Chapter 4

Convergence of composite slices

The goal of this chapter is to prove Theorem 2.6 on the convergence of slices to their limiting slices. To this end, we are first going to derive a "free" version of this result by finding slices with a free area and tilt within the uniform infinite half-planar quadrangulation. The latter is known to converge to the Brownian half-plane, which itself contains a "flow" of continuum slices with free areas and tilts; these are shown to be the scaling limits of the discrete slices. We conclude by a conditioning argument to pass from free to fixed area and tilt. First, let us start with deterministic considerations.

4.1 Metric spaces coded by real functions

Here we borrow some material from [13, Section 2.1], with however several slight differences, in order to describe in a unified fashion the various random metric spaces we will use. Let $\mathcal{C}$ (resp. $\mathcal{C}^{(2)}$) be the set of continuous functions of one variable (resp. of two variables) defined on some nonempty closed interval:

$$\mathcal{C} = \bigsqcup_{\substack{I \text{ closed interval} \\ I \neq \varnothing}} \mathcal{C}(I, \mathbb{R})$$

and

$$\mathcal{C}^{(2)} = \bigsqcup_{\substack{I \text{ closed interval} \\ I \neq \varnothing}} \mathcal{C}(I^2, \mathbb{R}).$$

For a function $f \in \mathcal{C}(I, \mathbb{R})$, we denote by $I(f) = I$ its interval of definition and by $\overline{\tau}(f) = \inf I$ and $\tau(f) = \sup I$ its extremities. The set $\mathcal{C}$ is naturally equipped with the topology of uniform convergence over compact subsets of $\mathbb{R}$; more precisely, the topology induced by the following metric:

$$\begin{aligned}\mathrm{dist}_{\mathcal{C}}(f, g) = &|\arctan(\overline{\tau}(f)) - \arctan(\overline{\tau}(g))| + |\arctan(\tau(f)) - \arctan(\tau(g))| \\ &+ \sum_{n \geq 1} \frac{1}{2^n} \sup_{t \in [-n,n]} |f(\overline{\tau}(f) \vee t \wedge \tau(f)) - g(\overline{\tau}(g) \vee t \wedge \tau(g))|.\end{aligned}$$

We also equip $\mathcal{C}^{(2)}$ with a straightforward adaptation $\mathrm{dist}_{\mathcal{C}^{(2)}}$.

4.1.1 $\mathbb{R}$-trees coded by functions

$\mathbb{R}$-trees. For $f \in \mathcal{C}$ and $s, t \in I = I(f)$ with $s \le t$, set

$$\underline{f}(s,t) = \inf_{[s,t]} f \tag{4.1}$$

and, for $s, t \in I$, set

$$d_f(s,t) = f(s) + f(t) - 2\underline{f}(s \wedge t, s \vee t). \tag{4.2}$$

This formula defines a pseudometric on I, which is continuous as a function from I^2 to $\mathbb{R}_{\ge 0}$, since $d_f(s,t) \le 2\omega(f; [s \wedge t, s \vee t])$, where $\omega(f; J) = \sup_J f - \inf_J f$. We let $\mathcal{T}_f = (I/\{d_f = 0\}, d_f)$ be the associated quotient space, and $\mathbf{p}_f : I \to \mathcal{T}_f$ be the canonical projection, which is continuous since d_f is. The space $\mathcal{T}_f$ is a so-called $\mathbb{R}$-*tree*, that is, satisfies the following:

- For every two points $a, b \in \mathcal{T}_f$, there exists a geodesic from a to b, that is, an isometric mapping $\chi_{a,b} : [0, d_f(a,b)] \to \mathcal{T}_f$ with $\chi_{a,b}(0) = a$ and $\chi_{a,b}(d_f(a,b)) = b$.
- The image of the path $\chi_{a,b}$, which we denote by $[\![a, b]\!]_f$, is the image of any injective path from a to b.

If I is compact, we let $a_*(f) = \mathbf{p}_f(t_*)$, where t_* is any point at which f attains its overall minimum. In this case, for $t \in I$ and $a = \mathbf{p}_f(t)$, the geodesic segment $[\![a, a_*(f)]\!]_f$ is given by

$$[\![a, a_*(f)]\!]_f = \mathbf{p}_f(\{s \in [t \wedge t_*, t \vee t_*] : f(s) \le f(u),\ \forall u \in [s \wedge t, s \vee t]\}).$$

In the case where I is unbounded, we will systematically make the extra assumption that

$$\begin{cases} \text{when } \overline{\tau}(f) = -\infty, & \inf_{t \le 0} f(t) = -\infty \text{ or } \lim_{t \to -\infty} f(t) = \infty, \\ \text{when } \tau(f) = \infty, & \inf_{t \ge 0} f(t) = -\infty \text{ or } \lim_{t \to \infty} f(t) = \infty. \end{cases} \tag{4.3}$$

In particular, it holds that

$$\forall s \in I, \qquad \lim_{|t| \to \infty, t \in I} d_f(s,t) = \infty, \tag{4.4}$$

which implies that $\mathcal{T}_f$ is locally compact, as the reader may easily check.

Gluing two $\mathbb{R}$-trees. Next, given two functions $f, g \in \mathcal{C}$ with common interval of definition

$$I = I(f) = I(g)$$

both satisfying (4.3), we define another pseudometric on I as the quotient pseudometric (defined by (3.2))

$$D_{f,g}(s,t) = d_g / \{d_f = 0\}, \tag{4.5}$$

and equip the quotient set $M_{f,g} = I/\{D_{f,g} = 0\}$ with the metric $D_{f,g}$. Note that $D_{f,g} \colon I^2 \to \mathbb{R}_{\geq 0}$ is continuous since

$$|D_{f,g}(s,t) - D_{f,g}(s',t')| \leq D_{f,g}(s,s') + D_{f,g}(t,t') \\ \leq 2\omega(g;[s \wedge s', s \vee s']) + 2\omega(g;[t \wedge t', t \vee t']).$$

For this reason, the canonical projection $\mathbf{p}_{f,g} \colon I \to M_{f,g}$ is continuous. We may view $(M_{f,g}, D_{f,g})$ as gluing the $\mathbb{R}$-tree $\mathcal{T}_g$ along the equivalence relation defining the $\mathbb{R}$-tree $\mathcal{T}_f$. In fact, since either of $d_f(s,t) = 0$ or $d_g(s,t) = 0$ implies $D_{f,g} = 0$, the canonical projection $\mathbf{p}_{f,g}$ factorizes as

$$\mathbf{p}_{f,g} = \pi_f \circ \mathbf{p}_f = \pi_g \circ \mathbf{p}_g,$$

where $\pi_f \colon \mathcal{T}_f \to M_{f,g}$ and $\pi_g \colon \mathcal{T}_g \to M_{f,g}$ are two surjective maps. Note that these functions are continuous: if $a_n = \mathbf{p}_f(t_n)$ converges to some point a, then, up to taking extractions (and using (4.4) if I is unbounded), we may assume that t_n converges to some limit t, and then $\mathbf{p}_f(t) = a$ by continuity of $\mathbf{p}_f$, while $\pi_f(a_n) = \mathbf{p}_{f,g}(t_n)$ converges to $\mathbf{p}_{f,g}(t) = \pi_f(a)$. As a consequence, every geodesic segment $[\![a,b]\!]_f$ in $\mathcal{T}_f$, and every geodesic $[\![c,d]\!]_g$ in $\mathcal{T}_g$ is "immersed" into $M_{f,g}$ via the mappings π_f, π_g.

4.1.2 Composite slices coded by two functions

Slice trajectory. We say that (f,g) is a *slice trajectory* if $f, g \in \mathcal{C}$ have common interval of definition I,

$$\forall s,t \in I, \quad d_f(s,t) = 0 \Rightarrow g(s) = g(t), \tag{4.6}$$

if $\inf I = -\infty$, then

$$\lim_{t \to -\infty} f(t) = +\infty \quad \text{and} \quad \inf_{t \leq 0} g(t) = -\infty,$$

and, if $\sup I = \infty$, then

$$\inf_{t \geq 0} f(t) = -\infty \quad \text{and} \quad \inf_{t \geq 0} g(t) = -\infty.$$

In particular, f and g both satisfy (4.3) in the case where I is noncompact, and the quantity $f(\inf I) \in \mathbb{R} \cup \{+\infty\}$ is always well defined.

In the remainder of this section, we fix a slice trajectory (f,g), and call the metric space

$$\mathsf{Sl}_{f,g} = (M_{f,g}, D_{f,g})$$

the *slice coded by* (f,g). For the moment, we focus on deterministic considerations; the functions f, g will be randomized in the following section.

Marks and measures. The slice $\mathsf{Sl}_{f,g}$ naturally comes with the following distinguished elements.

Geodesics sides. For every $t \in I$, we set

$$\Gamma_t(r) = \inf\{s \geq t : g(s) = g(t) - r\} \quad \text{for } r \in \mathbb{R}_{\geq 0} \text{ such that } \inf_{\substack{s \geq t \\ s \in I}} g(s) \leq g(t) - r,$$
$$\Xi_t(r) = \sup\{s \leq t : g(s) = g(t) - r\} \quad \text{for } r \in \mathbb{R}_{\geq 0} \text{ such that } \inf_{\substack{s \leq t \\ s \in I}} g(s) \leq g(t) - r.$$

In particular, we have $d_g(\Gamma_t(r), \Xi_t(r)) = 0$ for every $t \in I$ and every r satisfying both inequalities above.

We extend the definition given in Section 3.2 of geodesics to paths $\chi\colon [0,\infty) \to \mathcal{X}$ that satisfy (3.1) for every $s, t \in \mathbb{R}_{\geq 0}$. In this case, the point $\chi(0)$ is called the *origin* of χ, its length is set to $\mathrm{length}(\chi) = \infty$ by convention, and the range of χ is called a *geodesic ray*. The geodesic ray uniquely determines the geodesic χ by the same argument as for finite length, since the origin of a geodesic ray is the unique point a such that, for any $s > 0$, the number of points in the ray at distance s from a is one.

We observe that Γ_t and Ξ_t are geodesics (possibly of infinite length) from t for the pseudometrics d_g and $D_{f,g}$, in the sense that, for every r, r' such that $\Gamma_t(r)$ and $\Gamma_t(r')$ are defined,

$$d_g(\Gamma_t(r), \Gamma_t(r')) = D_{f,g}(\Gamma_t(r), \Gamma_t(r')) = |r' - r|, \tag{4.7}$$

and the same holds with Ξ_t in place of Γ_t. This fact is immediate for d_g by definition. In fact, when I is compact, one checks that the images of $\mathbf{p}_g \circ \Gamma_t$ and $\mathbf{p}_g \circ \Xi_t$ are the geodesic segments $[\![\mathbf{p}_g(t), a_*(g|_{[t,\sup I]})]\!]_g$ and $[\![\mathbf{p}_g(t), a_*(g|_{[\inf I,t]})]\!]_g$ in $\mathcal{T}_g$.

For $D_{f,g}$, this fact follows from the bound

$$|g(s) - g(s')| \leq D_{f,g}(s, s') \leq d_g(s, s'), \quad s, s' \in I,$$

where the first inequality is an easy consequence of the fact that (f, g) is a slice trajectory.

Therefore, for $t \in I$, the paths defined by

$$\gamma_t(r) = \mathbf{p}_{f,g}(\Gamma_t(r)), \quad 0 \leq r \leq g(t) - \underline{g}(t, \sup I), \ r \in \mathbb{R},$$
$$\xi_t(r) = \mathbf{p}_{f,g}(\Xi_t(r)), \quad 0 \leq r \leq g(t) - \underline{g}(\inf I, t), \ r \in \mathbb{R},$$

are two geodesics (possibly of infinite length) from $\mathbf{p}_{f,g}(t)$, sharing a common initial part. As mentioned in Section 3.2, we will often identify these paths with the pairs formed by their origins and image sets, the latter being $\pi_g([\![\mathbf{p}_g(t), a_*(g|_{[t,\sup I]})]\!]_g)$ and $\pi_g([\![\mathbf{p}_g(t), a_*(g|_{[\inf I,t]})]\!]_g)$ when I is compact.

The slice $M_{f,g}$ comes with zero, one, or two geodesic sides. If $\inf I > -\infty$, then the geodesic $\gamma = \gamma_{\inf I}$ is called the *maximal geodesic* of $M_{f,g}$, and, if $\sup I < \infty$, the geodesic $\xi = \xi_{\sup I}$ is called the *shuttle* of $M_{f,g}$. If $\inf I = -\infty$ (resp. $\sup I = \infty$), we let $\gamma_{-\infty}$ (resp. ξ_{∞}) be the empty set. If I is a bounded interval, then the paths $\gamma_{\inf I}$ and $\xi_{\sup I}$ have a common endpoint at the *apex*

$$x_* = \mathbf{p}_{f,g}(s_*) = \pi_g(a_*(g)),$$

where s_* denotes any point s in I such that $g(s) = \inf_I g$.

Base. For $x \in \mathbb{R}$, we define

$$T_x = \inf\{t \in I : f(t) = -x\} \in \mathbb{R} \cup \{\infty\},$$

the hitting time of level $-x$ by the function f, with the convention that $\inf \varnothing = \infty$. Note that, for $x \in \mathbb{R}$, $T_x \neq -\infty$ because of the fact that (f, g) is admissible. By convention, we also set $T_\infty = -T_{-\infty} = \infty$. The *base* of $\mathsf{Sl}_{f,g}$ is the set

$$\beta = \mathbf{p}_{f,g}\big(\{T_x : -f(\inf I) \le x \le -\inf_I f\} \cap \mathbb{R}\big).$$

Note that the set inside brackets projects via $\mathbf{p}_f$ to a geodesic in $\mathcal{T}_f$. When I is compact, the base is the path $\pi_f([\![\mathbf{p}_f(T_{-f(\inf I)}), \mathbf{p}_f(T_{-\inf_I f})]\!])$, and in general, it is the increasing union of the paths

$$\pi_f([\![\mathbf{p}_f(T_x), \mathbf{p}_f(T_y)]\!]_f), \quad -f(\inf I) \le x < y \le -\inf_I f,\ x, y \in \mathbb{R}.$$

Measures. Finally, denoting by Leb_J the Lebesgue measure on the interval J, the slice $\mathsf{Sl}_{f,g}$ is endowed with the following measures:

- the *area measure* $\mu = (\mathbf{p}_{f,g})_*\mathrm{Leb}_I$,
- the *base measure* ν, defined as the pushforward of $\mathrm{Leb}_{[-f(\inf I), -\inf_I f] \cap \mathbb{R}}$ by the mapping $x \mapsto \mathbf{p}_{f,g}(T_x)$.

Gluing slices. In what follows, we will make a slight abuse of notation and identify intervals of the form $[a, \infty]$, $[-\infty, a]$ for $a \in \mathbb{R}$ and $[-\infty, \infty]$ with the intervals $[a, \infty)$, $(-\infty, a]$ and $\mathbb{R}$, respectively. For L, L' in the extended line $\mathbb{R} \cup \{\pm\infty\}$ such that $-f(\inf I) \le L \le L' \le -\inf_I f$, we define the restrictions $f^{(L,L')}$ and $g^{(L,L')}$ of f and g to the interval $[T_L, T_{L'}] \cap I$, yielding also a slice trajectory. We may therefore define the slice coded by $(f^{(L,L')}, g^{(L,L')})$ and denote it by

$$\mathsf{Sl}^{(L,L')} = (M^{(L,L')}, D^{(L,L')}) = (M_{f^{(L,L')}, g^{(L,L')}}, D_{f^{(L,L')}, g^{(L,L')}}).$$

We let $\mathbf{p}^{(L,L')}\colon [T_L, T_{L'}] \to M^{(L,L')}$ be the canonical projection, $\gamma^{(L,L')}$, $\xi^{(L,L')}$, $\beta^{(L,L')}$ be the maximal geodesic, shuttle, and base, and $\mu^{(L,L')}$, $\nu^{(L,L')}$ be the area and base measures of $\mathsf{Sl}^{(L,L')}$.

This family of metric spaces is compatible with the gluing operation in the following sense, illustrated in Figure 4.1.

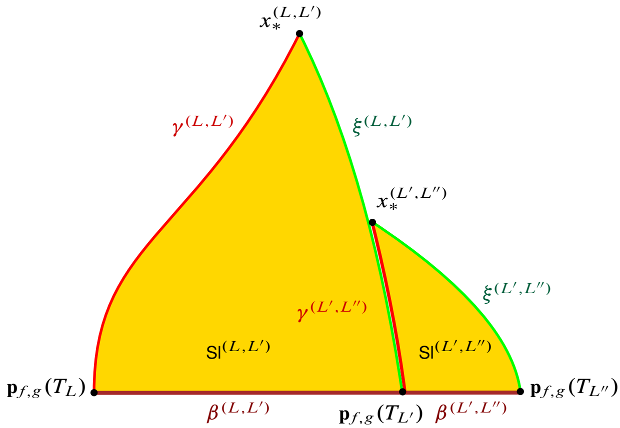

Figure 4.1. Gluing slices encoded by a slice trajectory: the gluing of $\mathsf{Sl}^{(L,L')}$ with $\mathsf{Sl}^{(L',L'')}$ results in $\mathsf{Sl}^{(L,L'')}$. Here, $T_L > -\infty$ and $T_{L''} < \infty$. We denoted by $x_*^{(L',L'')}$ the apex of $\mathsf{Sl}^{(L',L'')}$ and $x_*^{(L,L')}$ the apex of $\mathsf{Sl}^{(L,L')}$, which, on this example, is also the apex of $\mathsf{Sl}^{(L,L'')}$. Consequently, the shuttle $\xi^{(L,L'')}$ is obtained by the union of $\xi^{(L',L'')}$ and the part of $\xi^{(L,L')}$ that is not glued to $\gamma^{(L',L'')}$, whereas the maximal geodesic $\gamma^{(L,L'')} = \gamma^{(L,L')}$, as stated at the end of Proposition 4.1. The bases and measures simply add up. The fact that the slices depicted here are topological disks does not hold true in general; it will, however, be the case for the random processes we will consider in the upcoming sections.

Proposition 4.1. *Let $-f(\inf I) \le L < L' < L'' \le -\inf_I f$ be in the extended real line. Then*

$$\mathsf{Sl}^{(L,L'')} = G(\mathsf{Sl}^{(L,L')}, \mathsf{Sl}^{(L',L'')}; \xi^{(L,L')}, \gamma^{(L',L'')}).$$

Moreover, the marks and measures satisfy

$$\begin{aligned}
\gamma^{(L,L'')} &= \gamma^{(L,L')} \cup (\gamma^{(L',L'')} \setminus \xi^{(L,L')}),\\
\xi^{(L,L'')} &= \xi^{(L',L'')} \cup (\xi^{(L,L')} \setminus \gamma^{(L',L'')}),\\
\beta^{(L,L'')} &= \beta^{(L,L')} \cup \beta^{(L',L'')},\\
\mu^{(L,L'')} &= \mu^{(L,L')} + \mu^{(L',L'')},\\
\nu^{(L,L'')} &= \nu^{(L,L')} + \nu^{(L',L'')},
\end{aligned}$$

with the convention that, in the right-hand side, sets and measures are identified with their images and pushforwards by the canonical projections in $\mathsf{Sl}^{(L,L'')}$.

Proof. In the disjoint union $[T_L, T_{L'}] \sqcup [T_{L'}, T_{L''}]$, in order to avoid ambiguities due to the fact that the point $T_{L'}$ belongs to both intervals (thus should be duplicated), we use a superscript 0 for points in the first interval and a superscript 1 for points in the second interval. We observe that $d_{g^{(L,L'')}}$ can be seen as a quotient pseudometric d/R_1, where d is the disjoint union pseudometric on $[T_L, T_{L'}] \sqcup [T_{L'}, T_{L''}]$

given by $d(s,t) = d_g(s,t)$ if s, t belong to the same of the two intervals above and $d(s,t) = \infty$ otherwise, and R_1 is the coarsest equivalence relation containing

$$\{(\Xi_{T_{L'}}(r)^0, \Gamma_{T_{L'}}(r)^1),\ 0 \le r \le g(T_{L'}) - \underline{g}(T_L, T_{L'}) \vee \underline{g}(T_{L'}, T_{L''})\}.$$

Note also that, as $T_{L'}$ is a hitting time, the equivalence relation $\{d_{f(L,L'')} = 0\}$ factorizes over these two intervals, in the sense that if $d_{f(L,L'')}(s,t) = 0$ with $s \neq t$, then s, t must belong to the same interval $[T_L, T_{L'}]$ or $[T_{L'}, T_{L''}]$. So if R_2 is the equivalence relation on the above disjoint union given by $(s^i, t^j) \in R_2$ if and only if $d_f(s,t) = 0$ and $i = j \in \{0,1\}$, using (3.3), we have

$$D^{(L,L'')} = (d/R_1)/R_2 = (d/R_2)/R_1 = (D^{(L,L')} \sqcup D^{(L',L'')})/R_1,$$

which is precisely the quotient metric of $G(\mathsf{Sl}^{(L,L')}, \mathsf{Sl}^{(L',L'')}; \xi^{(L,L')}, \gamma^{(L',L'')})$.

Checking the claimed formulas for the marks and measures of $\mathsf{Sl}^{(L,L'')}$ is straightforward. ■

We finish this paragraph with a very strong identity, saying that the distances in a slice $\mathsf{Sl}^{(L,L')}$ encoded by a restriction of the slice trajectory (f, g) are in fact the restrictions of the distances in the "whole" slice $\mathsf{Sl}_{f,g}$.

Corollary 4.2. *Let (f, g) be a slice trajectory on the interval I, and $-f(\inf I) \le L \le L' \le -\inf_I f$. Then $D^{(L,L')}$ is the restriction of the function $D_{f,g}$ to $[T_L, T_{L'}]$.*

Proof. This is a direct consequence of the preceding proposition, which entails that $D_{f,g}$ is the pseudometric obtained by gluing $\mathsf{Sl}^{(L,L')}$ with $\mathsf{Sl}^{(L',\sup I)}$ along $\xi^{(L,L')}$ and $\gamma^{(L',\sup I)}$, and then by gluing the resulting space $\mathsf{Sl}^{(L,\sup I)}$ with $\mathsf{Sl}^{(\inf I,L)}$ along $\gamma^{(L,\sup I)}$ and $\xi^{(\inf I,L)}$. Since at each stage, the spaces that are glued together are isometrically embedded in the resulting gluing, we obtain that $\mathsf{Sl}^{(L,L')}$ is isometrically embedded in $\mathsf{Sl}^{(\inf I,\sup I)} = \mathsf{Sl}_{f,g}$. ■

4.2 Random continuum composite slices

We now randomize the functions f, g considered in the preceding section in various ways to construct random spaces of interest. For a fixed continuous function $f \in \mathcal{C}$ with $0 \in I(f)$, the *snake*[1] driven by f is a random centered Gaussian process $(Z^f_t, t \in I(f))$ with $Z^f_0 = 0$ and with covariance function specified by

$$\mathbb{E}[(Z^f_t - Z^f_s)^2] = d_f(s,t), \quad s,t \in I(f). \tag{4.8}$$

[1]Literally, this is rather called the "head of the snake driven by f"; see [63].

As soon as f is Hölder continuous, which will always be the case in this memoir, this process admits a continuous modification; we systematically consider this continuous modification of Z^f. If now Y is a (almost surely Hölder continuous) random function, then the random snake driven by Y is defined as the Gaussian process Z^Y conditionally given Y.

By (4.8), it holds that $Z^f_s = Z^f_t$ whenever $d_f(s,t) = 0$, so that, provided f satisfies the required limit conditions if $I(f)$ is noncompact, the pair (f, Z^f) is a slice trajectory. In what follows, we will let $(X, W)\colon (f,g) \mapsto (f,g)$ be the canonical process on $\mathcal{C}^2$.

Below and throughout this work, we use, for any process Y defined on an interval I, the piece of notation $\underline{Y}_t = \inf_{s\le t, s\in I} Y_s$.

Let us proceed to the definition of continuum slices, which arise in Theorem 2.6. Fix $A, L \in (0,\infty)$ and $\Delta \in \mathbb{R}$. We let $\mathbf{Slice}_{A,L,\Delta}$ be the probability distribution under which

- the process X is a first-passage bridge[2] of standard Brownian motion from 0 to $-L$ with duration A,
- conditionally given X, the process W has the same law as $(Z_t + \zeta_{-\underline{X}_t}, 0 \le t \le A)$, where Z is the random snake driven by $X - \underline{X}$, and $\zeta/\sqrt{3}$ is a standard Brownian bridge of duration L and terminal value $\Delta/\sqrt{3}$, independent of X and Z.

To be more precise, the process $\zeta = (\zeta_t, 0 \le t \le L)$ is Gaussian, with $\mathbb{E}[\zeta_t] = t\Delta/L$ for $t \in [0, L]$ and

$$\mathrm{Cov}(\zeta_s, \zeta_t) = 3\frac{s(L-t)}{L}, \quad 0 \le s \le t \le L.$$

With this definition, it is simple to see that $\mathbf{Slice}_{A,L,\Delta}$ is indeed carried by slice trajectories defined on the interval $I = [0, A]$.

Definition 4.3. The *(composite) slice with area A, width L and tilt Δ*, generically denoted by $\mathsf{Sl}_{A,L,\Delta}$, is the 5-marked[3] 2-measured metric space $\mathsf{Sl}_{X,W}$ under the law $\mathbf{Slice}_{A,L,\Delta}$, endowed with the marking

$$\partial \mathsf{Sl}_{A,L,\Delta} = (\beta, \gamma_0, \xi_A)$$

comprising its base and two geodesic marks, namely its maximal geodesic and its shuttle, as well as its area and base measures μ, ν.

The piece of notation $\partial \mathsf{Sl}_{A,L,\Delta}$ for the marking comes from the fact that the union of the three marks gives the topological boundary of $\mathsf{Sl}_{A,L,\Delta}$, as stated in Lemma 3.11.

[2] A *first-passage bridge of Brownian motion* can be defined from Brownian motion stopped when first hitting $-L$ by absolute continuity; see [25].

[3] Recall from Section 3.2 that each geodesic mark counts for 2 marks, the first one being its marked extremity.

4.3 The Brownian half-plane, and its embedded slices

There is a natural relation between slices and the Brownian half-plane [13,52], which we now introduce. Let $(B_t, t \geq 0)$, $(B'_t, t \geq 0)$ be two independent standard Brownian motions, and let $(\Pi_t = B'_t - 2\inf_{\{0\leq s\leq t\}} B'_s, t \geq 0)$ be the so-called *Pitman transform* of B', which is a three-dimensional Bessel process. Recall the piece of notation $\underline{X}_t = \inf_{s\leq t} X_s$. We let **Half** be the probability distribution on $\mathcal{C}^2$ under which

- the process X has same distribution as $(B_t \mathbf{1}_{\{t\geq 0\}} + \Pi_{-t}\mathbf{1}_{\{t<0\}}, t \in \mathbb{R})$, and
- conditionally given X, the process W has same distribution as $(Z_t + \zeta_{-\underline{X}_t}, t \in \mathbb{R})$, where Z is the random snake driven by $X - \underline{X}$, and $\zeta/\sqrt{3}$ is a two-sided standard Brownian motion,[4] independent of X and Z.

The measure **Half** is carried by slice trajectories defined on the interval $\mathbb{R}$. We note that we can make this definition more symmetric using standard excursion theory, in a way similar to the encoding triples of [75]. For this, we let $T_{L-} = \lim_{L'\uparrow L} T_{L'}$ and denote by

$$\begin{aligned} X^{(L)} &= (L + X_{T_{L-}+t}, 0 \leq t \leq T_L - T_{L-}), \\ W^{(L)} &= (W_{T_{L-}+t}, 0 \leq t \leq T_L - T_{L-}) \end{aligned}$$

the excursion of X above its past infimum at level $-L$ and the corresponding piece of W, respectively. Note first that the process $\zeta_L = W_{T_L}$, $L \in \mathbb{R}$, is under **Half** a standard two-sided Brownian motion multiplied by $\sqrt{3}$. Then, conditionally given ζ, the point measure on $\mathbb{R} \times \mathcal{C} \times \mathcal{C}$ given by

$$\mathcal{M}(\mathrm{d}L\,\mathrm{d}X\,\mathrm{d}W) = \sum_{L\in\mathbb{R}:T_L\neq T_{L-}} \delta_{(L,X^{(L)},W^{(L)})}$$

is a Poisson measure with intensity $2\mathrm{d}L\mathbb{N}_{\zeta_L}(\mathrm{d}(X,W))$, where $\mathbb{N}_x$ is the σ-finite "law" of the lifetime process and head of the Brownian snake (started at x) driven by the Itô measure of positive excursions of Brownian motion. The process (X, W) is then a measurable function of ζ and $\mathcal{M}$ by Itô's reconstruction theory of Brownian motion from its excursions.

Definition 4.4. The *Brownian half-plane*, which is generically denoted by BHP, is the 1-marked 2-measured metric space $\mathsf{Sl}_{X,W}$ considered under **Half**, endowed with the one mark $\partial\mathsf{BHP} = \beta$, its area measure μ and its base measure ν.

There is only one mark here, the base; there is no maximal geodesic nor shuttle since the interval of definition is $\mathbb{R}$. The name comes from the fact that BHP is

[4] This means that $(\zeta_x/\sqrt{3}, x \geq 0)$ and $(\zeta_{-x}/\sqrt{3}, x \geq 0)$ are independent (one-dimensional) standard Brownian motions issued from 0.

homeomorphic to the half-plane $\mathbb{R} \times \mathbb{R}_{\geq 0}$, its boundary as a topological manifold being equal to the base; see [13, Corollary 3.8].

In the light of Proposition 4.1, the Brownian half-plane can be seen to have a natural Markov property. First, let $\theta_t \colon f \mapsto f(t+\cdot) - f(t)$ be the translation operator on $\mathcal{C}$. We claim that **Half** is invariant under θ_{T_L}, since its action simply consists in translating by L the time in process ζ, and the first coordinate of $\mathcal{M}$, which leaves their laws invariant. For similar reasons, for every $L \in \mathbb{R}$, the processes $(X^{(0,L)}, W^{(0,L)})$, $(X^{(-\infty,0)}, W^{(-\infty,0)})$ and $(\theta_{T_L} X^{(L,+\infty)}, \theta_{T_L} W^{(L,+\infty)})$ are independent under **Half**, since they are respectively functionals of the independent random elements

$$\begin{array}{ll} (\zeta_x, 0 \leq x \leq L), & \mathcal{M}((0,L] \times \mathcal{C} \times \mathcal{C}), \\ (\zeta_x, x \leq 0), & \mathcal{M}((-\infty,0] \times \mathcal{C} \times \mathcal{C}), \\ (\zeta_{L+x} - \zeta_L, x \geq 0), & \mathcal{M}((L,\infty) \times \mathcal{C} \times \mathcal{C}). \end{array}$$

Free slices. Note that, under **Half**, the process $X^{(0,L)}$ is simply a standard Brownian motion killed at its first hitting time of $-L$, while the process $(W^{(0,L)}_{T_x}/\sqrt{3}, 0 \leq x \leq L)$ is a standard Brownian motion killed at time L. For this reason, the law of $(X^{(0,L)}, W^{(0,L)})$ under **Half** is the mixture

$$\mathbf{FSlice}_L = \int_0^\infty q_L(A)\,\mathrm{d}A \int_{\mathbb{R}} p_{3L}(\Delta)\,\mathrm{d}\Delta\, \mathbf{Slice}_{A,L,\Delta}, \tag{4.9}$$

where p_t, q_x are defined after (3.8). In what follows, a random metric space with same law as $\mathsf{Sl}^{(0,L)}$ under $\mathbf{FSlice}_L$ will be referred to as a *free (composite) slice* of width L. This, together with Proposition 4.1, yields the following result.

Proposition 4.5. *Fix $L < L' < L''$ in the extended line. Then, under* **Half***, it holds that $\mathsf{Sl}^{(L,L'')} = G(\mathsf{Sl}^{(L,L')}, \mathsf{Sl}^{(L',L'')}; \xi^{(L,L')}, \gamma^{(L',L'')})$, where the spaces $\mathsf{Sl}^{(L,L')}$, $\mathsf{Sl}^{(L',L'')}$ are independent. Moreover, if L and L' are finite, then $\mathsf{Sl}^{(L,L')}$ is a free slice of width $L' - L$.*

Recall that this result is illustrated in Figure 4.1, which can be completed by extending the brown segment into a line, letting the half-plane above be BHP, the line being its base $\beta = \beta^{(-\infty,\infty)}$. This also suggests that $\mathsf{Sl}^{(L,L')}$ is the bounded connected component of the complement of $\gamma^{(L,L')} \cup \xi^{(L,L')}$ in BHP. More precisely, the following holds.

Proposition 4.6. *For every $L < L'$ in $\mathbb{R}$, almost surely under* **Half***, the geodesics $\gamma^{(L,L')}$ and $\xi^{(L,L')}$ meet only at the apex $x_*^{(L,L')}$, and meet the base β only at their respective origins $\mathbf{p}_{X,W}(T_L)$ and $\mathbf{p}_{X,W}(T_{L'})$. Moreover, $\mathsf{Sl}^{(L,L')}$ is the closure of the bounded connected component of the complement of the union of these two paths in BHP. It is therefore homeomorphic to the closed unit disk, with boundary given by the union of the three sets $\beta^{(L,L')}$, $\gamma^{(L,L')}$ and $\xi^{(L,L')}$, which meet only at $\mathbf{p}_{X,W}(T_L)$, $\mathbf{p}_{X,W}(T_{L'})$ and $x_*^{(L,L')}$.*

Proof. This proposition is proved in the same way as [13, Lemma 6.15]. Let us recall briefly the ideas. For any point $t \in \mathbb{R}$, we let $\Sigma_t(r) = \inf\{s \geq t : X_s = X_t - r\}$ for $0 \leq r \leq X_t - \underline{X}_t$, so that the range of $\mathbf{p}_X \circ \Sigma_t$ is the geodesic path $[\![\mathbf{p}_X(t), \mathbf{p}_X(T_{-\underline{X}_t})]\!]_X$ in $\mathcal{T}_X$. Its image by π_X defines a path σ_t starting at $\mathbf{p}_{X,W}(t)$ and ending at the point $\mathbf{p}_{X,W}(T_{-\underline{X}_t})$ of the base. Moreover, almost surely, any path σ_t, $t \in \mathbb{R}$, do not intersect a geodesic γ_s, $s \in \mathbb{R}$, except possibly at the starting point of either $\mathbf{p}_{X,W}(s)$ or $\mathbf{p}_{X,W}(t)$. This implies that any point $\mathbf{p}_{X,W}(t)$ of $\mathsf{Sl}^{(L,L')}$ that is not in the union $\gamma^{(L,L')} \cup \xi^{(L,L')}$ can be linked to the bounded segment $\beta^{(L,L')}$ of the base of BHP by the path σ_t without intersecting $\gamma^{(L,L')} \cup \xi^{(L,L')}$ except perhaps at its endpoint. This latest possibility can be discarded by noting that, with probability 1, we have $T_{-\underline{X}_t} \notin \{T_L, T_{L'}\}$. Similarly, a point $\mathbf{p}_{X,W}(t)$ of BHP outside of $\mathsf{Sl}^{(L,L')}$ is linked to the unbounded set $\beta \setminus \beta^{(L,L')}$ of the base of BHP by the path σ_t, which does not intersect $\gamma^{(L,L')} \cup \xi^{(L,L')}$. This means that $\mathsf{Sl}^{(L,L')}$ is the closure of the bounded connected component of BHP minus $\gamma^{(L,L')} \cup \xi^{(L,L')}$. ■

The above discussion shows that the Brownian half-plane contains a natural "flow" of free slices. We can also link directly the slices of Section 4.2 with the Brownian half-plane via an absolute continuity argument. Recall the definitions of p_t and q_x after (4.9).

Lemma 4.7. *Let us fix $0 < K < L$, as well as $A > 0$ and $\Delta \in \mathbb{R}$. For every non-negative function G that is measurable with respect to the σ-algebra generated by $(X^{(0,K)}, W^{(0,K)})$, we have*

$$\mathbf{Slice}_{A,L,\Delta}[G] = \mathbf{Half}[\varphi_{A,L,\Delta}(T_K, K, W_{T_K}) \cdot G],$$

where

$$\varphi_{A,L,\Delta}(A', L', \Delta') = \frac{q_{L-L'}(A - A')}{q_L(A)} \frac{p_{3(L-L')}(\Delta - \Delta')}{p_{3L}(\Delta)}. \tag{4.10}$$

Proof. This comes from similar statements for Brownian bridges and first-passage bridges; see, for instance, [19, (18) and (19)]. For bounded measurable functions f, g on $\mathcal{C}$, for $0 < A' < A$ and $0 < K < L$,

$$\begin{aligned} &\mathbf{Slice}_{A,L,\Delta}[f(X|_{[0,A']}) \cdot g(\zeta|_{[0,K]})] \qquad (4.11)\\ &\quad = \mathbf{Half}\Big[f(X|_{[0,A']}) \frac{q_{L-X_{A'}}(A - A')}{q_L(A)} \mathbf{1}_{\{\underline{X}_{A'} > -L\}} \cdot g(\zeta|_{[0,K]}) \frac{p_{3(L-K)}(\Delta - \zeta_K)}{p_{3L}(\Delta)} \Big]. \end{aligned}$$

Here, the factor 3 in the index of the Gaussian density function comes from the fact that $\zeta/\sqrt{3}$ is a bridge of standard Brownian motion. We replace A' by T_K by a standard argument, writing

$$f(X^{(0,K)}) = \lim_{n \to \infty} \sum_{i \geq 0} \mathbf{1}_{\{(i-1)2^{-n} < T_K \leq i2^{-n}\}} f(X|_{[0,i2^{-n}]}),$$

using dominated convergence and applying the above equality (4.11) to $A' = i2^{-n}$, noting that $\mathbf{1}_{\{(i-1)2^{-n}<T_K\leq i2^{-n}\}} f(X|_{[0,i2^{-n}]})$ is a function of $X|_{[0,i2^{-n}]}$.

The result follows by noting that $W^{(0,K)}$ is built in the same way from $X^{(0,K)}$ and $\zeta|_{[0,K]}$ under $\mathbf{Slice}_{A,L,\Delta}$ as from $X^{(0,K)}$ and $\zeta|_{[0,K]}$ under **Half**. ■

We may now prove the statement about the topology and Hausdorff dimension of a slice.

Proof of Lemma 3.11 for slices. First, almost surely, the Brownian half-plane is homeomorphic to the half-plane [13, Corollary 3.8], is locally of Hausdorff dimension 4 and its boundary is locally of Hausdorff dimension 2. The latter facts are obtained from similar statements for Brownian disks [21] thanks to [13, Theorem 3.7] allowing us to couple arbitrary balls of BHP centered at the root $\mathbf{p}_{X,W}(0)$ with balls of large enough Brownian disks, centered at a point on the boundary.

Hence, under the probability distribution **Half**, for any $L < L'$, by Corollary 4.2, the metric space $\mathsf{Sl}^{(L,L')}$ is almost surely locally of Hausdorff dimension 4 and its base $\beta^{(L,L')}$ is locally of Hausdorff dimension 2. Furthermore, it is homeomorphic to the disk by Proposition 4.6 and its boundary is the union of its three marks $\beta^{(L,L')}$, $\gamma^{(L,L')}$ and $\xi^{(L,L')}$, whose pairwise intersections are identified singletons.

Now, arguing under $\mathbf{Slice}_{A,L,\Delta}$, we use the fact from Proposition 4.1 that $\mathsf{Sl}^{(0,L)} = G(\mathsf{Sl}^{(0,L/2)}, \mathsf{Sl}^{(L/2,L)}; \xi^{(0,L/2)}, \gamma^{(L/2,L)})$. Lemma 4.7 entails that, almost surely, under this probability distribution, the law of $\mathsf{Sl}^{(0,L/2)}$ is absolutely continuous with respect to that of the same random variable under **Half**, and so is homeomorphic to a disk. Now, we observe that, under $\mathbf{Slice}_{A,L,\Delta}$, the process $\theta_{T_{L/2}}(X^{(L/2,L)}, W^{(L/2,L)})$ has same distribution as $(X^{(0,L/2)}, W^{(0,L/2)})$, which we leave as an exercise to the reader. Therefore, under this law, $\mathsf{Sl}^{(L/2,L)}$ has same distribution as $\mathsf{Sl}^{(0,L/2)}$ and both are homeomorphic to a disk. We conclude that the same is true for $\mathsf{Sl}^{(0,L)}$ since it is obtained by gluing two topological disks along two segments of their boundaries. The identification of the marks given in Proposition 4.1 easily yields the desired property on the marks of $\mathsf{Sl}^{(0,L)}$. The facts on the local Hausdorff dimension are obtained similarly. ■

4.4 The uniform infinite half-planar quadrangulation

We now define a slight variant of the classical uniform infinite half-planar quadrangulation (UIHPQ) [13,14,33,43], the half-planar version of the uniform infinite random planar quadrangulation, in the following way. Let $F_\infty = (\mathbf{T}^k, k \in \mathbb{Z})$ be a two-sided sequence of independent Bienaymé–Galton–Watson trees with a geometric offspring distribution of parameter $1/2$. Conditionally on F_∞, we let λ^0_∞ be a uniformly chosen well-labeling condition function, meaning that every tree $\mathbf{T}^k$ is assigned a well-labeling condition function giving label 0 to its root vertex, independently, uniformly at random. Lastly, and independently of F_∞ and λ^0_∞, we let $(b_k, k \in \mathbb{Z})$ be a doubly-infinite walk with shifted geometric steps, meaning that $b_0 = 0$ almost surely, and that $b_k - b_{k-1}$, $k \in \mathbb{Z}$, are independent and identically distributed random variables with $\mathbb{P}(b_1 = r) = 2^{-r-2}$ for every $r \in \{-1, 0, 1, 2, \dots\}$. For a vertex $v \in \mathbf{T}^k$, we let $\lambda_\infty(v) = b_k + \lambda^0_\infty(v)$, and call $(F_\infty, \lambda_\infty)$ the *infinite random well-labeled forest*. We then embed F_∞ in the plane in such a way that all trees are contained in the upper half-plane, and the root ρ^k of $\mathbf{T}^k$ is located at the point $(k, 0) \in \mathbb{R}^2$. We also link consecutive roots ρ^k, ρ^{k+1} by a line segment. We then let $(c_i, i \in \mathbb{Z})$ be the sequence of corners of the upper half-plane part of the resulting map, in contour order from left to right, with origin the first corner c_0 incident to ρ^0. The UIHPQ is then the infinite map Q_∞ obtained by applying the CVS construction to $(F_\infty, \lambda_\infty)$, that is, by linking every corner to its successor as defined in Section 2.1, and removing all edges of the forest afterward. The root of Q_∞ is defined as the corner preceding the arc from c_0 to its successor. Note that, in this case, there is no need to add an extra vertex with a corner c_∞.

Remark 4.8. The difference between this definition of the UIHPQ and the one appearing in the mentioned references is a slight rooting bias. Indeed, the simplest way to obtain the usual definition is to consider a two-sided simple random walk $(z_i, i \in \mathbb{Z})$ and construct the sequence $(b_k, k \in \mathbb{Z})$ from it as follows. Let $S^\downarrow = \{i \in \mathbb{Z} : z_{i+1} - z_i = -1\}$ be the set of descending steps of $(z_i, i \in \mathbb{Z})$ and $i_0 = \sup(S^\downarrow \cap \mathbb{Z}_{\le 0})$ the index of the descending step preceding 0. Then we define the sequence $(b_k, k \in \mathbb{Z})$ by reindexing $(z_i - z_{i_0}, i \in S^\downarrow)$ with $\mathbb{Z}$ in such a way that i_0 corresponds to the index 0. The UIHPQ is then constructed as above with this bridge but rooted at the corner preceding the arc linking $s^{-i_0}(c_0)$ to its successor $s^{-i_0+1}(c_0)$ instead of the convention we presented. Apart from this slight root shift, the resulting law of $(b_k, k \in \mathbb{Z})$ is not exactly that of a doubly-infinite bridge with shifted geometric steps. The first step gets a size-biased distribution $\mathbb{P}(b_1 = r) = (r + 2)2^{-r-3}$, $r \ge -1$, whereas all other steps get the desired shifted geometric distribution. See the discussion in [13, Section 4.5.2] for more information.

The construction we use here has the advantage of making the law of the slices invariant by translation.

Convergence toward the Brownian half-plane. Denoting by v_i the vertex of F_∞ incident to c_i and by $\Upsilon(i) \in \mathbb{Z}$ the index of the tree to which v_i belongs, we define the *contour* and *label processes* on $\mathbb{R}$ by

$$C(i) = d_{\mathbf{T}^{\Upsilon(i)}}(v_i, \rho^{\Upsilon(i)}) - \Upsilon(i) \quad \text{and} \quad \Lambda(i) = \lambda_\infty(v_i), \quad i \in \mathbb{Z},$$

and by linear interpolation between integer values; see Figure 4.2. As is well known, the part of the contour process corresponding to $\mathbf{T}^k$ (counting the edge linking ρ^k to ρ^{k+1}) has the same distribution as a simple random walk started at $-k$ and killed

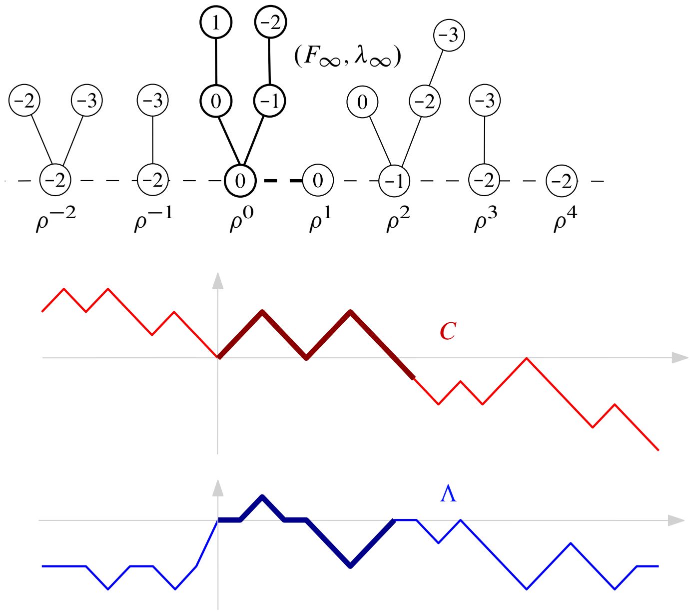

Figure 4.2. Contour and label processes associated with $(F_\infty, \lambda_\infty)$. The edges of the floor are represented with dashed lines. The tree $\mathbf{T}^0$ and the corresponding part of the encoding processes are highlighted (the corresponding floor edge and the final descending step of the contour function are also highlighted). For instance, $\tau_3 = 17$ on this example. The contour process can be thought of as recording the height of a particle moving at speed one around the forest. In this point of view, the root ρ^k should be at height $-k$ for each $k \in \mathbb{Z}$; this can be achieved, for instance, by vertically translating each tree in such a way that ρ^k is mapped to location $(k, -k)$ instead of $(k, 0)$.

when first hitting $-k-1$. Finally, for $k \geq 1$, we denote by τ_k the hitting time of $-k$ by C; its value is thus also equal to k plus twice the number of edges in the first k trees $\mathbf{T}^0, \dots, \mathbf{T}^{k-1}$.

As the vertices of the encoding objects are preserved through the CVS bijection, the vertex v_i can also be seen as a vertex of Q_∞. Let us define

$$D_\infty(i, j) = d_{Q_\infty}(v_i, v_j), \quad i, j \in \mathbb{Z}.$$

We extend D_∞ to a continuous function on $\mathbb{R}^2$ by “bilinear interpolation,” writing $\{s\} = s - \lfloor s \rfloor$ for the fractional part of s and then setting

$$\begin{aligned} D_\infty(s,t) = {} & (1-\{s\})(1-\{t\})D_\infty(\lfloor s \rfloor, \lfloor t \rfloor) + \{s\}(1-\{t\})D_\infty(\lfloor s \rfloor + 1, \lfloor t \rfloor) \\ & + (1-\{s\})\{t\}D_\infty(\lfloor s \rfloor, \lfloor t \rfloor + 1) \\ & + \{s\}\{t\}D_\infty(\lfloor s \rfloor + 1, \lfloor t \rfloor + 1). \end{aligned} \tag{4.12}$$

We then define the renormalized versions of C, Λ, and D_∞: for every $s, t \in \mathbb{R}$, we set

$$C_{(n)}(s) = \frac{C(2ns)}{\sqrt{2n}}, \quad \Lambda_{(n)}(s) = \frac{\Lambda(2ns)}{(8n/9)^{1/4}}, \quad D_{(n)}(s,t) = \frac{D_\infty(2ns, 2nt)}{(8n/9)^{1/4}}. \quad (4.13)$$

The next result can be seen as a reformulation of [52, Theorem 1.11] or [13, Theorem 3.6], proving the convergence of the UIHPQ to the Brownian half-plane defined in Section 4.3.

Proposition 4.9. *On $\mathcal{C} \times \mathcal{C} \times \mathcal{C}^{(2)}$, it holds that*

$$(C_{(n)}, \Lambda_{(n)}, D_{(n)}) \xrightarrow[n\to\infty]{(d)} (X, W, D_{X,W}), \quad (4.14)$$

where the limiting triple is understood under **Half**.

This statement does not appear in this exact form in the aforementioned references, which do not explicitly focus on the processes $C_{(n)}$, $\Lambda_{(n)}$, X, W. In [13, Remark 6.16], it was however mentioned how to extend the results therein in order to take into account these processes, so we will follow the line of reasoning sketched in that work.

Proof. The proof proceeds via established convergence results for random quadrangulations with one external face to Brownian disks. Fix some number $K > 0$. We will sample a quadrangulation with one external face, whose areas and perimeters are so large that, in a neighborhood of 0 of amplitude K, this rescaled large quadrangulation and its limit, a Brownian disk of large area and perimeter, are indistinguishable from the rescaled UIHPQ and the Brownian half-plane, in a sense to be made precise. In the following, we will use for all the objects related to the quadrangulation with one external face or the limiting Brownian disk a similar notation as for those related to the UIHPQ or the Brownian half-plane, only with a superscript prime symbol $'$.

Fix $L > 0$, which should be thought of as being large. For $n \geq 1$, we sample the aforementioned quadrangulation Q'_n with one external face as follows. First, consider a uniform random element (M'_n, λ'_n) of $\vec{\mathbf{M}}^{[0]}_{a_n,(l_n)}$, where

$$a_n = \lfloor nL \rfloor \quad \text{and} \quad l_n = \lfloor L\sqrt{2n} \rfloor.$$

We can view this as a labeled forest (F'_n, λ'_n) with l_n trees arranged in a circle, and rooted at $\rho^0, \dots, \rho^{l_n-1}$, where ρ^0 is the root of the tree containing the root corner of f_*. We let C'_n, Λ'_n be the contour and label process of this forest, defined as above, starting from the tree rooted at ρ^0. We let $Q'_n = \mathrm{CVS}(M'_n, \lambda'_n; f_*)$ be the rooted quadrangulation encoded by (M'_n, λ'_n), and we let $D'_n(i,j) = d_{Q'_n}(v'_i, v'_j)$ for $0 \leq i, j \leq 2a_n + l_n$, where v'_i is the i-th visited vertex of F'_n in contour order, viewed as a vertex of Q'_n. As usual, we extend D'_n into a continuous function on

$[0, 2a_n + l_n]^2$. Finally, we extend the definition of these processes to the interval $[-2a_n - l_n, 2a_n + l_n]$ by the simple translation formulas

$$C'_n(t) = C'_n(t + 2a_n + l_n) + l_n, \quad \Lambda'_n(t) = \Lambda'_n(t + 2a_n + l_n), \quad t \in [-2a_n - l_n, 0], \tag{4.15}$$

$$D'_n(s,t) = D'_n(s + (2a_n + l_n)\mathbf{1}_{\{s<0\}}, t + (2a_n + l_n)\mathbf{1}_{\{t<0\}}) \quad s,t \in [-2a_n - l_n, 2a_n + l_n]. \tag{4.16}$$

The idea behind this extension is that we are going to consider these processes in neighborhoods of 0, so that we are really interested in the behavior of these processes when the argument is close from 0 or from $2a_n + l_n$.

Define their rescaled versions: for $s, t \in [-2a_n - l_n, 2a_n + l_n]$,

$$C'_{(n)}(s) = \frac{C'_n(2ns)}{\sqrt{2n}}, \quad \Lambda'_{(n)}(s) = \frac{\Lambda'_n(2ns)}{(8n/9)^{1/4}}, \quad D'_{(n)}(s,t) = \frac{D'_n(2ns, 2nt)}{(8n/9)^{1/4}}. \tag{4.17}$$

Then by [25, (26) and Theorem 20], one has the joint convergence

$$(C'_{(n)}, \Lambda'_{(n)}, D'_{(n)}) \xrightarrow[n\to\infty]{(d)} (X', W', D') \tag{4.18}$$

in distribution in $\mathcal{C}([0,L]) \times \mathcal{C}([0,L]) \times \mathcal{C}([0,L]^2)$, where (X', W', D') is an explicit limiting process, which is the encoding process of the Brownian disk of area L and width $\sqrt{L}$. In particular, the process D' is a measurable function of the pair (X', W'). Due to the formulas in (4.15) and (4.16), this implies the convergence of these processes on $\mathcal{C}([-L,L]) \times \mathcal{C}([-L,L]) \times \mathcal{C}([-L,L]^2)$, where (X', W', D') are extended to functions on $[-L, L]$ or $[-L, L]^2$ in a similar way as above. Note that we choose to omit the dependence of (X', W', D') on L for lighter notation, but we will need later to choose L appropriately.

Now recall that $K > 0$ is a fixed number. The first crucial observation is that we may choose L large enough, so that with high probability, the laws of the restrictions of (X', W', D') and $(X, W, D_{X,W})$ to the interval $[-K, K]$ are very close. More precisely, given $\varepsilon \in (0,1)$, fix $r > 0$ and $A > 0$ such that

$$\mathbb{P}\Big(\max_{-K\le t\le K} D_{X,W}(0,t) > r\Big) < \frac{\varepsilon}{3}, \quad \mathbb{P}(T_{-A} < -K < K < T_A) \ge 1 - \frac{\varepsilon}{3}.$$

Then [13, Proposition 6.6] and its proof (Lemmas 6.7 and 6.8) show that there exists $L_0 > 0$ such that, for $L > L_0$, the two processes (X, W) and (X', W') can be coupled in such a way that on some event $\mathcal{F}$ of probability $\mathbb{P}(\mathcal{F}) \ge 1 - \varepsilon/3$, we have

$$X_t = X'_t, \quad W_t = W'_t, \quad D_{X,W}(s,t) = D'(s,t), \tag{4.19}$$

for every $s, t \in [T_{-A}, T_A]$ such that $\max(D_{X,W}(0,t), D_{X,W}(0,s)) \le r$. Given our choice of r, A, we see that with probability at least $1 - \varepsilon$, (4.19) holds for every $s, t \in [-K, K]$.

Our second important observation is that, still with K and ε fixed, and possibly up to choosing L even larger than the above, albeit in a way that does not depend on n, the laws of $(C_{(n)}, \Lambda_{(n)}, D_{(n)})$ and $(C'_{(n)}, \Lambda'_{(n)}, D'_{(n)})$ in restriction to the interval $[-K, K]$ are also close, in the sense that they can be coupled in such a way that these restrictions coincide with probability at least $1-\varepsilon$. This follows from the proof of [13, Theorem 3.6], a minor difference being that this proposition establishes that the balls of radius $(8n/9)^{1/4}r$ centered at the root in Q_∞ and Q'_n are isometric, rather than giving a statement on $D_{(n)}$ and $D'_{(n)}$. Therefore, in order to show that the latter coincide on $[-K, K]$, one again has to choose in the first place a radius $r > 0$ so that uniformly over n, with probability at least $1-\varepsilon/3$, the vertices v'_i for integers i lying in $[-2Kn, 2Kn]$ (where we naturally let $v'_i = v'_{i+2a_n+l_n}$ for $i \leq 0$), all belong to this ball. The existence of such an r is guaranteed by the convergence (4.18) and the continuity of D'. Finally, we see that both sides of (4.18) can be coupled in such a way that with probability at least $1-\varepsilon$, they coincide with both sides of (4.14). Since ε was arbitrary, we conclude that (4.14) holds in restriction to $[-K, K]$. Since K was arbitrary, this concludes the proof. ■

Seeing a slice as part of the UIHPQ. We consider a fixed $L > 0$ and a sequence $(l_n) \in \mathbb{N}^{\mathbb{N}}$ such that

$$\frac{l_n}{\sqrt{2n}} \underset{n\to\infty}{\longrightarrow} L$$

and, for each n, we let (F_n, λ_n) be the random well-labeled forest obtained by keeping only the labeled trees $\mathbf{T}^0, \ldots, \mathbf{T}^{l_n-1}$ of the infinite random well-labeled forest $(F_\infty, \lambda_\infty)$, as well as the root ρ^{l_n} of the tree $\mathbf{T}^{l_n}$. In particular, the forest F_n has l_n independent Bienaymé–Galton–Watson trees with Geometric$(1/2)$ offspring distribution, and the labels of the root vertices of the trees (including ρ^{l_n}) follow a random walk of length l_n whose step distribution is a shifted Geometric$(1/2)$ given by $\mathbb{P}(\cdot = r) = 2^{-r-2}$ for $r \geq -1$.

Recall that $(c_i, i \in \mathbb{Z})$ denotes the sequence of corners of the infinite random well-labeled forest $(F_\infty, \lambda_\infty)$ and that v_i is the vertex of F_∞ incident to c_i. According to the construction of Section 2.3, (F_n, λ_n) encodes a slice Q_n, which is part of the UIHPQ Q_∞ constructed from the whole infinite forest $(F_\infty, \lambda_\infty)$. More precisely, the maximal geodesic (resp. shuttle) can be read inside the UIHPQ as the chain of arcs linking c_0 (resp. $c_{\tau_{l_n}}$) to its subsequent successors[5] and the edges of the slice are given by the arcs from c_i to $s(c_i)$ for $0 \leq i < \tau_{l_n}$. As a consequence, the vertex v_i can be seen both as a vertex of Q_∞ and as a vertex of Q_n for $0 \leq i \leq \tau_{l_n}$.

Furthermore, we can check that Q_n is in fact isometrically embedded in Q_∞ in the sense that, whenever $0 \leq i, j \leq \tau_{l_n}$, it holds that $d_{Q_n}(v_i, v_j) = D_\infty(i, j)$.

[5] Recall Section 2.1.

Indeed, similarly to Proposition 4.1, Q_∞ can be obtained as the gluing of the infinite quadrangulation corresponding to the trees $\mathbf{T}^k$, $k < 0$, of $(F_\infty, \lambda_\infty)$, with Q_n and then with the infinite quadrangulation corresponding to the trees $\mathbf{T}^k$, $k \geq l_n$, of $(F_\infty, \lambda_\infty)$ along the proper shuttles and maximal geodesics. Alternatively, one may also argue that there are no shortcuts outside Q_n: for $0 \leq i, j \leq \tau_{l_n}$, any path linking v_i to v_j in Q_∞ may be shortened to a path that stays within Q_n since the maximal geodesic and shuttle are geodesics and since the path $c_0 \to s(c_0) \to s^2(c_0) \to \cdots$ is a geodesic ray that disconnects Q_∞.

The contour function, label function and pseudometric corresponding to Q_n are thus obtained by restricting to $[0, \tau_{l_n}]$ the analog functions corresponding to Q_∞. After rescaling, their joint limit is a direct consequence of Proposition 4.9.

Corollary 4.10. *On $\mathcal{C} \times \mathcal{C} \times \mathcal{C}^{(2)}$, it holds that*

$$(C_{(n)}|_{[0,\tau_{l_n}/2n]}, \Lambda_{(n)}|_{[0,\tau_{l_n}/2n]}, D_{(n)}|_{[0,\tau_{l_n}/2n]^2}) \xrightarrow[n\to\infty]{(d)} (X^{(0,L)}, W^{(0,L)}, D^{(0,L)}),$$

where we used the notation of Section 4.1, that is,

(a) *the pair $(X^{(0,L)}, W^{(0,L)})$ is the restriction to the interval $[0, T_L]$ of (X, W) distributed under* **Half**,

(b) $D^{(0,L)} = D_{X^{(0,L)},W^{(0,L)}}$ *is the random pseudometric on $\mathbb{R}$ defined by (4.5).*

Proof. By the Skorokhod representation theorem, we may and will assume that the convergence (4.14) holds almost surely. Classically, the almost sure path properties of X at time T_L, namely, the fact that X immediately visits the interval $(-L-\varepsilon, -L)$ after time T_L, yield that $\tau_{l_n}/2n$ almost surely converges to T_L. Corollary 4.2 then yields the result. ∎

4.5 Scaling limit of conditioned slices

We now derive Theorem 2.6 from the results of the previous section by standard conditioning arguments.

Convergence of the encoding processes. First, without loss of generality, we may assume that the contour and label processes (C, Λ) of the infinite random well-labeled forest defined in Section 4.4 are the canonical processes, considered under the probability distribution $\mathbb{P}_\infty$ on the canonical space. Next, for $a, l \in \mathbb{N}$ and $\delta \in \mathbb{Z}$, we denote by $\mathbb{P}_{a,l,\delta}$ the distribution of $(C|_{[0,2a+l]}, \Lambda|_{[0,2a+l]})$, where (C, Λ) is distributed under $\mathbb{P}_\infty[\,\cdot \mid \tau_l = 2a + l, \Lambda(\tau_l) = \delta]$. The corresponding forest encoded by this random process is thus composed of l Bienaymé–Galton–Watson trees with Geometric$(1/2)$ offspring distribution and uniform admissible labels, conditioned on the fact that the total number of edges in the trees is a and the label of the root of the last vertex-tree is δ. Similarly to the slice it encodes, we will say that the forest has *tilt* δ.

For every measurable nonnegative functional G, it thus holds that

$$\mathbb{E}_{a,l,\delta}[G] = \mathbb{E}_\infty[G((C(k), \Lambda(k)), 0 \le k \le \tau_l) \mid \tau_l = 2a + l,\ \Lambda(\tau_l) = \delta].$$

Let $(\mathcal{F}_k, k \ge 0)$ be the natural filtration associated with the canonical process (C, Λ). Note that $((C(k), \Lambda(k)), 0 \le k \le \tau_l)$ is the pair of contour and label processes of the first l trees in the forest, and that $\mathcal{F}_{\tau_l}$ is the σ-algebra generated by these first l trees (with their labels and that of the root ρ^l). Recall from Proposition 2.3 the definitions of Q_ℓ, P_ℓ.

Lemma 4.11. *Fix $0 < k < l$, as well as $a \in \mathbb{N}$ and $\delta \in \mathbb{Z}$. For every nonnegative functional G that is $\mathcal{F}_{\tau_k}$-measurable, we have*

$$\mathbb{E}_{a,l,\delta}[G] = \mathbb{E}_\infty[\Phi_{a,l,\delta}(\tau_k, k, \Lambda(\tau_k)) \cdot G],$$

where

$$\Phi_{a,l,\delta}(t, l', j) = \frac{Q_{l-l'}(2a + l - t)}{Q_l(2a + l)} \frac{P_{l-l'}(\delta - j)}{P_l(\delta)}.$$

Proof. It suffices to prove the result when G is the indicator of the contour and label processes of a given well-labeled forest with l' trees, $(t - l')/2$ edges, and tilt j. In this case, $\mathbb{E}_{a,l,\delta}[G]$ is equal to the number of ways in which one can complete this labeled forest into a well-labeled forest with l trees, a edges and tilt δ, which is the number of forests with $l - l'$ trees, $a + (l' - t)/2$ edges and tilt $\delta - j$, divided by the number of well-labeled forests with l trees, a edges and tilt δ. We conclude by Proposition 2.3. ∎

In addition to the already fixed sequence (l_n), we consider two sequences (a_n), (δ_n) satisfying (2.4). We will need the following direct consequence of the local limit theorem [26, Theorem 8.4.1]. Recall the definition of $\varphi_{A,L,\Delta}$ given in (4.10).

Lemma 4.12. *If the integer-valued sequence (l'_n) satisfies $l'_n/\sqrt{2n} \to L' \in (0, L)$, it holds that*

$$\sup_{0 \le t \le a_n,\, j \in \mathbb{Z}} \left| \Phi_{a_n,l_n,\delta_n}(t, l'_n, j) - \varphi_{A,L,\Delta}\left(\frac{t}{n}, L', \left(\frac{9}{8n}\right)^{1/4} j\right) \right| \underset{n\to\infty}{\longrightarrow} 0.$$

We start with the following conditioned version of Corollary 4.10.

Proposition 4.13. *On $\mathcal{C} \times \mathcal{C} \times \mathcal{C}^{(2)}$, the triple $(C_{(n)}, \Lambda_{(n)}, D_{(n)}|_{[0,\tau_{l_n}/2n]^2})$ considered under $\mathbb{P}_{a_n,l_n,\delta_n}$ converges in distribution to $(X, W, D_{X,W})$, considered under* $\mathbf{Slice}_{A,L,\Delta}$.

Proof. The joint convergence of the first two coordinates is standard; see, for example, [19, Corollary 16]. Let us fix $\varepsilon \in (0, L)$, define $l_n^\varepsilon = l_n - \lfloor \varepsilon\sqrt{2n} \rfloor$, so that $l_n^\varepsilon/\sqrt{2n} \to L - \varepsilon$, and set $D_{(n)}^\varepsilon = D_{(n)}|_{[0,\tau_{l_n^\varepsilon}/2n]^2}$ and $D_{(n)}^0 = D_{(n)}|_{[0,\tau_{l_n}/2n]^2}$.

By the usual bound (2.1), for every $i, j \in [0, \tau_{l_n}]$,

$$\begin{aligned}|D_\infty(i \wedge \tau_{l_n^\varepsilon}, j \wedge \tau_{l_n^\varepsilon}) - D_\infty(i, j)| &\le D_\infty(i, i \wedge \tau_{l_n^\varepsilon}) + D_\infty(j, j \wedge \tau_{l_n^\varepsilon}) \\ &\le 4(\omega(\Lambda_n; \tau_{l_n} - \tau_{l_n^\varepsilon}) + 1),\end{aligned}$$

where $\omega(f; \cdot)$ denotes the modulus of continuity of f. This implies that

$$\mathrm{dist}_{\mathcal{C}^{(2)}}(D^\varepsilon_{(n)}, D^0_{(n)}) \le \frac{\tau_{l_n} - \tau_{l_n^\varepsilon}}{2n} + 4\omega\Big(\Lambda_{(n)}, \frac{\tau_{l_n} - \tau_{l_n^\varepsilon}}{2n}\Big) + \mathcal{O}(n^{-1/4}).$$

From the joint convergence of the first two coordinates, we have, for every $\eta > 0$,

$$\begin{aligned}&\limsup_{n\to\infty} \mathbb{P}_{a_n,l_n,\delta_n}(\mathrm{dist}_{\mathcal{C}^{(2)}}(D^\varepsilon_{(n)}, D^0_{(n)}) \ge \eta) \\ &\qquad \le \mathbf{Slice}_{A,L,\Delta}(A - T_{L-\varepsilon} + 4\omega(W; A - T_{L-\varepsilon}) \ge \eta),\end{aligned}$$

which tends to 0 as $\varepsilon \to 0$ since $T_{L-\varepsilon} \to T_L = A$ almost surely under $\mathbf{Slice}_{A,L,\Delta}$. We next show that $D^\varepsilon_{(n)}$ under $\mathbb{P}_{a_n,l_n,\delta_n}$ converges in distribution to $D^{(0,L-\varepsilon)}$ under $\mathbf{Slice}_{A,L,\Delta}$, and use the principle of accompanying laws [92, Theorem 9.1.13] to conclude that, jointly with the convergence of $(C_{(n)}, \Lambda_{(n)})$ to (X, W), the process $D^0_{(n)}$ converges to the distributional limit of $D^{(0,L-\varepsilon)}$ as $\varepsilon \to 0$, which is none other than $D^{(0,L)}$, due to Corollary 4.2.

To prove the claimed convergence of $D^\varepsilon_{(n)}$ to $D^{(0,L-\varepsilon)}$, we denote by $C^\varepsilon_{(n)}$ and $\Lambda^\varepsilon_{(n)}$ the restrictions of $C_{(n)}$ and $\Lambda_{(n)}$ to $[0, \tau_{l_n^\varepsilon}/2n]$ and let F be a nonnegative bounded continuous function. Using Lemma 4.11, then Corollary 4.10 (for the choice of $L - \varepsilon$ instead of L) and Lemma 4.12 gives

$$\begin{aligned}\mathbb{E}_{a_n,l_n,\delta_n}[F(C^\varepsilon_{(n)}, \Lambda^\varepsilon_{(n)}, D^\varepsilon_{(n)})] &= \mathbb{E}_\infty[\Phi_{a_n,l_n,\delta_n}(\tau_{l_n^\varepsilon}, l_n^\varepsilon, \Lambda(\tau_{l_n^\varepsilon}))F(C^\varepsilon_{(n)}, \Lambda^\varepsilon_{(n)}, D^\varepsilon_{(n)})] \\ &\xrightarrow[n\to\infty]{} \mathbf{Half}[\varphi_{A,L,\Delta}(T_{L-\varepsilon}, L-\varepsilon, W_{T_{L-\varepsilon}})F(X^{(0,L-\varepsilon)}, W^{(0,L-\varepsilon)}, D^{(0,L-\varepsilon)})],\end{aligned}$$

the latter being equal to $\mathbf{Slice}_{A,L,\Delta}[F(X^{(0,L-\varepsilon)}, W^{(0,L-\varepsilon)}, D^{(0,L-\varepsilon)})]$ according to Lemma 4.7. ■

GHP convergence. We infer from Proposition 4.13 the GHP convergence of Theorem 2.6 by a standard method. First, by Skorokhod's representation theorem, we may assume that we are working on a probability space on which the convergence of Proposition 4.13 is almost sure. We let $\mathsf{Sl}_{A,L,\Delta}$ be the continuum slice coded by the limiting process, and Sl_n be the slice encoded by the forest whose rescaled contour and label processes make up the pair $(C_{(n)}, \Lambda_{(n)})$. As mentioned before Corollary 4.10, Sl_n is isometrically embedded in Q_∞, so that the process $D_{(n)}|_{[0,\tau_{l_n}/2n]^2}$ under $\mathbb{P}_{a_n,l_n,\delta_n}$ projects into the metric of $\Omega_n(\mathrm{Sl}_n)$.

Then, from this almost sure convergence, we easily deduce that the distortion of the correspondence $\mathcal{R}_n$ given by

$$\mathcal{R}_n = \{(v_{\lfloor(2a_n+l_n)s\rfloor}, \mathbf{p}_{X,W}(As)) : s \in [0, 1]\}$$

between $\Omega_n(\mathrm{Sl}_n)$ minus its shuttle and $\mathsf{Sl}_{A,L,\Delta}$ tends to 0 as $n \to \infty$. Forgetting the marks and measures, this gives the desired convergence in the 0-marked Gromov–Hausdorff topology.

In order to include the marking and measures, we use the technique of enlargement of correspondences already used in the proof of Lemma 3.9. Namely, we fix $\varepsilon > 0$ and let $\mathcal{R}_n^\varepsilon$ be the set of points of the form (v, x) in $\mathrm{Sl}_n \times \mathsf{Sl}_{A,L,\Delta}$ such that there exists $(w, y) \in \mathcal{R}_n$ satisfying $d_{\mathrm{Sl}_n}(v, w) < (8n/9)^{1/4}\varepsilon$ and $D(x, y) < \varepsilon$. As before, the distortion of $\mathcal{R}_n^\varepsilon$ is at most $\mathrm{dis}(\mathcal{R}_n) + 4\varepsilon$. Let us start with the marks.

Marks. For a function $f \in \mathcal{C}$ defined over the interval I, we say that $s \in I$ is a *left-minimum* of f if $f(t) \geq f(s)$ for every $t \leq s$ in I, and we call it strict if $f(t) > f(s)$ for $t < s$ in I. Note that the points of the form v_i and $\mathbf{p}_{X,W}(s)$, where i and s are left-minimums of Λ_n and W respectively belong to the maximal geodesics of Sl_n and $\mathsf{Sl}_{A,L,\Delta}$, and that all points in these sets are in fact of this form, where we can even require the stronger property that i and s are strict left-minimums.

By the uniform convergence of $\Lambda_{(n)}$ toward W, for every $\eta > 0$, the following holds provided $n \geq n_0$ for some n_0: every strict left-minimum of $\Lambda_{(n)}$ is at distance at most $\eta/2$ from some (not necessarily strict) left-minimum of W, and vice-versa, exchanging the roles of $\Lambda_{(n)}$ and W. Up to increasing n_0, we furthermore assume that $|(2a_n + l_n)/2n - A| < \eta/2$ as soon as $n \geq n_0$. Choosing η small enough so that $|D_{(n)}(s,t) - D_{(n)}(s',t')| \leq \varepsilon$ for every n, $|s - s'| \leq \eta$, $|t - t'| \leq \eta$, we deduce that the extended correspondence $\mathcal{R}_n^\varepsilon$ is compatible with the maximal geodesics for $n \geq n_0$.

The argument is similar for the shuttles. This time, we note that elements of the shuttle of $\mathsf{Sl}_{A,L,\Delta}$ are of the form $\mathbf{p}_{X,W}(s)$, where s is a right-minimum of the function W (with an obvious definition), while elements of the shuttle of Sl_n are at distance 1 from points of the form v_i, where i is a right-minimum of the function Λ_n.

The mark corresponding to the base is also treated similarly. Recall from Section 2.3 that vertices of the base are at distance at most $B_n = \max_{1 \leq i \leq l_n} |\Lambda_n(\rho^i) - \Lambda_n(\rho^{i-1})| + 1$ from some element of the floor $\{\rho^0, \ldots, \rho^{l_n}\}$ of the forest coding the slice. The process of labels $(\Lambda_n(\rho^i), 0 \leq i \leq l_n)$ forms a random walk with shifted Geometric$(1/2)$ increments conditioned to be equal to δ_n at time l_n, so, under our assumptions, it converges, after rescaling by $\sqrt{2n}$ in time and $(8n/9)^{1/4}$ in space, to a continuous process (which is easily checked to be the Brownian bridge $\zeta = (W_{T_x}, 0 \leq x \leq L)$), so that $B_n = o(n^{1/4})$ almost surely. Therefore, the base of Sl_n is at Hausdorff distance $o(n^{1/4})$ from the floor $\{\rho^i, 0 \leq i \leq l_n\}$. In turn, these vertices are exactly those of the form v_i, where i is a left-minimum of the contour process C_n. Moreover, by definition, the base of $\mathsf{Sl}_{A,L,\Delta}$ consists of the points $\mathbf{p}_{X,W}(s)$, where s is a left-minimum of the process X. Therefore, the same argument as for the maximal geodesic – replacing the processes Λ_n and W by C_n and X – shows that, almost surely, for every n large enough, the correspondence $\mathcal{R}_n^\varepsilon$ is also compatible with the bases of Sl_n and of $\mathsf{Sl}_{A,L,\Delta}$.

Measures. Finally, let us deal with the convergence of the measures, starting with the area measure. To this end, note that, for t in $[0, 2a_n + l_n]$, the contour process C_n at time t has either a left derivative equal to $+1$ or to -1. Letting $i_t^n = \lceil t \rceil$ in the former case and $i_t^n = \lfloor t \rfloor$ in the latter case, the image of $\mathrm{Leb}_{[0,(2a_n+l_n)/2n]}$ by $t \mapsto v_{i_{2nt}^n}$ is the counting measure on the set of all nonfloor vertices of the encoding forest, divided by n. Since the number of floor vertices is $\mathcal{O}(\sqrt{n})$, the counting measure on all vertices of Sl_n (except on the shuttle) divided by n is at vanishing Prokhorov distance from the counting measure on nonfloor vertices of the forest, divided by n. Let ω_n be the image of the Lebesgue measure on $[0, A \wedge ((2a_n + l_n)/2n)]$ by the mapping $t \mapsto (v_{i_{2nt}^n}, \mathbf{p}_{X,W}(t))$. Then ω_n is carried by the correspondence $\mathcal{R}_n^\varepsilon$ for every n large enough, and its image measures on Sl_n and $\mathsf{Sl}_{A,L,\Delta}$ by the coordinate projections are at vanishing Prokhorov distances from μ_{Sl_n}/n and $\mu^{(0,L)}$, respectively.

For the base measure, we let ω_n' be the image of the Lebesgue measure on $[0, L \wedge l_n/\sqrt{2n}]$ by the mapping $t \mapsto (v_{\tau_{\lfloor \sqrt{2n}t \rfloor}}, \mathbf{p}_{X,W}(T_t))$. Then ω_n' is carried by $\mathcal{R}_n^\varepsilon$, by the above discussion on the mark corresponding to the base. Moreover, the coordinate projections of ω_n' are at vanishing Prokhorov distance, respectively, from the counting measure on $\{\rho^0, \dots, \rho^{l_n}\}$ divided by $\sqrt{2n}$, and $\nu^{(0,L)}$. We now observe that, in turn, the counting measure on $\{\rho^0, \dots, \rho^{l_n}\}$ divided by $\sqrt{2n}$, is at vanishing Prokhorov distance from the renormalized counting measure (with multiplicities) $\nu_{\beta_n}/\sqrt{8n}$ of the base. To justify this, observe from Section 2.3 and the definition of the interval CVS bijection that the sequence $\Lambda_n(w_0), \dots, \Lambda_n(w_{2l_n+\delta_n})$ of labels of the vertices $w_0, \dots, w_{2l_n+\delta_n}$ of the base, taken in contour order, forms a simple random walk starting with a -1 step, and conditioned on hitting δ_n at time $2l_n + \delta_n$. Moreover, if we write the set $\{j \in \{0, \dots, 2l_n + \delta_n - 1\} : \Lambda_n(w_{j+1}) - \Lambda_n(w_j) = -1\}$ of down steps of this walk as $\{j_0, j_1, \dots, j_{l_n-1}\}$ with $0 = j_0 < j_1 < j_2 < \dots < j_{l_n-1}$, then the i-th root ρ^i is equal to w_{j_i} for $0 \le i < l_n$. Now consider a uniform random variable U in $[0, 1)$. Then $w_{j_{\lfloor l_n U \rfloor}}$ is a uniformly chosen forest root, while $w_{\lfloor (2l_n+\delta_n)U \rfloor}$ is a vertex of the base chosen with probability proportional to its multiplicity (and excluding ρ^{l_n} in both cases). Moreover, a standard large deviation estimate entails that $\max_{0 \le k < l_n} |j_k - 2k| = \mathcal{O}(\log n)$ in probability. In turn, this easily implies that $d_{\mathrm{Sl}_n}(w_{j_{\lfloor l_n U \rfloor}}, w_{\lfloor (2l_n+\delta_n)U \rfloor}) = \mathcal{O}(\log n)$ in probability, showing that the uniform measure on the $l_n \sim L\sqrt{2n}$ elements of $\{\rho^0, \dots, \rho^{l_n-1}\}$ is at vanishing Prokhorov distance from the law of the vertex incident to a corner uniformly chosen among the $2l_n + \delta_n \sim L\sqrt{8n}$ corners incident to the base.

Conclusion. By Lemma 3.3, we finally obtain that

$$\limsup_{n\to\infty} \mathrm{d}_{\mathrm{GHP}}^{(5,2)}(\Omega_n(\mathrm{Sl}_n), \mathsf{Sl}_{A,L,\Delta}) \le \varepsilon.$$

Since $\varepsilon > 0$ was arbitrary, this concludes the proof of Theorem 2.6.