ensemble, see Figure 1.4.

For $x \leq y$ and integers $m \leq n$, define the **last passage value**

$$(1.4) \quad f[(x, n) \rightarrow (y, m)] = \sup_{\pi} \|\pi\|_f,$$


where the supremum is over all paths from (x, n) to (y, m) . A function π that achieves this supremum is called a **geodesic**. When the function f is a collection of independent two-sided standard Brownian motions $B = \{B_i : i \in \mathbb{Z}\}$, this model is Brownian last passage percolation LPP

$$(x, n; y, m) \mapsto B[(x, n) \rightarrow (y, m)].$$

Going back to the work of Logan and Shepp [35] and Vershik and Kerov [50] on longest increasing subsequences, much of the progress on understanding integrable LPP models has come by understanding the Robinson-Schensted-Knuth (RSK) bijection. One direction of the classical RSK bijection maps an array of numbers to a pair of semistandard Young tableaux of the same shape. This pair of Young tableaux is built out of differences of certain *multi-point* last passage values. In the context of last passage percolation across a sequence of functions $f = (f_1, \dots, f_n)$ with domain $[0, t]$, these multi-point last passage values are precisely the data

$$(1.5) \quad \begin{aligned} f[(0^k, n) \rightarrow (y^k, m)] &:= \sup_{\pi} \sum_{i=1}^k \|\pi_i\|_f, \\ (y, m) &\in [0, t] \times \{1\} \cup \{t\} \times \llbracket 1, n \rrbracket, \quad k \in \llbracket 1, n - m + 1 \rrbracket. \end{aligned}$$

The supremum is over all k -tuples of disjoint paths $\pi = (\pi_1, \dots, \pi_k)$ from $(0, n)$ to (y, m) . Here and throughout the paper we write $x^k = (x, \dots, x) \in \mathbb{R}^k$ for $x \in \mathbb{R}$. In other words, one direction of this bijection records all multi-point last passage values from $(0, n)$, the bottom corner of the box $[0, t] \times \llbracket 1, n \rrbracket$, to points on the two far sides. It turns out that the whole function f can be reconstructed from this data. Given the importance of the RSK bijection, it is natural to ask what becomes of it in the directed landscape limit, and how it relates to the finite RSK bijection.

FIGURE 1.2. A disjoint optimizer for $k = 3$ from $((0, 0, 0.2), 5)$ to $((0.7, 0.9, 1), 1)$.

On the nonintegrable side, going back at least to the work of Licea and Newman [34] on first passage percolation, the joint structure of geodesics in random metric models has been an object of fruitful study. Questions about geodesic coalescence and disjointness are closely linked with questions about limit shapes, fluctuation exponents, and the structure of shocks in related growth models. More recently, geodesic coalescence and disjointness have been studied in the more tractable context of integrable last passage percolation by using probabilistic and geometric techniques, e.g. see Hammond [27]; Pimentel [42]; Basu, Sarkar, Sly and Zhang [7, 52]; Balázs, Busani, Georgiou, Rassoul-Agha, Seppäläinen, Shen [2, 23, 47]. Questions of geodesic coalescence and disjointness still make sense in the directed landscape, and studying these reveals interesting probabilistic structures, e.g. see Bates, Ganguly, and Hammond [4].

One way to think about problems of geodesic disjointness and coalescence is in terms of certain **multi-point last passage values** that generalize (1.5). For collections of points $\mathbf{x} = (x_1 \leq x_2 \leq \dots \leq x_k)$ and $\mathbf{y} = (y_1 \leq \dots \leq y_k)$, define

$$(1.6) \quad f[(\mathbf{x}, n) \rightarrow (\mathbf{y}, m)] := \sup_{\pi} \sum_{i=1}^k \|\pi_i\|_f,$$

where the supremum is over all k -tuples of disjoint paths $\pi = (\pi_1, \dots, \pi_k)$, where each π_i goes from (x_i, n) to (y_i, m) . We call a k -tuple π that achieves this supremum a **disjoint optimizer**, abbreviated as optimizer. See Figure 1.2 for an example of these definitions and Chapter 2.1 for a more precise setup. If there are disjoint geodesics π_i from x_i to y_i for $i = 1, \dots, k$, then $f[(\mathbf{x}, n) \rightarrow (\mathbf{y}, m)] = \sum_{i=1}^k f[(x_i, n) \rightarrow (y_i, m)]$. On the other hand, if for any collection of k geodesics from x_i to y_i , at least 2 must coalesce on some interval, then $f[(\mathbf{x}, n) \rightarrow (\mathbf{y}, m)] < \sum_{i=1}^k f[(x_i, n) \rightarrow (y_i, m)]$.

The following definition gives the analogue of multi-point last passage percolation in the directed landscape. This paper is devoted to studying this analogue, in

order to shed light on both the limit of RSK and the structure of geodesic disjointness and coalescence in $\mathcal{L}$.

DEFINITION 1.1. Let $\mathfrak{X}_\uparrow$ be the space of all points $(\mathbf{x}, s; \mathbf{y}, t)$, where $s < t$ and $\mathbf{x}, \mathbf{y}$ lie in the same space $\mathbb{R}_\leq^k = \{\mathbf{x} \in \mathbb{R}^k : x_1 \leq \dots \leq x_k\}$ for some $k \in \mathbb{N}$. For $(\mathbf{x}, s; \mathbf{y}, t) \in \mathfrak{X}_\uparrow$, define

$$(1.7) \quad \mathcal{L}(\mathbf{x}, s; \mathbf{y}, t) = \sup_{\pi_1, \dots, \pi_k} \sum_{i=1}^k \|\pi_i\|_{\mathcal{L}}.$$

Here and throughout we use the convention that k is such that $\mathbf{x}, \mathbf{y} \in \mathbb{R}_\leq^k$. The supremum is over all k -tuples of paths $\pi = (\pi_1, \dots, \pi_k)$ where each π_i is a path from (x_i, s) to (y_i, t) , and the paths satisfy the disjointness condition $\pi_i(r) \neq \pi_j(r)$ for all $i \neq j$ and $r \in (s, t)$. We call such a collection π a **disjoint k -tuple** from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$. We call the extension of $\mathcal{L}$ from $\mathfrak{X}_\uparrow \rightarrow \mathbb{R} \cup \{-\infty\}$ the **extended directed landscape**, abbreviated as extended landscape.

See Figure 1.3 for an illustration of Definition 1.1. Note that $\mathbb{R}_\uparrow^4 \subset \mathfrak{X}_\uparrow$, and since geodesics in $\mathcal{L}$ always exist, definition (1.7) on $\mathbb{R}_\uparrow^4$ coincides with the usual definition of $\mathcal{L}$. In the course of this paper, we will show that:

- (1) Just as the directed landscape is the limit of single-point Brownian LPP, the extended landscape is the scaling limit of multi-point Brownian LPP.
- (2) For any $s < t$, the function $(\mathbf{x}, \mathbf{y}) \mapsto \mathcal{L}(\mathbf{x}, s; \mathbf{y}, t)$ can be expressed in terms of a more tractable object: the parabolic Airy line ensemble. This makes $\mathcal{L}(\cdot, s; \cdot, t)$ more amenable to probabilistic analysis.
- (3) The supremum in (1.6) is always attained, and so

$$\mathcal{L}(\mathbf{x}, s; \mathbf{y}, t) = \sum_{i=1}^k \mathcal{L}(x_i, s; y_i, t)$$

if and only if there are geodesics π_i from (x_i, s) to (y_i, t) , $i = 1, \dots, k$ that are disjoint on (s, t) . When combined with point 2, this gives a formula for understanding geodesic disjointness and coalescence that uses only a single parabolic Airy line ensemble.

- (4) One direction of the RSK bijection passes to the limit.


1.4. Brownian LPP and the extended Airy sheet

To understand the extended landscape, we need to go back to understand multi-point LPP in the prelimit. We first focus on understanding the scaling limit of multi-point Brownian LPP from line n to line 1 as $n \rightarrow \infty$.

THEOREM 1.2. Let $B = \{B_i : i \in \mathbb{Z}\}$ be a collection of independent two-sided standard Brownian motions. Let $\mathfrak{X} = \bigcup_{k=1}^\infty \mathbb{R}_\leq^k \times \mathbb{R}_\leq^k$. For $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}_\leq^k \times \mathbb{R}_\leq^k$, define

$$\mathcal{S}^n(\mathbf{x}, \mathbf{y}) = n^{1/6} \left(B[(2n^{-1/3}\mathbf{x}, n) \rightarrow (1 + 2n^{-1/3}\mathbf{y}, 1)] - 2k\sqrt{n} - n^{1/6} \sum_{i=1}^k 2(y_i - x_i) \right),$$

Then $\mathcal{S}^n \xrightarrow{d} \mathcal{S}$ for some random continuous function $\mathcal{S} : \mathfrak{X} \rightarrow \mathbb{R}$. The underlying topology here is uniform convergence on compact subsets of $\mathfrak{X}$. The limit $\mathcal{S}$ is called the **extended Airy sheet**.

FIGURE 1.3. A disjoint k -tuple.

Certain marginals of the extended Airy sheet are familiar. Indeed, let $\mathcal{A} = \{\mathcal{A}_i : \mathbb{R} \rightarrow \mathbb{R}, i \in \mathbb{N}\}$ be the Airy line ensemble and let $\mathcal{B}_i(x) = \mathcal{A}_i(x) - x^2$ be the **parabolic Airy line ensemble**. Then the system $\mathcal{B} = \{\mathcal{B}_i : \mathbb{R} \rightarrow \mathbb{R}, i \in \mathbb{N}\}$ can be coupled with $\mathcal{S}$ so that

$$(1.8) \quad \sum_{i=1}^k \mathcal{B}_i(y) = \mathcal{S}(0^k, y^k)$$

for all $k \in \mathbb{N}, y \in \mathbb{R}$.

The usual Airy sheet, constructed in [15], is given by $\mathcal{S}|_{\mathbb{R}^2}$. It is the scaling limit of single-point last passage values from line n to line 1. The construction of the Airy sheet in [15] relies on showing that the **half-Airy sheet** $\mathcal{S}|_{[0, \infty) \times \mathbb{R}}$ is equal to $h(\mathcal{B})$ for an explicit function h . The function h is defined in terms of a last passage problem involving the parabolic Airy line ensemble, see Chapter 2.5 and Section 1.8 for some discussion of how this description arises from an identity in the prelimit. Our Theorem 1.2 also relies on characterizing $\mathcal{S}$ in terms of last passage percolation across $\mathcal{B}$. Doing so requires formalizing a notion of last passage percolation along infinite paths across $\mathcal{B}$.

For $x \in [0, \infty), z \in \mathbb{R}$, we say that a nonincreasing cadlag function $\pi : (-\infty, z] \rightarrow \mathbb{N}$ is a **parabolic path** from x to z if

$$\lim_{y \rightarrow -\infty} \frac{\pi(y)}{2y^2} = x.$$

See Figure 1.4 for an example. For a parabolic Airy line ensemble $\mathcal{B}$ with corresponding half-Airy sheet $h(\mathcal{B}) : [0, \infty) \times \mathbb{R} \rightarrow \mathbb{R}$, define the path length

$$\|\pi\|_{\mathcal{B}} = h(\mathcal{B})(x, z) + \lim_{y \rightarrow -\infty} \left(\|\pi|_{[y, z]}\|_{\mathcal{B}} - \mathcal{B}[(y, \pi(y)) \rightarrow (z, 1)] \right).$$

See Chapter 4 for more context regarding this definition. For $(\mathbf{x}, \mathbf{y}) \in \mathfrak{X}$ with $x_1 \geq 0$, we can then define the (multi-point) **last passage value**

$$(1.9) \quad \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] = \sup_{\pi_1, \dots, \pi_k} \sum_{i=1}^k \|\pi_i\|_{\mathcal{B}},$$




FIGURE 1.4. A parabolic path across the parabolic Airy line ensemble.

where the supremum is over k -tuples of parabolic paths from x_i to y_i that are disjoint away from the right endpoints y_i .

THEOREM 1.3. *The extended Airy sheet $\mathcal{S}$ satisfies the following properties:*

- $\mathcal{S}$ is shift invariant. More precisely, for $(\mathbf{x}, \mathbf{y}) \in \mathfrak{X}$ and $c \in \mathbb{R}$, let $T_c(\mathbf{x}, \mathbf{y}) = (x_1 + c, \dots, x_k + c, y_1 + c, \dots, y_k + c)$. Then $\mathcal{S} \stackrel{d}{=} \mathcal{S} \circ T_c$ for all $c \in \mathbb{R}$.
- $\mathcal{S}$ can be coupled with a parabolic Airy line ensemble $\mathcal{B}$ so that

$$(1.10) \quad \mathcal{S}(\mathbf{x}, \mathbf{y}) = \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$$

for all $(\mathbf{x}, \mathbf{y}) \in \mathfrak{X}$ with $x_1 \geq 0$.

Moreover, the law of $\mathcal{S}$ is the unique distribution on continuous functions on $\mathfrak{X}$ satisfying these properties.

The parabolic Airy line ensemble $\mathcal{B}$ in the coupling in Theorem 1.3 can be recovered from $\mathcal{S}$ via (1.8). While the definition of path length and last passage percolation across $\mathcal{B}$ are fairly involved, they are still workable. In Chapters 4 and 5 we prove basic properties of these structures that help make (1.10) a useful representation of the extended Airy sheet. As part of this work, we show that $\mathcal{S}(0^k, \mathbf{y})$ has a particularly accessible structure depending only on the top k lines of $\mathcal{B}$ in the compact set $[y_1, y_k]$ (see Proposition 5.9). We also prove certain symmetries of $\mathcal{S}$ (Lemma 5.5), a two-point tail bound (Lemma 5.6) that shows $\mathcal{S}$ is Hölder- $(1/2)^-$, and a metric composition law (Proposition 5.10).

1.5. The full scaling limit of multi-point Brownian LPP

In [15], the directed landscape is built out of independent Airy sheets via a metric composition law inherited from Brownian LPP. The authors then show that this describes the full scaling limit of single-point Brownian LPP. A similar procedure allows us to quickly construct the full scaling limit of multi-point Brownian LPP. For this next theorem, we say $\mathcal{S}_s$ is an extended Airy sheet of scale s if

$$\mathcal{S}_s(\mathbf{x}, \mathbf{y}) \stackrel{d}{=} s\mathcal{S}(s^{-2}\mathbf{x}, s^{-2}\mathbf{y})$$

jointly in all $\mathbf{x}, \mathbf{y}$.

THEOREM 1.4. *There is a unique (in law) random continuous function $\mathcal{L}^* : \mathfrak{X}_\uparrow \rightarrow \mathbb{R}$ such that*

- *For any $(\mathbf{x}, s; \mathbf{y}, t) \in \mathfrak{X}_\uparrow$ and $r \in (s, t)$, almost surely,*

$$\mathcal{L}^*(\mathbf{x}, s; \mathbf{y}, t) = \max_{\mathbf{z}} \mathcal{L}^*(\mathbf{x}, s; \mathbf{z}, r) + \mathcal{L}^*(\mathbf{z}, r; \mathbf{y}, t).$$

Here the maximum is over all $\mathbf{z} \in \mathbb{R}_{\leq}^k$, where k is the cardinality of $\mathbf{x}$ and $\mathbf{y}$.

- *For any finite collection of disjoint time intervals $(t_i, t_i + s_i^3)$, the functions $\mathcal{L}^*(\cdot, t_i; \cdot, t_i + s_i^3)$ are independent extended Airy sheets of scale s_i .*

THEOREM 1.5. *Let $(\mathbf{x}, s)_n = (s + 2\mathbf{x}n^{-1/3}, -\lfloor sn \rfloor)$, and define*

$$\mathcal{L}_n(\mathbf{x}, t; \mathbf{y}, s) = n^{1/6} \left(B[(\mathbf{x}, s)_n \rightarrow (\mathbf{y}, t)_n] - 2k(t - s)\sqrt{n} - n^{1/6} \sum_{i=1}^k 2(y_i - x_i) \right).$$

Then $\mathcal{L}_n \xrightarrow{d} \mathcal{L}^$, with $\mathcal{L}^*$ as in Theorem 1.4. Here the underlying topology is uniform convergence on compact subsets of $\mathfrak{X}_\uparrow$.*

We can think of Theorems 1.4 and 1.5 as an alternate way of constructing an extended directed landscape by first going back to the prelimit. The advantage of having done this is that the definition of the extended Airy sheet that underlies $\mathcal{L}^*$ is much more tractable than Definition 1.1 for $\mathcal{L}$. However, it is not clear from their constructions that $\mathcal{L}$ and $\mathcal{L}^*$ represent the same object. Much of the second half of the paper is devoted to showing this.

THEOREM 1.6. $\mathcal{L}^* = \mathcal{L}$.

The key difficulty in proving Theorem 1.6 is in showing that disjoint optimizers in (1.6) remain disjoint after passing to the limit. As an upshot of the proof of this fact, we show that the supremum (1.7) is always attained.

THEOREM 1.7. *Almost surely, the supremum in (1.7) is attained for every $\mathbf{u} = (\mathbf{x}, s; \mathbf{y}, t) \in \mathfrak{X}_\uparrow$ by some disjoint k -tuple π . We call π a **disjoint optimizer** for $\mathbf{u}$ in $\mathcal{L}$. Moreover, for any fixed $\mathbf{u} \in \mathfrak{X}_\uparrow$, almost surely there is a unique disjoint optimizer $\pi_{\mathbf{u}}$ for $\mathbf{u}$ in $\mathcal{L}$.*

Given that $\mathcal{L}^* = \mathcal{L}$, we can show that optimizers in the prelimit converge to optimizers in the limit. This theorem is the analogue of [15, Theorem 1.8].

THEOREM 1.8. *With $\mathcal{L}_n$ and $\mathcal{L}^* = \mathcal{L}$ as in Theorem 1.5, consider a coupling where $\mathcal{L}_n \rightarrow \mathcal{L}$ almost surely uniformly on compact subsets of $\mathfrak{X}_\uparrow$. For $\mathbf{u} = (\mathbf{x}, s; \mathbf{y}, t) \in \mathfrak{X}_\uparrow$, let $C_{\mathbf{u}}$ be the set of probability 1 where there is a unique disjoint optimizer $\pi = (\pi_1, \dots, \pi_k)$ for $\mathbf{u}$ in $\mathcal{L}$.*

In this coupling, there exists a set Ω of probability 1, such that the following holds. Consider any sequence of points $\mathbf{u}_n = (\mathbf{a}_n, m_n; \mathbf{b}_n, \ell_n)$ which rescale to $\mathbf{u}$ in the setup of Theorem 1.5. That is,

$$\left(\frac{n^{-2/3}m_n + n^{1/3}\mathbf{a}_n}{2}, -\frac{m_n}{n}; \frac{n^{-2/3}\ell_n + n^{1/3}\mathbf{b}_n}{2}, -\frac{\ell_n}{n} \right) \rightarrow \mathbf{u}.$$

Also consider any sequence of disjoint optimizers $\pi^{(n)} = (\pi_1^{(n)}, \dots, \pi_k^{(n)})$ for $\mathbf{u}_n$ across the Brownian motions that give rise to $\mathcal{L}_n$. Let $h_{n,i}$ be the order-preserving,

linear function mapping $[s, t]$ onto $[a_{n,i}, b_{n,i}]$. Then on $\Omega \cap C_{\mathbf{u}}$, for all $1 \leq i \leq k$, we have

$$\frac{\pi_i^{(n)} \circ h_{n,i} + nh_{n,i}}{2n^{2/3}} \rightarrow \pi_i$$

uniformly as functions from $[s, t]$ to $\mathbb{R}$.

In our exploration of the extended landscape, we also find continuity properties analogously to known properties for the directed landscape. The extended landscape is Hölder- $(1/3)^-$ in time (a consequence of Lemma 6.5), Hölder- $(1/2)^-$ in space (a consequence of Lemma 5.6) and its optimizers are Hölder- $(2/3)^-$ (a consequence of Lemma 6.8).

1.6. Consequences

The structure of the extended landscape established in the previous theorems allows us to use the object to understand the limiting analogue of the RSK bijection, and the structure of geodesic disjointness and coalescence.

We start with the RSK bijection. If we apply the RSK bijection to a random array or a sequence of continuous functions, then the KPZ scaling limit of the resulting pair of Young tableaux is a single parabolic Airy line ensemble $\mathcal{B}$. On the other hand, the KPZ scaling limit of the array itself is the directed landscape, with times restricted to the interval $[0, 1]$.

As a consequence of our work, we show that the limiting parabolic Airy line ensemble can be reconstructed from the directed landscape restricted to times in $[0, 1]$ via the natural limiting analogue of RSK. This shows that one direction of the RSK bijection survives into the limit.

COROLLARY 1.9. *Let $\mathcal{L}$ be the directed landscape restricted to the set $\{(x, s; y, t) : x, y \in \mathbb{R}, 0 \leq s < t \leq 1\} \subset \mathbb{R}_+^4$. Then there is a function f such that $f(\mathcal{L}) = \mathcal{B}$, where $\mathcal{B}$ is a parabolic Airy line ensemble. More precisely,*

$$\sum_{i=1}^k \mathcal{B}_i(y) = \mathcal{L}(0^k, 0; y^k, 1),$$

where the right-hand side is an extended landscape value defined from $\mathcal{L}$ as in Definition 1.1.

It is natural to ask whether the RSK map in Corollary 1.9 is still invertible in the limit. We believe that almost surely, this is the case.

CONJECTURE 1.10. *There is an analogue of the RSK bijection in the KPZ limit. More precisely, let f be as in Corollary 1.9, let $\mathcal{B}$ be a parabolic Airy line ensemble and let $\mathcal{L}$ be a directed landscape restricted to times in the interval $[0, 1]$. Then there exists a function g such that almost surely, $f \circ g(\mathcal{B}) = \mathcal{B}$ and $g \circ f(\mathcal{L}) = \mathcal{L}$.*

While we expect that such a function g exists, we do not expect it to resemble the inverse of the usual RSK bijection; this inverse no longer makes sense in the limit. Rather, we believe that such a g should exist because of certain almost sure probabilistic properties of $\mathcal{L}$ (e.g. laws of large numbers, 0–1 laws).

Our work on the extended landscape gives the following criterion for geodesic disjointness and coalescence.

COROLLARY 1.11. *Almost surely the following holds. For every $(\mathbf{x}, s; \mathbf{y}, t) \in \mathfrak{X}_\dagger$,*

$$(1.12) \quad \mathcal{L}(\mathbf{x}, s; \mathbf{y}, t) = \sum_{i=1}^k \mathcal{L}(x_i, s; y_i, t)$$

if and only if there exist $\mathcal{L}$ -geodesics $\pi_1, \dots, \pi_k$ where π_i goes from (x_i, s) to (y_i, t) , satisfying $\pi_i(r) < \pi_{i+1}(r)$ for all $i \in \llbracket 1, k-1 \rrbracket$ and $r \in (s, t)$.

For a fixed s, t , equation (1.12) is an equation about a single extended Airy sheet. In particular, by Theorem 1.3 it can be tackled by understanding a last passage problem across the parabolic Airy line ensemble $\mathcal{B}$. Because of the semi-discrete and locally Brownian nature of $\mathcal{B}$, understanding this problem is easier than understanding geodesic disjointness and coalescence in $\mathcal{L}$ directly.

1.7. More related work

We do not attempt to give a full history of work on the KPZ universality class, and will instead focus on the circle of ideas most closely related to the present work. For a gentle introduction to the KPZ universality class suitable for a newcomer to the area, see Romik [46]. Review articles and books focusing on more recent developments include Corwin [10]; Ganguly [20]; Ferrari and Spohn [19]; Quastel [45]; Weiss, Ferrari, and Spohn [51]; and Zygouras [53].

Many of the initial breakthroughs in the area of KPZ relied on understanding integrable models via the RSK bijection. These include the previously discussed papers of [3, 29, 43] that establish Tracy-Widom convergence and Airy process convergence. A connection between the RSK correspondence and systems of nonintersecting random walks or Brownian motions was later discovered and understood in a series of papers by O’Connell and coauthors [1, 32, 39, 40]. A description of the RSK correspondence in [1] leads to an identity for last passage percolation, see (1.13) below, which is the prelimiting version of the description of the Airy sheet in terms of the Airy line ensemble. Going beyond RSK, newer integrable ideas have yielded a richer set of formulas for limit objects, e.g. see [30, 33, 37, 44].

The works discussed above provide a strong integrable framework for understanding the directed landscape. More recently, probabilistic and geometric methods have been used in conjunction with a few key integrable inputs to prove regularity results, convergence statements, and exponent estimates in such models.

Corwin and Hammond [8] showed that the parabolic Airy line ensemble $\mathcal{B}$ satisfies a certain Brownian Gibbs property, making it amenable to probabilistic analysis. Hammond [25, 26, 28]; Dauvergne and Virág [17]; and Calvert, Hammond, and Hegde [9] used Brownian Gibbs analysis to quantitatively understand the Brownian nature of the parabolic Airy line ensemble. The parabolic Airy line ensemble plays a central role in our paper, and we will require several consequences of this research program. Having a strong understanding of the Brownian nature of $\mathcal{B}$ is what makes results like Theorem 1.2 and Corollary 1.11 useful in practice.

There are many other papers that use Brownian Gibbs analysis and related ideas to study the structure of geodesics, near geodesics, and disjoint optimizers in the directed landscape and other last passage models. Some prominent recent examples include Hammond [27]; Ganguly and Hammond [21, 22]; Basu, Ganguly, and Zhang [6]; Sarkar, Dauvergne, and Virág [16]; and Bates, Ganguly, and Hammond [4].

Beyond [15], perhaps the two papers most closely linked with our own are [48] and [5]. In [48], Sarkar and Virág show Brownian absolute continuity of the KPZ fixed point. One key idea in their work is to construct infinite last passage geodesics across the parabolic Airy line ensemble. Their setup for doing this is different than the setup we require for Theorem 1.2, but still based around the Airy sheet construction in [15]. In [5], Basu, Ganguly, Hammond, and Hegde study the geometry of disjoint optimizers between k identical start and endpoints for lattice last passage models, or “geodesic watermelons”. They find scaling exponents in k for the total length and transversal fluctuations of these optimizers.

Results, techniques, and frameworks developed in this paper have already been used to analyze the Airy sheet, and geodesics across the parabolic Airy line ensemble and in the directed landscape. For example, the work [13] which came out after the first version of this paper appeared online, uses the framework of this paper to relate marginals of the Airy sheet to marginals in Brownian last passage percolation. The work [24] (which also came out after the first version of this paper) about fractal geometry in the directed landscape requires an understanding of coalescence and disjointness between various pairs of landscape geodesics, which are equivalent to relations between infinite last passage geodesics across the parabolic Airy line ensemble. The analysis in [24] uses the framework of last passage percolation across the parabolic Airy line ensemble in Chapter 4, and the existence of disjoint optimizers across the parabolic Airy line ensemble proven in Chapter 5. Corollary 1.11 and related ideas will be used to analyze disjointness of $\mathcal{L}$ -geodesics in the forthcoming work [12].

1.8. Outline of the paper and a primer about the proofs

While the structure of the paper is similar to [15], the proofs are mostly distinct. Indeed, the main difficulties that were resolved in [15] yield lemmas that can be applied immediately here without need for generalization. As a consequence, the main difficulties in our work are unique to the multi-point setting and require different types of ideas. In this outline, we emphasize the differences between the two papers and some of the additional difficulties in multi-point setting. Generally, Chapters 6 and 9 follow a similar flow to corresponding sections in [15], and Chapters 4, 5, 7, and 8 contain the most novel ideas. Chapter 2 is a blend of background and new deterministic results for multi-point LPP, and Chapter 3 applies these multi-point LPP results to prove tightness for key objects.

The first half of the paper (Chapters 2 to 5) is devoted to constructing the extended Airy sheet. This is the part of the paper that leans most heavily on technical machinery from [15], and so to appreciate these chapters we recommend that the reader have some familiarity with the construction of the Airy sheet from [15]. See Chapter 2.5 for details on the exact inputs we use.

The starting point for the construction of the extended Airy sheet is a combinatorial identity about the RSK bijection. In essence, this identity shows that given a collection of functions $f = (f_1, \dots, f_n)$, we can construct a collection of ordered functions $Wf = (Wf_1 \geq \dots \geq Wf_n)$ with $Wf(0) = (0, \dots, 0)$ such that

$$(1.13) \quad f[(\mathbf{x}, n) \rightarrow (\mathbf{y}, 1)] = Wf[(\mathbf{x}, n) \rightarrow (\mathbf{y}, 1)],$$

for all $\mathbf{x}, \mathbf{y}$ with $x_1 \geq 0$. We refer to Wf as the melon of f , as ordered paths in Wf emanating from 0 resemble stripes on a watermelon. Versions of this identity

go back to [38] and [1]. When f is given by a collection of independent Brownian motions, then Wf is given by a collection of nonintersecting Brownian motions. In the scaling window we care about, the top lines of Wf converge to the parabolic Airy line ensemble $\mathcal{B}$. What Theorems 1.2 and 1.3 say (in particular, equation (1.10)) is that in this scaling, the identity (1.13) also passes to the limit. The left-hand side becomes the extended Airy sheet and the right-hand side becomes a last passage problem along parabolic paths in $\mathcal{B}$.

At the level of single points x, y , this limiting picture was developed in [15] to construct the usual Airy sheet. However, the construction of the Airy sheet does not require a well-developed notion of last passage percolation along infinite paths across the parabolic Airy line ensemble. We develop this theory in Chapters 4 and 5, expanding on the discussion prior to Theorem 1.3 above. Note that the theory of LPP along infinite paths has subtleties that are not present in the finite case. For example, it is not straightforward to show that the function $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$ is almost surely finite or continuous in $\mathbf{x}$ and $\mathbf{y}$, see Proposition 5.8.

To take advantage of this theory and prove Theorem 1.2, we need to prove tightness of both the extended sheets $\mathcal{S}^n$ and optimizers across the Brownian melon. To avoid obtaining new analytic estimates here, we take advantage of a variety of useful *quadrangle inequalities and monotonicity properties* for multi-point LPP that generalize corresponding properties for single-point LPP, see Chapter 2.2 and Lemma 5.7. These inequalities allow us to quickly deduce tightness and a modulus of continuity for the extended Airy sheet from bounds on the prelimiting Airy line ensembles and tightness of melon optimizers from tightness and coalescence properties of melon geodesics, see Chapter 3. These deterministic properties continue to appear as crucial tools throughout the paper. The construction of the extended Airy sheet (Theorems 1.2 and 1.3) is the culmination of Chapters 2-5.

The remainder of the paper proves Theorems 1.5, 1.6, 1.7, 1.8 and Corollaries 1.9 and 1.11. This part of the paper does not use technical machinery from [15], though as mentioned previously, Chapters 6 and 9 follow a similar flow of ideas to [15].

The limit $\mathcal{L}^*$ of Brownian LPP can be patched together from extended Airy sheets, just as the directed landscape can be built from Airy sheets. The procedure just requires a few technical estimates. We prove these along with Theorems 1.4 and 1.5, in Chapter 6.

Just as path length can be defined in the directed landscape by (1.3), we can define the length of a continuous multi-path $\pi : [s, t] \rightarrow \mathbb{R}_{\leq}^k$ in $\mathcal{L}^*$ by setting

$$\|\pi\|_{\mathcal{L}^*} = \inf_{m \in \mathbb{N}} \inf_{s=t_0 < t_1 < \dots < t_m=t} \sum_{i=1}^m \mathcal{L}^*(\pi(t_{i-1}), t_{i-1}; \pi(t_i), t_i).$$

We say that π is an optimizer from $(\pi(s), s)$ to $(\pi(t), t)$ if $\|\pi\|_{\mathcal{L}^*} = \mathcal{L}^*(\pi(s), s; \pi(t), t)$. Preliminary results about paths and length in $\mathcal{L}^*$ are developed in Chapter 7. Again, there are some subtleties that arise in the study of these objects that do not exist either in the prelimit or in the setting of single paths. For example, unlike for geodesics it is not straightforward that for any $(\mathbf{p}; \mathbf{q}) \in \mathfrak{X}_{\dagger}$ there is almost surely a unique $\mathcal{L}^*$ -optimizer from $\mathbf{p}$ to $\mathbf{q}$. This requires a resampling argument in the parabolic Airy line ensemble, see Chapter 7.2.

To show that the limit $\mathcal{L}^*$ can alternately be described by Definition 1.1, the key step is Proposition 8.1, which shows that almost surely, for every point in $\mathfrak{X}_{\dagger}$

there exists an optimizer in $\mathcal{L}^*$ consisting of *disjoint* paths. This is a three-step process, carried out in Chapter 8.

We first prove Proposition 8.1 for endpoints of the form $((x, x), s), ((y, y), t)$. This is an easier problem since the midpoint of such an optimizer can be characterized using only the top two lines of two independent parabolic Airy line ensembles $\mathcal{B}, \mathcal{B}'$. The key technical point that makes this observation useful is that for any compact set $K \subset \mathbb{R}$ and any $k \in \mathbb{N}$, on K the top k lines of $\mathcal{B}, \mathcal{B}'$ are absolutely continuous with respect to $2k$ independent Brownian motions with a well-controlled Radon-Nikodym derivative, see Theorem 8.10. At the level of any single Airy line, such a Radon-Nikodym derivative estimate was proven in [9]. The extension to multiple lines can be extracted by combining various intermediate lemmas in [9], see Appendix A.2.

Next, we move to endpoints of the more general form $((x_1, x_2), s), ((y_1, y_2), t)$. We do this with a resampling argument which shows that for any $[s', t'] \subset (s, t)$, there is optimizer from $((x_1, x_2), s)$ to $((y_1, y_2), t)$ that coincides on $[s', t']$ with the optimizer from $((0, 0), s - 1)$ to $((0, 0), t + 1)$. Finally, we treat the case of $k \geq 3$ endpoints by induction. The $k = 2$ case is both the base case and the key input for the inductive step.

Given Proposition 8.1, Theorems 1.6 and 1.7 and Corollaries 1.9 and 1.11 follow easily. In a final short chapter (Chapter 9) we give a deterministic argument to prove Theorem 1.8 from Theorems 1.5 and 1.6. This chapter is quite similar to Section 13 of [15], though the arguments have been simplified a bit.

Acknowledgments

The first author would like to thank Bálint Virág for many useful discussions about this project. The authors would also like to thank anonymous referees for their carefully reading of the paper and for providing many valuable comments that helped improve the exposition.

CHAPTER 2

Last passage percolation across lines

In this chapter, we recall and prove combinatorial properties of last passage percolation across lines, and gather necessary limiting results for Brownian LPP. Our presentation aligns with that of [15], where notation and coordinate orientation are set up so that last passage geodesics will rise from the bottom of the page to the top of the page in the Airy line ensemble limit, see Figure 1.4.

Recall from the introduction that a path from (x, n) to (y, m) is a cadlag, nonincreasing function $\pi : [x, y] \rightarrow \llbracket m, n \rrbracket$ with $\pi(y) = m$. We denote the left limit of π at a point t by $\pi(t^-)$. This is defined for all $t \in (x, y]$. We will also extend this to the point x by setting $\pi(x^-) = n$. For any path π , we can define a sequence of **jump times** $x = t_{n+1} \leq t_n \cdots \leq t_{m+1} \in [x, y]$, where

$$t_i = \inf \{t \in [x, y] : \pi(t) < i\}.$$

Typically, this is the jump when π jumps from line i to $i - 1$. We also set $t_m = y$. The **zigzag graph** of π is

$$\Gamma(\pi) = \{(t, k) \in [a, b] \times \llbracket m, n \rrbracket : \pi(t^-) \geq k \geq \pi(t)\}.$$

In other words, the zigzag graph of π connects up the graph of π by vertical lines at its jumps. We can make the set of paths into a topological space – path space – by specifying that $\pi_n \rightarrow \pi$ if $\Gamma(\pi_n) \rightarrow \Gamma(\pi)$ in the Hausdorff topology. Equivalently, $\pi_n \rightarrow \pi$ if the endpoints and jump times of π_n converge to the endpoints and jump times of π . With this definition, the space of all paths from p to q is compact.

We will also introduce a partial order on paths. Let $(p, q) = (x, n; y, m)$, $(p', q') = (x', n'; y', m')$ be such that $x \leq x'$, $y \leq y'$. Then for paths π, π' from p to q and p' to q' respectively, we say that $\pi \leq \pi'$ if for every $t \in [x, y] \cap [x', y']$, we have $\pi(t) \leq \pi'(t)$.

Now consider a sequence of continuous functions $f = (f_i : i \in I)$, where $I \subset \mathbb{Z}$ and each $f_i : \mathbb{R} \rightarrow \mathbb{R}$. We call the space of such functions $\mathcal{C}^I$. We will alternately think of f as a function from $\mathbb{R} \times I$ to $\mathbb{R}$, or as a function from $\mathbb{R} \times \mathbb{Z}$ to $\mathbb{R}$, where f is set equal to 0 outside of its natural domain. When $\llbracket m, n \rrbracket \subset I$, recall from the introduction that the f -length of a path π from (x, n) to (y, m) with jump times t_i is

$$\|\pi\|_f = \sum_{i=m}^n f_i(t_i) - f_i(t_{i+1}).$$

Observe that f -length is a continuous function in path space by the continuity of f . Now, for $(p, q) = (x, n; y, m)$ with $x \leq y$, $n \geq m$ we define the **last passage value**

$$f[p \rightarrow q] = \sup_{\pi} \|\pi\|_f,$$

where the supremum is over all paths π from p to q . Continuity of path length and compactness of the set of paths from p to q ensures that this supremum is always attained.

We call a path that attains the supremum a **geodesic** from $p = (x, n)$ to $q = (y, m)$. We say that π is a **rightmost geodesic** from p to q if $\pi(t) \geq \tau(t)$ for all $t \in [x, y]$ for any other geodesic τ from p to q . We similarly define the leftmost geodesic τ from p to q with the opposite inequality. Note our notion of rightmost and leftmost paths is with respect to the picture in Figure 1.2, where the line order is increasing as we go from top to bottom. Rightmost and leftmost geodesics between two points always exist by a basic compactness and continuity argument in path space, see [15, Lemma 3.5]. Moreover, these paths exhibit a particular tree structure and monotonicity, which can be straightforwardly deduced from their definitions.

PROPOSITION 2.1 ([15, Proposition 3.7]). *Take any $x_1 \leq x_2$ and $y_1 \leq y_2$, and let $\pi^+[x_i, y_i]$ denote the rightmost geodesic from (x_i, n) to $(y_i, 1)$ across a function f . Then $\pi^+[x_1, y_1] \leq \pi^+[x_2, y_2]$ and $\Gamma(\pi^+[x_1, y_1]) \cap \Gamma(\pi^+[x_2, y_2])$ is the zigzag graph of some path whenever this set is nonempty.*

In particular, if $x_1 = x_2$, then the rightmost geodesics to y_1 and y_2 are equal on some interval $[x_1, z]$, and $\pi^+[x_1, y_1](z') < \pi^+[x_2, y_2](z')$ whenever $z' \geq z$ is in the domain of both paths. We can think of the two paths as forming two branches in a tree. The same structure holds with rightmost paths replaced by leftmost paths.

Often, there will be a unique geodesic between two points across the functions that we consider. In this case, the tree structure in Proposition 2.1 will automatically hold; a unique geodesic is both a rightmost and leftmost geodesic.

2.1. Last passage with multiple paths

We can extend the definition of last passage percolation to multiple disjoint paths. We say that π and τ with domains $[a, b]$ and $[a', b']$ are **essentially disjoint** if

- $\pi(t) \neq \tau(t)$ for all $t \in (a, b) \cap (a', b')$
- Either $\pi \leq \tau$ or $\tau \leq \pi$.

Note that since all paths are cadlag, the first condition above is equivalent to the property that the intersection of the closed graphs $\Gamma(\pi) \cap \Gamma(\tau)$ is finite. This characterization will often be more useful for proofs. Essential disjointness is a closed condition: if π_n, τ_n are sequences of essentially disjoint paths converging to paths π, τ , then π and τ are essentially disjoint.

Now, consider vectors $\mathbf{p} = (p_1, \dots, p_k) = ((x_1, n_1), \dots, (x_k, n_k))$ and $\mathbf{q} = (q_1, \dots, q_k) = ((y_1, m_1), \dots, (y_k, m_k))$ in $(\mathbb{R} \times \mathbb{Z})^k$. We say that $(\mathbf{p}, \mathbf{q})$ is an **endpoint pair** of size k , if $n_i \geq m_i$ and $x_i \leq y_i, x_i \leq x_{i+1}, y_i \leq y_{i+1}$ for all i , and there is at least one **disjoint k -tuple (of paths)** from $\mathbf{p}$ to $\mathbf{q}$. Here a **disjoint k -tuple (of paths)** from $\mathbf{p}$ to $\mathbf{q}$ is a vector $\pi = (\pi_1, \dots, \pi_k)$, where

- π_i is a path from (x_i, n_i) to (y_i, m_i) ,
- π_i and π_j are essentially disjoint for all $i \neq j$,
- $\pi_i \leq \pi_j$ for $i < j$.

We put the product topology on the space of all k -tuples of paths: $\pi \rightarrow \tau$ if $\pi_i \rightarrow \tau_i$ for all i . The space of disjoint k -tuples is a closed subset of this space, since essential disjointness and all ordering requirements are closed conditions. As in the single path case, the set of all disjoint k -tuples from $\mathbf{p}$ to $\mathbf{q}$ is compact for any endpoint pair $(\mathbf{p}, \mathbf{q})$.

Now, for a disjoint k -tuple π and $f \in \mathcal{C}^I$, let $\|\pi\|_f = \sum_{i=1}^k \|\pi_i\|_f$. For any endpoint pair $(\mathbf{p}, \mathbf{q})$ and $f \in \mathcal{C}^I$ with $\llbracket m, n \rrbracket \subset I$, define the last passage value

$$f[\mathbf{p} \rightarrow \mathbf{q}] = \sup_{\pi} \|\pi\|_f,$$

where the supremum is over disjoint k -tuples π from $\mathbf{p}$ to $\mathbf{q}$. This supremum is always attained since length is a continuous function in path space and the set of all disjoint k -tuples from $\mathbf{p}$ to $\mathbf{q}$ is compact. A disjoint k -tuple that attains this supremum is a **disjoint optimizer**, abbreviated to optimizer.

For most parts of the paper, we will only be concerned with endpoint pairs where all the n_i are equal to some n , and all the m_i are equal to some m . As a slight abuse of notation we write $\mathbf{p} = (\mathbf{x}, n)$ and $\mathbf{q} = (\mathbf{y}, m)$ in this case.

2.2. Basic properties of disjoint optimizers and last passage values

Disjoint optimizers share certain features with geodesics. In particular, leftmost and rightmost optimizers still exist, and we have monotonicity and a useful quadrangle inequality. Throughout this section we take $f \in \mathcal{C}^I$, for some suitable $I \subset \mathbb{Z}$.

For two disjoint k -tuples of paths π, τ , we say that $\pi \leq \tau$ if $\pi_i \leq \tau_i$ for all i .

LEMMA 2.2. *For any endpoint pair $(\mathbf{p}, \mathbf{q})$, there exists an optimizer $\pi = (\pi_1, \dots, \pi_k)$ from $\mathbf{p}$ to $\mathbf{q}$ such that for any other optimizer τ from $\mathbf{p}$ to $\mathbf{q}$, $\tau \leq \pi$. We call π the **rightmost optimizer** from $\mathbf{p}$ to $\mathbf{q}$. Similarly, there always exists a **leftmost optimizer** from $\mathbf{p}$ to $\mathbf{q}$.*

PROOF. We first show that for any optimizers τ, π from $\mathbf{p}$ to $\mathbf{q}$, there exists optimizers ζ, ζ' from $\mathbf{p}$ to $\mathbf{q}$ such that $\zeta \geq \pi \geq \zeta'$ and $\zeta \geq \tau \geq \zeta'$. For each i, t , set $\zeta_i(t) = \max(\pi_i(t), \tau_i(t))$ and $\zeta'_i(t) = \min(\pi_i(t), \tau_i(t))$. We first check that ζ, ζ' are disjoint k -tuples from $\mathbf{p}$ to $\mathbf{q}$. The arguments are symmetric, so we just check ζ .

It is immediate from the definitions that each ζ_i is a path from p_i to q_i and that $\zeta_i \leq \zeta_j$ whenever $i < j$. Now, for $i \neq j$, if $\zeta_i(t) = \zeta_j(t)$ for some t , then the ordering properties for π, τ ensure that either $\pi_i(t) = \pi_j(t)$ or $\tau_i(t) = \tau_j(t)$. Also, if $\zeta_i(t) = \zeta_j(t)$ for some $t \in (x_i, y_i) \cap (x_j, y_j)$, then since both ζ_i, ζ_j are cadlag, $\zeta_i = \zeta_j$ on some interval $[t, t + \epsilon)$ for some $\epsilon > 0$. Therefore either $\pi_i(t) = \pi_j(t)$ or $\tau_i(t) = \tau_j(t)$ for infinitely many points in this interval, contradicting the essential disjointness of either π_i and π_j , or τ_i and τ_j . Therefore ζ_i, ζ_j must also be essentially disjoint, and so ζ is a disjoint k -tuple from $\mathbf{p}$ to $\mathbf{q}$. Now, by the construction of ζ_i, ζ'_i we have $\|\zeta_i\|_f + \|\zeta'_i\|_f = \|\tau_i\|_f + \|\pi_i\|_f$ for all i . Therefore

$$\|\zeta\|_f + \|\zeta'\|_f = \|\tau\|_f + \|\pi\|_f,$$

and so both ζ and ζ' must also be optimizers from $\mathbf{p}$ to $\mathbf{q}$.

We can complete the proof by appealing to Zorn's lemma. Indeed, the set of optimizers from $\mathbf{p}$ to $\mathbf{q}$ is a partially ordered set. Moreover, this set is compact by the continuity of length in path space, and the fact that the set of all disjoint k -tuples from $\mathbf{p}$ to $\mathbf{q}$ is compact. Therefore by Zorn's lemma, maximal optimizers exist. Finally, if τ, π are two maximal optimizers, then by the argument above there is an optimizer ζ with $\zeta \geq \tau, \zeta \geq \pi$. By maximality, this implies $\zeta = \tau = \pi$ is the unique maximal optimizer: the rightmost optimizer. By a symmetric argument there exists a leftmost optimizer. $\square$

In order to state the monotonicity lemma for multiple paths, we introduce a partial order on endpoint pairs starting on the same line n and ending on the same line m . For two endpoint pairs $(\mathbf{p}, \mathbf{q}) = (\mathbf{x}, n; \mathbf{y}, m)$ and $(\mathbf{p}', \mathbf{q}') = (\mathbf{x}', n; \mathbf{y}', m)$ of size $k = k'$, we say that $(\mathbf{p}, \mathbf{q}) \leq (\mathbf{p}', \mathbf{q}')$ if $x_i \leq x'_i$ and $y_i \leq y'_i$ for all i . If the sizes of the endpoint pairs differ or if we do not have an ordering between all endpoints, then we may still be able to compare the endpoint pairs. For two endpoint pairs $(\mathbf{p}, \mathbf{q})$ and $(\mathbf{p}', \mathbf{q}')$ of size k, k' that start and end on the same line, and $s \in \mathbb{Z}$, define

$$(\mathbf{p}, \mathbf{q}) \leq_s (\mathbf{p}', \mathbf{q}')$$

if $x_{i+s} \leq x'_i$ and $y_{i+s} \leq y'_i$ for all i such that either $i + s \in \llbracket 1, k \rrbracket$ or $i \in \llbracket 1, k' \rrbracket$. Here the coordinates x_j, y_j are defined to be equal to ∞ for $j > k$ and $-\infty$ for $j < 1$, and x'_j, y'_j are defined similarly in terms of k' .

This definition can be thought in the following way. First pad the endpoint pairs $(\mathbf{p}, \mathbf{q})$ and $(\mathbf{p}', \mathbf{q}')$ with points that are arbitrarily far to the right or left so that the indices $i + s$ in $(\mathbf{p}, \mathbf{q})$ and i in $(\mathbf{p}', \mathbf{q}')$ are now lined up and the new endpoint pairs have the same size. The ordering $\leq_s$ is then just the usual ordering $\leq$ on the padded endpoint pairs.

LEMMA 2.3. *Let $(\mathbf{p}, \mathbf{q})$ and $(\mathbf{p}', \mathbf{q}')$ be two endpoint pairs of sizes k, k' starting and ending on the same line. Let π be the rightmost optimizer from $\mathbf{p}$ to $\mathbf{q}$, and π' be the rightmost optimizer from $\mathbf{p}'$ to $\mathbf{q}'$.*

- (i) *Suppose that $k = k'$, and that $(\mathbf{p}, \mathbf{q}) \leq (\mathbf{p}', \mathbf{q}')$. Then $\pi \leq \pi'$.*
- (ii) *Suppose that $(\mathbf{p}, \mathbf{q}) \leq_s (\mathbf{p}', \mathbf{q}')$ for some $s \in \mathbb{Z}$. Then $\pi_i \leq \pi'_{i+s}$ for all $i \in \llbracket 1, k \rrbracket \cap \llbracket 1 - s, k' - s \rrbracket$.*

The same statements hold with leftmost optimizers in place of rightmost ones.

PROOF. We will just prove (i), as (ii) can be reduced to (i) by the padding procedure described above. We use a similar construction to Lemma 2.2. For each i , define paths ζ_i, ζ'_i as follows. On $[x_i, y_i] \cap [x'_i, y'_i]$, set $\zeta_i(t) = \min(\pi_i(t), \pi'_i(t))$ and set $\zeta'_i(t) = \max(\pi_i(t), \pi'_i(t))$. Extend ζ_i to all of $[x_i, y_i]$ by setting it equal to π_i on $[x_i, y_i] \setminus [x'_i, y'_i]$ and extend ζ'_i to all of $[x'_i, y'_i]$ by setting it equal to π'_i on $[x'_i, y'_i] \setminus [x_i, y_i]$.

With these definitions, because $x_i \leq x'_i$ and $y_i \leq y'_i$, ζ_i is a path from p_i to q_i and ζ'_i is a path from p'_i to q'_i . Moreover, exactly as in the proof of Lemma 2.2, we can check that $\zeta = (\zeta_1, \dots, \zeta_k)$ is a disjoint k -tuple from $\mathbf{p}$ to $\mathbf{q}$, ζ' is a disjoint k -tuple from $\mathbf{p}'$ to $\mathbf{q}'$, and

$$\|\zeta\|_f + \|\zeta'\|_f = \|\pi\|_f + \|\pi'\|_f.$$

Therefore ζ, ζ' must both be optimizers. Since $\pi' \leq \zeta'$ and π' is a rightmost optimizer, we have $\zeta' = \pi'$. Also, $\pi \leq \zeta'$ by construction, yielding (i). $\square$

We will also need two **quadrangle inequalities** for multi-point last passage values. These are generalizations of a commonly used quadrangle inequality for single-point last passage values, see for example, Proposition 3.8 in [15].

LEMMA 2.4. *Let $(\mathbf{p}, \mathbf{q}) = (\mathbf{x}, n; \mathbf{y}, m)$, $(\mathbf{p}', \mathbf{q}') = (\mathbf{x}', n; \mathbf{y}', m)$ be endpoint pairs of size k . Define $\mathbf{x}^\ell, \mathbf{y}^\ell, \mathbf{x}^r, \mathbf{y}^r \in \mathbb{R}_{\leq}^k$ by setting $x_i^\ell = x_i \wedge x'_i$, $y_i^\ell = y_i \wedge y'_i$, and $x_i^r = x_i \vee x'_i$, $y_i^r = y_i \vee y'_i$, for each $1 \leq i \leq k$, and let $\mathbf{p}^\ell = (\mathbf{x}^\ell, n)$, $\mathbf{p}^r = (\mathbf{x}^r, n)$, $\mathbf{q}^\ell = (\mathbf{y}^\ell, m)$, $\mathbf{q}^r = (\mathbf{y}^r, m)$. Then*

$$f[\mathbf{p} \rightarrow \mathbf{q}] + f[\mathbf{p}' \rightarrow \mathbf{q}'] \leq f[\mathbf{p}^\ell \rightarrow \mathbf{q}^\ell] + f[\mathbf{p}^r \rightarrow \mathbf{q}^r].$$

In particular, if $(\mathbf{p}, \mathbf{q}') \leq (\mathbf{p}', \mathbf{q})$, then

$$f[\mathbf{p} \rightarrow \mathbf{q}] + f[\mathbf{p}' \rightarrow \mathbf{q}'] \leq f[\mathbf{p} \rightarrow \mathbf{q}'] + f[\mathbf{p}' \rightarrow \mathbf{q}].$$

PROOF. Let π be an optimizer from $\mathbf{p}$ to $\mathbf{q}$, and let π' be an optimizer from $\mathbf{p}'$ to $\mathbf{q}'$. We can define disjoint k -tuples τ^ℓ, τ^r as follows. For each i , set $\tau_i^\ell = \min(\pi_i, \pi'_i)$ on $[x_i^r, y_i^\ell]$ and set $\tau_i^r = \max(\pi_i, \pi'_i)$ on $[x_i^r, y_i^\ell]$. On $[x_i^\ell, x_i^r]$, we set τ_i^ℓ to be either π_i or π'_i , depending on whether x_i^ℓ equals x_i or x'_i . Similarly, on $(y_i^\ell, y_i^r]$, set τ_i^r to be either π_i or π'_i . As in the proof of Lemma 2.2, one can check that τ^ℓ, τ^r are disjoint k -tuples from $\mathbf{p}^\ell$ to $\mathbf{q}^\ell$ and $\mathbf{p}^r$ to $\mathbf{q}^r$, respectively. Therefore

$$f[\mathbf{p} \rightarrow \mathbf{q}] + f[\mathbf{p}' \rightarrow \mathbf{q}'] = \|\pi\|_f + \|\pi'\|_f = \|\tau\|_f + \|\tau'\|_f \leq f[\mathbf{p}^\ell \rightarrow \mathbf{q}^\ell] + f[\mathbf{p}^r \rightarrow \mathbf{q}^r].$$

The second part of the theorem follows from the fact that if $(\mathbf{p}, \mathbf{q}') \leq (\mathbf{p}', \mathbf{q})$, then $\mathbf{p} = \mathbf{p}^\ell, \mathbf{p}' = \mathbf{p}^r, \mathbf{q} = \mathbf{q}^r$, and $\mathbf{q}' = \mathbf{q}^\ell$. $\square$

LEMMA 2.5. Let $(\mathbf{p}, \mathbf{q}), (\mathbf{p}, \mathbf{q}')$ be endpoint pairs of size $k \geq 2$ that start and end on the same line with $(\mathbf{p}, \mathbf{q}) \leq (\mathbf{p}, \mathbf{q}')$. Fix $1 \leq \ell < k$, and let $\mathbf{p}^L, \mathbf{q}^L, \mathbf{q}'^L$ be the first ℓ coordinates of $\mathbf{p}, \mathbf{q}, \mathbf{q}'$, and $\mathbf{p}^R, \mathbf{q}^R, \mathbf{q}'^R$ be the last $k - \ell$ coordinates of $\mathbf{p}, \mathbf{q}, \mathbf{q}'$. Suppose first that $\mathbf{q}^R = \mathbf{q}'^R$. Then

$$f[\mathbf{p} \rightarrow \mathbf{q}] + f[\mathbf{p}^L \rightarrow \mathbf{q}'^L] \geq f[\mathbf{p} \rightarrow \mathbf{q}'] + f[\mathbf{p}^L \rightarrow \mathbf{q}^L].$$

Similarly, suppose that $\mathbf{q}^L = \mathbf{q}'^L$. Then

$$f[\mathbf{p} \rightarrow \mathbf{q}] + f[\mathbf{p}^R \rightarrow \mathbf{q}'^R] \leq f[\mathbf{p} \rightarrow \mathbf{q}'] + f[\mathbf{p}^R \rightarrow \mathbf{q}^R],$$

PROOF. We prove the first inequality since the second one follows similarly. Let π be an optimizer from $\mathbf{p}$ to $\mathbf{q}'$, and let τ be an optimizer from $\mathbf{p}^L$ to $\mathbf{q}^L$. For $i \leq \ell$, we can define paths σ_i by setting $\sigma_i = \min(\pi_i, \tau_i)$ on $[x_i, y_i]$. We also set $\sigma_i = \pi_i$ for $i \in \{\ell + 1, \dots, k\}$. As in the proof of Lemma 2.2, one can check that σ is a disjoint k -tuple from $\mathbf{p}$ to $(\mathbf{q}^L, \mathbf{q}'^R)$. Moreover, $(\mathbf{q}^L, \mathbf{q}'^R) = \mathbf{q}$ since $\mathbf{q}^R = \mathbf{q}'^R$. Similarly set $\sigma'_i = \max(\pi_i, \tau_i)$ on $[x_i, y_i]$ and set $\sigma'_i = \pi_i$ on $(y_i, y'_i]$. Again as in the proof of Lemma 2.2, σ' is a disjoint k -tuple from $\mathbf{p}^L$ to $\mathbf{q}'^L$. Therefore

$$f[\mathbf{p} \rightarrow \mathbf{q}'] + f[\mathbf{p}^L \rightarrow \mathbf{q}^L] = \|\pi\|_f + \|\tau\|_f = \|\sigma\|_f + \|\sigma'\|_f \leq f[\mathbf{p} \rightarrow \mathbf{q}] + f[\mathbf{p}^L \rightarrow \mathbf{q}'^L]. \quad \square$$

We next record three deterministic bounds on multi-point last passage values which will be used to prove tightness. The first bound controls the difference between two last passage values. For this lemma, define the fluctuation of a function $f \in \mathcal{C}^I$ on a set $A \subset \mathbb{R} \times I$ by

$$\omega(f, A) = \sup_{(x, i), (y, i) \in A} |f_i(x) - f_i(y)|.$$

LEMMA 2.6. Let $(\mathbf{p}, \mathbf{q}), (\mathbf{p}, \mathbf{q}')$ be two endpoint pairs of size k with $\mathbf{p} = (\mathbf{x}, n)$ to $\mathbf{q} = (\mathbf{y}, m)$ and $\mathbf{q}' = (\mathbf{y}', m)$ differing only on a single coordinate $y_i < y'_i$. Let π, π' be optimizers from $\mathbf{p}$ to $\mathbf{q}$ and $\mathbf{p}$ to $\mathbf{q}'$. Then

$$f[\mathbf{p} \rightarrow \mathbf{q}'] - f[\mathbf{p} \rightarrow \mathbf{q}] \leq |\pi'_i(y_i) + 1 - m| \omega(f, [y_i, y'_i] \times \llbracket m, \pi'_i(y_i) \rrbracket),$$

$$f[\mathbf{p} \rightarrow \mathbf{q}] - f[\mathbf{p} \rightarrow \mathbf{q}'] \leq (2(k - i) + 1) \omega(f, [y_i, y'_i] \times \llbracket m, m + k - i \rrbracket).$$

PROOF. First observe that we can take the disjoint k -tuple π' and produce a disjoint k -tuple τ from $\mathbf{p}$ to $\mathbf{q}$ by restricting the path π'_i to the interval $[x_i, y_i]$ (and possibly redefining the value at the right endpoint y_i). The change in length

from doing this is $\|\pi'_i\|_{[y_i, y'_i]} \|f\|$. This is bounded above by the last passage value $f[(y_i, \pi'_i(y_i)) \rightarrow (y'_i, m)]$, which is bounded above by

$$|\pi'_i(y_i) + 1 - m| \omega(f, [y_i, y'_i] \times \llbracket m, \pi'_i(y_i) \rrbracket).$$

Since $\|\tau\|_f \leq f[\mathbf{p} \rightarrow \mathbf{q}]$, this yields the first bound in the lemma.

For the other bound, we can take the k -tuple π and extend the component π_i to a path π_i^* from (x_i, n) to (y'_i, m) by letting $\pi_i = m$ on the interval $[y_i, y'_i]$. This may break the essential disjointness with the path π_{i+1} , so we may need to redefine π_{i+1} on the interval $[y_i, y'_i]$. We can deal with this by defining a new path π_{i+1}^* so that $\pi_{i+1}^* = \max\{m+1, \pi_{i+1}\}$ on the intersection $[y_i, y'_i] \cap [x_{i+1}, y_{i+1})$, and setting $\pi_{i+1}^* = \pi_{i+1}$ elsewhere. Continuing in this way, we can redefine all of the paths $\pi_i, \dots, \pi_k$ to get functions π_{i+j}^* that are equal to $\max\{m+j, \pi_{i+j}\}$ on each of the intervals $[a_{i+j}, b_{i+j}) = [y_i, y'_i] \cap [x_{i+j}, y_{i+j})$, and are equal to π_{i+j} elsewhere.

We check that this process yields a disjoint k -tuple. The functions π_{i+j}^* are cadlag and nonincreasing on the interval $[a_{i+j}, b_{i+j})$ where the path was redefined. Since this interval is closed on the left and open on the right, this ensures that π_{i+j}^* is cadlag everywhere. Now, since $\pi_{i+j}^* \geq \pi_{i+j}$ on the interval $[a_{i+j}, b_{i+j})$, we have that π_{i+j}^* is nonincreasing on $[a_i, y_{i+j}]$. To check that π_{i+j}^* is nonincreasing everywhere it just remains to check the endpoint a_{i+j} , when $a_{i+j} = y_i$. For this, observe that the essential disjointness of $\pi_i, \pi_{i+1}, \dots, \pi_{i+j}$ implies that

$$\pi_i(y_i^-) < \pi_{i+1}(y_i^-) < \dots < \pi_{i+j}(y_i^-)$$

which forces $\pi_{i+j}(y_i^-) \geq m+j$. Since $\pi_{i+j}^*(y_i) = \max\{m+j, \pi_{i+j}(y_i)\}$, this implies that π_{i+j}^* is nonincreasing at y_i . Finally, observe that for any $j < j'$, the new definitions imply $\pi_{i+j}^* \leq \pi_{i+j'}^*$, and the two paths are essentially disjoint on the interval $[a_{i+j}, b_{i+j}) \cap [a'_{i+j}, b'_{i+j})$. Hence $\pi^* = (\pi_1, \dots, \pi_{i-1}, \pi_i^*, \dots, \pi_k^*)$ is a disjoint k -tuple from $\mathbf{p}$ to $\mathbf{q}'$.

Moreover, for each $j \geq 1$ we have

$$\|\pi_{i+j}\|_f - \|\pi_{i+j}^*\|_f \leq 2\omega(f, [y_i, y'_i] \times \llbracket m, m+j \rrbracket) \leq 2\omega(f, [y_i, y'_i] \times \llbracket m, m+k-i \rrbracket).$$

For $j = 0$ we have the same bound, except with the 2 removed since π_i is not defined on $[y_i, y'_i]$. Summing over $j \in \llbracket i, k \rrbracket$ and using that $\|\pi^*\|_f \leq f[\mathbf{p} \rightarrow \mathbf{q}']$ yields the second inequality. $\square$

The second lemma helps controls the weight of an individual path in a disjoint optimizer.

LEMMA 2.7. *For an endpoint pair (p, q) of single points, let (p^k, q^k) be an endpoint pair of size $k \geq 2$, where $p^k = (p, \dots, p)$ and $q^k = (q, \dots, q)$. Let $\pi = (\pi_1, \dots, \pi_k)$ be a disjoint optimizer for this endpoint pair. Then for all $i \in \llbracket 1, k \rrbracket$, we have*

$$\|\pi_i\|_f \geq f[p^k \rightarrow q^k] - f[p^{k-1} \rightarrow q^{k-1}].$$

PROOF. For each i , the collection $(\pi_j : j \neq i, j \in \llbracket 1, k \rrbracket)$ is a disjoint $(k-1)$ -tuple from p^{k-1} to q^{k-1} . Therefore

$$f[p^k \rightarrow q^k] = \|\pi\|_f = \|\pi_i\|_f + \sum_{j \neq i, j \in \llbracket 1, k \rrbracket} \|\pi_j\|_f \leq \|\pi_i\|_f + f[p^{k-1} \rightarrow q^{k-1}].$$

The lemma follows by rearranging the above inequality. $\square$

The next lemma gives naive bounds on the value of $f[\mathbf{p} \rightarrow \mathbf{q}]$ in terms of single-point last passage values and last passage values with clustered endpoints.

LEMMA 2.8. *Let $(\mathbf{p}, \mathbf{q}) = (\mathbf{x}, n; \mathbf{y}, m)$ be an endpoint pair of size $k \geq 2$. Then*

$$\sum_{i=1}^k (f[p_i^k \rightarrow q_i^k] - f[p_i^{k-1} \rightarrow q_i^{k-1}]) \leq f[\mathbf{p} \rightarrow \mathbf{q}] \leq \sum_{i=1}^k f[p_i \rightarrow q_i],$$

where the notation p^k is as in Lemma 2.7.

PROOF. The upper bound follows since any disjoint k -tuple from $\mathbf{p}$ to $\mathbf{q}$ gives rise to k paths from p_i to q_i . For the lower bound, we construct a disjoint k -tuple from $\mathbf{p}$ to $\mathbf{q}$ using a diagonal argument. For each $i \in \llbracket 1, k \rrbracket$, let τ^i be a disjoint optimizer from p_i^k to q_i^k . By the monotonicity established in Lemma 2.3, the components $\tau_1^1, \dots, \tau_k^k$ form k disjoint paths from $\mathbf{p}$ to $\mathbf{q}$. Finally, $\|\tau^i\|_f \geq f[p_i^k \rightarrow q_i^k] - f[p_i^{k-1} \rightarrow q_i^{k-1}]$ by Lemma 2.7. The conclusion follows. $\square$

We finish this section by recording a metric composition law, which can deduced from the definition of last passage values without much difficulty. (This is also recorded as [15, Lemma 4.4].)

LEMMA 2.9. *Let $(\mathbf{p}, \mathbf{q}) = (\mathbf{x}, n; \mathbf{y}, m)$ be an endpoint pair of size k and let $\ell \in \{m+1, \dots, n\}$. Then*

$$f[\mathbf{p} \rightarrow \mathbf{q}] = \max_{\mathbf{z}} f[\mathbf{p} \rightarrow (\mathbf{z}, \ell)] + f[(\mathbf{z}, \ell-1) \rightarrow \mathbf{q}],$$

where the maximum is taken over $\mathbf{z} \in \mathbb{R}_{\leq}^k$ such that both $(\mathbf{p}; \mathbf{z}, \ell)$ and $(\mathbf{z}, \ell-1; \mathbf{q})$ are endpoint pairs.

We note that the arg max of the right-hand side of the display above is precisely the location in $\mathbb{R}_{\leq}^k$ where an optimizer from $\mathbf{p}$ to $\mathbf{q}$ jumps from line ℓ to line $\ell-1$.

2.3. Melons

Let $f \in \mathcal{C}^{\llbracket 1, n \rrbracket}$. For any point $t \in \mathbb{R}$, the **melon of f opened up at t** is a sequence of functions $W_t f = (W_t f_1, \dots, W_t f_n)$ from $[t, \infty)$ to $\mathbb{R}$ defined as follows. Set $W_t f_1(s) = f[(t, n) \rightarrow (s, 1)]$ and for $k \in \llbracket 2, n \rrbracket$ let

$$W_t f_k(s) = f[(t, n)^k \rightarrow (s, 1)^k] - f[(t, n)^{k-1} \rightarrow (s, 1)^{k-1}].$$

The functions $W_t f_i$ satisfy $W_t f_i(t) = 0$ for all i and are ordered: $W_t f_1 \geq \dots \geq W_t f_n$, see the discussion in [15, Section 4] (before Proposition 4.1). Surprisingly, the melon operation preserves last passage values. This fact was essentially shown by Noumi and Yamada [38]. A version for single-point last passage values across continuous functions was proven by Biane, Bougerol, and O'Connell [1]. We quote a multi-point version from [15] which applies to our context.

THEOREM 2.10 ([15, Proposition 4.1]). *Let $f \in \mathcal{C}^{\llbracket 1, n \rrbracket}$, and let $(\mathbf{p}, \mathbf{q}) = (\mathbf{x}, n; \mathbf{y}, 1)$ be any endpoint pair. Then for all $t \leq x_1$, we have*

$$f[\mathbf{p} \rightarrow \mathbf{q}] = W_t f[\mathbf{p} \rightarrow \mathbf{q}].$$

A consequence of Theorem 2.10 is that disjointness of optimizers across the melon $W_t f$ is equivalent to disjointness across the original functions f . Let $(\mathbf{p}, \mathbf{q})$ and $(\mathbf{p}', \mathbf{q}')$ be endpoint pairs such that the concatenation $(\mathbf{p} \cup \mathbf{p}', \mathbf{q} \cup \mathbf{q}')$ remains an endpoint pair. For disjoint k -tuples π, τ from $\mathbf{p}$ to $\mathbf{q}$ and $\mathbf{p}'$ to $\mathbf{q}'$, we say that π and τ are **essentially disjoint** if (π, τ) is a disjoint k -tuple from $\mathbf{p} \cup \mathbf{p}'$ to $\mathbf{q} \cup \mathbf{q}'$.

For the next lemma, let $\pi_f^+[\mathbf{p}, \mathbf{q}]$ denote the rightmost optimizer from $\mathbf{p}$ to $\mathbf{q}$ across a function $f \in \mathcal{C}^{\llbracket 1, n \rrbracket}$, and let $\pi_f^-[\mathbf{p}, \mathbf{q}]$ denote the leftmost optimizer.

LEMMA 2.11. *Let $f \in \mathcal{C}^{\llbracket 1, n \rrbracket}$, and let $(\mathbf{p}, \mathbf{q}) = (\mathbf{x}, n; \mathbf{y}, 1)$, $(\mathbf{p}', \mathbf{q}') = (\mathbf{x}', n; \mathbf{y}', 1)$ be two endpoint pairs, such that the concatenation $(\mathbf{p} \cup \mathbf{p}', \mathbf{q} \cup \mathbf{q}')$ remains an endpoint pair. Fix $t \leq x_1$.*

Then $\pi_f^-[\mathbf{p}, \mathbf{q}]$ and $\pi_f^+[\mathbf{p}', \mathbf{q}']$ are essentially disjoint if and only if $\pi_{W_{tf}}^-[\mathbf{p}', \mathbf{q}']$ and $\pi_{W_{tf}}^+[\mathbf{u}_2, \mathbf{v}_2]$ are essentially disjoint.

Lemma 2.11 is essentially Lemma 4.5 from [15], but for paths with multiple starting and ending points. The proofs are identical up to trivial notational changes.

Optimizers across melons will often be simpler to analyze than optimizers across the original functions. For example, we have the following simple lemma from [15]. In this lemma, the function f takes the form of a melon opened up at 0.

LEMMA 2.12 ([15, Lemma 5.1]). *Let $f \in \mathcal{C}^{\llbracket 1, n \rrbracket}$ be such that $f_i(0) = 0$ for all $i \in \llbracket 1, n \rrbracket$ and $f_i \geq f_{i+1}$ for all $i \in \llbracket 1, n-1 \rrbracket$. Fix $j \leq k \leq n \in \mathbb{N}$. Let $(\mathbf{p}, \mathbf{q}) = (\mathbf{x}, n; \mathbf{y}, 1)$ be an endpoint pair of size k with $x_i = 0$ for all $i \in \llbracket 1, j \rrbracket$. Then there exists an optimizer π from $\mathbf{p}$ to $\mathbf{q}$ such that $\pi_i(t) = i$ for all $t \in (0, y_1)$, $i \in \llbracket 1, j \rrbracket$.*

In particular, Lemma 2.12 gives that the leftmost optimizer from $(t, n)^k$ to any $(\mathbf{y}, 1)$ (with $\mathbf{y} \in \mathbb{R}_{\leq}^k$) in any melon W_{tf} will only use the top k lines $W_{tf_1}, \dots, W_{tf_k}$.

2.4. Brownian melons and the parabolic Airy line ensemble

Melons have a remarkable probabilistic structure when the input function consists of n independent two-sided Brownian motions $B^n = (B_1^n, \dots, B_n^n)$. In this case, the **Brownian n -melon** $W^n := W_0 B^n$ is given by n Brownian motions started at 0, conditioned to never intersect. This was first shown in [41, Theorem 7]. This structure allows one to find the scaling limit of W^n at the edge. See Figure 2.1 for an illustration.

First tilt and rescale the melons $W^n = (W_1^n, \dots, W_n^n)$. Define $\mathcal{B}^n = (\mathcal{B}_1^n, \dots, \mathcal{B}_n^n)$ by

$$(2.1) \quad \mathcal{B}_i^n(y) = n^{1/6} (W_i^n(1 + 2yn^{-1/3}) - 2\sqrt{n} - 2yn^{1/6}).$$

Then the functions $\mathcal{B}^n$ converges in distribution to a continuous limit known as the parabolic Airy line ensemble.

THEOREM 2.13 ([8, Theorem 3.1]). *The sequence $\mathcal{B}^n$ converges in distribution to a continuous limit $\mathcal{B} : \mathbb{R} \times \mathbb{N} \rightarrow \mathbb{R}$, in the topology of uniform convergence on compact subsets of $\mathbb{R} \times \mathbb{N}$. The limit $\mathcal{B}$ is the **parabolic Airy line ensemble**.*

The qualifier parabolic comes from the fact that the process $\mathcal{A}(x) = \mathcal{B}(x) + x^2$ is stationary, so $\mathcal{B}$ has a parabolic shape. The process $\mathcal{A}$ is known as the **(stationary) Airy line ensemble**. Note that Corwin and Hammond technically worked with nonintersecting Brownian bridges (with diffusion parameter 1) from time 0 to time 2, rather than nonintersecting Brownian motions $\mathcal{B}^n$ (with diffusion parameter 1). The two objects are equivalent in the Airy line ensemble scaling limit by virtue of the standard transformation between Brownian bridge and Brownian motion.

Both Brownian melons and the parabolic Airy line ensemble are strictly ordered and satisfy a useful resampling property called the **Brownian Gibbs property**. This makes these objects useful in practice. The next theorem gathers results from [8], from Definition 2.13 and Theorem 1. We choose not to introduce the Brownian Gibbs property as formally as in that paper, since it only plays a tangential role in this paper.




FIGURE 2.1. A simulation of a Brownian melon (from [14]) and a window in which it converges to the parabolic Airy line ensemble.

THEOREM 2.14. Let W^n denote a Brownian n -melon, let $\mathcal{B}$ denote the parabolic Airy line ensemble and let $\tilde{\mathcal{B}} = 2^{-1/2}\mathcal{B}$. Almost surely,

$$(2.2) \quad \begin{aligned} W_i^n(t) &> W_{i+1}^n(t) \quad \text{for all } i \in \llbracket 1, n \rrbracket, t > 0, \quad \text{and} \\ \tilde{\mathcal{B}}_i(t) &> \tilde{\mathcal{B}}_{i+1}(t) \quad \text{for all } i \in \mathbb{N}, t \in \mathbb{R}. \end{aligned}$$

Moreover, for any box $S = \llbracket \ell, k \rrbracket \times [a, b]$ with $a > 0$ and $k \leq n$, the process $W^n|_S$ given $W^n|_S^c$ is just given by $k - \ell + 1$ Brownian bridges (with diffusion parameter 1) connecting up the points $W_i^n(a)$ and $W_i^n(b)$, conditioned so that the nonintersection conditions in (2.2) hold. This property is called the Brownian Gibbs property.

Similarly, for any box $S = \llbracket \ell, k \rrbracket \times [a, b]$, the process $\tilde{\mathcal{B}}|_S$ given $\tilde{\mathcal{B}}|_S^c$ is just given by $k - \ell + 1$ Brownian bridges (with diffusion parameter 1) connecting up the points $\tilde{\mathcal{B}}_i(a)$ and $\tilde{\mathcal{B}}_i(b)$, conditioned so that the nonintersection conditions in (2.2) hold.

We end this section by recording a few uniqueness results for Brownian last passage percolation. These results are stated for last passage percolation between multiple points on potentially different lines.

LEMMA 2.15. Let $(\mathbf{p}, \mathbf{q}) = (p_1, \dots, p_k, q_1, \dots, q_k) \in ([x, y] \times \llbracket m, n \rrbracket)^{2k}$ be an endpoint pair. Let $I \subset \mathbb{Z}$ be an integer interval containing $\llbracket m, n \rrbracket$, and let $B = \{B_i : i \in I\}$ be a sequence of random continuous functions with the following property. For $[a, b] \subset (x, y)$ and any $i \in I$, let $\mathcal{F}_{[a, b] \times \{i\}}$ be the σ -algebra generated by all increments $B_j(t) - B_j(s)$ with

$$[t, s] \times \{j\} \subset [x, y] \times \llbracket m, n \rrbracket \setminus ((a, b) \times \{i\}).$$

Suppose that for any $[a, b] \subset (x, y)$ and $i \in I$, the conditional distribution

$$(2.3) \quad \mathbb{P}(B_i(b) - B_i(a) \in \cdot \mid \mathcal{F}_{[a, b] \times \{i\}})$$

is a continuous distribution almost surely. Then there is almost surely a unique optimizer π from $\mathbf{p}$ to $\mathbf{q}$.

This lemma is due to Hammond, see [26, Lemma B.1]. However, since we have stated it in greater generality than in that paper, we include a brief proof using Hammond's method.

PROOF OF LEMMA 2.15. For any $\gamma = [a, b] \times \{i\}$ with $[a, b] \subset (x, y)$, $i \in I$, and $j \in \llbracket 1, k \rrbracket$, let

$$B_{\gamma,j}[\mathbf{p} \rightarrow \mathbf{q}] = \sup \|\pi\|_B,$$

where the supremum is taken over all disjoint k -tuples from $\mathbf{p}$ to $\mathbf{q}$ subject to the constraint that $\gamma \subset \Gamma(\pi_j)$. Define $B_{\gamma^c}[\mathbf{p} \rightarrow \mathbf{q}]$ similarly, but with the supremum taken over all disjoint k -tuples from $\mathbf{p}$ to $\mathbf{q}$ subject to the constraint that $\gamma \cap \Gamma(\pi_i) = \emptyset$ for all $i \in \llbracket 1, k \rrbracket$. We claim that almost surely,

$$(2.4) \quad B_{\gamma,j}[\mathbf{p} \rightarrow \mathbf{q}] \neq B_{\gamma^c}[\mathbf{p} \rightarrow \mathbf{q}]$$

for all j . Indeed, $B_{\gamma^c}[\mathbf{p} \rightarrow \mathbf{q}]$ is $\mathcal{F}_\gamma$ -measurable, and $B_{\gamma,j}[\mathbf{p} \rightarrow \mathbf{q}] = X + B_i(b) - B_i(a)$, where X is an $\mathcal{F}_\gamma$ -measurable random variable. Since $B_i(b) - B_i(a)$ has a continuous distribution, conditionally on $\mathcal{F}_\gamma$, this yields (2.4). Now, (2.4) holds simultaneously almost surely for all γ with rational endpoints and $j \in \llbracket 1, k \rrbracket$. On the other hand, if there were two optimizers π, π' from $\mathbf{p}$ to $\mathbf{q}$, then there would exist a $j \in \llbracket 1, k \rrbracket$ and a γ with rational endpoints such that $\gamma \cap \Gamma(\pi_i) = \emptyset$ for all $i \in \llbracket 1, k \rrbracket$ but $\gamma \subset \Gamma(\pi'_j)$. Therefore

$$B_{\gamma,j}[\mathbf{p} \rightarrow \mathbf{q}] = \|\pi'\|_B = \|\pi\|_B = B_{\gamma^c}[\mathbf{p} \rightarrow \mathbf{q}],$$

contradicting (2.4). $\square$

The conditions of the lemma are set up so that they apply to all the objects that we work with.

LEMMA 2.16. *The conditions of Lemma 2.15 are satisfied when B is a collection of independent Brownian motions for any I and $(\mathbf{p}, \mathbf{q})$, when $B = W^n$ is a Brownian melon with $x \geq 0$ and $I \subset \llbracket 1, n \rrbracket$, and when $B = \mathcal{B}$ is the parabolic Airy line ensemble and $I \subset \mathbb{N}$.*

PROOF. If B is a collection of independent Brownian motions, then (2.3) is a normal distribution almost surely, and hence is continuous. We treat the remaining two cases together by appealing to the Brownian Gibbs property in Theorem 2.14 for either B or $2^{-1/2}B$ (where either $B = W^n$ or $2^{1/2}B = \mathcal{B}$). By possibly increasing the size of I , we may assume $I = \llbracket 1, m \rrbracket$ for some m . Let $[a, b] \subset (x, y)$.

By the Brownian Gibbs property, conditionally on the σ -algebra $\mathcal{G}$ generated by $B_j(t)$ for all $(j, t) \notin \{i\} \times [a, y+1]$, the process $B_i(t) - B_i(a)$, $t \in [a, y+1]$ is a Brownian bridge connecting 0 and $B_i(y+1) - B_i(a)$, conditioned so that the ensemble B remains nonintersecting. In particular, conditionally on the σ -algebra $\mathcal{G}'$ generated by $\mathcal{G}$ and $B_i(y) - B_i(r)$, $r \in [b, y]$, almost surely the distribution of

$$B_i(b) - B_i(a)$$

is absolutely continuous with respect to Lebesgue measure on $\mathbb{R}$. Finally, $\mathcal{F}_{[a,b] \times \{i\}} \subset \mathcal{G}'$, giving the result. $\square$

Lemmas 2.15 and 2.16 together allow us to speak of a single optimizer or geodesic when considering last passage problems across these Brownian motions, Brownian melons, and the parabolic Airy line ensemble.

2.5. Melon geodesics and the Airy sheet

Recall the definition of the prelimiting extended Airy sheets (2.5)

$$\mathcal{S}^n(\mathbf{x}, \mathbf{y}) = n^{1/6} \left(B^n[(2n^{-1/3}\mathbf{x}, n) \rightarrow (1 + 2n^{-1/3}\mathbf{y}, 1)] - 2k\sqrt{n} - n^{1/6} \sum_{i=1}^k 2(y_i - x_i) \right),$$

from Theorem 1.2, where B^n is a collection of n independent two-sided standard Brownian motions. For thinking about the prelimiting sheets $\mathcal{S}^n$, it will be helpful to use an alternate formula for $\mathcal{S}^n$ in terms of the prelimiting Airy line ensembles $\mathcal{B}^n$ (defined in (2.1)). Recall the Brownian n -melon $W^n = W_0 B^n$. When $\mathbf{x} \in \mathbb{R}_{\leq}^k \cap [0, \infty)^k$, by Theorem 2.10 we have

$$\mathcal{S}^n(\mathbf{x}, \mathbf{y}) = n^{1/6} \left(W^n[(2n^{-1/3}\mathbf{x}, n) \rightarrow (1 + 2n^{-1/3}\mathbf{y}, 1)] - 2k\sqrt{n} - n^{1/6} \sum_{i=1}^k 2(y_i - x_i) \right).$$

Note that $\mathcal{B}^n$ is just an affine transformation of W^n , given by (2.1). Using the fact that last passage values commute with affine shifts, we have

$$(2.6) \quad \mathcal{S}^n(\mathbf{x}, \mathbf{y}) = \mathcal{B}^n[(\mathbf{x} - n^{1/3}/2, n) \rightarrow (\mathbf{y}, 1)] - kn^{2/3}.$$

One of the main insights of [15] was finding a way to take a limit of the right-hand side of (2.6) for single points x, y in order to define the Airy sheet $\mathcal{S} : \mathbb{R} \rightarrow \mathbb{R}$ in terms of the Airy line ensemble. The basic idea there was to carefully analyze the location and coalescence structure of geodesics across $\mathcal{B}^n$. The culmination of this analysis showed that if $\pi^n\{x, y\}, \pi^n\{x, z\}$ are the rightmost geodesics across $\mathcal{B}^n$ from $(x - n^{1/3}/2, n)$ to $(y, 1)$ and $(z, 1)$, then with high probability $\pi^n\{x, y\}$ and $\pi^n\{x, z\}$ coincide outside of an $O(1)$ region around the points $(y, 1), (z, 1)$. This suggests that difference

$$\mathcal{S}^n(x, y) - \mathcal{S}^n(x, z)$$

should converge to a difference of last passage problems in the Airy line ensemble. This, along with an estimate on the location of the paths $\pi^n\{x, z\}$, motivates the following definition of the Airy sheet on $[0, \infty) \times \mathbb{R}$.

DEFINITION 2.17. For a parabolic Airy line ensemble $\mathcal{B}$, we define the **half Airy sheet** of $\mathcal{B}$ to be the function $\mathcal{S}_{\mathcal{B}} : [0, \infty) \times \mathbb{R} \rightarrow \mathbb{R}$ specified by the formulas

- $\mathcal{S}_{\mathcal{B}}(0, y) = \mathcal{B}_1(y)$ for $y \in \mathbb{R}$.
- For $x > 0$ and $y, z \in \mathbb{R}$, we have

$$(2.7) \quad \begin{aligned} & \mathcal{S}_{\mathcal{B}}(x, y) - \mathcal{S}_{\mathcal{B}}(x, z) \\ &= \lim_{m \rightarrow \infty} \mathcal{B}[(-\sqrt{m/(2x)}, m) \rightarrow (y, 1)] - \mathcal{B}[(-\sqrt{m/(2x)}, m) \rightarrow (z, 1)]. \end{aligned}$$

- For any $x \in \mathbb{Q} \cap (0, \infty)$ and $y \in \mathbb{R}$, we have

$$(2.8) \quad \mathcal{S}_{\mathcal{B}}(x, y) = \lim_{a \rightarrow \infty} \frac{1}{a} \int_{-a}^0 (\mathcal{S}_{\mathcal{B}}(x, y) - \mathcal{S}_{\mathcal{B}}(x, z) - (x - z)^2 + \xi) dz,$$

where ξ is the expectation of the GUE Tracy-Widom distribution. Note that we could have integrated on the right-hand side of (2.8) over any interval of length a containing 0.

It turns out that almost surely, all the limits above exist, and the resulting function $\mathcal{S}_{\mathcal{B}}$ is continuous. The existence of such an object follows from [15, Theorem 8.3]. The first bullet is part of [15, Definition 8.1(ii)], the second bullet is

[15, Remark 8.1], and the third bullet is given by the second display in the proof of [15, Proposition 8.2].

The half-Airy sheet can be extended to all of $\mathbb{R}^2$ by a stationarity relationship, see [15, Definition 8.1 and Theorem 8.3].

DEFINITION 2.18. The **Airy sheet** is the unique (in law) random continuous function $\mathcal{S} : \mathbb{R}^2 \rightarrow \mathbb{R}$ satisfying

- $\mathcal{S}(\cdot, \cdot) \stackrel{d}{=} \mathcal{S}(t + \cdot, t + \cdot)$ for all $t \in \mathbb{R}$
- $\mathcal{S}|_{[0, \infty) \times \mathbb{R}}$ is a half Airy sheet.

Having defined the Airy sheet, we can now state the main convergence result from [15]. When stating this result, we also record convergence information for rightmost geodesics across $\mathcal{B}^n$ (parts (ii), (iii) below) which is the crucial input in defining the Airy sheet.

THEOREM 2.19. *For any subsequence $Y \subset \mathbb{N}$, there exists a further subsequence $Y' \subset Y$ and a coupling of $\mathcal{B}$, and $\{\mathcal{B}^n : n \in Y'\}$ such that the following statements all hold almost surely:*

- (i) *The pair $(\mathcal{B}^n, \mathcal{S}^n|_{[0, \infty) \times \mathbb{R}})$ converges in the uniform-on-compact topology to $(\mathcal{B}, \mathcal{S}_{\mathcal{B}})$. Here $\mathcal{B}$ is a parabolic Airy line ensemble, and $\mathcal{S}_{\mathcal{B}}$ is the half-Airy sheet of $\mathcal{B}$.*
- (ii) *Let $Z_m^n(x, y)$ denote the jump time from line $m+1$ to m for the rightmost geodesic $\pi^n\{x, y\}$ from $(x - n^{1/3}/2, n)$ to $(y, 1)$. For all $x \in \mathbb{Q} \cap (0, \infty)$, $y \in \mathbb{Q}$, $m \in \mathbb{N}$, the random variables $Z_m^n(x, y)$ converge almost surely to limits $Z_m(x, y)$. Moreover,*

$$\lim_{m \rightarrow \infty} \frac{Z_m(x, y)}{\sqrt{m}} = \frac{-1}{\sqrt{2x}}.$$

- (iii) *For every $x \in \mathbb{Q} \cap (0, \infty)$ and $y < z \in \mathbb{Q}$, there are points $X_1 < x < X_2$ with $X_1, X_2 \in \mathbb{Q} \cap (0, \infty)$ and $T < \min(y, z)$ such that for all large enough n , we have*

$$\Gamma(\pi^n\{X_1, y\}|_{[T, y]}) \cap \Gamma(\pi^n\{X_2, z\}|_{[T, z]}) \neq \emptyset.$$

Here recall that $\Gamma(\pi)$ denotes the zigzag graph of π .

This coupling is constructed in [15, Section 8] (after Lemma 8.4). The construction there shows that condition (ii) above is satisfied. Property (i) of the coupling is shown as [15, Lemma 8.5], and property (iii) of the coupling is shown in the proof of [15, Lemma 8.5]. Note that the notion of a point ‘lying along the path π ’ used in that proof means that a point is contained in the zigzag graph of π . We remark that while the rationals $\mathbb{Q}$ are used in the Theorem 2.19, they play no special role. The theorem would still hold with any other countable dense set D in place of $\mathbb{Q}$.

To prove convergence of $\mathcal{S}^n$ jointly over all k and $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$, we will similarly focus on understanding optimizers across $\mathcal{B}^n$. We use Theorem 2.19 of [15] as a starting point for our analysis. In the remainder of this section, we record a few auxiliary results from [15] that will also be needed in our analysis, along with some simple consequences of that paper. We start with two technical lemmas that are stepping stones along the path to Theorem 2.19.

For the first lemma, for a random array $\{R_{n,m} : n, m \in \mathbb{N}\}$, we write

$$(2.9) \quad R_{n,m} = o(r_m) \quad \text{if for all } \epsilon > 0 \quad \sum_{m=1}^{\infty} \limsup_{n \rightarrow \infty} \mathbb{P}(|R_{n,m}/r_m| > \epsilon) < \infty.$$

LEMMA 2.20 ([15, Lemma 7.1]). *Let K be a compact subset of $(0, \infty) \times \mathbb{R}$. Then we have*

$$\sup_{(x,y) \in K} \left| Z_m^n(x, y) + \sqrt{\frac{m}{2x}} \right| = o(\sqrt{m})$$

and $Z_m^n(x, y)$ is tight as a function of n for each fixed $m \in \mathbb{N}$, $(x, y) \in (0, \infty) \times \mathbb{R}$.

We also require a useful lemma about disjointness of geodesics.

LEMMA 2.21 ([15, Lemma 7.2]). *Fix $x > 0$ and $y_1 < y_2$. Then*

$$\lim_{\epsilon \rightarrow 0^+} \limsup_{n \rightarrow \infty} \mathbb{P}(\pi^n\{x - \epsilon, y_1\} \quad \text{and} \quad \pi^n\{x + \epsilon, y_2\} \text{ are essentially disjoint}) = 0.$$

Note that these two lemmas are quite non-trivial. In particular, Lemma 2.20 requires the full power of a difficult structural theorem for the Airy line ensemble from [17].

We also record a corollary of Lemma 2.21 that follows from symmetries of Brownian LPP.

COROLLARY 2.22. *Fix $y_1 < y_2$. Then*

$$\lim_{\epsilon \rightarrow 0^+} \limsup_{n \rightarrow \infty} \mathbb{P}(\pi^n\{0, y_1\} \quad \text{and} \quad \pi^n\{\epsilon, y_2\} \text{ are essentially disjoint}) = 0.$$

PROOF. We write $\pi^n[x, y]$ for the rightmost geodesic across the original Brownian motions from $(2n^{-1/3}x, n)$ to $(1 + 2n^{-1/3}y, 1)$. By Lemma 2.21 and Lemma 2.11, for any $x > 0$ we have

$$\lim_{\epsilon \rightarrow 0^+} \limsup_{n \rightarrow \infty} \mathbb{P}(\pi^n[x - \epsilon, y_1] \quad \text{and} \quad \pi^n[x + \epsilon, y_2] \text{ are essentially disjoint}) = 0.$$

for any $x > 0$. Translation invariance of Brownian increments implies that the above statement holds for any $x \in \mathbb{R}$, not just $x > 0$, and monotonicity of geodesics (Proposition 2.1) implies that the statement holds with $x - \epsilon$ replaced by x . Setting $x = 0$ and translating back to the melon environment via Lemma 2.11 yields the result. $\square$

Next, we record a few basic facts and symmetries about the Airy sheet.

LEMMA 2.23 (see [15, Lemma 9.1 and Remark 1.1.6]). *The process $(x, y) \mapsto \mathcal{S}(x, y) + (x - y)^2$ is translation invariant in both x and y . Also, $\mathcal{S}(x, y) \stackrel{d}{=} \mathcal{S}(-y, -x)$. Here the distributional equality is joint in all $x, y \in \mathbb{R}$.*

Moreover, $\mathcal{S}(0, 0)$ has GUE Tracy-Widom distribution, and hence satisfies the tail bound

$$\mathbb{P}(|\mathcal{S}(0, 0)| > m) \leq ce^{-dm^{3/2}}$$

for universal constants $c, d > 0$ and all $m > 0$.

Note that the tail bound in Lemma 2.23 on the GUE Tracy-Widom distribution goes back to [49].

We end this section by giving a more flexible description of the difference $\mathcal{S}_{\mathcal{B}}(x, y) - \mathcal{S}_{\mathcal{B}}(x, z)$ defined in (2.7).

LEMMA 2.24. *Almost surely the following is true. Take any $x \geq 0$ and $z_1 < z_2$. Let $\pi : (-\infty, z_1] \rightarrow \mathbb{N}$ be a nonincreasing cadlag function, such that $\lim_{y \rightarrow -\infty} \frac{\pi(y)}{2y^2} = x$. Then*

$$(2.10) \quad \mathcal{S}_{\mathcal{B}}(x, z_1) - \mathcal{S}_{\mathcal{B}}(x, z_2) = \lim_{y \rightarrow -\infty} \mathcal{B}[(y, \pi(y)) \rightarrow (z_1, 1)] - \mathcal{B}[(y, \pi(y)) \rightarrow (z_2, 1)].$$

PROOF. By applying Lemma 2.4 with the endpoint pairs $((y, \pi(y)), (z_2, 1))$ and $((-\sqrt{\pi(y)/(2x + \delta)}, \pi(y)), (z_1, 1))$, and sending $y \rightarrow -\infty$, we have that for any $\delta > 0$,

$$\begin{aligned} & \liminf_{y \rightarrow -\infty} \mathcal{B}[(y, \pi(y)) \rightarrow (z_1, 1)] - \mathcal{B}[(y, \pi(y)) \rightarrow (z_2, 1)] \\ & \geq \lim_{m \rightarrow \infty} \mathcal{B}[(-\sqrt{m/(2x + \delta)}, m) \rightarrow (z_1, 1)] - \mathcal{B}[(-\sqrt{m/(2x + \delta)}, m) \rightarrow (z_2, 1)] \\ & = \mathcal{S}_{\mathcal{B}}(x + \delta/2, z_1) - \mathcal{S}_{\mathcal{B}}(x + \delta/2, z_2). \end{aligned}$$

Therefore by continuity of $\mathcal{S}_{\mathcal{B}}$, the right-hand side of (2.10) is bounded below by the left-hand side. For $x > 0$, the opposite inequality holds by symmetric reasoning. For $x = 0$, the opposite inequality holds since

$$\mathcal{B}[p \rightarrow (z_1, 1)] - \mathcal{B}[p \rightarrow (z_2, 1)] \leq \mathcal{B}_1(z_1) - \mathcal{B}_1(z_2) = \mathcal{S}_{\mathcal{B}}(0, z_1) - \mathcal{S}_{\mathcal{B}}(0, z_2).$$

for any point $p \in (-\infty, z_1] \times \mathbb{N}$. Indeed, any path π from p to $(z_1, 1)$ can always be extended to a path from p to $(z_2, 1)$ by extending π to be equal to 1 on the interval $[z_1, z_2]$. This picks up the increment $\mathcal{B}_1(z_2) - \mathcal{B}_1(z_1)$. $\square$

CHAPTER 3

Tightness

3.1. Tightness of prelimiting sheets

Recall from Theorem 1.2 the space $\mathfrak{X} = \bigcup_{k=1}^{\infty} \mathbb{R}_{\leq}^k \times \mathbb{R}_{\leq}^k$. Topologically, $\mathfrak{X}$ is a disjoint union of certain subsets of $\mathbb{R}^{2k}$. Let $\mathcal{C}(\mathfrak{X}, \mathbb{R})$ be the space of functions from $\mathfrak{X}$ to $\mathbb{R}$ with the uniform-on-compact topology. The main goal of this section is to prove the following theorem.

THEOREM 3.1. *The functions $\mathcal{S}^n$ are tight in $\mathcal{C}(\mathfrak{X}, \mathbb{R})$.*

Note that $\mathcal{S}^n$ (from formula (2.5)) is not defined on all of $\mathfrak{X}$. To formally define $\mathcal{S}^n$ as a random element of $\mathcal{C}(\mathfrak{X}, \mathbb{R})$, we arbitrarily extend $\mathcal{S}$ to all of $\mathfrak{X}$ in a continuous way so that $\mathcal{S}^n \in \mathcal{C}(\mathfrak{X}, \mathbb{R})$. For any compact set $K \subset \mathfrak{X}$, $\mathcal{S}^n|_K$ is well-defined by (2.5) for all large enough n , so the arbitrary choice of extension does not affect any convergence or tightness statements.

Theorem 3.1 will follow from the deterministic bounds and inequalities in Chapter 2.2, and explicit tightness bounds for the prelimiting Airy line ensemble, which we quote from [17]. For this proposition, W^n is a Brownian melon.

PROPOSITION 3.2 ([17, Proposition 4.1]). *Fix $k \in \mathbb{N}$ and $c > 0$. There exist constants $c_k, d_k > 0$ such that for every $n \in \mathbb{N}$, $t > 0$, $s \in (0, ctn^{-1/3}]$, and $a > 0$ we have*

$$\mathbb{P}\left(\left|W_k^n(t) - W_k^n(t+s) + \frac{s\sqrt{n}}{\sqrt{t}}\right| > a\sqrt{s}\right) \leq c_k e^{-d_k a^{3/2}}.$$

We can translate Proposition 3.2 into a modulus of continuity on the prelimiting parabolic Airy line ensembles $\mathcal{B}^n$. To do this, we will employ a general lemma for establishing a modulus of continuity, also developed in [17]. This lemma will also be used later on when establishing a general modulus of continuity for the extended directed landscape.

LEMMA 3.3 ([17, Lemma 3.3]). *Let $T = I_1 \times \dots \times I_k$ be a product of bounded real intervals of length $b_1, \dots, b_k$. Let $c, d > 0$. Let $\mathcal{H}$ be a random continuous function from T taking values in a vector space V with norm $|\cdot|$. Assume that for every $i \in \llbracket 1, k \rrbracket$, that there exist $\alpha_i \in (0, 1)$, $\beta_i, r_i > 0$ such that*

$$(3.1) \quad \mathbb{P}(|\mathcal{H}(t + e_i u) - \mathcal{H}(t)| \geq au^{\alpha_i}) \leq ce^{-da^{\beta_i}}$$

for every coordinate vector e_i , every $a > 0$, and every $t, t + ue_i \in T$ with $u < r_i$. Set $\beta = \min_i \beta_i$, $\alpha = \max_i \alpha_i$, and $r = \max_i r_i$. Then with probability one we have

$$(3.2) \quad |\mathcal{H}(t+s) - \mathcal{H}(t)| \leq C \left(\sum_{i=1}^k |s_i|^{\alpha_i} \log^{1/\beta_i} \left(\frac{2r^{\alpha/\alpha_i}}{|s_i|} \right) \right),$$

for every $t, t+s \in T$ with $|s_i| \leq r_i$ for all i (here $s = (s_1, \dots, s_k)$). Here C is random constant satisfying

$$\mathbb{P}(C > a) \leq \left[\prod_{i=1}^k \frac{b_i}{r_i} \right] c c_0 e^{-c_1 a^\beta},$$

where c_0 and c_1 are constants that depend on $\alpha_1, \dots, \alpha_k, \beta_1, \dots, \beta_k, k$ and d . Notably, they do not depend on $b_1, \dots, b_k, c$ or $r_1, \dots, r_k$.

COROLLARY 3.4. Fix $k \in \mathbb{N}$ and $c > 0$. There exist positive constants c_k, d_k such that for every $n \in \mathbb{N}$, $t > 0$, $s \in [0, ctn^{-1/3}]$, and $a > 0$ we have

$$(3.3) \quad \mathbb{P}\left(\max_{t \leq x < y \leq t+s} \left| W_k^n(x) - W_k^n(y) + \frac{(y-x)\sqrt{n}}{\sqrt{t}} \right| > a\sqrt{s}\right) \leq c_k e^{-d_k a^{3/2}}.$$

PROOF. We apply Lemma 3.3 with $k = 1$, $\alpha_1 = 1/2$, $\beta_1 = 3/2$ and the function $W_k^n(x) - x\sqrt{n/t}$. The assumption (3.1) is implied by Proposition 3.2. Then we have

$$\left| W_k^n(x) - W_k^n(y) + \frac{(y-x)\sqrt{n}}{\sqrt{t}} \right| \leq C\sqrt{y-x} \log^{2/3}\left(\frac{2s}{y-x}\right)$$

for all $1 \leq x < y \leq 1+s$, where C is a random constant satisfying the tail bound on the right-hand side of (3.3) for some constants c_k, d_k . The right-hand side above is bounded above by $C\sqrt{s}$ for all $t \leq x < y \leq t+s$, yielding (3.3). $\square$

We are now in a position to prove a two-point tail bound for $\mathcal{S}^n$. We first define the stationary version $\mathcal{R}^n : \mathfrak{X} \rightarrow \mathbb{R}$ by

$$\mathcal{R}^n(\mathbf{x}, \mathbf{y}) = \mathcal{S}^n(\mathbf{x}, \mathbf{y}) + \sum_{i=1}^k (x_i - y_i)^2,$$

where $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$.

LEMMA 3.5. Take any $k, n \in \mathbb{N}$, $\mathbf{u} = (\mathbf{x}, \mathbf{y})$, $\mathbf{u}' = (\mathbf{x}', \mathbf{y}') \in \mathbb{R}_{\leq}^k \times \mathbb{R}_{\leq}^k$ with $\|\mathbf{u} - \mathbf{u}'\|_2 < 1$, $\|\mathbf{x}\|_2, \|\mathbf{y}\|_2, \|\mathbf{x}'\|_2, \|\mathbf{y}'\|_2 < n^{1/6}$ and $a > 0$. Then

$$\mathbb{P}(|\mathcal{R}^n(\mathbf{x}', \mathbf{y}') - \mathcal{R}^n(\mathbf{x}, \mathbf{y})| > a\sqrt{\|\mathbf{u} - \mathbf{u}'\|_2}) < ce^{-da^{3/2}},$$

for some constants $c, d > 0$ depending only on k .

In this proof and throughout the paper, for $\mathbf{x} \in \mathbb{R}_{\leq}^k$ we write $-\mathbf{x}$ for the unique element of $\mathbb{R}_{\leq}^k$ given by rearranging the coordinates of $-\mathbf{x}$. The basic idea of the proof of Lemma 3.5 is to reduce bounds on differences in $\mathcal{R}$ to bounds on differences of lines in W^n by using deterministic last passage inequalities from Chapter 2.2.

PROOF OF LEMMA 3.5. We first consider fixed n . Recall (from Theorem 1.2) that $\mathcal{S}^n$ is defined using multi-point last passage values across a collection of independent two-sided standard Brownian motions $B = \{B_i : i \in \mathbb{Z}\}$. By using Lemma 2.6 repeatedly and changing one coordinate at a time, we have

$$|\mathcal{S}^n(\mathbf{x}', \mathbf{y}') - \mathcal{S}^n(\mathbf{x}, \mathbf{y})| < 2k(2k+n)n^{1/6} \sup_{i,a,b} |B_i(a) - B_i(b)| + 2n^{1/3}(\|\mathbf{x} - \mathbf{y}\|_1 + \|\mathbf{x}' - \mathbf{y}'\|_1),$$

where the supremum is over all $1 \leq i \leq n$, and $a, b \in \mathbb{R}$, such that both a, b are in either $[2n^{-1/3}(x_i \wedge x'_i), 2n^{-1/3}(x_i \vee x'_i)]$ or $[1 + 2n^{-1/3}(y_i \wedge y'_i), 1 + 2n^{-1/3}(y_i \vee y'_i)]$. Therefore by using standard tail bounds on Brownian motion increments, for any fixed n the bound in the lemma holds by taking c large and d small (depending on n and k).

For the remainder of the proof it suffices to consider n sufficiently large (depending on k). By the triangle inequality, and by the symmetry $\mathcal{R}^n(\mathbf{x}, \mathbf{y}) \stackrel{d}{=} \mathcal{R}^n(-\mathbf{y}, -\mathbf{x})$, it suffices to prove the bound when $\mathbf{x} = \mathbf{x}'$ and $\mathbf{y}, \mathbf{y}'$ agree at all points except for a single coordinate $y_\ell < y'_\ell$. Moreover, if we let T_c be the map translating all coordinates in a vector by c , then $\mathcal{R}^n(T_c \mathbf{x}, T_c \mathbf{y}) \stackrel{d}{=} \mathcal{R}^n(\mathbf{x}, \mathbf{y})$ for all c , so we may assume $x_\ell = 0$ if we relax the norm bounds to $\|\mathbf{x}\|_2, \|\mathbf{y}\|_2, \|\mathbf{x}'\|_2, \|\mathbf{y}'\|_2 < 2n^{1/6}$. With these simplifications, the inequality is equivalent to

$$\mathbb{P}(|\mathcal{R}^n(\mathbf{x}, \mathbf{y}') - \mathcal{R}^n(\mathbf{x}, \mathbf{y})| > a(y'_\ell - y_\ell)^{1/2}) < ce^{-da^{3/2}}.$$

Now, by the representation (2.6) for $\mathcal{S}^n$, we can write

$$(3.4) \quad \mathcal{R}^n(\mathbf{x}, \mathbf{y}') - \mathcal{R}^n(\mathbf{x}, \mathbf{y}) = \mathcal{A}^n[(\mathbf{x} - n^{1/3}/2, n) \rightarrow (\mathbf{y}', 1)] - \mathcal{A}^n[(\mathbf{x} - n^{1/3}/2, n) \rightarrow (\mathbf{y}, 1)],$$

where $\mathcal{A}_i^n(x) = \mathcal{B}_i^n(x) + x^2$. As in Lemma 2.5, let $\mathbf{x}^L$ denote the first ℓ coordinates of $\mathbf{x}$ and let $\mathbf{y}^R$ denote the last $k - \ell$ coordinates of $\mathbf{x}$. By two applications of that lemma, and (3.4), we have

$$\mathcal{R}^n(\mathbf{x}^R, \mathbf{y}'^R) - \mathcal{R}^n(\mathbf{x}^R, \mathbf{y}^R) \leq \mathcal{R}^n(\mathbf{x}, \mathbf{y}') - \mathcal{R}^n(\mathbf{x}, \mathbf{y}) \leq \mathcal{R}^n(\mathbf{x}^L, \mathbf{y}'^L) - \mathcal{R}^n(\mathbf{x}^L, \mathbf{y}^L).$$

Now let $0^i \in \mathbb{R}^i$ denote the vector whose coordinates are all $x_\ell = 0$. We can bound the left- and right-hand sides above using Lemma 2.4 applied to the points $0^{k-\ell+1} \leq \mathbf{x}^R, \mathbf{y}^R \leq \mathbf{y}'^R$ and $\mathbf{x}^L \leq 0^\ell, \mathbf{y}^L \leq \mathbf{y}'^L$ to get that

$$\mathcal{R}^n(0^{k-\ell+1}, \mathbf{y}'^R) - \mathcal{R}^n(0^{k-\ell+1}, \mathbf{y}^R) \leq \mathcal{R}^n(\mathbf{x}, \mathbf{y}') - \mathcal{R}^n(\mathbf{x}, \mathbf{y}) \leq \mathcal{R}^n(0^\ell, \mathbf{y}'^L) - \mathcal{R}^n(0^\ell, \mathbf{y}^L).$$

By these inequalities and (3.4), it then suffices to bound

$$\begin{aligned} & \mathbb{P}(\mathcal{A}^n[(-n^{1/3}/2, n)^{k-\ell+1} \rightarrow (\mathbf{y}'^R, 1)] - \\ & \quad \mathcal{A}^n[(-n^{1/3}/2, n)^{k-\ell+1} \rightarrow (\mathbf{y}^R, 1)] < -a(y'_\ell - y_\ell)^{1/2}), \quad \text{and} \\ & \mathbb{P}(\mathcal{A}^n[(-n^{1/3}/2, n)^\ell \rightarrow (\mathbf{y}'^L, 1)] - \mathcal{A}^n[(-n^{1/3}/2, n)^\ell \rightarrow (\mathbf{y}^L, 1)] > a(y'_\ell - y_\ell)^{1/2}). \end{aligned}$$

By Lemma 2.12 applied to $\mathcal{A}^n$, for any endpoint pair starting $(-n^{1/3}/2, n)^i$ for some $i \leq k$, there is an optimizer that only uses the top k lines. By Lemma 2.6, the above two probabilities are bounded by

$$\mathbb{P}\left(2k \max_{1 \leq i \leq k, y_\ell \leq x < y \leq y'_\ell} |\mathcal{A}_i^n(x) - \mathcal{A}_i^n(y)| > a(y'_\ell - y_\ell)^{1/2}\right).$$

Rewriting this probability in terms of W^n gives

$$(3.5) \quad \mathbb{P}\left(2k \max_{1 \leq i \leq k, y_\ell \leq x < y \leq y'_\ell} \left| n^{1/6}(W_i^n(1 + 2n^{-1/3}x) - W_i^n(1 + 2n^{-1/3}y)) \right. \right. \\ \left. \left. + x^2 - y^2 - 2n^{1/3}(x - y) \right| > a(y'_\ell - y_\ell)^{1/2}\right).$$

Now, using that $|y_\ell| < 2n^{1/6}$ and $|y_\ell - y'_\ell| < 1$, for any $y_\ell \leq x < y \leq y'_\ell$ and n large enough, by a Taylor expansion we have

$$\left| \frac{2n^{1/3}(x - y)}{\sqrt{1 + 2n^{-1/3}y_\ell}} - 2n^{1/3}(x - y) + 2(x - y)y_\ell \right| \leq 3n^{-1/3}(y - x)y_\ell^2 \leq 12(y - x).$$

Therefore

$$(3.6) \quad \left| x^2 - y^2 - 2n^{1/3}(x - y) - n^{1/6} \frac{2n^{-1/3}(y - x)\sqrt{n}}{\sqrt{1 + 2n^{-1/3}y_\ell}} \right| < |x^2 - y^2 - 2(x - y)y_\ell| + 12(y - x) \\ \leq (y - x)(|x + y - 2y_\ell| + 12) \leq 14(y'_\ell - y_\ell)^{1/2}.$$

The conclusion then follows by combining (3.5), (3.6), Corollary 3.4, and a union bound. $\square$

PROOF OF THEOREM 3.1. It suffices to show that $\mathcal{S}^n|_K$ is tight for all compact sets $K \subset \mathbb{R}_\leq^k \times \mathbb{R}_\leq^k$. We may assume K contains 0. First, $\mathcal{S}^n(0^k, 0^k) = \sum_{i=1}^k \mathcal{B}_i^n(0)$, so $\mathcal{S}^n(0^k, 0^k)$ is tight by Theorem 2.13. Tightness of $\mathcal{S}^n|_K$ then follows from Lemma 3.5 and the Kolmogorov-Chentsov criterion, see Corollary 14.9 in [31]. $\square$

3.2. Tightness of melon optimizers

Here we prove tightness and asymptotic results about melon optimizers. For this, we extend the notions of geodesics and jump times from Chapter 2.5 to the case where the endpoints are not singletons. For vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}_\leq^k$ with $x_1 \geq 0$, we write

$$\pi^n\{\mathbf{x}, \mathbf{y}\} = (\pi_1^n\{\mathbf{x}, \mathbf{y}\}, \dots, \pi_k^n\{\mathbf{x}, \mathbf{y}\})$$

for the rightmost optimizer across $\mathcal{B}^n$ from $(\mathbf{x} - n^{1/3}/2, n)$ to $(\mathbf{y}, 1)$. We will write $\pi^n[\mathbf{x}, \mathbf{y}]$ for the rightmost optimizer across the original Brownian motions between the corresponding points $(2n^{-1/3}\mathbf{x}, n)$ and $(1 + 2n^{-1/3}\mathbf{y}, 1)$. We write $Z_{i,m}^n(\mathbf{x}, \mathbf{y})$ for the jump time from line $m + 1$ to m for the path $\pi_i^n\{\mathbf{x}, \mathbf{y}\}$.

We start with a weak tightness result.

LEMMA 3.6. *Let $k \in \mathbb{N}$, $x > 0$ and $y \in \mathbb{R}$, and set $x^k = (x, \dots, x)$, $y^k = (y, \dots, y) \in \mathbb{R}^k$. Then for every $i \in \llbracket 1, k \rrbracket$, the sequence of jump times $\{Z_{i,1}^n(x^k, y^k) : n \in \mathbb{N}\}$ is tight.*

PROOF. We write $\pi_i^n := \pi_i^n\{x^k, y^k\}$ for the i^{th} path in the disjoint optimizer $\pi^n\{x^k, y^k\}$. By Lemma 2.7 and (2.6), we have

$$\|\pi_i^n\|_{\mathcal{B}^n} - n^{2/3} \geq \mathcal{S}^n(x^k, y^k) - \mathcal{S}^n(x^{k-1}, y^{k-1}).$$

In particular, by Theorem 3.1, the random variables $Y_n := \|\pi_i^n\|_{\mathcal{B}^n} - n^{2/3}$ are tight. Now suppose that $Z_{i,1}^n(x^k, y^k) < r$ for some $r \in \mathbb{R}$. Then $\pi_i^n(z) = 1$ for all $z \in [r, y]$, so

$$\begin{aligned} \|\pi_i^n\|_{\mathcal{B}^n} &= \|\pi_i^n|_{[x - n^{1/3}/2, r]}\|_{\mathcal{B}^n} + \mathcal{B}_1^n(y) - \mathcal{B}_1^n(r) \\ &\leq \mathcal{B}^n[(x - n^{1/3}/2, n) \rightarrow (r, 1)] + \mathcal{B}_1^n(y) - \mathcal{B}_1^n(r). \end{aligned}$$

Therefore by (2.6) again, we have

$$(3.7) \quad Y_n \wedge 0 \leq \mathbf{1}(Z_{i,1}^n(x^k, y^k) < r) [\mathcal{S}^n(x, r) + \mathcal{S}^n(0, y) - \mathcal{S}^n(0, r)]$$

The term multiplying the indicator on right-hand side of (3.7) converges to

$$X(x, y, r) = \mathcal{S}(x, r) + \mathcal{S}(0, y) - \mathcal{S}(0, r).$$

Since $\mathcal{S}(x, y) + (x - y)^2$ has GUE Tracy-Widom distribution for all $x, y \in \mathbb{R}$ (Lemma 2.23), by a union bound, for all $x, y, r, m > 0$ we have

$$\mathbb{P}(X(x, y, r) > m - x^2 - y^2 + 2xr) \leq ce^{-dm^{3/2}},$$

for some constants $c, d > 0$. In particular, $X(x, y, r) \xrightarrow{d} -\infty$ as $r \rightarrow -\infty$ for fixed x, y . Combining this, (3.7), and the tightness of Y_n gives that

$$\lim_{r \rightarrow -\infty} \limsup_{n \rightarrow \infty} \mathbb{P}(Z_{i,1}^n(x^k, y^k) < r) = 0.$$

Since all the random variables $Z_{i,1}^n(x^k, y^k)$ are bounded above by y , this implies that the sequence $Z_{i,1}^n(x^k, y^k)$ is tight. $\square$

We can use Lemma 3.6 to prove a disjointness lemma for optimizers.

LEMMA 3.7. *Consider $\pi^n[x^k, y^k]$, the (almost surely unique) optimizer from $(2xn^{-1/3}, n)^k$ to $(1 + 2yn^{-1/3}, 1)^k$ across n independent standard Brownian motions B^n . Then for every $x, y \in \mathbb{R}, k \in \mathbb{N}$ and $\epsilon > 0$, we have that*

$$\lim_{r \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathbb{P}\left(\pi^n[x^k, y^k] \text{ is essentially disjoint from } \pi^n[x - \epsilon, y - r] \text{ and } \pi^n[x + \epsilon, y + r]\right) = 1.$$

PROOF. Throughout the proof we write $\pi^{n,k}[x, y] := \pi^n[x^k, y^k]$. First, by a union bound it suffices to show that for all $x, y \in \mathbb{R}, k \in \mathbb{N}$ and $\epsilon > 0$,

$$(3.8) \quad \lim_{r \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathbb{P}\left(\pi^{n,k}[x, y] \text{ is essentially disjoint from } \pi^n[x - \epsilon, y - r]\right) = 1,$$

and

$$(3.9) \quad \lim_{r \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathbb{P}\left(\pi^{n,k}[x, y] \text{ is essentially disjoint from } \pi^n[x + \epsilon, y + r]\right) = 1.$$

We first simplify (3.8) and (3.9). Translation invariance of Brownian increments and Brownian scaling gives that

$$B^n(t) \stackrel{d}{=} \alpha_n^{-1/2} \left(B^n(\alpha_n(t - 2n^{-1/3}(x - \epsilon))) - B^n(2\alpha_n n^{-1/3}(x - \epsilon)) \right),$$

$$\text{where } \alpha_n = \frac{1}{1 + 2(-x + \epsilon + y)n^{-1/3}},$$

and so (3.8) is equal to

$$\mathbb{P}\left(\pi^{n,k}[\epsilon + O(n^{-1/3}), 0] \text{ is essentially disjoint from } \pi^n[0, -r + O(n^{-1/3})]\right).$$

Here the $O(n^{-1/3})$ terms are small in the sense that for fixed r, x, y , there exists $c > 0$ such that $|O(n^{-1/3})| \leq cn^{-1/3}$. In particular, for large enough n , monotonicity of optimizers (Lemma 2.3) implies that this is bounded below by

$$(3.10) \quad \mathbb{P}\left(\pi^{n,k}[\epsilon/2, 0] \text{ is essentially disjoint from } \pi^n[0, -r/2]\right).$$

By applying translation invariance and Brownian scaling, we can similarly show that (3.9) is equal to

$$\mathbb{P}\left(\pi^{n,k}[0, -r + O(n^{-1/3})] \text{ is essentially disjoint from } \pi^n[\epsilon + O(n^{-1/3}), 0]\right),$$

which is again bounded below by

$$(3.11) \quad \mathbb{P}\left(\pi^{n,k}[0, -r/2] \text{ is essentially disjoint from } \pi^n[\epsilon/2, 0]\right)$$

for large enough n . Next, by Lemma 2.11 and Lemma 2.15, the probabilities (3.10) and (3.11) are the same as the corresponding probabilities with melon paths $\pi^n\{\cdot\}$ in place of the original Brownian paths $\pi^n[\cdot]$.

Now, for any $0 < b, a < c$ and $k, \ell \in \mathbb{N}$, since the melon path $\pi^n\{0^k, a^k\}$ only uses the top k lines by Lemma 2.12, $\pi^n\{0^k, a^k\}$ is disjoint from $\pi^n\{b^\ell, c^\ell\}$ whenever the jump time

$$Z_{1,k}^n(b^\ell, c^\ell) > a.$$

Therefore to prove (3.10) and (3.11), we just need to show that

$$(3.12) \quad \lim_{r \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathbb{P}(Z_{1,1}^n((\epsilon/2)^k, 0^k) > -r/2) = 1, \quad \text{and}$$

$$(3.13) \quad \lim_{r \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathbb{P}(Z_k^n(\epsilon/2, 0) > -r/2) = 1.$$

Equation (3.12) follows from the tightness of $Z_{1,1}^n(x^k, y^k)$ for fixed x, y in Lemma 3.6. Equation (3.13) follows from the tightness of $Z_k^n(\epsilon/2)$ for fixed k, ϵ in Lemma 2.20. $\square$

Lemma 3.7 can be combined with the asymptotics in Lemma 2.20 to give tightness and asymptotics for jump times on optimizers across the melon. For this next lemma, we set $(0, \infty)_{\leq}^k = (0, \infty)^k \cap \mathbb{R}_{\leq}^k$.

LEMMA 3.8. *For any $k \in \mathbb{N}$ and any compact set $K \subset (0, \infty)_{\leq}^k \times \mathbb{R}_{\leq}^k$, we have that*

$$(3.14) \quad \sup_{(\mathbf{x}, \mathbf{y}) \in K, i \in \llbracket 1, k \rrbracket} \left| Z_{i,m}^n(\mathbf{x}, \mathbf{y}) - \sqrt{\frac{m}{2x_i}} \right| = o(\sqrt{m}).$$

Moreover, for any fixed $\mathbf{x}, \mathbf{y}, m$, and i , the sequence $Z_{i,m}^n(\mathbf{x}, \mathbf{y})$ is tight in n .

PROOF. We first prove this for a single $(\mathbf{x}, \mathbf{y})$. In the notation of Lemma 2.3, for every i we have $(x_i^k, y_i^k)_{\leq 1-i}(\mathbf{x}, \mathbf{y})_{\leq k-i}(x_i^k, y_i^k)$. Therefore by that lemma, we have

$$Z_{1,m}^n(x_i^k, y_i^k) \leq Z_{i,m}^n(\mathbf{x}, \mathbf{y}) \leq Z_{k,m}^n(x_i^k, y_i^k),$$

so it suffices to prove bounds when $\mathbf{x}, \mathbf{y}$ consist only of repeated points. For this, observe that on the event $A_{\epsilon,r}$ where the melon optimizer $\pi^n\{x_i^k, y_i^k\}$ is essentially disjoint from $\pi^n\{x_i - \epsilon, -r\}$ and $\pi^n\{x_i + \epsilon, r\}$, that

$$(3.15) \quad Z_{j,m}^n(x_i^k, y_i^k) \in [Z_m^n(x_i - \epsilon, -r), Z_m^n(x_i + \epsilon, r)]$$

for all $m \in \mathbb{N}, j \in \llbracket 1, k \rrbracket$. By Lemma 2.11, essential disjointness of $\pi^n\{x_i^k, y_i^k\}$ from $\pi^n\{x_i - \epsilon, -r\}$ and $\pi^n\{x_i + \epsilon, r\}$ is equivalent to essential disjointness of the original Brownian optimizers $\pi^n[x_i^k, y_i^k]$ from $\pi^n[x_i - \epsilon, -r]$ and $\pi^n[x_i + \epsilon, r]$. Therefore by Lemma 3.7,

$$\lim_{r \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathbb{P}A_{\epsilon,r} = 1.$$

Moreover, the asymptotics of the interval on the right-hand side of (3.15) are given by Lemma 2.20. Putting these together proves (3.14) for a single point. The extension to the entire compact set follows again from monotonicity (Lemma 2.3).

Finally, the tightness claim for fixed k follows from (3.14), the definition of the notation o , and the fact that the $Z_{i,m}^n(\mathbf{x}, \mathbf{y})$ are nonincreasing in m : $Z_{i,1}^n(\mathbf{x}, \mathbf{y}) \geq Z_{i,2}^n(\mathbf{x}, \mathbf{y}) \geq \dots$ $\square$

For this next corollary, we extend the definition of path space to include paths with noncompact domains. Let $\mathcal{P}$ be the space of all nonincreasing cadlag functions from any closed interval $I \subset \mathbb{R}$ to $\mathbb{Z}$. For a sequence $\pi_n \in \mathcal{P}$, we say that $\pi^n \rightarrow \pi$ if

$$\Gamma(\pi_n) \cap [-n, n] \times \llbracket -n, n \rrbracket \rightarrow \Gamma(\pi) \cap [-n, n] \times \{-n, \dots, n\}$$

in the Hausdorff topology for all $n \in \mathbb{N}$. This is a Polish space, since the Hausdorff topology on paths whose zigzag graphs live in $[-n, n] \times \llbracket -n, n \rrbracket$ is Polish for all $n \in \mathbb{N}$.

COROLLARY 3.9. *For any $(\mathbf{x}, \mathbf{y}) \in (0, \infty)_{\leq}^k \times \mathbb{R}_{\leq}^k$, the paths $\pi^n\{\mathbf{x}, \mathbf{y}\}$ are tight in distribution in the product of k path spaces. Subsequential limits are k -tuples of nonincreasing paths $\pi_i : (-\infty, y_i] \rightarrow \mathbb{N}$. Moreover, any distributional subsequential limit $(\pi, \mathcal{B})$ of $(\pi^n\{\mathbf{x}, \mathbf{y}\}, \mathcal{B}^n)$ satisfies the following property:*

For any set of times $\mathbf{z} = (z_1, \dots, z_k)$ and $m \in \mathbb{N}$ such that $(z_i, m) \in \Gamma(\pi_i)$ for all i , the restricted paths $\{\pi_i|_{[z_i, y_i]} : i \in \llbracket 1, k \rrbracket\}$ form a disjoint optimizer in $\mathcal{B}$ from $(\mathbf{z}, k)$ to $(\mathbf{y}, 1)$.

PROOF. Tightness is immediate from the tightness of each of the jump time sequences $Z_{i,k}^n(\mathbf{x}, \mathbf{y})$ established in Lemma 3.8, and the definition of the topology on path space. Now, consider a subsequential limit $(\pi, \mathcal{B})$ of $(\pi^n\{\mathbf{x}, \mathbf{y}\}, \mathcal{B}^n)$, a coupling where $(\pi^n\{\mathbf{x}, \mathbf{y}\}, \mathcal{B}^n) \rightarrow (\pi, \mathcal{B})$ almost surely, and a set of times $\mathbf{z}$ as above. Since essential disjointness and path ordering are preversed under taking limits, on the almost sure set where this convergence holds, $\{\pi_i|_{[z_i, y_i]}\}$ is a disjoint k -tuple from $(\mathbf{z}, m)$ to $(\mathbf{y}, 1)$. Moreover, since $\mathcal{B}^n \rightarrow \mathcal{B}$ uniformly on compact sets, we have

$$(3.16) \quad \mathcal{B}^n[(\mathbf{z}, k) \rightarrow (\mathbf{y}, 1)] = \sum_{i=1}^m \|\pi_i^n\{\mathbf{x}, \mathbf{y}\}|_{[z_i, y_i]}\|_{\mathcal{B}^n} \rightarrow \sum_{i=1}^m \|\pi_i|_{[z_i, y_i]}\|_{\mathcal{B}}$$

as $n \rightarrow \infty$. Finally, $\mathcal{B}^n[(\mathbf{z}, k) \rightarrow (\mathbf{y}, 1)] \rightarrow \mathcal{B}[(\mathbf{z}, k) \rightarrow (\mathbf{y}, 1)]$ by uniform-on-compact convergence, so (3.16) implies that $\{\pi_i|_{[z_i, y_i]}\}$ is a disjoint optimizer in $\mathcal{B}$ from $(\mathbf{z}, k)$ to $(\mathbf{y}, 1)$. $\square$

NOT FOR DISTRIBUTION

CHAPTER 4

Last passage percolation across the parabolic Airy line ensemble

Having established tightness of melon optimizers and prelimiting sheets, our next goal is to construct the limits of these objects. To do this, we introduce a notion of length and last passage percolation for infinite paths in $\mathcal{B}$.

4.1. Parabolic paths, length, and geodesics in $\mathcal{B}$

A **parabolic path** across $\mathcal{B}$ from $x \geq 0$ to $z \in \mathbb{R}$ is a nonincreasing cadlag function $\pi : (-\infty, z] \rightarrow \mathbb{N}$ such that

$$(4.1) \quad \lim_{y \rightarrow -\infty} \frac{\pi(y)}{2y^2} = x.$$

For every $y < z$ define the **discrepancy of π at y** by

$$D_\pi(y) = \|\pi|_{[y,z]}\|_{\mathcal{B}} - \mathcal{B}[(y, \pi(y)) \rightarrow (z, 1)].$$

Note that $D_\pi(y) \leq 0$ for all y . We then define the length of π by

$$(4.2) \quad \|\pi\|_{\mathcal{B}} = \mathcal{S}(x, z) + \liminf_{y \rightarrow -\infty} D_\pi(y),$$

where $\mathcal{S}$ is the half-Airy sheet defined from $\mathcal{B}$ as in Definition 2.17. A parabolic path π is a **geodesic** from x to y if the length $\|\pi\|_{\mathcal{B}}$ is finite, and is maximal among all paths in $\mathcal{B}$ from x to y . A parabolic path π is **locally geodesic** if $\pi|_{[a,b]}$ is a geodesic for every compact interval $[a, b]$. We first record some basic properties of lengths and geodesics in $\mathcal{B}$. The first lemma records useful deterministic facts.

LEMMA 4.1. *Let $\mathcal{B}$ be a parabolic Airy line ensemble.*

- (i) *For any parabolic path π , the discrepancy $D_\pi(y)$ is increasing in y . In particular, the $\liminf$ on the right-hand side of (4.2) is actually a limit.*
- (ii) *A parabolic path π from x to z is a geodesic if and only if π is locally geodesic, or equivalently $\|\pi\|_{\mathcal{B}} = \mathcal{S}(x, z)$.*
- (iii) *If π_n is a sequence of parabolic paths from x_n to z_n converging to a parabolic path π from x to z , then $\limsup_{n \rightarrow \infty} \|\pi_n\|_{\mathcal{B}} \leq \|\pi\|_{\mathcal{B}}$.*

PROOF. For any parabolic path $\pi : (-\infty, z] \rightarrow \mathbb{N}$, for $y_1 < y_2 \leq z$ we have

$$\|\pi|_{[y_1,z]}\|_{\mathcal{B}} = \|\pi|_{[y_1,y_2]}\|_{\mathcal{B}} + \|\pi|_{[y_2,z]}\|_{\mathcal{B}}.$$

Combining this with the triangle inequality for last passage values in $\mathcal{B}$ between the points $(\pi(y_1), y_1)$, $(\pi(y_2), y_2)$, and $(\pi(z), z)$, we get that

$$D_\pi(y_2) - D_\pi(y_1) \geq \mathcal{B}[(y_1, \pi(y_1)) \rightarrow (y_2, \pi(y_2))] - \|\pi|_{[y_1,y_2]}\|_{\mathcal{B}} \geq 0,$$

so D_π is increasing, giving (i).

For part (ii), note that π is locally geodesic if and only if $D_\pi(y) = 0$ for all y , or equivalently, if $\|\pi\|_{\mathcal{B}} = \mathcal{S}(x, z)$. Noting that $D_\pi \leq 0$ for any path π , if $D_\pi = 0$, then π must be a geodesic. For the opposite direction, suppose that π is a path from x to y with

$$\lim_{y \rightarrow -\infty} D_\pi(y) = -c < 0.$$

Then $D_\pi(y) = -a \in [-c, 0)$ for some $y < z$. We could modify the path π by replacing $\pi|_{[y, z]}$ with a geodesic from $(y, \pi(y))$ to $(z, 1)$. The new path π' is also a parabolic path from x to z , and

$$D_{\pi'}(y') = D_\pi(y') + a$$

for all $y' < y$. Therefore $\|\pi'\|_{\mathcal{B}} > \|\pi\|_{\mathcal{B}}$, so π cannot be a geodesic.

For part (iii), observe that if $\pi_n \rightarrow \pi$, then the domains converge, and by continuity of $\mathcal{B}$, the last passage values on any compact interval $[x, y]$ also converge. In particular, $D_{\pi_n}(y) \rightarrow D_\pi(y)$ for all y . Combining this with the monotonicity from (i) and the continuity of the Airy sheet $\mathcal{S}$ (Definition 2.18) gives (iii). $\square$

Existence, uniqueness, and other basic structural results about geodesics across $\mathcal{B}$ are guaranteed by limiting results for Brownian melons.

LEMMA 4.2.

- (i) (*Uniqueness*) For any fixed $(x, y) \in [0, \infty) \times \mathbb{R}$, there exists a unique geodesic $\pi\{x, y\}$ in $\mathcal{B}$ from x to y almost surely.
- (ii) (*Existence*) Almost surely, for every $(x, y) \in [0, \infty) \times \mathbb{R}$, there exists a geodesic π in $\mathcal{B}$ from x to y . Moreover, almost surely for every $x, y \in [0, \infty) \times \mathbb{R}$, there are geodesics $\pi_L\{x, y\}$ and $\pi_R\{x, y\}$ from x to y satisfying $\pi_L\{x, y\}(t) \leq \pi(t) \leq \pi_R\{x, y\}(t)$ for any geodesic π from x to y and all $t \in (-\infty, y]$. We call $\pi_L\{x, y\}$ and $\pi_R\{x, y\}$ the **leftmost and rightmost geodesics from x to y** .
- (iii) (*Overlap in the trunk*) For a fixed $x \geq 0$ and $y, y' \in \mathbb{R}$ and any geodesics π and π' from x to y and x to y' , almost surely we have $\pi(z) = \pi'(z)$ for all sufficiently negative z .
- (iv) (*Disjointness Structure*) Let $x \geq 0$ and $y \in \mathbb{R}$. For any fixed $r > 0$, we have

$$(4.3) \quad \lim_{\epsilon \rightarrow 0^+} \mathbb{P}(\Gamma(\pi\{(x - \epsilon) \vee 0, y\}) \cap \Gamma(\pi\{x + \epsilon, y + r\}) = \emptyset) = 0.$$

Also, for any fixed $0 \leq x < x'$ and $y \in \mathbb{R}$, we have

$$(4.4) \quad \begin{aligned} \lim_{r \rightarrow \infty} \mathbb{P}(\pi\{x, y\} \text{ and } \pi\{x', y + r\} \text{ are essentially disjoint}) &= 1, & \text{and} \\ \lim_{r \rightarrow \infty} \mathbb{P}(\pi\{x, y - r\} \text{ and } \pi\{x', y\} \text{ are essentially disjoint}) &= 1. \end{aligned}$$

- (v) (*Monotonicity and tree structure*) Let Ω be the almost sure set where rightmost geodesics $\pi_R\{x, y\}$ from x to y exist for every $(x, y) \in [0, \infty) \times \mathbb{R}$. On Ω , for every $0 \leq x_1 \leq x_2$ and $y_1 \leq y_2$, we have

$$\pi_R\{x_1, y_1\}(t) \leq \pi_R\{x_2, y_2\}(t)$$

for all $t \leq y_1$, and the overlap of zigzag graphs

$$\Gamma(\pi_R\{x_1, y_1\}) \cap \Gamma(\pi_R\{x_2, y_2\})$$

is either empty, or else is the zigzag graph of a cadlag function π from a closed interval to $\mathbb{R}$.

While the proof of Lemma 4.2 is rather lengthy, the basic idea is just to use the limiting structure of Brownian melon geodesics in Theorem 2.19. This theorem guarantees that limits of Brownian melon geodesics are geodesics across the parabolic Airy line ensemble $\mathcal{B}$. Moreover, the coalescence claims from this theorem pass to coalescence results for geodesics across $\mathcal{B}$.

PROOF. We will work with a subsequence $Y \subset \mathbb{N}$ and a coupling of $\mathcal{B}^n$ and $\mathcal{B}$ so that the following conditions hold almost surely:

- (1) $\mathcal{B}^n \rightarrow \mathcal{B}$.
- (2) For all $(x, z) \in (\mathbb{Q} \cap (0, \infty)) \times \mathbb{Q}$, there exists a geodesic $\pi\{x, z\}$ across $\mathcal{B}$ from x to z such that $\pi^n\{x, z\} \rightarrow \pi\{x, z\}$.
- (3) For all $x \in \mathbb{Q} \cap (0, \infty), z < y \in \mathbb{Q}$, there exist $X_1 < x < X_2$ with $X_1, X_2 \in \mathbb{Q} \cap (0, \infty)$ such that for all large enough n , there is a point (W_n, R_n) in the zigzag graph of both $\pi^n\{X_1, z\}$ and $\pi^n\{X_2, y\}$. Moreover, $(W_n, R_n) \rightarrow (W, R)$ for some $(W, R) \in \mathbb{R} \times \mathbb{Z}$.
- (4) For any $w < y \in \mathbb{Q}$ and $n > m \in \mathbb{N}$, there is almost surely a unique geodesic in $\mathcal{B}$ from (w, n) to (y, m) .

The existence of a coupling satisfying conditions 1-3 follows from Theorem 2.19. Condition 1 is immediate from Theorem 2.19(i). Corollary 3.9 and the asymptotics in Theorem 2.19(ii) guarantees convergence of the finite geodesics $\pi^n\{x, z\}$ to a limiting parabolic path $\pi\{x, z\}$ from x to z . The second part of Corollary 3.9 guarantees that each $\pi\{x, z\}$ is locally geodesic, and hence is a geodesic by Lemma 4.1(ii). This gives condition 2.

For condition 3, Theorem 2.19(iii) guarantees that there exist $X_1 < x < X_2$ and $T \in \mathbb{R}$ such that for all large enough n , the zigzag graphs of $\pi^n\{X_1, z\}$ and $\pi^n\{X_2, y\}$ overlap on the interval $[T, z]$. Since the paths $\pi^n\{X_1, z\}|_{[T, z]}$ and $\pi^n\{X_2, y\}|_{[T, z]}$ both converge, the region of overlap also converges. Therefore we can find $(W_n, R_n) \in \Gamma(\pi^n\{X_1, z\}) \cap \Gamma(\pi^n\{X_2, y\})$ that converges to some (W, R) . Condition 4 follows from Lemma 2.15. For all proofs we work on the almost sure set where the four conditions above hold.

PROOF OF (I) FOR $x > 0$. Without loss of generality we can assume that $x, y \in \mathbb{Q}$; the general case can be dealt with by working on a version of the above coupling where $(\mathbb{Q} \cap (0, \infty)) \times \mathbb{Q}$ is replaced by $((\mathbb{Q} \cap (0, \infty)) \times \mathbb{Q}) \cup \{(x, y)\}$. The existence of such a coupling still holds in this context, see the discussion after Theorem 2.19. Suppose that π' is another geodesic from x to y , and let $z < y, z \in \mathbb{Q}$. It is enough to show that $\pi' = \pi\{x, y\}$ on the interval $[z, y]$.

Let X_1, X_2 be as in property 3 of the coupling for the triple $x, z < y$. The parabolic shape of the paths $\pi\{X_1, z\}, \pi\{x, y\}, \pi'$, and $\pi\{X_2, z\}$ ensures that for large enough $m \in \mathbb{N}$ we can find times $t_1 < t_2 < W$ and $s, s' \in (t_1, t_2)$ such that

$$\pi\{X_1, z\}(t_1) = \pi\{X_2, y\}(t_2) = \pi\{x, y\}(s) = \pi'(s') = m$$

Also, let r_1, r_2 be rational times with $s, s' \in (r_1, r_2) \subset (t_1, t_2)$. There are unique finite geodesics τ_1, τ_2 from (r_1, m) , (r_2, m) to $(y, 1)$. Since the paths $\pi\{X_1, z\}, \pi\{x, y\}, \pi'$, and $\pi\{X_2, z\}$ are locally geodesic by Lemma 4.1(ii), we can apply the monotonicity in Lemma 2.3(i) to get that

$$\pi\{X_1, z\}|_{[t_1, z]} \leq \tau_1 \leq \pi\{x, y\}|_{[s, y]} \leq \tau_2 \leq \pi\{X_2, z\}|_{[t_1, y]}.$$

The outer two inequalities imply that the point (W, R) is contained in the zigzag graphs of both τ_1 and τ_2 . Therefore by the tree structure of geodesics (Proposition

2.1) and the uniqueness of τ_1, τ_2 , the paths τ_1, τ_2 coincide on the interval $[W, y]$. The inner two inequalities above then imply that τ_1, τ_2 also coincide with $\pi\{x, y\}$ on this interval. The same holds for π' , and hence $\pi\{x, y\} = \pi'$ on $[W, y]$. Since $W \leq z$, this gives the desired claim.

PROOF OF (II) FOR $x > 0$. By Lemma 2.3, for any fixed z , the functions $(x, y) \mapsto \pi^n\{x, y\}(z)$ are nondecreasing in x and y . This property passes to the limits $\pi\{x, y\}$. Therefore for any $(x, y) \in (0, \infty) \times \mathbb{R}$, and any monotone decreasing sequences $x_n \downarrow x, y_n \downarrow y$ with (x_n, y_n) rational, the paths $\pi\{x_n, y_n\}$ have a limit in path space. This limit is a function $\pi_R\{x, y\} : (-\infty, y] \rightarrow \mathbb{N}$. Since $\pi_R\{x, y\} \leq \pi\{x_n, y_n\}$ for all n , and each π is a parabolic path from x_n to y_n , we have

$$(4.5) \quad \limsup_{z \rightarrow -\infty} \frac{\pi_R\{x, y\}(z)}{2z^2} \leq x.$$

Again by monotonicity, for any rational points $x' < x, y' < y$ and any n , we have $\pi\{x_n, y_n\} \geq \pi\{x', y'\}$. Therefore $\pi_R\{x, y\} \geq \pi\{x', y'\}$ as well, and so (4.5) is an equality with the limsup replaced by a limit, and hence $\pi_R\{x, y\}$ is a parabolic path from x to y . The fact that $\pi_R\{x, y\}$ is a geodesic follows from Lemma 4.1(iii) and continuity of the Airy sheet $\mathcal{S}$ (see Definition 2.18). This proves existence of geodesics for $x > 0$.

Next, we show that each $\pi_R\{x, y\}$ must be the rightmost geodesic from x to y . Suppose that there were another geodesic π' with $\pi'(t) > \pi_R\{x, y\}(t)$ for some $t \in (-\infty, y)$. Since both $\pi, \pi_R\{x, y\}$ are cadlag, there must exist $\epsilon > 0$ such that $\pi'(s) > \pi_R\{x, y\}(s)$ for all $s \in [t, t + \epsilon]$. Zigzag graph convergence of $\pi\{x_n, y_n\}$ to $\pi_R\{x, y\}$ implies pointwise convergence at all continuity points of $\pi_R\{x, y\}$. In particular, pointwise convergence holds for some $s \in [t, t + \epsilon]$. Therefore for all large enough n , we have

$$(4.6) \quad \pi\{x_n, y_n\}(s) = \pi_R\{x, y\}(s) < \pi'(s).$$

Now define a new function π^* on $(-\infty, x_n]$ by $\pi^*(t) = \max\{\pi\{x_n, y_n\}(t), \pi'(t)\}$ for $t \leq y$ and $\pi^* = \pi\{x_n, y_n\}$ on $(y, y_n]$. The function π^* is a parabolic path from x_n to y_n . Also, since geodesics are locally geodesic, π^* must also be locally geodesic, and hence is a geodesic from x_n to y_n . Since $\pi\{x_n, y_n\} \neq \pi^*$ by (4.6), this contradicts the uniqueness of $\pi\{x_n, y_n\}$ shown in (i). The existence of leftmost geodesics is similar.

PROOF OF (III) FOR $x > 0$. Let π, π' be two geodesics from x to two points $y < y'$. As in the proof of (i), for every $z < \min(y, y')$, we can find a time $W_z \leq z$ and a location R_z such that (W_z, R_z) lies on the zigzag graphs of both π and π' . Let $W = \{W_z : z < \min(y, y')\}$. Also, for any rational points $q < q' < \min(y, y')$, condition 4 of the coupling ensures that $\pi|_{[q, q']}$ is the unique geodesic from $(q, \pi(q))$ to $(q', \pi(q'))$, and $\pi'|_{[q, q']}$ is the unique geodesic from $(q, \pi'(q))$ to $(q', \pi'(q'))$. Therefore π and π' must agree on the half-open interval

$$[\inf(W \cap [q, q']), \sup(W \cap [q, q'])).$$

Since W is nonempty and unbounded below, and q, q' were arbitrary, this implies that $\pi(z) = \pi'(z)$ for all sufficiently negative z .

PROOF OF (IV) FOR $x > 0$. We start with (4.3). Since $x > 0$, we may replace $(x - \epsilon) \vee 0$ with $x - \epsilon$. Also, without loss of generality, we may assume that $x, y + r \in \mathbb{Q}$.

By the monotonicity of the paths $\pi\{x, z\}$ in x and z , to show the statement (4.3), it is enough to find random $X_1 < x < X_2$ with $X_1, X_2 \in \mathbb{Q}$ such that the zigzag graphs of

$$\pi\{X_1, y\}, \pi\{X_2, y + r\}$$

overlap. This follows from condition 3 of the coupling.

Equation (4.4) in the finite- n case follows from Lemma 3.7 with $k = 1$, and the translation of essential disjointness of optimizers across the original Brownian motions to essential disjointness of optimizers across the melon, Lemma 2.11. To pass to the limiting paths from the finite- n statement of Lemma 3.7, we use that essential disjointness is a closed property in path space.

PROOFS OF (I), (II), (III), AND (IV) FOR $x = 0$. The path $\pi_R\{0, y\} := 1$ is locally geodesic, and hence is always a geodesic from 0 to y by Lemma 4.1(ii). We next prove (iv) when $x = 0$ when $\pi\{0, y\}$ is replaced by the path $\pi_R\{0, y\}$. We will later show that $\pi_R\{0, y\}$ is almost surely the unique geodesic from 0 to y , proving (i).

The path $\pi_R\{0, y\}$ is the almost sure limit of the melon optimizers $\pi^n\{0, y\}$, which simply follow the top path by Lemma 2.12. In particular, (4.4) then follows from the exact same argument as in the $x \neq 0$ case.

For the first part of (iv), by Corollary 2.22 and the fact that $\pi^n\{0, y\}(z) = 1$ for all n, y, z we have $\pi\{\epsilon, y + r\} \rightarrow \pi_R\{0, y + r\}$ almost surely in path space as $\epsilon \rightarrow 0^+$. Equation (4.3) follows since $\pi^R\{0, y\} = 1$.

To prove (i), (ii) and (iii) for $x = 0$, we just need to show that almost surely for all $y \in \mathbb{R}$, $\pi_R\{0, y\}$ is the only geodesic from 0 to y . By the first part of (4.4) for the paths $\pi_R\{0, y\}$, we can work on an almost sure set where for every $z \in \mathbb{Q}$ we have $\pi\{\epsilon, z\} \rightarrow \pi_R\{0, z\}$ almost surely in path space as $\epsilon \rightarrow 0^+, \epsilon \in \mathbb{Q}$.

Now let $y \in \mathbb{R}$, and suppose that π' is any geodesic from 0 to y . For any $\epsilon \in \mathbb{Q} \cap (0, \infty)$ and $z \in \mathbb{Q} \cap (y, \infty)$, monotonicity of geodesics (Proposition 2.1) and the uniqueness of $\pi\{\epsilon, z\}$ implies that $\pi'(t) \leq \pi\{\epsilon, z\}(t)$ for $t \in (-\infty, y]$. Since $\pi\{\epsilon, z\} \rightarrow \pi_R\{0, z\}$ as $\epsilon \rightarrow 0^+$, this implies $\pi'(t) \leq \pi_R\{0, z\}(t) = 1$ for $t \in (-\infty, y]$ and hence $\pi'(t) = 1 = \pi_R\{0, y\}(t)$ for $t \in (-\infty, y]$.

PROOF OF (V). This follows from the fact that rightmost (and leftmost) geodesics are locally rightmost (and leftmost) geodesics, and the corresponding result in the finite case, Proposition 2.1. $\square$

For a parabolic path across $\mathcal{B}$, the definition of its length from (4.2) is not easy to work with, as it involves a $\liminf$ of the discrepancy. Thus we record the following lemma, which follows from (4.2) and Lemma 2.24, and says that for two parabolic paths that agree off of a compact set, their difference in length can be computed locally.

LEMMA 4.3. *The following statement holds almost surely. Let π_1, π_2 be any two parabolic paths across $\mathcal{B}$ from any point x to any points z_1, z_2 respectively, such that for some $z_0 < z_1 \wedge z_2$, we have $\pi_1(y) = \pi_2(y)$ for any $y \leq z_0$. Then*

$$\|\pi_1\|_{\mathcal{B}} - \|\pi_1|_{[z_0, z_1]}\|_{\mathcal{B}} = \|\pi_2\|_{\mathcal{B}} - \|\pi_2|_{[z_0, z_2]}\|_{\mathcal{B}}.$$

From this lemma we can deduce the following measurability result. Informally, this result says that for any parabolic path π without any ‘jump points’ inside a compact set $[a, b] \times \llbracket 1, k \rrbracket$, the length of π is determined by the values of $\mathcal{B}$ outside of $[a, b] \times \llbracket 1, k \rrbracket$ (see Figure 4.1).




FIGURE 4.1. An illustration of Lemma 4.4 and its proof: for a parabolic path π that does not have any jump point in the region $I \times \llbracket 1, k \rrbracket$, its length $\|\pi\|_{\mathcal{B}}$ is determined by $\mathcal{B}$ outside $I \times \llbracket 1, k \rrbracket$. The proof compares $\|\pi\|_{\mathcal{B}}$ with the length of π' , a parabolic path ending at z_I (the left endpoint of I). The difference is determined by $\mathcal{B}$ outside of $I \times \llbracket 1, k \rrbracket$ by Lemma 4.3. Using the definition of path length (4.2), $\|\pi'\|_{\mathcal{B}}$ is also determined by $\mathcal{B}$ outside of $I \times \llbracket 1, k \rrbracket$.

LEMMA 4.4. *Take any compact interval $I \subset \mathbb{R}$ and $k \in \mathbb{N}$. Let $\mathcal{F}$ be the σ -algebra generated by all null sets, all $\mathcal{B}_i$ for $i > k$, and $\{\mathcal{B}_i(x) : x \notin I\}$ for $1 \leq i \leq k$. Take any $x \geq 0$, and let Σ_x be the set of parabolic paths π from x to some $z \in \mathbb{R}$ such that either $\pi(y) > k$ for any $y \in (-\infty, z] \cap I$, or else $I \subset (-\infty, z]$ and π is constant on I . Let $F : \Sigma_x \rightarrow \mathbb{R}$ be the random function recording path length in $\mathcal{B}$: $F(\pi) = \|\pi\|_{\mathcal{B}}$. Then F is $\mathcal{F}$ -measurable.*

Before proceeding with the proof, let us briefly comment on the technical point that $\mathcal{F}$ contains all null sets. This is required to ensure that certain almost sure limits are $\mathcal{F}$ -measurable. Adding null sets into $\mathcal{F}$ does not give us very much extra information since the σ -algebra generated by null sets contains only the sets A with $\mathbb{P}(A) \in \{0, 1\}$.

PROOF. Let z_I be the left endpoint of I . We first show that the length of any parabolic path π' from x to z_I is $\mathcal{F}$ -measurable. From the definition of path length (4.2), it suffices to show that $\mathcal{S}(x, z_I)$ is $\mathcal{F}$ -measurable.

By (2.7), for every $y < z_I$ we have that $\mathcal{S}(x, z_I) - \mathcal{S}(x, y)$ is $\mathcal{F}$ -measurable. By (2.8) and translation invariance of $\mathcal{S}$ (Lemma 2.23), outside of a null set we have

$$\mathcal{S}(x, z_I) = \lim_{a \rightarrow \infty} \frac{1}{a} \int_{z_I - a}^{z_I} (\mathcal{S}(x, z_I) - \mathcal{S}(x, y) - (x - y)^2 + \xi) dy,$$

where ξ is the expectation of a GUE Tracy-Widom random variable. This implies that $\mathcal{S}(x, z_I)$ is $\mathcal{F}$ -measurable, thus the length of any parabolic path π' from x to z_I is $\mathcal{F}$ -measurable.

For a general parabolic path $\pi \in \Sigma_x$ (i.e., π starts at x and does not jump in $I \times \llbracket 1, k \rrbracket$), the idea is to compare the length of π with the length of some parabolic path ending at z_I using Lemma 4.3, and show that $\|\pi\|_{\mathcal{B}} - \|\pi'\|_{\mathcal{B}}$ is $\mathcal{F}$ -measurable. Indeed, one can construct a parabolic path π' from x to z_I such that $\pi(y) = \pi'(y)$ for all sufficiently negative y ; and by Lemma 4.3 applied to the paths π, π' we have that $\|\pi\|_{\mathcal{B}} - \|\pi'\|_{\mathcal{B}}$ is $\mathcal{F}$ -measurable. The conclusion follows. $\square$

This result will be used in the proof of Lemma 7.3 below, where we use the Brownian Gibbs property (Theorem 2.14) on sets of the form $I \times \llbracket 1, k \rrbracket$ to study distributional properties of parabolic path length.

4.2. Disjoint optimizers in $\mathcal{B}$

Now that we have a notion of length of parabolic paths in $\mathcal{B}$, we can define multi-point last passage values. For $(\mathbf{x}, \mathbf{y}) \in \mathfrak{X}$ with $x_1 \geq 0$, define

$$(4.7) \quad \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] = \sup_{\pi} \|\pi\|_{\mathcal{B}} := \sup_{\pi_1, \dots, \pi_k} \sum_{i=1}^k \|\pi_i\|_{\mathcal{B}}$$

where the supremum is over k -tuples of ordered, essentially disjoint parabolic paths from x_i to y_i . As in the finite case, we call such a collection π a **disjoint k -tuple** from $\mathbf{x}$ to $\mathbf{y}$, we refer to any disjoint k -tuple $\pi = (\pi_1, \dots, \pi_k)$ that attains the above supremum as a **disjoint optimizer**, as long as $\|\pi\|_{\mathcal{B}}$ is finite. We say that a k -tuple π is a **local optimizer** if for all $z \leq y_1$, the k -tuple consisting of the paths $\pi_i|_{[z, y_i]}$ is a disjoint optimizer. Note that the notation (4.7) is similar to the notation for finite last passage values. The two notations are distinguished by the lack of start and end lines in (4.7).

We first focus on understanding the structure of disjoint optimizers in $\mathcal{B}$ from distinct starting points $\mathbf{x} = (x_1 < x_2 < \dots < x_k)$ with $x_1 > 0$, as such paths are more easily related to geodesics.

PROPOSITION 4.5. *Take any $(\mathbf{x}, \mathbf{y}) \in \mathfrak{X}$ such that $\mathbf{x} = (x_1 < x_2 < \dots < x_k)$ with $x_1 > 0$.*

- (i) *Suppose that π is a disjoint optimizer from $\mathbf{x}$ to $\mathbf{y}$. Then π is locally optimal.*
- (ii) *Almost surely there is a unique optimizer $\pi = (\pi_1, \dots, \pi_k)$ from $\mathbf{x}$ to $\mathbf{y}$ in $\mathcal{B}$. Moreover, letting $\pi\{x_i, 0\}$ be the geodesic in $\mathcal{B}$ from x_i to 0, then for every i , there exists a (random) $Y \in \mathbb{R}$ such that $\pi\{x_i, 0\}(t) = \pi_i(t)$ for all $t \leq Y$.*
- (iii) *Almost surely, the only k -tuple π from $\mathbf{x}$ to $\mathbf{y}$ in $\mathcal{B}$ which is locally optimal is the unique optimizer from $\mathbf{x}$ to $\mathbf{y}$.*

The basic idea of the proof is to first look for pairs $(\mathbf{x}, \mathbf{y}^-)$ and $(\mathbf{x}, \mathbf{y}^+)$ with $y_i^- < y_i < y_i^+$ where the optimizers from $\mathbf{x}$ to $\mathbf{y}^\pm$ simply consists of the geodesics from x_i to $y_i^\pm$. Results about optimizers from $\mathbf{x}$ to $\mathbf{y}$ can then be inferred using monotonicity and coalescence results.

PROOF. Without loss of generality, we may assume that all x_i, y_i are rational. For (i), note that if any disjoint optimizer π were not locally optimal on some interval of lines $\llbracket 1, m \rrbracket$, then as in the $k = 1$ case of Lemma 4.1(ii), we can increase its length $\|\pi\|_{\mathcal{B}}$ by replacing π on those lines with an optimizer π' .

For (ii), we will work on the set where for all $x \in \mathbb{Q} \cap (0, \infty)$ and $y \in \mathbb{Q}$, there is a unique geodesic $\pi\{x, y\}$ from x to y in $\mathcal{B}$. We also assume that there is a unique optimizer in $\mathcal{B}$ from any rational starting location $\mathbf{q} = (q_1, \dots, q_k) \in (\mathbb{Q} \times \mathbb{N})^k$ to $(\mathbf{y}, 1)$. We can do this by Lemma 2.16. (In the more general case where x_i, y_i are not all rational, we just consider the set with $\mathbb{Q}$ replaced by $\mathbb{Q} \cup \{x_1, \dots, x_k, y_1, \dots, y_k\}$.)

First, by using both parts of (4.4) in Lemma 4.2 (iv), we can find rational $\epsilon > 0$ and rational points $y_i^\pm$ with

$$(4.8) \quad y_1^- \ll y_2^- \ll \dots \ll y_k^- < y_1 \leq y_k \ll y_1^+ \ll y_2^+ \dots \ll y_k^+$$

such that for any $i < j \in \llbracket 1, k \rrbracket$, the geodesics $\pi\{x_i, y_i^-\}$ and $\pi\{x_j - \epsilon, y_j^-\}$ are essentially disjoint, as are the geodesics $\pi\{x_i + \epsilon, y_i^+\}$ and $\pi\{x_j, y_j^+\}$.

Now, by monotonicity of geodesics, Lemma 4.2(v), for any rational $\delta \in (0, \epsilon)$ all of the geodesics $\gamma_\delta^- = \{\pi\{x_i - \delta, y_i^-\} : i \in \llbracket 1, k \rrbracket\}$ are essentially disjoint from each other. Similarly, the geodesics $\gamma_\delta^+ = \{\pi\{x_i + \delta, y_i^+\} : i \in \llbracket 1, k \rrbracket\}$ are also essentially disjoint. In particular, the disjoint k -tuples $\gamma_\delta^\pm$ are optimizers. They are also locally optimal by Lemma 4.1(ii).

We use these geodesics to prove the existence of an optimizer from $\mathbf{x}$ to $\mathbf{y}$. We start with a disjoint k -tuple $\pi^1 = (\pi_1^1, \dots, \pi_k^1)$ from $\mathbf{x}$ to $\mathbf{y}$. The k -tuple π^1 can be obtained from the geodesics $\pi_1, \dots, \pi_k$ from x_i to y_i in the following way. Since these paths have different asymptotic directions x_i , there exists $y^* \in \mathbb{R}$ such that $\pi_i(y) \neq \pi_j(y)$ for all $y < y^*$, and $i \neq j$. Modifying these geodesics in any way for $y > y^*$ to ensure essential disjointness and ordering gives a finite length k -tuple from $\mathbf{x}$ to $\mathbf{y}$. Each of the paths in this modification has finite length by Lemma 4.3.

Therefore $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] > -\infty$. Let $\pi^m = (\pi_1^m, \dots, \pi_k^m)$ be a sequence of k -tuples from $\mathbf{x}$ to $\mathbf{y}$ whose lengths converge to the supremum in (4.7). The asymptotic growth rate of parabolic paths guarantees that there exists a sequence $z_m \rightarrow -\infty$ such that

$$(4.9) \quad \pi\{x_i - 1/m, y_i^-\}(z_m) < \pi_i^m(z_m) < \pi\{x_i + 1/m, y_i^+\}(z_m)$$

for all $i \in \llbracket 1, k \rrbracket$. Next, for each m , modify π^m so that $\pi^m|_{(-\infty, z_m]}$ is an optimizer. Doing this can only increase the length, so the new path lengths still converge to the supremal value $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$. Moreover, since the k -tuples $\gamma_{1/m}^\pm$ are locally optimal, (4.9), (4.8), and monotonicity of optimizers (Lemma 2.3) implies that for $t \in [z_m, y_i^-]$, we have

$$\pi\{x_i - 1/m, y_i^-\}(t) \leq \pi_i^m(t) \leq \pi\{x_i + 1/m, y_i^+\}(t).$$

Now, as $m \rightarrow \infty$, each of the path collections $\pi\{x_i - 1/m, y_i^-\}$ converges to $\pi\{x_i, y_i^-\}$ in path space. Similarly each of the paths $\pi\{x_i + 1/m, y_i^+\}$ converges to $\pi\{x_i, y_i^+\}$. This, and monotonicity of geodesics implies that the sequence of k -tuples π^m is precompact in the product of k path spaces, with subsequential limits π that satisfy

$$(4.10) \quad \pi\{x_i, y_i^-\} \leq \pi_i \leq \pi\{x_i, y_i^+\}$$

for all i , and are locally optimal. This implies that π is also a k -tuple from $\mathbf{x}$ to $\mathbf{y}$. Since essential disjointness and ordering are preserved under limits, π is a disjoint k -tuple from $\mathbf{x}$ to $\mathbf{y}$. Also, Lemma 4.1(iii) implies that $\|\pi\|_{\mathcal{B}} \geq \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$ so π is a disjoint optimizer.

Next, we establish uniqueness of π by establishing (iii). This also completes the proof of (ii). Let γ be another k -tuple from $\mathbf{x}$ to $\mathbf{y}$ which is locally optimal, and let $\gamma^m = \gamma$ for all m . By a similar argument as above with γ^m used in place of π^m , the bounds (4.10) also hold with γ_i in place of π_i . Next, Lemma 4.2(iii) and (4.10) imply that there exists some y^* such that for all i and all $z < y^*$, $\pi_i(z) = \gamma_i(z) = \pi\{x_i, 0\}(z)$. Therefore γ, π are both locally optimal paths which are equal at their endpoints.

Also, we can find rational points $\{(z_i, m_i) : i \in \llbracket 1, k \rrbracket\} \in (\mathbb{Q} \cap (-\infty, y^*)) \times \mathbb{N}$ such that $\gamma_i(z_i) = \pi_i(z_i) = m_i$ for all i . Since we are working on an almost sure set where there are unique optimizers between all rational starting locations and $(\mathbf{y}, 1)$, this implies $\gamma = \pi$. $\square$

Understanding the structure of optimizers in $\mathcal{B}$ from general starting points is more difficult. We will wait until the construction of the extended Airy sheet to do this.

NOT FOR DISTRIBUTION

NOT FOR DISTRIBUTION

CHAPTER 5

Limits of Melon Optimizers and the Extended Airy Sheet

5.1. Convergence of $\mathcal{S}^n$

We now have the tools to obtain both the scaling limit of $\mathcal{S}^n$, and the joint scaling limit of melon optimizers. We will focus on first understanding the scaling limit of $\mathcal{S}^n$ on the set

$$\hat{\mathfrak{X}} = \{(\mathbf{x}, \mathbf{y}) \in \mathfrak{X} : 0 \leq x_1 \leq \cdots \leq x_k\}$$

Also, let

$$\hat{\mathbb{Q}} = \{(\mathbf{x}, \mathbf{y}) \in \hat{\mathfrak{X}} : x_i \in \mathbb{Q} \cap (0, \infty), y_i \in \mathbb{Q} \ \forall i \in \llbracket 1, k \rrbracket, \text{ and } 0 < x_1 < x_2 < \cdots < x_k\}.$$

Note that $\hat{\mathfrak{X}}$ is the closure of $\hat{\mathbb{Q}}$. By Theorems 2.13, 3.1, and Corollary 3.9, the functions $\mathcal{B}^n, \mathcal{S}^n|_{\hat{\mathfrak{X}}}$ and the paths $\{\pi^n\{\mathbf{x}, \mathbf{y}\} : (\mathbf{x}, \mathbf{y}) \in \hat{\mathbb{Q}}\}$ are jointly tight. Here tightness is with respect to the following topologies:

- For $\mathcal{B}^n$: uniform-on-compact convergence of functions from $\mathbb{R} \times \mathbb{Z}$ to $\mathbb{R}$.
- For $\mathcal{S}^n|_{\hat{\mathfrak{X}}}$: uniform-on-compact convergence of functions from $\mathfrak{X}$ to $\mathbb{R}$.
- For each of the paths $\pi^n\{\mathbf{x}, \mathbf{y}\}$: the path space topology defined at the beginning of Chapter 2.

Let

$$\mathcal{B}, \mathcal{S}, \{\pi\{\mathbf{x}, \mathbf{y}\} : (\mathbf{x}, \mathbf{y}) \in \hat{\mathbb{Q}}\}$$

be any joint distributional subsequential limit along some subsequence Y . In this chapter, we will understand the joint structure of these limiting objects. We start with a lemma and a proposition.

LEMMA 5.1. *There exists a subsequence $Y' \subset Y$ such that almost surely,*

$$(\mathcal{B}^n, \mathcal{S}^n, \{\pi^n\{\mathbf{x}, \mathbf{y}\} : (\mathbf{x}, \mathbf{y}) \in \hat{\mathbb{Q}}\}) \rightarrow (\mathcal{B}, \mathcal{S}, \{\pi\{\mathbf{x}, \mathbf{y}\} : (\mathbf{x}, \mathbf{y}) \in \hat{\mathbb{Q}}\}),$$

and for every $\mathbf{x} \in \mathbb{Q}^k$ with $0 < x_1 < \cdots < x_k$, there exist rational points $z_1 < z_2 < \cdots < z_k$ such that for all large enough n , the paths $\pi^n\{x_i, z_i\}$ are essentially disjoint.

PROOF. For $(\mathbf{x}, \mathbf{z}) \in \hat{\mathbb{Q}}$, define the indicator

$$D^n(\mathbf{x}, \mathbf{z}) = \mathbf{1}(\{\pi^n\{x_i, z_i\}\}_{i=1}^k \text{ are essentially disjoint}).$$

We can find a subsequence $Y' \subset Y$ such that the random variables

$$\mathcal{B}^n, \mathcal{S}^n, \{\pi^n\{\mathbf{x}, \mathbf{y}\} : (\mathbf{x}, \mathbf{y}) \in \hat{\mathbb{Q}}\}, \{D^n(\mathbf{x}, \mathbf{z}) : (\mathbf{x}, \mathbf{z}) \in \hat{\mathbb{Q}}\}$$

converge jointly in distribution. By Skorokhod's representation theorem, we can couple the environments along Y' so that this convergence takes place almost surely.

Finally, by Lemma 3.7 and Lemma 2.11, for every $\mathbf{x} \in \mathbb{Q}^k$ with $0 < x_1 < \dots < x_k$ and any $\epsilon > 0$, we can find $\mathbf{z}$ such that

$$\liminf_{n \rightarrow \infty} \mathbb{E} D_n(\mathbf{x}, \mathbf{z}) \geq 1 - \epsilon.$$

Therefore on this coupling, the paths $\pi^n\{x_i, z_i\}$ are essentially disjoint for all large enough n with probability at least $1 - \epsilon$. Since $\epsilon > 0$ was arbitrary, this holds almost surely for some rational $\mathbf{z}$. $\square$

PROPOSITION 5.2. *With notation as above, almost surely the following statements hold.*

- (1) *For all $(\mathbf{x}, \mathbf{y}) \in \hat{\mathbb{Q}}$, we have $\mathcal{S}(\mathbf{x}, \mathbf{y}) = \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$. In particular, by continuity $\mathcal{S}|_{\hat{\mathbb{Q}}}$ is a function of $\mathcal{B}$.*
- (2) *For all $(\mathbf{x}, \mathbf{y}) \in \hat{\mathbb{Q}}$, the k -tuple $\pi\{\mathbf{x}, \mathbf{y}\}$ is the unique optimizer in $\mathcal{B}$ from $\mathbf{x}$ to $\mathbf{y}$.*

PROOF. First, Corollary 3.9 and Lemma 3.8 ensure that each of the $\pi\{\mathbf{x}, \mathbf{y}\}$ for $(\mathbf{x}, \mathbf{y}) \in \hat{\mathbb{Q}}$ is a disjoint k -tuple in $\mathcal{B}$ from $\mathbf{x}$ to $\mathbf{y}$ which is locally optimal. Proposition 4.5(iii) then implies that $\pi\{\mathbf{x}, \mathbf{y}\}$ is the unique optimizer in $\mathcal{B}$ from $\mathbf{x}$ to $\mathbf{y}$, yielding statement 2.

For statement 1, we first observe that by Theorem 2.19(i), Lemma 4.1(ii), and the definition (4.7), we have $\mathcal{S}(x, y) = \mathcal{B}[x \rightarrow y]$ for $x > 0$ and $y \in \mathbb{R}$. Therefore to complete the proof it suffices to show that for every $(\mathbf{x}, \mathbf{y}) \in \hat{\mathbb{Q}}$ we can find rational points $\mathbf{z}$ such that

$$(5.1) \quad \mathcal{S}(\mathbf{x}, \mathbf{y}) - \sum_{i=1}^k \mathcal{S}(x_i, z_i) = \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] - \sum_{i=1}^k \mathcal{B}[x_i \rightarrow z_i].$$

To prove (5.1), we work with the subsequence Y' and the coupling in Lemma 5.1. On this coupling, there exists $\mathbf{z}$ with $(\mathbf{x}, \mathbf{z}) \in \hat{\mathbb{Q}}$ such that $\pi^n\{x_i, z_i\}$ are essentially disjoint for all large enough $n \in Y'$. Since essential disjointness is a closed condition, the paths $\pi\{x_i, z_i\}$ are also essentially disjoint. Moreover, by Proposition 4.5(ii), there exists some $T \in \mathbb{R}$ such that for all $t \leq T$, we have

$$(5.2) \quad \pi\{x_i, z_i\}(t) = \pi_i\{\mathbf{x}, \mathbf{y}\}(t).$$

In particular, by Lemma 4.3,

$$(5.3) \quad \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] - \sum_{i=1}^k \mathcal{B}[x_i \rightarrow z_i] = \sum_{i=1}^k \|\pi_i\{\mathbf{x}, \mathbf{y}\}|_{[T, y_i]}\|_{\mathcal{B}} - \|\pi\{x_i, z_i\}|_{[T, z_i]}\|_{\mathcal{B}},$$

Also, (5.2) and the convergence of paths in this coupling implies that there exists $T_n \rightarrow T$ such that for all large enough $n \in Y'$, we have

$$(5.4) \quad \pi^n\{x_i, z_i\}(T_n) = \pi_i^n\{\mathbf{x}, \mathbf{y}\}(T_n).$$

When the $\pi^n\{x_i, z_i\}$ are essentially disjoint, this equality also holds for all $y < T_n$. In particular, this holds for all large enough n , and so by (2.6),

$$\mathcal{S}^n(\mathbf{x}, \mathbf{y}) - \sum_{i=1}^k \mathcal{S}^n(x_i, z_i) = \sum_{i=1}^k \|\pi_i^n\{\mathbf{x}, \mathbf{y}\}|_{[Y_n, y_i]}\|_{\mathcal{B}^n} - \|\pi^n\{x_i, z_i\}|_{[Y_n, z_i]}\|_{\mathcal{B}^n}.$$

Since the paths $\pi^n\{x_i, z_i\}, \pi^n\{\mathbf{x}, \mathbf{y}\}$ converge to $\pi\{x_i, z_i\}, \pi\{\mathbf{x}, \mathbf{y}\}$ and $\mathcal{B}^n$ converges uniformly to $\mathcal{B}$, the right-hand side above converges to the right-hand side of (5.3). The left-hand side above converges to the left-hand side of (5.1), yielding (5.1). $\square$

Proposition 5.2 uniquely determines $\mathcal{S}$ on $\hat{\mathfrak{X}}$ by continuity. This uniquely determines the distribution of $\mathcal{S}$ by translation invariance.

DEFINITION 5.3. Let $\mathcal{C}(\mathfrak{X}, \mathbb{R})$ be the space of continuous functions from $\mathfrak{X}$ to $\mathbb{R}$ with the topology of uniform convergence. A random function $\mathcal{S} \in \mathcal{C}(\mathfrak{X}, \mathbb{R})$ is an **extended Airy sheet** if

- $\mathcal{S}$ can be coupled with a parabolic Airy line ensemble $\mathcal{B}$ so that

$$\mathcal{S}(\mathbf{x}, \mathbf{y}) = \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$$

for all $(\mathbf{x}, \mathbf{y}) \in \hat{\mathbb{Q}}$.

- For a vector $\mathbf{x} \in \mathbb{R}^k$ for some k , let $T_c \mathbf{x}$ denote the shifted vector $(x_1 + c, \dots, x_k + c)$. We can think of T_c as an operator acting on all of $\bigcup_{i=1}^{\infty} \mathbb{R}^k$. In particular, T_c acts on all of $\mathfrak{X}$ and $T_c(\mathfrak{X}) = \mathfrak{X}$. With this definition, for all $c \in \mathbb{R}$ we have

$$\mathcal{S} \stackrel{d}{=} \mathcal{S} \circ T_c.$$

The above definition clearly yields a unique distribution on $\mathcal{C}(\mathfrak{X}, \mathbb{R})$. Moreover, we have the following theorem. This theorem encompasses Theorem 1.2.

THEOREM 5.4. *The prelimits $\mathcal{S}^n$ converge in distribution to an extended Airy sheet $\mathcal{S}$.*

PROOF. Any subsequential limit $\mathcal{S}$ of $\mathcal{S}^n$ satisfies the first property of Definition 5.3 by Proposition 5.2. Moreover, for all c , translation invariance of Brownian increments guarantees that $\mathcal{S}^n \stackrel{d}{=} \mathcal{S}^n \circ T_c$, and so $\mathcal{S}$ also satisfies the second property of Definition 5.3. $\square$

5.2. Properties of the extended Airy sheet $\mathcal{S}$

In this section we record a few basic properties of the extended Airy sheet $\mathcal{S}$, and use these properties to better understand the structure of optimizers in $\mathcal{B}$. The culmination of this section will be a proof of (the remaining parts of) Theorem 1.3. We start with basic symmetries. For this lemma recall that for $\mathbf{x} = (x_1, \dots, x_k) \in \mathbb{R}_{\leq}^k$ with a slight abuse of notation we write $-\mathbf{x} = (-x_k, \dots, -x_1)$.

LEMMA 5.5. *The extended Airy sheet $\mathcal{S}$ satisfies $\mathcal{S}(\mathbf{x}, \mathbf{y}) \stackrel{d}{=} \mathcal{S}(-\mathbf{y}, -\mathbf{x})$, jointly in all $\mathbf{x}, \mathbf{y}$. Moreover, the parabolically shifted sheet*

$$\mathcal{R}(\mathbf{x}, \mathbf{y}) := \mathcal{S}(\mathbf{x}, \mathbf{y}) + \sum_{i=1}^k (x_i - y_i)^2$$

is stationary in the sense that for any $c_1, c_2 \in \mathbb{R}$, we have

$$\mathcal{R}(T_{c_1} \mathbf{x}, T_{c_2} \mathbf{y}) \stackrel{d}{=} \mathcal{R}(\mathbf{x}, \mathbf{y})$$

jointly in all $(\mathbf{x}, \mathbf{y}) \in \mathfrak{X}$. Here the shifts T_{c_i} are as in Definition 5.3.

PROOF. The first distributional equality follows from the distributional equality $B(\cdot) \stackrel{d}{=} B(1) - B(1 - \cdot)$ for Brownian motion. By the second part of Definition 5.3, it is enough to prove the second equality when $c_1 = 0$. Let $\mathcal{R}^n(\mathbf{x}, \mathbf{y}) = \mathcal{S}^n(\mathbf{x}, \mathbf{y}) + \sum_{i=1}^k (x_i - y_i)^2$ and $\alpha_n = 1 + 2c_2 n^{-1/3}$. By Brownian scaling,

$$\mathcal{R}^n(\mathbf{x}, \mathbf{y}) \stackrel{d}{=} \alpha_n^{-1/2} \mathcal{R}^n(\alpha_n \mathbf{x}, T_{c_2} \alpha_n \mathbf{y}) + e_n(\mathbf{x}, \mathbf{y})$$

jointly in $\mathbf{x}, \mathbf{y}$, where the error term $e_n(\mathbf{x}, \mathbf{y})$ term is deterministic and converges to 0 uniformly on compact sets. Therefore since $\mathcal{R}^n(\mathbf{x}, \mathbf{y}) \xrightarrow{d} \mathcal{R}$ in the uniform-on-compact topology, $\mathcal{R}$ is continuous, and $\alpha_n \rightarrow 1$, we also have $\mathcal{R}^n(\mathbf{x}, T_{c_2}\mathbf{y}) \xrightarrow{d} \mathcal{R}$. $\square$

Using Theorem 5.4 to pass Lemma 3.5 to the limit, we get the following result for the parabolically shifted sheet $\mathcal{R}$.

LEMMA 5.6. *Take any $k \in \mathbb{N}$, $\mathbf{u} = (\mathbf{x}, \mathbf{y}), \mathbf{u}' = (\mathbf{x}', \mathbf{y}') \in \mathbb{R}_{\leq}^k$ with $\|\mathbf{u} - \mathbf{u}'\|_2 < 1$, and $a > 0$. Then*

$$\mathbb{P}(|\mathcal{R}(\mathbf{x}', \mathbf{y}') - \mathcal{R}(\mathbf{x}, \mathbf{y})| > a\sqrt{\|\mathbf{u} - \mathbf{u}'\|_2}) < ce^{-da^{3/2}},$$

for some constants $c, d > 0$ depending only on k .

We will use this continuity bound to show that $\mathcal{S}(\mathbf{x}, \mathbf{y}) = \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$ for all $(\mathbf{x}, \mathbf{y}) \in \hat{\mathcal{X}}$. First, we record an analogue of Lemma 2.4 from $\mathcal{B}$.

LEMMA 5.7. *Take any $\mathbf{x}, \mathbf{y}, \mathbf{x}', \mathbf{y}' \in \mathbb{R}_{\leq}^k$ such that $x_1, x'_1 \geq 0$, and define $\mathbf{x}^\ell, \mathbf{y}^\ell, \mathbf{x}^r, \mathbf{y}^r \in \mathbb{R}_{\leq}^k$ by $x_i^\ell = x_i \wedge x'_i$, $y_i^\ell = y_i \wedge y'_i$, and $x_i^r = x_i \vee x'_i$, $y_i^r = y_i \vee y'_i$, for each $1 \leq i \leq k$. Then*

$$\mathcal{B}[\mathbf{x}^\ell \rightarrow \mathbf{y}^\ell] + \mathcal{B}[\mathbf{x}^r \rightarrow \mathbf{y}^r] \geq \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] + \mathcal{B}[\mathbf{x}' \rightarrow \mathbf{y}'].$$

PROOF. First, the inequality is trivial if either $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$ or $\mathcal{B}[\mathbf{x}' \rightarrow \mathbf{y}']$ is $-\infty$, so we may assume both are finite. Let π_n, π'_n be sequences of disjoint k -tuples from $\mathbf{x}$ to $\mathbf{y}$ and $\mathbf{x}'$ to $\mathbf{y}'$ whose weights converge to $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}], \mathcal{B}[\mathbf{x}' \rightarrow \mathbf{y}']$ as in (4.7). As in the proof of Lemma 2.4, we define disjoint k -tuples τ_n^ℓ, τ_n^r from $\mathbf{x}^\ell$ to $\mathbf{y}^\ell$, $\mathbf{x}^r$ to $\mathbf{y}^r$, by (for each $1 \leq i \leq k$) setting $\tau_{n,i}^\ell = \pi_{n,i} \wedge \pi'_{n,i}$, $\tau_{n,i}^r = \pi_{n,i} \vee \pi'_{n,i}$ on $(-\infty, y_i^\ell]$, and setting $\tau_{n,i}^r$ to be either $\pi_{n,i}$ or $\pi'_{n,i}$ on $(y_i^\ell, y_i^r]$, depending on whether y_i^r equals y_i or y'_i . Then as in the proof of Lemma 2.2, one can check that τ_n^ℓ, τ_n^r are disjoint k -tuples in $\mathcal{B}$ from $\mathbf{x}^\ell$ to $\mathbf{y}^\ell$ and from $\mathbf{x}^r$ to $\mathbf{y}^r$, respectively. To prove the lemma, we just need to show that

$$(5.5) \quad \|\tau_n^\ell\|_{\mathcal{B}} + \|\tau_n^r\|_{\mathcal{B}} = \|\pi_n\|_{\mathcal{B}} + \|\pi'_n\|_{\mathcal{B}}.$$

Indeed, for any parabolic path π from some $x \geq 0$ to $z \in \mathbb{R}$, and $y \leq z$, we denote

$$P(\pi, y) = \|\pi|_{[y, z]}\|_{\mathcal{B}} - \mathcal{B}[(y, \pi(y)) \rightarrow (z, 1)] + \mathcal{S}(x, z).$$

Then from the definition of the path length (4.2), we just need to verify that

$$\lim_{y \rightarrow -\infty} \sum_{i=1}^k P(\tau_{n,i}^\ell, y) + P(\tau_{n,i}^r, y) - P(\pi_{n,i}, y) - P(\pi'_{n,i}, y) = 0.$$

This follows from Lemma 2.24 and the fact that for any $y \leq y_1^\ell$, we have

$$\sum_{i=1}^k \|\tau_{n,i}^\ell|_{[y, y_i^\ell]}\|_{\mathcal{B}} + \|\tau_{n,i}^r|_{[y, y_i^r]}\|_{\mathcal{B}} - \|\pi_{n,i}|_{[y, y_i]}\|_{\mathcal{B}} - \|\pi'_{n,i}|_{[y, y'_i]}\|_{\mathcal{B}} = 0. \quad \square$$

We next study the extended Airy sheet and optimizers in $\mathcal{B}$ for all endpoints in $\hat{\mathcal{X}}$ (rather than just $\mathbb{Q}$).

PROPOSITION 5.8. *The function $(\mathbf{x}, \mathbf{y}) \mapsto \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$ is continuous on $\hat{\mathcal{X}}$. In particular,*

$$\mathcal{S}(\mathbf{x}, \mathbf{y}) = \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$$

for all $(\mathbf{x}, \mathbf{y}) \in \hat{\mathcal{X}}$. Moreover, almost surely, for any $(\mathbf{x}, \mathbf{y}) \in \hat{\mathcal{X}}$ there is an optimizer in $\mathcal{B}$ from $\mathbf{x}$ to $\mathbf{y}$.

Proposition 5.8 is the final piece of Theorem 1.3. The proof is lengthy and a bit nuanced, so we give a sketch here.

We know that $\mathcal{S}(\mathbf{x}, \mathbf{y}) = \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$ for $(\mathbf{x}, \mathbf{y}) \in \hat{\mathbb{Q}}$ and $\mathcal{S}$ is continuous. Moreover, it is not difficult to show that $\mathcal{B}$ is upper semicontinuous using Lemma 4.1(iii). Therefore the main goal is to show that $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] \leq \mathcal{S}(\mathbf{x}, \mathbf{y})$ for general $\mathbf{x}, \mathbf{y}$. We first prove this for rational $\mathbf{x}$ with $0 < x_1 < \dots < x_n$. In this case, we can upper bound $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$ by $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}'] + \epsilon$ for a nearby rational $\mathbf{y}'$ by making small modifications to a candidate optimizer from $\mathbf{x}$ to $\mathbf{y}$.

For general $\mathbf{x}$, by the quadrangle inequality (Lemma 5.7) and an approximation of $\mathbf{x}$ by a rational sequence, we have that $\mathcal{S}(\mathbf{x}, \mathbf{y}) - \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$ is non-decreasing in $\mathbf{y}$. Then it is enough to find a sequence $\mathbf{z}_n \rightarrow (-\infty, \dots, -\infty)$ with

$$(5.6) \quad \limsup_{n \rightarrow \infty} \mathcal{S}(\mathbf{x}, \mathbf{z}_n) - \mathcal{B}[\mathbf{x} \rightarrow \mathbf{z}_n] \geq 0.$$

The basic idea here is to choose the $\mathbf{z}_n$ so that with probability tending to 1 with n , we have

$$\mathcal{S}(\mathbf{x}, \mathbf{z}_n) = \sum_{i=1}^k \mathcal{S}(x_i, z_{n,i}) \geq \mathcal{B}[\mathbf{x} \rightarrow \mathbf{z}_n].$$

Note that there is a subtlety here because we need (5.6) to hold for all $\mathbf{x}$ simultaneously. We work around this with an approximation argument, making use of the strong modulus of continuity on $\mathcal{S}$ from Lemma 5.6.

PROOF OF PROPOSITION 5.8. Throughout the proof, we fix $k \in \mathbb{N}$ as the size of the set of points $\mathbf{x}, \mathbf{y}$ we work with. All points $\mathbf{x}$ have $x_1 \geq 0$. By Proposition 5.2, $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$ almost surely coincides with the continuous function $\mathcal{S}(\mathbf{x}, \mathbf{y})$ at all points in $\hat{\mathbb{Q}}$, and there are unique optimizers $\pi\{\mathbf{x}, \mathbf{y}\}$ in $\mathcal{B}$ for all these points by Proposition 4.5. Now consider an arbitrary point $(\mathbf{x}, \mathbf{y}) \in \hat{\mathcal{X}}$. We can approximate $(\mathbf{x}, \mathbf{y})$ by a sequence of points $(\mathbf{x}_n, \mathbf{y}_n) \in \hat{\mathbb{Q}}$ such that $x_{n,i} > x_i$ and $y_{n,i} > y_i$ for all i . Now, the collection of optimizers $\{\pi\{\mathbf{x}, \mathbf{y}\} : (\mathbf{x}, \mathbf{y}) \in \hat{\mathbb{Q}}\}$ is monotone in $\mathbf{x}$ and $\mathbf{y}$; this is inherited from the prelimiting monotonicity, which follows from Lemma 2.3. Therefore as in the proof of Lemma 4.2(ii), monotonicity of optimizers guarantees that the k -tuples $\pi\{\mathbf{x}_n, \mathbf{y}_n\}$ have a limit, which is itself a disjoint k -tuple $\pi_R\{\mathbf{x}, \mathbf{y}\}$ from $\mathbf{x}$ to $\mathbf{y}$. Lemma 4.1(iii) implies that

$$\|\pi_R\{\mathbf{x}, \mathbf{y}\}\|_{\mathcal{B}} \geq \mathcal{S}(\mathbf{x}, \mathbf{y}) = \lim_{n \rightarrow \infty} \mathcal{B}[\mathbf{x}_n \rightarrow \mathbf{y}_n].$$

Therefore $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] \geq \mathcal{S}(\mathbf{x}, \mathbf{y})$ for all $(\mathbf{x}, \mathbf{y}) \in \hat{\mathcal{X}}$. If we can show the opposite inequality, then the path $\pi_R\{\mathbf{x}, \mathbf{y}\}$ is an optimizer, and $\mathcal{S}(\mathbf{x}, \mathbf{y}) = \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$. Since $\mathcal{S}$ is continuous, this will complete the proof of the proposition.

For this, we first prove that for a fixed rational $\mathbf{x}$ with $0 < x_1 < \dots < x_k$, we have $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] = \mathcal{S}(\mathbf{x}, \mathbf{y})$ for all $\mathbf{y} \in \mathbb{R}_{\leq}^k$. Suppose that $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] > \mathcal{S}(\mathbf{x}, \mathbf{y})$ for some $\mathbf{y}$. Then by continuity of $\mathcal{S}$, there is an $\epsilon > 0$ and a disjoint k -tuple π in $\mathcal{B}$ from $\mathbf{x}$ to $\mathbf{y}$ with

$$(5.7) \quad \|\pi\|_{\mathcal{B}} > \mathcal{S}(\mathbf{x}, \mathbf{y}') + \epsilon$$

for all $\mathbf{y}'$ with $|\mathbf{y} - \mathbf{y}'| < \epsilon$. Now, Lemma 4.3 and the continuity of $\mathcal{B}$ and $\mathcal{S}$ ensures that there exists a rational k -tuple $\mathbf{y}' < \mathbf{y}$ (coordinatewise) and $|\mathbf{y} - \mathbf{y}'| < \epsilon$ such that the path π' from $\mathbf{x}$ to $\mathbf{y}'$ defined by $\pi'_i = \pi_i|_{(-\infty, y'_i]}$ satisfies

$$\|\pi'\|_{\mathcal{B}} > \|\pi\|_{\mathcal{B}} - \epsilon.$$

This is greater than $\mathcal{S}(\mathbf{x}, \mathbf{y}')$ by (5.7). On the other hand, $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}'] \geq \|\pi'\|_{\mathcal{B}}$ and $\mathcal{S}(\mathbf{x}, \mathbf{y}') = \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}']$, giving a contradiction.

Now consider general $\mathbf{x}$ and let $\mathbf{y}' \leq \mathbf{y} \in \mathbb{R}_{\leq}^k$. Consider a sequence of rational $\mathbf{x}_n$ with $0 < x_{n,1} < \dots < x_{n,k}$ such that $x_{n,i} \downarrow x_i$. By Lemma 5.7, we have

$$\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] - \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}'] \leq \mathcal{B}[\mathbf{x}_n \rightarrow \mathbf{y}] - \mathcal{B}[\mathbf{x}_n \rightarrow \mathbf{y}'].$$

Since the $\mathbf{x}_n$ have distinct positive rational entries, the right-hand side above is equal to the same difference with $\mathcal{S}(\cdot, \cdot)$ in place of $\mathcal{B}[\cdot \rightarrow \cdot]$. Therefore by continuity of $\mathcal{S}$, we have

$$(5.8) \quad \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] - \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}'] \leq \mathcal{S}(\mathbf{x}, \mathbf{y}) - \mathcal{S}(\mathbf{x}, \mathbf{y}')$$

for all $\mathbf{x}$ and $\mathbf{y}' \leq \mathbf{y}$. To complete the proof that $\mathcal{S}(\mathbf{x}, \mathbf{y}) \geq \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}]$ it just suffices to show that we can find a sequence $\mathbf{z}_n^-$ such that $z_{n,i}^- \rightarrow -\infty$ as $n \rightarrow \infty$ for all $i \in \llbracket 1, k \rrbracket$, and

$$(5.9) \quad \limsup_{n \rightarrow \infty} \mathcal{S}(\mathbf{x}, \mathbf{z}_n^-) - \mathcal{B}[\mathbf{x} \rightarrow \mathbf{z}_n^-] \geq 0$$

for every $\mathbf{x}$. Indeed, for any $\mathbf{x}, \mathbf{y}$, (5.8) gives that

$$\mathcal{S}(\mathbf{x}, \mathbf{y}) - \mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] \geq \mathcal{S}(\mathbf{x}, \mathbf{z}_n^-) - \mathcal{B}[\mathbf{x} \rightarrow \mathbf{z}_n^-]$$

for all large enough n , so (5.9) gives that $\mathcal{B}[\mathbf{x} \rightarrow \mathbf{y}] \leq \mathcal{S}(\mathbf{x}, \mathbf{y})$, as desired.

Let $\mathbf{x}_n \in \mathbb{R}_{\leq}^k$ be any sequence of points with distinct positive rational entries, such that any $\mathbf{x} \in \mathbb{R}_{\leq}^k$ satisfies $|\mathbf{x}_n - \mathbf{x}| < n^{-1/k}$ for infinitely many n . The fact that we can find such a sequence is a consequence of the fact that the Lebesgue measure of the ball $B(\mathbf{x}_n, n^{-1/k})$ is $O(1/n)$ as $n \rightarrow \infty$, and hence the sum over n is infinite. By equation (4.4) in Lemma 4.2, we can find a sequence of deterministic points $\mathbf{z}_n^-$ such that

$$(5.10) \quad \mathbb{P} \left(\mathcal{S}(\mathbf{x}_n, \mathbf{z}_n^-) = \mathcal{B}[\mathbf{x}_n \rightarrow \mathbf{z}_n^-] = \sum_{i=1}^k \mathcal{B}[x_{n,i} \rightarrow z_{n,i}^-] \right) \rightarrow 1$$

as $n \rightarrow \infty$, and for every $i \in \llbracket 1, k \rrbracket$ as $n \rightarrow \infty$ we have $z_{n,i}^- \rightarrow -\infty$. Moreover, the two-point estimate Lemma 5.6 and Lemma 3.3 gives that

$$(5.11) \quad \sup_{|\mathbf{y} - \mathbf{x}_n| \leq n^{-1/k}} |\mathcal{R}(\mathbf{x}_n, \mathbf{z}_n^-) - \mathcal{R}(\mathbf{y}, \mathbf{z}_n^-)| \leq C_n n^{-1/2k}$$

for a sequence of constants C_n satisfying $\mathbb{P}(C_n > a) \leq ce^{-da^{3/2}}$ for constants c, d that do not depend on n . This strong tail control on C_n ensures that the right-hand side of (5.11) converges to 0 almost surely as $n \rightarrow \infty$, and hence so does the left-hand side. Similarly,

$$(5.12) \quad \lim_{n \rightarrow \infty} \sup_{|\mathbf{y} - \mathbf{x}_n| \leq n^{-1/k}} \sum_{i=1}^k |\mathcal{R}(x_{n,i}, z_{n,i}^-) - \mathcal{R}(y_i, z_{n,i}^-)| = 0 \quad \text{almost surely.}$$

Combining (5.10), the convergence of (5.11), and (5.12) with the fact that any point $\mathbf{x} \in \mathbb{R}_{\leq}^k$ satisfies $|\mathbf{x}_n - \mathbf{x}| < n^{-1/k}$ infinitely often implies that for all $\mathbf{x} \in \mathbb{R}_{\leq}^k$,

$$\limsup_{n \rightarrow \infty} \mathcal{R}(\mathbf{x}, \mathbf{z}_n^-) - \sum_{i=1}^k \mathcal{R}(x_i, z_{n,i}^-) = 0,$$

and so after removing the parabolic correction,

$$\limsup_{n \rightarrow \infty} \mathcal{S}(\mathbf{x}, \mathbf{z}_n^-) - \sum_{i=1}^k \mathcal{S}(x_i, z_{n,i}^-) \geq 0.$$

Now, $\sum_{i=1}^k \mathcal{S}(x_i, z_{n,i}^-) \geq \mathcal{B}[\mathbf{x} \rightarrow \mathbf{z}_n^-]$ by the definition of parabolic path weight, yielding (5.9). $\square$

The relationship between $\mathcal{S}$ and $\mathcal{B}$ is particularly tractable when the start point $\mathbf{x} = 0^k$. This proposition immediately gives the relationship (1.8).

PROPOSITION 5.9. *Almost surely the following holds. For any $k \in \mathbb{N}$ and $\mathbf{y} \in \mathbb{R}_{\leq}^k$ we have*

$$(5.13) \quad \mathcal{S}(0^k, \mathbf{y}) = \sum_{i=1}^k \mathcal{B}_i(y_i) + \mathcal{B}[(y_1^k, k) \rightarrow (\mathbf{y}, 1)].$$

Moreover, there is a disjoint optimizer π in $\mathcal{B}$ from 0^k to $\mathbf{y}$ that only uses the top k lines.

PROOF. Equation (5.13) is true in the prelimit by Lemma 2.12, and hence holds in the limit as well. The ‘Moreover’ claim follows by an explicit construction. Let $\pi = (\pi_1, \dots, \pi_k)$ be given by $\pi_i|_{(-\infty, y_1)} = i$ and $\pi_i|_{[y_1, y_i]} = \tau_i$, where τ is a disjoint optimizer from (y_1^k, k) to $(\mathbf{y}, 1)$. We claim that $\|\pi\|_{\mathcal{B}}$ is equal to the right-hand side of (5.13). The result will then follow from Proposition 5.8.

By (4.2) and the fact that $\mathcal{S}(0, y_i) = \mathcal{B}_1(y_i)$ for all i , it is enough to show that

$$(5.14) \quad \mathcal{B}[(z, i) \rightarrow (y_i, 1)] - [\mathcal{B}_1(y_i) - \mathcal{B}_i(z)] \rightarrow 0$$

as $z \rightarrow -\infty$. For any interval $[n, n+1] \subset [z, y_i]$, the left-hand side above is always bounded between $I_n := \sup_{x \in [n, n+1]} \mathcal{B}_i(x) - \mathcal{B}_1(x)$ and 0. This is because the left-hand side of (5.14) can be written as

$$\sup_{z \leq z_{i-1} \leq \dots \leq z_1 \leq y_i} \sum_{j=1}^{i-1} \mathcal{B}_{j+1}(z_j) - \mathcal{B}_j(z_j).$$

The support of I_n contains 0 by the Brownian Gibbs property (Theorem 2.14), and I_n is a stationary, ergodic process by the main result of [11]. Hence $\liminf_{n \rightarrow -\infty} I_n = 0$, yielding (5.14). $\square$

5.3. Metric composition law

To construct the full scaling limit of multi-point Brownian LPP from the extended Airy sheet, a key property is the following metric composition law. Recall from the introduction that if $\mathcal{S}$ is an extended Airy sheet, then $s\mathcal{S}(s^{-2}\mathbf{x}, s^{-2}\mathbf{y})$ is an **extended Airy sheet of scale s** .

PROPOSITION 5.10. *For $s_1, s_2 > 0$, take independent extended Airy sheets $\mathcal{S}_1, \mathcal{S}_2$ of scale s_1, s_2 , respectively. Then almost surely, for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$ the maximum*

$$(5.15) \quad \mathcal{S}(\mathbf{x}, \mathbf{y}) = \max_{\mathbf{z} \in \mathbb{R}_{\leq}^k} \mathcal{S}_1(\mathbf{x}, \mathbf{z}) + \mathcal{S}_2(\mathbf{z}, \mathbf{y})$$

exists. Moreover, $\mathcal{S}$ is an extended Airy sheet of scale $(s_1^3 + s_2^3)^{1/3}$.

PROOF. Without loss of generality, we can assume that $s_1^3 + s_2^3 = 1$. We set up multi-point Brownian LPP converging to an extended Airy sheet $\mathcal{S}$ as in Theorem 5.4. Let B^n be a collection of n independent two-sided standard Brownian motions, and let $\mathcal{S}^n$ be the prelimiting extended Airy sheet defined using last passage

percolation across B^n (as in Theorem 1.2) so that $\mathcal{S}^n \xrightarrow{d} \mathcal{S}$ as $n \rightarrow \infty$. We now let

$$\mathcal{S}_1^n(\mathbf{x}, \mathbf{y}) = n^{1/6} \times \left(B[(2n^{-1/3}\mathbf{x}, n) \rightarrow (s_1^3 + 2n^{-1/3}\mathbf{y}, n - \lfloor s_1^3 n \rfloor)] - 2ks_1^3\sqrt{n} - n^{1/6} \sum_{i=1}^k 2(y_i - x_i) \right),$$

and

$$\mathcal{S}_2^n(\mathbf{x}, \mathbf{y}) = n^{1/6} \times \left(B[(s_1^3 + 2n^{-1/3}\mathbf{x}, n - \lfloor s_1^3 n \rfloor - 1) \rightarrow (1 + 2n^{-1/3}\mathbf{y}, 1)] - 2ks_2^3\sqrt{n} - n^{1/6} \sum_{i=1}^k 2(y_i - x_i) \right)$$

for any $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}_{\leq}^k \times \mathbb{R}_{\leq}^k$. Then $\mathcal{S}_1^n$ and $\mathcal{S}_2^n$ are independent. Using Theorem 5.4, Brownian scaling, and continuity of $\mathcal{S}$, we have $\mathcal{S}_1^n \xrightarrow{d} \mathcal{S}_1$ and $\mathcal{S}_2^n \xrightarrow{d} \mathcal{S}_2$ as $n \rightarrow \infty$. On the other hand, by Lemma 2.9 we have

$$\mathcal{S}^n(\mathbf{x}, \mathbf{y}) = \max_{\mathbf{z} \in \mathbb{R}_{\leq}^k} \mathcal{S}_1^n(\mathbf{x}, \mathbf{z}) + \mathcal{S}_2^n(\mathbf{z}, \mathbf{y})$$

for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$. This passes to the limit as long as the argmax for the right-hand side above is tight. This tightness for a fixed $\mathbf{x}, \mathbf{y}$ follows from Lemma 5.11 below. Uniform tightness when $(\mathbf{x}, \mathbf{y})$ are allowed to range over a compact subset of $\mathfrak{X}$ then follows from monotonicity of optimizers (Lemma 2.3). $\square$

LEMMA 5.11. *For any $k \in \mathbb{N}$, there exist constants $c, d > 0$ such that the following is true. Let $n, p, q \in \mathbb{N}$, with $p+q = n$, and denote $t = p/n$. Take independent standard Brownian motions $B^n = (B_1^n, \dots, B_n^n)$, and $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$ such that $\|\mathbf{x}\|_2, \|\mathbf{y}\|_2 < n^{1/6}$. For any $\mathbf{z} \in \mathbb{R}_{\leq}^k$ define*

$$(5.16) \quad A(\mathbf{z}) = B^n[(2n^{-1/3}\mathbf{x}, n) \rightarrow (t+2n^{-1/3}\mathbf{z}, q+1)] + B^n[(t+2n^{-1/3}\mathbf{z}, q) \rightarrow (1+2n^{-1/3}\mathbf{y}, 1)].$$

We set $A(\mathbf{z}) = -\infty$ if the right-hand side is not defined. Then for any $a > 0$, with probability at least $1 - ce^{-dr^{3/2}}$ the following is true: for any $\mathbf{z}^$ where A achieves its maximum, we must have $\|\mathbf{z}^* - t\mathbf{y} - (1-t)\mathbf{x}\|_2 < ca^2(t \wedge (1-t))^{1/3}$.*

This lemma is an analogue of [15, Lemma 9.3]. Its proof is also similar to the proof of [15, Lemma 9.3], involving some technical estimates on the Brownian n -melon W^n . We leave it to Appendix A.1.

CHAPTER 6

The Scaling Limit of Multi-point Brownian LPP

6.1. Tightness of the prelimiting extended landscape

Recall from the introduction that

$$\mathfrak{X}_\uparrow = \{(\mathbf{x}, s; \mathbf{y}, t) \in \bigcup_{k \in \mathbb{N}} (\mathbb{R}_{\leq}^k \times \mathbb{R})^2 : s < t\}.$$

Let $B = (B_i)_{i \in \mathbb{Z}}$ be an infinite sequence of independent two-sided standard Brownian motions. As in the introduction, let $(\mathbf{x}, s)_n = (s + 2\mathbf{x}n^{-1/3}, -\lfloor sn \rfloor)$, and define the prelimiting extended landscape

$$\mathcal{L}_n(\mathbf{x}, s; \mathbf{y}, t) = n^{1/6} \left(B[(\mathbf{x}, s)_n \rightarrow (\mathbf{y}, t)_n] - 2k(t-s)\sqrt{n} - n^{1/6} \sum_{i=1}^k 2(y_i - x_i) \right).$$

This is a random function on $\mathfrak{X}_\uparrow$. In this section we prove that $\mathcal{L}_n$ is tight in an appropriate function space. Given Lemma 3.5, it remains to prove a two-point tail bound on the deviation of $\mathcal{L}_n$ in the time direction. This is the analogue of [15, Lemma 11.2]. Let $\mathcal{K}_n$ be the stationary version of $\mathcal{L}_n$, defined as

$$\mathcal{K}_n(\mathbf{x}, s; \mathbf{y}, t) = \mathcal{L}_n(\mathbf{x}, s; \mathbf{y}, t) + \sum_{i=1}^k \frac{(x_i - y_i)^2}{t - s}.$$

LEMMA 6.1. *Take any $k \in \mathbb{N}$, $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$, and $t \in n^{-1}\mathbb{Z}$, such that $\|\mathbf{x}\|_2, \|\mathbf{y}\|_2 < n^{1/100}$, $1/2 \leq t < 1 - n^{-1/100}$. Also take $0 < a < n^{1/150}$. Letting $\mathbf{y}' = t\mathbf{x} + (1-t)\mathbf{y}$, we have*

$$\mathbb{P}(|\mathcal{K}_n(\mathbf{x}, 0; \mathbf{y}', t) - \mathcal{K}_n(\mathbf{x}, 0; \mathbf{y}, 1)| > a(1-t)^{1/3} |\log(1-t)|) < ce^{-da^{9/8}},$$

for some constants c, d depending only on k .

We sketch the idea of the proof of this lemma here. The complete proof is a technical computation and uses similar ideas to the proof of [15, Lemma 11.2], so we leave it to Appendix A.1 (with the proof of Lemma 5.11). For the upper tail of $\mathcal{K}_n(\mathbf{x}, 0; \mathbf{y}', t) - \mathcal{K}_n(\mathbf{x}, 0; \mathbf{y}, 1)$, via the triangle inequality it suffices to give a lower bound on $\mathcal{K}_n(\mathbf{y}', t + n^{-1}; \mathbf{y}, 1)$. This follows from Lemma 2.8 and tail bounds on points the Brownian melon, see Lemma A.4. For the lower tail, by the metric composition law we need to upper bound

$$\sup_{\mathbf{z} \in \mathbb{R}_{\leq}^k} (\mathcal{L}_n(\mathbf{x}, 0; \mathbf{z}, t) - \mathcal{L}_n(\mathbf{x}, 0; \mathbf{y}', t)) + \mathcal{L}_n(\mathbf{z}, t; \mathbf{y}, 1).$$

The term $\mathcal{L}_n(\mathbf{z}, t; \mathbf{y}, 1)$ can be bounded with a curvature estimate on the Brownian melon. When $\|\mathbf{z} - \mathbf{y}'\|_2$ is large, such a curvature estimate also works to bound $\mathcal{L}_n(\mathbf{x}, 0; \mathbf{z}, t)$, and the aforementioned Lemma A.4 can be used to bound

$\mathcal{L}_n(\mathbf{x}, 0; \mathbf{y}', t)$ below. When $\|\mathbf{z} - \mathbf{y}'\|_2$ is small we apply the more refined spatial continuity estimate on the prelimiting extended Airy sheets from Lemma 3.5 to bound the difference $\mathcal{L}_n(\mathbf{x}, 0; \mathbf{z}, t) - \mathcal{L}_n(\mathbf{x}, 0; \mathbf{y}', t)$. Putting together these bounds gives the desired result.

We now move to tightness. Let $\mathfrak{F}$ be the space of functions from $\mathfrak{X}_\uparrow$ to $\mathbb{R}$ that are either continuous, or of the form (6.1) for some n and some bi-infinite sequence of continuous functions f in place of B . This is a Polish space, and so all classical theorems about distributional convergence apply. All of the $\mathcal{L}_n$ are random functions on this space.

PROPOSITION 6.2. *The functions $\mathcal{L}_n$ are tight in $\mathfrak{F}$, and all subsequential limits are almost surely continuous.*

PROOF. Fix a compact set $K \subset (\mathbb{R}_{\leq}^k \times \mathbb{R})^2$ for some $k \in \mathbb{N}$. It suffices to show tightness of $\mathcal{L}_n|_K$. First, we replace $\mathcal{L}_n$ by a continuous version $\mathcal{J}_n$ on K . For each $(\mathbf{x}, \mathbf{y}) \in \mathfrak{X}$ and s with $s \in n^{-1}\mathbb{Z}$, define the function $\mathcal{J}_n(\mathbf{x}, s; \mathbf{y}, \cdot)$ by setting $\mathcal{J}_n(\mathbf{x}, s; \mathbf{y}, t) = \mathcal{L}_n(\mathbf{x}, s; \mathbf{y}, t)$ whenever $t \in n^{-1}\mathbb{Z}$ and by linear interpolation at times in between. Then for each $(\mathbf{x}, \mathbf{y}) \in \mathfrak{X}$ and $t \in \mathbb{R}$, we can define $\mathcal{J}_n(\mathbf{x}, s; \mathbf{y}, t)$ by linear interpolation between values when $s \in n^{-1}\mathbb{Z}$. This procedure gives a well-defined continuous function on K for large enough n . By Theorem 2.13, $\mathcal{J}_n(0^k, 0; 0^k, 1)$ is tight in n . Moreover, by Lemma 3.5, Lemma 6.1, and translation and scale invariance properties of $\mathcal{L}_n$ we get that for all $\mathbf{u}, \mathbf{u}' \in K$ and large enough n ,

$$\mathbb{P}(|\mathcal{J}_n(\mathbf{u}) - \mathcal{J}_n(\mathbf{u}')| > a \|\mathbf{u} - \mathbf{u}'\|_2^{1/3-\epsilon}) \leq ce^{-da^{9/8}}$$

for any $a, \epsilon > 0$. Here $c, d > 0$ are K -dependent constants. Using the Kolmogorov-Chentsov criterion, see [31, Corollary 14.9], we get that the sequence $\mathcal{L}_n$ is tight. $\square$

6.2. The explicit construction of $\mathcal{L}^*$

In this section, we construct the scaling limit $\mathcal{L}^*$ of multipoint Brownian LPP axiomatically and prove Theorems 1.4 and 1.5. We call this object an extended* directed landscape, or extended* landscape for brevity. Later we will show that this object coincides with the extended directed landscape as defined in Definition 1.1.

DEFINITION 6.3. An **extended* directed landscape** is a random continuous function $\mathcal{L}^*$ taking values in the space $\mathcal{C}(\mathfrak{X}_\uparrow, \mathbb{R}) \subset \mathfrak{F}$ of continuous functions from $\mathfrak{X}_\uparrow$ to $\mathbb{R}$ with the uniform-on-compact topology. It satisfies the following properties.

- I. (Independent extended Airy sheet marginals) For any disjoint time intervals $\{(s_i, t_i) : i \in \{1, \dots, k\}\}$, the random functions

$$(\mathbf{x}, \mathbf{y}) \mapsto \mathcal{L}^*(\mathbf{x}, s_i; \mathbf{y}, t_i), \quad i \in [1, k]$$

are independent extended Airy sheets of scale $(t_i - s_i)^{1/3}$.

- II. (Metric composition law) For any $r < s < t$, almost surely we have that

$$\mathcal{L}^*(\mathbf{x}, r; \mathbf{y}, t) = \max_{\mathbf{z} \in \mathbb{R}_{\leq}^k} \mathcal{L}^*(\mathbf{x}, r; \mathbf{z}, s) + \mathcal{L}^*(\mathbf{z}, s; \mathbf{y}, t),$$

for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$.

Note that $\mathcal{L}^*|_{\mathbb{R}_\uparrow^4}$ is the usual directed landscape, since extended Airy sheets are simply Airy sheets when restricted to $\mathbb{R}^2$.

While $\mathcal{L}^*$ can be constructed directly similarly to how the directed landscape was constructed in [15, Section 10], we will instead show its existence by proving that it is the scaling limit of $\mathcal{L}_n$. The next result encompasses Theorems 1.4 and 1.5.

THEOREM 6.4. *The extended* landscape $\mathcal{L}^*$ exists and is unique in law. Moreover, $\mathcal{L}_n \xrightarrow{d} \mathcal{L}^*$ as random functions in $\mathfrak{F}$.*

PROOF. The uniqueness of $\mathcal{L}^*$ follows since conditions I and II specify all finite dimensional distributions. Indeed, let $\mathbf{u}_1, \dots, \mathbf{u}_k \in \mathfrak{X}_\uparrow$ be any collection of points with time indices $S = \{s_i < t_i : i \in \llbracket 1, k \rrbracket\}$. Let $r_1 < \dots < r_\ell$ denote the order statistics of the set S , for some $2 \leq \ell \leq 2k$. Then the marginals $\mathcal{L}(\cdot, r_i; \cdot, r_{i+1}), i \in \llbracket 1, \ell - 1 \rrbracket$ are independent Airy sheets of scale $(r_{i+1} - r_i)^{1/3}$ by I. All the random variables $\mathcal{L}(\mathbf{u}_1), \dots, \mathcal{L}(\mathbf{u}_k)$ are measurable functions of $\mathcal{L}(\cdot, r_i; \cdot, r_{i+1}), i \in \llbracket 1, \ell - 1 \rrbracket$ by repeated applications of II.

Next, we know $\mathcal{L}_n$ is tight in $\mathcal{C}(\mathfrak{X}_\uparrow, \mathbb{R})$ by Proposition 6.2. Let $\mathcal{M} : \mathfrak{X} \rightarrow \mathbb{R}_\uparrow^4$ be any subsequential limit of $\mathcal{L}_n$. The function $\mathcal{M}$ has independent increments by the independence of the Brownian motions that give rise to $\mathcal{L}_n$. These increments must be rescaled extended Airy sheets by Theorem 5.4, and satisfy metric composition since $\mathcal{L}_n$ does, and maximizer locations are tight (Lemma 5.11). Therefore $\mathcal{M}$ is an extended* landscape. $\square$

In the remainder of this section, we gather continuity estimates for $\mathcal{L}^*$. Let $\mathcal{K}$ be the stationary extended landscape, defined as

$$\mathcal{K}(\mathbf{x}, s; \mathbf{y}, t) = \mathcal{L}^*(\mathbf{x}, s; \mathbf{y}, t) + \sum_{i=1}^k \frac{(x_i - y_i)^2}{t - s}.$$

By passing Lemma 6.1 to the limit, we have the following two-point bound on $\mathcal{K}$ (and hence on $\mathcal{L}$) in the time direction. Note that a two-point bound in the spatial direction follows from Lemma 5.6 and rescaling.

LEMMA 6.5. *Take any $k \in \mathbb{N}$, $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$, and $0 < t' < t$ with $2t' \geq t$. Letting $\mathbf{y}' = (t'/t)\mathbf{x} + (1 - t'/t)\mathbf{y}$, we have*

$$\mathbb{P}(|\mathcal{K}(\mathbf{x}, 0; \mathbf{y}', t') - \mathcal{K}(\mathbf{x}, 0; \mathbf{y}, t)| > a(t - t')^{1/3} |\log(1 - t'/t)|) < ce^{-da^{9/8}},$$

for all $a > 0$. Here $c, d > 0$ are constants depending only on k .

By Lemmas 5.6 and 6.5, and using Lemma 3.3, we have that in any compact subset of $\mathfrak{X}_\uparrow$, the function $\mathcal{L}$ is $(1/2 - \epsilon)$ -Hölder in the spatial coordinate and $(1/3 - \epsilon)$ -Hölder in the time coordinate, for any $\epsilon > 0$.

We also need uniform upper and lower bounds on $\mathcal{K}$ on $\mathfrak{X}_\uparrow$. We first give a one-point bound. This is obtained from passing the bound Lemma A.4 on Brownian last passage values to the limit.

LEMMA 6.6. *For any $k \in \mathbb{N}$, $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$, and $t > 0$, we have*

$$\mathbb{P}(|\mathcal{K}(\mathbf{x}, 0; \mathbf{y}, t)| > at^{1/3}) < ce^{-da^{3/2}}$$

for all $a > 0$. Here $c, d > 0$ are constants depending only on k .

Next we use Lemma 5.6 and Lemma 6.5 to upgrade Lemma 6.6 to a uniform bound that will be sufficient for our purposes.

LEMMA 6.7. *For any $\eta > 0$ and $k \in \mathbb{N}$, there is a random constant $R > 1$, such that for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$ and $s < t$, we have*

$$|\mathcal{K}(\mathbf{x}, s; \mathbf{y}, t)| < RG(\mathbf{x}, \mathbf{y}, s, t)^\eta (t-s)^{1/3}$$

where

$$G(\mathbf{x}, \mathbf{y}, s, t) = \left(1 + \frac{\|\mathbf{x}\|_1 + \|\mathbf{y}\|_1}{(t-s)^{2/3}}\right) \left(1 + \frac{|s|}{t-s}\right) (1 + |\log(t-s)|).$$

Also $\mathbb{P}(R > a) < ce^{-da}$ for any $a > 0$. Here $c, d > 0$ are constants depending on k, η .

PROOF. Fix $\eta > 0, k \in \mathbb{N}$. Throughout this proof we let c, d be constants depending on k, η , whose values can vary from line to line. For each $\ell \in \mathbb{Z}$, let $L_\ell \subset \mathbb{R}_{\leq}^k \times \mathbb{R}$ consist of all $(\mathbf{x}, s)$, where each coordinate of $\mathbf{x}$ is in $2^{2\ell}\mathbb{Z}$ and $s \in 2^{3\ell}\mathbb{Z}$. For any $(\mathbf{x}, s), (\mathbf{y}, t) \in L_\ell$, denote

$$F(\mathbf{x}, \mathbf{y}, s, t, \ell) = (1 + 2^{-2\ell}(\|\mathbf{x}\|_1 + \|\mathbf{y}\|_1))(1 + 2^{-3\ell}(|s| + |t|))(1 + |\ell|).$$

Take any $(\mathbf{x}, s), (\mathbf{y}, t), (\mathbf{y}', t') \in L_\ell$ with $s < t$, such that $\mathbf{y}'$ and $\mathbf{y}$ differ at exactly one coordinate, and by exactly $2^{2\ell}$. By Lemma 5.6 we have

$$(6.2) \quad \mathbb{P}(|\mathcal{K}(\mathbf{x}, s; \mathbf{y}, t) - \mathcal{K}(\mathbf{x}, s; \mathbf{y}', t')| > aF(\mathbf{x}, \mathbf{y}, s, t, \ell)^\eta 2^\ell) < ce^{-da^{3/2}F(\mathbf{x}, \mathbf{y}, s, t, \ell)^{3\eta/2}}.$$

We then consider $(\mathbf{x}, s), (\mathbf{y}, t), (\mathbf{y}', t') \in L_\ell$, such that $s < t', t = t' + 2^{3\ell}$, and so that $\mathbf{x}, \mathbf{y}, \mathbf{y}'$ satisfy the bound $|y'_i - ((t-t')x_i + (t'-s)y_i)/(t-s)| \leq 2^{2\ell}$ for $1 \leq i \leq k$. By Lemmas 5.6 and 6.5 we have

$$(6.3) \quad \mathbb{P}(|\mathcal{K}(\mathbf{x}, s; \mathbf{y}, t) - \mathcal{K}(\mathbf{x}, s; \mathbf{y}', t')| > aF(\mathbf{x}, \mathbf{y}, s, t, \ell)^\eta 2^\ell |\log(2^{3\ell}(t-s)^{-1})|) < ce^{-da^{9/8}F(\mathbf{x}, \mathbf{y}, s, t, \ell)^{9\eta/8}}.$$

We next consider any $(\mathbf{x}, s), (\mathbf{y}, t) \in L_\ell$ with $t-s = 2^{3\ell}$. By Lemma 6.6 we have

$$(6.4) \quad \mathbb{P}(|\mathcal{K}(\mathbf{x}, s; \mathbf{y}, t)| > aF(\mathbf{x}, \mathbf{y}, s, t, \ell)^\eta 2^\ell) < ce^{-da^{3/2}F(\mathbf{x}, \mathbf{y}, s, t, \ell)^{3\eta/2}}.$$

The right-hand sides of (6.2), (6.3), and (6.4) are summable over all allowable $\mathbf{x}, \mathbf{y}, s, t$ and ℓ with sums that decrease at least exponentially in a . In other words, we conclude that there exists a random number R , such that $\mathbb{P}(R > a) < ce^{-da}$ and the following is true.

- (1) For any $(\mathbf{x}, s), (\mathbf{y}, t), (\mathbf{x}', s), (\mathbf{y}', t) \in L_\ell$ with $s < t$, such that $\mathbf{x}'$, $\mathbf{x}$ differ at exactly one coordinate by $2^{2\ell}$ and $\mathbf{y}'$, $\mathbf{y}$ differ at exactly one coordinate by $2^{2\ell}$, we have

$$|\mathcal{K}(\mathbf{x}, s; \mathbf{y}, t) - \mathcal{K}(\mathbf{x}, s; \mathbf{y}', t)| < RF(\mathbf{x}, \mathbf{y}, s, t, \ell)^\eta 2^\ell,$$

$$|\mathcal{K}(\mathbf{x}', s; \mathbf{y}, t) - \mathcal{K}(\mathbf{x}, s; \mathbf{y}, t)| < RF(\mathbf{x}, \mathbf{y}, s, t, \ell)^\eta 2^\ell.$$

The second bound follows by a symmetric analogue of (6.2) where we vary $\mathbf{x}$ rather than $\mathbf{y}$.

- (2) For any $(\mathbf{x}, s), (\mathbf{y}, t), (\mathbf{y}', t') \in L_\ell$, such that $s < t', t = t' + 2^{3\ell}$, and $\mathbf{x}, \mathbf{y}, \mathbf{y}'$ satisfy the bound $|y'_i - ((t-t')x_i + (t'-s)y_i)/(t-s)| \leq 2^{2\ell}$ for $1 \leq i \leq k$, we have

$$|\mathcal{K}(\mathbf{x}, s; \mathbf{y}, t) - \mathcal{K}(\mathbf{x}, s; \mathbf{y}', t')| < RF(\mathbf{x}, \mathbf{y}, s, t, \ell)^\eta 2^\ell |\log(2^{3\ell}(t-s)^{-1})|.$$

- (3) For any $(\mathbf{x}, s), (\mathbf{y}, t) \in L_\ell$ with $t-s = 2^{3\ell}$, we have

$$|\mathcal{K}(\mathbf{x}, s; \mathbf{y}, t)| < RF(\mathbf{x}, \mathbf{y}, s, t, \ell)^\eta 2^\ell.$$

Now consider any $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$ and $s < t$ and let $\ell_0 = \lfloor \log_s(t-s) \rfloor - 1$. For each $\ell \leq \ell_0$, let $s_\ell = 2^{3\ell} \lceil 2^{-3\ell} s \rceil$, $t_\ell = 2^{3\ell} \lfloor 2^{-3\ell} t \rfloor$, and let $\mathbf{x}^{(\ell)}, \mathbf{y}^{(\ell)} \in \mathbb{R}_{\leq}^k$ be chosen such that $(\mathbf{x}^{(\ell)}, s_\ell), (\mathbf{y}^{(\ell)}, t_\ell) \in L_\ell$ and

$$(6.5) \quad |x_i^{(\ell)} - ((t-s_\ell)x_i + (s_\ell-s)y_i)/(t-s)| \leq 2^{2\ell}, \quad |y_i^{(\ell)} - ((t-t_\ell)x_i + (t_\ell-s)y_i)/(t-s)| \leq 2^{2\ell}$$

for each $1 \leq i \leq k$. As $\ell \rightarrow -\infty$ we have $(\mathbf{x}^{(\ell)}, \mathbf{y}^{(\ell)}, s_\ell, t_\ell) \rightarrow (\mathbf{x}, \mathbf{y}, s, t)$. Therefore by the continuity of $\mathcal{K}$ and the triangle inequality we have

$$(6.6) \quad \begin{aligned} |\mathcal{K}(\mathbf{x}, s; \mathbf{y}, t)| &\leq |\mathcal{K}(\mathbf{x}^{(\ell_0)}, s_{\ell_0}; \mathbf{y}^{(\ell_0)}, t_{\ell_0})| \\ &+ \sum_{\ell=\ell_0}^{-\infty} |\mathcal{K}(\mathbf{x}^{(\ell)}, s_\ell; \mathbf{y}^{(\ell)}, t_\ell) - \mathcal{K}(\mathbf{x}^{(\ell)}, s_\ell; \mathbf{y}^{(\ell-1)}, t_{\ell-1})| \\ &+ |\mathcal{K}(\mathbf{x}^{(\ell)}, s_\ell; \mathbf{y}^{(\ell-1)}, t_{\ell-1}) - \mathcal{K}(\mathbf{x}^{(\ell-1)}, s_{\ell-1}; \mathbf{y}^{(\ell-1)}, t_{\ell-1})|. \end{aligned}$$

Using the above bounds, we have

$$(6.7) \quad |\mathcal{K}(\mathbf{x}^{(\ell_0)}, s_{\ell_0}; \mathbf{y}^{(\ell_0)}, t_{\ell_0})| < RF(\mathbf{x}^{(\ell_0)}, \mathbf{y}^{(\ell_0)}, s_{\ell_0}, t_{\ell_0}, \ell_0)^\eta 2^{\ell_0} < cRG(\mathbf{x}, \mathbf{y}, s, t)^\eta (t-s)^{1/3},$$

where the last inequality is by the fact that $2^{-3\ell_0}(t-s)$ and $2^{-3\ell_0}(t_{\ell_0}-s_{\ell_0})$ are upper and lower bounded by constants, and (6.5). For each $\ell \leq \ell_0$, we can find a sequence $t_\ell = t_{\ell,1} \geq \dots \geq t_{\ell,m} = t_{\ell-1}$, and $\mathbf{y}^{(\ell)} = \mathbf{y}^{(\ell,1)}, \dots, \mathbf{y}^{(\ell,m)} = \mathbf{y}^{(\ell-1)} \in \mathbb{R}_{\leq}^k$ for some $m \leq c$, such that for each $1 \leq j < m$, one of the following two events happens:

- (1) $t_{\ell,j} = t_{\ell,j+1} + 2^{3\ell}$, and $|y_i^{(\ell,j+1)} - ((t_{\ell,j}-t_{\ell,j+1})x_i^{(\ell)} + (t_{\ell,j+1}-s_\ell)y_i^{(\ell,j)})/(t_{\ell,j}-s_\ell)| \leq 2^{2\ell}$ for each $1 \leq i \leq k$.
- (2) $t_{\ell,j} = t_{\ell,j+1}$ and $\mathbf{y}^{(\ell,j)}$ differ from $\mathbf{y}^{(\ell,j+1)}$ at exactly one coordinate by $2^{2\ell}$.

Thus we have

$$\begin{aligned} &|\mathcal{K}(\mathbf{x}^{(\ell)}, s_\ell; \mathbf{y}^{(\ell)}, t_\ell) - \mathcal{K}(\mathbf{x}^{(\ell)}, s_\ell; \mathbf{y}^{(\ell-1)}, t_{\ell-1})| \\ &< \sum_{j=1}^{m-1} RF(\mathbf{x}^{(\ell,j)}, \mathbf{y}^{(\ell,j)}, s_\ell, t_{\ell,j}, \ell)^\eta 2^{\ell} |\log(2^{3\ell}(t-s)^{-1})| \\ &< cRG(\mathbf{x}, \mathbf{y}, s, t)^\eta \left(\frac{1+|\ell|}{1+|\ell_0|} \right)^\eta 2^\ell (1+\ell_0-\ell). \end{aligned}$$

For $|\mathcal{K}(\mathbf{x}^{(\ell)}, s_\ell; \mathbf{y}^{(\ell-1)}, t_{\ell-1}) - \mathcal{K}(\mathbf{x}^{(\ell-1)}, s_{\ell-1}; \mathbf{y}^{(\ell-1)}, t_{\ell-1})|$, arguing similarly we get that the same bound holds. Combining these bounds, summed over all $\ell \leq \ell_0$, with the bound (6.7) and the triangle inequality (6.6), gives the result. $\square$

As a consequence of Lemma 6.7, we can estimate the location of the maximizer in the metric composition law for the extended landscape.

LEMMA 6.8. *For any small enough $\eta > 0$ and $k \in \mathbb{N}$, take the random variable $R > 1$ and the function G from Lemma 6.7. For any $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathbb{R}_{\leq}^k$ and $r < s < t$, if $\mathcal{L}^*(\mathbf{x}, r; \mathbf{z}, s) + \mathcal{L}^*(\mathbf{z}, s; \mathbf{y}, t) = \mathcal{L}^*(\mathbf{x}, r; \mathbf{y}, t)$, then*

$$\|\mathbf{z} - \tilde{\mathbf{z}}\|_2 < cRG(\mathbf{x}, \mathbf{y}, r, t)^\eta (t-s)^{1/3} (s-r)^{1/3},$$

where $\tilde{\mathbf{z}} = ((t-s)\mathbf{x} + (s-r)\mathbf{y})/(t-r)$ and c is a constant depending only on η, k .

PROOF. In this proof we let c denote a large constant depending on k, η , whose value may change from line to line. By Lemma 6.7 we have

$$\begin{aligned}
 (6.8) \quad 0 &= \mathcal{L}^*(\mathbf{x}, r; \mathbf{z}, s) + \mathcal{L}^*(\mathbf{z}, s; \mathbf{y}, t) - \mathcal{L}^*(\mathbf{x}, r; \mathbf{y}, t) \\
 &< \frac{\|\mathbf{x} - \mathbf{y}\|_2^2}{t-r} - \frac{\|\mathbf{z} - \mathbf{x}\|_2^2}{s-r} - \frac{\|\mathbf{z} - \mathbf{y}\|_2^2}{t-s} \\
 &\quad + RG(\mathbf{x}, \mathbf{z}, r, s)^\eta (s-r)^{1/3} + RG(\mathbf{z}, \mathbf{y}, s, t)^\eta (t-s)^{1/3} \\
 &\quad + RG(\mathbf{x}, \mathbf{y}, r, t)^\eta (t-r)^{1/3} \\
 &= -\frac{(t-r)\|\mathbf{z} - \tilde{\mathbf{z}}\|_2^2}{(t-s)(s-r)} + R(G(\mathbf{x}, \mathbf{z}, r, s)^\eta (s-r)^{1/3} \\
 &\quad + G(\mathbf{z}, \mathbf{y}, s, t)^\eta (t-s)^{1/3} + G(\mathbf{x}, \mathbf{y}, r, t)^\eta (t-r)^{1/3}).
 \end{aligned}$$

Now, $\frac{1+|\log(s-r)|}{1+|\log(t-r)|} \leq \frac{t-r}{s-r}$, and $\|\mathbf{x}\|_1 + \|\mathbf{y}\|_1 \geq \frac{s-r}{t-r}(\|\mathbf{x}\|_1 + \|\tilde{\mathbf{z}}\|_1)$, so from the definition of G we have

$$\begin{aligned}
 \frac{G(\mathbf{x}, \mathbf{z}, r, s)}{G(\mathbf{x}, \mathbf{y}, r, t)} &= \frac{1 + \frac{\|\mathbf{x}\|_1 + \|\mathbf{z}\|_1}{(s-r)^{2/3}}}{1 + \frac{\|\mathbf{x}\|_1 + \|\mathbf{y}\|_1}{(t-r)^{2/3}}} \times \frac{1 + \frac{|r|}{s-r}}{1 + \frac{|r|}{t-r}} \times \frac{1 + |\log(s-r)|}{1 + |\log(t-r)|} \\
 &\leq \frac{1 + \frac{\|\mathbf{x}\|_1 + \|\mathbf{z}\|_1}{(t-r)^{2/3}}}{1 + \frac{\|\mathbf{x}\|_1 + \|\tilde{\mathbf{z}}\|_1}{(t-r)^{2/3}}} \left(\frac{t-r}{s-r}\right)^{5/3} \times \left(\frac{t-r}{s-r}\right) \times \left(\frac{t-r}{s-r}\right) \\
 &\leq \left(1 + \left|\frac{\|\mathbf{z}\|_1 - \|\tilde{\mathbf{z}}\|_1}{(t-r)^{2/3}}\right|\right) \left(\frac{t-r}{s-r}\right)^{11/3}.
 \end{aligned}$$

Thus we have

$$\frac{G(\mathbf{x}, \mathbf{z}, r, s)}{G(\mathbf{x}, \mathbf{y}, r, t)} < c(1 + \|\mathbf{z} - \tilde{\mathbf{z}}\|_2(t-r)^{-2/3}) \left(\frac{t-r}{s-r}\right)^{10},$$

and similarly

$$\frac{G(\mathbf{z}, \mathbf{y}, s, t)}{G(\mathbf{x}, \mathbf{y}, r, t)} < c(1 + \|\mathbf{z} - \tilde{\mathbf{z}}\|_2(t-r)^{-2/3}) \left(\frac{t-r}{t-s}\right)^{10}.$$

Without loss of generality we assume that $t-s \geq s-r$; thus $(t-r)/2 \leq t-s \leq t-r$. We plug these two estimates into the inequality (6.8). By taking $\eta < 1/30$ we have

$$(6.9) \quad \|\mathbf{z} - \tilde{\mathbf{z}}\|_2^2(s-r)^{-1} < cRG(\mathbf{x}, \mathbf{y}, r, t)^\eta (1 + \|\mathbf{z} - \tilde{\mathbf{z}}\|_2(t-r)^{-2/3})^\eta (t-r)^{1/3}.$$

Now consider the function

$$f : Z \mapsto Z^2(s-r)^{-1} - cRG(\mathbf{x}, \mathbf{y}, r, t)^\eta (1 + Z(t-r)^{-2/3})^\eta (t-r)^{1/3}.$$

We have that $f(0) < 0$, and on $\mathbb{R}_+$ this function first decreases then increases. From (6.9), we have that $f(\|\mathbf{z} - \tilde{\mathbf{z}}\|_2) < 0$. Also, $f(cRG(\mathbf{x}, \mathbf{y}, r, t)^\eta (t-s)^{1/3}(s-r)^{1/3}) > 0$, since

$$\begin{aligned}
 &(cRG(\mathbf{x}, \mathbf{y}, r, t)^\eta (t-s)^{1/3}(s-r)^{1/3})^2(s-r)^{-1} \\
 &= (cRG(\mathbf{x}, \mathbf{y}, r, t)^\eta)^2(t-s)^{2/3}(s-r)^{-1/3} \\
 &\geq cRG(\mathbf{x}, \mathbf{y}, r, t)^\eta (t-r)^{1/3} + (cRG(\mathbf{x}, \mathbf{y}, r, t)^\eta)^2(t-s)^{1/3}(s-r)^{1/3}(t-r)^{-1/3} \\
 &\geq cRG(\mathbf{x}, \mathbf{y}, r, t)^\eta (1 + cRG(\mathbf{x}, \mathbf{y}, r, t)^\eta (t-s)^{1/3}(s-r)^{1/3}(t-r)^{-2/3})^\eta (t-r)^{1/3},
 \end{aligned}$$

where the first inequality is by $cRG(\mathbf{x}, \mathbf{y}, r, t)^\eta \geq c$ and taking c large enough, and the second inequality is by taking $\eta < 1$. These imply the conclusion. $\square$

We can now show that metric composition holds everywhere in $\mathcal{L}^*$.

PROPOSITION 6.9. *Almost surely, for every $r < s < t$ and $(\mathbf{x}, \mathbf{y}) \in \mathfrak{X}$ we have*

$$\mathcal{L}^*(\mathbf{x}, r; \mathbf{y}, t) = \max_{\mathbf{z} \in \mathbb{R}} \mathcal{L}^*(\mathbf{x}, r; \mathbf{z}, s) + \mathcal{L}^*(\mathbf{z}, s; \mathbf{y}, t).$$

Also, almost surely we have the triangle inequality

$$\mathcal{L}^*(\mathbf{x}, r; \mathbf{y}, t) \geq \mathcal{L}^*(\mathbf{x}, r; \mathbf{z}, s) + \mathcal{L}^*(\mathbf{z}, s; \mathbf{y}, t)$$

for every $r < s < t$, $\mathbf{x}, \mathbf{z}, \mathbf{y} \in \mathbb{R}_{\leq}^k$.

PROOF. By condition II in Definition 6.3, we can ensure that almost surely, metric composition holds at all rational times $r < s < t$. The triangle inequality then holds at all rational times. This extends to all times by continuity of $\mathcal{L}^*$.

Now, let $r < s < t$ and $(\mathbf{x}, \mathbf{y}) \in \mathfrak{X}$. Consider rational sequences $r_n \rightarrow r$, $s_n \rightarrow s$, $t_n \rightarrow t$. By the metric composition law at rational times, for every n we can find $\mathbf{z}_n$ such that

$$(6.10) \quad \mathcal{L}^*(\mathbf{x}, r_n; \mathbf{y}, t_n) = \mathcal{L}^*(\mathbf{x}, r_n; \mathbf{z}_n, s_n) + \mathcal{L}^*(\mathbf{z}_n, s_n; \mathbf{y}, t_n).$$

Lemma 6.8 ensures that all the points $\mathbf{z}_n$ are contained in a common compact set, and hence we can find a subsequential limit $\mathbf{z}$. Continuity of $\mathcal{L}^*$ ensures that Equation (6.10) then holds with the n 's removed. Combining this with the triangle inequality yields the metric composition law at $r < s < t$ and $(\mathbf{x}, \mathbf{y})$. $\square$

We finish this section by recording some symmetries of $\mathcal{L}^*$.

LEMMA 6.10. *Take $q > 0$, $r, c \in \mathbb{R}$, and let $T_c \mathbf{x}$ denote the shifted vector $(x_1 + c, \dots, x_k + c)$. We have the following equalities in distribution for $\mathcal{L}^*$ as functions in $\mathfrak{F}$.*

- (1) *Stationarity:* $\mathcal{L}^*(\mathbf{x}, s, \mathbf{y}, t) \stackrel{d}{=} \mathcal{L}^*(T_c \mathbf{x}, s + r, T_c \mathbf{y}, t + r).$
- (2) *Flip symmetry:* $\mathcal{L}^*(\mathbf{x}, s, \mathbf{y}, t) \stackrel{d}{=} \mathcal{L}^*(-\mathbf{y}, -t, -\mathbf{x}, -s).$
- (3) *Rescaling:* $\mathcal{L}^*(\mathbf{x}, s, \mathbf{y}, t) \stackrel{d}{=} q \mathcal{L}^*(q^{-2} \mathbf{x}, q^{-3} s, q^{-2} \mathbf{y}, q^{-3} t).$
- (4) *Skew symmetry:*

$$\mathcal{L}^*(\mathbf{x}, s, \mathbf{y}, t) + (t - s)^{-1} \|\mathbf{x} - \mathbf{y}\|_2^2 \stackrel{d}{=} \mathcal{L}^*(\mathbf{x}, s, T_c \mathbf{y}, t) + (t - s)^{-1} \|\mathbf{x} - T_c \mathbf{y}\|_2^2.$$

PROOF. The first three symmetries of $\mathcal{L}^*$ can be deduced by the convergence from $\mathcal{L}_n$ (Theorem 6.4), since finite versions hold for $\mathcal{L}_n$. The final symmetry follows from the corresponding symmetry in Lemma 5.5 and the characterization of $\mathcal{L}^*$ in Definition 6.3. $\square$

NOT FOR DISTRIBUTION

CHAPTER 7

Paths in the Extended Landscape

Having constructed $\mathcal{L}^*$, our next goal is to understand its optimizers. In this chapter we introduce both paths and optimizers in $\mathcal{L}^*$, and prove a selection of basic properties.

7.1. Path weights

We call a continuous function $\pi : [s, t] \rightarrow \mathbb{R}_{\leq}^k$ for some interval $[s, t]$ a **multi-path of size k** . For any multi-path $\pi : [s, t] \rightarrow \mathbb{R}_{\leq}^k$, define its **length** in $\mathcal{L}^*$ by

$$\|\pi\|_{\mathcal{L}^*} = \inf_{m \in \mathbb{N}} \inf_{s=t_0 < t_1 < \dots < t_m=t} \sum_{i=1}^m \mathcal{L}^*(\pi(t_{i-1}), t_{i-1}; \pi(t_i), t_i).$$

This is the $\mathcal{L}^*$ -analogue of the formula (1.3). For any $\pi : [r, t] \rightarrow \mathbb{R}_{\leq}^k$, and a sequence $\pi^{(i)} : [r_i, t_i] \rightarrow \mathbb{R}_{\leq}^k$ for $i \in \mathbb{N}$, we say that $\pi^{(i)} \rightarrow \pi$ in the **dyadic pointwise topology**, if $r_i \rightarrow r$, $t_i \rightarrow t$, and $\pi^{(i)}(s) \rightarrow \pi(s)$ for each $s \in \mathbb{Q}_2 \cap [r, t]$, where $\mathbb{Q}_2$ is the set of dyadic rational numbers. This is a Polish topology, making it easy to work with probabilistically. Note that the length above can also be defined for discontinuous functions π . However, almost surely all discontinuous functions will have length $-\infty$ by Lemma 6.7.

LEMMA 7.1. *For a sequence of multi-paths $\{\pi^{(i)}\}_{i \in \mathbb{N}}$ and a multi-path π , such that $\pi^{(i)} \rightarrow \pi$ in the dyadic pointwise topology, we have $\limsup_{i \rightarrow \infty} \|\pi^{(i)}\|_{\mathcal{L}^*} \leq \|\pi\|_{\mathcal{L}^*}$.*

PROOF. Suppose that π is on $[r, t]$ and each $\pi^{(i)}$ is on $[r_i, t_i]$. Take any $m \in \mathbb{N}$ and any sequence $r = s_0 < s_1 < \dots < s_m = t$, such that $s_j \in \mathbb{Q}_2$ for each $0 < j < m$. For each $i \in \mathbb{N}$ and $0 < j < m$ we denote $s_{i,j} = s_j$, and $s_{i,0} = r_i$, $s_{i,m} = t_i$. By the definition of $\|\pi^{(i)}\|_{\mathcal{L}^*}$, for all i large enough so that $r_i < s_1, s_{m-1} < t_i$, we have that

$$\|\pi^{(i)}\|_{\mathcal{L}^*} \leq \sum_{j=1}^m \mathcal{L}^*(\pi^{(i)}(s_{i,j-1}), s_{i,j-1}; \pi^{(i)}(s_{i,j}), s_{i,j}).$$

As $i \rightarrow \infty$ the right-hand side converges to $\sum_{j=1}^m \mathcal{L}^*(\pi(s_{j-1}), s_{j-1}; \pi(s_j), s_j)$, by the convergence of $\pi^{(i)}$ to π in the dyadic pointwise topology and the continuity of $\mathcal{L}^*$. Therefore

$$\limsup_{i \rightarrow \infty} \|\pi^{(i)}\|_{\mathcal{L}^*} \leq \sum_{j=1}^m \mathcal{L}^*(\pi(s_{j-1}), s_{j-1}; \pi(s_j), s_j).$$

By the continuity of $\mathcal{L}^*$ and of π , this inequality holds even when the points s_i are not in $\mathbb{Q}_2$. The conclusion then follows from the definition of $\|\pi\|_{\mathcal{L}^*}$. $\square$

7.2. Optimizers and transversal fluctuation

From the definition of $\|\cdot\|_{\mathcal{L}^*}$, and the triangle inequality for $\mathcal{L}^*$ (Proposition 6.9), for any $\pi : [s, t] \rightarrow \mathbb{R}_{\leq}^k$ we have that

$$(7.1) \quad \|\pi\|_{\mathcal{L}^*} \leq \mathcal{L}^*(\pi(s), s; \pi(t), t).$$

We call a multi-path π an **optimizer in $\mathcal{L}^*$** from $(\pi(s), s)$ to $(\pi(t), t)$, if equality holds in (7.1). If π is an optimizer, then

$$\mathcal{L}^*(\pi(s), s; \pi(t), t) = \sum_{i=1}^m \mathcal{L}^*(\pi(t_{i-1}), t_{i-1}; \pi(t_i), t_i)$$

for any partition $s = t_0 < t_1 < \dots < t_m = t$ of $[s, t]$. In the case where $k = 1$, this defines a geodesic in the directed landscape, since $\mathcal{L}^*|_{\mathbb{R}_+^4} = \mathcal{L}$. We next address the existence and uniqueness of optimizers. We start with a fixed pair of endpoints.

LEMMA 7.2. *Given $(\mathbf{x}, r; \mathbf{y}, t) \in \mathfrak{X}_\uparrow$, almost surely there is a unique optimizer in $\mathcal{L}^*$ from $(\mathbf{x}, r)$ to $(\mathbf{y}, t)$.*

We need the following result on the uniqueness of the maximum in the metric composition law.

LEMMA 7.3. *Given $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$ and $r < s < t$, almost surely the function $A_s(\mathbf{z}) = \mathcal{L}^*(\mathbf{x}, r; \mathbf{z}, s) + \mathcal{L}^*(\mathbf{z}, s; \mathbf{y}, t)$ has a unique maximum.*

We leave the proof of this lemma to the end of this section, and continue our discussion of existence and uniqueness of optimizers.

PROOF OF LEMMA 7.2. By Lemma 7.3, almost surely for each rational $s \in (r, t)$, the function $A_s(\mathbf{z}) = \mathcal{L}^*(\mathbf{x}, r; \mathbf{z}, s) + \mathcal{L}^*(\mathbf{z}, s; \mathbf{y}, t)$ has a unique maximum. Therefore the value of any optimizer from $(\mathbf{x}, r)$ to $(\mathbf{y}, t)$ is uniquely determined at all rational times, and hence the optimizer itself is uniquely determined by continuity.

For existence, we can construct an optimizer π as follows. Let $\pi(r) = \mathbf{x}$, $\pi(t) = \mathbf{y}$, and for any rational $s \in (r, t)$ let $\pi(s) \in \mathbb{R}_{\leq}^k$ be the unique maximum of A_s . Then for any triple $s_1 < s_2 < s_3 \in ((r, t) \cap \mathbb{Q}) \cup \{r, t\}$ we claim that

$$(7.2) \quad \mathcal{L}^*(\pi(s_1), s_1; \pi(s_2), s_2) + \mathcal{L}^*(\pi(s_2), s_2; \pi(s_3), s_3) = \mathcal{L}^*(\pi(s_1), s_1; \pi(s_3), s_3).$$

Indeed, by the metric composition law there exists $\mathbf{z}^{(1)}, \mathbf{z}^{(2)}, \mathbf{z}^{(3)} \in \mathbb{R}_{\leq}^k$, such that

$$\begin{aligned} \mathcal{L}^*(\mathbf{x}, r; \mathbf{y}, t) &= \mathcal{L}^*(\mathbf{x}, r; \mathbf{z}^{(1)}, s_1) + \mathcal{L}^*(\mathbf{z}^{(1)}, s_1; \mathbf{z}^{(2)}, s_2) \\ &\quad + \mathcal{L}^*(\mathbf{z}^{(2)}, s_2; \mathbf{z}^{(3)}, s_3) + \mathcal{L}^*(\mathbf{z}^{(3)}, s_3; \mathbf{y}, t). \end{aligned}$$

The triangle inequality for $\mathcal{L}^*$ (Proposition 6.9) and the uniqueness of maxima for the functions A_s ensures that $\mathbf{z}^{(i)} = \pi(s_i)$ for $i = 1, 2, 3$ and enforces equation (7.2). By Lemma 6.8, π is continuous at rational points and at r, t . Therefore we can extend π to a continuous function on $[r, t]$.

Finally we check that for any $r = t_0 < t_1 < \dots < t_m = t$, we have

$$\sum_{i=1}^m \mathcal{L}^*(\pi(t_{i-1}), t_{i-1}; \pi(t_i), t_i) = \mathcal{L}^*(\mathbf{x}, r; \mathbf{y}, t).$$

If for all $0 < i < m$, t_i is rational, this follows by (7.2). This extends to general times t_i by the continuity of $\mathcal{L}^*$. We conclude that π is an optimizer. $\square$

We can upgrade the existence of optimizers to hold simultaneously for all pairs of endpoints, although the same cannot be achieved for uniqueness.

LEMMA 7.4. *Almost surely, for any $(\mathbf{x}, s; \mathbf{y}, t) \in \mathfrak{X}_\uparrow$, there is an optimizer in the extended landscape $\mathcal{L}^*$ from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$.*

PROOF. By Lemma 7.2, almost surely there is a unique optimizer between any pair of rational endpoints. For any $\mathbf{u} = (\mathbf{x}, s; \mathbf{y}, t) \in \mathfrak{X}_\uparrow$, we can take a sequence of rational points $(\mathbf{x}^{(i)}, s_i; \mathbf{y}^{(i)}, t_i) \in \mathfrak{X}_\uparrow$ converging to $\mathbf{u}$, and let $\pi^{(i)}$ be the unique optimizer from $(\mathbf{x}^{(i)}, s_i)$ to $(\mathbf{y}^{(i)}, t_i)$. All these optimizers are Hölder- $(2/3)^-$ with a common Hölder constant c by Lemma 6.8. Therefore the sequence $\pi^{(i)}$ has a subsequential limit π in the dyadic pointwise topology which is itself Hölder- $(2/3)^-$ continuous. By Lemma 7.1 and the continuity of $\mathcal{L}^*$, we have that π is an optimizer from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$. $\square$

We finish this section with the proof of Lemma 7.3. We first reduce the problem to understanding the sum of two multi-point last passage values across two independent parabolic Airy line ensembles $\mathcal{B}, \mathcal{B}'$. This is an optimization problem involving parabolic paths across $\mathcal{B}, \mathcal{B}'$. The remainder of the proof is essentially a resampling argument for parabolic paths, similar in spirit to the proof of Lemma 2.15.

The general idea here is to fix a rational interval I , and consider two restricted versions of the optimization problem: one which forces exactly one of the parabolic paths to jump in the interval I either on or off of the line $\mathcal{B}_1$, and one which does not allow any jumps on or off of $\mathcal{B}_1$ in the interval I . Using the Brownian Gibbs property for $\mathcal{B}$ (Theorem 2.14), we conclude that almost surely these two restricted optimization problems have different maxima. Now if the function A_s has two different maxima $\mathbf{z}, \mathbf{z}'$, then we show that almost surely we can find a rational interval I where the two optimizations problems discussed above have the same maximum, yielding a contradiction.

PROOF OF LEMMA 7.3. We now implement the arguments summarized above.

Step 1: Reducing to multi-point last passage values across $\mathcal{B}$. By symmetries of $\mathcal{L}^*$ (Lemma 6.10), we can assume that $r = -1, s = 0, t > 0$ and $x_1 > 0 > y_k$. Since $\mathcal{L}^*$ has independent extended sheet marginals, using symmetries of extended sheets (Lemma 5.5) we have

$$(\mathcal{L}^*(\mathbf{x}, -1; \mathbf{z}, 0), \mathcal{L}^*(\mathbf{z}, 0; \mathbf{y}, t)) \stackrel{d}{=} (\mathcal{S}(\mathbf{x}, \mathbf{z}), \mathcal{S}_t(-\mathbf{y}, -\mathbf{z})),$$

where $\mathcal{S}, \mathcal{S}_t$ are independent extended sheets of scale 1 and $t^{1/3}$. Therefore by Proposition 5.8, we have

$$(7.3) \quad A_s(\cdot) \stackrel{d}{=} \mathcal{B}[\mathbf{x} \rightarrow \cdot] + \mathcal{B}'[-\mathbf{y} \rightarrow -\cdot],$$

where $\mathcal{B} = \{\mathcal{B}_i\}_{i \in \mathbb{N}}$ is a parabolic Airy line ensemble, $\mathcal{B}' = \{\mathcal{B}'_i\}_{i \in \mathbb{N}}$ is independent of $\mathcal{B}$, and $\{x \mapsto t^{-1/3} \mathcal{B}'_i(t^{2/3}x)\}_{i \in \mathbb{N}}$ is a parabolic Airy line ensemble. We will show that the right-hand side of (7.3) has a unique maximum $\mathbf{z}$. While the argument is similar in spirit to the one used in Lemma 2.15, there are extra complexities coming from the definition of length for parabolic paths.

Step 2: Setup of paths intersecting/avoiding an interval. We next define certain collections of parabolic paths that do or not jump in the given interval.

We will later show that almost surely, these collections have different maximum weights.

We start by setting up notation for jump times. For a disjoint k -tuple of parabolic paths $\pi = (\pi_1, \dots, \pi_k)$ from $\mathbf{x}$ to $\mathbf{z}$, let

$$(7.4) \quad z[\pi]_{i,m} = \sup\{w \leq z_i : \pi_i(w) \geq m+1\},$$

for any $i \in \llbracket 1, k \rrbracket$ and non-negative integer m . In words, $z[\pi]_{i,m}$ is the jump time from line $m+1$ to line m for π_i (when $m \geq 1$), and $z[\pi]_{i,0} = z_i$.

For any interval $I \subset \mathbb{R}$, $1 \leq i \leq k$, and $j \in \{0, 1\}$, let $\mathcal{P}_{I,i}^{(j)}$ be the collection of all k -tuples of essentially disjoint parabolic paths π in $\mathcal{B}$ such that $z[\pi]_{i,j} \in I$, and $z[\pi]_{i',m} \notin I$ for any $(i', m) \neq (i, j)$. Also let $\mathcal{P}_I^c$ be the collection of all k -tuples of essentially disjoint parabolic paths π in $\mathcal{B}$ such that $z[\pi]_{i',m} \notin I$ for any (i', m) .

In other words, $\mathcal{P}_{I,i}^{(0)}$, $\mathcal{P}_{I,i}^{(1)}$ and $\mathcal{P}_I^c$ contain k -tuples of essentially disjoint parabolic paths π , satisfying the following conditions:

- $\mathcal{P}_{I,i}^{(0)}$: the endpoint of π_i is in I , and no other endpoint or jump point from line 2 to line 1 is in I .
- $\mathcal{P}_{I,i}^{(1)}$: the jump point of π_i from line 2 to line 1 is in I , and no endpoint or other jump point from line 2 to line 1 is in I .
- $\mathcal{P}_I^c$: no endpoint or jump point from line 2 to line 1 is in I .

Now we fix a compact interval $I \subset \mathbb{R}$ and $1 \leq i \leq k$. Define M_0, M_1, M_c using the same expression

$$\sup \|\pi\|_{\mathcal{B}} + \|\pi'\|_{\mathcal{B}'},$$

where the supremums are over different sets of pairs of disjoint k -tuples π, π' from $\mathbf{x}$ to $\mathbf{z}$ and $-\mathbf{y}$ to $-\mathbf{z}$ for some $\mathbf{z} \in \mathbb{R}_{\leq}^k$. For M_0 , we require that $\pi \in \mathcal{P}_{I,i}^{(0)}$. For M_1 , we require $\pi \in \mathcal{P}_{I,i}^{(1)}$. For M_c , we require $\pi \in \mathcal{P}_I^c$. We have no additional restriction on π' . See Figure 7.1 for an illustration of the paths.

Step 3: Almost surely $M_0 \neq M_c$ and $M_1 \neq M_c$. The general idea for this step is to show that conditional on M_c , the random variable M_0 (or M_1) has a continuous distribution, using the Brownian Gibbs property for the interval I .

We let $\mathcal{F}$ be the σ -algebra generated by null sets, $\mathcal{B}'$, all $\mathcal{B}_m$ for $m \geq 2$, and $\{\mathcal{B}_1(x) : x \notin I\}$. Then M_c is $\mathcal{F}$ -measurable, since the function recording all lengths of paths $\pi \in \mathcal{P}_{I,i}^c$ is $\mathcal{F}$ -measurable by Lemma 4.4.

We then deduce that, given $\mathcal{F}$, the random variable M_0 has continuous distribution. For this, we investigate how M_0 depends on $\mathcal{B}_1$ in I (see the left panel of Figure 7.1). By Lemma 4.3 we can write

$$M_0 = \sup_{\pi, \pi'} \|\pi\|_{\mathcal{B}} + \|\pi'\|_{\mathcal{B}'} + \sup_{x \in I} (\mathcal{B}_1(x) + \mathcal{B}'_1(x)) - \mathcal{B}_1(z_I) - \mathcal{B}'_1(-z'_I).$$

Here z_I, z'_I are the left and right endpoints of I , i.e., $I = [z_I, z'_I]$. The first supremum above is taken over all $\pi \in \mathcal{P}_I^c$ and $\pi' \in \mathcal{P}_{-I}^c$, where π is from $\mathbf{x}$ to $\mathbf{z}$ and π' is from $-\mathbf{y}$ to $\mathbf{z}'$, and such that $z_i = z_I, z'_{k+1-i} = -z'_I$, and $z_j = -z'_{k+1-j}$ for any $j \neq i$. Therefore $M_0 - \sup_{x \in I} (\mathcal{B}_1(x) + \mathcal{B}'_1(x))$ is $\mathcal{F}$ -measurable by Lemma 4.4. On the other hand, the Brownian Gibbs property for $\mathcal{B}$ (Theorem 2.14) implies that conditioned on $\mathcal{F}$ the law of $\mathcal{B}_1$ on I is absolutely continuous to a Brownian bridge (with diffusion parameter 2). Therefore conditioned on $\mathcal{F}$, the random variable $\sup_{x \in I} (\mathcal{B}_1(x) + \mathcal{B}'_1(x))$ almost surely has a continuous distribution, and hence so does M_0 . Since M_c is $\mathcal{F}$ -measurable, $M_0 \neq M_c$ almost surely. The argument to show that $M_1 \neq M_c$ almost surely is similar. Moreover, these inequalities hold almost




FIGURE 7.1. An illustration of the paths used in defining M_0 (upper left), M_1 (upper right), and M_c (below). The parabolic Airy line ensemble $\mathcal{B}'$ is rotated by 180 degrees for the picture. The green regions indicate the interval I and the lines where there are constraints on jumps. For M_0 , here we have $\pi \in \mathcal{P}_{I,2}^{(0)}$; for M_1 , we have $\pi \in \mathcal{P}_{I,2}^{(1)}$; for M_c , we have $\pi \in \mathcal{P}_I^c$.

surely simultaneously for all compact rational intervals I and all $i \in \llbracket 1, k \rrbracket$. Below we assume this probability one event.

Step 4: Contradiction with two maximums. Now we consider the function $\mathbf{z} \mapsto \mathcal{B}[\mathbf{x} \rightarrow \mathbf{z}] + \mathcal{B}'[-\mathbf{y} \rightarrow -\mathbf{z}]$. Note that by the metric composition law (Proposition 5.10) and symmetry of the extended Airy sheet (Lemma 5.5) this function attains its maximum. Suppose that the maximum is attained at two points $\mathbf{z}^{(1)} \neq \mathbf{z}^{(2)}$. We take any disjoint optimizers $\pi^{(1)}$ and $\pi^{(2)}$ in $\mathcal{B}$, from $\mathbf{x}$ to $\mathbf{z}^{(1)}$ and $\mathbf{z}^{(2)}$; and $\pi'^{(1)}$ and $\pi'^{(2)}$ in $\mathcal{B}'$, from $-\mathbf{y}$ to $-\mathbf{z}^{(1)}$ and $-\mathbf{z}^{(2)}$. Such optimizers exist by Proposition 5.8. Consider the jump times $z[\pi^{(1)}]_{i,0} = z_i^{(1)}$, $z[\pi^{(1)}]_{i,1}$, and $z[\pi^{(2)}]_{i,0} = z_i^{(2)}$, $z[\pi^{(2)}]_{i,1}$. By Lemma 7.5 below (and see Figure 7.2), almost surely we have

$$z[\pi^{(1)}]_{1,1} < z[\pi^{(1)}]_{1,0} < z[\pi^{(1)}]_{2,1} < \cdots < z[\pi^{(1)}]_{k,1} < z[\pi^{(1)}]_{k,0},$$

and

$$z[\pi^{(2)}]_{1,1} < z[\pi^{(2)}]_{1,0} < z[\pi^{(2)}]_{2,1} < \cdots < z[\pi^{(2)}]_{k,1} < z[\pi^{(2)}]_{k,0}.$$

Since $\mathbf{z}^{(1)} \neq \mathbf{z}^{(2)}$, we can find a number in the first sequence which does not appear in the second sequence. There are two cases:

- (1) We can find $i \in \llbracket 1, k \rrbracket$, such that $z[\pi^{(1)}]_{i,0} = z_i^{(1)}$ is not in the second sequence. Then we can find a rational interval I that contains $z_i^{(1)}$ and does not contain any other number in the two sequences. Thus $\pi^{(1)} \in \mathcal{P}_{I,i}^{(0)}$ and $\pi^{(2)} \in \mathcal{P}_{I,i}^c$. Then we have that $M_0 = M_c$ for such I and i .
- (2) We can find $i \in \llbracket 1, k \rrbracket$, such that $z[\pi^{(1)}]_{i,1}$ is not in the second sequence. Then we can find a rational interval I that contains $z[\pi^{(1)}]_{i,1}$ and does not contain any other number in the two sequences. Thus $\pi^{(1)} \in \mathcal{P}_{I,i}^{(1)}$ and $\pi^{(2)} \in \mathcal{P}_{I,i}^c$. Then we have that $M_1 = M_c$ for such I and i .

In either case, we get a contradiction with the assumption that $M_0 \neq M_c$, $M_1 \neq M_c$ for all compact rational intervals I and all $i \in \llbracket 1, k \rrbracket$. This means that A_s cannot have two different maximums. $\square$

It remains to prove the following lemma, which says that almost surely, consecutive jumps happen at different locations.

LEMMA 7.5. *Almost surely the following statement is true. Fix $t > 0$, and as in the proof of Lemma 7.3, let $\mathcal{B}$ be a parabolic Airy line ensemble, $\mathcal{B}' = \{\mathcal{B}'_i\}_{i \in \mathbb{N}}$ be independent of $\mathcal{B}$ such that $\{x \mapsto t^{-1/3}\mathcal{B}'_i(t^{2/3}x)\}_{i \in \mathbb{N}}$ is a parabolic Airy line ensemble.*

Take any $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$ such that $x_1 > 0 > y_k$. Let $\mathbf{z}^$ be any maximum of*

$$\mathbf{z} \mapsto \mathcal{B}[\mathbf{x} \rightarrow \mathbf{z}] + \mathcal{B}'[-\mathbf{y} \rightarrow -\mathbf{z}],$$

and let π^, π'^* be optimizers in $\mathcal{B}$ from $\mathbf{x}$ to $\mathbf{z}^*$, and in $\mathcal{B}'$ from $-\mathbf{y}$ to $-\mathbf{z}^*$, respectively. Recall the notation from (7.4), and let $z_{i,m}^* = z[\pi^*]_{i,m}$ and $-z_{k+1-i,-m}^* = z[\pi'^*]_{i,m}$ for any $i \in \llbracket 1, k \rrbracket$ and non-negative integer m (see Figure 7.2). Then for all $m \in \mathbb{Z}$, we have $z_{i,m}^* \neq z_{i,m+1}^*$ for $1 \leq i \leq k$, and $z_{i-1,m}^* \neq z_{i,m+1}^*$ for $2 \leq i \leq k$.*




FIGURE 7.2. An illustration of the jump times in the statement of Lemma 7.5.

Since $\llbracket 1, k \rrbracket \times \mathbb{Z}$ is countable, it suffices to show that for fixed (i, m) , almost surely $z_{i,m}^* \neq z_{i,m+1}^*$ and $z_{i-1,m}^* \neq z_{i,m+1}^*$ (if $i \geq 2$). The idea is again to use the Brownian Gibbs property, reducing this to the following statement: on a compact

interval, two Brownian motions that are correlated (but not identical or trivially related) almost surely attain their maxima at different points.

PROOF. We note that there are interlacing relations $z_{i,m+1}^* \leq z_{i,m}^*$ and $z_{i-1,m}^* \leq z_{i,m+1}^*$, from the definition and essential disjointness of the parabolic paths.

Now suppose that there is some $\hat{z} \in \mathbb{R}$, such that $\hat{z} = z_{i,m}^* = z_{i,m+1}^*$ or $\hat{z} = z_{i-1,m}^* = z_{i,m+1}^*$ for some i, m . The plan is to use the Brownian Gibbs property to reduce this problem to studying the arg max of Brownian motion. The interlacing relations impose constraints on the domain of the arg max, so we need to consider all jump times that equal $\hat{z}$.

By looking at all $z_{i,m}^*$ that equal $\hat{z}$, we can find some integers m_- and m_+ with $m_+ > m_-$, such that for each $m \in [m_-, m_+]$, there is at least one jump time equal to $\hat{z}$ on line m , and there is no jump time equal $\hat{z}$ on the lines $m_- - 1$ and $m_+ + 1$. We then let $\Phi \subset [1, k] \times \mathbb{Z}$ be the set of indices for all jump times equal to $\hat{z}$ amongst the lines $[m_-, m_+]$. In other words, we define

$$\Phi = \{(i, m) : z_{i,m}^* = \hat{z}, \quad m_- \leq m \leq m_+\},$$

such that $z_{i,m_- - 1}^* \neq \hat{z}$, $z_{i,m_+ + 1}^* \neq \hat{z}$ for any $i \in [1, k]$.

We note that the number of Φ of this form is countable. It now suffices to prove the following claim.

CLAIM. For any fixed $\Phi \subset [1, k] \times \mathbb{Z}$ such that $\Phi \cap ([1, k] \times \{m\})$ is nonempty if and only if $m \in [m_-, m_+]$ for some $m_- < m_+$, almost surely we cannot find a number $\hat{z}$ so that $\Phi = \{(i, m) : z_{i,m}^* = \hat{z}, \quad m_- - 1 \leq m \leq m_+ + 1\}$.

To prove this claim, the general strategy is as follows:

- (1) translate this event to properties of the arg max of some linear combination of lines of $\mathcal{B}$;
- (2) replace $\mathcal{B}$ by independent Brownian motions (with diffusion parameter 2);
- (3) analyze the probability of certain events involving Brownian motions.

CLAIM PROOF. It suffices to show that, for any such Φ and any rational interval I , almost surely the following event does not happen: we can find $\hat{z} \in I'$, where I' is the middle $1/3$ of I , such that

- for any $(i, m) \in \Phi$, we have $z_{i,m}^* = \hat{z}$,
- for any $(i, m) \in [m_- - 1, m_+ + 1] \times \mathbb{Z} \setminus \Phi$, we have $z_{i,m}^* \notin I$.

By comparing path lengths using Lemma 4.3, this implies that $(\mathcal{B}_{m^-}, \dots, \mathcal{B}_{m^+ + 1}) \in A$, where A is the following event defined on sequences of continuous functions $f_{m^-}, \dots, f_{m^+ + 1} : I \rightarrow \mathbb{R}$. The function

$$g_f(\{z_{i,m}\}_{(i,m) \in \Phi}) = \sum_{(i,m) \in \Phi} f_{m+1}(z_{i,m}) - f_m(z_{i,m})$$

defined on the domain where $z_{i,m} \in I$, $z_{i,m+1} \leq z_{i,m}$ and $z_{i-1,m} \leq z_{i,m+1}$ for all i and m attains its maximum when $\{z_{i,m}\}_{(i,m) \in \Phi} = \{\hat{z}\}_{(i,m) \in \Phi}$ for some $\hat{z} \in I'$. Here when defining $(\mathcal{B}_{m^-}, \dots, \mathcal{B}_{m^+ + 1})$ we let $\mathcal{B}_m = \mathcal{B}'_{1-m}$ for $m \leq 0$.

We next apply the Brownian Gibbs property to $\mathcal{B}$ and $\mathcal{B}'$, which enables us to replace $(\mathcal{B}_{m^-}, \dots, \mathcal{B}_{m^+ + 1})$ by independent Brownian motions (with diffusion parameter 2). Namely, it now suffices to prove that almost surely $B = (B_{m^-}, \dots, B_{m^+ + 1}) \in A$, where the B_i are independent Brownian motions (with diffusion parameter 2).

The last step is to prove $\mathbb{P}(B \in A) = 0$, by studying the Brownian motions and using inequalities given by the maximum condition. Consider the functions

$$S_1 = \sum_{(i,m) \in \Phi} B_{m+1} - B_m, \quad S_2 = B_{m_0+1} - B_{m_0},$$

where $(i_0, m_0) \in \Phi$ is chosen so that $(i_0, m_0 - 1), (i_0 + 1, m_0 + 1) \notin \Phi$. Under the event A , for any $z \in I$ we have

$$g_B(\{\hat{z}\}_{(i,m) \in \Phi}) \geq g_B(\{z\}_{(i,m) \in \Phi}),$$

so $S_1(\hat{z}) \geq S_1(z)$. We can also just deviate the (i_0, m_0) coordinate, but now only in one direction since we must preserve the interlacing conditions. Namely, for any $z \in I$, $z \geq \hat{z}$, under the event A we have

$$g_B(\{\hat{z}\}_{(i,m) \in \Phi}) \geq g_B(\{\hat{z} + \mathbf{1}((i, m) = (i_0, m_0))(z - \hat{z})\}_{(i,m) \in \Phi}),$$

so $S_2(\hat{z}) \geq S_2(z)$. It remains to show that, for any $\hat{z} \in I'$,

$$(7.5) \quad \mathbb{P}(S_2(\hat{z}) \geq S_2(z), \forall z \geq \hat{z}, z \in I \mid \hat{z} = \arg \max_I S_1) = 0.$$

Note that S_1 and S_2 are Brownian motions, and we can write $S_2 = \alpha S_1 + \beta S_3$ for some $\alpha, \beta \in \mathbb{R}$, $\beta > 0$, where S_3 is a Brownian motion independent of S_1 . However, for any $\hat{z} \in I'$, we have

$$(7.6) \quad \mathbb{P}\left(\limsup_{z \searrow \hat{z}} \frac{S_3(z) - S_3(\hat{z})}{S_1(\hat{z}) - S_1(z)} > \alpha\beta^{-1} \mid \hat{z} = \arg \max_I S_1\right) = 1.$$

This is because, conditional on $\hat{z} = \arg \max_I S_1$ and as $z \searrow \hat{z}$, a rescaling of $S_1(\hat{z}) - S_1(z)$ converges in distribution to a Bessel process of order 3. By Blumenthal's zero-one law, the left-hand side of (7.6) is either zero or one. It cannot be zero since the distribution of the ratio of a Brownian motion over an independent Bessel process of order 3 is unbounded.

From (7.6) we get (7.5), which means that $\mathbb{P}(B \in A) = 0$ and the conclusion follows. $\square$

Given the above claim, since the number of possible Φ is countable, the conclusion follows. $\square$

7.3. Monotonicity of optimizers

In this section we aim to establish monotonicity for optimizers in the extended landscape. Some arguments are in parallel to those in Chapter 2.2 for last passage across lines. We first establish that leftmost and rightmost optimizers are well-defined. For this lemma we write $\pi \leq \pi'$ for two multi-paths π and π' if the weak inequality holds pointwise and coordinatewise.

LEMMA 7.6. *The following statement holds almost surely for $\mathcal{L}^*$. For any $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$ and $s < t$, there are optimizers π^ℓ, π^r from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$, such that $\pi^\ell \leq \pi \leq \pi^r$ for any other optimizer π from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$. We call π^ℓ and π^r the leftmost and rightmost optimizers from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$, respectively.*

PROOF. Let $\pi^{(1)}$ and $\pi^{(2)}$ be two optimizers from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$. Let $\pi^{(3)} = \pi^{(1)} \wedge \pi^{(2)}$ and $\pi^{(4)} = \pi^{(1)} \vee \pi^{(2)}$, where $\wedge$ and $\vee$ are defined pointwise and coordinatewise. Since all optimizers are continuous by Lemma 6.8, $\pi^{(3)}$ and $\pi^{(4)}$ are

continuous functions from $[t, s]$ to $\mathbb{R}_{\leq}^k$. By the definition of $\|\cdot\|_{\mathcal{L}^*}$, the fact that $\mathcal{L}^*$ has extended Airy sheet marginals, and Lemma 5.7, we have that

$$\|\pi^{(3)}\|_{\mathcal{L}^*} + \|\pi^{(4)}\|_{\mathcal{L}^*} \geq \|\pi^{(1)}\|_{\mathcal{L}^*} + \|\pi^{(2)}\|_{\mathcal{L}^*}.$$

Since $\pi^{(1)}, \pi^{(2)}$ are optimizers, this must be an equality. Thus $\pi^{(3)}, \pi^{(4)}$ are also optimizers.

Now consider any monotone sequence of optimizers $\pi^{(1)} \leq \pi^{(2)}, \dots$. By Lemma 6.8, this sequence has a bounded pointwise limit π' on dyadic rationals, and π' is continuous. This limit is also an optimizer by Lemma 7.1. Thus by Zorn's lemma, there is an optimizer π^ℓ , such that for any optimizer π , the condition $\pi \leq \pi^\ell$ implies $\pi = \pi^\ell$. Thus for any optimizer π , since the multi-path $\pi \wedge \pi^\ell$ is an optimizer satisfying $\pi \wedge \pi^\ell \leq \pi^\ell$, we must have $\pi \wedge \pi^\ell = \pi^\ell$, implying that $\pi^\ell \leq \pi$. Therefore π^ℓ is the leftmost optimizer. The existence of the rightmost optimizer follows similarly. $\square$

LEMMA 7.7. *The following statements hold almost surely. For any $\mathbf{x} \leq \mathbf{x}', \mathbf{y} \leq \mathbf{y}' \in \mathbb{R}_{\leq}^k$ and $s < t$, let π be the leftmost (resp. rightmost) optimizer from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$ and π' be the leftmost (resp. rightmost) optimizer from $(\mathbf{x}', s)$ to $(\mathbf{y}', t)$. Then $\pi \leq \pi'$.*

PROOF. We prove for the case where π, π' are the leftmost optimizers. The rightmost case follows similarly. Define $\pi^\ell = \pi \wedge \pi'$ and $\pi^r = \pi \vee \pi'$. Then π^ℓ and π^r are both continuous multi-paths from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$, and from $(\mathbf{x}', s)$ to $(\mathbf{y}', t)$, respectively.

We claim that $\|\pi^\ell\|_{\mathcal{L}^*} + \|\pi^r\|_{\mathcal{L}^*} \geq \|\pi\|_{\mathcal{L}^*} + \|\pi'\|_{\mathcal{L}^*}$. Indeed, this follows by the definition of $\|\cdot\|_{\mathcal{L}^*}$ and Lemma 5.7. However, we also have $\|\pi\|_{\mathcal{L}^*} \geq \|\pi^\ell\|_{\mathcal{L}^*}$ and $\|\pi'\|_{\mathcal{L}^*} \geq \|\pi^r\|_{\mathcal{L}^*}$, by the definition of optimizers. Therefore $\|\pi\|_{\mathcal{L}^*} = \|\pi^\ell\|_{\mathcal{L}^*}$ and $\|\pi'\|_{\mathcal{L}^*} = \|\pi^r\|_{\mathcal{L}^*}$, and hence π^ℓ, π^r are also optimizers. As π is the leftmost optimizer, we have $\pi \leq \pi^\ell$. On the other hand, $\pi^\ell \leq \pi, \pi'$ from the definition of π^ℓ . Thus $\pi = \pi^\ell \leq \pi'$. $\square$

7.4. Sums of disjoint paths

The goal of this section is to show the following proposition.

PROPOSITION 7.8. *Almost surely the following statement is true. Take any $s < t$ and any multi-paths $\pi : [s, t] \rightarrow \mathbb{R}_{\leq}^k$ and $\pi' : [s, t] \rightarrow \mathbb{R}_{\leq}^{k'}$, such that for any $r \in (s, t)$, we have $\pi_k(r) \leq \pi'_1(r)$. Let $\pi'' : [s, t] \rightarrow \mathbb{R}_{\leq}^{k+k'}$ be such that $\pi''(r) = (\pi(r), \pi'(r))$ for all $r \in (s, t)$. Then*

$$\|\pi''\|_{\mathcal{L}^*} \leq \|\pi\|_{\mathcal{L}^*} + \|\pi'\|_{\mathcal{L}^*}.$$

Moreover, if $\pi_k(r) < \pi'_1(r)$ for all $r \in (s, t)$, then $\|\pi''\|_{\mathcal{L}^} = \|\pi\|_{\mathcal{L}^*} + \|\pi'\|_{\mathcal{L}^*}$.*

The inequality in Proposition 7.8 is immediate from the definition of $\|\cdot\|_{\mathcal{L}^*}$, and the fact that

$$\mathcal{L}^*(\mathbf{x}, s, \mathbf{y}, t) + \mathcal{L}^*(\mathbf{x}', s, \mathbf{y}', t) \geq \mathcal{L}^*((\mathbf{x}, \mathbf{x}'), s, (\mathbf{y}, \mathbf{y}'), t)$$

for any $\mathbf{x}, \mathbf{y}, \mathbf{x}', \mathbf{y}', s, t$ for which both sides above make sense. This inequality is inherited from the prelimit $\mathcal{L}_n$, where it is clear. To prove the claimed equality in Proposition 7.8, we require a few lemmas.

LEMMA 7.9. *Let $k, \ell \in \mathbb{N}$ and $h, t > 0$. Then*

$$\mathbb{P}(\mathcal{L}^*(0^k, 0; 0^k, t) + \mathcal{L}^*(h^\ell, 0; h^\ell, t) > \mathcal{L}^*((0^k, h^\ell), 0; (0^k, h^\ell), t)) < ce^{-dh^3t^{-2}},$$

for c, d depending only on k, ℓ .

PROOF. By rescaling we can assume that $t = 1$. Since $\mathcal{L}^*$ has extended Airy sheet marginals, by Proposition 5.8 the probability in question is the same as the probability of the event

$$A = \{\mathcal{B}[0^k \rightarrow 0^k] + \mathcal{B}[h^\ell \rightarrow h^\ell] > \mathcal{B}[(0^k, h^\ell) \rightarrow (0^k, h^\ell)]\}.$$

By Proposition 5.9, the event A implies that every optimizer π from h^ℓ to h^ℓ in $\mathcal{B}$ intersects the first k lines of $\mathcal{B}$ in the interval $(-\infty, 0]$. That is, $\pi_1(0) \leq k$, and so

$$(7.7) \quad \mathcal{B}[h^\ell \rightarrow h^\ell] \leq \mathcal{B}[h^{\ell-1} \rightarrow h^{\ell-1}] + \mathcal{B}[h \rightarrow 0] + \mathcal{B}_1(h) - \mathcal{B}_k(0).$$

Now by translation invariance of the extended Airy sheet $\mathcal{S}$, $\mathcal{B}[h^\ell \rightarrow h^\ell] - \mathcal{B}[h^{\ell-1} \rightarrow h^{\ell-1}] \stackrel{d}{=} \mathcal{B}_\ell(0)$ and $\mathcal{B}[h \rightarrow 0] \stackrel{d}{=} \mathcal{B}_1(-h)$. Finally, since $\mathcal{B}(x) + x^2$ is stationary, (7.7) is equivalent to an inequality of the form

$$(7.8) \quad X_1 + X_2 + 2h^2 \leq X_3 + X_4,$$

where each of the random variables X_i are equal in distribution to $\mathcal{B}_i(0)$ for some $i \leq k \wedge \ell$. The points $\mathcal{B}_i(0)$ are points in the Airy point process, which are known to have well-controlled tails. For example, we can pass Theorem A.1 to the limit to get that $\mathbb{P}(|X_i| > a) \leq ce^{-da^{3/2}}$ for all $1 \leq i \leq 4$ and constants c, d that depend only on k, ℓ . Therefore by a union bound, the probability of (7.8) is bounded above by $ce^{-dh^3t^{-2}}$, completing the proof. $\square$

LEMMA 7.10. *For any $s < t$, almost surely the following statement holds. For any $\mathbf{x}^{(1)} \leq \mathbf{x}^{(2)} \leq \mathbf{x}^{(3)} \leq \mathbf{x}^{(4)} \in \mathbb{R}_{\leq}^k$, and $\mathbf{y}^{(1)} \leq \mathbf{y}^{(2)} \leq \mathbf{y}^{(3)} \leq \mathbf{y}^{(4)} \in \mathbb{R}_{\leq}^k$, if*

$$\mathcal{L}^*((\mathbf{x}^{(2)}, \mathbf{x}^{(3)}), s; (\mathbf{y}^{(2)}, \mathbf{y}^{(3)}), t) = \mathcal{L}^*(\mathbf{x}^{(2)}, s; \mathbf{y}^{(2)}, t) + \mathcal{L}^*(\mathbf{x}^{(3)}, s; \mathbf{y}^{(3)}, t),$$

we then have

$$\mathcal{L}^*((\mathbf{x}^{(1)}, \mathbf{x}^{(4)}), s; (\mathbf{y}^{(1)}, \mathbf{y}^{(4)}), t) = \mathcal{L}^*(\mathbf{x}^{(1)}, s; \mathbf{y}^{(1)}, t) + \mathcal{L}^*(\mathbf{x}^{(4)}, s; \mathbf{y}^{(4)}, t).$$

PROOF. Without loss of generality we assume that $s = 0$, $t = 1$. With these choices, the lemma is a fact about an extended Airy sheet of scale 1. We claim that

$$(7.9) \quad \begin{aligned} & \mathcal{S}(\mathbf{x}^{(2)}, \mathbf{y}^{(2)}) + \mathcal{S}(\mathbf{x}^{(3)}, \mathbf{y}^{(3)}) - \mathcal{S}((\mathbf{x}^{(2)}, \mathbf{x}^{(3)}), (\mathbf{y}^{(2)}, \mathbf{y}^{(3)})) \\ & \geq \mathcal{S}(\mathbf{x}^{(1)}, \mathbf{y}^{(1)}) + \mathcal{S}(\mathbf{x}^{(4)}, \mathbf{y}^{(4)}) - \mathcal{S}((\mathbf{x}^{(1)}, \mathbf{x}^{(4)}), (\mathbf{y}^{(1)}, \mathbf{y}^{(4)})) \\ & \geq 0. \end{aligned}$$

For this we study the prelimiting sheet $\mathcal{S}^n$. By Lemma 2.5, for n large enough we have

$$\begin{aligned} & \mathcal{S}^n(\mathbf{x}^{(2)}, \mathbf{y}^{(2)}) + \mathcal{S}^n((\mathbf{x}^{(2)}, \mathbf{x}^{(3)}), (\mathbf{y}^{(1)}, \mathbf{y}^{(3)})) \\ & \geq \mathcal{S}^n(\mathbf{x}^{(2)}, \mathbf{y}^{(1)}) + \mathcal{S}^n((\mathbf{x}^{(2)}, \mathbf{x}^{(3)}), (\mathbf{y}^{(2)}, \mathbf{y}^{(3)})), \end{aligned}$$

and

$$\begin{aligned} & \mathcal{S}^n(\mathbf{x}^{(3)}, \mathbf{y}^{(3)}) + \mathcal{S}^n((\mathbf{x}^{(2)}, \mathbf{x}^{(3)}), (\mathbf{y}^{(1)}, \mathbf{y}^{(4)})) \\ & \geq \mathcal{S}^n(\mathbf{x}^{(3)}, \mathbf{y}^{(4)}) + \mathcal{S}^n((\mathbf{x}^{(2)}, \mathbf{x}^{(3)}), (\mathbf{y}^{(1)}, \mathbf{y}^{(3)})). \end{aligned}$$

Adding up these two inequalities, and passing to the limit via Theorem 5.4, we get

$$\begin{aligned} \mathcal{S}(\mathbf{x}^{(2)}, \mathbf{y}^{(2)}) + \mathcal{S}(\mathbf{x}^{(3)}, \mathbf{y}^{(3)}) - \mathcal{S}((\mathbf{x}^{(2)}, \mathbf{x}^{(3)}), (\mathbf{y}^{(2)}, \mathbf{y}^{(3)})) \\ \geq \mathcal{S}(\mathbf{x}^{(2)}, \mathbf{y}^{(1)}) + \mathcal{S}(\mathbf{x}^{(3)}, \mathbf{y}^{(4)}) - \mathcal{S}((\mathbf{x}^{(2)}, \mathbf{x}^{(3)}), (\mathbf{y}^{(1)}, \mathbf{y}^{(4)})). \end{aligned}$$

Similarly, we also have

$$\begin{aligned} \mathcal{S}(\mathbf{x}^{(2)}, \mathbf{y}^{(1)}) + \mathcal{S}(\mathbf{x}^{(3)}, \mathbf{y}^{(4)}) - \mathcal{S}((\mathbf{x}^{(2)}, \mathbf{x}^{(3)}), (\mathbf{y}^{(1)}, \mathbf{y}^{(4)})) \\ \geq \mathcal{S}(\mathbf{x}^{(1)}, \mathbf{y}^{(1)}) + \mathcal{S}(\mathbf{x}^{(4)}, \mathbf{y}^{(4)}) - \mathcal{S}((\mathbf{x}^{(1)}, \mathbf{x}^{(4)}), (\mathbf{y}^{(1)}, \mathbf{y}^{(4)})). \end{aligned}$$

Thus adding up the above two inequalities we get the first inequality in (7.9). The second inequality in (7.9) is obvious for $\mathcal{S}^n$, so by passing to the limit via Theorem 5.4 it also holds for $\mathcal{S}$. Finally, when the first line in (7.9) equals zero, so does the second line. The conclusion follows. $\square$

For this next lemma and vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^k$, we write

$$\min(\mathbf{x}, \mathbf{y}) = \min\{x_i \wedge y_i : i \in \llbracket 1, k \rrbracket\}.$$

We similarly define $\max(\mathbf{x}, \mathbf{y})$.

LEMMA 7.11. *For each $M, h > 0$ and $k, k' \in \mathbb{N}$, there is a random number $P > 0$ such that the following is true. For any $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k, \mathbf{x}', \mathbf{y}' \in \mathbb{R}_{\leq}^{k'}$ and $s, t \in \mathbb{R}$ with $\|\mathbf{x}\|_2, \|\mathbf{y}\|_2, |s|, |t| < M$, $s < t$, $t - s < P$, and $\min(\mathbf{x}', \mathbf{y}') - \max(\mathbf{x}, \mathbf{y}) > h$, we have*

$$\mathcal{L}^*((\mathbf{x}, \mathbf{x}'), s; (\mathbf{y}, \mathbf{y}'), t) = \mathcal{L}^*(\mathbf{x}, s; \mathbf{y}, t) + \mathcal{L}^*(\mathbf{x}', s; \mathbf{y}', t).$$

PROOF. For each $\ell \in \mathbb{Z}$, let

$$J_\ell = \{(x, x', s, t) : |x|, |s| < M, x \in 2^\ell \mathbb{Z}, x' = x + 2^\ell, s \in 2^{2\ell} \mathbb{Z}, t = s + 2^{2\ell-1}\}.$$

By Lemma 7.9, for any $\ell < 0$, with probability at least $1 - cM^2 2^{-3\ell} e^{-d2^{-\ell}}$ (for some constants c, d depending on k, k'), for any $(x, x', s, t) \in J_\ell$, we have

$$\mathcal{L}^*(x^k, s; x^k, t) + \mathcal{L}^*(x'^{k'}, s; x'^{k'}, t) = \mathcal{L}^*((x^k, x'^{k'}), s; (x^k, x'^{k'}), t).$$

Then almost surely, there is a random $L_0 \in \mathbb{Z}_-$, such that this event happens for all $\ell \leq L_0$. Now, we can choose P small enough such that for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k, \mathbf{x}', \mathbf{y}' \in \mathbb{R}_{\leq}^{k'}$ and $s, t \in \mathbb{R}$ satisfying the conditions of the lemma, we can find $\ell \leq L_0$ and $(\tilde{x}, \tilde{x}', \tilde{s}, \tilde{t}) \in J_\ell$ such that the following holds:

- $\tilde{s} < s < t < \tilde{t}$ and $\min(\mathbf{x}', \mathbf{y}') > \tilde{x} + h/3$, $\max(\mathbf{x}, \mathbf{y}) < \tilde{x}' - h/3$
- There exist optimizers π, π' from $(\tilde{x}^k, \tilde{s})$ to $(\tilde{x}^k, \tilde{t})$ and from $(\tilde{x}'^{k'}, \tilde{s})$ to $(\tilde{x}'^{k'}, \tilde{t})$, such that $\mathbf{x} \leq \pi(s) \leq \pi'(s) \leq \mathbf{x}'$ and $\mathbf{y} \leq \pi(t) \leq \pi'(t) \leq \mathbf{y}'$.

To ensure the second condition, we have used the transversal fluctuation bound on optimizers from Lemma 6.8. The fact that $\ell \leq L_0$ ensures that

$$\mathcal{L}^*(\tilde{x}^k, \tilde{s}; \tilde{x}^k, \tilde{t}) + \mathcal{L}^*(\tilde{x}'^{k'}, \tilde{s}; \tilde{x}'^{k'}, \tilde{t}) = \mathcal{L}^*((\tilde{x}^k, \tilde{x}'^{k'}), s; (\tilde{x}^k, \tilde{x}'^{k'}), t),$$

which implies that

$$\mathcal{L}^*(\pi(s), s; \pi(t), t) + \mathcal{L}^*(\pi'(s), s; \pi'(t), t) = \mathcal{L}^*((\pi(s), \pi'(s)), s; (\pi(t), \pi'(t)), t).$$

Finally, assuming that for any $s < t \in \mathbb{Q}$ the event in Lemma 7.10 holds, we have that $\mathcal{L}^*((\mathbf{x}, \mathbf{x}'), s; (\mathbf{y}, \mathbf{y}'), t) = \mathcal{L}^*(\mathbf{x}, s; \mathbf{y}, t) + \mathcal{L}^*(\mathbf{x}', s; \mathbf{y}', t)$, if $s, t \in \mathbb{Q}$. By continuity of $\mathcal{L}^*$ the conclusion follows. $\square$

PROOF OF THE EQUALITY IN PROPOSITION 7.8. First assume that $\|\pi''\|_{\mathcal{L}^*} > -\infty$. Let $\delta > 0$. Take $s < t_0 < t_1 < \dots < t_m < t$, such that $\sum_{i=1}^m \mathcal{L}^*(\pi''(t_{i-1}), t_{i-1}; \pi''(t_i), t_i) < \|\pi''\|_{\mathcal{L}^*} + \delta$, and $t_0 - t, s - t_m < \delta$. We next choose parameters to apply Lemma 7.11. Let

$$h = \min_{r \in [t_0, t_m]} \{\min \pi'(r) - \max \pi(t)\}, \quad M = \max\{|s|, |t|, \max_{r \in [t, s]} \|\pi''(r)\|_2\}.$$

Observe that $h > 0$ by the assumptions of the proposition. Let P be as in Lemma 7.11 for this M, h . Then we choose $\bar{m} \geq m$, and $t_0 = \bar{t}_0 < \bar{t}_1 < \dots < \bar{t}_{\bar{m}} = t_m$ such that $\{t_0, t_1, \dots, t_m\} \subset \{\bar{t}_0, \bar{t}_1, \dots, \bar{t}_{\bar{m}}\}$, and $\bar{t}_i - \bar{t}_{i-1} < P$ for each $1 \leq i \leq \bar{m}$. Then we have

$$\begin{aligned} & \|\pi\|_{\mathcal{L}^*} - \mathcal{L}^*(\pi(s), s; \pi(t_0), t_0) - \mathcal{L}^*(\pi(t_m), t_m; \pi(t), t) \\ & \quad + \|\pi'\|_{\mathcal{L}^*} - \mathcal{L}^*(\pi'(s), s; \pi'(t_0), t_0) - \mathcal{L}^*(\pi'(t_m), t_m; \pi'(t), t) \\ & \leq \sum_{i=1}^{\bar{m}} \mathcal{L}^*(\pi(\bar{t}_{i-1}), \bar{t}_{i-1}; \pi(\bar{t}_i), \bar{t}_i) + \mathcal{L}^*(\pi'(\bar{t}_{i-1}), \bar{t}_{i-1}; \pi'(\bar{t}_i), \bar{t}_i) \\ & = \sum_{i=1}^{\bar{m}} \mathcal{L}^*(\pi''(\bar{t}_{i-1}), \bar{t}_{i-1}; \pi''(\bar{t}_i), \bar{t}_i) < \|\pi''\|_{\mathcal{L}^*} + \delta. \end{aligned}$$

Now we send $\delta \rightarrow 0$. By Lemma 6.7 we have

$$\limsup_{\delta \rightarrow 0} \mathcal{L}^*(\pi(s), s; \pi(t_0), t_0) = \limsup_{\delta \rightarrow 0} \mathcal{L}^*(\pi(t_m), t_m; \pi(t), t) \leq 0,$$

$$\limsup_{\delta \rightarrow 0} \mathcal{L}^*(\pi'(s), s; \pi'(t_0), t_0) = \limsup_{\delta \rightarrow 0} \mathcal{L}^*(\pi'(t_m), t_m; \pi'(t), t) \leq 0.$$

Therefore $\|\pi\|_{\mathcal{L}^*} + \|\pi'\|_{\mathcal{L}^*} \leq \|\pi''\|_{\mathcal{L}^*}$, and our conclusion follows. In the case when $\|\pi''\|_{\mathcal{L}^*} = -\infty$, we can apply the same argument with an arbitrary $b \in \mathbb{R}$ in place of $\|\pi''\|_{\mathcal{L}^*} + \delta$ to get the result. $\square$

CHAPTER 8

Disjointness of Optimizers

The main goal of this chapter is to prove the following disjointness result.

PROPOSITION 8.1. *Almost surely, for any $(\mathbf{x}, r; \mathbf{y}, t) \in \mathfrak{X}_\dagger$, there exists an optimizer π in $\mathcal{L}^*$ from $(\mathbf{x}, r)$ to $(\mathbf{y}, t)$ such that $\pi_i(s) < \pi_j(s)$ for all $i < j$ and $s \in (r, t)$.*

Proposition 8.1 essentially says that disjoint optimizers in Brownian LPP remain disjoint even as we pass to the limit. In other words, it implies that constituent paths in disjoint optimizers in Brownian LPP are separated from each other by an amount that remains visible in the limiting scaling. This is not an obvious fact, and we will prove it in stages. Note also that Proposition 8.1 still leaves open the possibility that at exceptional points $(\mathbf{x}, r; \mathbf{y}, t) \in \mathfrak{X}_\dagger$ where there are multiple $\mathcal{L}^*$ -optimizers, some of these optimizers may have overlapping constituent paths; we do not believe that such optimizers exist but we do not attempt to resolve this issue here.

By using Proposition 8.1 and checking the definition of the extended directed landscape from Definition 1.1, we will show that $\mathcal{L}^* = \mathcal{L}$. This is the content of Theorem 1.6. We will also use Proposition 8.1 to prove Theorem 1.7 and Corollaries 1.9 and 1.11.

8.1. Convergence in the overlap topology

We start with the following weaker result, which says that all optimizers are disjoint at fixed time.

LEMMA 8.2. *For any fixed s the following holds almost surely for $\mathcal{L}^*$. For any $(\mathbf{x}, \mathbf{y}) \in \mathfrak{X}$, $r < s < t$, and any optimizer π from $(\mathbf{x}, r)$ to $(\mathbf{y}, t)$, we have that $\pi_1(s) < \pi_2(s) < \dots < \pi_k(s)$.*

PROOF. As any optimizer restricted to a smaller interval of time is also an optimizer, it suffices to prove the result for fixed $r = s - \delta$ and $t = s + \delta$ with a fixed small δ , and for $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$ with each coordinate in a compact interval. Since $\mathcal{L}^*$ has extended Airy sheet marginals, the conclusion follows from Lemma 7.5 for compact sets $K \subset \mathfrak{X}$ such that every point $(\mathbf{x}, \mathbf{y}) \in K$ satisfies $x_1 > 0 > y_k$. For more general compact sets, the conclusion follows by skew symmetry of $\mathcal{L}^*$ (Lemma 6.10). $\square$

For any continuous paths $\pi_n : [r_n, t_n] \rightarrow \mathbb{R}_{\leq}^k$, $n \in \mathbb{N}$, and $\pi : [r, t] \rightarrow \mathbb{R}_{\leq}^k$, we say that $\pi_n \rightarrow \pi$ **in the overlap topology**, if for all large enough n , $O_n = \{s \in [r, t] \cap [r_n, t_n] : \pi_n(s) = \pi(s)\}$ is an interval, and the endpoints of O_n converge to r and t . Overlap was first introduced in [16, Section 3], and is particularly useful for studying geodesics or optimizers in the directed or extended landscape.

We aim to prove an overlap convergence result for optimizers. We will require two closely related results for $\mathcal{L}$ -geodesics from [16]. To state them, for any path

$\pi : [r, t] \rightarrow \mathbb{R}$, define the **graph** of π by

$$\mathbf{g}\pi := \{(\pi(s), s) : s \in [r, t]\}.$$

This is the usual graph of a function with coordinates reversed.

LEMMA 8.3 ([16, Lemma 3.1]). *Almost surely the following is true. Let $(p_n; q_n) \rightarrow (p; q) \in \mathbb{R}_+^4$, and let π_n be any sequence of geodesics from p_n to q_n . Then the sequence $\mathbf{g}\pi_n$ is precompact in the Hausdorff metric, and any subsequential limit is the graph of a geodesic from p to q .*

LEMMA 8.4 ([16, Lemma 3.3]). *Almost surely the following is true. Let $(p_n; q_n) \rightarrow (p; q) \in \mathbb{R}_+^4$, and let π_n be any sequence of geodesics from p_n to q_n . Suppose that $(p_n; q_n) \in \mathbb{Q}^4$, and $\mathbf{g}\pi_n \rightarrow \mathbf{g}\pi$ in the Hausdorff metric, for some geodesic π from p to q . Then $\pi_n \rightarrow \pi$ in the overlap topology.*

From these we can deduce the following result.

LEMMA 8.5. *Almost surely the following is true. Let $(p_n; q_n) = (x_n, s; y_n, t) \rightarrow (p; q) = (x, s; y, t) \in \mathbb{R}_+^4$, and suppose that $x_n \geq x, y_n \geq y$ for all n . Let π_n be the sequence of rightmost geodesics from p_n to q_n , and let π be the rightmost geodesic from p to q . Then $\pi_n \rightarrow \pi$ in the overlap topology.*

The existence of rightmost and leftmost geodesics follows from [15, Lemma 13.2]; alternately, it follows from Lemma 7.6.

PROOF OF LEMMA 8.5. First, by Lemma 8.3 the sequence $\mathbf{g}\pi_n$ is precompact in the Hausdorff metric and any subsequential limit is the graph of a geodesic from p to q . Consider such a subsequential limit $\mathbf{g}\pi'$. Since the π_n are rightmost geodesics, by Lemma 7.7 we have $\pi_n \geq \pi$ for all n , and hence $\pi' \geq \pi$. Since π is a rightmost geodesic, this implies $\pi' = \pi$, and therefore $\mathbf{g}\pi_n \rightarrow \mathbf{g}\pi$. Lemma 8.4 then completes the proof. $\square$

We can now upgrade the above lemma to optimizers in $\mathcal{L}^*$.

LEMMA 8.6. *Almost surely the following statement is true. Take any $r < t$, $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$, and two sequences $\mathbf{x}^{(i)}, \mathbf{y}^{(i)} \in \mathbb{R}_{\leq}^k$, for $i \in \mathbb{N}$. Suppose that $x_j^{(i)} > x_j$, $y_j^{(i)} > y_j$ for $j \in [1, k]$, and that $\mathbf{x}^{(i)} \rightarrow \mathbf{x}$, $\mathbf{y}^{(i)} \rightarrow \mathbf{y}$ as $i \rightarrow \infty$. Let $\pi^{(i)}$ be the rightmost optimizer from $(\mathbf{x}^{(i)}, r)$ to $(\mathbf{y}^{(i)}, t)$, and π be the rightmost optimizer from $(\mathbf{x}, r)$ to $(\mathbf{y}, t)$. Then $\pi^{(i)} \rightarrow \pi$ in the overlap topology.*

PROOF. First, the graphs of all the optimizer paths $\pi_j^{(i)}$ are Hölder- $(2/3)^-$ with a common Hölder constant by Lemma 6.8. Therefore along any subsequence we can take a further subsequence so that $\pi^{(i)}$ converges to a continuous limit in the dyadic pointwise topology. This limit must be an optimizer from $(\mathbf{x}, r)$ to $(\mathbf{y}, t)$, by Lemma 7.1 and the continuity of $\mathcal{L}^*$. Thus the limit must be π since $\pi \leq \pi^{(i)}$ for each i , by Lemma 7.7. We conclude that $\pi^{(i)} \rightarrow \pi$ in the dyadic pointwise topology.

We take any $s \in \mathbb{Q}_2$ with $r < s < t$ (recall that $\mathbb{Q}_2$ is the set of dyadic rational numbers). By Lemma 8.2 we can assume that $\pi_1(s), \dots, \pi_k(s)$ are pairwise distinct. Then in a small neighborhood of s , the paths $\pi_1, \dots, \pi_k$ are mutually disjoint. By Lemma 6.8, we can find a (random) $\delta > 0$, such that $\delta \in \mathbb{Q}_2$, and for any $0 < \delta_1^-, \dots, \delta_k^-, \delta_1^+, \dots, \delta_k^+ < \delta$, and any geodesics from $(\pi_j(s - \delta) + \delta_j^-, s - \delta)$ to $(\pi_j(s +$

$\delta) + \delta_j^+, s + \delta)$, $1 \leq j \leq k$, these geodesics are disjoint. By the dyadic convergence established above and Lemma 7.7, for all large enough i we have

$$(8.1) \quad \pi_j(s \pm \delta) \leq \pi_j^{(i)}(s \pm \delta) < \pi_j(s \pm \delta) + \delta$$

for all $j \in \llbracket 1, k \rrbracket$. From now on, we work with i such that (8.1) holds. For each $1 \leq j \leq k$, let $\tau_j^{(i)}$ be the rightmost geodesic from $(\pi_j^{(i)}(s - \delta), s - \delta)$ to $(\pi_j^{(i)}(s + \delta), s + \delta)$. We claim that $\pi_j^{(i)} = \tau_j^{(i)}$ on the interval $I_\delta = [s - \delta, s + \delta]$ as long as i is sufficiently large. Indeed, letting $\tau^{(i)} = (\tau_1^{(i)}, \dots, \tau_k^{(i)})$, we have

$$\|\tau^{(i)}\|_{\mathcal{L}^*} = \sum_{j=1}^k \|\tau_j^{(i)}\|_{\mathcal{L}^*} \geq \sum_{j=1}^k \|\pi_j^{(i)}|_{I_\delta}\|_{\mathcal{L}^*} \geq \|\pi^{(i)}|_{I_\delta}\|_{\mathcal{L}^*}.$$

Here the equality follows from Proposition 7.8 and (8.1), the first inequality uses that each $\tau_j^{(i)}$ is a geodesic, and the second inequality uses Proposition 7.8 again. Since $\pi^{(i)}$ is an optimizer, all inequalities above must be equalities, so $\tau^{(i)}$ must also be an optimizer, and all the paths $\pi_j^{(i)}|_{I_\delta}$ must be geodesics. Since $\pi^{(i)}$ is a rightmost optimizer, we have $\pi^{(i)}|_{I_\delta} \geq \tau^{(i)}$. Since each of the $\pi_j^{(i)}|_{I_\delta}$ are geodesics and each of the $\tau_j^{(i)}$ are rightmost geodesics, this implies that $\tau_j^{(i)} = \pi_j^{(i)}|_{I_\delta}$ for all i large enough.

The same argument shows that each $\pi_j|_{I_\delta}$ is also a rightmost geodesic. Therefore by Lemma 8.5, and the fact that for any $\delta > 0$, (8.1) holds for all large enough i shows that for i large enough we have $\pi_j^{(i)}(s) = \pi_j(s)$ for all $j \in \llbracket 1, k \rrbracket$.

Next we take $s_1, s_2 \in \mathbb{Q}_2$ with $r < s_1 < s_2 < t$. For i large enough we have $\pi^{(i)}(s_1) = \pi(s_1)$ and $\pi^{(i)}(s_2) = \pi(s_2)$. Thus $\pi^{(i)}(s) = \pi(s)$ for any $s_1 < s < s_2$ since on $[s_1, s_2]$ both $\pi^{(i)}$ and π are the rightmost optimizer from $(\pi(s_1), s_1)$ to $(\pi(s_2), s_2)$. By sending $s_1 \rightarrow r$ and $s_2 \rightarrow t$ we get the conclusion. $\square$

REMARK 8.7. A similar statement to Lemma 8.6 holds for convergence to leftmost optimizers.

8.2. Two paths

To prove Proposition 8.1, we start with the two-path case with fixed endpoints.

LEMMA 8.8. *Fix $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^2$ and $s < t$. Then almost surely, the unique optimizer in $\mathcal{L}^*$ from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$ consists of two paths that are disjoint, except possibly at the endpoints.*

The proof of Lemma 8.8 consists of two steps. First, we will handle the common endpoint case when $x_1 = x_2, y_1 = y_2$. This can be dealt with by proving a quantitative estimate on the location of optimizers from $\mathbf{x}$ to $\mathbf{y}$ at a fixed time, Lemma 8.9 below, and then appealing to Hölder continuity of optimizers. We will then compare the case of general $\mathbf{x} = (x_1 < x_2), \mathbf{y} = (y_1 < y_2)$ to the common endpoint case by applying a resampling argument to show that optimizers from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$ can be approximated in overlap by optimizers from $((0, 0), s - 1)$ to $((0, 0), t + 1)$ in a sequence of extended landscapes defined with different, independent noise on $[s - 1, s]$ and $[t, t + 1]$.

LEMMA 8.9. *Let $\pi = (\pi_1, \pi_2) : [0, 1] \rightarrow \mathbb{R}_{\leq}^2$ be the almost surely unique optimizer from $((0, 0), 0)$ to $((0, 0), 1)$ in $\mathcal{L}^*$. Take any $\delta, \eta, d > 0$. There exists $c > 0$ depending on δ, d, η , such that for any $t \in [\delta, 1 - \delta]$, and $\epsilon > 0$, we have*

$$\mathbb{P}(|\pi_1(t) - \pi_2(t)| < \epsilon, |\pi_1(t)|, |\pi_2(t)| < d) < c\epsilon^{2-\eta}.$$

The basic strategy for the proof of Lemma 8.9 is to use Proposition 5.8 to relate the joint distribution of $\pi_1(t), \pi_2(t)$ to a certain optimization problem across the first two lines of two parabolic Airy line ensembles. This optimization problem is amenable to analysis since the parabolic Airy line ensemble withstands a strong comparison to independent Brownian motions, and is further reduced to the same optimization problem for several independent Brownian motions (see Lemma 8.11 below).

It is worth mentioning that exponents for disjointness of *geodesics* in the directed landscape have previously been analyzed by Hammond [27] by similarly appealing to a particular statistic in the parabolic Airy line ensemble. Hammond showed that the probability that k geodesics from time 0 to time 1 that all start and end within ϵ of each other are mutually disjoint is bounded above by $\epsilon^{(k^2-1)/2+o(1)}$. This exponent is expected to be sharp. One thing that is new in the analysis in Lemma 8.9 is the use of the precise relationship between the Airy sheet and the parabolic Airy line ensemble, which was not yet known when [27] was written.

The optimization problem (Lemma 8.11) that arises in Lemma 8.9 is not straightforward to analyze, even heuristically. Indeed, we do not expect that the exponent 2 in Lemma 8.9 is sharp and do not have a prediction for the true exponent. To set up the proof, we need the following result giving a strong comparison between the parabolic Airy line ensemble $\mathcal{B}$ on a compact set and a sequence of independent Brownian motions.

THEOREM 8.10. *For $d > 0$, let $\mathcal{C}_d$ be the space of continuous functions on $[-d, d]$ which vanish at $-d$. Let μ_d denote the law of a standard Brownian motion on $[-d, d]$, and for $k \in \mathbb{N}$ let $\mu_d^{\otimes k}$ denote the law of k -tuples of functions in $\mathcal{C}_d^k$ given by the product of k copies of μ_d . For any measurable set $A \subset \mathcal{C}_d^k, k \in \mathbb{N}$ and $d \geq 1$ we have*

$$\mathbb{P}(\hat{\mathcal{B}}^k \in A) \leq \mu_d^{\otimes k}(A) \exp\left(bkd^6 + de^{bk}(\log[\mu_d^{\otimes k}(A)]^{-1})^{5/6}\right),$$

where $b > 0$ is a universal constant, $\hat{\mathcal{B}}^k = (\hat{\mathcal{B}}_1^k, \dots, \hat{\mathcal{B}}_k^k)$, and each $\hat{\mathcal{B}}_i^k$ is given by

$$\hat{\mathcal{B}}_i^k(x) = 2^{-1/2}(\mathcal{B}_i^k(x) - \mathcal{B}_i^k(-d)).$$

The main result in [9] (Theorem 3.11 therein) shows that each of the marginals $\hat{\mathcal{B}}_i^k$ satisfy the above Radon-Nikodym derivative bound with μ_d in place of $\mu_d^{\otimes k}$. While Theorem 8.10 is stronger than [9, Theorem 3.11], it can nonetheless be proven by combining the same key technical ingredients developed in Sections 4 and 5 of [9]. We do this in Appendix A.2.

PROOF OF LEMMA 8.9. We define

$$\tilde{\mathbf{z}} = \arg \max_{-2d \leq z_1 \leq z_2 \leq 2d} \mathcal{L}^*((0, 0), 0; \mathbf{z}, t) + \mathcal{L}^*(\mathbf{z}, t; (0, 0), 1).$$

Since $\mathcal{L}^*$ has extended Airy sheet marginals, by Proposition 5.8 and the symmetry $\mathcal{S}(\mathbf{x}, \mathbf{y}) = \mathcal{S}(-\mathbf{y}, -\mathbf{x})$ (Lemma 5.5), we could alternatively define $\tilde{\mathbf{z}}$ as

$$(8.2) \quad \tilde{\mathbf{z}} = \arg \max_{-2d \leq z_1 \leq z_2 \leq 2d} \mathcal{B}[(0, 0) \rightarrow (z_1, z_2)] + \mathcal{B}'[(0, 0) \rightarrow (-z_2, -z_1)],$$

where $t^{1/3}\mathcal{B}(t^{-2/3} \cdot)$ and $(1-t)^{1/3}\mathcal{B}'((1-t)^{-2/3} \cdot)$ are independent parabolic Airy line ensembles. Uniqueness of the argmax follows the same arguments as in the proof of Lemma 7.3. It suffices to prove

$$\mathbb{P}(|\tilde{z}_1 - \tilde{z}_2| < \epsilon, |\tilde{z}_1|, |\tilde{z}_2| < d) < c\epsilon^{2-\eta},$$

since if $|\pi_1(t)|, |\pi_2(t)| < d$, we must have that $\pi_1(t) = \tilde{z}_1$ and $\pi_2(t) = \tilde{z}_2$. By Proposition 5.9,

$$\begin{aligned} \mathcal{B}[(0,0) \rightarrow (z_1, z_2)] &= \max_{z_1 \leq w \leq z_2} \mathcal{B}_1(z_1) + \mathcal{B}_1(z_2) - \mathcal{B}_1(w) + \mathcal{B}_2(w), \\ \mathcal{B}'[(0,0) \rightarrow (-z_2, -z_1)] &= \max_{-z_2 \leq w \leq -z_1} \mathcal{B}'_1(-z_2) + \mathcal{B}'_1(-z_1) - \mathcal{B}'_1(w) + \mathcal{B}'_2(w). \end{aligned}$$

Therefore by Theorem 8.10 applied to the interval $[-2\delta^{-2/3}d, 2\delta^{-2/3}d]$ and Brownian scaling and time-reversal symmetry of Brownian motion, it suffices to study the same problem when $\mathcal{B}_1(\cdot), \mathcal{B}_2(\cdot), \mathcal{B}'_1(\cdot), \mathcal{B}'_2(\cdot)$ are replaced by independent Brownian motions. This is done in Lemma 8.11, implying the desired result. $\square$

LEMMA 8.11. *Take four independent two-sided Brownian motions $B_1, B_2, B'_1, B'_2 : \mathbb{R} \rightarrow \mathbb{R}$, with diffusion parameter 2. Let $(\tilde{z}_1, \tilde{z}_2, \tilde{w}, \tilde{w}')$ be*

$$\begin{aligned} \arg \max_{-2 \leq z_1 \leq w, w' \leq z_2 \leq 2} & (B_1(z_1) + B_1(z_2) - B_1(w) + B_2(w)) \\ & + (B'_1(z_1) + B'_1(z_2) - B'_1(w') + B'_2(w')). \end{aligned}$$

Note that a priori we do not assume the argmax is unique, and just take an arbitrary one. Then given any small $\eta > 0$, for any small enough $\epsilon > 0$ we have $\mathbb{P}(|\tilde{z}_1|, |\tilde{z}_2| < 1, |\tilde{z}_1 - \tilde{z}_2| < \epsilon) < \epsilon^{2-\eta}$.

While now we only work with Brownian motions, analyzing the argmax formula is still involved. Our general strategy is to simplify the problem by restricting the choice of parameters (for example, we will usually take $w = w'$). We expect that these simplifications do not capture the full picture and so we do not expect the exponent of $2 - \eta$ to be sharp.

PROOF OF LEMMA 8.11. Throughout the proof we assume that $\eta > 0$ is small and that $\epsilon > 0$ is sufficiently small given η .

Step 1: Conditioning on the location of the argmax. We first split $[-1, 1]$ into $\lceil \epsilon^{-1} \rceil$ intervals, each of length at most 2ϵ . We just need to show that, for each interval I , we have $\mathbb{P}(\tilde{z}_1, \tilde{z}_2 \in I) < \epsilon^{3-\eta}$.

Let $F(z_1, z_2, w, w')$ denote the function inside the argmax. Denote the center of I by z_I , and $F_I = F(z_I, z_I, z_I, z_I)$. If $\tilde{z}_1, \tilde{z}_2 \in I$, then one of $\mathcal{E}$ and $\mathcal{E}_1 \cap \mathcal{E}_2$ happens, where

$$\begin{aligned} \mathcal{E} : & \max_{z_1 \leq w, w' \leq z_2, z_1, z_2 \in I} F(z_1, z_2, w, w') \geq F_I + \epsilon^{1/2-\eta/10}, \\ \mathcal{E}_1 : & \max_{z_I \leq w \leq z_2 \leq 2} F(z_I, z_2, w, w) < F_I + \epsilon^{1/2-\eta/10}, \\ \mathcal{E}_2 : & \max_{-2 \leq z_1 \leq w \leq z_I} F(z_1, z_I, w, w) < F_I + \epsilon^{1/2-\eta/10}. \end{aligned}$$

We have taken $w' = w$ in defining $\mathcal{E}_1$ and $\mathcal{E}_2$ to simplify arguments below. From the tail of Brownian motions in an interval of length 2ϵ we have $\mathbb{P}(\mathcal{E}) < e^{-\epsilon^{-\eta/6}}$. Also note that $\mathcal{E}_1$ and $\mathcal{E}_2$ are independent, since they depend only on $B_1 + B'_1 - B_1(z_I) - B'_1(z_I)$ and $B_2 + B'_2 - B_2(z_I) - B'_2(z_I)$, to the right and left of z_I , respectively. Thus

$\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2) = \mathbb{P}(\mathcal{E}_1)\mathbb{P}(\mathcal{E}_2)$. It remains to show that $\mathbb{P}(\mathcal{E}_1) < \epsilon^{3/2-\eta/3}$, since similarly we will also have $\mathbb{P}(\mathcal{E}_2) < \epsilon^{3/2-\eta/3}$.

Step 2: Reducing to two Brownian motions. To bound $\mathbb{P}(\mathcal{E}_1)$, we rewrite $F(z_I, z_2, w, w)$ using two independent Brownian motions. Consider two processes on $[0, 3]$, defined as $\tilde{B}_1(z) = 2^{-1}(B_1(z_I + z) + B'_1(z_I + z) - B_1(z_I) - B'_1(z_I))$ and $\tilde{B}_2(z) = 2^{-1}(B_2(z_I + z) + B'_2(z_I + z) - B_2(z_I) - B'_2(z_I))$, respectively. These are two independent standard Brownian motions. Letting $h = 2^{-1}\epsilon^{1/2-\eta/10}$, we have

$$(8.3) \quad \mathbb{P}(\mathcal{E}_1) \leq \mathbb{P}\left(\max_{0 \leq w \leq z \leq 1} \tilde{B}_1(z) - \tilde{B}_1(w) + \tilde{B}_2(w) < h\right).$$

Now let $\hat{B}(w) = \max\{\tilde{B}_1(z) - \tilde{B}_1(1-w) : 1-w \leq z \leq 1\}$. By [36, Theorem 2.34] and the independence of $\tilde{B}_1$ and $\tilde{B}_2$, we have $(\hat{B}, \tilde{B}_2) \stackrel{d}{=} (|B|, \tilde{B}_2)$ on $[0, 1]$, where B is another Brownian motion, also independent of $\tilde{B}_2$. Therefore the right-hand side of (8.3) equals

$$(8.4) \quad \mathbb{P}\left(\max_{w \in [0,1]} |B(1-w)| + \tilde{B}_2(w) < h\right).$$

Our goal is to show that (8.4) is $O(h^3)$.

Step 3: Computation using the reflection principle. Up to a constant factor, (8.4) further equals

$$(8.5) \quad \iint \mathbb{P}\left(\max_{w \in [0,1]} |B(1-w)| + \tilde{B}_2(w) < h \mid B(1) = a, \tilde{B}_2(1) = b\right) e^{-(a^2+b^2)/2} da db.$$

Conditioned on $B(1) = a, \tilde{B}_2(1) = b$, the processes $B(1-w) - a(1-w)$ and $\tilde{B}_2(w) - bw$ are independent Brownian bridges. Thus we can write the probability in (8.5) as

$$\mathbb{P}\left(\max_{w \in [0,1]} \max\{G_1(w) + G_2(w) + a(1-w) + bw, -G_1(w) + G_2(w) - a(1-w) + bw\} < h\right),$$

where $G_1, G_2 : [0, 1] \rightarrow \mathbb{R}$ are two independent Brownian bridges. Using that $H_1 := 2^{-1/2}(G_1 + G_2)$ and $H_2 := 2^{-1/2}(G_1 - G_2)$ are independent Brownian bridges, this probability can be further written as

$$\mathbb{P}\left(\max_{w \in [0,1]} \sqrt{2}H_1(w) + a(1-w) + bw < h\right) \mathbb{P}\left(\max_{w \in [0,1]} \sqrt{2}H_2(w) - a(1-w) + bw < h\right).$$

These two probabilities can be computed using the reflection principle (see e.g. [36, Theorem 2.19]). The first one equals (for B being a standard Brownian motion on $[0, 1]$)

$$\begin{aligned} & \mathbb{P}\left(\max_{w \in [0,1]} B(w) < 2^{-1/2}(h-a) \mid B(1) = 2^{-1/2}(b-a)\right) \\ &= \frac{\mathbb{P}(B(1) = 2^{-1/2}(b-a)) - \mathbb{P}(B(1) = 2^{-1/2}(2h-a-b))}{\mathbb{P}(B(1) = 2^{-1/2}(b-a))} \\ &= e^{(a-b)^2/4} (e^{-(a-b)^2/4} - e^{-(2h-a-b)^2/4}), \end{aligned}$$

and similarly for the second one. Therefore we can write (8.5) as

$$\begin{aligned}
& \iint_{|a|, b \leq h} e^{-(a^2+b^2)/2} e^{(a-b)^2/4} (e^{-(a-b)^2/4} - e^{-(2h-a-b)^2/4}) \\
& \quad \times e^{(-a-b)^2/4} (e^{-(-a-b)^2/4} - e^{-(2h+a-b)^2/4}) dadb \\
& = \iint_{|a|, b \leq h} e^{-(a^2+b^2)/2} (1 - e^{-(h-a)(h-b)}) (1 - e^{-(h+a)(h-b)}) dadb \\
& < \iint_{|a|, b \leq h} e^{-(a^2+b^2)/2} (h-a)(h+a)(h-b)^2 dadb < 2h^3 \int_{b \leq h} e^{-b^2/2} (h-b)^2 db.
\end{aligned}$$

We note that the integral in the last line is uniformly bounded for $h < 1$. Thus we conclude that $\mathbb{P}(\mathcal{E}_1) < e^{3/2-\eta/3}$, and our conclusion follows. $\square$

Before moving to the proof of Lemma 8.8, we need one more result.

LEMMA 8.12. *Let $s < t$, and let $\mathcal{F}$ denote the σ -algebra generated by $\mathcal{L}^*$ restricted to time increments $[r, r'] \subset [s, t]$. Let π denote the almost surely unique optimizer from $((0, 0), s-1)$ to $((0, 0), t+1)$. Then the conditional law of $(\pi(s), \pi(t))$ given $\mathcal{F}$ almost surely has full support $\mathbb{R}_{\leq}^2 \times \mathbb{R}_{\leq}^2$.*

PROOF. Let $(\mathbf{x}_*, \mathbf{y}_*) = \arg \max_{(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in \mathbb{R}_{\leq}^2} F(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$, where

$$F(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) = \mathcal{L}^*((0, 0), s-1; \tilde{\mathbf{x}}, s) + \mathcal{L}^*(\tilde{\mathbf{x}}, s; \tilde{\mathbf{y}}, t) + \mathcal{L}^*(\tilde{\mathbf{y}}, t; (0, 0), t+1).$$

Then $(\pi(s), \pi(t)) = (\mathbf{x}^*, \mathbf{y}^*)$ and by Lemma 7.2, the argmax is almost surely unique. Now, the outer two functions are independent of $\mathcal{F}$. Moreover, by Proposition 5.9, we have

$$(8.6) \quad \mathcal{L}^*((0, 0), s-1; \tilde{\mathbf{x}}, s) = \max_{\tilde{x}_1 \leq w \leq \tilde{x}_2} \mathcal{B}_1(\tilde{x}_1) + \mathcal{B}_1(\tilde{x}_2) - \mathcal{B}_1(w) + \mathcal{B}_2(w),$$

where $\mathcal{B}$ is a parabolic Airy line ensemble. A similar decomposition exists for $\mathcal{L}^*(\tilde{\mathbf{y}}, t; (0, 0), t+1)$ in terms of an independent parabolic Airy line ensemble $\mathcal{B}'$. Now let $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}_{\leq}^2 \times \mathbb{R}_{\leq}^2$. Conditionally on $\mathcal{F}$, we can apply the Brownian Gibbs property to resample the first two lines of $\mathcal{B}, \mathcal{B}'$ on an interval $[-m, m]$ containing x_1, x_2, y_1, y_2 . Let F' denote the analogue of the original function F after resampling. By (8.6), for any $M, \delta > 0$, with positive probability we have

$$F'(\mathbf{x}, \mathbf{y}) - F(\mathbf{x}, \mathbf{y}) > M,$$

$$|F(\mathbf{u}) - F'(\mathbf{u})| \leq \delta \text{ for all } \mathbf{u} \text{ such that } \|\mathbf{u} - (\mathbf{x}, \mathbf{y})\|_2 > \delta.$$

Since F achieves its argmax, this implies that F' can achieve its argmax arbitrarily close to $(\mathbf{x}, \mathbf{y})$. Since $F \stackrel{d}{=} F'$, this gives the result. $\square$

PROOF OF LEMMA 8.8. Let $\pi = (\pi_1, \pi_2)$ be the optimizer from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$. By Lemma 7.2 we assume that it is the unique one.

Step 1. We first prove the case where $x_1 = x_2, y_1 = y_2$. By the symmetries of $\mathcal{L}^*$ (Lemma 6.10) we may assume $x_1 = x_2 = y_1 = y_2 = 0$, and $s = 0, t = 1$.

Fix some small δ with $0 < \delta < 1$. Take a large $N \in \mathbb{N}$, and let $t_i = \delta + (1 - 2\delta)i/N$ for $i = 0, \dots, N$. By Lemma 8.9, for any fixed d and η there is some constant $c > 0$ such that

$$\mathbb{P}(\exists i, |\pi_1(t_i) - \pi_2(t_i)| < N^{\eta-2/3}, |\pi_1(t_i)|, |\pi_2(t_i)| < d) < cN^{3\eta-1/3}.$$

By Lemma 6.8, each π_i is Hölder $2/3^-$. Therefore taking $N \rightarrow \infty$ we have

$$\mathbb{P}(\exists t' \in [\delta, 1 - \delta], \pi_1(t') = \pi_2(t'), |\pi_1(t')| < d) = 0.$$

Since d and δ are arbitrary, we have $\pi_1(t') \neq \pi_2(t'), \forall t' \in (0, 1)$.

Step 2. Now we prove the general case by a resampling argument.

Let π' be the optimizer from $((0, 0), s - 1)$ to $((0, 0), t + 1)$, which is assumed to be unique by Lemma 7.2. Setting

$$(8.7) \quad (\mathbf{x}_*, \mathbf{y}_*) = \arg \max_{(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in \mathbb{R}_\leq^2} \mathcal{L}^*((0, 0), s - 1; \tilde{\mathbf{x}}, s) + \mathcal{L}^*(\tilde{\mathbf{x}}, s; \tilde{\mathbf{y}}, t) + \mathcal{L}^*(\tilde{\mathbf{y}}, t; (0, 0), t + 1),$$

then π' is the concatenation of the optimizers from $((0, 0), t - 1)$ to $(\mathbf{x}_*, t)$, from $(\mathbf{x}_*, t)$ to $(\mathbf{y}_*, s)$, and from $(\mathbf{y}_*, s)$ to $((0, 0), s + 1)$. Each of these three optimizers must be unique, otherwise π' is not unique.

Now we take a series of independent samples of $\mathcal{L}^*$, denoted as $\mathcal{L}^{*,i}$ for $i \in \mathbb{N}$. Using these samples, we can define landscapes $\hat{\mathcal{L}}^{*,i}$ by setting $\hat{\mathcal{L}}^{*,i}(\cdot, r; \cdot, r')$ equal to $\mathcal{L}^{*,i}$ when $[r, r'] \subset (s, t)^c$ and equal to $\mathcal{L}^*$ when $[r, r'] \subset [s, t]$. Defining $\hat{\mathcal{L}}^{*,i}$ at all other time increments via metric composition yields an extended landscape.

We denote by $\pi^{(i)}$ the optimizer from $((0, 0), s - 1)$ to $((0, 0), t + 1)$ in $\hat{\mathcal{L}}^{*,i}$, and define $(\mathbf{x}_*^{(i)}, \mathbf{y}_*^{(i)})$ as in (8.7) with $\hat{\mathcal{L}}^{*,i}$ in place of $\mathcal{L}^*$, so that arguing as before, $\pi^{(i)}$ is a concatenation of the unique optimizer from $((0, 0), s - 1)$ to $(\mathbf{x}_*^{(i)}, s)$ in $\mathcal{L}^{*,i}$, the unique optimizer from $(\mathbf{x}_*^{(i)}, s)$ to $(\mathbf{y}_*^{(i)}, t)$ in $\mathcal{L}^*$, and the unique optimizer from $(\mathbf{y}_*^{(i)}, t)$ to $((0, 0), t + 1)$ in $\mathcal{L}^{*,i}$. In addition, from the first step, we have that each $\pi^{(i)}$ consists of disjoint paths, except for the endpoints.

Conditioned on $\mathcal{L}^*$, for any fixed $\mathbf{x}, \mathbf{y} \in \mathbb{R}_\leq^2$ and any $\epsilon > 0$, by Lemma 8.12 there is a positive probability that $x_{*,j}^{(i)} > x_j$ and $y_{*,j}^{(i)} > y_j$ for all $j \in \llbracket 1, k \rrbracket$, and that $\|\mathbf{x}_*^{(i)} - \mathbf{x}\|_2 < \epsilon$ and $\|\mathbf{y}_*^{(i)} - \mathbf{y}\|_2 < \epsilon$. Thus almost surely, we can find a sequence $i_1 < i_2 < \dots$, such that $\mathbf{x}_*^{(i_\ell)} \rightarrow \mathbf{x}$ and $\mathbf{y}_*^{(i_\ell)} \rightarrow \mathbf{y}$ as $\ell \rightarrow \infty$, and $x_{*,j}^{(i_\ell)} > x_j$ and $y_{*,j}^{(i_\ell)} > y_j$ for all ℓ . Then by Lemma 8.6, the optimizer from $(\mathbf{x}_*^{(i_k)}, s)$ to $(\mathbf{y}_*^{(i_k)}, t)$ converges to the rightmost optimizer from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$, in the overlap topology. Since for each i , the optimizer from $(\mathbf{x}_*^{(i)}, s)$ to $(\mathbf{y}_*^{(i)}, t)$ consists of disjoint paths, the optimizer from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$ must also consist of disjoint paths, except possibly at the endpoints. $\square$

We upgrade Lemma 8.8 to all endpoints simultaneously.

LEMMA 8.13. *Almost surely the following statement is true. For any $\mathbf{x}, \mathbf{y} \in \mathbb{R}_\leq^2$ and $s < t$, there exists an optimizer in the extended landscape from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$, that consists of two paths that are disjoint, except possibly at the endpoints.*

PROOF. By Lemma 7.2 and Lemma 8.8, almost surely for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}_\leq^2 \cap \mathbb{Q}^2$ and $s < t \in \mathbb{Q}$ the above statement is true, and there is a unique optimizer from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$. For any general $\mathbf{x}, \mathbf{y} \in \mathbb{R}_\leq^2$ and $s < t$, we take a sequence $\mathbf{x}^{(i)}, \mathbf{y}^{(i)} \in \mathbb{R}_\leq^2$ and $s^{(i)} < t^{(i)}$ consisting of rational numbers, and satisfying the following conditions:

- $s^{(i)} < s < t < t^{(i)}$, and $s^{(i)} \rightarrow s, t^{(i)} \rightarrow t$ as $i \rightarrow \infty$;
- for each $1 \leq j \leq k$ and $i \in \mathbb{N}$, $x_j^{(i)} - x_j > (s - s^{(i)})^{1/5}$, $y_j^{(i)} - y_j > (t^{(i)} - t)^{1/5}$;
- $\mathbf{x}^{(i)} \rightarrow \mathbf{x}, \mathbf{y}^{(i)} \rightarrow \mathbf{y}$ as $i \rightarrow \infty$.

Let $\pi^{(i)}$ be the unique optimizer from $(\mathbf{x}^{(i)}, s^{(i)})$ to $(\mathbf{y}^{(i)}, t^{(i)})$. By Lemma 6.8, for i large enough we have $\pi^{(i)}(s) > \mathbf{x}$ and $\pi^{(i)}(t) > \mathbf{y}$, while $\pi^{(i)}(s) \rightarrow \mathbf{x}$ and $\pi^{(i)}(t) \rightarrow \mathbf{y}$

as $i \rightarrow \infty$. Then by Lemma 8.6, as $i \rightarrow \infty$ the (unique) optimizer from $(\pi^{(i)}(s), s)$ to $(\pi^{(i)}(t), t)$ converges to the rightmost optimizer from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$ in the overlap topology. This means that the rightmost optimizer from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$ consists of two paths that are disjoint, except possibly at the endpoints. $\square$

8.3. Multiple paths

We now extend from the two-path case to the general k -path case. The basic idea is that if two adjacent paths in k -path optimizer overlap at a point (z, r) , then a similar situation must occur with two paths in a two-path optimizer whose endpoints are chosen close to (z, r) .

PROOF OF PROPOSITION 8.1. We show that for each $k \geq 2$, almost surely, for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$ and $r < t$, there exists an optimizer from $(\mathbf{x}, r)$ to $(\mathbf{y}, t)$ consisting of paths that are disjoint, except possibly at the endpoints. We prove this by induction on k . The $k = 2$ case is Lemma 8.13. Now suppose that $k > 2$ and that the statement is true for $k - 1$.

We first prove the fixed endpoint version, i.e., for any fixed $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$ and $r < t$, almost surely the unique optimizer π from $(\mathbf{x}, r)$ to $(\mathbf{y}, t)$ is disjoint, except possibly at the endpoints.

Take $\mathbf{x}', \mathbf{y}' \in \mathbb{R}_{\leq}^{k-1}$ consisting of the first $k - 1$ coordinates of $\mathbf{x}, \mathbf{y}$, respectively. We assume that the optimizer from $(\mathbf{x}', r)$ to $(\mathbf{y}', t)$ is unique, and denote it as π' . We then have that almost surely, these optimizers interlace; i.e. for each $1 \leq i \leq k - 1$ and $r \leq s \leq t$ we have

$$\pi_i(s) \leq \pi'_i(s) \leq \pi_{i+1}(s).$$

This follows from Lemma 7.7. Indeed, by Lemma 6.8 one can find large enough $x'', y'' \in \mathbb{R}$, such that the geodesic from (x'', r) to (y'', t) (denoted as π'') is disjoint from π' . Then (π', π'') is a optimizer from $((\mathbf{x}', x''), r)$ to $((\mathbf{y}', y''), t)$, by Proposition 7.8; and by Lemma 7.7 applied to π and (π', π'') the first inequality is obtained. The second inequality follows similarly by taking x'', y'' sufficiently negative.

By the inductive hypothesis, π' consists of paths that are disjoint, except possibly at the endpoints. This implies that π_i and π_{i+2} are disjoint except possibly at the endpoints, for each $1 \leq i \leq k - 2$.

Now suppose that $\pi_i(s) = \pi_{i+1}(s)$ for some $r < s < t$ and $1 \leq i \leq k - 1$. Then $\pi_{i-1}(s)$ (if $i > 1$) and $\pi_{i+2}(s)$ (if $i + 1 < k$) are different from $\pi_i(s) = \pi_{i+1}(s)$. Then there exists some $\epsilon > 0$, such that

$$\begin{aligned} \max_{s' \in [s-\epsilon, s+\epsilon]} \pi_{i-1}(s') + \epsilon &< \min_{s' \in [s-\epsilon, s+\epsilon]} \pi_i(s'), \\ \min_{s' \in [s-\epsilon, s+\epsilon]} \pi_{i+2}(s') - \epsilon &> \max_{s' \in [s-\epsilon, s+\epsilon]} \pi_{i+1}(s'). \end{aligned}$$

Now for small enough $\delta > 0$, let $\mathbf{x}^\delta = (\pi_i(s - \delta), \pi_{i+1}(s - \delta))$ and let $\mathbf{y}^\delta = (\pi_i(s + \delta), \pi_{i+1}(s + \delta))$. By Lemma 8.13, we can find an optimizer π^δ from $(\mathbf{x}^\delta, s - \delta)$ to $(\mathbf{y}^\delta, s + \delta)$ with $\pi_1^\delta(s) < \pi_2^\delta(s)$; in particular, $\pi^\delta \neq (\pi_i, \pi_{i+1})$. By Lemma 6.8, for small enough $\delta > 0$, the optimizer π^δ is disjoint from π_{i-1} (if $i > 1$) and π_{i+2} (if $i + 1 < k$). Therefore letting τ^δ denote π with π^δ in place of (π_i, π_{i+1}) on the interval $[s - \delta, s + \delta]$, Proposition 7.8 ensures that $\|\tau^\delta\|_{\mathcal{L}^*} \geq \|\pi\|_{\mathcal{L}^*}$. Thus this new path is also an optimizer from $(\mathbf{x}, r)$ to $(\mathbf{y}, t)$, contradicting the uniqueness assumption.

To upgrade this to hold for all endpoints simultaneously, we use the arguments in the proof of Lemma 8.13, essentially verbatim. $\square$

PROOF OF THEOREMS 1.6 AND 1.7 AND COROLLARIES 1.9 AND 1.11. First, we can couple $\mathcal{L}^*, \mathcal{L}$ so that $\mathcal{L}^*|_{\mathbb{R}_\uparrow^4} = \mathcal{L}|_{\mathbb{R}_\uparrow^4}$. By Proposition 7.8, for any multi-path $\pi : [s, t] \rightarrow \mathbb{R}_{\leq}^k$ with $\pi_i(r) < \pi_{i+1}(r)$ for all $r \in (s, t)$, we have

$$(8.8) \quad \|\pi\|_{\mathcal{L}^*} = \sum_{i=1}^k \|\pi_i\|_{\mathcal{L}^*} = \sum_{i=1}^k \|\pi_i\|_{\mathcal{L}}.$$

Moreover, almost surely for all $(\mathbf{x}, s; \mathbf{y}, t) \in \mathfrak{X}_\uparrow$, Proposition 8.1 guarantees that $\mathcal{L}^*(\mathbf{x}, s; \mathbf{y}, t) = \sup_{\pi} \|\pi\|_{\mathcal{L}^*}$, where the supremum is over all multi-paths π from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$ that are disjoint away from the endpoints. Comparing this with Definition 1.1 gives that $\mathcal{L}^* = \mathcal{L}$, proving Theorem 1.6. Theorem 1.7 then follows from (8.8), Proposition 8.1, and Lemma 7.2. Corollary 1.9 follows from (1.8) and Theorem 1.6. For Corollary 1.11, (1.12) follows from the existence of disjoint geodesics by definition. The opposite direction uses Theorem 1.7. $\square$

CHAPTER 9

Convergence of Optimizers

In this chapter we prove Theorem 1.8, which shows that disjoint optimizers in Brownian LPP converge to disjoint optimizers in $\mathcal{L}$. The convergence for geodesics was shown in [15]. The argument in [15] is purely deterministic, relying only the metric composition law for $\mathcal{L}$ and a few basic regularity properties. We will adopt a similar strategy here. In this chapter, we will work in a coupling where the following conditions hold on some set Ω of probability 1.

- (i) $\mathcal{L}_n \rightarrow \mathcal{L}$ uniformly on compact subsets of $\mathfrak{X}_\uparrow$.
- (ii) For every bounded set $K = [-b, b]^4 \cap \mathbb{R}_\uparrow^4$, there exists some finite C_b such that for all $\epsilon \in (0, 1)$ we have

$$\limsup_{n \rightarrow \infty} \sup_{(x, s; y, t) \in K} \mathcal{L}_n(x, s; y, t) + \frac{(x - y)^2}{t - s + \epsilon} \leq C_b.$$

- (iii) For any $\eta > 0$, there is a constant $R > 0$ such that

$$\left| \mathcal{L}(x, s; y, t) + \frac{(x - y)^2}{t - s} \right| \leq R(t - s)^{1/3} G(x, s; y, t)^\eta.$$

Here the function G is as in Lemma 6.7.

The fact that such a coupling exists follows from Theorem 1.5 for the first statement, [15, Lemma 13.3] for the second statement, and Lemma 6.7 for the third statement (or alternately, [15, Corollary 10.7]). We let $B^n = (B_i^n : i \in \mathbb{Z})$ denote the collection of standard Brownian motions that give rise to $\mathcal{L}_n$ in this coupling. We work on Ω for all statements and proofs in this chapter.

Most of this chapter is focused on proving Hausdorff convergence of rescaled zigzag graphs; we translate to the language of Theorem 1.8 at the end. For a path $\pi : [a, b] \rightarrow \mathbb{Z}$, recall from Chapter 2 that its zigzag graph is

$$\Gamma(\pi) = \{(c, y) \in \mathbb{R} \times \mathbb{Z} : c \in [a, b], \pi(r) \leq y \leq \pi(r^-)\}.$$

Note that we write $\pi(r^-)$ for the left-hand limit at r , and that $\pi(r^-)$ is always defined, see Chapter 2. Also let A_n be the linear transformation of $\mathbb{R}^2$ given by the matrix

$$(9.1) \quad A_n = \begin{bmatrix} n^{1/3}/2 & n^{-2/3}/2 \\ 0 & -n^{-1} \end{bmatrix}.$$

For any path π , its transformed zigzag graph $A_n \Gamma(\pi)$ is contained in $\mathbb{R} \times n^{-1}\mathbb{Z}$. Moreover, the restriction $A_n|_{\mathbb{R} \times \mathbb{Z}} : \mathbb{R} \times \mathbb{Z} \rightarrow \mathbb{R} \times n^{-1}\mathbb{Z}$ is the inverse of the map $(x, s) \mapsto (x, s)_n$ used in the construction of $\mathcal{L}_n$ in Theorem 1.5. Therefore for any path π from (a, m) to (b, ℓ) and any $n \in \mathbb{N}$, after tracing through the definitions we

get that

$$(9.2) \quad \begin{aligned} \|\pi\|_{\mathcal{L}_n} &= \inf \sum_{i=1}^k \mathcal{L}_n(p_{i-1}; p_i), & \text{where} \\ \|\pi\|_{\mathcal{L}_n} &:= \|\pi\|_{B^n} + 2\sqrt{n}(b-a) + n^{1/6}(A_n(b, \ell)_1 - A_n(a, m)_1). \end{aligned}$$

Here the infimum is over all finite sequences $p_0, \dots, p_k \subset A_n\Gamma(\pi)$ such that

$$a = (A_n^{-1}p_0)_1 < (A_n^{-1}p_1)_1 < \dots < (A_n^{-1}p_k)_1 = b.$$

Here and in (9.2), $(A_n^{-1}p)_1$ denotes the first coordinate of $A_n^{-1}p$. We begin with a tightness statement for zigzag graphs.

LEMMA 9.1. *Let π_n be a sequence of paths from (a_n, m_n) to (b_n, ℓ_n) such that*

$$(9.3) \quad A_n(a_n, m_n) \rightarrow (x, r) \quad \text{and} \quad A_n(b_n, \ell_n) \rightarrow (y, t)$$

as $n \rightarrow \infty$. Suppose also that

$$(9.4) \quad \liminf_{n \rightarrow \infty} \|\pi_n\|_{\mathcal{L}_n} > -\infty$$

almost surely. Then on Ω , the sequence $A_n\Gamma(\pi_n)$ is precompact in the Hausdorff metric. Moreover, any subsequential limit of $A_n\Gamma(\pi_n)$ is equal to $\mathfrak{g}\pi = \{(\pi(s), s) : s \in [r, t]\}$ for some continuous function $\pi : [r, t] \rightarrow \mathbb{R}$ with $\pi(r) = x$ and $\pi(t) = y$.

PROOF. First, let

$$\Gamma'_n(\pi_n) = \{x \in \mathbb{R}^2 : d(x, A_n\Gamma(\pi_n)) \leq n^{-2/3}\}.$$

Here $d(x, A)$ denotes the Euclidean distance between a point and a set. The definitions of A_n and $\Gamma_n(\pi)$ ensure that the sets $\Gamma'_n(\pi_n)$ are all connected. Moreover, the Hausdorff distance $d_H(\Gamma'_n(\pi_n), \Gamma_n(\pi_n))$ is at most $n^{-2/3}$, so it suffices to prove all statements in the lemma for $\Gamma'_n(\pi_n)$. Next, fix an interval $[-b, b] \subset \mathbb{R}$. The definition of the scaling matrix A_n and the limiting statements (9.3) guarantee that $\Gamma'_n(\pi_n) \cap ([-b, b] \times \mathbb{R})$ is precompact in the Hausdorff topology, with subsequential limits contained in $[-b, b] \times [r, t]$.

Take b large enough so that $x, y \in (-b, b)$. Connectedness of the sets $\Gamma'_n(\pi_n)$ implies either there is a subsequential limit of $\Gamma'_n(\pi_n) \cap ([-b, b] \times \mathbb{R})$ that intersects the boundary $\{-b, b\} \times \mathbb{R}$, or else the sequence $\Gamma'_n(\pi_n)$ is precompact, and all subsequential limits are contained in $(-b, b) \times \mathbb{R}$.

Suppose that some subsequential limit of $\Gamma'_n(\pi_n) \cap [-b, b] \times \mathbb{R}$ intersects the boundary $\{-b, b\} \times \mathbb{R}$ at a point $p \in \mathbb{R}^2$. Then there exists a sequence of points $p_n \in A_n\Gamma(\pi_n)$ that converge to p . By the triangle inequality for $\mathcal{L}_n$ and (9.2) we have

$$\|\pi\|_{\mathcal{L}_n} \leq \mathcal{L}_n(A_n(a_n, m_n), p_n) + \mathcal{L}_n(p_n, A_n(b_n, \ell_n)).$$

If $p_n \rightarrow (z, s)$ for some $(z, s) \in \{-b, b\} \times \{r, t\}$, then the right-hand side above converges to $-\infty$ by condition (ii) above, contradicting (9.4). If $p_n \rightarrow (z, s)$ for some $s \in (r, t)$ and $z = \pm b$, then uniform-on-compact convergence of $\mathcal{L}_n$ to $\mathcal{L}$ guarantees that the right-hand side above converges to

$$\mathcal{L}(x, r; z, s) + \mathcal{L}(z, s; y, t).$$

For b large enough, condition (iii) above guarantees that this quantity can become arbitrarily large and negative, contradicting (9.4). Therefore the sequence $\Gamma'_n(\pi_n)$ is precompact, and all subsequential limits are contained in $(-B, B) \times [r, t]$ for some random $B > 0$. Since all subsequential limits of $\Gamma'_n(\pi_n)$ are connected and contain

the points (x, r) and (y, t) , to show that any subsequential limit Γ is of the form $\{(\pi(s), s) : s \in [r, t]\}$ for some continuous function $\pi : [r, t] \rightarrow \mathbb{R}$ with $\pi(r) = x$ and $\pi(t) = y$, we just need to show that Γ intersects each horizontal line at most once.

Suppose that this is not the case, and that $p = (z, s), p' = (z', s) \in \Gamma$ for some $z \neq z'$. Then there are sequences $p_n \in A_n \Gamma(\pi_n)$ and $p'_n \in A_n \Gamma(\pi_n)$ converging to p, p' , respectively. Without loss of generality, we may assume that $(A_n^{-1} p_n)_1 < (A_n^{-1} p'_n)_1$ infinitely often, so that by (9.2) we have

$$\|\pi_n\|_{\mathcal{L}_n} \leq \mathcal{L}_n(A_n(a_n, m_n); p_n) + \mathcal{L}_n(p_n; p'_n) + \mathcal{L}_n(p'_n; A_n(b_n, \ell_n))$$

for infinitely many n . Condition (ii) guarantees that almost surely, the middle term on the right-hand side above converges to $-\infty$, whereas the first and third terms are bounded above. Again, this contradicts (9.4). $\square$

THEOREM 9.2. *Fix $\mathbf{u} = (\mathbf{x}, s; \mathbf{y}, t) \in \mathfrak{X}_\uparrow$, and let $C_{\mathbf{u}}$ be the almost sure set where there is a unique disjoint optimizer π in $\mathcal{L}$ from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$. Let $\pi^{(n)}$ be any sequence of $\mathcal{L}_n$ -optimizers from $(\mathbf{a}_n, m_n)$ to $(\mathbf{b}_n, \ell_n)$ where $A_n(a_{n,i}, m_n) \rightarrow (x_i, s)$ and $A_n(b_{n,i}, \ell_n) \rightarrow (y_i, t)$.*

Then on $\Omega \cap C_{\mathbf{u}}$, $A_n \Gamma(\pi_i^{(n)}) \rightarrow \mathfrak{g}\pi_i = \{(\pi_i(r), r) : r \in [s, t]\}$ in the Hausdorff metric for all i . Moreover, letting $h_{n,i} : [s, t] \rightarrow [a_{n,i}, b_{n,i}]$ be the linear function satisfying $h_{n,i}(s) = a_{n,i}, h_{n,i}(t) = b_{n,i}$, on $\Omega \cap C_{\mathbf{u}}$ we have the uniform convergence

$$\tilde{\pi}_i^{(n)} := \frac{\pi_i^{(n)} \circ h_{n,i} + nh_{n,i}}{2n^{2/3}} \rightarrow \pi_i,$$

as functions from $[s, t]$ to $\mathbb{R}$.

The ‘Moreover’ in Theorem 9.2 is Theorem 1.8.

PROOF OF THEOREM 9.2. In the proof, we work on the set $C_{\mathbf{u}} \cap \Omega$. Let k be such that $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$. Since the $\pi^{(n)}$ are optimizers, and $\mathcal{L}_n \rightarrow \mathcal{L}$ uniformly on compact sets, we have

$$(9.5) \quad \sum_{i=1}^k \|\pi_i^{(n)}\|_{\mathcal{L}_n} = \mathcal{L}_n(A_n(\mathbf{a}_n, m_n), A_n(\mathbf{b}_n, \ell_n)) \rightarrow \mathcal{L}(\mathbf{x}, s; \mathbf{y}, t).$$

Also, for each i , we have

$$(9.6) \quad \|\pi_i^{(n)}\|_{\mathcal{L}_n} \leq \mathcal{L}_n(A_n(a_{n,i}, m_n), A_n(a_{n,i}, \ell_n)).$$

The right-hand side above converges to $\mathcal{L}(x_i, s; y_i, t)$, so for all i , by (9.5) and (9.6), we have

$$\liminf_{n \rightarrow \infty} \|\pi_i^{(n)}\|_{\mathcal{L}_n} \geq \mathcal{L}(\mathbf{x}, s; \mathbf{y}, t) - \sum_{1 \leq j \leq k, j \neq i} \mathcal{L}(x_j, s; y_j, t) > -\infty.$$

Hence by Lemma 9.1, each of the sequences $\{A_n \Gamma(\pi_i^{(n)}) : n \in \mathbb{N}\}$ is precompact, with subsequential limits that are of the form $(\mathfrak{g}\gamma_1, \dots, \mathfrak{g}\gamma_k)$ for some continuous multi-path γ from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$. Now, let $P_n \subset \mathbb{R}_{\leq}^k \times \mathbb{Z}$ be the set of all points $(\mathbf{z}, j)$ such that

$$(9.7) \quad B^n[(\mathbf{a}_n, m_n) \rightarrow (\mathbf{b}_n, \ell_n)] = B^n[(\mathbf{b}_n, m_n) \rightarrow (\mathbf{z}, j)] + B^n[(\mathbf{z}, j-1) \rightarrow (\mathbf{b}_n, \ell_n)],$$

and such that $(z_i, j), (z_i, j-1) \in \Gamma(\pi_i^{(n)})$ for all i . Metric composition (Lemma 2.9) and the fact that $\pi^{(n)}$ is an optimizer guarantees that for every $j \in \{\ell_n + 1, \dots, m_n\}$,

there exists $(\mathbf{z}, j) \in P_n$. In particular, this implies that along a subsequence where $A_n \Gamma(\pi_i^{(n)}) \rightarrow \mathfrak{g} \gamma_i$ for all i , we have

$$A_n P_n \rightarrow \mathfrak{g} \gamma = \{(\gamma(r), r) : r \in [s, t]\}$$

and so (9.7) passes to the limit to give that

$$\mathcal{L}(\mathbf{x}, s; \mathbf{y}, t) = \mathcal{L}(\mathbf{x}, s; \gamma(r), r) + \mathcal{L}(\gamma(r), r; \mathbf{y}, t)$$

for all $r \in (s, t)$. This can only occur if γ is the unique optimizer in $\mathcal{L}$ from $(\mathbf{x}, s)$ to $(\mathbf{y}, t)$, yielding the first part of the theorem.

For the ‘Moreover’, it is enough to show that $\mathfrak{g} \tilde{\pi}_i^{(n)} \rightarrow \mathfrak{g} \pi_i$ for all i in the Hausdorff metric, since Hausdorff convergence of graphs implies uniform convergence of functions when the limit is continuous. For this, by the first part of the theorem we just need to show that the Hausdorff distance $d_H(\mathfrak{g} \tilde{\pi}_i^{(n)}, A_n \Gamma(\pi_i^{(n)}))$ converges to 0 with n .

Since $A_n(a_{n,i}, m_n) \rightarrow (x_i, s)$ and $A_n(b_{n,i}, \ell_n) \rightarrow (y_i, t)$ we have $a_{n,i} \rightarrow s$ and $b_{n,i} \rightarrow t$. Then the function $h_{n,i}$ converges to the identity, so $d_H(\mathfrak{g} \tilde{\pi}_i^{(n)}, \mathfrak{g} \hat{\pi}_i^{(n)}) \rightarrow 0$, where

$$\hat{\pi}_i^{(n)}(x) = \frac{\pi_i^n(x) + nx}{2n^{2/3}}.$$

Moreover, letting $\Lambda \pi_i^{(n)} = \{(c, \pi_i^{(n)}(c)) : c \in [a, b]\}$ denote the graph of $\pi_i^{(n)}$, the first part of the theorem guarantees that $d_H(A_n \Lambda \pi_i^{(n)}, A_n \Gamma \pi_i^{(n)}) \rightarrow 0$. Therefore it suffices to show that $d_H(A_n \Lambda \pi_i^{(n)}, \mathfrak{g} \hat{\pi}_i^{(n)}) \rightarrow 0$ with n . This boils down to a matrix computation. We have $\mathfrak{g} \hat{\pi}_i^{(n)} = D_n \mathfrak{g} \pi_i^{(n)}$ and $\mathfrak{g} \pi_i^{(n)} = R \Lambda \pi_i^{(n)}$, where

$$(9.8) \quad D_n = \begin{bmatrix} n^{-2/3}/2 & n^{1/3}/2 \\ 0 & 1 \end{bmatrix}, \quad R = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Therefore $D_n R A_n^{-1}(A_n \Lambda \pi_i^{(n)}) = \mathfrak{g} \hat{\pi}_i^{(n)}$. A quick computation shows that $D_n R A_n^{-1} \rightarrow I$, yielding the result. $\square$

APPENDIX A

Appendix

A.1. Brownian melon estimates

In this appendix we prove Lemma 5.11 and 6.1, using some Brownian melon estimates from the literature. We start by quoting these results.

THEOREM A.1 ([17, Theorem 3.1]). *There exist positive constants $c_k, d_k, k \in \mathbb{N}$ such that the following holds. For all $m \in (0, 5n^{2/3})$ and $n \geq 1$ we have*

$$\mathbb{P}(W_k^n(1) - 2\sqrt{n} \geq mn^{-1/6}) \leq c_1 e^{-d_1 m^{3/2}},$$

$$\mathbb{P}(W_k^n(1) - 2\sqrt{n} \leq -mn^{-1/6}) \leq c_k e^{-d_k m^3}.$$

Also, for all $m \geq 5n^{2/3}$ and $n \geq 1$ we have

$$\mathbb{P}(|W_k^n(1) - 2\sqrt{n}| \geq mn^{-1/6}) \leq c_1 e^{-d_1 n^{-1/3} m^2}.$$

For any $n \in \mathbb{N}$, $x, a, b, w > 0$, denote

$$\mathcal{N}_{b,w}(n, x, a) = 2\sqrt{nx} + \sqrt{x}n^{-1/6}(a + b \log^{2/3}(n^{1/3}|\log(x/w)| + 1)).$$

Note that for any $\alpha > 0$, we have

$$(A.1) \quad \mathcal{N}_{b,\alpha w}(n, \alpha x, a) = \sqrt{\alpha} \mathcal{N}_{b,w}(n, x, a).$$

PROPOSITION A.2 ([17, Proposition 4.3]). *There exist positive constants b, c and d such that for all $w, a > 0$ and $n \geq 1$, the probability that*

$$W_1^n(x) \leq \mathcal{N}_{b,w}(n, x, a), \quad \forall x \in (0, \infty)$$

is greater than or equal to $1 - ce^{-da^{3/2}}$.

The following estimate is also necessary. This estimate is simply a deterministic inequality and does not involve any probabilistic objects.

LEMMA A.3 ([15, Lemma 9.4]). *Let $b > 0$ be a fixed constant. Then there exists a constant c such that for all $n \in \mathbb{N}, t \in \{1/n, 2/n, \dots, (n-1)/n\}, a > 1$ and*

$$z \in [0, t - c(t \wedge (1-t))^{1/3} a^2 n^{-1/3}] \cup [t + c(t \wedge (1-t))^{1/3} a^2 n^{-1/3}, 1],$$

we have that

$$\mathcal{N}_{b,t}(nt, z, a) + \mathcal{N}_{b,1-t}(n(1-t), 1-z, a) \leq 2\sqrt{n} - an^{-1/6}.$$

We first use Theorem A.1 to deduce an estimate on last passage values across Brownian motions. To clean up the notation in this lemma, its proof, and in the subsequent proof of Lemma 5.11, for a vector $\mathbf{x}$, we let

$$\hat{\mathbf{x}} = (2n^{-1/3}\mathbf{x}, n) \quad \text{and} \quad \tilde{\mathbf{x}} = (1 + 2n^{-1/3}\mathbf{x}, 1).$$

The dependence on n in the notation is implicit.

LEMMA A.4. *Take independent standard Brownian motions $B^n = (B_1^n, \dots, B_n^n)$, and $\mathbf{x}, \mathbf{y} \in \mathbb{R}_{\leq}^k$ such that $\|\mathbf{x}\|_2, \|\mathbf{y}\|_2 < n^{1/6}$. For any $a > 0$ we have*

$$\mathbb{P}\left(\left|B^n[\hat{\mathbf{x}} \rightarrow \tilde{\mathbf{y}}] - \sum_{i=1}^k 2\sqrt{n(1+2n^{-1/3}(y_i - x_i))}\right| > an^{-1/6}\right) < ce^{-da^{3/2}},$$

where c, d are constants depending only on k .

PROOF. By Lemma 2.8 we have

$$\sum_{i=1}^k (B^n[\hat{x}_i^k \rightarrow \tilde{y}_i^k] - B^n[\hat{x}_i^{k-1} \rightarrow \tilde{y}_i^{k-1}]) \leq B^n[\hat{\mathbf{x}} \rightarrow \tilde{\mathbf{y}}] \leq \sum_{i=1}^k B^n[\hat{x}_i \rightarrow \tilde{y}_i].$$

If W^n is the n dimensional Brownian melon, then by Theorem 2.10,

$$B^n[\hat{x}_i \rightarrow \tilde{y}_i] \stackrel{d}{=} W_1^n(1 + 2n^{-1/3}(y_i - x_i)) \quad \text{and} \\ B^n[\hat{x}_i^k \rightarrow \tilde{y}_i^k] - B^n[\hat{x}_i^{k-1} \rightarrow \tilde{y}_i^{k-1}] \stackrel{d}{=} W_k^n(1 + 2n^{-1/3}(y_i - x_i)).$$

Then the conclusion follows from Theorem A.1, using that $\|\mathbf{x}\|_2, \|\mathbf{y}\|_2 < n^{1/6}$, and scale invariance of the Brownian melon: $\sqrt{\alpha}W^n(\cdot) \stackrel{d}{=} W^n(\alpha \cdot)$ for any $\alpha > 0$. $\square$

PROOF OF LEMMA 5.11. Throughout this proof we let c, d denote constants depending on k , whose values may change from line to line. We also assume that n is large enough, since otherwise the conclusion follows by taking c large and d small.

By Lemma A.4, we have

$$(A.2) \quad \mathbb{P}\left(\max_{\mathbf{z} \in \mathbb{R}_{\leq}^k} A(\mathbf{z}) < \sum_{i=1}^k 2\sqrt{n(1+2n^{-1/3}(y_i - x_i))} - an^{-1/6}\right) < ce^{-da^{3/2}}.$$

For each $\mathbf{z} \in \mathbb{R}_{\leq}^k$ for which $A(\mathbf{z})$ is not equal to $-\infty$ (i.e. when it is defined by (5.16)), by Lemma 2.8 we also have

$$(A.3) \quad A(\mathbf{z}) \leq \sum_{i=1}^k (B^n[\hat{x}_i \rightarrow (t + 2n^{-1/3}z_i, q + 1)] + B^n[(t + 2n^{-1/3}z_i, q) \rightarrow \tilde{y}_i]).$$

By Proposition A.2, with probability at least $1 - ce^{-da^{3/2}}$, the i th summand on the right-hand side of (A.3) is bounded above by

$$\mathcal{N}_{b, t(1+2n^{-1/3}(y_i - x_i))}(p, t + 2n^{-1/3}(z_i - x_i), a) \\ + \mathcal{N}_{b, (1-t)(1+2n^{-1/3}(y_i - x_i))}(q, 1 - t + 2n^{-1/3}(y_i - z_i), a),$$

where b is a universal constant. By (A.1) this equals

$$\sqrt{1 + 2n^{-1/3}(y_i - x_i)} \times \left(\mathcal{N}_{b, t}\left(nt, \frac{t + 2n^{-1/3}(z_i - x_i)}{1 + 2n^{-1/3}(y_i - x_i)}, a\right) \right. \\ \left. + \mathcal{N}_{b, (1-t)}\left(n(1-t), \frac{1-t + 2n^{-1/3}(y_i - z_i)}{1 + 2n^{-1/3}(y_i - x_i)}, a\right) \right).$$

Recall that we require $\|\mathbf{x}\|_2, \|\mathbf{y}\|_2 < n^{1/6}$. By Lemma A.3, for any $a > 1$ the above can be bounded by $\sqrt{1 + 2n^{-1/3}(y_i - x_i)}(2\sqrt{n} - an^{-1/6})$, when $|z_i - ty_i - (1-t)x_i| >$

$ca^2(t \wedge (1-t))^{1/3}$. Thus we conclude that, for any $a > 0$, with probability at least $1 - ce^{-da^{3/2}}$ we have

$$A(\mathbf{z}) < \sum_{i=1}^k 2\sqrt{n(1+n^{-1/3}(y_i - x_i))} - an^{-1/6}$$

for any $\mathbf{z}$ with $\|\mathbf{z} - t\mathbf{y} - (1-t)\mathbf{x}\|_2 > ca^2(t \wedge (1-t))^{1/3}$. This with (A.2) finishes the proof. $\square$

Now we complete proving Lemma 6.1, following the outline in Chapter 6.1.

PROOF OF LEMMA 6.1. In this proof we let c, d denote large and small constants depending on k , whose values may change from line to line.

We first upper bound $\mathcal{K}_n(\mathbf{x}, 0; \mathbf{y}', t) - \mathcal{K}_n(\mathbf{x}, 0; \mathbf{y}, 1)$. By the triangle inequality we have

$$\mathcal{L}_n(\mathbf{x}, 0; \mathbf{y}', t) - \mathcal{L}_n(\mathbf{x}, 0; \mathbf{y}, 1) \leq -\mathcal{L}_n(\mathbf{y}', t + n^{-1}; \mathbf{y}, 1).$$

Thus we have

$$\begin{aligned} \mathbb{P}(\mathcal{K}_n(\mathbf{x}, 0; \mathbf{y}', t) - \mathcal{K}_n(\mathbf{x}, 0; \mathbf{y}, 1) > a(1-t)^{1/3}) \\ \leq \mathbb{P}(\mathcal{K}_n(\mathbf{y}', t + n^{-1}; \mathbf{y}, 1) < -a(1-t)^{1/3}) \leq ce^{-da^{3/2}}. \end{aligned}$$

Here the last inequality follows by applying Lemma A.4 to $(1-t)n - 1$ Brownian motions, and elementary calculations. We next lower bound $\mathcal{K}_n(\mathbf{x}, 0; \mathbf{y}', t) - \mathcal{K}_n(\mathbf{x}, 0; \mathbf{y}, 1)$. For any $\mathbf{z} \in \mathbb{R}_{\leq}^k$ we denote $A(\mathbf{z}) = (\mathcal{L}_n(\mathbf{x}, 0; \mathbf{z}, t) - \mathcal{L}_n(\mathbf{x}, 0; \mathbf{y}', t)) + \mathcal{L}_n(\mathbf{z}, t; \mathbf{y}, 1)$. It remains to bound the probability of this event

$$\sup_{\mathbf{z} \in \mathbb{R}_{\leq}^k} A(\mathbf{z}) > -\|\mathbf{y} - \mathbf{x}\|_2^2(1-t) + a(1-t)^{1/3}|\log(1-t)|.$$

To bound $A(\mathbf{z})$, we collect some estimates on $\mathcal{L}_n(\mathbf{z}, t; \mathbf{y}, 1)$ and $\mathcal{L}_n(\mathbf{x}, 0; \mathbf{z}, t) - \mathcal{L}_n(\mathbf{x}, 0; \mathbf{y}', t)$. For this, take any $1 < \hat{a} < n^{1/100}$.

Estimate 1. By Lemma 2.8 we have $\mathcal{L}_n(\mathbf{z}, t; \mathbf{y}, 1) \leq \sum_{i=1}^k \mathcal{L}_n(z_i, t; y_i, 1)$. By Proposition A.2 and using the notation there, for some constant $b > 0$, with probability $> 1 - ce^{-d\hat{a}^{3/2}}$ we have

$$\begin{aligned} (A.4) \quad \mathcal{L}_n(\mathbf{z}, t; \mathbf{y}, 1) &< \sum_{i=1}^k \left[n^{1/6} \mathcal{N}_{b, 1-t+2(y_i - y'_i)n^{-1/3}}((1-t)n, 1-t+2(y_i - z_i)n^{-1/3}, \hat{a}) \right. \\ &\quad \left. - 2(1-t)n^{2/3} - 2n^{1/3}(y_i - z_i) \right], \end{aligned}$$

for any $\mathbf{z} \in \mathbb{R}_{\leq}^k$ such that $1-t+2(y_i - z_i)n^{-1/3} > 0$ for each i . We now give a more explicit bound for the i^{th} summand in the right-hand side of (A.4), when $|z_i - y'_i| < cn^{1/20}$. Note that $\|\mathbf{x}\|_2, \|\mathbf{y}\|_2 < n^{1/100}$, $1-t > n^{-1/100}$, so in this case we would have $(y_i - z_i)n^{-1/3}, (y_i - y'_i)n^{-1/3} < d(1-t)$. Further, recall that

$$\begin{aligned} &\mathcal{N}_{b, 1-t+2(y_i - y'_i)n^{-1/3}}((1-t)n, 1-t+2(y_i - z_i)n^{-1/3}, \hat{a}) \\ &= 2\sqrt{(1-t)n(1-t+2(y_i - z_i)n^{-1/3})} + \sqrt{1-t+2(y_i - z_i)n^{-1/3}} \\ &\quad \times ((1-t)n)^{-1/6} \left(\hat{a} + b \log^{2/3} \left(((1-t)n)^{1/3} \left| \log \left(\frac{1-t+2(y_i - z_i)n^{-1/3}}{1-t+2(y_i - y'_i)n^{-1/3}} \right) \right| + 1 \right) \right). \end{aligned}$$

By Taylor expansion of $y_i - z_i$, we can bound the first term in the right-hand side by

$$2(1-t)\sqrt{n} + 2n^{1/6}(y_i - z_i) - n^{-1/6} \frac{(y_i - z_i)^2}{1-t} + cn^{-1/2} \frac{(y_i - z_i)^3}{(1-t)^2}.$$

For the second term, we use that $\sqrt{1-t+2(y_i - z_i)n^{-1/3}} < c\sqrt{1-t}$, and that

$$\left| \log \left(\frac{1-t+2(y_i - z_i)n^{-1/3}}{1-t+2(y_i - y'_i)n^{-1/3}} \right) \right| < cn^{-1/3} \frac{|y'_i - z_i|}{1-t},$$

to bound it by

$$cn^{-1/6} \hat{a}(1-t)^{1/3} + cn^{-1/6} \frac{|y'_i - z_i|}{(1-t)^{1/3}}.$$

Thus when $|z_i - y'_i| < cn^{1/20}$ we can bound the i^{th} summand in the right-hand side of (A.4) by

$$(A.5) \quad -\frac{(y_i - z_i)^2}{1-t} + c\hat{a}(1-t)^{1/3} + c \frac{|y'_i - z_i|}{(1-t)^{1/3}}.$$

Estimate 2. For $\mathcal{L}_n(\mathbf{x}, 0; \mathbf{z}, t) - \mathcal{L}_n(\mathbf{x}, 0; \mathbf{y}', t)$ we give two different bounds. The first of these bounds $\mathcal{L}_n(\mathbf{x}, 0; \mathbf{z}, t)$ and $\mathcal{L}_n(\mathbf{x}, 0; \mathbf{y}', t)$ separately.

By Lemma 2.8 we have $\mathcal{L}_n(\mathbf{x}, 0; \mathbf{z}, t) \leq \sum_{i=1}^k \mathcal{L}_n(x_i, 0; z_i, t)$. By Proposition A.2 (for $\mathcal{L}_n(\mathbf{x}, 0; \mathbf{z}, t)$) and Lemma A.4 (for $\mathcal{L}_n(\mathbf{x}, 0; \mathbf{y}', t)$), with probability $> 1 - ce^{-d\hat{a}^{3/2}}$ we have

$$(A.6) \quad \begin{aligned} \mathcal{L}_n(\mathbf{x}, 0; \mathbf{z}, t) - \mathcal{L}_n(\mathbf{x}, 0; \mathbf{y}', t) &< \sum_{i=1}^k \left[n^{1/6} \mathcal{N}_{b, t+2(y'_i - x_i)n^{-1/3}}(tn, t+2(z_i - x_i)n^{-1/3}, \hat{a}) \right. \\ &\quad \left. - 2n^{2/3} \sqrt{t(t+2(y'_i - x_i)n^{-1/3})} + \hat{a} \sqrt{t+2(y'_i - x_i)n^{-1/3}t^{-1/6} - 2n^{1/3}(z_i - y'_i)} \right], \end{aligned}$$

for any $\mathbf{z} \in \mathbb{R}_{\leq}^k$ such that $t+2(z_i - x_i)n^{-1/3} > 0$ for each i . Similar to Estimate 1, when $|z_i - y'_i| < cn^{1/20}$, we can bound the i^{th} summand in the right-hand side of (A.6) by

$$(A.7) \quad -\frac{(z_i - x_i)^2}{t} + c\hat{a}t^{1/3} + c \frac{|y'_i - z_i|}{t^{1/3}} + \frac{(y'_i - x_i)^2}{t}.$$

Estimate 3. The second bound for $\mathcal{L}_n(\mathbf{x}, 0; \mathbf{z}, t) - \mathcal{L}_n(\mathbf{x}, 0; \mathbf{y}', t)$ is from the continuity of the prelimiting extended Airy sheet (Lemma 3.5). It is more refined when $\|\mathbf{z} - \mathbf{y}'\|_2$ is small.

By Lemma 3.5 and using Lemma 3.3, we also have that with probability $> 1 - ce^{-d\hat{a}^{3/2}}$,

$$(A.8) \quad \begin{aligned} \mathcal{L}_n(\mathbf{x}, 0; \mathbf{z}, t) - \mathcal{L}_n(\mathbf{x}, 0; \mathbf{y}', t) &< \sum_{i=1}^k \left[\hat{a} \log^{2/3}(2|z_i - y'_i|^{-1}) \sqrt{|z_i - y'_i|} \right] \\ &\quad - \frac{\|\mathbf{z} - \mathbf{x}\|_2^2 - \|\mathbf{y}' - \mathbf{x}\|_2^2}{t}, \end{aligned}$$

for any $\mathbf{z} \in \mathbb{R}_{\leq}^k$ with $\|\mathbf{y}' - \mathbf{z}\|_2 < 1$.

Below we shall bound $A(\mathbf{z})$ assuming that the above three estimates (A.4), (A.6), (A.8) hold.

Upper bound $A(\mathbf{z})$ for $\|\mathbf{y}' - \mathbf{z}\|_2 < 1$. In this case, by (A.5) and (A.8) we have

$$\begin{aligned}
A(\mathbf{z}) &< \sum_{i=1}^k \left[-\frac{(y_i - z_i)^2}{1-t} + c\hat{a}(1-t)^{1/3} + c\frac{|y'_i - z_i|}{(1-t)^{1/3}} \right. \\
&\quad \left. + \hat{a} \log^{2/3}(2|z_i - y'_i|^{-1})\sqrt{|z_i - y'_i|} - \frac{(z_i - x_i)^2 - (y'_i - x_i)^2}{t} \right] \\
&= -(1-t)\|\mathbf{y} - \mathbf{x}\|_2^2 + \sum_{i=1}^k \left[-\frac{(y'_i - z_i)^2}{t(1-t)} + \hat{a} \log^{2/3}(2|z_i - y'_i|^{-1})\sqrt{|z_i - y'_i|} \right. \\
&\quad \left. + c\hat{a}(1-t)^{1/3} + c\frac{|y'_i - z_i|}{(1-t)^{1/3}} \right] \\
&< -(1-t)\|\mathbf{y} - \mathbf{x}\|_2^2 + c\hat{a}^{4/3}(1-t)^{1/3}|\log(1-t)|,
\end{aligned}$$

where the last inequality uses that

$$\frac{(y'_i - z_i)^2}{2t(1-t)} + c^2(1-t)^{1/3} > c\frac{|y'_i - z_i|}{(1-t)^{1/3}},$$

$$\frac{(y'_i - z_i)^2}{2t(1-t)} + c\hat{a}^{4/3}(1-t)^{1/3}|\log(1-t)| > \hat{a} \log^{2/3}(2|z_i - y'_i|^{-1})\sqrt{|z_i - y'_i|}.$$

Upper bound $A(\mathbf{z})$ for $\|\mathbf{y}' - \mathbf{z}\|_2 \geq 1$. In this case we use (A.4) and (A.6). Letting A_i be the sum of the i^{th} term in the right-hand side of (A.4) and (A.6), we have $A(\mathbf{z}) \leq \sum_{i=1}^k A_i$. By (A.1) and Lemma A.3, for each $1 \leq i \leq k$ such that $|z_i - y'_i| > c\left(\frac{\hat{a}}{1-t}\right)^2(1-t)^{1/3}$ we have

$$\begin{aligned}
&\mathcal{N}_{b,1-t+2(y_i-y'_i)n^{-1/3}}((1-t)n, 1-t+2(y_i-z_i)n^{-1/3}, \hat{a}) \\
&\quad + \mathcal{N}_{b,t+2(y'_i-x_i)n^{-1/3}}(tn, t+2(z_i-x_i)n^{-1/3}, \hat{a}) \\
&\leq \sqrt{1+2(y_i-x_i)n^{-1/3}} \left(\mathcal{N}_{b,1-t} \left((1-t)n, \frac{1-t+2(y_i-z_i)n^{-1/3}}{1+2(y_i-x_i)n^{-1/3}}, \frac{\hat{a}}{1-t} \right) \right. \\
&\quad \left. + \mathcal{N}_{b,t} \left(tn, \frac{t+2(z_i-x_i)n^{-1/3}}{1+2(y_i-x_i)n^{-1/3}}, \frac{\hat{a}}{1-t} \right) \right) \\
&\leq \sqrt{1+2(y_i-x_i)n^{-1/3}} \left(2\sqrt{n} - \left(\frac{\hat{a}}{1-t} \right) n^{-1/6} \right).
\end{aligned}$$

So using that $\|\mathbf{x}\|, \|\mathbf{y}\|_2 < n^{1/100}$, $1-t > n^{-1/100}$, we have

$$\begin{aligned}
A_i &\leq 2n^{2/3} \sqrt{1+2(y_i-x_i)n^{-1/3}} - \frac{\hat{a}}{1-t} \sqrt{1+2(y_i-x_i)n^{-1/3}} - 2n^{1/3}(y_i-y'_i) \\
&\quad - 2(1-t)n^{2/3} - 2n^{2/3} \sqrt{t(t+2(y'_i-x_i)n^{-1/3})} + \hat{a} \sqrt{t+2(y'_i-x_i)n^{-1/3}} t^{-1/6} \\
&\leq 2(1-t)n^{2/3} \sqrt{1+2(y_i-x_i)n^{-1/3}} - 2(1-t)n^{1/3}(y_i-x_i) - 2(1-t)n^{2/3} - \frac{d\hat{a}}{1-t} \\
&< -(1-t)(y_i-x_i)^2 + c(1-t)^{1/3} - \frac{d\hat{a}}{1-t}
\end{aligned}$$

When $|z_i - y'_i| \leq c \left(\frac{\hat{a}}{1-t} \right)^2 (1-t)^{1/3} < cn^{1/20}$, using (A.5) and (A.7) we have

$$\begin{aligned} A_i &\leq -\frac{(y_i - z_i)^2}{1-t} + c\hat{a}(1-t)^{1/3} + c\frac{|y'_i - z_i|}{(1-t)^{1/3}} - \frac{(z_i - x_i)^2}{t} \\ &\quad + c\hat{a}t^{1/3} + c\frac{|y'_i - z_i|}{t^{1/3}} + \frac{(y'_i - x_i)^2}{t} \\ &< -\frac{(y'_i - z_i)^2}{t(1-t)} - (1-t)(y_i - x_i)^2 + c\hat{a} + \frac{c|y'_i - z_i|}{(1-t)^{1/3}} \\ &< -\frac{(y'_i - z_i)^2}{2(1-t)} - (1-t)(y_i - x_i)^2 + c\hat{a} + c(1-t)^{1/3}. \end{aligned}$$

Thus by the above two inequalities, and using that $\|\mathbf{y}' - \mathbf{z}\|_2 \geq 1$, we get

$$A(\mathbf{z}) < c\hat{a} + c(1-t)^{1/3} - \frac{d}{1-t} - (1-t)\|\mathbf{y} - \mathbf{x}\|_2^2 < c\hat{a}^{4/3}(1-t)^{1/3} - (1-t)\|\mathbf{y} - \mathbf{x}\|_2^2.$$

Finally, from these bounds on $A(\mathbf{z})$ in each case, we conclude that

$$\begin{aligned} &\mathbb{P}(\mathcal{K}_n(\mathbf{x}, 0; \mathbf{y}, 1) - \mathcal{K}_n(\mathbf{x}, 0; \mathbf{y}', t) > a(1-t)^{1/3} |\log(1-t)|) \\ &< \mathbb{P}(\sup_{\mathbf{z} \in \mathbb{R}_{\leq}^k} A(\mathbf{z}) > -\|\mathbf{y} - \mathbf{x}\|_2^2(1-t) + a(1-t)^{1/3} |\log(1-t)|) < ce^{-da^{9/8}}. \end{aligned}$$

The conclusion follows. $\square$

A.2. Proof of Theorem 8.10

In this appendix, we extend the main result of Calvert, Hegde, and Hammond [9] to prove Theorem 8.10. For brevity, we don't give full context for the paper [9] here and refer the interested reader to that paper. The paper [28] may also be a useful reference, as the work [9] builds on results from that paper. We strive to use the same notation as [9] so the interested reader can refer back easily. The main exception to this is that we use the notation $\tilde{\mathcal{B}}_i = 2^{-1/2}\mathcal{B}_i$ for lines in the (rescaled) parabolic Airy line ensemble. In [9], the authors use the notation $\mathcal{L}(i, \cdot)$ for these lines, which conflicts with our notation for the directed landscape. The factor of $2^{-1/2}$ is introduced in [9] so that comparison statements can be made with Brownian motions with diffusion parameter 1, rather than 2.

Throughout this section, we let $b > 0$ be a large constant and $b' > 0$ be a small constant, whose values may change from line to line but do not depend on any parameters. Other constants will retain the definitions used in [9].

First, fix an interval $[-d, d]$ with $d \geq 1$ and a collection of line indices $\llbracket 1, k \rrbracket$. For universal positive constants c, C , as in [9] we define

$$\begin{aligned} c_k &= ((3 - 2^{3/2})^{3/2} 2^{-1} 5^{-3/2})^{k-1} (2^{-5/2} c \wedge 1/8), \\ C_k &= \max \left\{ 10 \cdot 20^{k-1} 5^{k/2} \left(\frac{10}{3 - 2^{3/2}} \right)^{k(k-1)/2} C, e^{c/2} \right\}, \\ D_k &= \max \left\{ k^{1/3} c_k^{-1/3} (2^{-9/2} - 2^{-5})^{-1/3}, 36(k^2 - 1), 2 \right\}. \end{aligned}$$

The precise values of these constants are not important for our purposes here, but we will record the bounds

$$(A.9) \quad b' \leq D_k \leq e^{bk}, \quad b' \leq C_k \leq e^{bk^2}.$$

Next, let $\epsilon > 0$ satisfy the (k, d) -dependent upper bound

$$\epsilon < e^{-1} \wedge (17)^{-1/k} C_k^{-1/k} D_k^{-1} \wedge \exp(-(24)^6 d^6 / D_k^3).$$

This simplifies to

$$(A.10) \quad \epsilon < e^{-bd^6 - bk}.$$

Finally, we set $T = D_k(\log \epsilon^{-1})^{1/3}$.

Now, for a function $f : [a, b] \rightarrow \mathbb{R}$, define its **bridge version** $f^{[a,b]} = f - L$, where L is the linear function satisfying $L(a) = f(a)$ and $L(b) = f(b)$. Next, with all parameters d, k, ϵ, T fixed as above, let $\mathcal{F}_k$ be the σ -algebra generated by

- all the lower curves $\tilde{\mathcal{B}}_i : \mathbb{R} \rightarrow \mathbb{R}, i \geq k+1$,
- the top k curves $\tilde{\mathcal{B}}_i$ restricted to the set $\{x \in \mathbb{R} : |x| \geq 2T\}$,
- certain $\sigma(\tilde{\mathcal{B}}_{k+1})$ -measurable random variables $\mathfrak{l} \leq \mathfrak{r} \in [-T, T]$, and
- the $2k$ bridges $\tilde{\mathcal{B}}_i^{[-2T, \mathfrak{l}], i = 1, \dots, k}$ and $\tilde{\mathcal{B}}_i^{[\mathfrak{r}, 2T]}, i = 1, \dots, k$.

Here we use the notation $\sigma(X)$ for the σ -algebra generated by X . We let $\mathbb{P}_{\mathcal{F}_k}(\cdot) = \mathbb{P}(\cdot \mid \mathcal{F}_k)$ be the conditional law given $\mathcal{F}_k$. The precise nature of the random variables $\mathfrak{l}$ and $\mathfrak{r}$ is not important for us here, only their potential ranges and that they are functions of the $(k+1)^{\text{st}}$ curve $\tilde{\mathcal{B}}_{k+1}$. For precise definitions, see the beginning of Section 4.1.5 in [9].

With parameters d, k, ϵ, T fixed as above, in [9] and previously in [28], the authors define a collection of random functions $J = \{J_i : [-2T, 2T] \rightarrow \mathbb{R}, i \in \llbracket 1, k \rrbracket\}$ known as the **jump ensemble**. First, for any sequence of functions $X = \{X_i : [-2T, 2T] \rightarrow \mathbb{R}, i \in \llbracket 1, k \rrbracket\}$, we can define a **resampled ensemble** $\tilde{\mathcal{B}}^{\text{re}, X} = \{\tilde{\mathcal{B}}_i^{\text{re}, X} : \mathbb{R} \rightarrow \mathbb{R}, i \in \mathbb{N}\}$. For this definition we let $L(x, a; b, y)$ denote the affine function with $L(x) = a, L(y) = b$.

(A.11)

$$\tilde{\mathcal{B}}_i^{\text{re}, X}(x) = \begin{cases} \tilde{\mathcal{B}}_i(x), & (i, x) \notin \llbracket 1, k \rrbracket \times [-2T, 2T] \\ \tilde{\mathcal{B}}_i^{[-2T, \mathfrak{l}]}(x) + L(-2T, \tilde{\mathcal{B}}_i(-2T); \mathfrak{l}, X_i(\mathfrak{l}))(x), & x \in [-2T, \mathfrak{l}], i \leq k \\ \tilde{\mathcal{B}}_i^{[\mathfrak{r}, 2T]}(x) + L(\mathfrak{r}, X_i(\mathfrak{r}); 2T, \tilde{\mathcal{B}}_i(2T))(x), & x \in [\mathfrak{r}, 2T], i \leq k \\ X_i(x), & x \in [\mathfrak{l}, \mathfrak{r}], i \leq k. \end{cases}$$

Note that in [9], the same object is only defined for the top k lines. Next, in [9], the authors define an $\mathcal{F}_k$ -measurable finite set $P \subset [\mathfrak{l}, \mathfrak{r}]$ called a **pole set**, see the discussion in [9, Section 4.1.5]. The precise nature of this set is not important for us. Let $B = \{B_i : [-2T, 2T] \rightarrow \mathbb{R}, i \in \llbracket 1, k \rrbracket\}$ be a collection of Brownian bridges with $B_i(\pm 2T) = 0$ that are independent of $\tilde{\mathcal{B}}$ and each other. Finally, we define the ensemble J in the following way.

- First let $J' = \{J'_i : [-2T, 2T] \rightarrow \mathbb{R}, i \in \llbracket 1, k \rrbracket\}$ be given by connecting up the points $\tilde{\mathcal{B}}_i(\pm 2T)$ with the Brownian bridges B_i . That is, for all $i \in \llbracket 1, k \rrbracket$,

$$J'_i = B_i + L(-2T, \tilde{\mathcal{B}}_i(-2T); 2T, \tilde{\mathcal{B}}_i(2T)).$$

- Next, let J be given by the ensemble J' , conditionally on the events

$$\begin{aligned} \tilde{\mathcal{B}}_1^{\text{re}, J'}(x) &> \tilde{\mathcal{B}}_2^{\text{re}, J'}(x) > \dots > \tilde{\mathcal{B}}_{k+1}^{\text{re}, J'}(x) \quad \text{for } x = [-2T, \mathfrak{l}] \cup [\mathfrak{r}, 2T], \quad \text{and} \\ \tilde{\mathcal{B}}_i^{\text{re}, J'}(x) &\geq \tilde{\mathcal{B}}_{k+1}^{\text{re}, J'}(x), \quad \text{for all } i \in \llbracket 1, k \rrbracket, x \in P. \end{aligned}$$

This is the same as the definition given at the beginning of Section 4.1.6 in [9]. Next, let $\text{Pass}(J)$ be the indicator of event where

$$\tilde{\mathcal{B}}_i^{\text{re},J}(x) > \tilde{\mathcal{B}}_{i+1}^{\text{re},J}(x) \quad \text{for all } x \in [-2T, 2T], i \in \llbracket 1, k \rrbracket.$$

The relevance of the jump ensemble J lies in the following lemma..

LEMMA A.5. *We have*

$$\mathbb{P}_{\mathcal{F}_k}(\tilde{\mathcal{B}}_i^{\text{re},J}|_{\llbracket 1, k \rrbracket \times [-2T, 2T]} \in \cdot \mid \text{Pass}(J) = 1) = \mathbb{P}_{\mathcal{F}_k}(\tilde{\mathcal{B}}|_{\llbracket 1, k \rrbracket \times [-2T, 2T]} \in \cdot).$$

Here the restriction to $\llbracket 1, k \rrbracket \times [-2T, 2T]$ is a restriction to the top k lines and the interval $[-2T, 2T]$.

This is a special case of Lemma 4.5 in [9]. To compare with that lemma, we take $X' = J$, and replace the deterministic values ℓ and r and the set A with random $\mathcal{F}_k$ -measurable values $\mathfrak{l}$ and $\mathfrak{r}$ and the $\mathcal{F}_k$ -measurable set P . As noted in [9] in the discussion immediately following Equation (17) (at the end of Section 4.1.6), this replacement with $\mathcal{F}_k$ -measurable random variables follows does not affect the lemma since the claim is about $\mathcal{F}_k$ -conditional distributions.

The usefulness of Lemma A.5 in practice comes from the following four facts.

- (I) There exists an $\mathcal{F}_k$ -measurable event $\text{Fav}_{k,\epsilon}$ such that

$$\mathbb{P}_{\mathcal{F}_k}(\text{Pass}(J) = 1) \geq \exp(-3973k^{7/2}(d_{ip})^2 D_k^2(\log \epsilon^{-1})^{2/3}) \mathbf{1}(\text{Fav}_{k,\epsilon}).$$

This is Proposition 4.2 in [28], quoted as Proposition 4.9 in [9]. For use in [28], the quantity d_{ip} above is a parameter related to the pole set P , but in [9] and for our purposes, we take $d_{ip} = 5d$ (see Equation (18) in [9, Section 4.1.8] and surrounding discussions). Moving forward, we work with the simplified version of the above bound given by

$$\mathbb{P}_{\mathcal{F}_k}(\text{Pass}(J) = 1) \geq \exp(-d^2 e^{bk} (\log \epsilon^{-1})^{2/3}) \mathbf{1}(\text{Fav}_{k,\epsilon}).$$

- (II) The event $\text{Fav}_{k,\epsilon}$ satisfies

$$\mathbb{P}(\text{Fav}_{k,\epsilon}^c) \leq \epsilon.$$

This bound uses Lemma 4.1 in [28], cited as Lemma 4.10 in [9].

- (III) On $\text{Fav}_{k,\epsilon}$, we have $[-d, d] \subset [\mathfrak{l}, \mathfrak{r}]$. This follows from the paragraph before Lemma 4.10 in [9], which states that $[-T/2, T/2] \subset [\mathfrak{l}, \mathfrak{r}]$, and the discussions in [9] after Lemma 4.10, which shows that $[-d, d] \subset [-T/2, T/2]$.

- (IV) Let $\mathcal{C}_d$ denote the space of continuous functions from $[-d, d]$ to $\mathbb{R}$ that vanish at $-d$, equipped with the Borel σ -algebra in the topology of uniform convergence. Let $\mu_{0,*}^{[-d,d]}$ denote the law of a standard Brownian motion on $[-d, d]$ started from the initial condition $B(-d) = 0$. This is a measure on $\mathcal{C}_d$. Then there exists an absolute positive constant G such that if $\mu_{0,*}^{[-d,d]}(A) = \epsilon$, then

$$\begin{aligned} \mathbb{P}_{\mathcal{F}_k}(J_k(\cdot) - J_k(-d) \in A) \mathbf{1}(\text{Fav}_{k,\epsilon}) \\ \leq \epsilon G d^{1/2} D_k^4 (\log \epsilon^{-1})^{4/3} \exp(792d D_k^{5/2} (\log \epsilon^{-1})^{5/6}). \end{aligned}$$

This is Theorem 4.11 in [9]. The most important term to keep in mind here is the $(\log \epsilon^{-1})^{5/6}$ in the exponent. Moving forward, we will work with the simplified version of the bound given by

$$(A.12) \quad \mathbb{P}_{\mathcal{F}_k}(J_k(\cdot) - J_k(-d) \in A) \mathbf{1}(\text{Fav}_{k,\epsilon}) \leq \epsilon \exp(de^{bk} (\log \epsilon^{-1})^{5/6}).$$

Observe also that the inequality (A.12) for all A with $\mu_{0,*}^{[-d,d]}(A) = \epsilon$ implies that

$$(A.13) \quad \mathbb{P}_{\mathcal{F}_k}(J_k(\cdot) - J_k(-d) \in B) \mathbf{1}(\text{Fav}_{k,\epsilon}) \leq \mu_{0,*}^{[-d,d]}(B) \exp(de^{bk}(\log \epsilon^{-1})^{5/6})$$

for all B with $\mu_{0,*}^{[-d,d]}(B) \geq \epsilon$.

In [9], the authors use these three bounds with Lemma A.5 to find explicit Radon-Nikodym derivative estimates for individual parabolic Airy lines versus Brownian motion. With only slightly more work, we can upgrade these estimates to give bounds for multiple parabolic Airy lines versus several independent Brownian motions. We start with a lemma that translates the bounds on the jump ensemble to a conditional bound on parabolic Airy lines. For this lemma, we will also need to define $\text{Fav}_{k,\epsilon}$ when ϵ does not satisfy the bound in (A.10). In this case, we set $\text{Fav}_{k,\epsilon}$ to be the whole space.

LEMMA A.6. *With $\mu_{0,*}^{[-d,d]}$ as above, for every $d \geq 1, k \in \mathbb{N}$, and $\epsilon \in (0, 1]$, for every Borel measurable set A in $\mathcal{C}_d$ with $\mu_{0,*}^{[-d,d]}(A) \geq \epsilon$, we have*

$$(A.14) \quad \mathbb{P}_{\mathcal{F}_k}(\tilde{\mathcal{B}}_k(\cdot) - \tilde{\mathcal{B}}_k(-d) \in A) \mathbf{1}(\text{Fav}_{k,\epsilon}) < \mu_{0,*}^{[-d,d]}(A) \exp(bd^6 + de^{bk}(\log \epsilon^{-1})^{5/6}).$$

Moreover, if we let $f_{\mathcal{F}_k}$ denote the (random) Radon-Nikodym derivative of the random measure $\mathbb{P}_{\mathcal{F}_k}(\tilde{\mathcal{B}}_k(\cdot) - \tilde{\mathcal{B}}_k(-d) \in \cdot)$ with respect to $\mu_{0,*}^{[-d,d]}$, then almost surely,

$$(A.15) \quad \mathbb{E}_{\mathcal{F}_k} f_{\mathcal{F}_k} \mathbf{1}(f_{\mathcal{F}_k} \geq \exp(bd^6 + de^{bk}(\log \epsilon^{-1})^{5/6})) \mathbf{1}(\text{Fav}_{k,\epsilon}) \leq \epsilon \exp(bd^6 + de^{bk}(\log \epsilon^{-1})^{5/6}).$$

PROOF. First, the bound (A.14) holds trivially whenever ϵ does not satisfy (A.10) as long as b is taken large enough. Therefore we may assume that (A.10) holds.

In this case, we can let J be the jump ensemble defined with parameters d, k , and ϵ . Then by Lemma A.5, we can write

$$\mathbb{P}_{\mathcal{F}_k}(\tilde{\mathcal{B}}_k(\cdot) - \tilde{\mathcal{B}}_k(-d) \in A) = \mathbb{P}_{\mathcal{F}_k}(\tilde{\mathcal{B}}_k^{\text{re},J}(\cdot) - \tilde{\mathcal{B}}_k^{\text{re},J}(-d) \in A \mid \text{Pass}(J) = 1).$$

Now, by assertion (III) above, $[-d, d] \subset [l, r]$ on $\text{Fav}_{k,\epsilon}$. Therefore by the definition (A.11) of the resampled ensemble $\tilde{\mathcal{B}}_k^{\text{re},J}$, on $\text{Fav}_{k,\epsilon}$ we have

$$\begin{aligned} \mathbb{P}_{\mathcal{F}_k}(\tilde{\mathcal{B}}_k^{\text{re},J}(\cdot) - \tilde{\mathcal{B}}_k^{\text{re},J}(-d) \in A \mid \text{Pass}(J) = 1) &= \mathbb{P}_{\mathcal{F}_k}(J_k(\cdot) - J_k(-d) \in A \mid \text{Pass}(J) = 1) \\ &\leq \frac{\mathbb{P}_{\mathcal{F}_k}(J_k(\cdot) - J_k(-d) \in A)}{\mathbb{P}_{\mathcal{F}_k}(\text{Pass}(J) = 1)}. \end{aligned}$$

By assertion (I) and (A.13) above, on $\text{Fav}_{k,\epsilon}$ the right-hand side above is bounded by

$$\mu_{0,*}^{[-d,d]}(A) \exp(de^{bk}(\log \epsilon^{-1})^{5/6} + d^2 e^{bk}(\epsilon^{-1})^{2/3}).$$

The bound (A.10) on ϵ implies that this is bounded above by the right-hand side of (A.14). It remains to show (A.15).

We first claim that $\mathbb{P}_{\mathcal{F}_k}(\tilde{\mathcal{B}}_k(\cdot) - \tilde{\mathcal{B}}_k(-d) \in \cdot)$ is absolutely continuous with respect to $\mu_{0,*}^{[-d,d]}$ and so the Radon-Nikodym derivative $f_{\mathcal{F}_k}$ is well-defined. To see this, observe that $\mathcal{F}_k$ is contained in the σ -algebra $\mathcal{G}_k$ generated by $\tilde{\mathcal{B}}|_{S^c}$, where $S = \llbracket 1, k \rrbracket \times [-T, T]$, so it suffices to prove the same absolute continuity for $\mathbb{P}_{\mathcal{G}_k}(\tilde{\mathcal{B}}_k(\cdot) - \tilde{\mathcal{B}}_k(-d) \in \cdot)$. Conditional on $\mathcal{G}_k$, the Brownian Gibbs property for

$\tilde{\mathcal{B}}$ (Theorem 2.14) ensures that the process $\tilde{\mathcal{B}}_k|_{[-T,T]}$ is absolutely continuous with respect to a Brownian bridge between the $\mathcal{G}_k$ -measurable endpoints $\tilde{\mathcal{B}}_k(\pm T)$ at times $\pm T$. Since Brownian bridge increments are absolutely continuous with respect to Brownian motion increments away from the endpoints of the bridge interval, $\mathbb{P}_{\mathcal{G}_k}(\tilde{\mathcal{B}}_k(\cdot) - \tilde{\mathcal{B}}_k(-d) \in \cdot)$ is absolutely continuous with respect to $\mu_{0,*}^{[-d,d]}$.

Next, let

$$A = \{f_{\mathcal{F}_k} \geq \exp(bd^6 + de^{bk}(\log \epsilon^{-1})^{5/6})\},$$

so that the left-hand side of (A.15) is equal to

$$\mathbb{P}_{\mathcal{F}_k}(\tilde{\mathcal{B}}_k(\cdot) - \tilde{\mathcal{B}}_k(-d) \in A) \mathbf{1}(\text{Fav}_{k,\epsilon}).$$

By the definition of A , this is bounded below by

$$\mu_{0,*}^{[-d,d]}(A) \exp(bd^6 + de^{bk}(\log \epsilon^{-1})^{5/6}) \mathbf{1}(\text{Fav}_{k,\epsilon}).$$

By (A.14), this implies that $\mu_{0,*}^{[-d,d]}(A) \leq \epsilon$. Therefore we can find a set S such that $\mu_{0,*}^{[-d,d]}(A \cup S) = \epsilon$. Then by (A.14), we have

$$\begin{aligned} \mathbb{P}_{\mathcal{F}_k}(\tilde{\mathcal{B}}_k(\cdot) - \tilde{\mathcal{B}}_k(-d) \in A) \mathbf{1}(\text{Fav}_{k,\epsilon}) &\leq \mathbb{P}_{\mathcal{F}_k}(\tilde{\mathcal{B}}_k(\cdot) - \tilde{\mathcal{B}}_k(-d) \in A \cup S) \mathbf{1}(\text{Fav}_{k,\epsilon}) \\ &\leq \epsilon \exp(bd^6 + de^{bk}(\log \epsilon^{-1})^{5/6}), \end{aligned}$$

giving (A.15). $\square$

The next theorem is a restatement of Theorem 8.10.

THEOREM A.7. *Let $\mu_d^{\otimes k}$ denote the law of k -tuples of functions in $\mathcal{C}_d^k$ given by the product of k copies of $\mu_{0,*}^{[-d,d]}$. Define $\hat{\mathcal{B}}^k = (\hat{\mathcal{B}}_1, \dots, \hat{\mathcal{B}}_k)$ by letting $\hat{\mathcal{B}}_i = \tilde{\mathcal{B}}_i(\cdot) - \tilde{\mathcal{B}}_i(-d)$, restricted to the interval $[-d, d]$. Then for any set $A, k \in \mathbb{N}$ and $d \geq 1$, we have*

$$\mathbb{P}(\hat{\mathcal{B}}^k \in A) \leq \mu_d^{\otimes k}(A) \exp(bkd^6 + de^{bk}(\log[\mu_d^{\otimes k}(A)]^{-1})^{5/6}),$$

where $b > 0$ is a universal constant.

PROOF. Fix $d \geq 1$. We will first show that for every $k \in \mathbb{N}$, and $\epsilon \in (0, 1]$, that there exists an $\mathcal{F}_k$ -measurable set $\text{Fav}!_{k,\epsilon}$ with $\mathbb{P}(\text{Fav}!_{k,\epsilon}) \leq k\epsilon$ such that for every $\mathcal{C}_d^k$ -measurable set A with $\mu_d^{\otimes k}(A) \geq \epsilon$, we have

$$(A.16) \quad \mathbb{P}_{\mathcal{F}_k}(\hat{\mathcal{B}}^k \in A) \mathbf{1}(\text{Fav}!_{k,\epsilon}) \leq \mu_d^{\otimes k}(A) (4 \exp(bd^6 + de^{bk}(\log \epsilon^{-1})^{5/6}))^k.$$

The set $\text{Fav}!_{k,\epsilon}$ is not from [9], and is contained in $\text{Fav}_{k,\epsilon}$ but does not have an explicit construction. Think of the ‘!’ as indicating that it is an extra favourable version of the set $\text{Fav}_{k,\epsilon}$. To prove (A.16) we use induction on k . For the proof of (A.16), we fix the constant b , since increasing b during the inductive step would be problematic. The $k = 1$ case for all $\epsilon \in (0, 1]$ is given in Lemma A.6 with the set $\text{Fav}!_{k,\epsilon} = \text{Fav}_{k,\epsilon}$. Now suppose that the claim holds for $k - 1$ and all $\epsilon \in (0, 1]$. Let $A \subset \mathcal{C}_d^k$ be a Borel measurable set with $\mu_d^{\otimes k}(A) = \epsilon$. For every $x \in \mathcal{C}_d$, define the fibre

$$A_x = \{y \in \mathcal{C}_d^{k-1} : (y, x) \in A\}.$$

Then we can write

$$(A.17) \quad \mathbb{P}_{\mathcal{F}_k}(\hat{\mathcal{B}}^k \in A) = \mathbb{P}_{\mathcal{F}_k}(\hat{\mathcal{B}}^{k-1} \in A_{\hat{\mathcal{B}}_k}) = \mathbb{E}_{\mathcal{F}_k}(\mathbb{P}_{\mathcal{F}_{k-1}}(\hat{\mathcal{B}}^{k-1} \in A_{\hat{\mathcal{B}}_k})).$$

where the last equality uses that $\mathcal{F}_k \subset \mathcal{F}_{k-1}$. We use the inductive hypothesis to estimate $\mathbb{P}_{\mathcal{F}_{k-1}}(\hat{\mathcal{B}}^{k-1} \in A_{\hat{\mathcal{B}}_k})$. First, let S be any set with $\mu_d^{\otimes k-1}(S) = \epsilon^2$. Then we can write

$$\begin{aligned}
 & \mathbb{P}_{\mathcal{F}_{k-1}}(\hat{\mathcal{B}}^{k-1} \in A_{\hat{\mathcal{B}}_k}) \\
 & \leq \mathbb{P}_{\mathcal{F}_{k-1}}(\hat{\mathcal{B}}^{k-1} \in A_{\hat{\mathcal{B}}_k} \cup S) \\
 & \leq \mathbb{P}_{\mathcal{F}_{k-1}}(\hat{\mathcal{B}}^{k-1} \in A_{\hat{\mathcal{B}}_k} \cup S) \mathbf{1}(\text{Fav!}_{k-1, \epsilon^2}) + \mathbf{1}(\text{Fav!}_{k-1, \epsilon^2}^c) \\
 (A.18) \quad & \leq \mu_d^{\otimes k-1}(A_{\hat{\mathcal{B}}_k} \cup S) \left(4 \exp(bd^6 + de^{b(k-1)}(\log \epsilon^{-1})^{5/6}) \right)^{k-1} + \mathbf{1}(\text{Fav!}_{k-1, \epsilon^2}^c).
 \end{aligned}$$

Here the final inequality uses the inductive hypothesis, and the fact that $\mu_d^{\otimes k-1}(A_{\hat{\mathcal{B}}_k} \cup S)$ is always greater than ϵ^2 . Next, we want to apply $\mathbb{E}_{\mathcal{F}_k}$ to the right-hand side of (A.18). We start with the term $\mu_d^{\otimes k-1}(A_{\hat{\mathcal{B}}_k} \cup S)$. As in Lemma A.6, let $f_{\mathcal{F}_k}$ denote the Radon-Nikodym derivative of $\mathbb{P}_{\mathcal{F}_k}(\hat{\mathcal{B}}_k \in \cdot)$ with respect to $\mu_{0,*}^{[-d,d]}$. Letting W be an independent Brownian motion drawn from the distribution $\mu_{0,*}^{[-d,d]}$, we can write

$$\begin{aligned}
 & \mathbb{E}_{\mathcal{F}_k} \mu_d^{\otimes k-1}(A_{\hat{\mathcal{B}}_k} \cup S) = \mathbb{E}_{\mathcal{F}_k} f_{\mathcal{F}_k}(W) \mu_d^{\otimes k-1}(A_W \cup S) \\
 (A.19) \quad & \leq \mathbb{E}_{\mathcal{F}_k} f_{\mathcal{F}_k}(W) \mathbf{1}(f_{\mathcal{F}_k}(W) \geq \exp(bd^6 + de^{bk}(\log \epsilon^{-1})^{5/6})) \\
 & \quad + \exp(bd^6 + de^{bk}(\log \epsilon^{-1})^{5/6}) \mathbb{E}_{\mathcal{F}_k} \mu_d^{\otimes k-1}(A_W \cup S).
 \end{aligned}$$

Now, by the definition of the sets A_x and a union bound, we have

$$\mathbb{E}_{\mathcal{F}_k} \mu_d^{\otimes k-1}(A_W \cup S) \leq \mathbb{E}_{\mathcal{F}_k} [\mu_d^{\otimes k-1}(A_W) + \mu_d^{\otimes k-1}(S)] = \mu_d^{\otimes k}(A) + \epsilon^2 \leq 2\epsilon.$$

Also, on the event $\text{Fav}_{k,\epsilon}$ in Lemma A.6, we can bound the first term on the right-hand side of (A.19) above using (A.15). Therefore

$$(A.20) \quad \mathbb{E}_{\mathcal{F}_k} \mu_d^{\otimes k-1}(A_{\hat{\mathcal{B}}_k} \cup S) \mathbf{1}(\text{Fav}_{k,\epsilon}) \leq 3\epsilon \exp(bd^6 + de^{bk}(\log \epsilon^{-1})^{5/6}).$$

We now bound the second term on the right-hand side of (A.18). We have

$$\mathbb{E} \mathbb{E}_{\mathcal{F}_k} (\mathbf{1}(\text{Fav!}_{k-1, \epsilon^2}^c)) = \mathbb{P}(\text{Fav!}_{k-1, \epsilon^2}^c) \leq (k-1)\epsilon^2,$$

by the inductive hypothesis, so by Markov's inequality, we have

$$(A.21) \quad \mathbb{E}_{\mathcal{F}_k} (\mathbf{1}(\text{Fav!}_{k-1, \epsilon^2}^c)) \leq \epsilon$$

on a set B of probability $1 - (k-1)\epsilon$. We now set $\text{Fav!}_{k,\epsilon} = B \cap \text{Fav}_{k,\epsilon}$. Assertion (II) and a union bound shows that $\mathbb{P}(\text{Fav!}_{k,\epsilon}^c) \leq k\epsilon$. Finally, gathering the inequalities (A.18), (A.20), and (A.21), we have

$$\begin{aligned}
 & \mathbb{E}_{\mathcal{F}_k} \mathbb{P}_{\mathcal{F}_{k-1}}(\hat{\mathcal{B}}^{k-1} \in A_{\hat{\mathcal{B}}_k}) \mathbf{1}(\text{Fav!}_{k,\epsilon}) \\
 & \leq \left(4 \exp(bd^6 + de^{b(k-1)}(\log \epsilon^{-1})^{5/6}) \right)^{k-1} 3\epsilon \exp(bd^6 + de^{bk}(\log \epsilon^{-1})^{5/6}) + \epsilon.
 \end{aligned}$$

This is bounded above by (A.16) when $\mu_d^{\otimes k}(A) = \epsilon$. As in (A.13), the extension of (A.16) to all A with $\mu_d^{\otimes k}(A) \geq \epsilon$ is immediate.

The theorem then follows by averaging over $\mathcal{F}_k$. More precisely, let A be any set, and define $\epsilon = \mu_d^{\otimes k}(A)$. Then

$$\begin{aligned}
 \mathbb{P}(\hat{\mathcal{B}} \in A) & \leq \mathbb{E} \mathbb{P}_{\mathcal{F}_k}(\hat{\mathcal{B}} \in A) \mathbf{1}(\text{Fav!}_{k,\epsilon}) + \mathbb{P}(\text{Fav!}_{k,\epsilon}^c) \\
 & \leq \epsilon \left(4 \exp(bd^6 + de^{bk}(\log \epsilon^{-1})^{5/6}) \right)^k + k\epsilon,
 \end{aligned}$$

This gives the desired bound after increasing b . □

NOT FOR DISTRIBUTION

Bibliography

- [1] Philippe Biane, Philippe Bougerol, and Neil O’Connell, *Littellmann paths and Brownian paths*, Duke Math. J. **130** (2005), no. 1, 127–167, DOI 10.1215/S0012-7094-05-13014-9. MR2176549
- [2] Márton Balázs, Ofer Busani, and Timo Seppäläinen, *Local stationarity in exponential last-passage percolation*, Probab. Theory Related Fields **180** (2021), no. 1-2, 113–162, DOI 10.1007/s00440-021-01035-7. MR4265019
- [3] Jinho Baik, Percy Deift, and Kurt Johansson, *On the distribution of the length of the longest increasing subsequence of random permutations*, J. Amer. Math. Soc. **12** (1999), no. 4, 1119–1178, DOI 10.1090/S0894-0347-99-00307-0. MR1682248
- [4] Erik Bates, Shirshendu Ganguly, and Alan Hammond, *Hausdorff dimensions for shared endpoints of disjoint geodesics in the directed landscape*, Electron. J. Probab. **27** (2022), Paper No. 1, 44, DOI 10.1214/21-ejp706. MR4361743
- [5] Riddhipratim Basu, Shirshendu Ganguly, Alan Hammond, and Milind Hegde, *Interlacing and scaling exponents for the geodesic watermelon in last passage percolation*, Comm. Math. Phys. **393** (2022), no. 3, 1241–1309, DOI 10.1007/s00220-022-04388-9. MR4453234
- [6] Riddhipratim Basu, Shirshendu Ganguly, and Lingfu Zhang, *Temporal correlation in last passage percolation with flat initial condition via Brownian comparison*, Comm. Math. Phys. **383** (2021), no. 3, 1805–1888, DOI 10.1007/s00220-021-03958-7. MR4244262
- [7] Riddhipratim Basu, Sourav Sarkar, and Allan Sly, *Coalescence of geodesics in exactly solvable models of last passage percolation*, J. Math. Phys. **60** (2019), no. 9, 093301, 22, DOI 10.1063/1.5093799. MR4002528
- [8] Ivan Corwin and Alan Hammond, *Brownian Gibbs property for Airy line ensembles*, Invent. Math. **195** (2014), no. 2, 441–508, DOI 10.1007/s00222-013-0462-3. MR3152753
- [9] Jacob Calvert, Alan Hammond, and Milind Hegde, *Brownian structure in the KPZ fixed point*, Astérisque **441** (2023), vi+119. MR4612529
- [10] I. Corwin, *Kardar-Parisi-Zhang universality*, Notices Amer. Math. Soc. **63** (2016), no. 3, 230–239, DOI 10.1090/noti1334. MR3445162
- [11] Ivan Corwin and Xin Sun, *Ergodicity of the Airy line ensemble*, Electron. Commun. Probab. **19** (2014), no. 49, 11, DOI 10.1214/ECP.v19-3504. MR3246968
- [12] Duncan Dauvergne, *In preparation*.
- [13] Duncan Dauvergne, *Last passage isometries for the directed landscape*, Probab. Theory Related Fields **186** (2023), no. 1-2, 391–437, DOI 10.1007/s00440-022-01173-6. MR4586223
- [14] Duncan Dauvergne, Mihai Nica, and Bálint Virág, *Uniform convergence to the Airy line ensemble* (English, with English and French summaries), Ann. Inst. Henri Poincaré Probab. Stat. **59** (2023), no. 4, 2220–2256, DOI 10.1214/22-aihp1314. MR4663521
- [15] Duncan Dauvergne, Janosch Ortmann, and Bálint Virág, *The directed landscape*, Acta Math. **229** (2022), no. 2, 201–285, DOI 10.4310/acta.2022.v229.n2.a1. MR4554223
- [16] Duncan Dauvergne, Sourav Sarkar, and Bálint Virág, *Three-halves variation of geodesics in the directed landscape*, Ann. Probab. **50** (2022), no. 5, 1947–1985, DOI 10.1214/22-aop1574. MR4474505
- [17] Duncan Dauvergne and Bálint Virág, *Bulk properties of the Airy line ensemble*, Ann. Probab. **49** (2021), no. 4, 1738–1777, DOI 10.1214/20-aop1492. MR4260467
- [18] Duncan Dauvergne and Virág Bálint, *The scaling limit of the longest increasing subsequence*, arXiv preprint, [arXiv:2104.08210](https://arxiv.org/abs/2104.08210), (2021).
- [19] P. L. Ferrari and H. Spohn, *Random growth models*, The Oxford handbook of random matrix theory, Oxford Univ. Press, Oxford, 2011, pp. 782–801. MR2932658

- [20] Shirshendu Ganguly, *Random metric geometries on the plane and Kardar-Parisi-Zhang universality*, Notices Amer. Math. Soc. **69** (2022), no. 1, 26–35, DOI 10.1090/noti2404. MR4353350
- [21] Shirshendu Ganguly and Alan Hammond, *The geometry of near ground states in Gaussian polymer models*, Electron. J. Probab. **28** (2023), Paper No. 60, 80, DOI 10.1214/22-ejp895. MR4583067
- [22] Shirshendu Ganguly and Alan Hammond, *Stability and chaos in dynamical last passage percolation*, Commun. Am. Math. Soc. **4** (2024), 387–479, DOI 10.1090/cams/35. MR4755868
- [23] Nicos Georgiou, Firas Rassoul-Agha, and Timo Seppäläinen, *Stationary cocycles and Busemann functions for the corner growth model*, Probab. Theory Related Fields **169** (2017), no. 1-2, 177–222, DOI 10.1007/s00440-016-0729-x. MR3704768
- [24] Shirshendu Ganguly and Lingfu Zhang, *Fractal geometry of the space-time difference profile in the directed landscape via construction of geodesic local times*, arXiv preprint, arXiv:2204.01674, 2022.
- [25] Alan Hammond, *Modulus of continuity of polymer weight profiles in Brownian last passage percolation*, Ann. Probab. **47** (2019), no. 6, 3911–3962, DOI 10.1214/19-aop1350. MR4038045
- [26] Alan Hammond, *A patchwork quilt sewn from Brownian fabric: regularity of polymer weight profiles in Brownian last passage percolation*, Forum Math. Pi **7** (2019), e2, 69, DOI 10.1017/fmp.2019.2. MR3987302
- [27] Alan Hammond, *Exponents governing the rarity of disjoint polymers in Brownian last passage percolation*, Proc. Lond. Math. Soc. (3) **120** (2020), no. 3, 370–433, DOI 10.1112/plms.12292. MR4008374
- [28] Alan Hammond, *Brownian regularity for the Airy line ensemble, and multi-polymer watermelons in Brownian last passage percolation*, Mem. Amer. Math. Soc. **277** (2022), no. 1363, v+133, DOI 10.1090/memo/1363. MR4403929
- [29] Kurt Johansson, *Shape fluctuations and random matrices*, Comm. Math. Phys. **209** (2000), no. 2, 437–476, DOI 10.1007/s002200050027. MR1737991
- [30] Kurt Johansson and Mustazee Rahman, *Multitime distribution in discrete polynuclear growth*, Comm. Pure Appl. Math. **74** (2021), no. 12, 2561–2627, DOI 10.1002/cpa.21980. MR4373163
- [31] Olav Kallenberg, *Foundations of modern probability*, 2nd ed., Probability and its Applications (New York), Springer-Verlag, New York, 2002, DOI 10.1007/978-1-4757-4015-8. MR1876169
- [32] Wolfgang König, Neil O’Connell, and Sébastien Roch, *Non-colliding random walks, tandem queues, and discrete orthogonal polynomial ensembles*, Electron. J. Probab. **7** (2002), no. 5, 24, DOI 10.1214/EJP.v7-104. MR1887625
- [33] Zhipeng Liu, *Multipoint distribution of TASEP*, Ann. Probab. **50** (2022), no. 4, 1255–1321, DOI 10.1214/21-aop1557. MR4420420
- [34] Cristina Licea and Charles M. Newman, *Geodesics in two-dimensional first-passage percolation*, Ann. Probab. **24** (1996), no. 1, 399–410, DOI 10.1214/aop/1042644722. MR1387641
- [35] B. F. Logan and L. A. Shepp, *A variational problem for random Young tableaux*, Advances in Math. **26** (1977), no. 2, 206–222, DOI 10.1016/0001-8708(77)90030-5. MR1417317
- [36] Peter Mörters and Yuval Peres, *Brownian motion*, Cambridge Series in Statistical and Probabilistic Mathematics, vol. 30, Cambridge University Press, Cambridge, 2010. With an appendix by Oded Schramm and Wendelin Werner, DOI 10.1017/CBO9780511750489. MR2604525
- [37] Konstantin Matetski, Jeremy Quastel, and Daniel Remenik, *The KPZ fixed point*, Acta Math. **227** (2021), no. 1, 115–203, DOI 10.4310/acta.2021.v227.n1.a3. MR4346267
- [38] Masatoshi Noumi and Yasuhiko Yamada, *Tropical Robinson-Schensted-Knuth correspondence and birational Weyl group actions*, Representation theory of algebraic groups and quantum groups, Adv. Stud. Pure Math., vol. 40, Math. Soc. Japan, Tokyo, 2004, pp. 371–442, DOI 10.2969/aspm/04010371. MR2074600
- [39] Neil O’Connell, *Random matrices, non-colliding processes and queues*, Séminaire de Probabilités, XXXVI, Lecture Notes in Math., vol. 1801, Springer, Berlin, 2003, pp. 165–182, DOI 10.1007/978-3-540-36107-7_3. MR1971584
- [40] Neil O’Connell and Marc Yor, *Brownian analogues of Burke’s theorem*, Stochastic Process. Appl. **96** (2001), no. 2, 285–304, DOI 10.1016/S0304-4149(01)00119-3. MR1865759
- [41] Neil O’Connell and Marc Yor, *A representation for non-colliding random walks*, Electron. Comm. Probab. **7** (2002), 1–12, DOI 10.1214/ECP.v7-1042. MR1887169

- [42] Leandro P. R. Pimentel, *Duality between coalescence times and exit points in last-passage percolation models*, Ann. Probab. **44** (2016), no. 5, 3187–3206, DOI 10.1214/15-AOP1044. MR3551194
- [43] Michael Prähofer and Herbert Spohn, *Scale invariance of the PNG droplet and the Airy process*, J. Statist. Phys. **108** (2002), no. 5-6, 1071–1106, DOI 10.1023/A:1019791415147. Dedicated to David Ruelle and Yasha Sinai on the occasion of their 65th birthdays. MR1933446
- [44] Jeremy Quastel and Daniel Remenik, *KP governs random growth off a 1-dimensional substrate*, Forum Math. Pi **10** (2022), Paper No. e10, 26, DOI 10.1017/fmp.2021.9. MR4413350
- [45] Jeremy Quastel, *Introduction to KPZ*, Current developments in mathematics, 2011, Int. Press, Somerville, MA, 2012, pp. 125–194. MR3098078
- [46] Dan Romik, *The surprising mathematics of longest increasing subsequences*, Institute of Mathematical Statistics Textbooks, vol. 4, Cambridge University Press, New York, 2015. MR3468738
- [47] Timo Seppäläinen and Xiao Shen, *Coalescence estimates for the corner growth model with exponential weights*, Electron. J. Probab. **25** (2020), Paper No. 85, 31, DOI 10.1214/20-ejp489. MR4125790
- [48] Sourav Sarkar and Bálint Virág, *Brownian absolute continuity of the KPZ fixed point with arbitrary initial condition*, Ann. Probab. **49** (2021), no. 4, 1718–1737, DOI 10.1214/20-aop1491. MR4260466
- [49] Craig A. Tracy and Harold Widom, *Level-spacing distributions and the Airy kernel*, Comm. Math. Phys. **159** (1994), no. 1, 151–174. MR1257246
- [50] A. M. Vershik and S. V. Kerov, *Asymptotic behavior of the Plancherel measure of the symmetric group and the limit form of Young tableaux* (Russian), Dokl. Akad. Nauk SSSR **233** (1977), no. 6, 1024–1027. MR0480398
- [51] Thomas Weiss, Patrik Ferrari, and Herbert Spohn, *Reflected Brownian motions in the KPZ universality class*, SpringerBriefs in Mathematical Physics, vol. 18, Springer, Cham, 2017, DOI 10.1007/978-3-319-49499-9. MR3585775
- [52] Lingfu Zhang, *Optimal exponent for coalescence of finite geodesics in exponential last passage percolation*, Electron. Commun. Probab. **25** (2020), Paper No. 74, 14, DOI 10.1214/20-ecp354. MR4167385
- [53] Nikos Zygouras, *Some algebraic structures in KPZ universality*, Probab. Surv. **19** (2022), 590–700, DOI 10.1214/19-ps335. MR4524523