The Brownian web and net

J.M. Swart

August 11, 2024

Contents

Chapter 1

One-dimensional voter models

1.1 The voter model

Let Λ be a countable set and let $\{0,1\}^\Lambda$ be the space of all configurations $x = (x(i))_{i\in\mathbb{Z}}$ of zeros and ones on Λ. We will be interested in continuous-time Markov processes $(X_t)_{t\geq 0}$ with state space $\{0,1\}^\Lambda$. We call $X_t(i)$ the *type* of the *site* $i \in \Lambda$ at time $t \geq 0$. Let $(\Pi_i)_{i\in\Lambda}$ an i.i.d. collection of rate one Poisson subsets of $\mathbb{R}$ and let p be a probability kernel on Λ. The *voter model* on Λ with kernel p is the continuous-time Markov process $(X_t)_{t\geq 0}$ taking values in $\{0,1\}^\Lambda$ with the following informal description:

> Each site i adapts at each time $t \in \Pi_i$ the type of a random neighbour, chosen according to the probability law $p(i,\cdot\,)$.

More formally, such a process can be constructed as follows. For each $i, j \in \Lambda$, we define a *voter model map* $\texttt{vot}_{ji} : \{0,1\}^\Lambda \to \{0,1\}^\Lambda$ by the formula:

$$\texttt{vot}_{ji}(x)(k) := \begin{cases} x(j) & \text{if } k = i, \\ x(k) & \text{otherwise.} \end{cases} \tag{1.1}$$

The effect of this map is that the site i adapts the type of the site j. We set $\mathcal{G} := \big\{\texttt{vot}_{ji} : i,j \in \Lambda\big\}$ and define a measure ρ on $\mathcal{G}\times\mathbb{R}$ by

$$\rho\big(\{\texttt{vot}_{ji}\}\times[s,t]\big) := p(i,j)(t-s) \qquad (i,j\in\Lambda,\ s,t\in\mathbb{R},\ s\leq t). \tag{1.2}$$

We let ω be a Poisson point set of $\mathcal{G}\times\mathbb{R}$ with intensity measure ρ, i.e., ω is a random subset of $\mathcal{G}\times\mathbb{R}$ such that the number of elements of $\omega\cap A$ is Poisson distributed with mean $\rho(A)$ for each measurable $A\subset\mathcal{G}\times\mathbb{R}$ such that $\rho(A)<\infty$, and if $A_1,\ldots,A_n$ are disjoint, then the random variables $\omega\cap A_1,\ldots,\omega\cap A_n$ are independent. Note that elements of ω are of the form

(m,t) with $m \in \mathcal{G}$ and $t \in \mathbb{R}$. We can now apply [Swa22, Thm 4.19] to conclude that almost surely, for each $x \in \{0,1\}^\Lambda$ and $s \in \mathbb{R}$, there exists a unique function $[s,\infty) \ni t \mapsto X_t \in \{0,1\}^\Lambda$ such that $t \mapsto X_t(i)$ is piecewise constant and right continuous for each $i \in \Lambda$, and

$$X_s = x \quad \text{and} \quad X_t = \begin{cases} m(X_{t-}) & \text{if } (m,t) \in \omega \text{ for some } m \in \mathcal{G}, \\ X_{t-} & \text{otherwise,} \end{cases} \tag{1.3}$$

where $X_{t-}(i) := \lim_{r \uparrow t} X_r(i)$ $(i \in \Lambda)$ denotes the state at the site i just before time t, and $X_{t-} = (X_{t-}(i))_{i\in\Lambda}$. Since the Lebesgue measure on $\mathbb{R}$ is atomless, it almost surely never happens that two elements of ω have the same time coordinate, so this equation is well-defined. By [Swa22, Thm 4.19] it almost surely has a unique solution for each $x \in \{0,1\}^\Lambda$ and $s \in \mathbb{R}$ simultaneously, so we can define random maps $(\mathbf{X}_{s,t})_{s\le t}$ from $\{0,1\}^\Lambda$ into itself by

$$\mathbf{X}_{s,t}(x) := X_t \quad \text{where } (X_t)_{t\ge s} \text{ solves (1.3).} \tag{1.4}$$

These random maps form a *stochastic flow*, which means that

$$\mathbf{X}_{s,s} = 1 \quad \text{and} \quad \mathbf{X}_{t,u} \circ \mathbf{X}_{s,t} = \mathbf{X}_{s,u} \qquad (s \le t \le u),$$

where 1 denotes the identity map. If X_0 is a random variable with values in $\{0,1\}^\Lambda$, independent of ω, then by [Swa22, Thm 4.19], setting

$$X_t := \mathbf{X}_{0,t}(X_0) \qquad (t \ge 0) \tag{1.5}$$

defines a Markov process $(X_t)_{t\ge 0}$ with values in $\{0,1\}^\Lambda$. We will call this the *voter model* with kernel p. Let

$$P_t(x,\,\cdot\,) := \mathbb{P}\big[\mathbf{X}_{0,t}(x) \in \,\cdot\,\big] \qquad \big(x \in \{0,1\}^\Lambda,\ t \ge 0\big) \tag{1.6}$$

denote its transition kernels. We equip $\{0,1\}^\Lambda$ with the product topology and we equip the space $\mathcal{M}_1(\{0,1\}^\Lambda)$ of probability measures on $\{0,1\}^\Lambda$ with the topology of weak convergence. Then $\{0,1\}^\Lambda$ is compact by Tychonov's theorem and consequently $\mathcal{M}_1(\{0,1\}^\Lambda)$ is compact by Prohorov's theorem. We let $\mathcal{C}(\{0,1\}^\Lambda)$ denote the space of all continuous functions $f : \{0,1\}^\Lambda \to \mathbb{R}$, equipped with the supremumnorm. We associate a probability kernel K on $\{0,1\}^\Lambda$ with the linear operator $K : \mathcal{C}(\{0,1\}^\Lambda) \to \mathcal{C}(\{0,1\}^\Lambda)$ defined by

$$Kf(x) := \int K(x,\mathrm{d}x) f(y).$$

Now [Swa22, Thm 4.19] tells us that

(i) the map $(x,t) \mapsto P_t(x,\,\cdot\,)$ from $\{0,1\}^\Lambda \times [0,\infty)$ to $\mathcal{M}_1(\{0,1\}^\Lambda)$ is continous,

(ii) $P_0 = 1$ and $P_sP_t = P_{s+t}$ $(s,t \geq 0)$,

where 1 denotes the identity map from $\mathcal{C}(\{0,1\}^\Lambda)$ into itself and P_sP_t denotes the composition of P_s and P_t, viewed as linear operators. The conditions (i) and (ii) say that the transition kernels $(P_t)_{t\geq 0}$ form a *Feller semigroup*. General theory tells us that each Feller semigroup is uniquely characterised by its *generator*, which is the linear operator $G : \mathcal{D}(G) \to \mathcal{C}(\{0,1\}^\Lambda)$ defined by

$$Gf(x) := \lim_{t\to 0} t^{-1}\big(P_t f - f\big), \tag{1.7}$$

where by definition, the domain $\mathcal{D}(G)$ of G is the set of all functions $f \in \mathcal{C}(\{0,1\}^\Lambda)$ for which the limit in (1.7) exists with respect to the supremumnorm. By [Swa22, Thm 4.30], the generator of the voter model is given by

$$Gf(x) = \sum_{i,j\in\Lambda} p(i,j)\big\{f\big(\mathtt{vot}_{ji}(x)\big) - f\big(x\big)\big\} \qquad \big(x \in \{0,1\}^\Lambda\big), \tag{1.8}$$

which is defined first for functions f that depend on finitely many coordinates, and then for more general functions by taking the closure of the operator whose domain are the functions depending on finitely many coordinates. We refer to [Swa22, Section 4.4] for details.

What is important for us is that the evolution equation (1.3) makes our informal description of the voter model at the beginning of this section rigorous. More precisely, let $(\Pi_i)_{i\in\Lambda}$ as before be an i.i.d. collection of rate one Poisson subsets of $\mathbb{R}$. Conditional on $(\Pi_i)_{i\in\Lambda}$, independently for each (i,t) with $t \in \Pi_i$, we can choose a random $j \in \Lambda$ according to the probability law $p(i,\,\cdot\,)$. Then one can check that the collection of all pairs $(\mathtt{vot}_{ji}, t)$ with $t \in \Pi_i$ and j random as just described, forms a Poisson point process on $\mathcal{G} \times \mathbb{R}$ with intensity ρ as in (1.2). Thus, our formal construction of the voter model coincides completely with the informal description given before.

1.2 One-dimensional voter models

We will exclusively be interested in the case that $\Lambda = \mathbb{Z}$, the one-dimensional integer lattice. In pictures, we draw space $\mathbb{Z}$ horizontally, we draw time $\mathbb{R}$ vertically, and for each element $(\mathtt{vot}_{ji}, t)$ of the Poisson set ω we draw an arrow from the space-time point (j,t) to the space-time point (i,t). We will later need to distinguish several types of arrows that represent different sort of maps. For reasons that will become more clear later, we will represent the voter map by an arrow with a black rectangle at its tip: $\longrightarrow\!\blacksquare$

With this convention, a piece of the Poisson set ω could look like this:

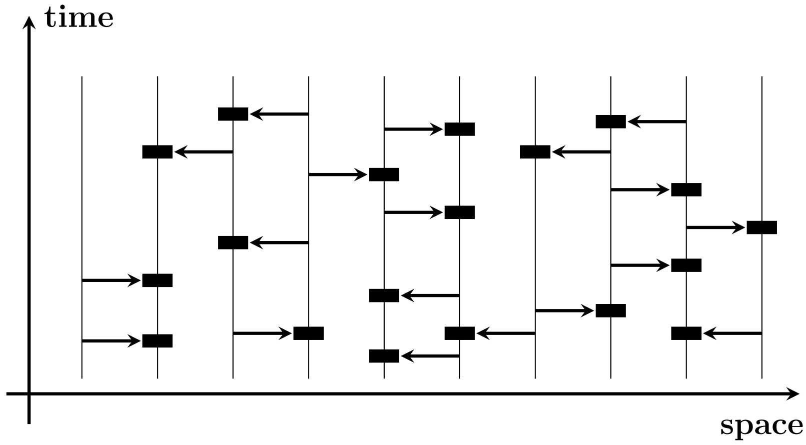

In this example, the kernel p is the *nearest-neighbour kernel*

$$p(i,j) := \begin{cases} \frac{1}{2} & \text{if } |i-j| = 1, \\ 0 & \text{otherwise,} \end{cases} \tag{1.9}$$

which has the effect that arrows only join sites at distance one from each other. Starting from an initial state $X_0 \in \{0,1\}^{\mathbb{Z}}$, we can find the solution $(X_t)_{t\geq 0}$ of the evolution equation (1.3) by applying the right maps at the right times:

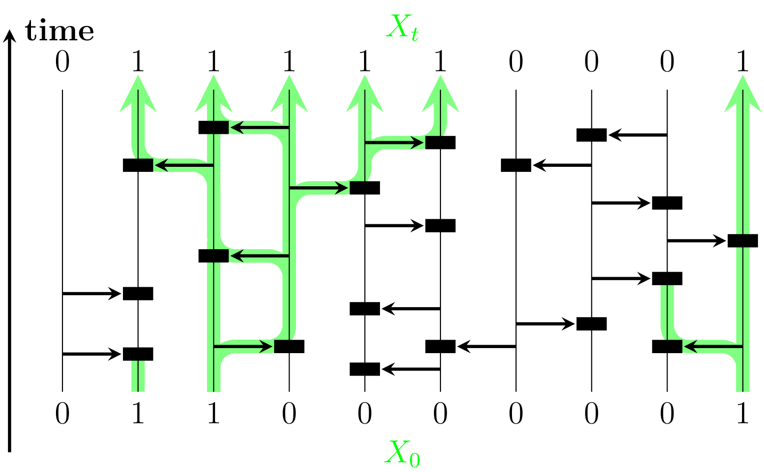

This construction is called the *graphical representation* of the voter model. More abstractly, we also refer to the Poisson set ω as the graphical representation of the voter model. It is quite easy to simulate a voter model on a computer. In the following picture, a space-time point (i, t) is white or black depending on whether $X_t(i) = 0$ or $= 1$. The initial state X_0 is chosen such that $X_0(i) = 1$ for all $i \leq 0$ and $X_0(i) = 0$ for $i > 0$.

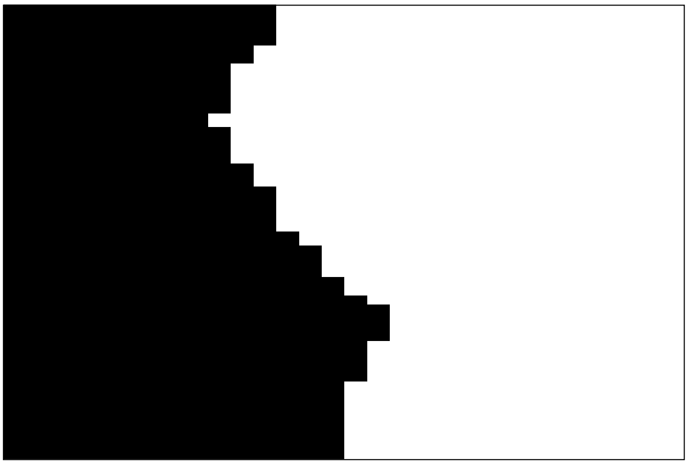

It is easy to see that the boundary between the ones and zeros evolves like a continuous-time random walk that jumps with Poisson rate $\frac{1}{2}$ one step to the left and with Poisson rate $\frac{1}{2}$ one step to the right. Therefore, by Donsker's invariance principle, if we rescale space by a factor ε, time by a factor ε^2, and send $\varepsilon \to 0$, then the boundary between the ones and zeros should converge to a standard Brownian motion, as can already be seen a bit from the following larger picture.

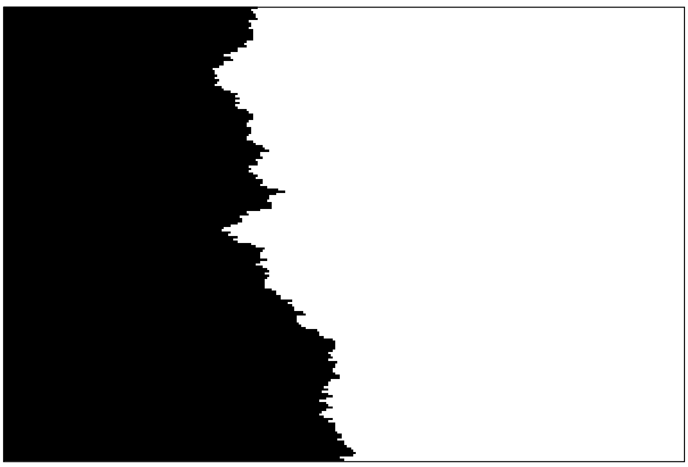

Things get a bit more complicated when we allow for more general initial states. In the following picture, the random variables $(X_0(i))_{i\in\mathbb{Z}}$ are i.i.d. and uniformly distributed on $\{0,1\}$. We have used periodic boundary conditions, i.e., we have replaced $\mathbb{Z}$ by $\mathbb{Z}/N$ for some large value of N (in this picture, $N = 300$).

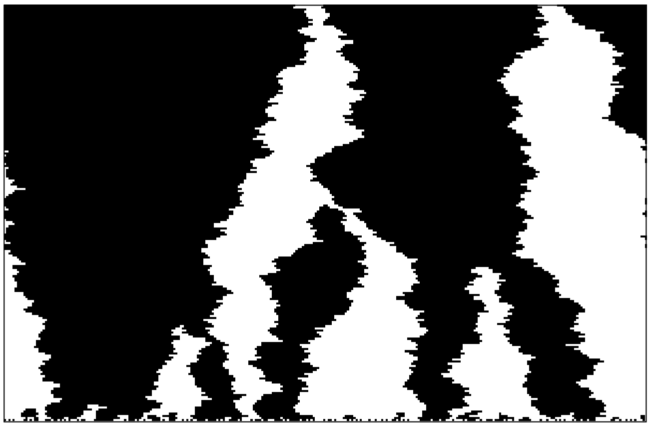

In this picture, the boundaries between zeros and ones evolve like *annihilating* random walks, which in the limit should converge to annihilating Brownian motions. A slight complication is that in the limit, these annihilating Brownian motions start from every point in space, which raises the question whether the process is well-defined. The simulations suggest this is the case, and the process "comes down from infinity", in the sense that at each positive time, the density of boundaries is already finite. We can also define voter models with more than two types. In the following picture, each site in the lattice originally has a different colour. The boundaries between these colours now evolve like *coalescing* random walks, or in the limit, as coalescing Brownian motions.

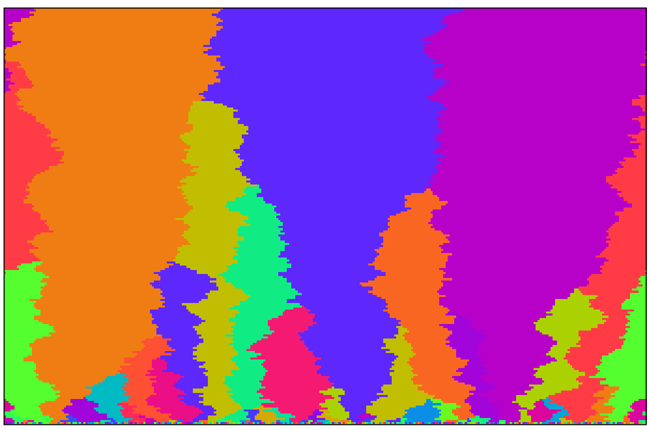

In all these simulations, we used the nearest-neighbour kernel. If instead we use the *range two kernel*

$$p(i,j) := \begin{cases} \frac{1}{4} & \text{if } 1 \leq |i-j| \leq 2, \\ 0 & \text{otherwise,} \end{cases}$$

then the picture gets more messy:

Nevertheless, the simulations suggest that on a sufficiently large scale, the limit should be the same as before, namely annihilating Brownian motions starting from each point in space. We will see that this is indeed true. In fact, it has been proved that the limit is *universal*, as long as the kernel p has mean zero and a finite $(3+\varepsilon)$-th moment.

1.3 Dual coalescing random walks

We have already seen that there is a close relation between one-dimensional voter models and systems of annihilating or coalescing random walks, because the latter describe the boundaries between intervals in which all sites have the same type. In the present section, we will see that voter models are related to coalescing random walks in yet another way, that is not restricted to one dimension. If we want to know the state of a site i at a time t, then the obvious thing to do is to look back in the graphical representation how this site got its type:

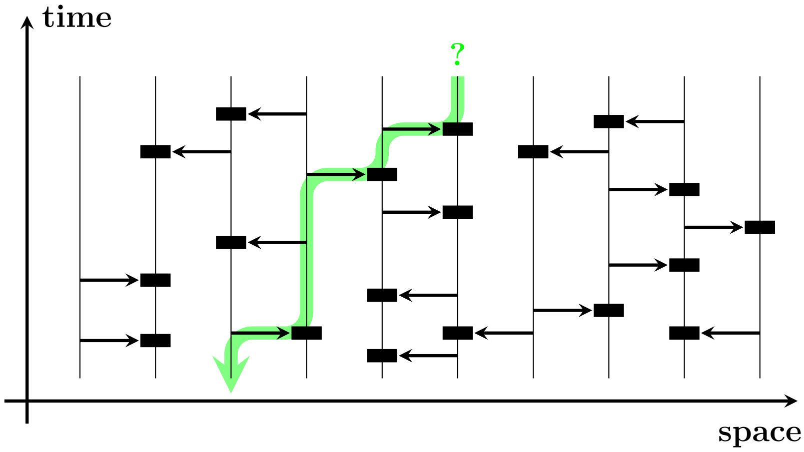

This means that starting from the space-time point (i,t), we walk till the last time when the tip of an arrow indicates that the site i copied the type of another site j. From that time on, we follow the site j back in time till the last time it changed its type and so on. Paths started from different space-time points *coalesce* as soon as they meet:

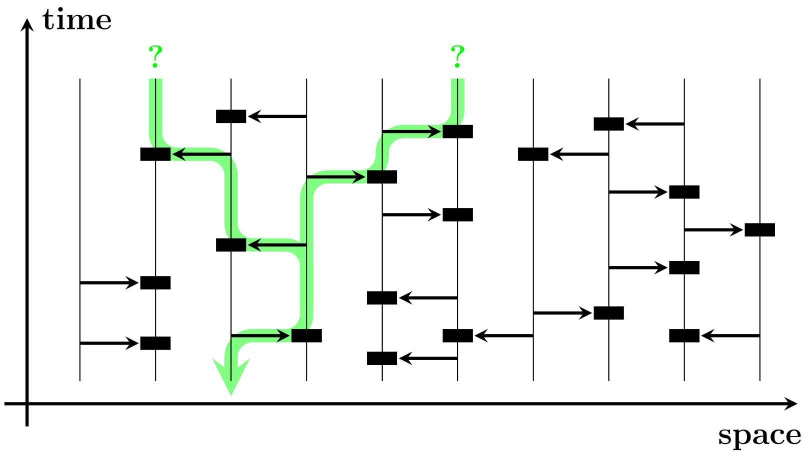

In this way, the graphical representation of the voter model can be used to construct a system of coalescing random walks, where each individual path is a continuous-time random walk that jumps with Poisson rate $p(i,j)$ from a site i to another site j.

It will be useful to view this system of coalescing random walkers as an interacting particle system in its own right. To this aim, we turn the graphical representation of the voter model upside down and reverse the direction of all arrows. We interpret an arrow from (i,t) to (j,t) of the form ■⟶ as saying that at this time, the *coalescing random walk map* should be applied, which is defined as

$$\texttt{rw}_{ij}(y)(k) := \begin{cases} 0 & \text{if } k = i \\ y(i) \vee y(j) & \text{if } k = j \\ y(k) & \text{otherwise.} \end{cases} \tag{1.10}$$

Here, we interpret sites of type 1 as *occupied* and sites of type 0 as *empty*. Then the map $\texttt{rw}_{ij}$ has the effect that if there is a particle at i, then this particle moves to j, coalescing with any particle that may already be present on that site. We obtain a graphical representation $\hat{\omega}$ for a Markov process $(Y_t)_{t\geq 0}$ with values in $\{0,1\}^{\mathbb{Z}}$ by setting

$$\hat{\omega} := \big\{ \big(\texttt{rw}_{ij}, t\big) : \big(\texttt{vot}_{ji}, -t\big) \in \omega \big\}, \tag{1.11}$$

which corresponds to reversing time and replacing voter model maps by coalescing random walk maps. In a picture, the construction looks like this:

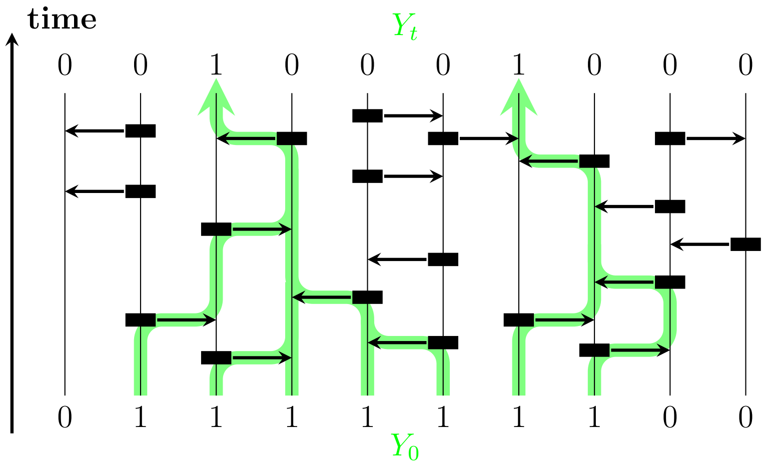

This construction is well-defined by exactly the same theorems that we cited in case of the voter model. One can check that $\hat{\omega}$ is a Poisson set and that the generator H of the Markov process $(Y_t)_{t\geq 0}$ is given by

$$Hf(y) := \sum_{i,j\in\Lambda} p(i,j)\big\{ f\big(\texttt{rw}_{ij}(y)\big) - f\big(y\big)\big\} \qquad \big(y \in \{0,1\}^{\Lambda}\big).$$

Here is a simulation of the process started form the fully occupied initial state:

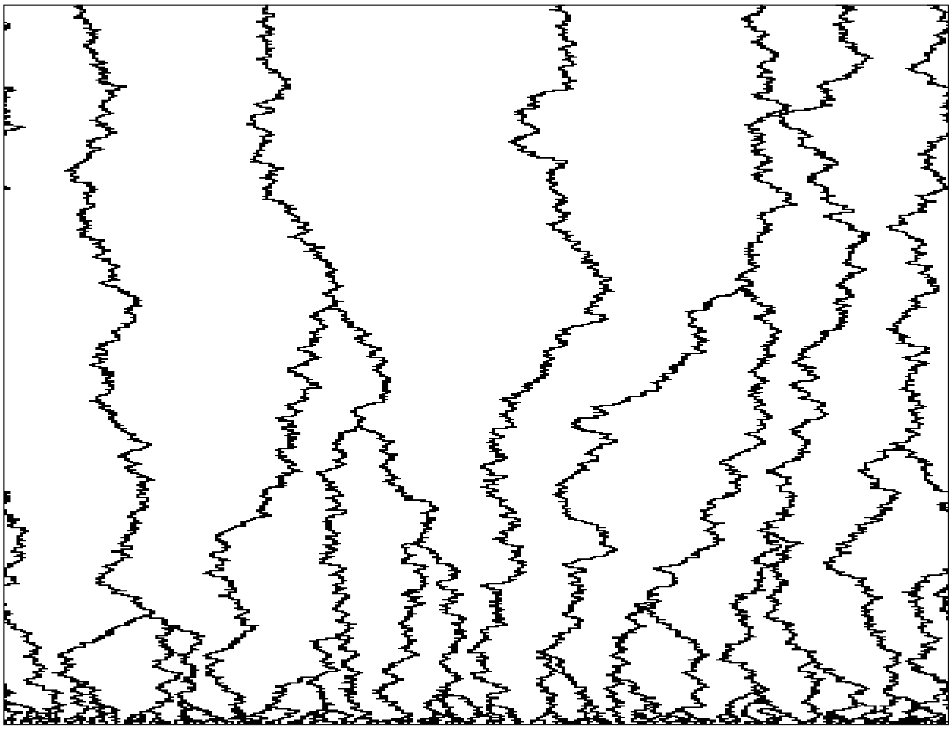

1.4 Adding branching and deaths

We can make our processes more interesting, but also more complicated, by adding two additional maps. We define a *branching map* $\mathtt{bra}_{ij}$ and *death map* $\mathtt{dth}_i$ by

$$\mathtt{bra}_{ij}(x)(k) := \begin{cases} x(i) \vee x(j) & \text{if } k = j, \\ x(k) & \text{otherwise,} \end{cases} \tag{1.12}$$

and

$$\mathtt{dth}_i(x)(k) := \begin{cases} 0 & \text{if } k = i, \\ x(k) & \text{otherwise,} \end{cases} \tag{1.13}$$

with $i, j \in \mathbb{Z}$. In graphical representations, we represent the application of the branching map $\mathtt{bra}_{ij}$ at a time t by a normal arrow $\longrightarrow$ from (i,t) to (j,t), and we represent the application of the death map $\mathtt{dth}_i$ at some time t by a "blocking symbol" ▬ . A graphical representation that contains voter model maps, branching maps, ad death maps then could look like this:

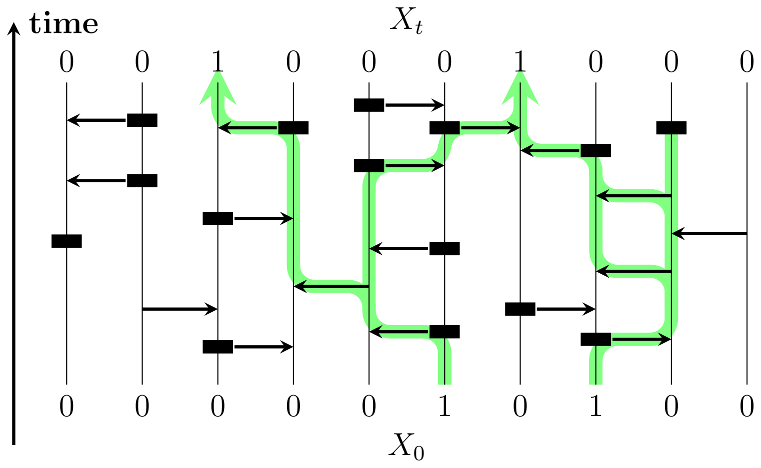

This is a good point to explain the reason why we have depicted voter model maps, coalescing random walk maps, branching maps, and death maps as we did. By definition, an *open path* in a graphical representation is a piecewise constant, right-continuous function $\gamma : [s,u] \to \mathbb{Z}$ such that:

(i) If $\gamma(t) \neq \gamma(t-)$, then there is an arrow from $(\gamma(t-),t)$ to $(\gamma(t),t)$.

(ii) If $\gamma(t) = \gamma(t-)$, then there is no blocking symbol at $(\gamma(t),t)$.

We call $(\gamma(s),s)$ the *starting point* of the path γ and $(\gamma(t),t)$ its *endpoint*. We also say that γ is an open path *from* $(\gamma(s),s)$ *to* $(\gamma(t),t)$. With these conventions, one can check that for a voter model $(X_t)_{t\geq 0}$ with additional branching and deaths

$$X_t(j) = 1 \quad \Leftrightarrow \quad \begin{array}{l}\text{there is an open path from a point}\\ (i,0) \text{ with } X_0(i) = 1 \text{ to } (j,t),\end{array} \tag{1.14}$$

and a similar statement holds for systems of coalescing random walks $(Y_t)_{t\geq 0}$ with additional branching and deaths. Given a graphical representation ω for a voter model with additional branching and deaths, we can construct a graphical representation $\hat{\omega}$ for a system of coalescing random walks with additional branching and deaths according to the recipe:

(i) If ω contains an arrow from (i,t) to (j,t), then $\hat{\omega}$ contains an arrow from $(j,-t)$ to $(i,-t)$.

(ii) If ω contains a blocking symbol at (i,t), then $\hat{\omega}$ contains a blocking symbol at $(i,-t)$.

Even more briefly, this can be described as: reverse time, reverse the direction of all arrows, and keep the blocking symbols. This generalises out earlier definition in (1.11). We can use the graphical representation ω to construct a stochastic flow $(\mathbf{X}_{s,u})_{s\leq u}$ that describes how the Markov process X evolves between two given times, and similarly we can use the graphical representation $\hat{\omega}$ to construct a stochastic flow $(\mathbf{Y}_{s,u})_{s\leq u}$ that can be used to construct the Markov process Y. We claim that these two stochastic flows are *dual* in the sense that

$$\mathbf{X}_{s,u}(x)\wedge y\neq\underline{0}\quad\Leftrightarrow\quad x\wedge\mathbf{Y}_{-u,-s}(y)\neq\underline{0}, \tag{1.15}$$

where $\underline{0}\in\{0,1\}^{\mathbb{Z}}$ denotes the configuration that is constantly zero, and $x\wedge y$ denotes the pointwise minimum of two configurations x and y. Indeed,

$$\begin{aligned}
&\mathbf{X}_{s,u}(x)\wedge y\neq\underline{0}\quad\Leftrightarrow\\
&\quad\exists i,j\in\mathbb{Z}\text{ s.t. } x(i)=1,\ y(j)=1,\text{ and there}\\
&\quad\text{exists an open path in }\omega\text{ from }(i,s)\text{ to }(j,u)\\
&\qquad\qquad\qquad\Leftrightarrow\\
&\quad\exists i,j\in\mathbb{Z}\text{ s.t. } x(i)=1,\ y(j)=1,\text{ and there}\\
&\quad\text{exists an open path in }\hat{\omega}\text{ from }(j,-u)\text{ to }(i,-s)\\
&\qquad\qquad\qquad\qquad\Leftrightarrow\quad x\wedge\mathbf{Y}_{-u,-s}(y)\neq\underline{0}.
\end{aligned}$$

There is a slight complication: formula (1.15) holds almost surely for given (deterministic) times $s\leq u$, but it does not hold almost surely for all $s\leq u$ simultaneously, due to our convention that Markov processes and our open paths are right-continuous. If we want (1.15) to hold for all $s\leq u$ simultaneously, then we have to modify our definitions of $\mathbf{X}_{s,u}$ and $\mathbf{Y}_{s,u}$ so that one is right-continuous and the other is left-continuous.

Let us first look at the case that there are no deaths. We will be interested in the interacting particle system with generator

$$\begin{aligned}
Gf(x)=&(1-\varepsilon)\sum_{i,j\in\Lambda}p(i,j)\big\{f\big(\mathtt{vot}_{ji}(x)\big)-f\big(x\big)\big\}\\
&+\varepsilon\sum_{i,j\in\Lambda}p(i,j)\big\{f\big(\mathtt{bra}_{ji}(x)\big)-f\big(x\big)\big\}\qquad\big(x\in\{0,1\}^{\Lambda}\big),
\end{aligned}\tag{1.16}$$

where p is the nearest-neighbour kernel defined in (1.9). In other words, this is the interacting particle system where voter model maps $\mathtt{vot}_{ji}$ occur with Poisson rate $(1-\varepsilon)p(i,j)$ and branching maps $\mathtt{bra}_{ji}$ occur with Poisson rate $\varepsilon p(i,j)$. If we start in an initial state such that all sites $i\leq 0$ have type 1 and all sites $i>0$ have type zero, then it is easy to see that the boundary between the ones and zeros evolves like a continuous-time random walk that jumps with Poisson rate $(1-\varepsilon)\frac{1}{2}$ one step to the left and with Poisson rate $\frac{1}{2}$ one step to the right. If we rescale all spatial distances by ε and time by ε^2, then such a random walk converges to a Brownian motion with drift $+1$. Similarly, if initially sites left of the origin have type 0 and right of the origin type 1, then the limiting Brownian motion has drift -1. The inclusion of even a little bit of branching gives the ones an advantage, so that starting with a finite interval of ones, there is a positive probability that the ones eventually take over the whole lattice:

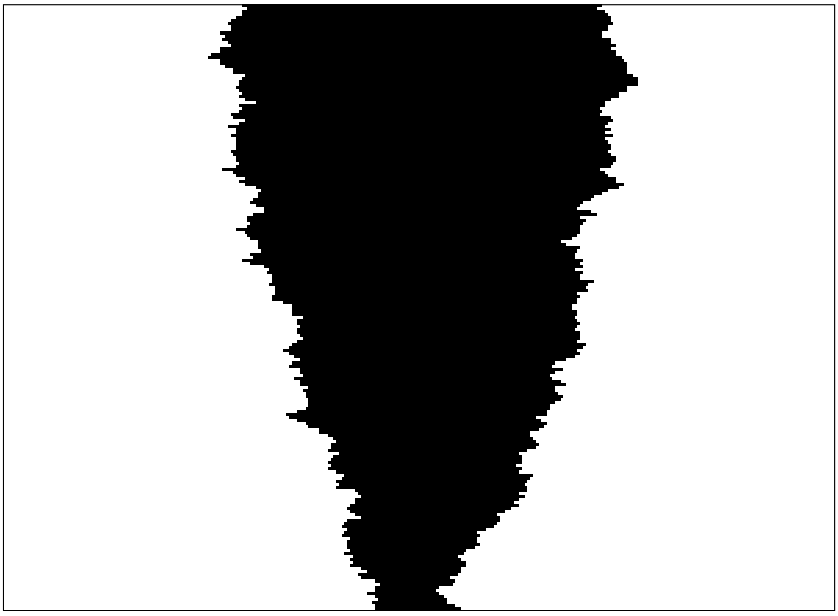

The dual process, in the sense of (1.15), is the process $(Y_t)_{t\geq 0}$ with generator

$$\begin{aligned} Gf(y) = (1-\varepsilon) \sum_{i,j\in\Lambda} p(i,j)\big\{f\big(\mathtt{rw}_{ij}(y)\big) - f\big(y\big)\big\} \\ +\varepsilon \sum_{i,j\in\Lambda} p(i,j)\big\{f\big(\mathtt{bra}_{ij}(y)\big) - f\big(y\big)\big\} \qquad \big(y \in \{0,1\}^\Lambda\big). \end{aligned} \tag{1.17}$$

In other words, this is the interacting particle system where coalescing random walk maps $\mathtt{rw}_{ij}$ occur with Poisson rate $(1-\varepsilon)p(i,j)$ and branching maps $\mathtt{bra}_{ij}$ occur with Poisson rate $\varepsilon p(i,j)$. Here is a simulation of such a process, started in the fully occupied initial state:

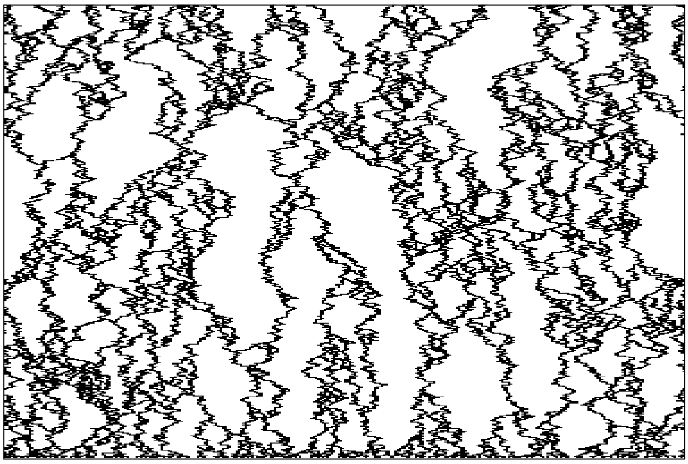

One can check that product measure with intensity ε is a reversible invariant law for this process. Moreover, if the process is started with finitely many 1's, then the position of the right-most one evolves like a continuous-time random walk that jumps with Poisson rate $(1-\varepsilon)\frac{1}{2}$ one step to the left and with Poisson rate $\frac{1}{2}$ one step to the right. These observations suggest that the process should have a diffusive scaling limit as we rescale all spatial distances by ε and time by ε^2. The simulations suggest that just like coalescing Brownian motions, this limiting process comes down from infinity. The limit process is not simply branching and coalescing Brownian motions, however. Indeed, if we follow the right-most one, then this one branches to the right with rate ε, while we rescale time by a factor ε^2. This means that in the rescaled process, the number of branchings per time unit is ε^{-1} and hence tends to infinity as $\varepsilon \to 0$.

We next add deaths as well. We will be interested in voter models with branching and deaths and generator of the form

$$\begin{aligned} Gf(x) = (1-\varepsilon) &\sum_{i,j\in\Lambda} p(i,j)\big\{f\big(\mathtt{vot}_{ji}(x)\big) - f\big(x\big)\big\} \\ +\varepsilon &\sum_{i,j\in\Lambda} p(i,j)\big\{f\big(\mathtt{bra}_{ji}(x)\big) - f\big(x\big)\big\} \\ +\delta\varepsilon^2 &\sum_{i\in\Lambda} \big\{f\big(\mathtt{dth}_i(x)\big) - f\big(x\big)\big\} \qquad \big(x\in\{0,1\}^\Lambda\big), \end{aligned} \tag{1.18}$$

where $\delta > 0$ is a fixed constant and we rescale space by ε and time by ε^2 and send $\varepsilon \to 0$. In this case, the limit looks much more nontrivial than in the case without deaths:

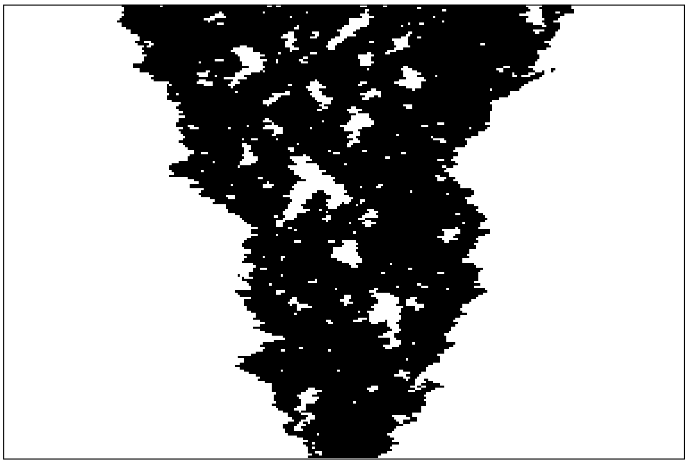

For the dual process, the picture looked already complicated in the case without deaths, and adding deaths does not complicate things much further:

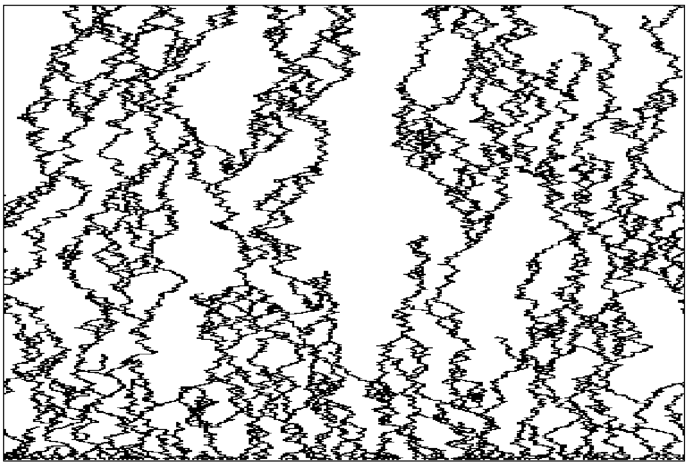

Note that since deaths occur only with rate $\delta\varepsilon^2$, and we rescale time by a factor ε^2, for the rescale process the death rate is δ. So compared to the picture with only branching and coalescing, we have just added deaths with rate δ.

1.5 Outline

In the next chapters, we will develop a mathematical theory for diffusive scaling limits of (biased) voter models and their dual systems (branching) coalescing random walks. It turns out that mathematically, the collection of all open paths in a graphical representation is a good object to work with.

Our first aim will be to construct the *Brownian web*, which can informally be described as an infinite collection of paths of coalescing Brownian motions, starting from every point in space and time. In a second step, we will then add branching, which yields the *Brownian net*. Adding also deaths yields the *Brownian net with killing*.

The Brownian web arose from the work of Arratia [Arr79, Arr81], was further developed by Tóth and Werner [TW98], and then further by Fontes, Isopi, Newman, and Ravishankar [FINR04]. The Brownian net was invented by Rongfeng Sun and myself in [SS08] and independently by Newman, Ravishankar, and Schertzer in [NRS10]. The Brownian net with killing was introduced by the same authors in [NRS15].

The Brownian web and net are believed to be universal scaling limits, that occur in a wide range of problems. For the Brownian web, there are results that show this is the limit even when the kernel is not nearest-neighbour [NRS05], but the analogue result for the Brownian net is at the moment still open (though being investigated right now). Other models that have been shown to be related to the Brownian web and net are self-repellent random walks in one dimension [TW98] and one-dimensional stochastic Potts models at low temperatures [NRS17]. We conclude this chapter by showing, as an illustration, a picture of such a Potts model.

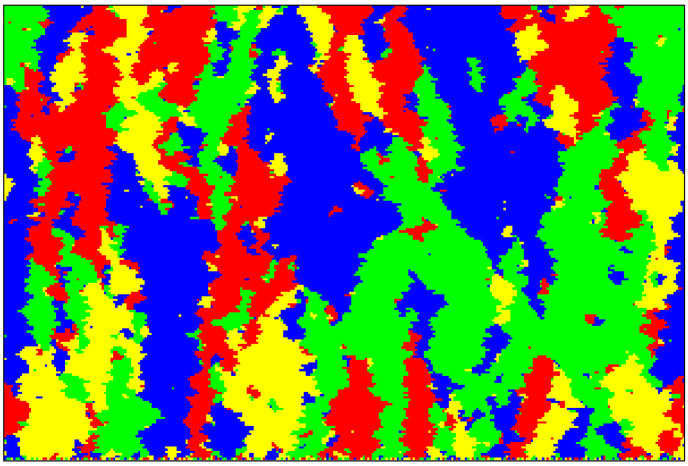

Chapter 2

Topological prerequisites

2.1 Topological spaces

We are interested in diffusive scaling limits of systems of branching and coalescing particles with small branching rate. In order to to be able to formulate the convergence, in the present chapter, we introduce the right spaces. In particular, we will need a space of paths, introduced in Section 2.7, and the space of all compact sets of paths, equipped with the Hausdorff metric, introduced in Section 2.5.

A *topological space* is a set $\mathcal{X}$ equipped with a collection $\mathcal{O}$ of subsets of $\mathcal{X}$ that are called *open* sets, such that

(i) If $(O_\gamma)_{\gamma\in\Gamma}$ is any collection of (possibly uncountably many) sets $O_\gamma \in \mathcal{O}$, then $\bigcup_{\gamma\in\Gamma} O_\gamma \in \mathcal{O}$.

(ii) If $O_1, O_2 \in \mathcal{O}$, then $O_1 \cap O_2 \in \mathcal{O}$.

(iii) $\emptyset, \mathcal{X} \in \mathcal{O}$.

Any such collection of sets is called a *topology.* It is fairly standard to also assume the *Hausdorff* property

(iv) For each $x_1, x_2 \in \mathcal{X}$, $x_1 \neq x_2$ $\exists O_1, O_2 \in \mathcal{O}$ s.t. $O_1 \cap O_2 = \emptyset$, $x_1 \in O_1$, $x_2 \in O_2$.

A set $V \subset \mathcal{X}$ is a *neighbourhood* of a point $x \in \mathcal{X}$ if $x \in O \subset V$ for some $O \in \mathcal{O}$. We let $\mathcal{V}_x$ denote the set of all neighbourhoods of x. A *fundamental system* of neighbourhoods of x is a set $\mathcal{V}'_x \subset \mathcal{V}_x$ such that

$$\forall V \in \mathcal{V}_x \ \exists V' \in \mathcal{V}'_x \text{ s.t. } V' \subset V.$$

For example, the set of all $O \in \mathcal{O}$ such that $x \in O$ is a fundamental system of neighbourhoods of x. A sequence of points $x_n \in \mathcal{X}$ converges to a limit x in a given topology $\mathcal{O}$ if for each $V \in \mathcal{V}_x$ there is an n such that $x_m \in V$ for all $m \geq n$. It suffices to check this condition for a fundamental system of neighbourhoods $\mathcal{V}'_x$. If the topology is Hausdorff, then limits are unique, i.e., $x_n \to x$ and $x_n \to x'$ implies $x = x'$.

If $(\mathcal{X}, \mathcal{O})$ is a topological space (with $\mathcal{O}$ the collection of open subsets of $\mathcal{X}$) and $\mathcal{X}' \subset \mathcal{X}$ is any subset of $\mathcal{X}$, then $\mathcal{X}'$ is also naturally equipped with a topology given by the collection of open subsets $\mathcal{O}' := \{O \cap \mathcal{X}' : O \in \mathcal{O}\}$. This topology is called the *induced* topology from $\mathcal{X}$. If $x_n, x \in \mathcal{X}'$, then $x_n \to x$ in the induced topology on $\mathcal{X}'$ if and only if $x_n \to x$ in $\mathcal{X}$.

A *basis* of a topology is a subset $\mathcal{O}' \subset \mathcal{O}$ such that each element of $\mathcal{O}$ can be written as the union of (possibly uncountably many) elements of $\mathcal{O}'$. Equivalently, this says that

$$\mathcal{O} = \{O \subset \mathcal{X} : \forall x \in O\ \exists O' \in \mathcal{O}' \text{ s.t. } x \in O' \subset O\}.$$

If $\mathcal{O}'$ is a basis for $\mathcal{O}$, then $\mathcal{V}'_x := \{O \in \mathcal{O}' : x \in O\}$ is a fundamental system of neighbourhoods of x. A topology is *first countable* if every $x \in \mathcal{X}$ has a countable fundamental system of neighbourhoods. A topology is *second countable* if there exists a countable basis of the topology.

A set $C \subset \mathcal{X}$ is called *closed* if its complement is open. Because of property (i) in the definition of a topology, for each $A \subset \mathcal{X}$, the union of all open sets contained in A is itself an open set. We call this the *interior* of A, denoted as $\text{int}(A) := \bigcup\{O : O \subset A,\ O \text{ open}\}$. Then clearly $\text{int}(A)$ is the largest open set contained in A. Similarly, by taking complements, for each set $A \subset \mathcal{X}$ there exists a smallest closed set containing A. We call this the *closure* of A, denoted as $\overline{A} := \bigcap\{C : C \supset A,\ C \text{ closed}\}$. If the topology is first countable, then

$$\overline{A} = \{x \in \mathcal{X} : \exists x_n \in \mathcal{X} \text{ s.t. } x_n \to x\}, \tag{2.1}$$

i.e., $\overline{A}$ is the set of all limits of sequences in A. A similar statement holds for general topological spaces if we replace sequences by the more general concept of a *net*, that we will not discuss here. Since a set is closed if and only if it coincides with its closure, it follows from (2.1) that in a first countable topological space, knowing all convergent sequences and their limits uniquely determines the closed sets and their complements, the open sets, and hence the whole topology.

A topological space is called *separable* if there exists a countable set $D \subset \mathcal{X}$ such that D is dense in $\mathcal{X}$, where we say that a set $D \subset \mathcal{X}$ is *dense* if its closure is $\mathcal{X}$, or equivalently, if every nonempty open subset of $\mathcal{X}$ has a nonempty intersection with D.

A *metric* on a set $\mathcal{X}$ is a function $d : \mathcal{X} \times \mathcal{X} \to [0, \infty)$ such that for all $x, y, z \in \mathcal{X}$,

(i) $d(x, y) = d(y, x)$,

(ii) $d(x, z) \leq d(x, y) + d(y, z)$,

(iii) $d(x, y) = 0$ implies $x = y$.

A *metric space* is a space with a metric defined on it. If d is a metric on $\mathcal{X}$, and $B_\varepsilon(x) := \{y \in \mathcal{X} : d(x, y) < \varepsilon\}$ denotes the open ball around x of radius ε, then

$$\mathcal{O} := \big\{O \subset \mathcal{X} : \forall x \in O \ \exists \varepsilon > 0 \text{ s.t. } B_\varepsilon(x) \subset O\big\}$$

defines a Hausdorff topology on $\mathcal{X}$ such that convergence $x_n \to x$ in this topology is equivalent to $d(x_n, x) \to 0$. Note that the open balls form a basis for this topology. Since open balls of radius $1/n$ around a point x form a fundamental system of neighbourhoods, metric spaces are first countable. We say that the metric d *generates* the topology $\mathcal{O}$. If for a given topology $\mathcal{O}$ there exists a metric d that generates $\mathcal{O}$, then we say that the topological space $(\mathcal{X}, \mathcal{O})$ is *metrisable*. Such a metric, if it exist, can always be chosen such that it is bounded. For example, if d is any metric on $\mathcal{X}$, then $d'(x, y) := d(x, y) \wedge 1$ is a bounded metric that generates the same topology. A metrisable space is always first countable. It is second countable if and only if it is separable.

A sequence x_n in a metric space $(\mathcal{X}, d)$ is a *Cauchy sequence* if for all $\varepsilon > 0$ there is an n such that $d(x_k, x_l) \leq \varepsilon$ for all $k, l \geq n$. A metric space is *complete* if every Cauchy sequence converges. Every metric space $(\mathcal{X}, d)$ has a *completion*, i.e., there exists a complete metric space $(\overline{\mathcal{X}}, \overline{d})$ such that $\mathcal{X} \subset \overline{\mathcal{X}}$ is dense and the metric on $\mathcal{X}$ is the *induced metric* from $\overline{\mathcal{X}}$, i.e., $d(x, y) = \overline{d}(x, y)$ for all $x, y \in \mathcal{X}$. Such a completion is unique up to isometries.

A *Polish space* is a separable topological space $(\mathcal{X}, \mathcal{O})$ such that there exists a metric d on $\mathcal{X}$ with the property that $(\mathcal{X}, d)$ is complete and d generates $\mathcal{O}$. *Warning:* there may be many different metrics on $\mathcal{X}$ that generate the same topology. It may even happen that $\mathcal{X}$ is not complete in some of these metrics, and complete in others (in which case $\mathcal{X}$ is still Polish).[1] Example: $\mathbb{R}$ is separable and complete in the usual metric $d(x, y) =$

[1]The use of the term "Polish space" has a long history and there is some variation in its definition. While our use of the term is in line with most of the modern literature,

$|x - y|$, and therefore $\mathbb{R}$ is a Polish space. But $d'(x,y) := |\arctan(x) - \arctan(y)|$ is another metric that generates the same topology, while $(\mathbb{R}, d')$ is not complete. Indeed, the completion of $\mathbb{R}$ w.r.t. the metric d' is $[-\infty, \infty]$.

2.2 Compactness

A subset K of a general topological space $\mathcal{X}$ (with collection of open sets $\mathcal{O}$) is called *compact* if every open cover has a finite subcover, i.e., if for any collection $(O_\gamma)_{\gamma\in\Gamma}$ of open subsets of $\mathcal{X}$ such that $\bigcup_{\gamma\in\Gamma} O_\gamma \supset K$, there exists a finite $\Delta \subset \Gamma$ such that $\bigcup_{\gamma\in\Delta} O_\gamma \supset K$. Using this definition, it is easy to see that the image of a compact set under a continuous function is again compact. Compact subsets of Hausdorff topological spaces are closed. A subset K of a metric space $\mathcal{X}$ is compact if and only if it is complete and *totally bounded*, which means that for every $\varepsilon > 0$ there exists a finite collection $\{B_\varepsilon(x_1), \ldots, B_\varepsilon(x_n)\}$ of open balls such that

$$B_\varepsilon(x_1) \cup \cdots \cup B_\varepsilon(x_n) \supset K.$$

From this, it is not hard to see that compact metrisable spaces are always separable. If $(x_n)_{n\in\mathbb{N}}$ is a sequence and $m : \mathbb{N} \to \mathbb{N}$ is a function such that $m(n) \to \infty$ as $n \to \infty$, then setting $x'_n := x_{m(n)}$ $(n \in \mathbb{N})$ defines a new sequence. Such a sequence is called a *subsequence* of the original sequence. A *cluster point* of a sequence is a limit of a subsequence.

Theorem 2.1 (Bolzano-Weierstrass) *Let $\mathcal{X}$ be a metrisable space and let $K \subset \mathcal{X}$. Then K is compact if and only if every sequence in K has a subsequence that converges to a limit in K.*

The Bolzano-Weierstrass theorem also holds for second countable spaces. (Note that metrisable spaces need in general not be second countable, and conversely, not every second countable space is metrisable.) There is also a version of the Bolzano-Weierstrass theorem that holds in general topological spaces but in this case one has to replace sequences by the more general nets. A set A is *precompact* if its closure is compact. In metrisable spaces, this is equivalent to the statement that each sequence of points $x_n \in A$ has a convergent subsequence. Note that in this case we do not require that the limit is an element of A. The following simple lemma is often useful.

Lemma 2.2 (Convergence and compactness) *Let $\mathcal{X}$ be a metrisable space and let $x, x_n \in \mathcal{X}$. Then $x_n \to x$ if and only if the following two conditions are satisfied.*

---

some authors use "topologically Polish" for what we call "Polish" and reserve the latter term for the more restricted setting of a complete and separable metric space.

(i) *The set* $\{x_n : n \in \mathbb{N}\}$ *is precompact.*

(ii) *For every subsequence* $x_{n(m)}$ *such that* $x_{n(m)} \underset{m\to\infty}{\longrightarrow} x'$ *for some* $x' \in \mathcal{X}$*, one has* $x' = x$*.*

If $(\mathcal{X}, \mathcal{O})$ is a topological space, then a *compactification* of $\mathcal{X}$ is a compact topological space $\overline{\mathcal{X}}$ such that $\mathcal{X}$ is a dense subset of $\overline{\mathcal{X}}$ and the topology on $\mathcal{X}$ is the induced topology from $\overline{\mathcal{X}}$. If $\overline{\mathcal{X}}$ is metrisable, then we say that $\overline{\mathcal{X}}$ is a *metrisable compactification* of $\mathcal{X}$. It turns out that each separable metrisable space $\mathcal{X}$ has a metrisable compactification [Cho69, Theorem 6.3].

A topological space $\mathcal{X}$ is called *locally compact* if for every $x \in \mathcal{X}$ there exists a compact neighbourhood of x. We cite the following proposition from [Eng89, Thms 3.3.8 and 3.3.9].

Proposition 2.3 (Compactification of locally compact spaces) *Let* $\mathcal{X}$ *be a metrisable topological space. Then the following statements are equivalent.*

(i) $\mathcal{X}$ *is locally compact and separable.*

(ii) *There exists a metrisable compactification* $\overline{\mathcal{X}}$ *of* $\mathcal{X}$ *such that* $\mathcal{X}$ *is an open subset of* $\overline{\mathcal{X}}$*.*

(iii) *For each metrisable compactification* $\overline{\mathcal{X}}$ *of* $\mathcal{X}$*,* $\mathcal{X}$ *is an open subset of* $\overline{\mathcal{X}}$*.*

We note that if $\mathcal{X}$ satisfies the equivalent conditions of Proposition 2.3, then it is possible to find a metrisable compactification $\overline{\mathcal{X}}$ of $\mathcal{X}$ such that $\overline{\mathcal{X}}\backslash\mathcal{X}$ consists of just one point, usually denoted by ∞. In this case, $\overline{\mathcal{X}} = \mathcal{X} \cup \{\infty\}$ is called the *one-point compactification* of $\mathcal{X}$. The open sets of $\mathcal{X} \cup \{\infty\}$ are all open sets of $\mathcal{X}$ plus all sets of the form $\{\infty\} \cup O$ where $\mathcal{X}\backslash O$ is a compact subset of $\mathcal{X}$.

A subset $A \subset \mathcal{X}$ of a topological space $\mathcal{X}$ is called a G_δ*-set* if A is a countable intersection of open sets (i.e., there exist $O_i \in \mathcal{O}$ such that $A = \bigcap_{i=1}^\infty O_i$. If $\mathcal{X}$ is metrisable, then every closed set $A \subset \mathcal{X}$ is a G_δ-set, since it is the intersection of the open sets $\{x \in \mathcal{X} : d(x, A) < 1/n\}$. The following result can be found in [Bou58, §6 No. 1, Theorem. 1]. See also [Oxt80, Thms 12.1 and 12.3].

Proposition 2.4 (Compactification of Polish spaces) *Let* $\mathcal{X}$ *be a metrisable topological space. Then the following statements are equivalent.*

(i) $\mathcal{X}$ *is Polish.*

(ii) *There exists a metrisable compactification* $\overline{\mathcal{X}}$ *of* $\mathcal{X}$ *such that* $\mathcal{X}$ *is a* G_δ*-subset of* $\overline{\mathcal{X}}$*.*

(iii) *For each metrisable compactification* $\overline{\mathcal{X}}$ *of* $\mathcal{X}$, $\mathcal{X}$ *is a* G_δ*-subset of* $\overline{\mathcal{X}}$.

Moreover, a subset $\mathcal{Y} \subset \mathcal{X}$ *of a Polish space* $\mathcal{X}$ *is Polish in the induced topology if and only if* $\mathcal{Y}$ *is a* G_δ*-subset of* $\mathcal{X}$.

We note that if $\overline{\mathcal{X}}$ is a compactification of a Polish space $\mathcal{X}$, equipped with a concrete metric, then $\overline{\mathcal{X}}$ is also the completion of $\mathcal{X}$ in this metric. Thus, unless $\mathcal{X}$ is itself compact, it will never be complete in such a metric (even though, by the definition of a Polish space, there exists metrics generating the same topology with respect to which $\mathcal{X}$ is complete).

2.3 Weak convergence

Let $\mathcal{X}$ be a metrisable space. We let $\mathcal{B}(\mathcal{X})$ denote *Borel-*σ*-field* on $\mathcal{X}$, i.e., the σ-field generated by the open sets. We let $\mathcal{C}(\mathcal{X})$ denote the space of all continuous functions $f : \mathcal{X} \to \mathbb{R}$. We let $B_{\rm b}(\mathcal{X})$ denote the space of all bounded Borel-measurable real functions on $\mathcal{X}$ and we let $\mathcal{C}_{\rm b}(\mathcal{X}) := \mathcal{C}(\mathcal{X}) \cap B_{\rm b}(\mathcal{X})$ denote the space of all bounded continuous real functions on $\mathcal{X}$. We equip $\mathcal{C}_{\rm b}(\mathcal{X})$ with the *supremumnorm*

$$\|f\|_\infty := \sup_{x \in \mathcal{X}} |f(x)|.$$

With this norm, $\mathcal{C}_{\rm b}(\mathcal{X})$ is a Banach space [Dud02, Theorem 2.4.9]. We let $\mathcal{M}(\mathcal{X})$ denote the space of all finite measures on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$ and write $\mathcal{M}_1(\mathcal{X})$ for the subspace of all probability measures. We cite the following well-known fact from [EK86, Theorems 3.1.7 and 3.3.1].

Proposition 2.5 (Weak convergence) *Let* $\mathcal{X}$ *be a separable metrisable space. Then it is possible to equip* $\mathcal{M}_1(\mathcal{X})$ *with a metric* $d_{\rm P}$ *such that*

(i) $(\mathcal{M}_1(\mathcal{X}), d_{\rm P})$ *is a separable metric space,*

(ii) $d_{\rm P}(\mu_n, \mu) \to 0$ *if and only if* $\int f \mathrm{d}\mu_n \to \int f \mathrm{d}\mu$ *for all* $f \in \mathcal{C}_{\rm b}(\mathcal{X})$.

If $\mathcal{X}$ *is a Polish space, then* $d_{\rm P}$ *can be chosen such that* $(\mathcal{M}_1(\mathcal{X}), d_{\rm P})$ *is moreover complete.*

In many applications, we are not interested in the precise choice of $d_{\rm P}$ (there are several canonical ways to define such a metric). Since a metrisable topology is uniquely characterized by its convergent sequences, property (ii) uniquely characterizes the topology generated by $d_{\rm P}$ in terms of the topology on $\mathcal{X}$. We call this topology the *topology of weak convergence* and denote convergence in this topology as

$$\mu_n \Rightarrow \mu.$$

Proposition 2.5 shows in particular that if $\mathcal{X}$ is a Polish space, then so is $\mathcal{M}_1(\mathcal{X})$, equipped with the topology of weak convergence.

One possible choice for a metric $d_{\rm P}$ as in Proposition 2.5 is the Prohorov metric. For each subset $A \subset \mathcal{X}$ and $\varepsilon > 0$, we set

$$A^\varepsilon := \big\{x \in \mathcal{X} : d(x, A) < \varepsilon\big\} \quad \text{with} \quad d(x, A) := \inf_{y \in A} d(x, y).$$

If $(\mathcal{X}, d)$ is a metric space, then the *Prohorov metric* is the metric $d_{\rm P}$ on $\mathcal{M}_1(\mathcal{X})$ defined as

$$d_{\rm P}(\mu, \nu) := \inf\{\varepsilon > 0 : \mu(A) \leq \nu(A^\varepsilon) + \varepsilon \ \forall A \in \mathcal{B}(\mathcal{X})\big\}.$$

It follows from [EK86, Lemma 3.1.1] that $d_{\rm P}$ is a metric. It is possible to give an alternative characterisation of $d_{\rm P}$ in terms of coupling. Let $C(\mu, \nu)$ denote the space of all probability measures η on $X \times X$ whose first and second marginals are given by μ and ν, respectively. We cite the following lemma from [EK86, Thm 3.1.2].

Lemma 2.6 (Prohorov metric and coupling) *Let $(\mathcal{X}, d)$ be a separable metric space and let $\mu, \nu \in \mathcal{M}_1(\mathcal{X})$. Then*

$$\begin{aligned} &d_{\rm P}(\mu, \nu) = \\ &\quad \inf\big\{\varepsilon > 0 : \exists \eta \in C(\mu, \nu) \ s.t. \ \eta(\{(x, y) \in \mathcal{X}^2 : d(x, y) \geq \varepsilon\}) \leq \varepsilon\big\}. \end{aligned} \tag{2.2}$$

In words, (2.2) says that $d_{\rm P}(\mu, \nu)$ is the infimum of all $\varepsilon > 0$ for which it is possible to couple random variables X, Y with laws μ, ν such that $\mathbb{P}[d(X, Y) \geq \varepsilon] \leq \varepsilon$. We cite the following lemmas from [EK86, Thms 3.1.7 and 3.3.1].

Lemma 2.7 (Properties of Prohorov metric) *Let $(\mathcal{X}, d)$ be a separable metric space and let $d_{\rm P}$ be the Prohorov metric. Then $(\mathcal{M}_1(\mathcal{X}), d_{\rm P})$ is a separable metric space. If $(\mathcal{X}, d)$ is complete, then so is $(\mathcal{M}_1(\mathcal{X}), d_{\rm P})$.*

Lemma 2.8 (Prohorov metric and weak convergence) *Let $(\mathcal{X}, d)$ be a separable metric space and let $d_{\rm P}$ be the Prohorov metric. Then $\mu_n, \mu \in \mathcal{M}_1(\mathcal{X})$ satisfy $d_{\rm P}(\mu_n, \mu) \to 0$ if and only if $\int f \mathrm{d}\mu_n \to \int f \mathrm{d}\mu$ for all $f \in \mathcal{C}_{\rm b}(\mathcal{X})$.*

In particular, Lemmas 2.7 and 2.8 imply Proposition 2.5. The following well-known alternative characterisation of weak convergence [EK86, Theorem 3.3.1] is sometimes useful.

Lemma 2.9 (Characterization with open and closed sets) *Let μ_n and μ be probability measures on a metrisable space $\mathcal{X}$. Then the following statements are equivalent.*

(i) $\mu_n \Rightarrow \mu$.

(ii) $\limsup_{n\to\infty} \mu_n(C) \leq \mu(C)$ *for all closed* $C \subset \mathcal{X}$.

(iii) $\liminf_{n\to\infty} \mu_n(O) \geq \mu(O)$ *for all open* $O \subset \mathcal{X}$.

Exercise 2.10 (Measures concentrated on a subset) *Let $\mathcal{X}$ be a Polish space and let $\mathcal{X}' \subset \mathcal{X}$ be a G_δ-set, equipped with the induced topology. We naturally identify $\mathcal{M}_1(\mathcal{X}')$ with the subset of $\mathcal{M}_1(\mathcal{X})$ consisting of all $\mu \in \mathcal{M}_1(\mathcal{X})$ such that $\mu(\mathcal{X}') = 1$. Show that the topology on $\mathcal{M}_1(\mathcal{X}')$ coincides with the induced topology from its embedding in $\mathcal{M}_1(\mathcal{X})$. (Hint: Lemma 2.9.) Use this to conclude that $\mathcal{M}_1(\mathcal{X}')$ is a G_δ-subset of $\mathcal{M}_1(\mathcal{X})$. (Hint: Proposition 2.4).*

A very useful characterization of weak convergence in terms of coupling is given by the next theorem [EK86, Cor 3.1.6 and Thm 3.1.8].

Theorem 2.11 (Skorohod representation) *Let μ_n and μ be probability measures on a Polish space $\mathcal{X}$. Then $\mu_n \Rightarrow \mu$ if and only if it is possible to couple random variables X_n, X with laws μ_n, μ, respectively, in such a way that $X_n \to X$ a.s.*

The next result is known as Prohorov's theorem (see, e.g., [EK86, Theorem 3.2.2] or [Bil99, Theorems 5.1 and 5.2]).

Theorem 2.12 (Prohorov) *Let $\mathcal{X}$ be a Polish space. Let $\mathcal{M}_1(\mathcal{X})$ be equipped with the topology of weak convergence. Then a subset $\mathcal{C} \subset \mathcal{M}_1(\mathcal{X})$ is precompact if and only if $\mathcal{C}$ is* tight, *i.e.,*

$$\forall \varepsilon > 0 \ \exists K \subset \mathcal{X} \text{ compact, s.t. } \sup_{\mu \in \mathcal{C}} \mu(\mathcal{X} \backslash K) \leq \varepsilon.$$

2.4 Locally uniform convergence

Let $(\mathcal{X}, d)$ be a metric space and let $I \subset \mathbb{R}$ be a closed interval. We let $\mathcal{C}_I(\mathcal{X})$ denote the space of all continuous functions $f : I \to \mathcal{X}$.

Lemma 2.13 (Locally uniform convergence) *For $f_n, f \in \mathcal{C}_I(\mathcal{X})$, the following conditions are equivalent:*

(i) $\sup_{t\in C} d\big(f_n(t), f(t)\big) \underset{n\to\infty}{\longrightarrow} 0$ *for all compact* $C \subset I$,

(ii) $f_n(t_n) \underset{n\to\infty}{\longrightarrow} f(t)$ *for all* $t_n, t \in I$ *such that* $t_n \underset{n\to\infty}{\longrightarrow} t$.

Proof Assume (i) and let $t_n, t \in I$ satisfy $t_n \underset{n\to\infty}{\longrightarrow} t$. By Lemma 2.2 (i), there exists a compact set $C \subset I$ such that $t_n \in C$ for all n (and hence also $t \in C$). Then for each $\varepsilon > 0$, there exists an $N < \infty$ such that $d(f_n(t), f(t)) \leq \varepsilon$ for all $n \geq N$. Now

$$d\big(f_n(t_n), f(t)\big) \leq d\big(f_n(t_n), f(t_n)\big) + d\big(f(t_n), f(t)\big) \leq \varepsilon + d\big(f(t_n), f(t)\big)$$

for all $n \geq N$, and hence

$$\limsup_{n\to\infty} d\big(f_n(t_n), f(t)\big) \leq \varepsilon$$

by the continuity of f. Since $\varepsilon > 0$ is arbitrary, this shows that (i) implies (ii). On the other hand, if (i) fails for some compact $C \subset I$, then for s suitable subsequence we can choose $t_n \in C$ and $\varepsilon > 0$ such that

$$d\big(f_n(t_n), f(t_n)\big) \geq \varepsilon \quad \forall n.$$

Since C is compact, by going to a further subsequence, we can without loss of generality assume that $t_n \to t$ for some $t \in C$. Since

$$d\big(f_n(t_n), f(t)\big) \geq d\big(f_n(t_n), f(t_n)\big) - d\big(f(t_n), f(t)\big) \geq \varepsilon + d\big(f(t_n), f(t)\big),$$

using the continuity of f, we see that for our chosen subsequence

$$\liminf_{n\to\infty} d\big(f_n(t_n), f(t)\big) \geq \varepsilon,$$

which contradicts (ii). ∎

There exists a metrisable topology on $\mathcal{C}_I(\mathcal{X})$ such that a $f_n \in \mathcal{C}_I(\mathcal{X})$ converges to a limit f if and only if the equivalent conditions of Lemma 2.13 are satisfied. Note that by (2.1) and the remarks below it, these conditions uniquely determine the topology. Note also that by condition (ii) of Lemma 2.13, the topology on $\mathcal{C}_I(\mathcal{X})$ depends only on the topology on $\mathcal{X}$ and not on the precise choice of the metric on $\mathcal{X}$. A possible choice of a metric on $\mathcal{C}_I(\mathcal{X})$ is

$$\rho(g, f) := \sum_{n=1}^{\infty} 2^{-n} \sup_{t\in[-n,n]\cap I} d\big(g(t), f(t)\big),$$

where d is a bounded metric that generates the topology on $\mathcal{X}$. Such a metric can always be found: if d is any metric generating the topology on $\mathcal{X}$, then $d'(x,y) := d(x,y) \wedge 1$ is a bounded metric that generates the same topology. Usually, we do not care about the precise choice of the metric on $\mathcal{C}_I(\mathcal{X})$; apart from ρ, there are many other possible choices. We call this the topology on $\mathcal{C}_I(\mathcal{X})$ the *topology of locally uniform convergence.*

2.5 The Hausdorff metric

Let $(\mathcal{X}, d)$ be a metric space, let $\mathcal{K}(\mathcal{X})$ be the space of all compact subsets of $\mathcal{X}$ and set $\mathcal{K}_+(\mathcal{X}) := \{K \in \mathcal{K}(\mathcal{X}) : K \neq \emptyset\}$. Then the *Hausdorff metric* $d_{\rm H}$ on $\mathcal{K}_+(\mathcal{X})$ is defined as

$$\begin{aligned} d_{\rm H}(K_1, K_2) &:= \sup_{x_1 \in K_1} d(x_1, K_2) \vee \sup_{x_2 \in K_2} d(x_2, K_1) \\ &= \inf\big\{\varepsilon > 0 : K_1 \subset K_2^\varepsilon \text{ and } K_2 \subset K_1^\varepsilon\big\}, \end{aligned} \tag{2.3}$$

where as before $d(x, A) := \inf_{y \in A} d(x, y)$ denotes the distance between a point $x \in \mathcal{X}$ and a set $A \subset \mathcal{X}$ and $A^\varepsilon := \big\{x \in \mathcal{X} : d(x, A) < \varepsilon\big\}$. The corresponding topology is naturally called the *Hausdorff topology.* Note the subtle difference between "the Hausdorff topology" (the topology generated by the Hausdorff metric) and "a Hausdorff topology" (any topology satisfying condition (iv) of Section 2.1). We extend this topology to $\mathcal{K}(\mathcal{X})$ by adding $\emptyset$ as an isolated point.

A *correspondence* between two sets A and B is a set $R \subset A \times B$ such that

$$\forall a \in A \ \exists b \in B \text{ s.t. } (a, b) \in R \quad \text{and} \quad \forall b \in B \ \exists a \in A \text{ s.t. } (a, b) \in R.$$

In words, this says that for each element of A, there is a "corresponding" element of B and conversely for each element of B, there is a corresponding element of A. We let $\mathrm{Corr}(A, B)$ denote the set of all correspondences between A and B. The following exercise relates the Hausdorff metric to correspondences.

Exercise 2.14 *Let $(\mathcal{X}, d)$ be a metric space. Show that*

$$d_{\rm H}(K_1, K_2) = \inf_{R \in \mathrm{Corr}(K_1, K_2)} \sup_{(x_1, x_2) \in R} d(x_1, x_2) \qquad \big(K_1, K_2 \in \mathcal{K}_+(\mathcal{X})\big). \tag{2.4}$$

A good source for the Hausdorff topology is [SSS14, Appendix B]. Some more information can be found in [BBI01, Chapter 7]. The basic properties of the Hausdorff topologies can be summarised in three lemmas. We first state all three lemmas, and then give proofs. The first lemma shows that the Hausdorff topology depends only on the topology on $\mathcal{X}$, and not on the choice of the metric.

Lemma 2.15 (Convergence criterion) *Let* $K_n, K \in \mathcal{K}_+(\mathcal{X})$ $(n \geq 1)$. *Then* $K_n \to K$ *in the Hausdorff topology if and only if the following three conditions are satisfied*

(i) *There exists a compact* $C \subset \mathcal{X}$ *such that* $K_n \subset C$ *for all* n.

(ii) $K = \big\{x \in \mathcal{X} : \exists x_n \in K_n \text{ s.t. } x_n \to x\big\}$.

(iii) $K = \big\{x \in \mathcal{X} : \exists x_n \in K_n \text{ s.t. } x \text{ is a cluster point of } (x_n)_{n\in\mathbb{N}}\big\}$.

The next lemma shows that $\mathcal{K}_+(\mathcal{X})$ is Polish if $\mathcal{X}$ is.

Lemma 2.16 (Properties of the Hausdorff metric)

(a) *If* $(\mathcal{X}, d)$ *is separable, then so is* $(\mathcal{K}_+(\mathcal{X}), d_{\rm H})$.

(b) *If* $(\mathcal{X}, d)$ *is complete, then so is* $(\mathcal{K}_+(\mathcal{X}), d_{\rm H})$.

The final and third lemma implies in particular that $\mathcal{K}_+(\mathcal{X})$ is compact if $\mathcal{X}$ is compact.

Lemma 2.17 (Compactness in the Hausdorff topology) *A set* $\mathcal{A} \subset \mathcal{K}_+(\mathcal{X})$ *is precompact if and only if there exists a compact* $C \subset \mathcal{X}$ *such that* $K \subset C$ *for each* $K \in \mathcal{A}$.

We now set out to prove Lemmas 2.15–2.17. We start by giving an alternative formulation of conditions (ii) and (iii) of Lemma 2.15.

Lemma 2.18 (Cluster and limit points) *Conditions (ii) and (iii) of Lemma 2.15 are equivalent to*

(ii)' $K = \big\{x \in \mathcal{X} : \lim_{n\to\infty} d(x, K_n) = 0\big\}$,

(iii)' $K = \big\{x \in \mathcal{X} : \liminf_{n\to\infty} d(x, K_n) = 0\big\}$.

Proof If for some $x \in \mathcal{X}$ there exist $(x_n)_{n\in\mathbb{N}}$ with $x_n \in K_n$ for all n and $x_n \to x \in K$, then $d(x, K_n) \leq d(x, x_n) \to 0$, and conversely, if $d(x, K_n) \to 0$, then we can choose $x_n \in K_n$ such that $d(x, x_n) \leq 2d(x, K_n) \to 0$, proving the first equality. Similarly, if for some $x \in \mathcal{X}$ there exist $(x_n)_{n\in\mathbb{N}}$ with $x_n \in K_n$ for all n and an infinite set $N \subset \mathbb{N}$ such that the subsequence $(x_n)_{n\in N}$ satisfies

$$\lim_{N\ni n\to\infty} d(x_n, x) = 0,$$

then clearly $\liminf_{n\to\infty} d(x, K_n) = 0$, and conversely, if this holds, then we can choose an infinite set $N \subset \mathbb{N}$ such that

$$\lim_{N\ni n\to\infty} d(x, K_n) = 0,$$

and for each $n \in N$ we can choose $x_n \in K_n$ such that $d(x, x_n) \leq 2d(x, K_n)$. Using the fact that $K_n \neq \emptyset$ for all n, we can extend $(x_n)_{n\in N}$ to a sequence $(x_n)_{n\in\mathbb{N}}$ with $x_n \in K_n$ for all $n \in \mathbb{N}$ that has x as a cluster point. ∎

Proof of Lemma 2.15 We claim that condition (ii) of Lemma 2.15 implies that

$$\sup_{x\in K} d(x, K_n) \underset{n\to\infty}{\longrightarrow} 0. \tag{2.5}$$

Indeed, if this is not the case, then there exist an $\varepsilon > 0$ and an infinite $N \subset \mathbb{N}$ such that $\sup_{x\in K} d(x, K_n) \geq 2\varepsilon$ for all $n \in N$. Then we can choose $x_n \in K$ such that $d(x_n, K_n) \geq \varepsilon$ for all $n \in N$. Since K is compact, the sequence $(x_n)_{n\in N}$ has a subsequence that converges to a limit in K, i.e., there exists an infinite $N' \subset N$ and $x \in K$ such that

$$x = \lim_{N'\ni n\to\infty} x_n.$$

Since $d(x, K_n) \geq d(x_n, K) - d(x, x_n) \geq \varepsilon - d(x, x_n)$ for all $n \in N'$, we have

$$\liminf_{N'\ni n\to\infty} d(x, K_n) \geq \varepsilon,$$

which contradicts condition (ii)' of Lemma 2.18. We next claim that conditions (i) and (iii) of Lemma 2.15 imply that

$$\sup_{x\in K_n} d(x, K) \underset{n\to\infty}{\longrightarrow} 0. \tag{2.6}$$

Indeed, if this is not the case, then there exist an $\varepsilon > 0$ and an infinite $N \subset \mathbb{N}$ such that $\sup_{x\in K_n} d(x, K) \geq 2\varepsilon$ for all $n \in N$. Then we can choose $x_n \in K_n$ such that $d(x_n, K) \geq \varepsilon$ for all $n \in N$. By condition (i), there exists a compact C such that $K_n \subset C$ for all n. It follows that there exists an infinite $N' \subset N$ and $x \in C$ such that

$$x = \lim_{N'\ni n\to\infty} x_n.$$

Since $d(x, K) \geq d(x_n, K) - d(x, x_n) \geq \varepsilon - d(x, x_n)$ for all $n \in N'$, taking the limit, we see that $d(x, K) \geq \varepsilon$ and hence $x \notin K$. On the other hand, x is a cluster point of the sequence $(x_n)_{n\in\mathbb{N}}$, so we arrive at a contradiction with condition (iii). Together, (2.5) and (2.6) show that (i)–(iii) imply that $d_{\rm H}(K_n, K) \to 0$ as $n \to \infty$.

Assume, conversely, that $d_{\rm H}(K_n, K) \to 0$ as $n \to \infty$. Set $K_\infty := K$ and $\overline{\mathbb{N}} := \mathbb{N} \cup \{\infty\}$. We claim that $C := \bigcup_{n \in \overline{\mathbb{N}}} K_n$ is compact. To see this, let $(x_k)_{k\in\mathbb{N}}$ be a sequence with $x_k \in C$ for all k. We claim that $(x_k)_{k\in\mathbb{N}}$ has a subsequence that converges to a limit in C. For each $k \in \mathbb{N}$, we can choose $n(k) \in \overline{\mathbb{N}}$ such that $x_k \in K_{n(k)}$. If there exists an $n \in \overline{\mathbb{N}}$ such that $n(k) = n$ for infinitely many values of k, then the claim follows from the compactness of K_n. In the opposite case, there exists an infinite set $N \subset \mathbb{N}$ such that $n(k) \in \mathbb{N}$ for each $k \in N$ and $n(k) \neq n(k')$ for each $k, k' \in N$ with $k \neq k'$. For each $k \in N$, we can find $x'_k \in K$ such that $d(x'_k, x_k) \leq 2d(K, K_{n(k)})$. Since K is compact, we can find an infinite set $N' \subset N$ such that the sequence $(x'_k)_{k\in N'}$ converges to a limit $x \in K$. Since $d(x, x_k) \leq d(x, x'_k) + 2d(K, K_{n(k)})$ tends to zero as $N' \ni k \to \infty$, we conclude that $(x'_k)_{k\in N'}$ converges to $x \in K \subset C$, proving the compactness of C. In particular, this proves that the sets K_n satisfy condition (i).

To see that conditions (ii) and (iii) hold too, we observe that for each $x \in K$, one has $d(x, K_n) \leq d(K, K_n) \to 0$ as $n \to \infty$, while for $x \notin K$, one has $d(x, K_n) \geq d(x, K) - d(K, K_n) \to d(x, K) > 0$ as $n \to \infty$. This shows that

$$\big\{x \in \mathcal{X} : \liminf_{n\to\infty} d(x, K_n) = 0\big\} \subset K \subset \big\{x \in \mathcal{X} : \lim_{n\to\infty} d(x, K_n) = 0\big\},$$

so by Lemma 2.18 we conclude that conditions (ii) and (iii) are satisfied. ∎

The proofs of Lemmas 2.16 and 2.17 need a little preparation. Recall that in any metric space $(\mathcal{X}, d)$, by definition, a set $A \subset \mathcal{X}$ is *totally bounded* if for every $\varepsilon > 0$ there exists a finite collection of points $x_1, \ldots, x_n \in \mathcal{X}$ such that $A \subset \bigcup_{i=1}^n B_\varepsilon(x_i)$, where $B_\varepsilon(x)$ denotes the open ball of radius ε around x. It is well-known that total boundedness is equivalent to the statement that every sequence $x_n \in A$ has a Cauchy subsequence. As a consequence, a set $A \subset \mathcal{X}$ is compact if and only if it is complete and totally bounded.

Lemma 2.19 (Totally bounded sets in the Hausdorff metric) *A set $\mathcal{A} \subset \mathcal{K}_+(\mathcal{X})$ is totally bounded in the metric $d_{\rm H}$ if and only if the set $A := \{x \in \mathcal{X} : \exists K \in \mathcal{A} \text{ s.t. } x \in K\}$ is totally bounded in the metric d.*

Proof Let $B_\varepsilon(x)$ denote the open ball in $\mathcal{X}$ of radius ε around a point $x \in \mathcal{X}$, and let $\mathcal{B}_\varepsilon(x)$ denote the open ball in $\mathcal{K}_+(\mathcal{X})$ of radius ε around a point $K \in \mathcal{K}_+(\mathcal{X})$.

Assume that A is totally bounded. Let $\varepsilon > 0$ and let $\Delta \subset \mathcal{X}$ be a finite set such that $A = \bigcup_{x\in\Delta} B_\varepsilon(x)$. Let $K \in \mathcal{A}$ and set $\Delta' := \{x \in \Delta : B_\varepsilon(x) \cap K \neq$

$\emptyset\}$. Then for all $y \in K$ there is an $x \in \Delta'$ such that $d(x,y) < \varepsilon$ and for all $x \in \Delta'$ there is a $y \in K$ such that $d(x,y) < \varepsilon$ proving that $d_{\rm H}(\Delta', K) < \varepsilon$. This shows that

$$\mathcal{A} \subset \bigcup_{\Delta'} \mathcal{B}_\varepsilon(\Delta'),$$

where we take the union over all nonempty subsets $\Delta' \subset \Delta$. Since there are finitely many such Δ', and each Δ' is a compact subset of $\mathcal{X}$, this proves that $\mathcal{A}$ is totally bounded.

Conversely, if $\mathcal{A}$ is totally bounded, then for each $\varepsilon > 0$ we can find $K_1, \ldots, K_n \in \mathcal{K}_+(\mathcal{X})$ such that $\mathcal{A} \subset \bigcup_{k=1}^n \mathcal{B}_{\varepsilon/2}(K_n)$, where $\mathcal{B}_\varepsilon(K)$ denotes the open ball in the Hausdorff metric of radius ε centered around a compact set K. Since compact sets are totally bounded, for each k we can find finitely many points $x_{k,1}, \ldots, x_{k,m_k} \in \mathcal{X}$ such that $K_k \subset \bigcup_{j=1}^{m_k} B_{\varepsilon/2}(x_{k,j})$. It follows that $A \subset \bigcup_{k=1}^n \bigcup_{j=1}^{m_k} B_\varepsilon(x_{k,j})$, showing that A is totally bounded. ∎

Proof of Lemma 2.16 To prove part (a), let $\mathcal{D}$ be a countable dense subset of $(\mathcal{X}, d)$, and let $\mathcal{A}$ be the set of all finite nonempty subsets of $\mathcal{D}$. Then $\mathcal{A} \subset \mathcal{K}_+(\mathcal{X})$ and $\mathcal{A}$ is countable, so it suffices to prove that $\mathcal{A}$ is dense in $(\mathcal{K}_+(\mathcal{X}), d_{\rm H})$. We will show that for each $K \in \mathcal{K}_+(\mathcal{X})$ and $\varepsilon > 0$, there exists a set $A \in \mathcal{A}$ such that $d_{\rm H}(A, K) \leq \varepsilon$. Since K is compact, it is totally bounded, so we can find a finite set $B \subset \mathcal{X}$ such that $K \subset \bigcup_{x \in B} B_{\varepsilon/2}(x)$. Without loss of generality, we can assume that $d(x, K) < \varepsilon/2$ for all $x \in B$. Since $\mathcal{D}$ is dense, for each $x \in B$ we can find an $x' \in \mathcal{D}$ such that $d(x, x') < \varepsilon/2$. Then $A := \{x' : x \in B\}$ is a finite subset of $\mathcal{D}$ such that for all $y \in K$, there exists a $z \in A$ such that $d(y,z) < \varepsilon$ and conversely, for all $z \in A$, there exists a $y \in K$ such that $d(y,z) < \varepsilon$, proving that $d_{\rm H}(A, K) \leq \varepsilon$.

To prove part (b), let $K_n \in \mathcal{K}_+(\mathcal{X})$ be a Cauchy sequence and let

$$A := \big\{x \in \mathcal{X} : \lim_{n\to\infty} d(x, K_n) = 0\big\}, \quad B := \big\{x \in \mathcal{X} : \liminf_{n\to\infty} d(x, K_n) = 0\big\}.$$

We claim that $A = B$. Indeed, if there exists some $x \in B \backslash A$, then there is some $\varepsilon > 0$ such that for each $k \geq 1$ we can find $n, m \geq k$ such that $d(x, K_n) \leq \varepsilon$ and $d(x, K_m) \geq 2\varepsilon$, hence $d_{\rm H}(K_n, K_m) \geq \varepsilon$, contradicting the assumption that the K_n form a Cauchy sequence.

Let $K := A = B$. We claim that K is closed. To prove this, we will show that if $x_k \in A$ satisfy $x_k \to x$ for some $x \in \mathcal{X}$, then $x \in B$. Since $x_k \in A$ we can find $x_{k,n} \in K_n$ such that $x_{k,n} \to x_k$ as $n \to \infty$. For each k, we can choose $n(k) \geq k$ such that $d(x_{k,n(k)}, x_k) \leq d(x_k, x)$. Then $n(k) \to \infty$ and $d(x, K_{n(k)}) \leq d(x_{k,n(k)}, x) \leq 2d(x_k, x) \to 0$ as $k \to \infty$ and hence $x \in B$.

We next claim that K is compact. Since each sequence in the set $\{K_n : n \geq 1\}$ contains a Cauchy subsequence, the set $\{K_n : n \geq 1\}$ is totally bounded, hence by Lemma 2.19, there exists some totally bounded set containing all of the K_n. Let C denote its closure. Then C is compact since $\mathcal{X}$ is complete, hence also $K \subset C$ is compact since K is closed.

Since the sets K_n are contained in the compact set C and since $A = B$, we see that conditions (i)–(iii) of Lemma 2.15 are satisfied, so we conclude that $d_{\rm H}(K_n, K) \to 0$ as $n \to \infty$. We have shown that every Cauchy sequence in $\big(\mathcal{K}_+(\mathcal{X}), d_{\rm H}\big)$ is converges, i.e., this space is complete. ∎

Proof of Lemma 2.17 Assume that $\mathcal{A} \subset \mathcal{K}_+(\mathcal{X})$ and that there exists a compact $C \subset \mathcal{X}$ such that $K \subset C$ for all $K \in \mathcal{A}$. Then C is totally bounded and complete, so by Lemmas 2.19 and 2.16 (b), the same is true for $\{K \in \mathcal{K}_+(\mathcal{X}) : K \subset \mathcal{X}\}$, implying the latter is compact and hence its subset $\mathcal{A}$ is precompact.

Conversely, if $\mathcal{A} \subset \mathcal{K}_+(\mathcal{X})$ is precompact, then its closure $\overline{\mathcal{A}}$ in the metric $d_{\rm H}$ is compact. To complete the proof, it suffices to show that $C = \{x \in \mathcal{X} : \exists K \in \overline{\mathcal{A}} \text{ s.t. } x \in K\}$ is compact. Since $\overline{\mathcal{A}}$ is compact, it is totally bounded, so Lemma 2.19 implies that C is totally bounded too. It therefore suffices to show that C is complete. Any cluster point of a Cauchy sequence must necessarily be a limit point. Therefore, to show that C is complete, it suffices to show that each Cauchy sequence $x_n \in C$ has a cluster point $x \in C$. Choose $K_n \in \overline{\mathcal{A}}$ such that $x_n \in K_n$. Since $\overline{\mathcal{A}}$ is compact, by going to a subsequence if necessary, we may assume that $K_n \to K$ for some $K \in \overline{\mathcal{A}}$. Since $d_{\rm H}(K_n, K) \to 0$, we can choose $x'_n \in K$ such that $d(x_n, x'_n) \to 0$. Since K is compact, by going to a further subsequence if necessary, we may assume that $x'_n \to x$ for some $x \in K$. Since $d(x_n, x) \leq d(x_n, x'_n) + d(x'_n, x) \to 0$ this proves that the original sequence x_n has a cluster point $x \in K \subset C$. ∎

We conclude this section with two more lemmas. The first lemma is useful when proving convergence of $\mathcal{K}_+(\mathcal{X})$-valued random variables.

Lemma 2.20 (Tightness criterion) *Assume that $\mathcal{X}$ is a Polish space and let K_n $(n \in \mathbb{N})$ be $\mathcal{K}_+(\mathcal{X})$-valued random variables. Then the collection of laws $\{\mathbb{P}[K_n \in \cdot\,] : n \in \mathbb{N}\}$ is tight if and only if for each $\varepsilon > 0$ there exists a compact $C \subset \mathcal{X}$ such that $\mathbb{P}[K_n \subset C] \geq 1 - \varepsilon$ for all $n \in \mathbb{N}$.*

Proof This is an immediate consequence of Lemma 2.17. Indeed, if a $C \subset \mathcal{X}$ is compact, then by Lemma 2.17, the set $\mathcal{C} := \{K \in \mathcal{K}_+(\mathcal{X}) : K \subset C\}$ is compact, so it is clear that the conditions of the lemma imply tightness of the laws $\{\mathbb{P}[K_n \in \cdot\,] : n \in \mathbb{N}\}$. Conversely, if these laws are tight, then for each $\varepsilon > 0$ there exists a compact $\mathcal{C} \subset \mathcal{K}_+(\mathcal{X})$ such that $\mathbb{P}[K_n \in \mathcal{C}] \geq 1 - \varepsilon$ for all $n \in \mathbb{N}$, which by Lemma 2.17 implies the existence of a compact $C \subset \mathcal{X}$ such that $K \subset C$ for all $K \in \mathcal{C}$ and hence also $\mathbb{P}[K_n \subset C] \geq 1 - \varepsilon$. ∎

The next and final lemma of this section says that if $\psi : \mathcal{X} \to \mathcal{Y}$ is a continuous map between metrisable topological spaces, then $\mathcal{K}_+(\mathcal{X}) \ni K \mapsto \psi(K) \in \mathcal{K}_+(\mathcal{Y})$ is also a continuous map.

Lemma 2.21 (Map acting on compact sets) *Let $\mathcal{X}$ and $\mathcal{Y}$ be metrisable topological spaces, let $\psi : \mathcal{X} \to \mathcal{Y}$ be a continuous map, and let*

$$\hat{\psi}(K) := \big\{\psi(x) : x \in K\big\} \qquad \big(K \in \mathcal{K}_+(\mathcal{X})\big)$$

denote the image of a compact set $K \subset \mathcal{X}$ under ψ. Then $\psi(K) \in \mathcal{K}_+(\mathcal{Y})$ for all $K \in \mathcal{K}_+(\mathcal{X})$, and the map $\hat{\psi} : \mathcal{K}_+(\mathcal{X}) \to \mathcal{K}_+(\mathcal{X})$ is continuous with respect to the Hausdorff topology.

Proof The well-known fact that the continuous image of a compact set is itself a compact set has already been mentioned at the beginning of Section 2.2. To see that $\hat{\psi} : \mathcal{K}_+(\mathcal{X}) \to \mathcal{K}_+(\mathcal{X})$ is continuous, assume that $K_n \to K$. Then by Lemma 2.15,

$$\exists C \in \mathcal{K}_+(\mathcal{X}) \text{ s.t. } K_n \subset C \ \forall n \geq 1 \tag{2.7}$$

and

$$\begin{aligned} K &= \{x \in \mathcal{X} : \exists x_n \in K_n \text{ s.t. } x_n \to x\} \\ &= \{x \in \mathcal{X} : \exists x_n \in K_n \text{ s.t. } x \text{ is a cluster point of } (x_n)_{n\in\mathbb{N}}\}. \end{aligned} \tag{2.8}$$

Since $\hat{\psi}(C)$ is compact and $\hat{\psi}(K_n) \subset \hat{\psi}(C)$ for all $n \geq 1$, by Lemma 2.15, to prove that $\hat{\psi}(K_n) \to \hat{\psi}(K)$, it suffices to show that

$$\begin{aligned} \hat{\psi}(K) &= \{y \in \mathcal{X} : \exists y_n \in \hat{\psi}(K_n) \text{ s.t. } y_n \to y\} \\ &= \{y \in \mathcal{X} : \exists y_n \in \hat{\psi}(K_n) \text{ s.t. } y \text{ is a cluster point of } (y_n)_{n\in\mathbb{N}}\}. \end{aligned}$$

The latter condition can be rewritten as

$$\begin{aligned} \big\{\psi(x) : x \in K\big\} &= \{y \in \mathcal{X} : \exists x_n \in K_n \text{ s.t. } \psi(x_n) \to y\} \\ &= \{y \in \mathcal{X} : \exists x_n \in K_n \text{ s.t. } y \text{ is a cluster point of } \big(\psi(x_n)\big)_{n\in\mathbb{N}}\}. \end{aligned}$$

It therefore suffices to prove that

(i) $\big\{\psi(x) : x \in K\big\} \subset \{y \in \mathcal{X} : \exists x_n \in K_n \text{ s.t. } \psi(x_n) \to y\}$,

(ii) $\{y \in \mathcal{X} : \exists x_n \in K_n \text{ s.t. } y \text{ is a cluster point of } \big(\psi(x_n)\big)_{n\in\mathbb{N}}\} \subset \big\{\psi(x) : x \in K\big\}$.

To prove (i), we use that by (2.8), for each $x \in K$ there exist $x_n \in K_n$ such that $x_n \to x$, and hence $\psi(x_n) \to \psi(x)$ by the continuity of ψ. To prove (ii), assume that $x_n \in K_n$ $(n \in \mathbb{N})$ and there exists a sequence $(n(m))_{m\geq 1}$

with $\lim_{m\to\infty} n(m) = \infty$ such that $y = \lim_{m\to\infty} \psi(x_{n(m)})$. By (2.7) and the compactness of C, by going to a further subsequence if necessary, we can assume without loss of generality that $\lim_{m\to\infty} x_{n(m)} = x$ for some $x \in C$. Then $x \in K$ by (2.8) and $\lim_{m\to\infty} \psi(x_{n(m)}) = \psi(x)$ by the continuity of ψ which shows that $y = \psi(x)$. ∎

2.6 Squeezed space

Let $(\mathcal{X}, d)$ be a metric space, let $\{*\}$ be a set containing a single element called $*$, which we assume is not an element of $\mathcal{X}$, and let

$$\mathcal{R}(\mathcal{X}) := \big(\mathcal{X} \times \mathbb{R}) \cup \big\{(*, -\infty), (*, +\infty)\big\}. \tag{2.9}$$

See Figure 2.1 for a picture of $\mathcal{R}(\overline{\mathbb{R}})$, with $\overline{\mathbb{R}} := [-\infty, \infty]$. We will show that it is possible to equip $\mathcal{R}(\mathcal{X})$ with a metrisable topology such that $(x_n, t_n) \to (x, t) \in \mathcal{X} \times \mathbb{R}$ if and only if $x_n \to x$ and $t_n \to t$, while $(x_n, t_n) \to (*, \pm\infty)$ if and only if $t_n \to \pm\infty$ (with no conditions on x_n). We can think of the space $\mathcal{R}(\mathcal{X})$ as being obtained from $\mathcal{X} \times \overline{\mathbb{R}}$ by squeezing the sets $\mathcal{X} \times \{\pm\infty\}$ into the single points $(*, \pm\infty)$. For this reason, we call $\mathcal{R}(\mathcal{X})$ the *squeezed space*.

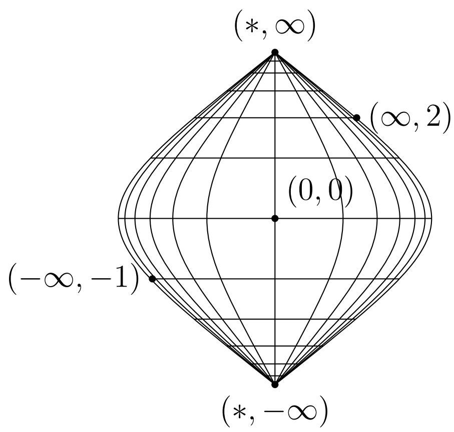

Figure 2.1: The squeezed space $\mathcal{R}(\overline{\mathbb{R}})$.

To define a metric on $\mathcal{R}(\mathcal{X})$ with the desired properties, we first extend d to $\mathcal{X} \cup \{*\}$ by setting $d(x, *) = d(*, x) := \infty$ if $x \neq *$ and $:= 0$ otherwise. Let $\overline{\mathbb{R}} := [-\infty, \infty]$ denote the usual two-point compactification of the real line. We fix a continuous function $\phi : \overline{\mathbb{R}} \to [0, \infty)$ such that $\phi(t) > 0$ for all

$t \in \mathbb{R}$ and $\phi(\pm\infty) = 0$, we choose a metric $d_{\overline{\mathbb{R}}}$ that generates the topology on $\overline{\mathbb{R}}$, and we define $d_{\rm sqz} : \mathcal{R}(\mathcal{X})^2 \to [0, \infty)$ by

$$d_{\rm sqz}\big((x,s),(y,t)\big) := \big(\phi(s)\wedge\phi(t)\big)\big(d(x,y)\wedge 1\big) + \big|\phi(s)-\phi(t)\big| + d_{\overline{\mathbb{R}}}(s,t) \quad (2.10)$$

Lemma 2.22 (Metric on squeezed space) *The function $d_{\rm sqz}$ is a metric on $\mathcal{R}(\mathcal{X})$.*

Proof For brevity, we write $d'(x,y) := d(x,y) \wedge 1$. Then d' is a metric on $\mathcal{X}$. The only nontrivial statement that we have to prove is the triangle inequality, and it suffices to prove this for the function

$$d'_{\rm sqz}\big((x,s),(y,t)\big) := \big(\phi(s) \wedge \phi(t)\big) d'(x,y) + \big|\phi(s) - \phi(t)\big|.$$

We estimate

$$d'_{\rm sqz}\big((x,s),(z,u)\big) \leq \big(\phi(s)\wedge\phi(u)\big)\big(d'(x,y)+d'(y,z)\big) + \big|\phi(s)-\phi(u)\big|. \quad (2.11)$$

If $\phi(t) \geq \phi(s) \wedge \phi(u)$, then $\phi(s) \wedge \phi(u)$ is less than $\phi(s) \wedge \phi(t)$ and also less than $\phi(t) \wedge \phi(u)$, so we can simply estimate the expression in (2.11) from above by

$$\big(\phi(s) \wedge \phi(t)\big) d'(x,y) + \big(\phi(t) \wedge \phi(u)\big) d'(y,z)\big) + \big|\phi(s) - \phi(t)\big| + \big|\phi(t) - \phi(u)\big|$$

and we are done. On the other hand, if $\phi(t) < \phi(s) \wedge \phi(u)$, then

$$\big|\phi(s) - \phi(t)\big| + \big|\phi(t) - \phi(u)\big| = \big|\phi(s) - \phi(u)\big| + 2\big(\phi(s) \wedge \phi(u) - \phi(t)\big).$$

Using the fact that $d' \leq 1$, we can now estimate the right-hand side of (2.11) from above by

$$\begin{aligned} &\phi(t)\big(d'(x,y) + d'(y,z)\big) + 2\big(\phi(s) \wedge \phi(u) - \phi(t)\big) + \big|\phi(s) - \phi(u)\big| \\ &\quad = \big(\phi(s) \wedge \phi(t)\big) d'(x,y) + \big(\phi(t) \wedge \phi(u)\big) d'(y,z) \\ &\qquad + \big|\phi(s) - \phi(t)\big| + \big|\phi(t) - \phi(u)\big|, \end{aligned}$$

and again we are done. ∎

The following lemma shows that the topology generated by the metric $d_{\rm sqz}$ has the desired properties we stated earlier. In particular, this lemma shows that the topology generated by the metric $d_{\rm sqz}$ depends only on the topology on $\mathcal{X}$ and not on the choice of the metric on $\mathcal{X}$. Recall that by (2.1), a metrisable topology is uniquely characterised by its convergent sequences, so the topology on $\mathcal{R}(\mathcal{X})$ is uniquely characterised by the conditions (i) and (ii) below.

Lemma 2.23 (Topology on squeezed space) *A sequence* $(x_n, t_n) \in \mathcal{R}(\mathcal{X})$ *converges to a limit* (x, t) *in the metric* $d_{\rm sqz}$ *defined in (2.10) if and only if the following two conditions are satisfied:*

(i) $t_n \to t$ *in the topology on* $\overline{\mathbb{R}}$,

(ii) *if* $t \in \mathbb{R}$, *then* $x_n \to x$ *in the topology on* $\mathcal{X}$.

Proof This is immediate from the definition of $d_{\rm sqz}$. ∎

The following lemma shows that $\mathcal{R}(\mathcal{X})$ is a Polish space if $\mathcal{X}$ is Polish.

Lemma 2.24 (Properties of squeezed space)

(a) *If* $(\mathcal{X}, d)$ *is separable, then so is* $(\mathcal{R}(\mathcal{X}), d_{\rm sqz})$.

(b) *If* $(\mathcal{X}, d)$ *is complete, then so is* $(\mathcal{R}(\mathcal{X}), d_{\rm sqz})$.

Proof If D is a countable dense subset of $(\mathcal{X}, d)$, then $D \times \mathbb{Q}$ is a countable dense subset of $(\mathcal{R}(\mathcal{X}), d_{\rm sqz})$, proving (a).

To prove (b), let (x_n, t_n) be a Cauchy sequence in $(\mathcal{R}(\mathcal{X}), d_{\rm sqz})$. Then by (2.10) t_n is a Cauchy sequence in $\overline{\mathbb{R}}$ and hence $t_n \to t$ for some $t \in \overline{\mathbb{R}}$. If $t \in \mathbb{R}$, then by (2.10) x_n is a Cauchy sequence in $(\mathcal{X}, d)$ so by the completeness of the latter, $x_n \to x$ for some $x \in \mathcal{X}$. By Lemma 2.23, it follows that (x_n, t_n) converges, proving the completeness of $(\mathcal{R}(\mathcal{X}), d_{\rm sqz})$. ∎

The following lemma identifies the compact subsets of $\mathcal{R}(\mathcal{X})$. In particular, the lemma shows that $\mathcal{R}(\mathcal{X})$ is compact if $\mathcal{X}$ is compact.

Lemma 2.25 (Compactness criterion) *A set* $A \subset \mathcal{R}(\mathcal{X})$ *is precompact if and only if for each* $T < \infty$, *there exists a compact set* $K \subset \mathcal{X}$ *such that* $\{x \in \mathcal{X} : (x, t) \in A,\ t \in [-T, T]\} \subset K$.

Proof Assume that $A \subset \mathcal{R}(\mathcal{X})$ has the property that for each $T < \infty$, there exists a compact set $K \subset \mathcal{X}$ such that $\{x \in \mathcal{X} : (x, t) \in A,\ t \in [-T, T]\} \subset K$. To show that A is precompact, we will show that each sequence $(x_n, t_n) \in A$ has a convergent subsequence. By the compactness of $\overline{\mathbb{R}}$, we can select a subsequence (x'_n, t'_n) such that $t'_n \to t$ for some $t \in \overline{\mathbb{R}}$. If $t = \pm\infty$, then by Lemma 2.23 $(x'_n, t'_n) \to (*, \pm\infty)$ and we are done. Otherwise, there exists a $T < \infty$ such that $t'_n \in [-T, T]$ for all n large enough. By assumption, there then exists a compact set $K \subset \mathcal{X}$ such that $x'_n \in K$ for all n large enough, so we can select a further subsequence such that (x''_n, t''_n) converges to a limit $(x, t) \in \mathcal{X} \times \mathbb{R}$.

Assume, on the other hand, that $A \subset \mathcal{R}(\mathcal{X})$ has the property that for some $T < \infty$, there does not exist a compact set $K \subset \mathcal{X}$ such that $\{x \in \mathcal{X} : (x, t) \in A,\ t \in [-T, T]\} \subset K$. Set

$$B := \big\{x \in \mathcal{X} : (x, t) \in A \text{ for some } t \in [-T, T]\big\}$$

The closure of B cannot be compact, since this would contradict our assumption. It follows that there exists a sequence $x_n \in B$ that does not contain a convergent subsequence, and there exist $t_n \in [-T, T]$ such that $(x_n, t_n) \in A$. But then, in view of Lemma 2.23, the sequence (x_n, t_n) cannot contain a convergent subsequence either, proving that A is not precompact. ∎

2.7 Path space

Let $\mathcal{X}$ be a metrisable space and let $\mathcal{R}(\mathcal{X})$ be the squeezed space defined in Section 2.6. By definition, a *path* in $\mathcal{X}$ is a nonempty compact subset $\pi \subset \mathcal{R}(\mathcal{X})$ such that $\{x \in \mathcal{X} : (x,t) \in \pi\}$ has at most one element for each given $t \in \mathbb{R}$ and the set

$$\overline{I}_\pi := \{t \in \overline{\mathbb{R}} : \exists x \in \mathcal{X} \cup \{*\} \text{ s.t. } (x,t) \in \pi\} \tag{2.12}$$

is a closed subinterval of $\overline{\mathbb{R}}$. We call $I_\pi := \overline{I}_\pi \cap \mathbb{R}$ the *domain* of π and we call

$$\sigma_\pi := \inf \overline{I}_\pi \quad \text{and} \quad \tau_\pi := \sup \overline{I}_\pi \tag{2.13}$$

the *starting time* and *final time* of the path π. For each $t \in \overline{I}_\pi$, we define $\pi(t) \in \mathcal{X} \cup \{*\}$ by $\{\pi(t)\} := \{x \in \mathcal{X} : (x,t) \in \pi\}$. Then $I_\pi \ni t \mapsto \pi(t)$ is a function from I_π to $\mathcal{X}$. We let $\Pi(\mathcal{X})$ denote the set of all paths in $\mathcal{X}$.

Lemma 2.26 (Path viewed as a function) *The domain I_π of a path $\pi \in \Pi(\mathcal{X})$ is a closed subinterval of $\mathbb{R}$, and $t \mapsto \pi(t)$ is a continuous function from I_π to $\mathcal{X}$. Conversely, if $I \subset \mathbb{R}$ is a closed interval and $t \mapsto f(t)$ is a continuous function from I to $\mathcal{X}$, then there exists a path $\pi \in \Pi(\mathcal{X})$ such that $I_\pi = I$ and $\pi(t) = f(t)$ $(t \in I)$. The path π is uniquely determined by the interval I and function f, except in the trivial case when $I = \emptyset$, in which case there are two possible choices for π.*

Proof We first show that for each $\pi \in \Pi(\mathcal{X})$, the function $I_\pi \ni t \mapsto \pi(t)$ is continuous. Assume that $t_n, t \in I_\pi$ and $t_n \to t$. Since π is compact, the sequence $(\pi(t_n), t_n)$ is precompact. Since $\pi(t)$ is the only element of $\{x \in \mathcal{X} : (x,t) \in \pi\}$, each subsequence of the $(\pi(t_n), t_n)$ must converge to $(\pi(t), t)$. By Lemma 2.2, we conclude that $(\pi(t_n), t_n) \to (\pi(t), t)$. Since $t \in \mathbb{R}$, by Lemma 2.23, we conclude that $\pi(t_n) \to \pi(t)$, which shows that $I_\pi \ni t \mapsto \pi(t)$ is continuous on I as claimed.

Let $I \subset \mathbb{R}$ be a closed interval and let $f : I \to \mathcal{X}$ be continuous. Assume that I is nonempty. Let $\overline{I}$ be the closure of I in $\overline{\mathbb{R}}$. Extend f to $\overline{I}$ by setting

$f(t) := *$ if $t = \pm\infty$. Let $\pi := \{(f(t), t) : t \in \overline{I}\}$. It follows from Lemma 2.23 and the continuity of f that the map

$$\overline{I} \ni t \mapsto \big(f(t), t\big) \in \mathcal{R}(\mathcal{X}) \tag{2.14}$$

is continuous. Since $\overline{I}$ is compact and since π is the image of $\overline{I}$ under the continuous map (2.14), we conclude that π is compact. Clearly, $\{x \in \mathcal{X} : (x, t) \in \pi\}$ has precisely one element for $t \in \overline{I}$, and is empty for $t \notin \overline{I}$. This shows that $\pi \in \Pi(\mathcal{X})$. Since $\overline{I}$ is the only closed subinterval of $\overline{\mathbb{R}}$ such that $\overline{I} \cap \mathbb{R} = I$, we see that π is uniquely determined by the interval I and function f.

In the special case that $I = \emptyset$, it is easy to see that there exist precisely two paths π such that $I_\pi = I$ (the condition $\pi(t) = f(t)$ $(t \in I)$ is void in this case). These are the trivial paths with $\overline{I}_\pi = \{-\infty\}$ or $= \{\infty\}$, respectively. ∎

In view of Lemma 2.26, we often view a path $\pi \in \Pi(\mathcal{X})$ as a continuous function defined on a closed interval $I_\pi \subset \mathbb{R}$. If $I \subset \mathbb{R}$ is a closed nonempty interval, then we identify the space $\mathcal{C}_I(\mathcal{X})$ defined in Section 2.4 with the subset of $\Pi(\mathcal{X})$ defined as $\{\pi \in \Pi(\mathcal{X}) : I_\pi = I\}$.

Let $\mathcal{K}_+(\mathcal{R}(\mathcal{X}))$ be the set of nonempty compact subsets of the squeezed space $\mathcal{R}(\mathcal{X})$. We equip $\mathcal{K}_+(\mathcal{R}(\mathcal{X}))$ with the Hausdorff topology. We observe that $\Pi(\mathcal{X})$ is a subset of $\mathcal{K}_+(\mathcal{R}(\mathcal{X}))$. We naturally equip $\Pi(\mathcal{X})$ with the induced topology from its embedding in $\mathcal{K}_+(\mathcal{R}(\mathcal{X}))$.

Lemma 2.27 (Paths with a fixed domain) *Let $I \subset \mathbb{R}$ be a closed nonempty interval. The induced topology on $\mathcal{C}_I(\mathcal{X})$ from its embedding in $\Pi(\mathcal{X})$ is the topology of locally uniform convergence.*

Proof Assume that $\pi_n, \pi \in \mathcal{C}_I(\mathcal{X})$, viewed as functions, satisfy $\pi_n \to \pi$ locally uniformly. We need to show that viewed as compact subsets of $\mathcal{R}(\mathcal{X})$, the sets π_n, π satisfy $\pi_n \to \pi$ in the Hausdorff topology on $\mathcal{K}_+(\mathcal{R}(\mathcal{X}))$. Let $\overline{I}$ denote the closure of I in $\overline{\mathbb{R}}$. By Lemma 2.15, we need to show that $\bigcup_n \pi_n$ is precompact and

$$\begin{aligned}
&\pi \subset \big\{(x, t) \in \mathcal{R}(\mathcal{X}) : \exists t_n \in \overline{I} \text{ s.t. } \big(\pi_n(t_n), t_n\big) \to (x, t)\big\},\\
&\big\{(x, t) \in \mathcal{R}(\mathcal{X}) : (x, t) \text{ is a cluster}\\
&\qquad\qquad \text{point of } \big(\pi_n(t_n), t_n\big) \text{ for some } t_n \in \overline{I}\big\} \subset \pi.
\end{aligned} \tag{2.15}$$

To see that $\bigcup_n \pi_n$ is precompact, we need to show that each sequence of the form $(\pi_{n(m)}(t_m), t_m)_{m\geq 1}$ has a convergent subsequence. If $n(m)$ infinitely often takes the same value n, then the claim is obvious from the compactness of π_n, so without loss of generality we may assume that $n(m) \to \infty$. Going to a subsequence if necessary, we may assume that $t_m \to t$ for some $t \in \overline{I}$. If $t = \pm\infty$, then the claim is again obvious so we may assume that $t \in I$. In this case Lemma 2.13 (ii) tells us that $\pi_{n(m)}(t_m) \to \pi(t)$ so we have found a convergent subsequence as required.

To prove the first inclusion in (2.15), let $(\pi(t), t) \in \pi$ and set $t_n := t$ for all n. If $t \in I$, then $\pi_n(t) \to \pi(t)$ since locally uniform convergence implies pointwise convergence, and if $t = \pm\infty$ then trivially $(*, t) \to (*, t)$ as $n \to \infty$. To prove the second inclusion, assume that $(\pi_{n(m)}(t_{n(m)}), t_{n(m)}) \to (x, t)$ as $m \to \infty$ for some $(x, t) \in \mathcal{R}(\mathcal{X})$, $t_n \in \overline{I}$, and $n(m) \to \infty$. If $t \in I$, then we can use Lemma 2.13 (ii) which tells us that $\pi_{n(m)}(t_{n(m)}) \to \pi(t)$ and hence $(x, t) = (\pi(t), t) \in \pi$. If $t = \pm\infty$, then trivially $x = *$ and $(*, t) \in \pi$.

Assume, conversely, that $\pi_n \to \pi$ in the Hausdorff topology on $\mathcal{K}_+(\mathcal{R}(\mathcal{X}))$. We need to show that $\pi_n, \pi \in \mathcal{C}_I(\mathcal{X})$ and that $\pi_n \to \pi$ locally uniformly. Assume that $t_n, t \in I$ such that $t_n \to t$. By Lemma 2.13 (ii), it suffices to show that $\pi_n(t_n) \to \pi(t)$ for all such t_n, t. Equivalently, we may show that $(\pi_n(t_n), t_n) \to (\pi(t), t)$. By Lemma 2.2, it suffices to show that the set $\{(\pi_n(t_n), t_n) : n \in \mathbb{N}\}$ is precompact and $(\pi(t), t)$ is the only cluster point of the sequence $(\pi_n(t_n), t_n)$. By Lemma 2.15, there exists a compact set $C \subset \mathcal{R}(\mathcal{X})$ such that $\pi_n \subset C$ for all n, so $\{(\pi_n(t_n), t_n) : n \in \mathbb{N}\}$ is precompact as required. Let (x, t) be any cluster point. By Lemma 2.15 (ii), $(x, t) \in \pi$ and hence $x = \pi(t)$, which shows that $\pi_n(t_n) \to \pi(t)$ as required. ∎

Our next proposition says that the space of paths in $\mathcal{X}$ is Polish provided $\mathcal{X}$ has this property.

Proposition 2.28 (Polish space) *If $\mathcal{X}$ is a Polish space, then so is $\Pi(\mathcal{X})$.*

The proof of Proposition 2.28 needs some preparations. Let d be a metric generating the topology on $\mathcal{X}$ and let $\pi \in \Pi(\mathcal{X})$. For each $\pi \in \Pi(\mathcal{X})$, $\delta > 0$ and $T < \infty$, we define

$$m_{T,\delta}(\pi) := \sup\big\{d\big(\pi(s), \pi(t)\big) : s, t \in I_\pi,\ -T \le s \le t \le T,\ t - s \le \delta\big\}. \tag{2.16}$$

The quantity $m_{T,\delta}(\pi)$ is called the *modulus of continuity* of the path π. More generally, for any compact subset $K \subset \mathcal{R}(\mathcal{X})$, we can define

$$m_{T,\delta}(K) := \sup\big\{d\big(x, y\big) : (x, s), (y, t) \in K,\ -T \le s \le t \le T,\ t - s \le \delta\big\},$$

which coincides with our previous definition if π is a path. In analogy with (2.12), we also define

$$\overline{I}_K := \big\{t \in \overline{\mathbb{R}} : \exists x \in \mathcal{X} \cup \{*\} \text{ s.t. } (x, t) \in K\big\}.$$

Lemma 2.29 (Characterisation of paths) *A compact subset* $\pi \subset \mathcal{R}(\mathcal{X})$ *is an element of the path space* $\Pi(\mathcal{X})$ *if and only if* $\overline{I}_K$ *is a closed subinterval of* $\overline{\mathbb{R}}$ *and* $\lim_{\delta\to 0} m_{T,\delta}(\pi) = 0$ *for all* $T < \infty$.

Proof Assume that $\pi \in \mathcal{K}_+(\mathcal{R}(\mathcal{X}))$ and $\limsup_{\delta\to 0} m_{T,\delta}(\pi) > 0$ for some $T < \infty$. Then we can find $(x_n, s_n), (y_n, t_n) \in \pi$ and $\delta > 0$ with $d(x_n, y_n) \geq \delta$, $-T \leq s_n \leq t_n \leq T$, and $t_n - s_n \leq 1/n$. Since π is compact, by going to a subsequence, we can assume that $(x_n, s_n) \to (x, s)$ and $(y_n, t_n) \to (y, t)$ for some $(x, s), (y, t) \in \pi$ with $d(x, y) \geq \delta > 0$, $-T \leq s \leq t \leq T$, and $t - s = 0$. This shows that $\pi \notin \Pi(\mathcal{X})$.

Conversely, if $\pi \notin \Pi(\mathcal{X})$, then either $\overline{I}_\pi$ is not a closed subinterval of $\overline{\mathbb{R}}$ or there exist $(x, t), (y, t) \in \pi$ with $x \neq y$. In the latter case, since $(*, \pm\infty)$ are the only points in $\mathcal{R}(\mathcal{X})$ with time coordinate $\pm\infty$ we must have $t \in \mathbb{R}$. But then $m_{T,\delta}(\pi) \geq d(x, y) > 0$ for all $T \geq |t|$, which shows that $\limsup_{\delta\to 0} m_{T,\delta}(\pi) > 0$ for some $T < \infty$. ∎

Proof of Proposition 2.28 If $\mathcal{X}$ is a Polish space, then by Lemma 2.24 so is $\mathcal{R}(\mathcal{X})$ and hence by Lemma 2.16 so is $\mathcal{K}_+(\mathcal{R}(\mathcal{X}))$. Let us set

$$\mathcal{K}' := \big\{K \in \mathcal{K}_+(\mathcal{R}(\mathcal{X})) : \overline{I}_K \text{ is a closed subinterval of } \overline{\mathbb{R}}\big\}. \qquad (2.17)$$

Then $\mathcal{K}'$ is a closed subset of $\mathcal{K}_+(\mathcal{R}(\mathcal{X}))$ and hence Polish in the induced topology by Proposition 2.4. For each $\varepsilon, \delta > 0$ and $T < \infty$, the set

$$A_{T,\varepsilon,\delta} := \{K \in \mathcal{K}' : m_{T,\delta}(K) \geq \varepsilon\}$$

is a closed subset of $\mathcal{K}'$ and hence its complement $A^{\mathrm{c}}_{T,\varepsilon,\delta}$ is open. By Lemma 2.29,

$$\Pi(\mathcal{X}) = \bigcap_{n,m} \bigcup_k A^{\mathrm{c}}_{n,1/m,1/k},$$

which is a countable intersection of open sets, i.e., a G_δ-set. By Proposition 2.4, it follows that $\Pi(\mathcal{X})$ is a Polish space. ∎

A set $\mathcal{A} \subset \Pi(\mathcal{X})$ is called *equicontinuous* if

$$\lim_{\delta\to 0} \sup_{\pi\in\mathcal{A}} m_{T,\delta}(\pi) = 0 \qquad (T < \infty).$$

The following theorem identifies the compact subsets of $\Pi(\mathcal{X})$. Condition (ii) is called the *compact containment* condition. If $I \subset \mathbb{R}$ is a closed nonempty interval, then $\mathcal{C}_I(\mathcal{X})$ is a closed subset of Π and hence the following theorem can also be used to identify the precompact subsets of $\mathcal{C}_I(\mathcal{X})$. In this context, the result is known as the *Arzela-Ascoli theorem.* Note that while the definition of equicontinuity depends (at least a priori) on the choice of the metric d on $\mathcal{X}$, whether a set $\mathcal{A} \subset \Pi(\mathcal{X})$ is precompact only depends on the topology on $\mathcal{X}$, so when verifying conditions (i) and (ii) below, we are free to choose any metric d that generates the topology on $\mathcal{X}$.

Theorem 2.30 (Arzela-Ascoli) *A set $\mathcal{A} \subset \Pi(\mathcal{X})$ is precompact if and only if*

(i) *$\mathcal{A}$ is equicontinuous,*

(ii) *for each $T < \infty$, there exists a compact set $C \subset \mathcal{X}$ such that $\pi(t) \in C$ for all $\pi \in \mathcal{A}$, $t \in [-T, T]$.*

Proof Let $\mathcal{K}'$ be the space defined in (2.17), equipped with the Hausdorff topology. Let $\overline{\mathcal{A}}$ denote the closure of $\mathcal{A}$, viewed as a subset of the space $\mathcal{K}'$. Then $\mathcal{A}$ is a precompact subset of $\Pi(\mathcal{X})$ if and only if $\overline{\mathcal{A}}$ is a compact subset of $\mathcal{K}'$ and $\overline{\mathcal{A}} \subset \Pi(\mathcal{X})$. By Lemmas 2.17 and 2.25, $\overline{\mathcal{A}}$ is a compact subset of $\mathcal{K}'$ if and only if condition (ii) holds. To complete the proof, it suffices to show that assuming that (ii) holds, one has $\overline{\mathcal{A}} \subset \Pi(\mathcal{X})$ if and only if (i) holds.

We first show that (i) implies $\overline{\mathcal{A}} \subset \Pi(\mathcal{X})$. Assume that $\pi_n \in \mathcal{A}$ converge in the Hausdorff topology to a compact subset $\pi \subset \mathcal{R}(\mathcal{X})$. To show that $\pi \in \Pi(\mathcal{X})$, will apply Lemma 2.29. If $(x, s), (y, t) \in \pi$, then by Lemma 2.15, there exist $(x_n, s_n), (y_n, t_n) \in \pi_n$ such that $(x_n, s_n) \to (x, s)$ and $(y_n, t_n) \to (y, t)$. If $s, t \in [-T, T]$ and $|t - s| \leq \delta$, then for n large enough we have $s_n, t_n \in [-T - 1, T + 1]$ and $|t_n - s_n| \leq 2\delta$. Since $d(x_n, y_n) \to d(x, y)$, it follows that

$$\limsup_{\delta \to 0} m_{T,\delta}(\pi) \leq \limsup_{\delta \to 0} \sup_n m_{T+1,2\delta}(\pi_n) = 0 \qquad (\delta > 0,\ T < \infty),$$

which by Lemma 2.29 implies that $\pi \in \Pi(\mathcal{X})$.

Assume now that (ii) holds but (i) fails. Then there exist $T < \infty$ and $\varepsilon > 0$ such that for each $n \geq 1$, we can find $\pi_n \in \mathcal{A}$ with $m_{T,1/n}(\pi_n) \geq \varepsilon$. This means that there exist $-T \leq s_n \leq t_n \leq T$ such that $d(\pi_n(s_n), \pi_n(t_n)) \geq \varepsilon$ and $t_n - s_n \leq 1/n$. By (ii), $\overline{\mathcal{A}}$ is a compact subset of $\mathcal{K}'$, so by going a subsequence we may assume that $\pi_n \to \pi \in \mathcal{K}'$. By going to a further subsequence, we may assume that $s_n \to s$ and $t_n \to t$ for some $s, t \in [-T, T]$. But then $s = t$ since $t_n - s_n \leq 1/n$. Let $x_n := \pi_n(s_n)$ and $y_n := \pi_n(t_n)$. By (ii), we can select a further subsequence such that $x_n \to x$ and $y_n \to y$ for some x, y with $d(x, y) \geq \varepsilon$. By Lemma 2.15, we have $(x, t), (y, t) \in \pi$ which shows that $\pi \notin \Pi(\mathcal{X})$ and hence $\overline{\mathcal{A}}$ is not a subset of $\Pi(\mathcal{X})$. ∎

2.8 Tightness

In this section, we use the general results from the previous section to derive a tightness criterion for sequences of random variables with values in the space $\Pi(\overline{\mathbb{R}})$.

Lemma 2.31 (Precompactness) *Let $\mathcal{A}$ be a subset of $\Pi(\overline{\mathbb{R}})$. Then $\mathcal{A}$ is precompact if and only if for all $T<\infty$ and $\varepsilon>0$, there exists a $\delta>0$ such that*

$$\begin{aligned}\big|\pi(u)-\pi(t)\big| \leq \varepsilon \text{ for all } \pi\in\mathcal{A} \text{ and } \sigma_\pi\leq t\leq u\\ \text{s.t. } (\pi(t),t),(\pi(u),u)\in[-T,T]^2,\ u-t\leq\delta.\end{aligned}$$

This statement remains true if we drop one of the conditions $(\pi(t),t)\in[-T,T]^2$ or $(\pi(u),u)\in[-T,T]^2$.

Proof Let $\phi:\overline{\mathbb{R}}\to[-1,1]$ be strictly increasing and continuous with $\phi(\pm\infty)=\pm1$. Then

$$d(x,y):=\big|\phi(x)-\phi(y)\big| \qquad (x,y\in\overline{\mathbb{R}}).$$

is a metric generating the topology on $\overline{\mathbb{R}}$. Since $\overline{\mathbb{R}}$ is compact, by the Arzela-Ascoli theorem (Theorem 2.30), $\mathcal{A}$ is precompact if and only if it is equicontinuous, i.e.,

$$\begin{aligned}\sup\big\{d\big(\pi(t),\pi(u)\big):\pi\in\mathcal{A},\ \sigma_\pi\leq t\leq u\leq\tau_\pi,\\ t,u\in[-T,T],\ u-t\leq\delta\big\}\underset{\delta\to0}{\longrightarrow}0 \qquad \forall T<\infty.\end{aligned}$$

In other words, $\mathcal{A}$ is *not* precompact if and only if

$$\begin{aligned}\exists T<\infty \text{ and } \varepsilon>0 \text{ s.t. } \forall\delta>0\ \exists\pi\in\mathcal{A} \text{ and } \sigma_\pi\leq t\leq u\leq\tau_\pi\\ \text{s.t. } t,u\in[-T,T],\ u-t\leq\delta \text{ and } d\big(\pi(t),\pi(u)\big)>\varepsilon.\end{aligned} \tag{2.18}$$

We claim that this is equivalent to

$$\begin{aligned}\exists S,T<\infty \text{ and } \varepsilon>0 \text{ s.t. } \forall\delta>0\ \exists\pi\in\mathcal{A} \text{ and } \sigma_\pi\leq t\leq u\leq\tau_\pi \text{ s.t.}\\ t,u\in[-T,T],\ \pi(t),\pi(u)\in[-S,S],\ u-t\leq\delta \text{ and } d\big(\pi(t),\pi(u)\big)>\varepsilon/3.\end{aligned} \tag{2.19}$$

The implication (2.19)⇒(2.18) is trivial. To prove the converse, assume that (2.18) holds for some $T<\infty$ and $\varepsilon>0$. Making ε smaller if necessary, we can without loss of generality assume that $0<\varepsilon<1$. We can choose the function ϕ in the definition of our metric d on $\overline{\mathbb{R}}$ to be symmetric and then define $S>0$ by $d(\pm S,\pm\infty)=\varepsilon/3$. Now fix $\delta>0$ and let π be as in (2.18). If $\pi(t),\pi(u)\in[-S,S]$ already holds we are done. If $\pi(t)\notin$

$[-S,S]$, then either 1. $\pi(t)\in[-\infty,-S)$ or 2. $\pi(t)\in(S,\infty]$. Assume that we are in case 1. Since $d\big(\pi(t),\pi(u)\big)>\varepsilon$, we must have $\pi(u)\in(-S,\infty]$. Therefore, by continuity, there must be some $t'\in[t,u]$ such that $\pi(t')=-S$. Then $d\big(\pi(t'),\pi(u)\big)>(2/3)\varepsilon$. If $\pi(u)\leq S$ we are done. Otherwise, by continuity, there must be some $u'\in[t',u]$ such that $\pi(u')=S$ and now $d\big(\pi(t'),\pi(u')\big)=d(-S,S)>\varepsilon/3$. Case 2 is similar, by symmetry, and the case that $\pi(t)\in[-S,S]$ but $\pi(u)\notin[-S,S]$ can also be treated in the same way.

If we drop one of the conditions $\pi(t)\in[-S,S]$ or $\pi(u)\in[-S,S]$ from (2.19), then this condition is weaker than (2.19) but stronger than (2.18). However, we have already shown that (2.18) implies (2.19), so all conditions are equivalent. In other words, $\mathcal{A}$ is precompact if and only if for all $S,T<\infty$ and $\varepsilon>0$, there exists a $\delta>0$ such that

$$\begin{aligned}&d\big(\pi(t),\pi(u)\big)\leq\varepsilon/3\text{ for all }\pi\in\mathcal{A}\text{ and }\sigma_\pi\leq t\leq u\text{ s.t.}\\&\qquad\pi(t),\pi(u)\in[-S,S],\ t,u\in[-T,T],\ u-t\leq\delta,\end{aligned}$$

and the same is true if we drop one of the conditions $\pi(t)\in[-S,S]$ or $\pi(u)\in[-S,S]$. Replacing S and T by $S\vee T+\delta$ if necessary, we can simplify this by saying that $\mathcal{A}$ is precompact if and only if for all $T<\infty$ and $\varepsilon>0$, there exists a $\delta>0$ such that

$$\begin{aligned}&d\big(\pi(t),\pi(u)\big)\leq\varepsilon\text{ for all }\pi\in\mathcal{A}\text{ and }\sigma_\pi\leq t\leq u\text{ s.t.}\\&\qquad\big(\pi(t),t\big),\big(\pi(u),u\big)\in[-T,T]^2,\ u-t\leq\delta,\end{aligned}$$

and the same is true if we drop one of the conditions $\big(\pi(t),t\big)\in[-T,T]^2$ or $\big(\pi(u),u\big)\in[-T,T]^2$. We can choose the function ϕ that we used to define the metric d on $\overline{\mathbb{R}}$ to be Lipschitz continuous with Lipschitz constant one; then d has the property that $d(x,y)\leq|x-y|$ for all $x,y\in\overline{\mathbb{R}}$. Conversely, as long as at least one of $\pi(t)$ and $\pi(u)$ lies inside $[-T,T]$, by making $d\big(\pi(t),\pi(u)\big)$ as small as we wish, we can also make $\big|\pi(t)-\pi(u)\big|$ as small as we wish. From these observations, the claim of the lemma follows. ∎

Proposition 2.32 (Almost sure precompactness) *Let $\mathcal{A}$ be a random subset of $\Pi(\overline{\mathbb{R}})$. Then $\mathcal{A}$ is almost surely a precompact subset of $\Pi(\overline{\mathbb{R}})$ if and only if*

$$\begin{aligned}&\mathbb{P}\big[\,\big|\pi(u)-\pi(t)\big|\geq\varepsilon\text{ for some }\pi\in\mathcal{A}\text{ and }\sigma_\pi\leq t\leq u\leq\tau_\pi\\&\qquad\text{s.t. }(\pi(t),t),(\pi(u),u)\in[-T,T]^2,\ u-t\leq\delta\big]\underset{\delta\to0}{\longrightarrow}0\quad\forall T<\infty,\ \varepsilon>0.\end{aligned}$$

The same is true if we drop one of the conditions $\big(\pi(t),t\big)\in[-T,T]^2$ or $\big(\pi(u),u\big)\in[-T,T]^2$.

Proof We only prove the claim with both conditions $\big(\pi(t),t\big) \in [-T,T]^2$ and $\big(\pi(u),u\big) \in [-T,T]^2$ in place. If we drop one of these conditions, then the argument goes precisely in the same way. Let $A^\delta_{T,\varepsilon}$ denote the event that

$$\begin{array}{c}\big|\pi(u)-\pi(t)\big| \geq \varepsilon \text{ for some } \pi\in\mathcal{A} \text{ and } \sigma_\pi \leq t \leq u \leq \tau_\pi \\ \text{s.t. } (\pi(t),t),(\pi(u),u) \in [-T,T]^2,\ u-t\leq\delta.\end{array}$$

Then $\delta \leq \delta'$ implies $A^\delta_{T,\varepsilon} \subset A^{\delta'}_{T,\varepsilon}$ and $A_{T,\varepsilon} := \bigcap_{\delta>0} A^\delta_{T,\varepsilon}$ is the event that

$$\begin{array}{c}\forall\delta>0\ \exists\pi\in\mathcal{A} \text{ and } \sigma_\pi\leq t\leq u\leq\tau_\pi \\ \text{s.t. } (\pi(t),t),(\pi(u),u)\in[-T,T]^2,\ u-t\leq\delta, \text{ and } \big|\pi(u)-\pi(t)\big|\geq\varepsilon.\end{array}$$

The assumption of the proposition says that $\lim_{\delta\to 0} P(A^\delta_{T,\varepsilon}) = 0$, which implies $P(A_{T,\varepsilon}) = 0$. Since this holds for all $T<\infty$ and $\varepsilon>0$, it follows that

$$\mathbb{P}\big(\bigcup_{n\geq 1}\bigcup_{m\geq 1} A_{n,1/m}\big) = 0,$$

which shows that almost surely, for all $n\geq 1$ and $m\geq 1$, there exists a $\delta>0$ such that

$$\begin{array}{c}\big|\pi(u)-\pi(t)\big| < 1/m \text{ for all } \pi\in\mathcal{A} \text{ and } \sigma_\pi\leq t\leq u\leq\tau_\pi \\ \text{s.t. } (\pi(t),t),(\pi(u),u)\in[-n,n]^2,\ u-t\leq\delta.\end{array}$$

By Lemma 2.31, it follows that $\mathcal{A}$ is almost surely precompact.

On the other hand, if the assumption of the proposition does not hold, then the event $A_{T,\varepsilon}$ has positive probability for some $T<\infty$ and $\varepsilon>0$, which by Lemma 2.31 implies that $\mathcal{A}$ is with positive probability not precompact. ∎

Proposition 2.33 (Tightness of random compact sets of paths) *Let $\mathcal{K}_+(\Pi(\overline{\mathbb{R}}))$ be the set of nonempty compact subsets of $\Pi(\overline{\mathbb{R}})$, equipped with the Hausdorff topology. Let $(\mathcal{A}_n)_{n\geq 1}$ be a sequence of random variables with values in $\mathcal{K}_+(\Pi(\overline{\mathbb{R}}))$. Then the probability laws $\big(\mathbb{P}[\mathcal{A}_n \in \cdot\,]\big)_{n\geq 1}$ are tight if and only if*

$$\begin{array}{c}\sup_{n\geq 1}\mathbb{P}\big[\,\big|\pi(u)-\pi(t)\big|\geq\varepsilon \textit{ for some } \pi\in\mathcal{A}_n \textit{ and } \sigma_\pi\leq t\leq u\leq\tau_\pi \\ \textit{s.t. } (\pi(t),t),(\pi(u),u)\in[-T,T]^2,\ u-t\leq\delta\big] \underset{\delta\to 0}{\longrightarrow} 0\end{array}$$

for all $T<\infty$ and $\varepsilon>0$. The same is true if we drop one of the conditions $\big(\pi(t),t\big)\in[-T,T]^2$ or $\big(\pi(u),u\big)\in[-T,T]^2$.

Proof We only prove the claim with both conditions $\big(\pi(t),t\big)\in[-T,T]^2$ and $\big(\pi(u),u\big)\in[-T,T]^2$ in place. If we drop one of these conditions, then the argument goes precisely in the same way. By Theorem 2.12, the probability laws $\big(\mathbb{P}[\mathcal{A}_n\in\,\cdot\,]\big)_{n\geq 1}$ are tight if and only if for each $\eta>0$, there exists a compact set $C\subset\mathcal{K}_+(\Pi(\overline{\mathbb{R}}))$ such that

$$\inf_{n\geq 1}\mathbb{P}[\mathcal{A}_n\in C]\geq 1-\eta.$$

Equivalently, we may show that there exists a precompact set $C\subset\mathcal{K}_+(\Pi(\overline{\mathbb{R}}))$ with this property, because its closure $\overline{C}$ is then compact with $\mathbb{P}[\mathcal{A}_n\in\overline{C}]\geq\mathbb{P}[\mathcal{A}_n\in C]$. By Lemma 2.17, a subset $C\subset\mathcal{K}_+(\Pi(\overline{\mathbb{R}}))$ is precompact if and only there exists a compact $\mathcal{C}\subset\Pi(\overline{\mathbb{R}})$ such that $\mathcal{A}\subset\mathcal{C}$ for all $\mathcal{A}\in C$. It follows that the probability laws $\big(\mathbb{P}[\mathcal{A}_n\in\,\cdot\,]\big)_{n\geq 1}$ are tight if and only if[2]

$$\forall\eta>0\ \exists\text{ compact }\mathcal{C}\subset\Pi(\overline{\mathbb{R}})\text{ s.t. }\inf_{n\geq 1}\mathbb{P}[\mathcal{A}_n\subset\mathcal{C}]\geq 1-\eta. \tag{2.20}$$

Let

$$\begin{aligned}\mathcal{A}^\delta_{T,\varepsilon}:=\big\{\pi\in\Pi(\overline{\mathbb{R}}):\big|\pi(u)-\pi(t)\big|<\varepsilon\ \forall\sigma_\pi\leq t\leq\tau_\pi\\ \text{s.t. }\big(\pi(t),t\big),\big(\pi(u),u\big)\in[-T,T]^2,\ u-t\leq\delta\big\}.\end{aligned}$$

The assumption of the proposition then says that

$$\inf_{n\geq 1}\mathbb{P}\big[\mathcal{A}_n\subset\mathcal{A}^\delta_{T,\varepsilon}\big]\underset{\delta\to 0}{\longrightarrow}1\qquad(T<\infty,\ \varepsilon>0).$$

Let $(\eta_{k,m})_{k,m\geq 1}$ be positive constants. Then we can choose $\delta(k,m)>0$ such that

$$\inf_{n\geq 1}\mathbb{P}\big[\mathcal{A}_n\subset\mathcal{A}^{\delta(k,m)}_{k,1/m}\big]\geq 1-\eta_{k,m}\qquad(k,m\geq 1).$$

Then

$$\inf_{n\geq 1}\mathbb{P}\big[\mathcal{A}_n\subset\bigcap_{k,m\geq 1}\mathcal{A}^{\delta(k,m)}_{k,1/m}\big]\geq 1-\sum_{k,m\geq 1}\eta_{k,m}.$$

By Lemma 2.31, the set $\bigcap_{k,m\geq 1}\mathcal{A}^{\delta(k,m)}_{k,1/m}$ is precompact. Since the positive constants $(\eta_{k,m})_{k,m\geq 1}$ are arbitrary, we can make $\sum_{k,m\geq 1}\eta_{k,m}$ as small as we wish. Taking for $\mathcal{C}$ the closure of $\bigcap_{k,m\geq 1}\mathcal{A}^{\delta(k,m)}_{k,1/m}$, this proves (2.20) and shows that the assumption of the proposition implies tightness of the laws $\big(\mathbb{P}[\mathcal{A}_n\in\,\cdot\,]\big)_{n\geq 1}$.

[2] Indeed, the existence of such a $\mathcal{C}$ is necessary by our previous condition and Lemma 2.17, and conversely, if such a $\mathcal{C}$ exists, then by Lemma 2.17 $C:=\{\mathcal{A}:\mathcal{A}\subset\mathcal{C}\}$ is compact so we can apply our previous condition.

Assume, conversely, that these laws are tight and hence (2.20) holds. Fix $T < \infty$ and $\varepsilon > 0$. By (2.20), for each $\eta > 0$, there exists a compact $\mathcal{C}$ such that

$$\sup_{n\geq 1} \mathbb{P}[\mathcal{A}_n \not\subset \mathcal{C}] \leq \eta.$$

Since $\mathcal{C}$ is compact, by Lemma 2.31, there exists a $\delta > 0$ such that $\mathcal{C} \subset \mathcal{A}^{\delta}_{T,\varepsilon}$. This shows that for each $\eta > 0$, there exists a $\delta > 0$ such that

$$\sup_{n\geq 1} \mathbb{P}\big[\mathcal{A}_n \not\subset \mathcal{A}^{\delta}_{T,\varepsilon}\big] \leq \eta,$$

which implies that the assumption of the proposition is satisfied. ∎

2.9 Cadlag paths

Paths, as we have defined them in Section 2.7, correspond to continuous functions defined on a closed time interval and taking values in a metrisable space $\mathcal{X}$. For some purposes, it will occasionally be necessary to generalise the concept of a path so that paths can make jumps. We recall that a function, defined on a closed real interval and taking values in a metrisable space $\mathcal{X}$, is called *cadlag* (from the French "continue à droit, limite à gauche") if it is right-continuous and has left limits at each time. It is possible to define a space of cadlag paths that extends the space of continuous paths defined in Section 2.7. This has been done in detail in [NS22]. Since the technicalities are somewhat involved, we will only cite some of the results of that paper without proof. Our set-up differs slightly from the set-up in the main body of [NS22], but the results we cite below can be translated into our set-up using [NS22, Lemma 5.3].

Let $\mathcal{X}$ be a metrisable space and let $\preceq$ be a partial order on $\mathcal{X}$. By definition, the partial order $\preceq$ is *compatible with the topology* if the set

$$\mathcal{X}^{\langle 2\rangle} := \big\{(x,y) \in \mathcal{X}^2 : x \preceq y\big\}$$

is a closed subset of $\mathcal{X}^2$, equipped with the product topology. Equivalently, this says that if $x_n \to x$, $y_n \to y$, and $x_n \preceq y_n$ for all n, then it should always be true that $x \preceq y$.

By definition, a *cadlag path* in $\mathcal{X}$ is a nonempty compact subset $\pi \subset \mathcal{R}(\mathcal{X})$ that is equipped with a relation $\preceq$ such that:

(i) for each $t \in \mathbb{R}$, the set $\big\{x \in \mathcal{X} : (x,t) \in \pi\big\}$ has at most two elements,

(ii) $\overline{I}_\pi := \big\{t \in \overline{\mathbb{R}} : \exists x \in \mathcal{X} \cup \{*\} \text{ s.t. } (x,t) \in \pi\big\}$ is a closed subset of $\overline{\mathbb{R}}$,

(iii) $\preceq$ is a total order on π that is compatible with the topology on π and $(x,s) \preceq (y,t)$ for all $(x,s),(y,t) \in \pi$ with $s < t$.

We call $\sigma_\pi := \inf \overline{I}_\pi$ and $\tau_\pi := \sup \overline{I}_\pi$ the *starting time* and *final time* of π, and we set $I_\pi := \overline{I}_\pi \cap \mathbb{R}$. For each $t \in \overline{I}_\pi$, we define $\pi(t-), \pi(t+) \in \mathcal{X} \cup \{*\}$ by

$$\big\{\pi(t-), \pi(t+)\big\} = \big\{x \in \mathcal{X} : (x,t) \in \pi\big\} \quad \text{with} \quad \pi(t-) \preceq \pi(t+).$$

The total order $\preceq$ on a cadlag path π corresponds to the time order and helps us, when the set $\big\{x \in \mathcal{X} : (x,t) \in \pi\big\}$ contains two elements, to find out which element is the "first" and which the "second" in the time order. It is shown in [NS22, Lemma 3.2] that cadlag paths, as we have just defined them, correspond to cadlag functions in the usual sense of the word, in the sense that the function $I_\pi \ni t \mapsto \pi(t+)$ is right-continuous and the function $I_\pi \ni t \mapsto \pi(t-)$ satisfies

$$\pi(t-) = \lim_{s \uparrow t} \pi(t+) \qquad (t \in I_\pi,\ \sigma_\pi < t).$$

The limit from the left is not defined at $t = \sigma_\pi$, however, and contrary to the usual conventions for cadlag functions, it may happen that $\pi(\sigma_\pi -) \neq \pi(\sigma_\pi +)$. Likewise, it is possible that a cadlag path π makes a jump at its final time τ_π.

Recall the definition of a correspondence from Section 2.5. We write $z_1 \prec z_2$ as a shorthand for $z_1 \preceq z_2$ and $z_1 \neq z_2$, and say that a correspondence R between two cadlag paths π, π' is *monotone* if

there are no $(z_1, z_1'), (z_2, z_2') \in R$ such that $z_1 \prec z_2$ in π and $z_2' \prec z_1'$ in π'.

We let $\mathrm{Corr}_+(\pi_1, \pi_2)$ denote the set of all monotone correspondences between two cadlag paths π_1 and π_2. We denote the space of cadlag paths in $\mathcal{X}$ by $\Pi_{\mathrm{S}}(\mathcal{X})$ and in analogy with (2.4), we define a metric d_{S} on $\Pi_{\mathrm{S}}(\mathcal{X})$ by

$$d_{\mathrm{S}}(\pi, \pi') := \inf_{R \in \mathrm{Corr}_+(\pi,\pi')} \sup_{(z,z') \in R} d_{\mathrm{sqz}}(z, z') \qquad \big(\pi, \pi' \in \Pi_{\mathrm{S}}(\mathcal{X})\big), \tag{2.21}$$

where d_{sqz} is the metric on the squeezed space $\mathcal{R}(\mathcal{X})$. It is shown in [NS22, Prop 3.3] that if $\mathcal{X}$ is a Polish space, then so is $\Pi_{\mathrm{S}}(\mathcal{X})$, equipped with the topology generated by d_{S}. By [NS22, Lemma 3.5], the space of (continuous) paths $\Pi(\mathcal{X})$ is a closed subset of $\Pi_{\mathrm{S}}(\mathcal{X})$, and by [NS22, Prop 3.4], the topology on $\Pi(\mathcal{X})$ coincides with the induced topology from its embedding in $\Pi_{\mathrm{S}}(\mathcal{X})$. It is shown in [NS22, Section 3.4] for sequences of paths that are all defined on the same time interval and that do not jump at the endpoints of this interval, convergence in the topology on $\Pi_{\mathrm{S}}(\mathcal{X})$ corresponds to correspondence in the classical Skorohod topology.

For cadlag paths, there is an analogue of the Arzela-Ascoli theorem (Theorem 2.30). The *Skorohod modulus of continuity* of a cadlag path π is defined as

$$\begin{aligned} m^{\rm S}_{T,\delta}(\pi) := \sup\big\{\, d\big(x_2,\{x_1,x_3\}\big) : (x_i,t_i)\in\pi,\ t_i\in[-T,T]\ (i=1,2,3),\\ (x_1,t_1)\preceq(x_2,t_2)\preceq(x_3,t_3),\ t_3-t_1\leq\delta\big\}. \end{aligned} \tag{2.22}$$

We say that a set $\mathcal{A}\subset\Pi_{\rm S}(\mathcal{X})$ is *Skorohod-equicontinuous* if

$$\lim_{\delta\to 0}\sup_{\pi\in\mathcal{A}} m^{\rm S}_{T,\delta}(\pi)=0 \qquad (T<\infty).$$

The following theorem, that we cite from [NS22, Thm 3.7], is very similar to Theorem 2.30.

Theorem 2.34 (Compactness criterion) *A set $\mathcal{A}\subset\Pi_{\rm S}(\mathcal{X})$ is precompact if and only if*

(i) *$\mathcal{A}$ is Skorohod-equicontinuous,*

(ii) *for each $T<\infty$, there exists a compact set $C\subset\mathcal{X}$ such that $\pi(t)\in C$ for all $\pi\in\mathcal{A}$, $t\in[-T,T]$.*

Chapter 3

The Brownian web

3.1 Arrow configurations

In this chapter, we are interested in the diffusive scaling limit of the (unperturbed, standard) voter model, and its dual system of coalescing random walks. We will focus on the systems coalescing random walks, and more precisely on the collection of all open paths in their graphical representation. Thus, we will be interested in collections of coalescing random walks, starting from each point in space-time, and their diffusive scaling limit, which can informally be described as coalescing Brownian motions, start starting from each point in space-time.

In the context of the voter model, it is natural to consider coalescing random walks in discrete space and continuous time. It is sometimes more convenient to consider coalescing random walks in discrete space and time. We will therefore start by studying the latter, and prove that they have a diffusive scaling limit. In Section 3.7 below, we will indicate how the arguments can be adapted to the continuous-time setting of the voter model.

By definition, we call

$$\mathbb{Z}^2_{\rm even} := \big\{(x,t) \in \mathbb{Z}^2 : x+t \text{ is even}\big\}$$

the *even sublattice* of $\mathbb{Z}^2$. Let $\omega = (\omega_z)_{z\in\mathbb{Z}^2_{\rm even}}$ be an i.i.d. collection of random variables that are uniformly distributed on $\{-1,+1\}$. We can use ω to define a random directed graph with vertex set $\mathbb{Z}^2_{\rm even}$ and set of oriented edges

$$\vec{E} := \big\{\big((x,t),(x+\omega_{(x,t)},t+1)\big) : (x,t) \in \mathbb{Z}^2_{\rm even}\big\}.$$

We call the random directed graph $(\mathbb{Z}^2_{\rm even},\vec{E})$ an *arrow configuration*. See Figure 3.1 for a picture.

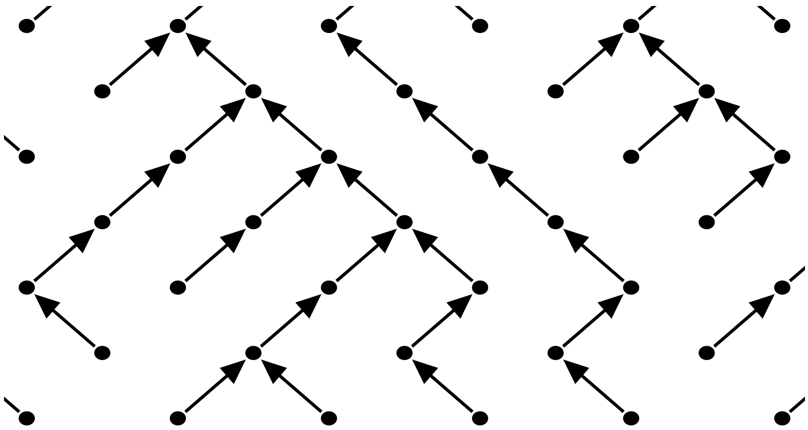

Figure 3.1: An arrow configuration.

In Section 2.7, for any metrisable space $\mathcal{X}$, we gave a definition of the path space $\Pi(\mathcal{X})$. Recall that I_π denotes the domain of a path $\pi \in \Pi(\mathcal{X})$ and that σ_π, τ_π denote its starting time and final time, respectively. We will especially be interested in the case that the metrisable space $\mathcal{X}$ is $\overline{\mathbb{R}} := [-\infty, \infty]$, the extended real line. We let

$$\Pi^\uparrow := \big\{\pi \in \Pi(\overline{\mathbb{R}}) : \tau_\pi = \infty\big\}.$$

We call $\Pi^\uparrow$ the space of all *upward paths*. In view of Lemma 2.26, elements of $\Pi^\uparrow$ correspond to continuous functions $\pi : I_\pi \to \overline{\mathbb{R}}$, where I_π is an interval of the form $[\sigma_\pi, \infty)$ if the starting time σ_π is finite, and

$$I_\pi = \mathbb{R} \text{ if } \sigma_\pi = -\infty \quad \text{and} \quad I_\pi = \emptyset \text{ if } \sigma_\pi = +\infty.$$

We will call the point

$$z_\pi := \big(\pi(\sigma_\pi), \sigma_\pi\big)$$

the *starting point* of the path π. Note that in general z_π is an element of $\mathcal{R}(\overline{\mathbb{R}})$, the squeezed space defined in Section 2.6. By definition, a *open path in the arrow configuration* $(\mathbb{Z}^2_{\rm even}, \vec{E})$, or simply a *open path in* ω, is a path $\pi \in \Pi^\uparrow$ with the following properties:

(i) $\sigma_\pi \in \mathbb{Z} \cup \{-\infty, +\infty\}$ and $\big(\pi(t), t\big) \in \mathbb{Z}^2_{\rm even}$ $(t \in \mathbb{Z},\ t \geq \sigma_\pi)$,

(ii) $\pi(t+1) = \pi(t) + \omega_{(\pi(t),t)}$ $(t \in \mathbb{Z},\ t \geq \sigma_\pi)$,

(iii) $\pi(t+s) = (1-s)\pi(t) + s\pi(t+1)$ $(0 \leq s \leq 1,\ t \in \mathbb{Z},\ t \geq \sigma_\pi)$.

In words, these are upward paths that visit points in the even sublattice at integer times and follow the arrows, with linear interpolation between integer times. We let

$$\mathcal{U} = \mathcal{U}(\omega) := \big\{\pi \in \Pi^{\uparrow} : \pi \text{ is a open path in } \omega\big\}. \tag{3.1}$$

We let $\overline{\mathcal{U}}$ denote the closure of $\mathcal{U}$ in the topology on $\Pi^{\uparrow}$. The following proposition says that $\overline{\mathcal{U}}$ is a.s. compact and compared to $\mathcal{U}$ only contains a few extra trivial paths. Below, we use the notation $\overline{\mathbb{Z}} := \mathbb{Z} \cup \{-\infty, \infty\}$, i.e., this is the closure of $\mathbb{Z}$ in $\overline{\mathbb{R}}$.

Proposition 3.1 (Compact set of paths) *The closure $\overline{\mathcal{U}}$ of the random set of upward paths $\mathcal{U}$ defined in (3.1) is almost surely a compact subset of $\Pi^{\uparrow}$. Moreover, almost surely, the set $\overline{\mathcal{U}} \backslash \mathcal{U}$ consists of all paths $\pi \in \Pi^{\uparrow}$ with $\sigma_\pi \in \overline{\mathbb{Z}}$ and either $\pi(t) = -\infty$ for all $t \in I_\pi$ or $\pi(t) = +\infty$ for all $t \in I_\pi$.*

Proof Since paths in $\mathcal{U}$ are Lipschitz continuous with Lipschitz constant one, equicontinuity is obvious so $\mathcal{U}$ is precompact by Proposition 2.32 and hence $\overline{\mathcal{U}}$ is compact.

Let $s \in \mathbb{Z}$ and let $\pi \in \Pi^{\uparrow}$ be defined by $\sigma_\pi := s$ and $\pi(t) := -\infty$ for all $\sigma_\pi \leq t < \infty$. To see that $\pi \in \overline{\mathcal{U}}$, choose $x_n \in \mathbb{Z}$ such that $(x_n, s) \in \mathbb{Z}^2_{\rm even}$ and $x_n \to -\infty$. Let $\pi_n \in \mathcal{U}$ be the unique open path started at (x_n, t). Since $\overline{\mathcal{U}}$ is compact, by going to a subsequence if necessary, we can assume that $\pi_n \to \pi'$ for some $\pi' \in \Pi^{\uparrow}$. Since π_n is a random walk starting from (x_n, t) and $x_n \to -\infty$, the law of $\pi_n(t)$ converges weakly to the delta measure on $-\infty$ for each $t \geq s$, from which we conclude that $\pi' = \pi$ and hence $\pi \in \overline{\mathcal{U}}$. In the same way, we see that $\overline{\mathcal{U}}$ contains all trivial paths π with $\sigma_\pi \in \mathbb{Z}$ and $\pi(t) = \infty$ for all $\sigma_\pi \leq t < \infty$. Since $\overline{\mathcal{U}}$ is closed, it also contains all limits of such paths, so letting $\sigma_\pi \to \infty$ or $\sigma_\pi \to -\infty$ we see that $\overline{\mathcal{U}}$ also contains all trivial paths with $\sigma_\pi = -\infty$ and either $\pi(t) = -\infty$ for all $t \in \mathbb{R}$ or $\pi(t) = +\infty$ for all $t \in \mathbb{R}$, as well as the trivial path with $\sigma_\pi = +\infty$.

To complete the proof, we must show that if $\pi \in \overline{\mathcal{U}}$ satisfies $\pi(t) \in \mathbb{R}$ for some $t \geq \sigma_\pi$, then $\pi(t) \in \mathbb{R}$ for all $t \geq \sigma_\pi$. We first note that paths in $\mathcal{U}$ are *noncrossing* in the sense that there do not exist $\pi, \pi' \in \mathcal{U}$ and $\sigma_\pi \vee \sigma_{\pi'} \leq s < t < \infty$ such that $\pi(s) < \pi'(s)$ while $\pi'(t) < \pi(t)$. It is easy to see that this property is preserved in the limit so paths in $\overline{\mathcal{U}}$ are noncrossing too. Now assume that $\pi \in \overline{\mathcal{U}}$ satisfies $\pi(t) \in \mathbb{R}$ for some $t \geq \sigma_\pi$. Choose $z_n = (x_n, s_n) \in \mathbb{Z}^2_{\rm even}$ with $s_n < \sigma_\pi$ such that $z_n \to (\infty, s)$ for some $s \in \mathbb{R}$, and let $\pi_n \in \mathcal{U}$ denote the open path started from z_n. Then π_n is a random walk started from z_n. By our previous arguments, $\pi_n(t) \to \infty$ a.s. so $\pi(t) < \pi_n(t)$ for all n large enough. Since paths in $\overline{\mathcal{U}}$ are noncrossing it follows that there exists an n such that $\pi(t) \leq \pi_n(t) < \infty$ for all $t \geq \sigma_\pi$. In the same way, by symmetry, we see that $-\infty < \pi(t)$ for all $t \geq \sigma_\pi$. ∎

We now turn to what we are mainly interested in, which is the diffusive scaling limit of arrow configurations. For each $\varepsilon > 0$, we define a *diffusive scaling* map $\theta_\varepsilon : \mathbb{R}^2 \to \mathbb{R}^2$ by

$$\theta_\varepsilon(x, t) := (\varepsilon x, \varepsilon^2 t) \qquad \big((x, t) \in \mathbb{R}^2\big). \tag{3.2}$$

Let $\mathcal{R}(\overline{\mathbb{R}})$ be the squeezed space defined in Section 2.6. We extend θ_ε continuously to $\mathcal{R}(\overline{\mathbb{R}})$ in the obvious way, by setting

$$\theta_\varepsilon(\pm\infty, t) := (\pm\infty, \varepsilon^2 t) \quad (t \in \mathbb{R}) \quad \text{and} \quad \theta_\varepsilon(*, \pm\infty) := (*, \pm\infty).$$

For any subset $A \subset \mathcal{R}(\overline{\mathbb{R}})$, we let

$$\theta_\varepsilon(A) := \big\{\theta_\varepsilon(z) : z \in A\big\}$$

denote the image of A under θ_ε. In particular, this notation applies to paths $\pi \in \Pi(\overline{\mathbb{R}})$, which according to their defininition in Section 2.7 correspond to compact subsets of $\mathcal{R}(\overline{\mathbb{R}})$. It is easy to see that $\theta_\varepsilon(\pi) \in \Pi^\uparrow$ for all $\pi \in \Pi^\uparrow$, so the diffusive scaling map $\theta_\varepsilon : \mathcal{R}(\overline{\mathbb{R}}) \to \mathcal{R}(\overline{\mathbb{R}})$ naturally gives rise to a diffusive scaling map from $\Pi^\uparrow$ to $\Pi^\uparrow$ which by a slight abuse of notation we also denote by θ_ε. Going one step further, for any subset $\mathcal{A} \subset \Pi^\uparrow$, we let

$$\theta_\varepsilon(\mathcal{A}) := \big\{\theta_\varepsilon(\pi) : \pi \in \mathcal{A}\big\}$$

denote the image of $\mathcal{A}$ under this map.

In Section 2.5, we equipped the space $\mathcal{K}(\mathcal{X})$ of all compact subsets of a metrisable topological space $\mathcal{X}$ with the Hausdorff topology. As an immediate consequence of Lemma 2.21, we obtain:

Lemma 3.2 (Scaling of paths) *For each $\varepsilon > 0$, the map $\theta_\varepsilon : \Pi^\uparrow \to \Pi^\uparrow$ is continuous.*

Proof Immediate from Lemma 2.21, the continuity of the map $\theta_\varepsilon : \mathcal{R}(\overline{\mathbb{R}}) \to \mathcal{R}(\overline{\mathbb{R}})$, and the fact that in Section 2.7 we viewed the path space $\Pi(\overline{\mathbb{R}})$ as a subset of $\mathcal{K}(\mathcal{R}(\overline{\mathbb{R}}))$ and equipped it with the induced topology from this embedding. ∎

Let $\mathcal{U}$ be the set of all open paths in an arrow configuration and let $\overline{\mathcal{U}}$ be its closure, which by Proposition 3.1 is a random compact subset of $\Pi^\uparrow$. Then, since the continuous image of a compact set is compact, by Lemma 3.2, for each $\varepsilon > 0$, the diffusively rescaled set of paths $\theta_\varepsilon(\overline{\mathcal{U}})$ is a random compact subset of $\Pi^\uparrow$. Our aim is to prove that

$$\mathbb{P}\big[\theta_\varepsilon(\overline{\mathcal{U}}) \in \cdot\,\big] \underset{\varepsilon\to 0}{\Longrightarrow} \mathbb{P}\big[\mathcal{W} \in \cdot\,\big] \tag{3.3}$$

where $\Rightarrow$ denotes weak convergence of probability laws on the space $\mathcal{K}(\Pi^\uparrow)$, equipped with the Hausdorff topology, and $\mathcal{W}$ is a random compact subset of $\Pi^\uparrow$ that will be called the *Brownian web*.

3.2 Coalescing Brownian motions

As a first step towards proving (3.3), we start by proving something like convergence of finite dimensional distributions. More precisely, for each $\varepsilon > 0$, we choose finitely many points $z^\varepsilon_1, \ldots, z^\varepsilon_n$ in the diffusively rescaled lattice $\theta_\varepsilon(\mathbb{Z}^2_{\rm even})$, in such a way that

$$(z^\varepsilon_1, \ldots, z^\varepsilon_n) \underset{\varepsilon\to 0}{\longrightarrow} (z_1, \ldots, z_n)$$

for some $z_1, \ldots, z_n \in \mathbb{R}^2$. Letting $\pi^\varepsilon_1, \ldots, \pi^\varepsilon_n$ denote the paths in $\mathcal{U}$ with starting points $z^\varepsilon_1, \ldots, z^\varepsilon_n$, we will argue that $(\pi^\varepsilon_1, \ldots, \pi^\varepsilon_n)$ converges in distribution to a collection of coalescing Brownian motions.

Let $B^1 = (B^1_t)_{t\geq 0}$ and $B^2 = (B^2_t)_{t\geq 0}$ be two independent standard one-dimensional Brownian motions started from initial states $B^i_0 = x_i$ $(i = 1, 2)$, and let

$$\tau := \inf\{t \geq 0 : B^1_t = B^2_t\},$$

which is a.s. finite since $(B^1_t - B^2_t)_{t\geq 0}$ is a Brownian motion (with double the quadratic variation of a standard Brownian motion), and one-dimensional Brownian motion is point recurrent. Let $\tilde{B}^2 = (\tilde{B}^2_t)_{t\geq 0}$ be defined by

$$\tilde{B}^2_t := \begin{cases} B^2_t & \text{if } t \leq \tau, \\ B^1_t & \text{if } \tau \leq t. \end{cases}$$

Then it is possible to check[1] that $\tilde{B}^2$ is a standard Brownian motion, and our definition is symmetric in the sense that if we define $\tilde{B}^1_t := B^1_t$ $(t \leq \tau)$ and $:= B^2_t$ $(\tau \leq t)$, then $(\tilde{B}^1_t, B^2_t)_{t\geq 0}$ is equally distributed with $(B^1_t, \tilde{B}^2_t)_{t\geq 0}$. The processes B^1 and $\tilde{B}^2$ are of course not independent. The process $(B^1_t, \tilde{B}^2_t)_{t\geq 0}$ is a Markov process that is known as *coalescing Brownian motions*.

We can carry out the same construction for any finite number of Brownian motions, that can moreover start at different times. See Figure 3.2 for an

[1] In fact, one way to prove this is to derive it from Proposition 3.3 below.

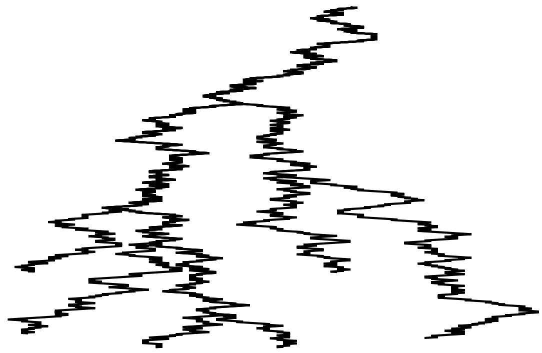

Figure 3.2: Coalescing Brownian motions.

illustration. Let $z_1, \ldots, z_n \in \mathbb{R}^2$ with $z_i = (x_i, s_i)$ $(i = 1, \ldots, n)$, and let $B^1, \ldots, B^n$ be independent Brownian motions such that $B^i = (B^i_t)_{t \geq s_i}$ starts at time s_i in $B^i_{s_i} = x_i$. We set $\tau_1 := \infty$, $A_1 := \{(B^1_t, t) : s_1 \leq t < \infty\}$ and define inductively for $j = 2, \ldots, n$

$$\begin{aligned} \tau_j &:= \inf \big\{ t \geq s_j : (B^j_t, t) \in A_1 \cup \cdots \cup A_{j-1} \big\}, \\ A_j &:= \big\{ (B^j_t, t) : s_j \leq t < \tau_j \big\}. \end{aligned}$$

By the recurrence of one-dimensional Brownian motion, almost surely $\tau_j < \infty$ for all $2 \leq j \leq n$. Note that the sets $A_1, \ldots, A_n$ are disjoint. In view of this, we can uniquely define $\iota(j) \in \{1, \ldots, j-1\}$ by the requirement that

$$(B^j_{\tau_j}, \tau_j) \in A_{\iota(j)}.$$

Using this, we define inductively $\tilde{B}^1 := B^1$ and

$$\tilde{B}^j_t := \begin{cases} B^j_t & \text{if } s_i \leq t \leq \tau_j, \\ \tilde{B}^{\iota(j)}_t & \text{if } \tau_j \leq t. \end{cases}$$

We call $\tilde{B}^1, \ldots, \tilde{B}^n$ *coalescing Brownian motions* starting from the space-time points $z_1, \ldots, z_n \in \mathbb{R}^2$.

We are now ready to formulate a result about the convergence in law of finitely many open paths in an arrow configuration. We have already become used (hopefully!) to the slight abuse of notation by which θ_ε can denote both a diffusive scaling map acting on space-time points, or on sets of space-time points such as paths, or even sets of paths. Taking this one step further, we also denote

$$\theta_\varepsilon(z_1, \ldots, z_n) := \big(\theta_\varepsilon(z_1), \ldots, \theta_\varepsilon(z_n)\big), \quad \theta_\varepsilon(\pi_1, \ldots, \pi_n) := \big(\theta_\varepsilon(\pi_1), \ldots, \theta_\varepsilon(\pi_n)\big)$$

when $z_1, \ldots, z_n$ are space-time points and $\pi_1, \ldots, \pi_n$ are paths.

Proposition 3.3 (Convergence of finite dimensional distributions) *Let $\varepsilon_k > 0$ satisfy $\varepsilon_k \to 0$. Fix $n \geq 1$ and for each k, let $z_1^k, \ldots, z_n^k \in \mathbb{Z}^2_{\rm even}$. Assume that*

$$\theta_{\varepsilon_k}(z_1^k, \ldots, z_n^k) \underset{k\to\infty}{\longrightarrow} (z_1, \ldots, z_n) \in (\mathbb{R}^2)^n.$$

Fix an arrow configuration and for each k, let $\pi_1^k, \ldots, \pi_n^k$ be the unique open paths in the arrow configuration with starting points $z_1^k, \ldots, z_n^k$. Then

$$\mathbb{P}\big[\theta_{\varepsilon_k}(\pi_1^k, \ldots, \pi_n^k) \in \cdot\,\big] \underset{k\to\infty}{\Longrightarrow} \mathbb{P}\big[(\pi_1, \ldots, \pi_n) \in \cdot\,\big],$$

where $\Rightarrow$ denotes weak convergence of probability measures on $(\Pi^\uparrow)^n$ and $\pi_1, \ldots, \pi_n$ are coalescing Brownian motions starting from $z_1, \ldots, z_n$.

Proof Our definition of coalescing Brownian motions involved a procedure that started with n independent Brownian motions $(B^1, \ldots, B^n)$ and used them to construct n coalescing Brownian motions $(\tilde{B}^1, \ldots, \tilde{B}^n)$. More formally, we can view $(\tilde{B}^1, \ldots, \tilde{B}^n)$ as the image of $(B^1, \ldots, B^n)$ under a map

$$(\pi_1, \ldots, \pi_n) \mapsto (\tilde{\pi}_1, \ldots, \tilde{\pi}_n) \tag{3.4}$$

that takes n paths $\pi_1, \ldots, \pi_n$ in $\Pi^\uparrow$ with starting points in $\mathbb{R}^2$ and maps them into n new paths $\tilde{\pi}_1, \ldots, \tilde{\pi}_n$ with the same starting points.

For each k, let $(R^{k,1}, \ldots, R^{k,n})$ be a collection of independent random walks started from $z_1^k, \ldots, z_n^k$, and let $(\tilde{R}^{k,1}, \ldots, \tilde{R}^{k,n})$ be its image under the map from (3.4). Then $(\tilde{R}^{k,1}, \ldots, \tilde{R}^{k,n})$ are coalescing random walks. It is easy to see that they are equal in law with $(\pi_1^k, \ldots, \pi_n^k)$. We want to show that

$$\mathbb{P}\big[\theta_{\varepsilon_k}(\tilde{R}^{k,1}, \ldots, \tilde{R}^{k,n}) \in \cdot\,\big] \underset{k\to\infty}{\Longrightarrow} \mathbb{P}\big[(\tilde{B}^1, \ldots, \tilde{B}^n) \in \cdot\,\big].$$

It is easy to see that the diffusive scaling map commutes with the map in (3.4), i.e., the random variable in the left-hand side of our equation is the same as what we would obtain if we first diffusively rescale the independent random walk paths and then apply the map from (3.4).

Weak convergence in law of diffusively rescaled independent random walks to independent Brownian motions follows from Donsker's invariance principle. Using Skorohod's representation theorem (Theorem 2.11), we can couple our random variables such that

$$\theta_{\varepsilon_k}(R^{k,1}, \ldots, R^{k,n}) \underset{k\to\infty}{\longrightarrow} (B^1, \ldots, B^n) \quad \text{a.s.}$$

in the topology on $(\Pi^\uparrow)^n$. If the map in (3.4) would be continuous with respect to the topology on $(\Pi^\uparrow)^n$, then the rest of the proof would now be easy, since we would just apply this map to both sides of our last equation and we would be done.

Things are not quite so simple, however, since it is easy to check (even for $n = 2$) that the map in (3.4) is not continuous with respect to the topology on $(\Pi^\uparrow)^n$. It turns out, however, that $(B^1, \ldots, B^n)$ is almost surely a point of continuity of this map, which is just as good. Here, with a point of continuity of the map in (3.4) we mean, of course, a collection of paths $(\pi_1, \ldots, \pi_n)$ with the property that for each $(\pi_1^k, \ldots, \pi_n^k)$ such that

$$(\pi_1^k, \ldots, \pi_n^k) \underset{k\to\infty}{\longrightarrow} (\pi_1, \ldots, \pi_n),$$

one also has

$$(\tilde\pi_1^k, \ldots, \tilde\pi_n^k) \underset{k\to\infty}{\longrightarrow} (\tilde\pi_1, \ldots, \tilde\pi_n).$$

That $(B^1, \ldots, B^n)$ is almost surely a point of continuity follows quite easily from our definitions and from Lemma 3.4 and Exercise 3.5 below. We leave the details to the reader. ∎

Lemma 3.4 (Brownian paths cross when they meet) *Let $B^i = (B^i_t)_{t\geq s_i}$ $(i = 1, 2)$ be independent Brownian motions started from deterministic space-time points $z_i = (x_i, s_i)$ $(i = 1, 2)$, respectively, and let*

$$\tau := \inf\{t \geq s_1 \vee s_2 : B^1_t = B^2_t\}.$$

Then almost surely, for each $\varepsilon > 0$, there exist times $t_-, t_+ \in [\tau, \tau + \varepsilon]$ such that

$$B^1_{t_-} < B^2_{t_-} \quad \textit{and} \quad B^1_{t_+} > B^2_{t_+}.$$

Proof By the strong Markov property, $(B^1_{\tau+t} - B^2_{\tau+t})_{t\geq 0}$ is a Brownian motion, so it suffices to prove that for a Brownian motion $(B_t)_{t\geq 0}$ started in zero both $\tau_- := \inf\{t \geq 0 : B_t < 0\}$ and $\tau_+ := \inf\{t \geq 0 : B_t > 0\}$ are a.s. zero. Since $(B_t)_{t\geq 0}$ is equally distributed with $(\sqrt{\lambda} B_{\lambda^{-1}t})_{t\geq 0}$, we see that $\tau_\pm$ is equally distributed with $\lambda^{-1}\tau_\pm$, for each $\lambda > 0$. It follows that the function $\lambda \mapsto \mathbb{P}[\tau_\pm \geq \lambda]$ is constant on $(0, \infty)$. However, if $\mathbb{P}[\tau_\pm \geq 1] > 0$, then it is easy to see that $\mathbb{P}[\tau_\pm \geq 2]$ must be strictly smaller than $\mathbb{P}[\tau_\pm \geq 1]$, so we conclude that $\mathbb{P}[\tau_\pm \geq \lambda] = 0$ for all $\lambda > 0$. ∎

Exercise 3.5 (Convergence of meeting times) *Let $\pi_1, \pi_2 \in \Pi^\uparrow$ have starting points $z_i = (x_i, s_i)$ $(i = 1, 2)$, respectively, and assume that their first meeting time*

$$\tau := \inf\{t \geq s_1 \vee s_2 : \pi_1(t) = \pi_2(t)\}$$

satisfies $\tau < \infty$. *Assume moreover that for each* $\varepsilon > 0$, *there exist times* $t_-, t_+ \in [\tau - \varepsilon, \tau + \varepsilon]$ *such that*

$$\pi_1(t_-) < \pi_2(t_-) \quad \textit{and} \quad \pi_1(t_+) > \pi_2(t_+).$$

Let $\pi_1^k, \pi_2^k \in \Pi^\uparrow$ *satisfy* $\pi_i^k \to \pi_i$ $(i = 1, 2)$. *Then the first meeting times* τ_k *of* π_1^k *and* π_2^k *satisfy* $\tau_k \to \tau$. Hint: *First show that generally* $\tau \leq \liminf_{k\to\infty} \tau_k$. *Then use the assumption about crossing to prove that* $\limsup_{k\to\infty} \tau_k \leq \tau_k$.

3.3 The Brownian web

Let $\mathcal{D} \subset \mathbb{R}^2$ be countable. Since $\mathcal{D}$ is countable, we can enumerate it as $\mathcal{D} := \{z_i : i \geq 1\}$ where $(z_i)_{i\geq 1}$ be a sequence of space-time points $z_i \in \mathbb{R}^2$. Then for each $n \geq 1$, we can construct a collection of random paths $(\pi_1, \ldots, \pi_n)$ that are distributed as coalescing Brownian motions starting from $(z_1, \ldots, z_n)$. Since these laws are consistent, by Kolmogorov's extension theorem, we can construct a random collection of paths $(\pi_z)_{z\in\mathcal{D}}$ such that for each finite set $\Delta \subset \mathcal{D}$, the paths $(\pi_z)_{z\in\Delta}$ that are distributed as coalescing Brownian motions starting from the points in Δ. We call $(\pi_z)_{z\in\mathcal{D}}$ a *collection of coalescing Brownian motions* started from the countable set $\mathcal{D}$.

Proposition 3.6 (Precompactness) *Let* $(\pi_z)_{z\in\mathcal{D}}$ *be a collection of coalescing Brownian motions started from a countable set* $\mathcal{D} \subset \mathbb{R}^2$. *Then* $\{\pi_z : z \in \mathcal{D}\}$ *is almost surely a precompact subset of* $\Pi^\uparrow$.

Proof (sketch) We apply Proposition 2.32 to $\mathcal{A} := \{\pi_z : z \in \mathcal{D}\}$. Fix $T < \infty$ and $\varepsilon, \delta > 0$ and consider the grid

$$\mathcal{G}_{\varepsilon,\delta} := \big\{(\tfrac{1}{3}k\varepsilon, l\delta) : k, l \in \mathbb{Z}\big\}.$$

Let $\mathcal{A}' = \{\pi'_z : z \in \mathcal{D} \cup \mathcal{G}_{\varepsilon,\delta}\}$ be a collection of coalescing Brownian motions started from the countable set $\mathcal{D} \cup \mathcal{G}_{\varepsilon,\delta}$. We can couple $\mathcal{A}'$ to $\mathcal{A}$ such that $\pi'_z = \pi_z$ for each $z \in \mathcal{D}$. Since paths in $\mathcal{A}$ cannot cross paths in $\{\pi'_z : z \in \mathcal{G}_{\varepsilon,\delta}\}$, it is not hard to see (see Figure 3.3) that almost surely on the event

$$\begin{aligned}&\big|\pi(u) - \pi(t)\big| \geq \varepsilon \text{ for some } \pi \in \mathcal{A} \text{ and } \sigma_\pi \leq t \leq u\\ &\qquad \text{s.t. } (\pi(t), t) \in [-T, T]^2,\ u - t \leq \delta\end{aligned}$$

one has that

$$\begin{aligned}&\big|\pi'_{(x,s)}(s + r) - x\big| \geq \tfrac{1}{3}\varepsilon\\ &\qquad \text{for some } (x, s) \in \mathcal{G}_{\varepsilon,\delta} \cap [-T - \varepsilon, T + \varepsilon]^2 \text{ and } r \in [0, 2\delta].\end{aligned} \tag{3.5}$$

By Lemma 3.11 below, if B is a standard Brownian motion, then

$$\mathbb{P}\big[\sup_{r\in[0,2\delta]}|B_r|\geq\tfrac{1}{3}\varepsilon\big]\leq Ce^{-c\varepsilon^2/\delta},$$

for some $C<\infty$ and $c>0$. A simple union bound then tells us that the probability of the event in (3.5) can be estimated from above by

$$C_T\varepsilon^{-1}\delta^{-1}e^{-c\varepsilon^2/\delta}$$

for some $C_T<\infty$ and $c>0$. This quantity goes to zero as $\delta\to 0$ for fixed $T<\infty$ and $\varepsilon>0$, so by Proposition 2.32 we conclude that $\{\pi_z : z\in\mathcal{D}\}$ is almost surely precompact. ∎

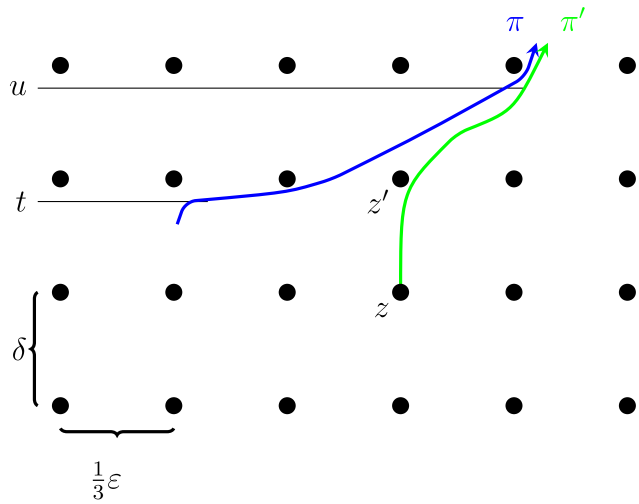

Figure 3.3: The tightness argument. The blue path π moves a distance $\geq\varepsilon$ during a time interval $[t,u]$ of length $\leq\delta$, forcing the green path π' starting from the point $z\in\mathcal{G}_{\varepsilon,\delta}$ to move a distance $\geq\varepsilon/3$ from its starting position during a time interval of length $2d$. Note that the blue path could have passed below the point $z'\in\mathcal{G}_{\varepsilon,\delta}$ that lies just above z.

We adopt the following notation. If $\mathcal{A}\subset\mathcal{K}(\Pi^\uparrow)$ is a collection of paths and $D\subset\mathcal{R}(\overline{\mathbb{R}})$ is a set, then we let

$$\mathcal{A}(D):=\big\{\pi\in\mathcal{A}:z_\pi\in D\big\} \tag{3.6}$$

denote the subset of $\mathcal{A}$ consisting of all paths that have their starting points in D. In particular, for $z\in\mathcal{R}(\overline{\mathbb{R}})$, we write $\mathcal{A}(z):=\mathcal{A}(\{z\})$. As before, we let $\overline{\mathcal{A}}$ denote the closure of a set $\mathcal{A} \subset \Pi^\uparrow$. The following theorem introduces the main object of interest of these lecture notes. See Figure 3.4 for an illustration.

Theorem 3.7 (The Brownian web) *There exists a random compact set $\mathcal{W} \subset \mathcal{K}(\Pi^\uparrow)$ whose distribution is uniquely determined by the following properties.*

(i) *For each $z \in \mathbb{R}^2$, almost surely there exists a unique $\pi_z \in \Pi^\uparrow$ such that $\mathcal{W}(z) = \{\pi_z\}$.*

(ii) *For each $z_1, \ldots, z_n \in \mathbb{R}^2$, the paths $(\pi_{z_1}, \ldots, \pi_{z_n})$ are distributed as coalescing Brownian motions starting from $z_1, \ldots, z_n$.*

(iii) *For each countable dense set $\mathcal{D} \subset \mathbb{R}^2$, almost surely $\mathcal{W} = \overline{\mathcal{W}(\mathcal{D})}$.*

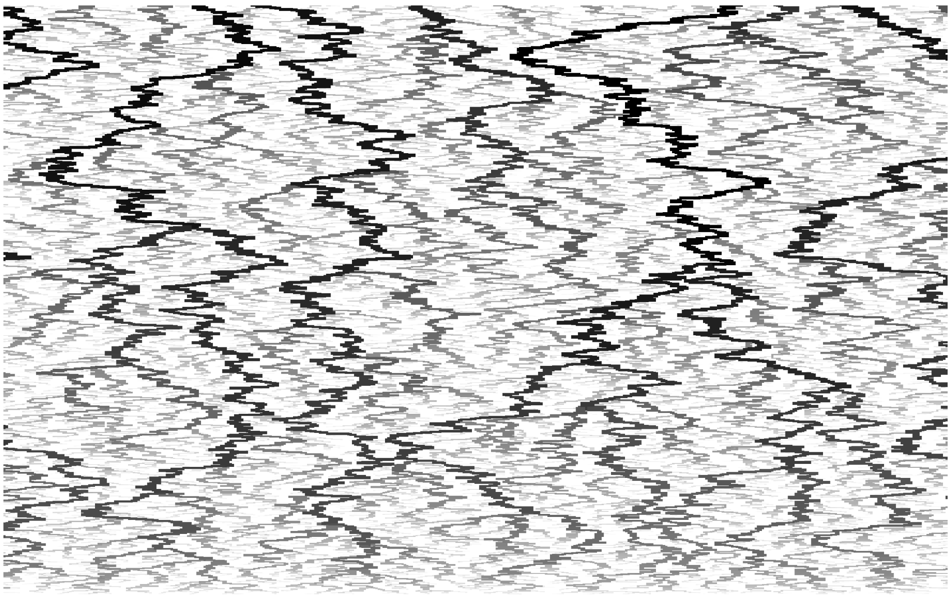

Figure 3.4: Artist's impression of the Brownian web.

Remark 1 In Section 4.4, we will see that in point (i) of Theorem 3.7, the order of the "for all" and "almost surely" statements cannot be interchanged. Although for a fixed, deterministic $z \in \mathbb{R}^2$, it is true that almost surely, $\mathcal{W}(z)$ consists of a single path, there exist random points $z \in \mathbb{R}^2$ in which $\mathcal{W}(z)$ has two, or even three elements.

Remark 2 A collection of paths of the form $\mathcal{W}(\mathcal{D})$, where $\mathcal{D}$ is any deterministic countable dense subset of $\mathbb{R}^2$, is called a *skeleton* of the Brownian web.

Proof of Theorem 3.7 Let $\mathcal{D} \subset \mathbb{R}^2$ be countable and dense and let $(\pi_z)_{z\in\mathcal{D}}$ be a collection of coalescing Brownian motions started from $\mathcal{D}$. Then $\{\pi_z : z \in \mathcal{D}\}$ is precompact by Proposition 3.6 and hence

$$\mathcal{W} := \overline{\{\pi_z : z \in \mathcal{D}\}} \tag{3.7}$$

is a random compact subset of $\Pi^{\uparrow}$. We claim that paths in $\mathcal{W}$ do not cross, in the sense that there do not exist $\pi, \pi' \in \mathcal{W}$ and $\sigma_\pi \vee \sigma_{\pi'} \leq s < t$ such that $\pi(s) < \pi'(s)$ but $\pi'(t) < \pi(t)$. Indeed, if such paths would exist, then they would be limits of paths π_n, π'_n in $\{\pi_z : z \in \mathcal{D}\}$ that would also have to cross for n large enough, which is not possible.

We will now show that $\mathcal{W}$ has the properties (i)–(iii) from the theorem. Fix $z = (x, s) \in \mathbb{R}^2$. Let ε_n be positive constants converging to zero, let $z_n^\pm := (x, s \pm \varepsilon_n)$, and let $\mathcal{D}' := \mathcal{D} \cup \{z_n^\pm : n \geq 1\}$. We can couple $(\pi_z)_{z\in\mathcal{D}}$ to a collection of coalescing Brownian motions $(\pi'_z)_{z\in\mathcal{D}'}$ started from $\mathcal{D}'$ such that $\pi_z = \pi'_z$ for all $z \in \mathcal{D}$. Let

$$\tau_n := \inf\{t \geq 0 : \pi'_{z_n^-}(t) = \pi'_{z_n^+}(t)\}. \tag{3.8}$$

Since paths cannot cross, we see that $\tau_1 \geq \tau_2 \geq \cdots$ and hence $\tau_n \to \tau_\infty$ a.s. for some random variable τ_∞. Using Lemma 3.4, it is easy to see that if we start two independent Brownian motions from z_n^- and z_n^+, then their first meeting time converges to zero in probability as $n \to \infty$. Together with our earlier observation, this implies that $\tau_\infty = s$ a.s. Since paths in $\mathcal{W}$ do not cross the paths $\pi'_{z_n^\pm}$, any path $\pi \in \mathcal{W}$ that starts in $(\pi(\sigma_\pi), \sigma_\pi) = z$ must satisfy

$$\pi'_{z_n^-}(t) \leq \pi(t) \leq \pi'_{z_n^+}(t) \qquad (t \geq s). \tag{3.9}$$

Since $\tau_\infty = s$ a.s., there can be at most one such path, proving property (i).

Property (ii) now follows from the fact that we can couple $(\pi_z)_{z\in\mathcal{D}}$ to a collection of coalescing Brownian motions $(\pi'_z)_{z\in\mathcal{D}\cup\{z_1,\ldots,z_n\}}$ such that $\pi_z = \pi'_z$ for all $z \in \mathcal{D}$. To prove property (iii), we ust show that our construction does not depend on the choice of the countable dense set $\mathcal{D}$. Let $\mathcal{D}$ and $\mathcal{D}'$ be countable dense subsets of $\mathbb{R}^2$, let $(\pi_z)_{z\in\mathcal{D}\cup\mathcal{D}'}$ be coalescing Brownian motions started from $\mathcal{D} \cup \mathcal{D}'$, and let

$$\begin{aligned} \mathcal{W} := \overline{\{\pi_z : z \in \mathcal{D}\}}, \quad \mathcal{W}' := \overline{\{\pi_z : z \in \mathcal{D}'\}}, \\ \text{and} \quad \mathcal{W}'' := \overline{\{\pi_z : z \in \mathcal{D} \cup \mathcal{D}'\}}. \end{aligned} \tag{3.10}$$

To prove (iii), it suffices to show that $\mathcal{W} = \mathcal{W}'$. By symmetry, it suffices to show that $\mathcal{W} \subset \mathcal{W}'$. Since both $\mathcal{W}$ and $\mathcal{W}'$ are closed, it suffices to show that for each $z \in \mathcal{D}$, the path π_z satisfies $\pi_z \in \mathcal{W}'$. By what we have already proved, there exists unique paths $\pi' \in \mathcal{W}'$ and $\pi'' \in \mathcal{W}''$ with starting points $z_{\pi'} = z_{\pi''} = z$. Since $\pi' \in \mathcal{W}''$ we must have $\pi' = \pi''$ and since $\pi_z \in \mathcal{W}''$ we must have $\pi'' = \pi_z$, so we conclude that $\pi_z = \pi'' = \pi' \in \mathcal{W}'$. ∎

For the next lemma, we let

$$\Pi^{\uparrow}_{\rm triv} := \big\{\pi \in \Pi^{\uparrow} : \pi(t) = -\infty\ \forall t \geq \sigma_\pi\big\} \cup \big\{\pi \in \Pi^{\uparrow} : \pi(t) = +\infty\ \forall t \geq \sigma_\pi\big\}$$

denote the set of trivial paths that are constantly $-\infty$ or ∞.

Lemma 3.8 (Trivial paths) *Let $\mathcal{W}$ be a Brownian web. Then $\Pi^{\uparrow}_{\rm triv} \subset \mathcal{W}$ a.s. and each $\pi \in \mathcal{W}\backslash\Pi^{\uparrow}_{\rm triv}$ satisfies $\pi(t) \in \mathbb{R}$ for all $\sigma_\pi \leq t < \infty$.*

Proof This follows from the same argument as in the proof of Proposition 3.1. ∎

We still need to provide an estimate that we have used in the proof of Proposition 3.6.

Lemma 3.9 (Reflection principle) *Let $(B_t)_{t\geq 0}$ be Brownian motion. Then*

$$\mathbb{P}\big[\sup_{s\in[0,t]} B_s < a\big] = \mathbb{P}\big[|B_t| \leq a\big] \qquad (t,a > 0). \tag{3.11}$$

Proof Let $\tau := \inf\{t > 0 : B_t = a\}$. By the strong Markov property and the symmetry of Brownian motion, conditional on the event $\{\tau < t\}$, the events $\{B_t > a\}$ and $\{B_t < a\}$ have equal probabilities (see Figure 3.5). Since $\mathbb{P}[B_t = a] = 0$ and the event $\{B_t > a\}$ almost surely implies $\{\tau < t\}$, it follows that

$$\mathbb{P}\big[\sup_{s\in[0,t]} B_s < a\big] = 1 - 2\mathbb{P}[B_t > a] = \mathbb{P}\big[|B_t| \leq a\big]. \tag{3.12}$$

∎

Lemma 3.10 (Tail estimate) *Let N be a standard normal random variable. Then*

$$\mathbb{P}[N \geq a] \leq \tfrac{1}{2} e^{-a^2/2}. \tag{3.13}$$

Proof This follows by writing

$$\begin{aligned}\mathbb{P}[N \geq a] &= \frac{1}{\sqrt{2\pi}}\int_a^\infty e^{-x^2/2}\mathrm{d}x = \frac{1}{\sqrt{2\pi}}\int_0^\infty e^{-(x+a)^2/2}\mathrm{d}x\\ &= e^{-a^2/2}\frac{1}{\sqrt{2\pi}}\int_0^\infty e^{-x^2/2 - ax}\mathrm{d}x\\ &\leq e^{-a^2/2}\frac{1}{\sqrt{2\pi}}\int_0^\infty e^{-x^2/2}\mathrm{d}x = \tfrac{1}{2}e^{-a^2/2}.\end{aligned} \tag{3.14}$$

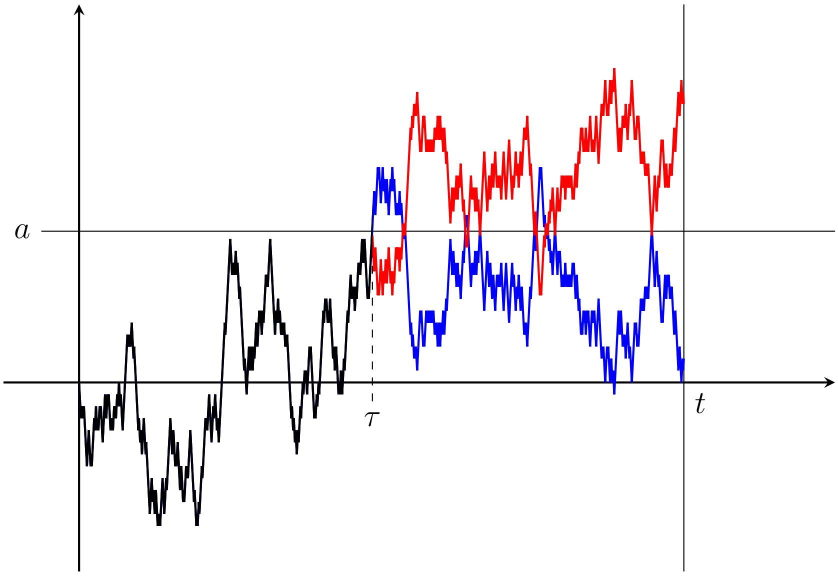

Figure 3.5: The reflection principle: $\mathbb{P}[\tau < t] = 2\mathbb{P}[B_t > a]$.

∎

Lemma 3.11 (Large displacements) *Let* $(B_t)_{t\geq 0}$ *be Brownian motion. Then*

$$\mathbb{P}\big[\sup_{s\in[0,t]} |B_s| \geq a\big] \leq 2e^{a^2/(2t)}. \tag{3.15}$$

Proof Let N denote a standard normal random variable. We estimate, using Lemmas 3.9 and 3.10,

$$\begin{aligned}\mathbb{P}\big[\sup_{s\in[0,t]} |B_s| \geq a\big] &\leq 2\mathbb{P}\big[\sup_{s\in[0,t]} B_s \geq a\big] = 2\mathbb{P}\big[|B_t| > a\big] = 4\mathbb{P}[B_t > a] \\ &= 4\mathbb{P}[\sqrt{t}N > a] \leq 2e^{a^2/(2t)}.\end{aligned} \tag{3.16}$$

∎

3.4 Dual arrow configurations

By definition, we call

$$\mathbb{Z}^2_{\rm odd} := \big\{(x,t) \in \mathbb{Z}^2 : x + t \text{ is odd}\big\}$$

the *odd sublattice* of $\mathbb{Z}^2$. In Section 3.1, we showed how an i.i.d. collection $\omega = (\omega_z)_{z\in\mathbb{Z}^2_{\rm even}}$ of uniformly distributed $\{-1,+1\}$-valued random variables defines a random directed graph $(\mathbb{Z}^2_{\rm even}, \vec{E})$ that we called an arrow configuration. Given ω, we define $\hat\omega = (\hat\omega_z)_{z\in\mathbb{Z}^2_{\rm odd}}$ by

$$\hat\omega_{(x,t+1)} = \omega_{(x,t)} \qquad \big((x,t)\in\mathbb{Z}^2_{\rm even}\big). \tag{3.17}$$

We can use $\hat\omega$ to define a random directed graph with vertex set $\mathbb{Z}^2_{\rm odd}$ and set of oriented edges

$$\vec{F} := \big\{\big((x,t),(x-\hat\omega_{(x,t)},t-1)\big) : (x,t)\in\mathbb{Z}^2_{\rm odd}\big\}.$$

We call the random directed graph $(\mathbb{Z}^2_{\rm odd}, \vec{F})$ the *dual arrow configuration* associated with the original ("forward") arrow configuration $(\mathbb{Z}^2_{\rm even}, \vec{E})$. The dual arrows are uniquely characterised in terms of the forward arrows by the property that dual arrows and forward arrows do not cross. See Figure 3.6 for a picture.

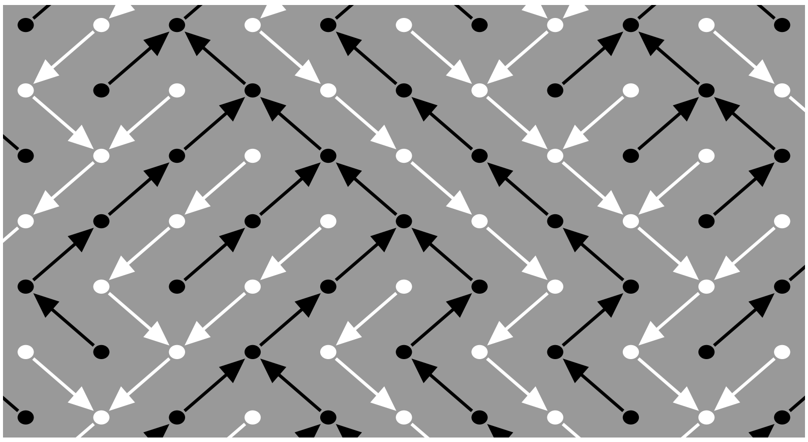

Figure 3.6: An arrow configuration (black) and its dual (white).

Recall that in general, σ_π and τ_π denote the starting and final time of a path $\pi\in\Pi(\overline{\mathbb{R}})$. In particular, we define

$$\Pi^\downarrow := \big\{\pi\in\Pi(\overline{\mathbb{R}}) : \sigma_\pi = -\infty\big\}.$$

We call $\Pi^\downarrow$ the space of all *downward paths.* Clearly, $\Pi^\downarrow$ is equal to $\Pi^\uparrow$ after a rotation over 180 degrees. When no confusion can arrive,[2] we will call the

[2]We have to be careful since the intersection of $\Pi^\uparrow$ and $\Pi^\downarrow$ is not empty, but consists of all bi-infinite paths for which $\sigma_\pi = -\infty$ and $\tau_\pi = \infty$. As we will see in a moment, however, there are no nontrivial bi-infinite paths in an arrow configuration.

point

$$z_\pi := \big(\pi(\tau_\pi), \tau_\pi\big)$$

the *starting point* of a downward path $\pi \in \Pi^\downarrow$. We define a *downward open path in the dual arrow configuration* $(\mathbb{Z}^2_{\rm odd}, \vec{F})$, or simply a *open path in* $\hat{\omega}$ in exactly the same way as we defined upward open paths in the forward arrow configuration. We let

$$\mathcal{U}^* = \mathcal{U}^*(\hat{\omega}) := \big\{\pi \in \Pi^\downarrow : \pi \text{ is a open path in } \hat{\omega}\big\} \tag{3.18}$$

denote the set of all downward open paths in the dual arrow configuration and we let $\overline{\mathcal{U}}^*$ denote the closure of $\mathcal{U}^*$ in the topology on $\Pi^\downarrow$.

3.5 The dual Brownian web

We have already introduced notation for the diffusive scaling map θ_ε which may be applied to points $z = (x,t)$ in space-time $\mathcal{R}(\overline{\mathbb{R}})$, to subsets of space-time such as paths, and even to sets of paths. We will use similar notation for the map

$$\mathcal{R}(\overline{\mathbb{R}}) \ni (x,t) \mapsto -(x,t) = (-x,-t) \in \mathcal{R}(\overline{\mathbb{R}}).$$

Thus, for any set $A \subset \mathcal{R}(\overline{\mathbb{R}})$, we set $-A := \{-z : z \in A\}$. In particular, this applies to the case that $A = \pi \in \Pi^\uparrow$. Then $\Pi^\uparrow \ni \pi \mapsto -\pi \in \Pi^\downarrow$ is a bijection from $\Pi^\uparrow$ to $\Pi^\downarrow$. Also, if $\mathcal{A} \subset \Pi^\uparrow$ is a sets whose elements are paths, then we set $-\mathcal{A} := \{-\pi : \pi \in \mathcal{A}\}$. Using this notation, we say that $\hat{\pi}_1, \ldots, \hat{\pi}_n$ are *downward* coalescing Brownian motions starting from space-time points $z_1, \ldots, z_n$ if $-\hat{\pi}_1, \ldots, -\hat{\pi}_n$ are (usual, forward) coalescing Brownian motions starting from space-time points $-z_1, \ldots, -z_n$. In the same way, we define countable collections of downward coalescing Brownian motions.

Let $\hat{\pi}_1, \hat{\pi}_2 \in \Pi^\downarrow$ be two downward paths started from space-time points $(x_i, s_i) \in \mathbb{R}^2$ $(i = 1,2)$, and let

$$\tau = \tau(\hat{\pi}_1, \hat{\pi}_2) := \sup\big\{t < s_1 \wedge s_2 : \hat{\pi}_1(t) = \hat{\pi}_2(t)\big\}$$

be their first meeting time (in the downward direction), which may be $-\infty$. The open set

$$W(\hat{\pi}_1, \hat{\pi}_2) := \big\{(x,t) : \tau < t < s_1 \wedge s_2 : \hat{\pi}_1(t) < x < \hat{\pi}_2(t)\big\}$$

is called the *wedge* defined by $\hat{\pi}_1, \hat{\pi}_2$. See Figure 3.7 for an illustration. We say that a (forward) path $\pi \in \Pi^\uparrow$ *enters* the wedge $W(\hat{\pi}_1, \hat{\pi}_2)$ if there exist times $\sigma_\pi \leq s < t$ such that

$$\big(\pi(s), s\big) \not\in \overline{W}(\hat{\pi}_1, \hat{\pi}_2) \quad \text{and} \quad \big(\pi(t), t\big) \in W(\hat{\pi}_1, \hat{\pi}_2),$$

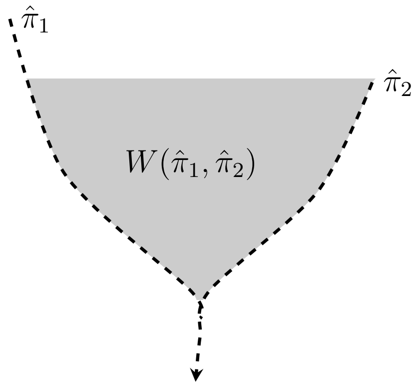

Figure 3.7: The wedge $W(\hat\pi_1,\hat\pi_2)$ defined by the dual paths $\hat\pi_1$ and $\hat\pi_2$.

where $\overline{W}(\hat\pi_1,\hat\pi_2)$ denotes the closure of $W(\hat\pi_1,\hat\pi_2)$. In a completely analogous way, we define the first meeting time of two forward paths, the wedge defined by two forward paths, and what it means for a downward path to enter such a wedge. We make the following simple observation.

Lemma 3.12 (Limits of wedges) *Let $(\hat\pi^n_i)_{n\geq 1}$ $(i=1,2)$ be sequences of downward paths and let $(\pi^n)_{n\geq 1}$ be a sequence of forward paths. Assume that there exist $\hat\pi_i\in\Pi^\downarrow$ $(i=1,2)$ and $\pi\in\Pi^\uparrow$ such that*

$$\hat\pi^n_i \underset{n\to\infty}{\longrightarrow} \hat\pi_i \quad (i=1,2) \quad \textit{and} \quad \pi^n \underset{n\to\infty}{\longrightarrow} \pi$$

in the topologies on $\Pi^\downarrow$ and $\Pi^\uparrow$, and that moreover

$$\tau(\hat\pi^n_1,\hat\pi^n_2) \underset{n\to\infty}{\longrightarrow} \tau(\hat\pi_1,\hat\pi_2).$$

Assume that for each n, the path π^n does not enter the wedge $W(\hat\pi^n_1,\hat\pi^n_2)$. Then the path π does not enter the wedge $W(\hat\pi_1,\hat\pi_2)$.

Proof By definition, if π enters the wedge $W(\hat\pi_1,\hat\pi_2)$, then there exist times $\sigma_\pi\leq s<t$ such that

$$\big(\pi(s),s\big)\not\in\overline{W}(\hat\pi_1,\hat\pi_2) \quad \text{and} \quad \big(\pi(t),t\big)\in W(\hat\pi_1,\hat\pi_2).$$

Since $\pi_n\to\pi$, there exist times $\sigma_{\pi_n}\leq s_n<t_n$ such that $\big(\pi_n(s_n),s_n\big)\to\big(\pi(s),s\big)$ and $\big(\pi_n(t_n),t_n\big)\to\big(\pi(t),t\big)$. We claim that for n sufficiently large,

$$\big(\pi^n(s_n),s_n\big)\not\in\overline{W}(\hat\pi^n_1,\hat\pi^n_2)$$

Indeed, if $\big(\pi^n(s_n), s_n\big) \in \overline{W}(\hat\pi_1^n, \hat\pi_2^n)$ for infinitely many values of n, then going to a subsequence and taking the limit, using the convergence of the paths and meeting times, we would find that $\big(\pi(s), s\big) \in \overline{W}(\hat\pi_1, \hat\pi_2)$, which contradicts our assumptions. In the same way, we see that

$$\big(\pi^n(t_n), t_n\big) \in W(\hat\pi_1^n, \hat\pi_2^n)$$

for n sufficiently large, so we arrive at a contradiction with the assumption that π^n does not enter $W(\hat\pi_1^n, \hat\pi_2^n)$. ∎

Proposition 3.13 (Dual coalescing Brownian motions) *Let $\mathcal{D}, \hat{\mathcal{D}}$ be countable dense subsets of $\mathbb{R}^2$. Then it is possible to construct a collection $(\pi_z)_{z\in\mathcal{D}}$ of coalescing Brownian motions together with a collection $(\hat\pi_z)_{z\in\hat{\mathcal{D}}}$ of downward coalescing Brownian motions in such a way that:*

- *For each $z \in \mathcal{D}$ and $z_1, z_2 \in \hat{\mathcal{D}}$, the path π_z does not enter the wedge $W(\hat\pi_{z_1}, \hat\pi_{z_2})$.*
- *For each $z \in \hat{\mathcal{D}}$ and $z_1, z_2 \in \mathcal{D}$, the downward path $\hat\pi_z$ does not enter the wedge $W(\pi_{z_1}, \pi_{z_2})$.*

The proof of Proposition 3.13 makes use of the following simple lemma.

Lemma 3.14 (Tightness of joint law) *Let $\mathcal{X}, \mathcal{Y}$ be Polish spaces, let $(X_n, Y_n)_{n\geq 1}$ be a sequence of random variables with values in $\mathcal{X} \times \mathcal{Y}$, and let X and Y be random variables with values in $\mathcal{X}$ and $\mathcal{Y}$, respectively. Assume that*

$$\mathbb{P}[X_n \in \cdot\,] \underset{n\to\infty}{\Longrightarrow} \mathbb{P}[X \in \cdot\,] \quad \textit{and} \quad \mathbb{P}[Y_n \in \cdot\,] \underset{n\to\infty}{\Longrightarrow} \mathbb{P}[Y \in \cdot\,]$$

Then the probability laws

$$\big(\mathbb{P}\big[(X_n, Y_n) \in \cdot\,\big]\big)_{n\geq 1}$$

are tight.

Proof The convergence of the marginal laws implies that the probability laws

$$\big(\mathbb{P}[X_n \in \cdot\,]\big)_{n\geq 1} \quad \text{and} \quad \big(\mathbb{P}[Y_n \in \cdot\,]\big)_{n\geq 1}$$

are tight, so for each $\varepsilon > 0$, there exist compact sets $C \subset \mathcal{X}$ and $K \subset \mathcal{Y}$ such that

$$\sup_{n\geq 1} \mathbb{P}[X_n \notin C] \leq \varepsilon \quad \text{and} \quad \sup_{n\geq 1} \mathbb{P}[Y_n \notin K] \leq \varepsilon$$

Then $C \times K$ is compact and

$$\sup_{n\geq 1} \mathbb{P}\big[(X_n, Y_n) \notin C \times K\big] \leq 2\varepsilon.$$

Since $\varepsilon > 0$ is arbitrary, it follows that the laws of (X_n, Y_n) are tight. ∎

Proof of Proposition 3.13 (sketch) Let $\mathcal{U}$ be the collection of open paths in an arrow configuration and let $\mathcal{U}^*$ be the collection of downward open paths in the associated dual arrow configuration. Let ε_n be positive constants tending to zero. For each $z \in \mathcal{D}$, choose $z_n \in \mathbb{Z}^2_{\rm even}$ such that $\theta_{\varepsilon_n}(z_n) \to z$, and for each $z \in \hat{\mathcal{D}}$, choose $z^n \in \mathbb{Z}^2_{\rm odd}$ such that $\theta_{\varepsilon_n}(z^n) \to z$. For each $z \in \mathcal{D}$ and $n \geq 1$, let $R^n_z \in \mathcal{U}$ be the unique forward open path starting at z_n, let $\hat{R}^n_z \in \mathcal{U}^*$ be the unique downward open path starting at z^n, and let

$$\pi^n_z := \theta_{\varepsilon_n}(R^n_z) \quad \text{and} \quad \hat{\pi}^n_z := \theta_{\varepsilon_n}(\hat{R}^n_z)$$

denote the associated diffusively rescaled paths. We claim that

$$\begin{aligned} \mathbb{P}\big[(\pi^n_z)_{z\in\mathcal{D}} \in \cdot\,\big] &\underset{n\to\infty}{\Longrightarrow} \mathbb{P}\big[(\pi_z)_{z\in\mathcal{D}} \in \cdot\,\big], \\ \mathbb{P}\big[(\hat{\pi}^n_z)_{z\in\hat{\mathcal{D}}} \in \cdot\,\big] &\underset{n\to\infty}{\Longrightarrow} \mathbb{P}\big[(\hat{\pi}_z)_{z\in\mathcal{D}} \in \cdot\,\big], \end{aligned}$$

where $\Rightarrow$ denotes weak convergence of probability laws on the spaces $(\Pi^\uparrow)^{\mathcal{D}}$ and $(\Pi^\downarrow)^{\mathcal{D}}$, respectively, which are equipped with the product topology, and $(\pi_z)_{z\in\mathcal{D}}$ is a collection of coalescing Brownian motions while $(\hat{\pi}_z)_{z\in\hat{\mathcal{D}}}$ is a collection of downward coalescing Brownian motions. Indeed, to prove this, by the definition of the product topology, it suffices to prove convergence of finite dimensional distributions. But this has already been done in Proposition 3.3.

In fact, using Exercise 3.5, we can strengthen our previous claim in a sense that also includes convergence of meeting times. More precisely, one can show that

$$\begin{aligned} &\mathbb{P}\Big[\big((\pi^n_z)_{z\in\mathcal{D}}, (\tau(\pi^n_{z_1}, \pi^n_{z_2}))_{(z_1,z_2)\in\mathcal{D}^2}\big) \in \cdot\,\Big] \\ &\quad \underset{n\to\infty}{\Longrightarrow} \mathbb{P}\Big[\big((\pi_z)_{z\in\mathcal{D}}, (\tau(\pi_{z_1}, \pi_{z_2}))_{(z_1,z_2)\in\mathcal{D}^2}\big) \in \cdot\,\Big], \end{aligned} \tag{3.19}$$

and similarly for the collection of downward paths.

By Lemma 3.14, going to a subsequence if necessary, we can assume that the joint law of the random variables

$$(\pi^n_z)_{z\in\mathcal{D}}, \quad (\tau(\pi^n_{z_1}, \pi^n_{z_2}))_{(z_1,z_2)\in\mathcal{D}^2}, \quad (\hat{\pi}^n_z)_{z\in\mathcal{D}}, \quad (\tau(\hat{\pi}^n_{z_1}, \hat{\pi}^n_{z_2}))_{(z_1,z_2)\in\mathcal{D}^2}$$

converges weakly. Then we can use Skorohod's representation theorem (Theorem 2.11) to couple our random variables so that the convergence is almost sure, i.e., we can find a coupling such that

$$\pi^n_z \underset{n\to\infty}{\longrightarrow} \pi_z \text{ a.s.} \quad \text{and} \quad \tau(\pi^n_{z_1}, \pi^n_{z_2}) \underset{n\to\infty}{\longrightarrow} \tau(\pi_{z_1}, \pi_{z_2}) \text{ a.s.}$$

for all $z, z_1, z_2 \in \mathcal{D}$, and likewise for downward paths. Since paths of $\mathcal{U}$ do not enter wedges of $\mathcal{U}^*$ and vice versa, we can use Lemma 3.12 to conclude that the same is true for the limit object. ∎

Theorem 3.15 (Wedge characterisation of the Brownian web) *Let $\mathcal{D}, \hat{\mathcal{D}}$ be countable dense subsets of $\mathbb{R}^2$, let $(\pi_z)_{z\in\mathcal{D}}$ be a collection of coalescing Brownian motions started from $\mathcal{D}$, and let $(\hat{\pi}_z)_{z\in\hat{\mathcal{D}}}$ be a collection of downward coalescing Brownian motions started from $\hat{\mathcal{D}}$. Assume that paths in $(\pi_z)_{z\in\mathcal{D}}$ do not enter wedges of $(\hat{\pi}_z)_{z\in\hat{\mathcal{D}}}$. Let*

$$\begin{aligned}\mathcal{W}_- &:= \overline{\{\pi_z : z \in \mathcal{D}\}},\\ \mathcal{W}_+ &:= \big\{\pi \in \Pi^\uparrow : \pi \text{ does not enter wedges of } (\hat{\pi}_z)_{z\in\hat{\mathcal{D}}}\big\}.\end{aligned}$$

Then $\mathcal{W}_- = \mathcal{W}_+$.

Proof (sketch) To prove the inclusion $\mathcal{W}_- \subset \mathcal{W}_+$, let $\pi \in \mathcal{W}_-$. Then there exists $z_n \in \mathcal{D}$ such that $\pi_{z_n} \to \pi$ as $n \to \infty$. Let $z^1, z^2 \in \hat{\mathcal{D}}$. By assumption, π_{z_n} does not enter the wedge $W(\hat{\pi}_{z^1}, \hat{\pi}_{z^2})$ for any $n \geq 1$. By Lemma 3.12, it follows that π does not enter $W(\hat{\pi}_{z^1}, \hat{\pi}_{z^2})$. This completes the proof that $\mathcal{W}_- \subset \mathcal{W}_+$.

Before we continue, we note that our assumptions imply that the forward paths do not cross downward paths, in the sense that if $z = (x, s) \in \mathcal{D}$ and $z' = (y, u) \in \hat{\mathcal{D}}$ satisfy $s < u$, then $\pi_z(s) < \hat{\pi}_{z'}(s)$ implies $\pi_z(t) \leq \hat{\pi}_{z'}(t)$ for all $t \in [s, u]$. Indeed, we can always choose some $z'' = (y', u') \in \hat{\mathcal{D}}$ with $u \leq u'$ such that $\hat{\pi}_{z'}(u) < \hat{\pi}_{z''}(u)$ and the meeting time $\tau(\hat{\pi}_{z'}, \hat{\pi}_{z''})$ is less than s. Then $\pi_z(t) > \hat{\pi}_{z'}(t)$ for some $t \in (s, u]$ would imply that π_z enters the wedge $W(\hat{\pi}_{z'}, \hat{\pi}_{z''})$, contradicting our assumptions.

We now prove that $\mathcal{W}_+ \subset \mathcal{W}_-$. Let $\pi \in \mathcal{W}_+$. By Lemma 3.8 we can without loss of generality assume that $\pi(t) \in \mathbb{R}$ for all $t \in I_\pi$. Fix $\sigma_\pi < t_1 < \cdots < t_m$ and $\varepsilon > 0$. We claim that there exists a $z = (x, s) \in \mathcal{D}$ such that $\sigma_\pi < s < t_1$ and $|\pi_z(t_i) - \pi(t)| \leq \varepsilon$ for all $i = 1, \ldots, m$. To prove this, we will use a "fish trap" construction illustrated in Figure 3.8. For each $i = 1, \ldots, m$, we choose $z^i_\pm = (x^i_\pm, t^i_\pm) \in \hat{\mathcal{D}}$ such that $t^i_\pm > t_i$ and

$$\pi(t_i) - \varepsilon < \hat{\pi}_{z_-}(t_i) < \pi(t_i) < \hat{\pi}_{z_+}(t_i) < \pi(t_i) + \varepsilon.$$

Since π does not enter the wedge $W(\hat{\pi}_{z^i_-}, \hat{\pi}_{z^i_+})$, the meeting time of $\hat{\pi}_{z^i_-}$ and $\hat{\pi}_{z^i_+}$ must satisfy

$$\tau(\hat{\pi}_{z^i_-}, \hat{\pi}_{z^i_+}) \leq \sigma_\pi,$$

and we have $\hat{\pi}_{z^i_-}(t) \leq \pi(t) \leq \hat{\pi}_{z^i_+}(t)$ for all $t \in [\sigma_\pi, t_i]$. We can now choose $z = (x, s) \in \mathcal{D}$ such that $\sigma_\pi < s < t_1$ and

$$\sup_{1\leq i\leq m} \hat{\pi}_{z^i_-}(t_1) < \pi_z(t_1) < \inf_{1\leq i\leq m} \hat{\pi}_{z^i_+}(t_1).$$

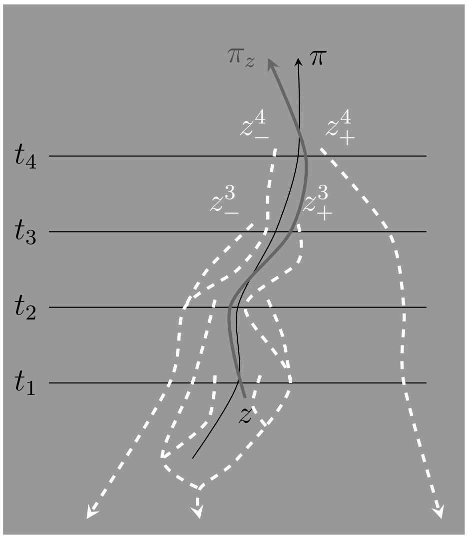

Figure 3.8: "Fish trap" construction showing that the path π that does not enter wedges can be approximated by a path π_z starting at a point z taken from a deterministic countable dense set of space-time points.

Since the path π_z cannot cross any of the downward paths $\hat{\pi}_{z^i_\pm}$, we must have

$$\hat{\pi}_{z_-}(t_i) < \pi_z(t_i) < \hat{\pi}_{z_+}(t_i) \qquad (1 \leq i \leq m)$$

and hence $|\pi_z(t_i) - \pi(t)| \leq \varepsilon$ for all $i = 1, \ldots, m$, proving our claim.

Now let $\varepsilon_n > 0$ satisfy $\varepsilon_n \to 0$ and let $\sigma_\pi < t_1 < \cdots < t_m$. By what we have just proved, for each n there exists a $z_n \in \mathcal{D}$ such that $|\pi_{z_n}(t_i) - \pi(t)| \leq \varepsilon$ for all $i = 1, \ldots, m$. By Proposition 3.6, the closure of $\{\pi_z : z \in \mathcal{D}\}$ is compact, so we can find a convergent subsequence. It follows that there exists a $\pi' \in \mathcal{W}_-$ such that $\pi'(t_i) = \pi(t_i)$ for all $i = 1, \ldots, m$. Now let $\{t_i : i \in \mathbb{N}\} \subset (\sigma_\pi, \infty)$ be countable and dense. By what we have just proved, for each m, there exists a $\pi_m \in \mathcal{W}_-$ such that $\pi_m(t_i) = \pi(t_i)$ for all $i = 1, \ldots, m$. Since $\mathcal{W}_+$ is compact, we can find a convergent subsequence, the limit of which must be the path π. This proves that $\mathcal{W}_+ \subset \mathcal{W}_-$. ∎

3.6 Convergence to the Brownian web

Proposition 3.16 (Tightness of rescaled arrow configurations) *Let $\mathcal{U}$ be the set of all open paths in an arrow configurations and let $\overline{\mathcal{U}}$ be its closure. Let $\varepsilon_n > 0$ be positive constants such that $\varepsilon_n \to 0$. The the probability laws*

$$\big(\mathbb{P}[\theta_{\varepsilon_n}(\overline{\mathcal{U}}) \in \cdot\,]\big)_{n \geq 1}$$

on $\mathcal{K}(\Pi^\uparrow)$ are tight.

Proof (crude sketch) One needs to check the tightness criterion of Proposition 2.33. This is very similar to the proof of Proposition 3.6. One uses convergence of finite dimensional distributions (Proposition 3.3) and then uses a grid as in the proof of Proposition 3.6 to estimate the event in Proposition 2.33. We refer to [FINR04, Prop. B2] and [SSS16, Prop. 6.6.4] for details. ∎

Let $\mathcal{D}, \hat{\mathcal{D}}$ be countable dense subsets of $\mathbb{R}^2$. By Proposition 3.13, we can construct a collection $(\pi_z)_{z\in\mathcal{D}}$ of coalescing Brownian motions starting from $\mathcal{D}$ and a collection $(\hat\pi_z)_{z\in\hat{\mathcal{D}}}$ of downward coalescing Brownian motions starting from $\hat{\mathcal{D}}$ such that paths in $(\pi_z)_{z\in\mathcal{D}}$ do not enter wedges of $(\hat\pi_z)_{z\in\hat{\mathcal{D}}}$ and vice versa. We call the pair $(\mathcal{W}, \hat{\mathcal{W}})$ defined as

$$\mathcal{W} := \overline{\{\pi_z : z \in \mathcal{D}\}} \quad \text{and} \quad \hat{\mathcal{W}} := \overline{\{\hat\pi_z : z \in \hat{\mathcal{D}}\}} \tag{3.20}$$

the *double Brownian web* and we call $\hat{\mathcal{W}}$the *dual Brownian web*

Lemma 3.17 (Double Brownian web) *The law of the random variable $(\mathcal{W}, \hat{\mathcal{W}})$ does not depend on the choice of the countable dense sets $\mathcal{D}, \hat{\mathcal{D}} \subset \mathbb{R}^2$.*

Proof The analogue statement for the Brownian web has already been proved as part of the proof of Theorem 3.7, around (3.10). The statement for a single web does, as far as I can see, not automatically imply the statement for the double Brownian web, but one can adapt the argument given at (3.10). Here we give an alternative argument that also reproofs the statement for a single web and does not depend on the earlier argument.

Let $\mathcal{D}, \mathcal{D}', \hat{\mathcal{D}}$ be countable dense subsets of $\mathbb{R}^2$. Let $(\pi_z)_{z\in\mathcal{D}}$ be a collection of coalescing Brownian motions starting from $\mathcal{D}$, let $(\pi'_z)_{z\in\mathcal{D}'}$ be a collection of coalescing Brownian motions starting from $\mathcal{D}'$, and let $(\hat\pi_z)_{z\in\hat{\mathcal{D}}}$ be a collection $(\hat\pi_z)_{z\in\hat{\mathcal{D}}}$ of downward coalescing Brownian motions starting from $\hat{\mathcal{D}}$. By Proposition 3.13, we can couple $(\pi_z)_{z\in\mathcal{D}}$ to $(\hat\pi_z)_{z\in\hat{\mathcal{D}}}$ in such a way that paths in $(\pi_z)_{z\in\mathcal{D}}$ do not enter wedges of $(\hat\pi_z)_{z\in\hat{\mathcal{D}}}$ and vice versa. Similarly, we can couple $(\pi'_z)_{z\in\mathcal{D}'}$ to $(\hat\pi_z)_{z\in\hat{\mathcal{D}}}$ in such a way that paths in $(\pi'_z)_{z\in\mathcal{D}'}$ do not enter wedges of $(\hat\pi_z)_{z\in\hat{\mathcal{D}}}$ and vice versa. We can then couple all three collections $(\pi_z)_{z\in\mathcal{D}}$, $(\pi'_z)_{z\in\mathcal{D}'}$, and $(\hat\pi_z)_{z\in\hat{\mathcal{D}}}$ in such a way that the joint law of $(\pi_z)_{z\in\mathcal{D}}$ and $(\hat\pi_z)_{z\in\hat{\mathcal{D}}}$ is as before and the joint law of $(\pi'_z)_{z\in\mathcal{D}'}$ and $(\hat\pi_z)_{z\in\hat{\mathcal{D}}}$ is also as before. For example, this can be achieved by making $(\pi_z)_{z\in\mathcal{D}}$ and $(\pi'_z)_{z\in\mathcal{D}'}$ conditionally indepenent given $(\hat\pi_z)_{z\in\hat{\mathcal{D}}}$, and with the same conditional laws as before.

For this coupling, let $(\mathcal{W}, \hat{\mathcal{W}})$ be defined using $\mathcal{D}, \hat{\mathcal{D}}$ and let $(\mathcal{W}', \hat{\mathcal{W}})$ be defined using $\mathcal{D}', \hat{\mathcal{D}}$. Then Theorem 3.15 tells us that

$$\mathcal{W} = \big\{\pi \in \Pi^\uparrow : \pi \text{ does not enter wedges of } (\hat{\pi}_z)_{z\in\hat{\mathcal{D}}}\big\} = \mathcal{W}' \quad \text{a.s.}$$

It follows that the joint law of $(\mathcal{W}, \hat{\mathcal{W}})$ is the same as the joint law of $(\mathcal{W}', \hat{\mathcal{W}})$. In the same way, we can also replace $\hat{\mathcal{D}}$ by another countable dense subset of $\mathbb{R}^2$ without changing the law of the double Brownian web. ∎

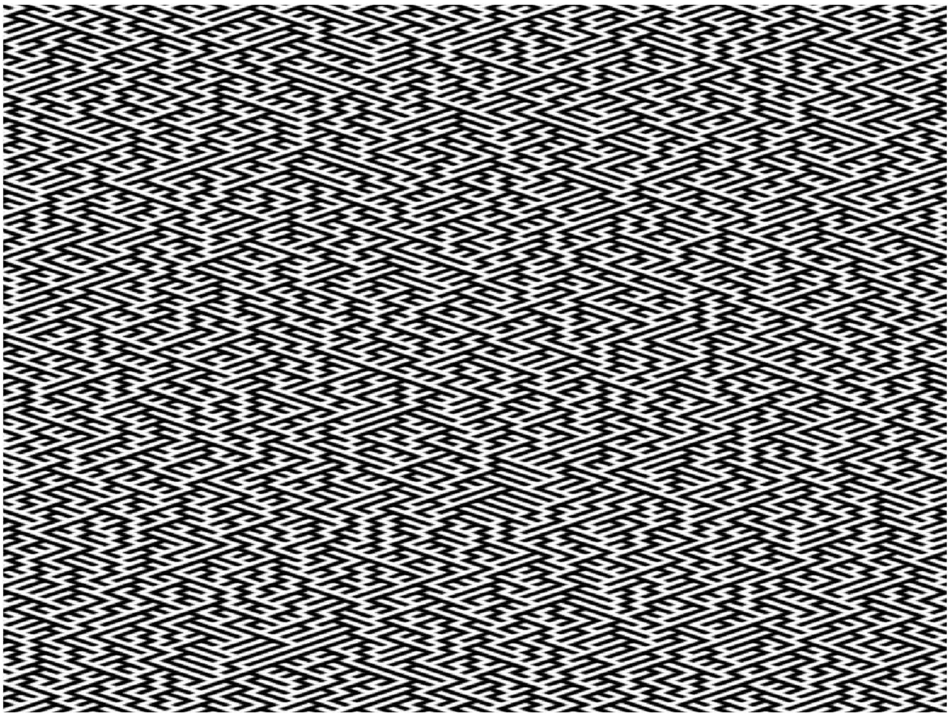

Figure 3.9: A rescaled discrete web and its dual.

The following theorem, which is the main result of this chapter, implies in particular the convergence in (3.3). See Figure 3.9 for an illustration.

Theorem 3.18 (Approximation of the double Brownian web) *Let $\mathcal{U}$ be the set of open paths in an arrow configuration and let $\mathcal{U}^*$ be the set of downward open paths in the associated dual arrow configuration. Then*

$$\mathbb{P}\big[\theta_\varepsilon(\overline{\mathcal{U}},\overline{\mathcal{U}}^*) \in \cdot\,\big] \underset{\varepsilon\to 0}{\Longrightarrow} \mathbb{P}\big[(\mathcal{W}, \hat{\mathcal{W}}) \in \cdot\,\big], \tag{3.21}$$

where $\Rightarrow$ denotes weak convergence of probability laws on the space $\mathcal{K}(\Pi^\uparrow) \times \mathcal{K}(\Pi^\downarrow)$, and $(\mathcal{W}, \hat{\mathcal{W}})$ is the double Brownian web.

Proof Fix countable dense sets $\mathcal{D},\hat{\mathcal{D}}\subset\mathbb{R}^2$ and define $(\mathcal{W},\hat{\mathcal{W}})$ as in (3.20). It suffices to prove convergence along any sequence ε_n of positive constants tending to zero. It follows from Proposition 3.16 (compare Lemma 3.14) that the laws

$$\big(\mathbb{P}\big[\theta_{\varepsilon_n}(\overline{\mathcal{U}},\overline{\mathcal{U}}^*)\in\cdot\,\big]\big)_{n\geq 1}$$

are tight, so by going to a subsequence, we may assume that they converge to some limit law $\mathbb{P}[(\mathcal{V},\hat{\mathcal{V}})\in\cdot\,]$. By Lemma 2.2, it suffices to show that each such subsequential limit is equal to $\mathbb{P}[(\mathcal{W},\hat{\mathcal{W}})\in\cdot\,]$.

As in the proof of Proposition 3.13, for each $z\in\mathcal{D}$, we choose $z_n\in\mathbb{Z}^2_{\rm even}$ such that $\theta_{\varepsilon_n}(z_n)\to z$, and for each $z\in\hat{\mathcal{D}}$, we choose $z^n\in\mathbb{Z}^2_{\rm odd}$ such that $\theta_{\varepsilon_n}(z^n)\to z$. For each $z\in\mathcal{D}$ and $n\geq 1$, we let $R^n_z\in\mathcal{U}$ be the unique forward open path starting at z_n, we let $\hat{R}^n_z\in\mathcal{U}^*$ be the unique downward open path starting at z^n, and we let

$$\pi^n_z:=\theta_{\varepsilon_n}(R^n_z)\quad\text{and}\quad\hat{\pi}^n_z:=\theta_{\varepsilon_n}(\hat{R}^n_z)$$

denote the associated diffusively rescaled paths. In the proof of Proposition 3.13, we have shown that

$$\begin{aligned}&\mathbb{P}\big[\big((\pi^n_z)_{z\in\mathcal{D}},(\tau(\pi^n_{z_1},\pi^n_{z_2}))_{(z_1,z_2)\in\mathcal{D}^2}\big)\in\cdot\,\big]\\&\quad\underset{n\to\infty}{\Longrightarrow}\mathbb{P}\big[\big((\pi_z)_{z\in\mathcal{D}},(\tau(\pi_{z_1},\pi_{z_2}))_{(z_1,z_2)\in\mathcal{D}^2}\big)\in\cdot\,\big],\end{aligned}$$

and similarly for the collection of downward paths. We argued there that going to a subsequence if necessary and using Skorohod's representation theorem, we can couple our random variables such that

$$\pi^n_z\underset{n\to\infty}{\longrightarrow}\pi_z\ \text{a.s.}\quad\text{and}\quad\tau(\pi^n_{z_1},\pi^n_{z_2})\underset{n\to\infty}{\longrightarrow}\tau(\pi_{z_1},\pi_{z_2})\ \text{a.s.}$$

for all $z,z_1,z_2\in\mathcal{D}$, and likewise for downward paths. We can extend this argument to obtain that moreover

$$\theta_{\varepsilon_n}(\overline{\mathcal{U}},\overline{\mathcal{U}}^*)\underset{n\to\infty}{\longrightarrow}(\mathcal{V},\hat{\mathcal{V}})\quad\text{a.s.}$$

in the topology on $\mathcal{K}(\Pi^\uparrow)\times\mathcal{K}(\Pi^\downarrow)$ for some random compact sets $\mathcal{V}\subset\Pi^\uparrow$ and $\hat{\mathcal{V}}\subset\Pi^\downarrow$. We will show that for this particular coupling, $(\mathcal{V},\hat{\mathcal{V}})=(\mathcal{W},\hat{\mathcal{W}})$ a.s., where the latter is defined in terms of $(\pi_z)_{z\in\mathcal{D}}$ and $(\hat{\pi}_z)_{z\in\hat{\mathcal{D}}}$. This shows that all subsequential limit laws are the same and hence by Lemma 2.2 that the original sequence converges.

By symmetry between forward and dual webs, it suffices to prove that $\mathcal{V}=\mathcal{W}$. We will prove that $\mathcal{W}_-\subset\mathcal{V}\subset\mathcal{W}_+$, where $\mathcal{W}_-$ and $\mathcal{W}_+$ are defined as in Theorem 3.15. Since $\mathcal{W}=\mathcal{W}_-=\mathcal{W}_+$, the claim then follows.

Since $\mathcal{V}$ is closed, to prove that $\mathcal{W}_- \subset \mathcal{V}$, it suffices to prove that $\pi_z \in \mathcal{V}$ for all $z \in \mathcal{D}$. Recalling Lemma 2.15, this is obvious since $\pi^n_z \in \theta_{\varepsilon_n}(\mathcal{U})$ for all n while $\pi^n_z \to \pi_z$ a.s. and $\theta_{\varepsilon_n}(\mathcal{U}) \to \mathcal{V}$ a.s.

To prove that $\mathcal{V} \subset \mathcal{W}_+$, we need to show that paths $\pi \in \mathcal{V}$ do not enter wedges of $(\hat{\pi}_z)_{z\in\hat{\mathcal{D}}}$. By Lemma 2.15, for each $\pi \in \mathcal{V}$, there exist $\pi^n \in \theta_{\varepsilon_n}(\mathcal{U})$ such that $\pi^n \to \pi$. To see that π does not enter any wedge $W(\hat{\pi}_{z_1}, \hat{\pi}_{z_2})$ of $(\hat{\pi}_z)_{z\in\hat{\mathcal{D}}}$, we use that for each n, the path π^n does not enter the wedge $W(\hat{\pi}^n_{z_1}, \hat{\pi}^n_{z_2})$. By our assumptions, the discrete paths $\hat{\pi}^n_{z_i}$ $(i = 1, 2)$ converge a.s. to $\hat{\pi}_{z_i}$ $(i = 1, 2)$ and moreover their meeting times converge a.s., so we can use Lemma 3.12 to conclude that π does not enter $W(\hat{\pi}_{z_1}, \hat{\pi}_{z_2})$. ∎

Exercise 3.19 (Shortened paths) *Let $\mathcal{W}$ be a Brownian web. Show that almost surely, for each $\pi \in \mathcal{W}$ and $s \geq \sigma_\pi$, the path π' defined by $\sigma_{\pi'} := s$ and $\pi'(t) := \pi(t)$ $(t \geq s)$ satisfies $\pi' \in \mathcal{W}$.*

3.7 Continuous time random walks

In this section we return to the one-dimensional nearest-neighbour voter model and its dual system of coalescing random walks, introduced in Sections 1.1–1.3. We let ω denote the graphical representation of a nearest-neighbour voter model on $\mathbb{Z}$, i.e., ω is a Poisson point set with intensity measure as in (1.2), where $\Lambda = \mathbb{Z}$ and $p(i, j)$ is the nearest-neighbour kernel defined in (1.9). Elements of ω are of the form $(\mathtt{vot}_{ji}, t)$, where $\mathtt{vot}_{ji}$ is a nearest-neighbour voter map (i.e., $|i - j| = 1$) that has to be applied at time t. In pictures, we draw space $\mathbb{Z}$ horizontally, time vertically, and we represent an element $(\mathtt{vot}_{ji}, t)$ of the graphical representation ω by an arrow from the space-time point (j, t) to (i, t) of the form: $\longrightarrow\!\blacksquare$.

Slightly deviating from our conventions in Chapter 1, we define

$$\omega^{\downarrow} := \big\{ \big(\mathtt{rw}_{ij}, t\big) : \big(\mathtt{vot}_{ji}, t\big) \in \omega \big\}, \tag{3.22}$$

i.e., this is the graphical representation of the dual system of coalescing random walks defined in (1.11), except that we have not reversed time. In pictures, we represent an element $(\mathtt{rw}_{ij}, t)$ of the graphical representation $\omega^{\downarrow}$ by an arrow from (i, t) to (j, t) of the form: $\blacksquare\!\longrightarrow$. In other words, $\omega^{\downarrow}$ is obtained from ω by reversing the direction of all arrows, but not turning the picture upside down, as we did earlier.

Let $\mathbb{Z} + \frac{1}{2} := \{i + \frac{1}{2} : i \in \mathbb{Z}\}$. We can define voter maps $\mathtt{vot}_{ji}$ and coalescing random walk maps $\mathtt{rw}_{ij}$ with $i, j \in \mathbb{Z} + \frac{1}{2}$ in the same way as we did for $i, j \in \mathbb{Z}$. The difference is that these maps now act on configurations in $\{0,1\}^{\mathbb{Z}+\frac{1}{2}}$ instead of $\{0,1\}^{\mathbb{Z}}$. In pictures, we represent these maps by arrows just as we are used to. With these conventions, we define

$$\begin{aligned}\omega^{\uparrow} := &\big\{\big(\mathtt{rw}_{i-\frac{1}{2},i+\frac{1}{2}},t\big) : \big(\mathtt{vot}_{i-1,i},t\big)\in\omega\big\}\\ &\cup\big\{\big(\mathtt{rw}_{i+\frac{1}{2},i-\frac{1}{2}},t\big) : \big(\mathtt{vot}_{i+1,i},t\big)\in\omega\big\}.\end{aligned} \tag{3.23}$$

The reason behind this definition is that if we apply the voter map $\mathtt{vot}_{i-1,i}$ to a configuration that has a boundary between the ones and zeros at the position $i-\frac{1}{2}$, then this boundary moves to $i+\frac{1}{2}$. Likewise, an application of $\mathtt{vot}_{i+1,i}$ moves a boundary from $i+\frac{1}{2}$ to $i-\frac{1}{2}$. In particular, if we start the voter model with each site in a different colour, then the boundaries between these colours evolve as coalescing random walks described by the graphical representation $\omega^{\uparrow}$.

We next define downward open paths in the graphical representation $\omega^{\downarrow}$ and upward open paths in the graphical representation $\omega^{\uparrow}$. A technical issue that we have to deal with is that because we work in continuous time, these open paths will have jumps. We use the formalism of cadlag paths described in Section 2.9. Recall that $\Pi_{\rm S}(\overline{\mathbb{R}})$, defined there, is the space of cadlag paths in $\overline{\mathbb{R}}$, equipped with a topology that (at least for paths defined on fixed domains) corresponds to the Skorohod topology. We let

$$\Pi_{\rm S}^{\uparrow} := \big\{\pi\in\Pi_{\rm S}(\overline{\mathbb{R}}) : \tau_\pi=\infty\big\} \quad\text{and}\quad \Pi_{\rm S}^{\downarrow} := \big\{\pi\in\Pi_{\rm S}(\overline{\mathbb{R}}) : \sigma_\pi=-\infty\big\}$$

denote the spaces of cadlag half-infinite upward and downward paths. We say that a path $\pi\in\Pi_{\rm S}^{\uparrow}$ is an *open upward path* if

(i) $\pi(t-),\pi(t+)\in\mathbb{Z}+\frac{1}{2}$ for all $t\in I_\pi$,

(ii) if $\pi(t+)\neq\pi(t-)$, then $(\mathtt{rw}_{\pi(t-),\pi(t+)},t)\in\omega^{\uparrow}$,

(iii) if $t>\sigma_\pi$ and $(\mathtt{rw}_{\pi(t-),j},t)\in\omega^{\uparrow}$ for some $j\in\mathbb{Z}+\frac{1}{2}$, then $\pi(t+)=j$.

Similarly, we say that a path $\pi\in\Pi_{\rm S}^{\downarrow}$ is a *open downward path* if

(i) $\pi(t-),\pi(t+)\in\mathbb{Z}$ for all $t\in I_\pi$,

(ii) if $\pi(t+)\neq\pi(t-)$, then $(\mathtt{rw}_{\pi(t+),\pi(t-)},t)\in\omega^{\downarrow}$,

(iii) if $t<\tau_\pi$ and $(\mathtt{rw}_{\pi(t+),j},t)\in\omega^{\uparrow}$ for some $j\in\mathbb{Z}$, then $\pi(t-)=j$.

See Figure 3.10 for an illustration. We let $\mathcal{U}^{\uparrow}$ and $\mathcal{U}^{\downarrow}$ denote the sets of all open upward and open downward paths, respectively. The reason for the conditions $t>\sigma_\pi$ and $t<\tau_\pi$ in point (iii) of each definition is that we want to work with compact sets of paths. If $(\mathtt{rw}_{i,j},t)\in\omega^{\uparrow}$, then for each $n\geq 1$, there exists an open upward path π_n in $\omega^\uparrow$ that starts at $(i, t + n^{-1})$, and these paths converge to a limit path π that starts at (i, t) but does not jump at its starting time. Thus, if we would require (iii) also for $t = \sigma_\pi$, then the set of open upward paths would not be closed. On the other hand, with our present definition, one can prove the following statement, that is similar to Proposition 3.1. For brevity, we skip its proof.

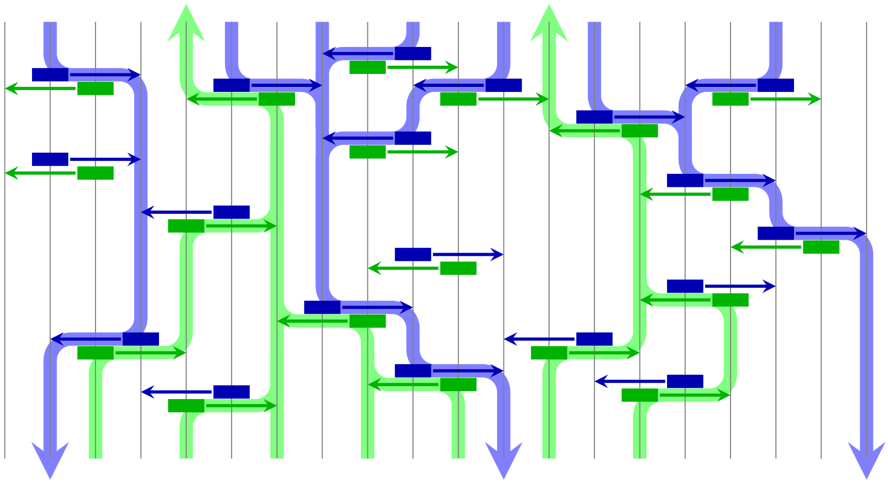

Figure 3.10: Open upward and downward paths in the graphical representation of a one-dimensional nearest-neighbour voter model.

Proposition 3.20 (Compact set of paths) *The closure $\overline{\mathcal{U}}^\uparrow$ of the random set of open upward paths $\mathcal{U}^\uparrow$ is almost surely a compact subset of $\Pi^\uparrow_{\rm S}$. Moreover, almost surely, the set $\overline{\mathcal{U}}^\uparrow \backslash \mathcal{U}^\uparrow$ consists of all paths $\pi \in \Pi^\uparrow_{\rm S}$ with either $\pi(t) = -\infty$ for all $t \in I_\pi$ or $\pi(t) = +\infty$ for all $t \in I_\pi$. An analogue statement holds for $\mathcal{U}^\downarrow$.*

In the remainder of this section, we sketch the proof of the following theorem, that is similar to Theorem 3.18.

Theorem 3.21 (Convergence to the double Brownian web) *Let $\mathcal{U}^\uparrow$ be the set of open upward paths in $\omega^\uparrow$ and let $\mathcal{U}^\downarrow$ be the set of open downward paths in $\omega^\downarrow$. Then*

$$\mathbb{P}\big[\theta_\varepsilon(\overline{\mathcal{U}}^\uparrow, \overline{\mathcal{U}}^\downarrow) \in \cdot\,\big] \underset{\varepsilon \to 0}{\Longrightarrow} \mathbb{P}\big[(\mathcal{W}, \hat{\mathcal{W}}) \in \cdot\,\big], \tag{3.24}$$

where $\Rightarrow$ denotes weak convergence of probability laws on the space $\mathcal{K}(\Pi^\uparrow_{\rm S}) \times \mathcal{K}(\Pi^\downarrow_{\rm S})$, and $(\mathcal{W}, \hat{\mathcal{W}})$ is the double Brownian web.

Proof (crude sketch) The proof is completely analogous to the proof of Theorem 3.18. Convergence of rescaled continuous-time random walks to Brownian motion is standard, so one obtains convergence of finite dimensional distributions precisely as in the discrete time setting (Proposition 3.3). Using Exercise 3.5, one sees that this convergence can be strengthened so that it also includes convergence of meeting times in the sense of (3.19). Adapting the argument of Proposition 3.16 (which itself is an adaptation of the proof of Proposition 3.6), using also Lemma 3.14, one can moreover show that if ε_n are positive constants tending to zero, then the laws

$$\big(\mathbb{P}[\theta_{\varepsilon_n}(\overline{\mathcal{U}}^{\uparrow},\overline{\mathcal{U}}^{\downarrow}) \in \cdot\,]\big)_{n\geq 1} \tag{3.25}$$

on $\mathcal{K}(\Pi^{\uparrow}_{\rm S}) \times \mathcal{K}(\Pi^{\downarrow}_{\rm S})$ are tight. For brevity, we are rather sloppy on this part. To fill in the details, one would need to work with Skorohod-equicontinuity and Theorem 2.34 to prove a tightness criterion for laws on $\mathcal{K}(\Pi^{\uparrow}_{\rm S})$, which is similar to Proposition 2.33, but the conditions of which are weaker, since Skorohod-equicontinuity is a weaker concept than the usual equicontinuity.

To complete the proof, one needs to show that if the laws in (3.25) converge weakly along a subsequence to a limit law on $\mathcal{K}(\Pi^{\uparrow}_{\rm S}) \times \mathcal{K}(\Pi^{\downarrow}_{\rm S})$, then this limit law must be the law of a double Brownian web. We fix a countable dense set $\mathcal{D} \subset \mathbb{R}^2$ and for each $z \in \mathcal{D}$, we choose $z_n \in \mathbb{Z} \times \mathbb{R}$ such that $\theta_{\varepsilon_n}(z_n) \to z$. Since the points z_n are deterministic, at each z_n there almost surely start a unique upward and downward open path. By what we have already proved, these paths converge in law to coalescing Brownian motions, and also their meeting times converge in law. Using Skorohod's representation, we can find a coupling for which the convergence is almost sure. Let $(\pi_z)_{z\in\mathcal{D}}$ and $(\hat{\pi}_z)_{z\in\mathcal{D}}$ be the almost sure limits of the chosen upward and downward open paths, respectively. Using the argument of Lemma 3.12, we see that paths in $(\pi_z)_{z\in\mathcal{D}}$ do not enter wedges of $(\hat{\pi}_z)_{z\in\mathcal{D}}$. Theorem 3.15 now tells us that the sets $\mathcal{W}_-$ and $\mathcal{W}_+$ defined there satisfy $\mathcal{W}_- = \mathcal{W}_+$, and both are distributed as a Brownian web. Letting $\mathcal{V}$ denote the almost sure limit of $\theta_{\varepsilon_n}(\overline{\mathcal{U}}^{\uparrow})$, as in the proof of Theorem 3.18, it then suffices to show that $\mathcal{W}_- \subset \mathcal{V} \subset \mathcal{W}_+$.

The inclusion $\mathcal{W}_- \subset \mathcal{V}$ is straightforward. If we would know that $\mathcal{V} \subset \Pi^{\uparrow}$ almost surely, then the inclusion $\mathcal{V} \subset \mathcal{W}_+$ would follow by exactly the same argument as in the proof of Theorem 3.18, but for the moment we have only indicated how one can obtain the weaker statement that $\mathcal{V} \subset \Pi^{\uparrow}_{\rm S}$, where $\Pi^{\uparrow}_{\rm S}$ is the space of cadlag upward paths. Therefore, one way to complete the argument is to show that $\mathcal{V} \subset \Pi^{\uparrow}$ almost surely, which can probably be done by using arguments similar to those used in the proof of Proposition 3.6. An alternative argument, that is probably easier, is to strengthen Theorem 3.15

by showing that the set $\mathcal{W}_+$ there can be replaced by the (a priori larger) set

$$\mathcal{W}'_+ := \big\{\pi \in \Pi^\uparrow_{\mathrm{S}} : \pi \text{ does not enter wedges of } (\hat{\pi}_z)_{z\in\mathcal{D}}\big\}.$$

It seems this should follow from the same arguments as those used in the proof of Theorem 3.15, but because of time restrictions, we skip the details.

Though there is no doubt among experts in the field that Theorem 3.21 holds, it seems nobody so far has bothered to write down a detailed proof. In fact, the only published paper that I am aware of that shows convergence of collections of cadlag paths is [EFS17], which however deals with the Brownian net instead of the web and shows convergence for a different approximating model. ∎

3.8 Some historical notes

The Brownian web originated from Arratia's PhD thesis [Arr79] and a subsequential unfinished manuscript [Arr81]. The topic remained dormant until the work of Tóth and Werner [TW98] who used the Brownian web to study a form of one-dimensional self-repellent random walk. They classified all types of special points. Together with Soucaliuc [STW00] they also proved that forward and dual paths interact through Skorohod reflection. Fontes, Isopi, Newman and Stein got interested in the Brownian web motivated by a one-dimensional model in mathematical physics [FINS01], which led Fontes, Isopi, Newman and Ravishankar [FINR04] to study this object in more detail. In particular, they were the first to give the Brownian web its name, view it as a compact set of paths, and prove convergence with respect to the Hausdorff topology. Wedges were first introduced in the framework of the Brownian net in [SS08]. A more detailed account of the history of the Brownian web can be found in [SSS16].

Chapter 4

Properties of the Brownian web

4.1 The coalescing point set

Let $(\mathcal{W}, \hat{\mathcal{W}})$ be a double Brownian web, i.e., a Brownian web and its dual. For each $s, t \in \mathbb{R}$ with $s \leq t$ and closed $A \subset \overline{\mathbb{R}}$, we define

$$\begin{aligned}
\mathcal{X}_{s,t}(A) &:= \big\{x \in \overline{\mathbb{R}} : \exists \hat{\pi} \in \hat{\mathcal{W}}(x,t) \text{ s.t. } \hat{\pi}(s) \in A\big\}, \\
\mathcal{Y}_{s,t}(A) &:= \big\{\pi(t) : \pi \in \mathcal{W}(A \times \{s\})\big\}, \\
\hat{\mathcal{X}}_{t,s}(A) &:= \big\{x \in \overline{\mathbb{R}} : \exists \pi \in \mathcal{W}(x,s) \text{ s.t. } \pi(t) \in A\big\}, \\
\hat{\mathcal{Y}}_{t,s}(A) &:= \big\{\hat{\pi}(s) : \hat{\pi} \in \hat{\mathcal{W}}(A \times \{t\})\big\}.
\end{aligned} \tag{4.1}$$

We can think of the maps $(\mathcal{X}_{s,t})_{s \leq t}$ as a continuum analogue of the stochastic flow $(\mathbf{X}_{s,t})_{s \leq t}$ defined in Section 1.1. Let us fix closed sets $A, B \subset \overline{\mathbb{R}}$ and define, in analogy with (1.5),

$$A_t := \mathcal{X}_{0,t}(A) \quad \text{and} \quad B_t := \mathcal{Y}_{0,t}(B) \qquad (t \geq 0). \tag{4.2}$$

Then we can think of the process $(A_t)_{t \geq 0}$ as of some sort of continuum version of the voter model and similarly, we can think of $(B_t)_{t \geq 0}$ as a continuum version of coalescing random walks, i.e., this process should correspond to coalescing Brownian motions. We call $(A_t)_{t \geq 0}$ the *continuum voter model* and $(B_t)_{t \geq 0}$ the *coalescing point set*. In Section 4.5 below, we will prove that $(A_t)_{t \geq 0}$ and $(B_t)_{t \geq 0}$ are indeed Markov processes. In the present section, we start by proving some elementary properties of the maps $(\mathcal{X}_{s,t})_{s \leq t}$ and $(\mathcal{Y}_{s,t})_{s \leq t}$.

Lemma 4.1 (Additivity) *One has $\mathcal{X}_{s,t}(A) \in \mathcal{K}(\overline{\mathbb{R}})$ for each $A \in \mathcal{K}(\overline{\mathbb{R}})$ and $s \leq t$. Moreover,*

$$\mathcal{X}_{s,t}(A \cup B) = \mathcal{X}_{s,t}(A) \cup \mathcal{X}_{s,t}(B) \qquad \big(A, B \in \mathcal{K}(\overline{\mathbb{R}}),\ s \leq t\big). \tag{4.3}$$

Analogue statements hold with $\mathcal{X}_{s,t}$ replaced by $\mathcal{Y}_{s,t}$.

Proof If $x_n \in \mathcal{X}_{s,t}(A)$ satisfy $x_n \to x$ for some $x \in \overline{\mathbb{R}}$, then there exist $\hat{\pi}_n \in \hat{\mathcal{W}}(x_n, t)$ such that $\hat{\pi}_n(s) \in A$. Since $\hat{\mathcal{W}}$ is compact, by going to a subsequence, we can assume that $\hat{\pi}_n \to \hat{\pi}$ for some $\hat{\pi} \in \hat{\mathcal{W}}$. Then $\hat{\pi} \in \hat{\mathcal{W}}(x, t)$. Since A is closed moreover $\pi(t) \in A$. Together, the last two observations imply that $x \in \mathcal{X}_{s,t}(A)$, proving that $\mathcal{X}_{s,t}(A)$ is closed. The proof for $\mathcal{Y}_{s,t}$ is the same. Formula (4.3) follows immediately from the definitions of $\mathcal{X}_{s,t}$ and $\mathcal{Y}_{s,t}$. ∎

The following lemma is the continuum analogue of the duality relation (1.15), for the moment restricted to coalescing Brownian motions without branching or deaths.

Lemma 4.2 (Continuum duality) *For each $A, B \in \mathcal{K}(\overline{\mathbb{R}})$ and $s, t \in \mathbb{R}$ with $s \leq t$, one has*

$$1_{\{\mathcal{X}_{s,t}(A) \cap B \neq \emptyset\}} = 1_{\{A \cap \hat{\mathcal{Y}}_{t,s}(B) \neq \emptyset\}}.$$

Proof This is a straightforward consequence of our definitions, since

$$\mathcal{X}_{s,t}(A) \cap B \neq \emptyset \;\Leftrightarrow\; \exists \hat{\pi} \in \hat{\mathcal{W}}\big(B \times \{t\}\big) \text{ s.t. } \hat{\pi}(s) \in A \;\Leftrightarrow\; \hat{\mathcal{Y}}_{t,s}(B) \cap A \neq \emptyset.$$

∎

Our next aim is to show that the coalescing point set $(B_t)_{t \geq 0}$ defined in (4.2) "comes down from infinity" in the sense that regardless of the initial state B, for each $t > 0$ the set B_t is locally finite. Since clearly $B \subset B'$ implies $\mathcal{Y}_{0,t}(B) \subset \mathcal{Y}_{0,t}(B')$, it suffices to prove the claim for $B = \overline{\mathbb{R}}$. Roughly speaking, the following result says that if we start particles performing coalescing Brownian motions from each point on the real line, then at each positive time there are only locally finitely many particles left. This is sometimes described by saying that coalescing Brownian motions *come down from infinity.*

Proposition 4.3 (Density of the coalescing point set) *One has*

$$\mathbb{E}\big[\big|\mathcal{Y}_{0,t}(\overline{\mathbb{R}}) \cap [a,b]\big|\big] = \frac{b-a}{\sqrt{\pi t}} \qquad (a, b \in \mathbb{R},\ a < b,\ t > 0).$$

Proof We first calculate the probability that $\mathcal{Y}_{0,t}(\mathbb{R}) \cap [a,b] \neq \emptyset$. We construct $(\mathcal{W}, \hat{\mathcal{W}})$ from collections $(\pi_z)_{z \in \mathcal{D}}$ and $(\hat{\pi}_z)_{z \in \hat{\mathcal{D}}}$ of forward and downward coalescing Brownian motions, so that paths in $(\pi_z)_{z \in \mathcal{D}}$ do not enter wedges of $(\hat{\pi}_z)_{z \in \hat{\mathcal{D}}}$ and vice versa. We choose $\hat{\mathcal{D}}$ such that $(a,t), (b,t) \in \hat{\mathcal{D}}$. Let

$$\tau_{a,b} = \tau\big(\hat{\pi}_{(a,t)}, \hat{\pi}_{(b,t)}\big)$$

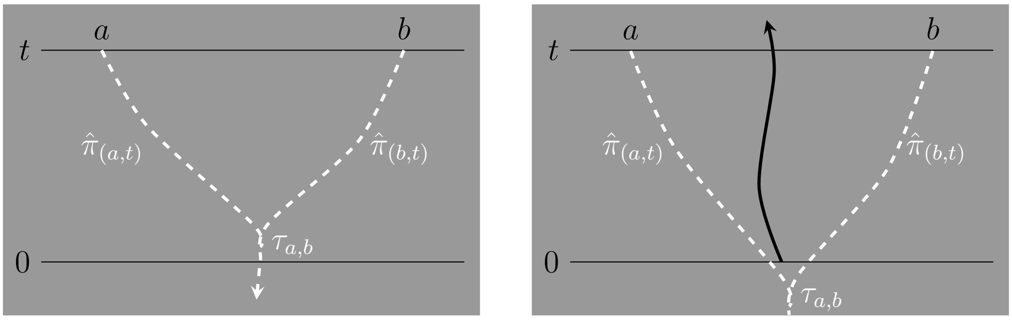

Figure 4.1: Illustration of formula (4.4). If $\tau_{a,b} > 0$, then no path in $\mathcal{W}$ starting at time 0 can pass through (a,b) at time t. On the other hand, if $\tau_{a,b} < 0$, then any path in $\mathcal{W}$ starting at time 0 between $\hat{\pi}_{(a,t)}(0)$ and $\hat{\pi}_{(b,t)}(0)$ must pass through $[a,b]$ at time t.

be the first meeting time of the downward paths started at (a,t) and (b,t). We claim that (see Figure 4.1)

$$\mathcal{Y}_{0,t}(\mathbb{R}) \cap (a,b) \neq \emptyset \quad \text{implies} \quad \tau_{a,b} \leq 0 \quad \text{implies} \quad \mathcal{Y}_{0,t}(\mathbb{R}) \cap [a,b] \neq \emptyset. \tag{4.4}$$

Indeed, if $\tau_{a,b} > 0$, then the paths $\hat{\pi}_{(a,t)}$ and $\hat{\pi}_{(b,t)}$ form a wedge that prevents paths in $\mathcal{W}$ starting at time zero from passing between (a,t) and (b,t), proving the first implication. On the other hand, if $\tau_{a,b} \leq 0$, then for each time $s > 0$ we can find some x such that $\hat{\pi}_{(a,t)}(s) < x < \hat{\pi}_{(b,t)}(s)$. The web $\mathcal{W}$ must contain a path π starting at (x,s) and since such a path cannot cross the downward paths $\hat{\pi}_{(a,t)}$ and $\hat{\pi}_{(b,t)}$, it must satisfy $a \leq \pi(t) \leq b$. We can construct such a path π^s with starting time s for each $s > 0$, so using the compactness of $\mathcal{W}$, we see that $\mathcal{W}$ must also contain a path π^0 starting at time zero such that $a \leq \pi(t) \leq b$, proving the second implication.

The difference $(B_1(s) - B_2(s))_{s\geq 0}$ of two Brownian motions is equally distributed with $(\sqrt{2}B(s))_{s\geq 0}$, where $(B(s))_{t\geq 0}$ is a single Brownian motion. Therefore, using the reflection principle,

$$\begin{aligned}\mathbb{P}\big[\tau_{a,b} \leq 0\big] &= \mathbb{P}\big[\sup_{0\leq s\leq t}\big(B_2(s) - B_1(s)\big) \leq b-a\big] \\ &= \mathbb{P}\big[\sup_{0\leq s\leq t} B(s) \leq \frac{b-a}{\sqrt{2}}\big] = \frac{1}{\sqrt{2\pi t}}\int_{-\frac{b-a}{\sqrt{2}}}^{\frac{b-a}{\sqrt{2}}} e^{-x^2/2t}\,\mathrm{d}x.\end{aligned}$$

In particular, this implies that

$$\mathbb{P}[x \in \mathcal{Y}_{0,t}(\mathbb{R})] = \lim_{\varepsilon\to 0}\mathbb{P}\big[\mathcal{Y}_{0,t}(\mathbb{R}) \cap (x-\varepsilon, x+\varepsilon) \neq \emptyset\big] = 0 \qquad (x \in \mathbb{R},\ t > 0),$$

and hence

$$\mathbb{P}\big[\mathcal{Y}_{0,t}(\mathbb{R}) \cap (a,b) \neq \emptyset\big] = \mathbb{P}\big[\mathcal{Y}_{0,t}(\mathbb{R}) \cap [a,b] \neq \emptyset\big] = \mathbb{P}[\tau_{a,b} \leq 0].$$

Now

$$\begin{aligned}\mathbb{E}\big[\big|\mathcal{Y}_{0,t}(\mathbb{R}) \cap [0,1]\big|\big] &= \lim_{n\to\infty} \sum_{i=1}^{2^n} \mathbb{P}\big[\mathcal{Y}_{0,t}(\mathbb{R}) \cap [(i-1)2^{-n}, i2^{-n}] \neq \emptyset\big] \\ = \lim_{\varepsilon\to 0} \varepsilon^{-1} \frac{1}{\sqrt{2\pi t}} \int_{-\varepsilon/\sqrt{2}}^{\varepsilon/\sqrt{2}} e^{-x^2/2t}\, \mathrm{d}x &= \frac{1}{\sqrt{\pi t}}.\end{aligned}$$

A similar formula holds for the expectation of $\big|\mathcal{Y}_{0,t}(\mathbb{R}) \cap [0,r]\big|$ for any $r > 0$ and the general result follows by translation invariance. ∎

We conclude this section with some useful consequences of Proposition 4.3. In the following lemma, we let $\Pi^{\updownarrow} := \Pi^{\uparrow} \cap \Pi^{\downarrow}$ denote the space of all bi-infinite paths and we let $\Pi^{\updownarrow}(\mathbb{R}) := \Pi^{\updownarrow} \cap \Pi(\mathbb{R})$ denote the space bi-infinite paths with values in $\mathbb{R}$ (as opposed to $\overline{\mathbb{R}}$).

Lemma 4.4 (No bi-infinite paths) *Let $\mathcal{W}$ be a Brownian web. Then $\mathcal{W} \cap \Pi^{\updownarrow}(\mathbb{R}) = \emptyset$ a.s.*

Proof We start by observing that

$$\mathbb{P}\big[\mathcal{W} \cap \Pi^{\updownarrow}(\mathbb{R}) \neq \emptyset\big] \leq \lim_{n\to\infty} \mathbb{P}\big[\exists \pi \in \mathcal{W} \text{ s.t. } \sigma_\pi = -\infty,\ \pi(0) \in [-n,n]\big],$$

where by Lemma 3.8 the inequality is in fact an equality. Now Proposition 4.3 gives

$$\begin{aligned}&\mathbb{P}\big[\exists \pi \in \mathcal{W} \text{ s.t. } \sigma_\pi = -\infty,\ \pi(0) \in [-n,n]\big] \\ &\quad\leq \lim_{t\to\infty} \mathbb{P}\big[\exists \pi \in \mathcal{W} \text{ s.t. } \sigma_\pi \leq -t,\ \pi(0) \in [-n,n]\big] = \lim_{t\to\infty} \frac{2n}{\sqrt{\pi t}} = 0,\end{aligned}$$

Here again, with a bit of extra work, one can show that the inequality is in fact an equality, but we do not presently need this. ∎

Lemma 4.5 (Coalescence of paths) *Almost surely, for all paths $\pi, \pi' \in \mathcal{W}$, if $\pi(t) = \pi'(t)$ for some $t > \sigma_\pi \vee \sigma_{\pi'}$, then $\pi(u) = \pi'(u)$ for all $u \geq t$.*

Proof By Lemma 3.8, it suffices to prove the statement under the additional assumption that $\pi(t) = \pi'(t) \in \mathbb{R}$. Let $\mathcal{T} \subset \mathbb{R}$ be countable and dense. If $t > \sigma_\pi \vee \sigma_{\pi'}$, then there exist $r, s \in \mathcal{T}$ with $\sigma_\pi \vee \sigma_{\pi'} < r < s \leq t$ and the paths obtained from π and π' by cutting off the piece before time r are also paths in the Brownian web. Therefore, it suffices to prove for deterministic $r < s$ that if two paths $\pi, \pi' \in \mathcal{W}$ with $\sigma_\pi = \sigma_{\pi'} = r$ satisfy $\pi(t) = \pi'(t)$ for some $t \geq s$, then $\pi(u) = \pi'(u)$ for all $u \geq t$.

By Proposition 4.3, the random set

$$A := \big\{\pi(s) : \pi \in \mathcal{W},\ \sigma_\pi = r\big\} \cap \mathbb{R}$$

is locally finite. We claim that for each $x \in A$, there exists a unique path $\pi_{(x,s)} \in \mathcal{W}(x, s)$, and conditional on A, the collection of paths

$$(\pi_{(x,s)})_{x \in A}$$

is distributed as coalescing Brownian motions. Indeed, this follows from the fact (which can easily be proved using discrete approximation) that restrictions of the Brownian web to disjoint parts of space-time are independent. As a result, the random set A is independent of $\mathcal{W}\big(\mathbb{R} \times [s, \infty)\big)$, so after we condition on A, paths started from a countable collection of fixed points (x, s) with $x \in \mathbb{R}$ will be distributed as coalescing Brownian motions.

The statement we want to prove now follows from the fact that if two coalescing Brownian motions meet at some random time, then they coalesce, i.e., the two paths are equal from that time onwards. ∎

Lemma 4.6 (Minimal and maximal paths) *Let $\mathcal{W}$ be a Brownian web. Then almost surely, for each $z = (x, s) \in \mathbb{R}^2$, there exist paths $\pi_z^-, \pi_z^+ \in \mathcal{W}(z)$ such that $\pi_z^-(t) \leq \pi(t) \leq \pi^+(t)$ for all $\pi \in \mathcal{W}(z)$ and $t \geq 0$. If $x_n^- < x < x_n^+$ satisfy $x_n^\pm \to x$ and $\pi_{(n)}^\pm \in \mathcal{W}(x_n^\pm)$, then*

$$\pi_{(n)}^\pm \underset{n \to \infty}{\longrightarrow} \pi_z^\pm \tag{4.5}$$

in the topology on $\Pi^\uparrow$.

Proof By symmetry, it suffices to prove the statements for π_z^+. Let $x < x_n^+$ satisfy $x_n^+ \to x$ and choose $\pi_{(n)}^+ \in \mathcal{W}(x_n^+)$. By the compactness of $\mathcal{W}$, the set $\{\pi_{(n)}^+ : n \geq 1\}$ is precompact in the topology on $\Pi^\uparrow$. Let π_z^+ be any subsequential limit of the sequence $\big(\pi_{(n)}^+\big)_{n \geq 1}$. Then clearly $\pi \in \mathcal{W}(z)$. By Lemma 4.5, each $\pi \in \mathcal{W}(z)$ satisfies $\pi(t) \leq \pi_{(n)}^+(t)$ for all $n \geq 1$ and $t \geq 0$, so taking the limit, we see that $\pi(t) \leq \pi_z^+(t)$ for all $t \geq 0$. This proves that the set $\mathcal{W}(z)$ has a maximal element π_z^+. Such a maximal element is clearly unique, so the sequence $\big(\pi_{(n)}^+\big)_{n \geq 1}$ has a unique cluster point, which is π_z^+. Since $\{\pi_{(n)}^+ : n \geq 1\}$ is precompact, it follows that the sequence $\big(\pi_{(n)}^+\big)_{n \geq 1}$ converges to π_z^+. ∎

The following lemma says that if a sequence of paths in the web converges (such as for example in (4.5)), then this convergence actually takes place in a rather strong sense.

Lemma 4.7 (Strong convergence of paths) *Let $\mathcal{W}$ be a Brownian web. Then almost surely, for all $\pi_n, \pi \in \mathcal{W}$ such that $\pi_n \to \pi$, there exist times $t_n > \sigma_{\pi_n} \vee \sigma_\pi$ such that $t_n \to \sigma_\pi$ and $\pi_n(t) = \pi(t)$ for all $t \geq t_n$.*

Remark We recall that the Brownian web has the property that almost surely, for any countable dense set $\mathcal{D} \subset \mathbb{R}^2$, one has $\mathcal{W} = \overline{\mathcal{W}(\mathcal{D})}$, where $\mathcal{W}(\mathcal{D})$ is called a skeleton of $\mathcal{W}$. Lemma 4.7 implies that for each path in the web $\pi \in \mathcal{W}$ and for each $\varepsilon > 0$, there exists a skeletal path $\pi' \in \mathcal{W}(\mathcal{D})$ such that $\pi(t) = \pi'(t)$ for all $t \geq \sigma_\pi + \varepsilon$.

Proof By Lemma 3.8, it suffices to prove the statement under the additional assumption that $\pi(t) \in \mathbb{R}$. Proposition 4.3 tells us that for each deterministic $s < t$, the set

$$A_{s,t} := \big\{\pi(t) : \pi \in \mathcal{W},\ \sigma_\pi \leq s\big\} \cap \mathbb{R}$$

is a.s. a locally finite subset of $\mathbb{R}$. Let $\mathcal{T}$ be a countable dense subset of $\mathbb{R}$. Then almost surely, $A_{s,t}$ is locally finite for all $s, t \in \mathcal{T}$ with $s < t$. Now if $\pi_n, \pi \in \mathcal{W}$ satisfy $\pi_n \to \pi$, then for each $s, t \in \mathcal{T}$ with $\sigma_\pi < s < t$, we have for n sufficiently large that $\sigma_{\pi_n} < s$ and hence $\pi_n(t), \pi(t) \in A_{s,t}$. Since $\pi_n(t) \to \pi(t)$ and since $A_{s,t}$ is locally finite, it follows that $\pi_n(t) = \pi(t)$ for n sufficiently large. By Lemma 4.5, $\pi_n(t) = \pi(t)$ implies $\pi_n(u) = \pi(u)$ for all $u \geq t$. Since $\mathcal{T}$ is dense, we can choose t as close to σ_π as we wish, and hence the statement of the lemma follows. ∎

Exercise 4.8 *Show that formula (4.4) can be strengthened (for deterministic a, b, and t) in the sense that $\mathcal{Y}_{0,t}(\mathbb{R}) \cap [a, b] \neq \emptyset$ almost surely implies $\mathcal{Y}_{0,t}(\mathbb{R}) \cap (a, b) \neq \emptyset$.*

4.2 Brownian local time

In the next section, we will study the interaction between paths in the forward and dual Brownian web. As a preparation, in the present section, we collect some well-known facts about Skorohod reflection and Brownian local time. Let $\mathcal{C} := \mathcal{C}_{[0,\infty)}(\mathbb{R})$ denote the space of continuous functions $f : [0, \infty) \to \mathbb{R}$. We set

$$\begin{aligned}
\mathcal{C}_0 &:= \big\{f \in \mathcal{C} : f_0 = 0\big\}, \\
\mathcal{C}^+ &:= \big\{f \in \mathcal{C} : f \text{ is nondecreasing}\big\}, \\
\mathcal{C}_{\text{pos}} &:= \big\{f \in \mathcal{C} : f \geq 0\big\},
\end{aligned}$$

and we write $\mathcal{C}_0^+ := \mathcal{C}^+ \cap \mathcal{C}_0$. We set

$$m_t(f) := 0 \wedge \inf_{0 \leq s \leq t} f_s \qquad \big(t \geq 0,\ f \in \mathcal{C}\big). \tag{4.6}$$

In particular, if $f \in \mathcal{C}_0$, then this is the *running minimum* of the function f. Assume that $h \in \mathcal{C}$ satisfies $h_0 \geq 0$, and let

$$g_t := h_t - m_t(h) \qquad (t \geq 0). \tag{4.7}$$

We observe that $g_t \geq 0$ and that $\psi_t := -m_t(h)$ is a nondecreasing function that increases only at times when $g_t = 0$. The following definition makes this precise. Let $g_0 \in [0, \infty)$ and $f \in \mathcal{C}_0$ be given. By definition, a solution to the *Skorohod reflection equation*

$$\mathrm{d}g_t = \mathrm{d}f_t + \mathrm{d}\psi_t \quad (t \geq 0) \tag{4.8}$$

is a pair (g, ψ) of functions $g \in \mathcal{C}_{\rm pos}$ and $\psi \in \mathcal{C}_0^+$ such that

(i) $g_t = g_0 + f_t + \psi_t$ $(t \geq 0)$,

(ii) $\displaystyle\int_0^\infty 1_{\{g_t > 0\}} \mathrm{d}\psi_t = 0$.

Here, the notation in point (ii) means that we integrate the function $t \mapsto 1_{\{g_t>0\}}$ with respect to the measure whose distribution function is ψ, i.e., this is the unique measure μ on $[0, \infty)$ such that $\mu([0, t]) = \psi_t$ $(t \geq 0)$. Condition (ii) makes the intuitive concept precise that ψ_t increases only at times when $g_t = 0$. See Figure 4.2 for an illustration. The following lemma says that solutions to (4.8) are unique and given by (4.7).

Lemma 4.9 (Skorohod reflection) *For each $g_0 \in [0, \infty)$ and $f \in \mathcal{C}_0$, the Skorohod reflection equation (4.8) has a unique solution (g, ψ) with initial condition g_0. This solution is given by*

$$g_t = \tilde f_t - m_t(\tilde f) \quad \textit{and} \quad \psi_t = -m_t(\tilde f) \qquad (t \geq 0), \tag{4.9}$$

where $\tilde f_t := g_0 + f_t$ $(t \geq 0)$.

Proof It is not hard to check that if we define g and ψ by (4.9), then (g, ψ) is a solution to the Skorohod reflection equation (4.8). To prove uniqueness, assume, conversely, that (4.8) has two solutions (g, ψ) and (g', ψ') that are not equal but which satisfy $g_0 = g'_0$. Then, by condition (i) of the definition of a solution to the Skorohod reflection equation, there must be a $u > 0$ such that $g_u \neq g'_u$. By symmetry, we can without loss of generality assume that $g_u < g'_u$. Let $s := \sup\{t \in [0, u] : g_t = g'_t\}$. By continuity, $g_s = g'_s$ and $g_t < g'_t$ for all $t \in (s, u]$. Setting $\hat f_t := f_t - f_s$ and $\hat\psi_t := \psi_t - \psi_s$, we observe that $(g, \hat\psi)$ solves the Skorohod reflection equation

$$\mathrm{d}g_t = \mathrm{d}\hat f_t + \mathrm{d}\hat\psi_t \quad (t \geq s)$$

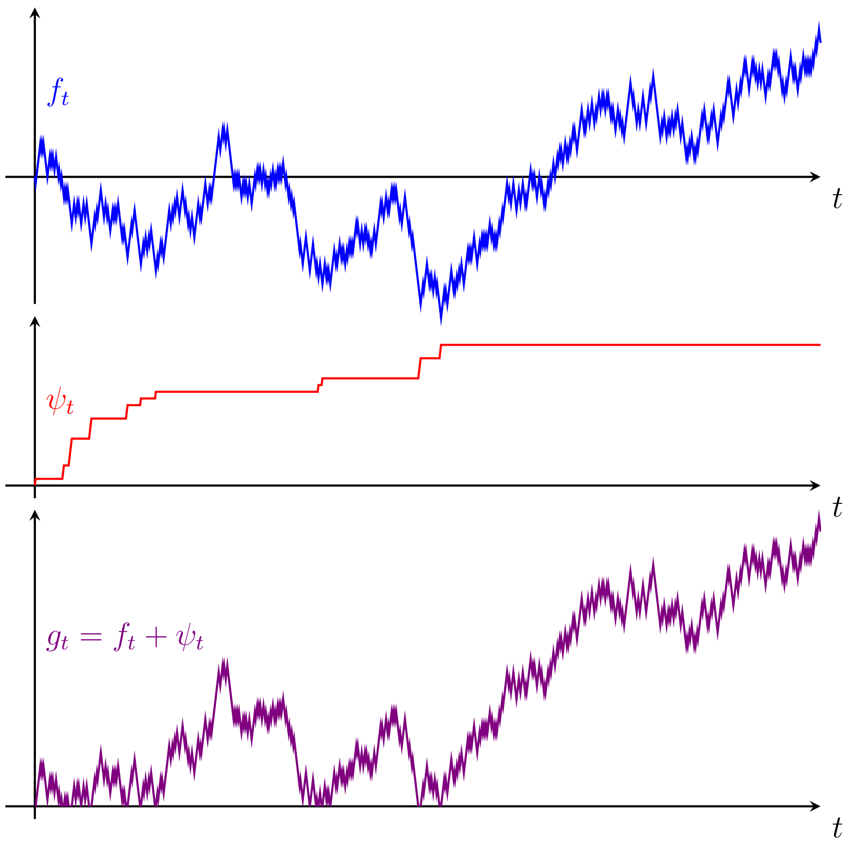

Figure 4.2: Reflected Brownian motion: the functions f, g, and ψ from Lemma 4.9 in the case that f is a Brownian path.

on the time interval $[s, \infty)$, and an analogue statement holds for (g', ψ'). In view of this, shifting the time s to zero, if uniqueness does not hold, then we can without loss of generality assume that we are in the following situation. We have two solutions (g, ψ) and (g', ψ') to the Skorohod reflection equation (4.8) with initial states $g_0 = g'_0$ for which there exists a $u > 0$ such that $g_t < g'_t$ for all $t \in (0, u]$.

Since $0 \leq g_t < g'_t$ for all $t \in (0, u]$, condition (ii) of the definition of a solution implies that $\psi'_t = 0$ for all $t \in [0, u]$. Together with the fact that $g_0 = g'_0$ and condition (i) of the definition of a solution, this implies that $g_t \geq g'_t$ for all $t \in [0, u]$, contradicting our assumption that $g_t < g'_t$ for all $0 < t \leq u$. ∎

Exercise 4.10 *Let $f \in \mathcal{C}_0$ be given and assume that (g, ψ) and (g', ψ') are solutions to the Skorohod reflection equation (4.8) with possibly different* *initial states g_0 and g'_0. Show that*

$$g_0 \leq g'_0 \quad \textit{implies} \quad g_t \leq g'_t \quad \textit{and} \quad g'_t - g_t \leq g'_0 - g_0 \qquad (t \geq 0). \tag{4.10}$$

We will especially be interested in the case that the function f from Lemma 4.9 is Brownian motion. In this case, the function g is reflected Brownian motion, and ψ is its local time at the origin. To explain this in a bit more detail, we need to take a small detour.

If $(B_t)_{t\geq 0}$ is a d-dimensional Brownian motion, then we can define a stochastic process $(\ell_t)_{t\geq 0}$ taking values in the space $\mathcal{M}(\mathbb{R}^d)$ of finite measures on $\mathbb{R}^d$ by

$$\int_{\mathbb{R}^d} \ell_t(\mathrm{d}x)\, f(x) := \int_0^t \mathrm{d}s\, f(B_s) \qquad \big(t \geq 0,\ f \in B_{\mathrm{b}}(\mathbb{R}^d)\big).$$

The random measure ℓ_t is called the *occupation local measure* of the Brownian motion $(B_t)_{t\geq 0}$. In particular

$$\ell_t(A) = \int_0^t \mathrm{d}s\, 1_A(B_s) \qquad \big(A \in \mathcal{B}(\mathbb{R}^d)\big)$$

is the amount of time the Brownian motion has spent inside a measurable set A up to time t. In one dimension, it is well-known that ℓ_t has a density with respect to the Lebesgue measure. The following theorem is originally due to Trotter. The process $(L_t)_{t\geq 0}$ below is called *Brownian local time*.

Theorem 4.11 (Brownian local time) *Let $(B_t)_{t\geq 0}$ be a one-dimensional Brownian motion. Then almost surely, there exists a random continuous function $L : [0,\infty) \times \mathbb{R} \to [0,\infty)$ such that*

$$\int_{\mathbb{R}} \mathrm{d}x\, L_t(x) f(x) = \int_0^t \mathrm{d}s\, f(B_s) \qquad \big(t \geq 0,\ f \in B_{\mathrm{b}}(\mathbb{R})\big).$$

Modern proofs of Theorem 4.11 are based on *Tanaka's formula*, which says that

$$|B_t| = \int_0^t \mathrm{sgn}(B_s)\, \mathrm{d}B_s + L_t(0) \qquad (t \geq 0), \tag{4.11}$$

where the integral is an Itô stochastic integral. Tanaka's formula can be used as a definition of Brownian local time, for which one then proves the properties described in Theorem 4.11. For details, we refer to [McK69, Mey76, RW87]. In fact, in the remainder of this chapter, we will mostly work with Tanaka's formula as the definition of $L_t(0)$ and do not really need its interpretation as local time in the sense of Theorem 4.11.

Proposition 4.12 (Reflected Brownian motion) *Let $B = (B_t)_{t\geq 0}$ be a standard Brownian motion and let $(L_t(0))_{t\geq 0}$ be its local time at 0. Let $W = (W_t)_{t\geq 0}$ be another standard Brownian motion and let*

$$A_t := W_t - m_t(W) \quad \textit{and} \quad L_t := -m_t(W) \qquad (t \geq 0). \tag{4.12}$$

Then

$$\mathbb{P}\big[\big(|B_t|, L_t(0)\big)_{t\geq 0} \in \cdot\,\big] = \mathbb{P}\big[(A_t, L_t)_{t\geq 0} \in \cdot\,\big].$$

Proof (sketch) Let $(B_t)_{t\geq 0}$ be a Brownian motion and let

$$W_t := -\int_0^t \mathrm{sgn}(B_s)\,\mathrm{d}B_s \qquad (t \geq 0).$$

It is not hard to show that $W = (W_t)_{t\geq 0}$ is a Brownian motion. We will show that $A_t = |B_t|$ and $L_t = L_t(0)$ $(t \geq 0)$. We apply Lemma 4.9. Tanaka's formula (4.11) says that $|B_t| = L_t(0) - W_t$ $(t \geq 0)$. Clearly $|B_t|$ is nonnegative and $L_t(0)$ is nondecreasing and increases only when $|B_t| = 0$. For the details, we refer to [KS91, Thm 3.6.17]. ∎

4.3 Law of a forward and dual path

Let $(\mathcal{W}, \hat{\mathcal{W}})$ be a double Brownian web, let $(x,s),(y,u) \in \mathbb{R}^2$, and let $\pi \in \mathcal{W}(x,s)$ and $\hat{\pi} \in \hat{\mathcal{W}}(y,u)$ be the almost surely unique paths in the web and dual web starting at these points. If $s \geq u$, then it is easy to see (for example using discrete approximation) that the paths π and $\hat{\pi}$ are independent, but this cannot be the case when $s < u$, since we have seen in Section 3.5 that π and $\hat{\pi}$ do not cross, which would have a positive probability for independent forward and backward Brownian paths. In this section, we give a precise description of the joint law of π and $\hat{\pi}$. This goes back to [STW00].

Lemma 4.13 (Path reflected to the right off a dual path) *Let $\pi \in \Pi^\uparrow$ and $\hat{\pi} \in \Pi^\downarrow$ satisfy $\sigma_\pi < \tau_{\hat{\pi}}$ and $\hat{\pi}(\sigma_\pi) \leq \pi(\sigma_\pi)$. Then there exists a unique path $\pi' \in \Pi^\uparrow$ and continuous function $\psi : \mathbb{R} \to \mathbb{R}$ such that*

(i) *$z_{\pi'} = z_\pi$ and $\pi'(t) = \pi(t) + \psi(t)$ $(t \geq \sigma_\pi)$,*

(ii) *$\hat{\pi}(t) \leq \pi'(t)$ for all $s \leq t \leq u$,*

(iii) *ψ is nondecreasing with $\psi(\sigma_\pi) = 0$ and the measure $\mathrm{d}\psi$ is concentrated on $\big\{t \in [\sigma_\pi, \tau_{\hat{\pi}}] : \hat{\pi}(t) = \pi'(t)\big\}$.*

Proof Let $(x, s) := \big(\pi(\sigma_\pi), \sigma_\pi\big)$ and $(y, u) := \big(\hat\pi(\tau_{\hat\pi}), \tau_{\hat\pi}\big)$ denote the starting points of π and $\hat\pi$. Define $\tilde f : [s, u] \to \mathbb{R}$ by $\tilde f_t := \pi(t) - \hat\pi(t)$ and set $f_t := \tilde f_t - \tilde f_0$ $(t \in [s, u])$. Then Lemma 4.9 tells us that the Skorohod reflection equation

$$\mathrm{d}g_t = \mathrm{d}f_t + \mathrm{d}\chi_t \qquad (s \leq t \leq u) \tag{4.13}$$

has a unique solution with initial state $g_s = \pi(s) - \hat\pi(s)$. (Lemma 4.9 is stated for unbounded time domains but the statement and proof carry over without a change if the time domain is bounded.) Define $\psi : \mathbb{R} \to \mathbb{R}$ by $\psi(t) := \chi_t$ $(t \in [s, u])$, $\psi(t) := \chi_s$ $(t \leq s)$ and $\psi(t) := \chi_u$ $(t \geq u)$, and define $\pi' \in \Pi^\uparrow(x, s)$ by $\pi'(t) = \hat\pi(t) + g_t$ $(t \in [s, u])$ and $\pi'(t) = \pi'(u) + \big(\pi(t) - \pi(u)\big)$ $(t \geq u)$. Then π' and ψ satisfy conditions (i)–(iii) of the lemma if and only if g and χ satisfy the Skorohod reflection equation (4.13) with initial state $g_s = \pi(s) - \hat\pi(s)$. Thus, existence and uniqueness of π' and χ follow directly from Lemma 4.9. ∎

By symmetry, in complete analogy to Lemma 4.13, we can also define a path by reflection to the left off a dual path.

Lemma 4.14 (Path reflected to the left off a dual path) *Let $\pi \in \Pi^\uparrow$ and $\hat\pi \in \Pi^\downarrow$ satisfy $\sigma_\pi < \tau_{\hat\pi}$ and $\hat\pi(\sigma_\pi) \geq \pi(\sigma_\pi)$. Then there exists a unique path $\pi' \in \Pi^\uparrow$ and continuous function $\psi : \mathbb{R} \to \mathbb{R}$ such that*

(i) $z_{\pi'} = z_\pi$ *and* $\pi'(t) = \pi(t) - \psi(t)$ $(t \geq \sigma_\pi)$,

(ii) $\hat\pi(t) \geq \pi'(t)$ *for all* $s \leq t \leq u$,

(iii) ψ *is nondecreasing with* $\psi(\sigma_\pi) = 0$ *and the measure* $\mathrm{d}\psi$ *is concentrated on* $\big\{t \in [\sigma_\pi, \tau_{\hat\pi}] : \hat\pi(t) = \pi'(t)\big\}$.

Let $\mathcal{C}(\mathbb{R})$ denote the space of all continuous functions $\psi : \mathbb{R} \to \mathbb{R}$, equipped with the topology of uniform convergence. Then we can define a map Φ by

$$\Pi^\uparrow \times \Pi^\downarrow \ni (\pi, \hat\pi) \overset{\Phi}{\longmapsto} (\pi', \psi) \in \Pi^\uparrow \times \mathcal{C}(\mathbb{R}), \tag{4.14}$$

where:

- $\pi' := \pi$ and $\psi(t) := 0$ $(t \in \mathbb{R})$ if $\sigma_\pi \geq \tau_{\hat\pi}$,
- π' and ψ are defined as in Lemma 4.13 if $\sigma_\pi < \tau_{\hat\pi}$ and $\hat\pi(\sigma_\pi) < \pi(\sigma_\pi)$,
- π' and ψ are defined as in Lemma 4.14 if $\sigma_\pi \leq \tau_{\hat\pi}$ and $\pi(\sigma_\pi) \leq \hat\pi(\sigma_\pi)$.

Lemma 4.15 (Continuity of the reflection map) *Assume that* $\pi_n, \pi \in \Pi^\uparrow$ *and* $\hat\pi_n, \hat\pi \in \Pi^\downarrow$ *satisfy* $\pi_n \to \pi$ *and* $\hat\pi_n \to \hat\pi$ *in the topologies on* $\Pi^\uparrow$ *and* $\Pi^\downarrow$. *Assume moreover that either* $\sigma_\pi \geq \tau_{\hat\pi}$ *or* $\sigma_\pi < \tau_{\hat\pi}$ *and* $\hat\pi(\sigma_\pi) \neq \pi(\sigma_\pi)$. *Then*

$$\Phi(\pi_n, \hat\pi_n) \underset{n\to\infty}{\longrightarrow} \Phi(\pi, \hat\pi)$$

in the topology on $\Pi^\uparrow \times \mathcal{C}(\mathbb{R})$.

Remark Our proof of Lemma 4.15 will show that the conclusion remains valid even when $\sigma_\pi < \tau_{\hat\pi}$ and $\hat\pi(\sigma_\pi) = \pi(\sigma_\pi)$, provided that $\pi_n(\sigma_{\pi_n}) \leq \hat\pi_n(\sigma_{\pi_n})$ for all n.

Proof of Lemma 4.15 In the proof of Lemma 4.13 we have seen that there exists a one-to-one correspondence between solutions to conditions (i)–(iii) of the lemma and solutions of the Skorohod reflection equation (4.13). By Lemma 4.9, solutions of such a Skorohod reflection equation are of the form (4.9). Thus, the claim follows from the fact that if a sequence of functions $\tilde f^n \in \mathcal{C}$ with $\tilde f^n_0 \geq 0$ converges locally uniformly to a limit $\tilde f$, then also the functions $t \mapsto m_t(\tilde f^n)$ converge locally uniformly to $t \mapsto m_t(\tilde f)$.

Note that if $\sigma_\pi < \tau_{\hat\pi}$, then we have to assume that $\hat\pi(\sigma_\pi) < \pi(\sigma_\pi)$ or $\pi_n(\sigma_{\pi_n}) \leq \hat\pi_n(\sigma_{\pi_n})$ for all n because of the obvious discontinuity in our definition of the map Φ in points $(\pi, \hat\pi)$ with $\hat\pi(\sigma_\pi) = \pi(\sigma_\pi)$. ∎

Theorem 4.16 (Interaction between forward and dual paths) *Let* $(\mathcal{W}, \hat{\mathcal{W}})$ *be a double Brownian web, let* $(x, s), (y, u) \in \mathbb{R}^2$, *and let* $\pi, \hat\pi$ *be the almost surely unique paths such that* $\pi \in \mathcal{W}(x, s)$ *and* $\hat\pi \in \hat{\mathcal{W}}(y, u)$. *Let* $B = (B_t)_{t\geq s}$ *be a Brownian motion started at* $B_s = x$, *independent of* $\hat\pi$, *and let* $(\pi', \psi) := \Phi(B, \hat\pi)$, *where* Φ *is the map in (4.14). Then* $(\pi, \hat\pi)$ *is equal in law to* $(\pi', \hat\pi)$.

Proof We use discrete approximation. We first prove an analogue statement for open paths in an arrow configuration and then take the limit, using Skorohod's representation theorem and the continuity properties of the map Φ defined in (4.14).

Let $\mathcal{U}$ be the set of open paths in an arrow configuration and let $\mathcal{U}^*$ be the set of downward open paths in the associated dual arrow configuration. Fix $(y, u) \in \mathbb{Z}^2_{\rm odd}$ and let $\hat P$ be the unique element of $\mathcal{U}^*(y, u)$. Fix $(x, s) \in \mathbb{Z}^2_{\rm even}$ and let $(X_k)_{k\geq s+1}$ be i.i.d. uniformly distributed $\{-1, +1\}$-valued random variables, independent of $\hat P$. Let P be the random walk that is defined for integer times by

$$P_t := x + \sum_{k=s+1}^{t} X_k \qquad (t \geq s),$$

and then for general $t \geq s$ by linear interpolation. We can then define a reflected random walk $P' = (P'_t)_{t \geq s}$ started at $P'_s = x$ first for integer times by

$$P'_t := \begin{cases} P'_{t-1} - 1 & \text{if } t \leq u,\ \hat{P}_t = P'_{t-1},\ \text{and } \hat{P}_{t-1} = P'_{t-1} + 1, \\ P'_{t-1} + 1 & \text{if } t \leq u,\ \hat{P}_t = P'_{t-1},\ \text{and } \hat{P}_{t-1} = P'_{t-1} - 1, \\ P'_{t-1} + X_t & \text{otherwise,} \end{cases}$$

and then for general $t \geq s$ by linear interpolation. Then it is easy to see that the conditional law of P' given $\hat{P}$ is precisely the conditional law of the the unique element of $\mathcal{U}(x, s)$ given $\hat{P}$.

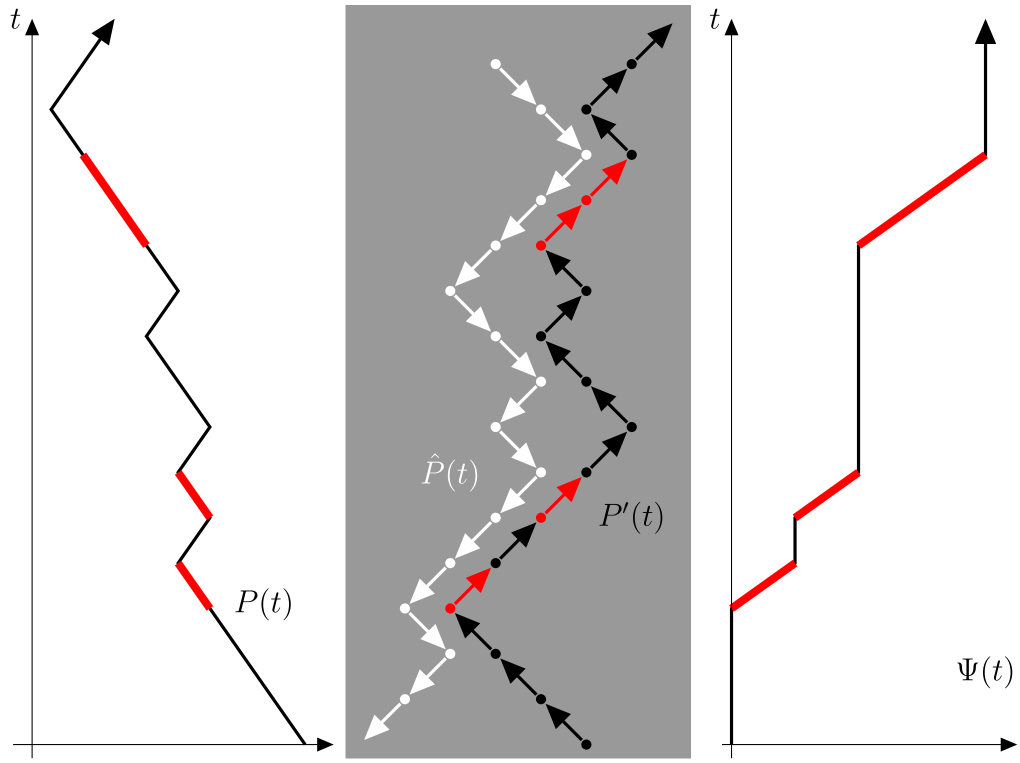

Figure 4.3: A random walk reflected to the right off a dual random walk path $\hat{P}$. The reflected path is $P'(t) = P(t) + \Psi(t)$, where P is an independent random walk and Ψ is a reflection term. The steps where reflection takes place are indicated in red.

If $u \leq s$, then clearly $P' = P'$. In the opposite case, we claim that P' is a reflected version of P, very similar to the reflected paths defined in Lemmas

4.13 and 4.14. To see this, let us first focus on the case that $\hat{P}(s) < x$. On this event, we define $\Psi : [s, t] \to \mathbb{R}$ by

$$\Psi(t) := \int_0^t \big(\tfrac{\partial}{\partial r}\hat{P}(r) - \tfrac{\partial}{\partial t}P(r)\big)1_{\{\hat{P}(r) + 1 = P'(r)\}}\mathrm{d}r \qquad \big(t \in [s, u]\big),$$

with the convention that the derivatives are zero at integer times. We extend Ψ to $\mathbb{R}$ by setting $\Psi(t) := 0$ for $t \le s$ and $\Psi(t) := \Psi(u)$ for $t \ge u$. Then it is straightforward to check (see Figure 4.3) that:

(i) $P'(t) = P(t) + \Psi(t)$ $(t \ge s)$,

(ii) $\hat{P}(t) \le P'(t)$ for all $s \le t \le u$,

(iii) Ψ is nondecreasing with $\Psi(s) = 0$ and the measure $\mathrm{d}\Psi$ is concentrated on $\big\{t \in [s, u] : \hat{P}(t) + 1 = P'(t)\big\}$.

In other words, this says that P' is precisely the random walk path P, reflected to the right off the path $\big(\hat{P}(t) + 1\big)_{t \in [s,u]}$ in the sense of Lemma 4.13. In a similar way, we see that on the event that $x < \hat{P}(s)$, the path P' is the random walk P, reflected to the left off the path $\big(\hat{P}(t) - 1\big)_{t \in [s,u]}$ in the sense of Lemma 4.14.

We can now prove the statement of the theorem. We first treat the case that $s < u$. We choose positive constants ε_n, tending to zero, and points $(x_n, s_n) \in \mathbb{Z}^2_{\mathrm{even}}$ and $(y_n, u_n) \in \mathbb{Z}^2_{\mathrm{odd}}$ such that

$$\theta_{\varepsilon_n}(x_n, s_n) \underset{n \to \infty}{\longrightarrow} (x, s) \quad \text{and} \quad \theta_{\varepsilon_n}(y_n, u_n) \underset{n \to \infty}{\longrightarrow} (y, u).$$

Note that $s_n < u_n$ for all n large enough. We let $\hat{P}_n$ be the unique element of $\mathcal{U}(y_n, u_n)$, we let P_n be an independent random walk started at $P_n(s) = x$, and we define reflected paths P'_n in terms of $\hat{P}_n$ and P_n as above. We let Ψ_n denote the associated reflection functions. It will also be handy to introduce notation for the dual path $\hat{P}_n$ shifted by $+1$ or -1 depending on whether $\hat{P}_n(s) < x_n$ or $x_n < \hat{P}_n(s)$. We denote this modified dual path by $\tilde{P}_n$, i.e.,

$$\tilde{P}_n(t) := \begin{cases} \hat{P}_n(t) + 1 & \text{if } \hat{P}_n(s) < x_n, \\ \hat{P}_n(t) - 1 & \text{if } x_n < \hat{P}_n(s) \end{cases} \qquad \big(t \in (-\infty, u_n]\big).$$

We denote the corresponding diffusively rescaled paths and functions by

$$(\tilde{\pi}_n, B_n, B'_n, \psi_n) := \theta_{\varepsilon_n}(\tilde{P}_n, P_n, P'_n, \Psi_n).$$

Note that we rescale the reflection function in the same way as the paths. The properties (i)–(iii) above are preserved under rescaling, so on the event that $\tilde{\pi}_n(s_n) \neq x_n$, we have that

$$(B'_n, \psi_n) = \Phi(B_n, \tilde{\pi}_n), \tag{4.15}$$

where Φ is the map in (4.14). In the case when $u \leq s$, we can without loss of generality choose our approximating points such that $u_n \leq s_n$ for all n. Then, taking (4.15) as a definition, we also have that $(B'_n, \hat{\pi}_n)$ is distributed as the pair of rescaled forward and dual paths started from (x_n, s_n) and (y_n, u_n).

The rest of the proof is now easy. It follows from Theorem 3.18 that the pair of rescaled paths $(B'_n, \tilde{\pi}_n)$ converges in law to $(\pi, \hat{\pi})$, the forward and dual Brownian web paths mentioned in the theorem. On the other hand, the pair of rescaled paths $(\tilde{\pi}_n, B_n)$ converges in law to a pair $(\hat{\pi}, B)$ where $\hat{\pi}$ is as before and B is an independent Brownian motion started at (x, s). Using Skorhod's representation theorem, we can couple our random variables such that the latter convergence is almost sure. Since the event $\tilde{\pi}(s) \neq x$ has probability one, we can use (4.15) and the continuity property of the map Φ stated in Lemma 4.15, to conclude that for this coupling, (B'_n, ψ_n) converge almost surely to $(\pi', \psi) := \Phi(B, \hat{\pi})$. In particular, this implies that $(\pi, \hat{\pi})$ is equal in law to $(\pi', \hat{\pi})$, as claimed. ∎

Figure 4.4 shows a numerical simulation of the set of all Brownian web paths started at a fixed time s, together with the set of all dual Brownian web paths started at a fixed time $t > s$, where one can (with a bit of imagination) see the forward path being reflected off the dual paths, and vice versa.

4.4 Special points

We have defined the Brownian web $\mathcal{W}$ as the closure of $\{\pi_z : z \in \mathcal{D}\}$, where $(\pi_z)_{z \in \mathcal{D}}$ is a collection of coalescing Brownian motions started from a countable dense set $\mathcal{D} \subset \mathbb{R}^2$. Here $\{\pi_z : z \in \mathcal{D}\}$ is precompact by Proposition 3.6 and hence $\mathcal{W}$ is a compact subset of $\Pi^\uparrow$. Using compactness and the fact that $\mathcal{D}$ is dense, we see that for each $z \in \mathbb{R}^2$, there exists at least one path $\pi \in \mathcal{W}$ that starts at z. For each $z \in \mathbb{R}^2$, we let

$$m_{\text{out}}(z) := \big|\mathcal{W}(z)\big|$$

denote the number of paths in $\mathcal{W}$ that start at z. In Theorem 3.7, we have proved that $m_{\text{out}}(z) = 1$ a.s. for each deterministic $z \in \mathbb{R}^2$. In this section, we will prove that in spite of this, almost surely, there exist points z with

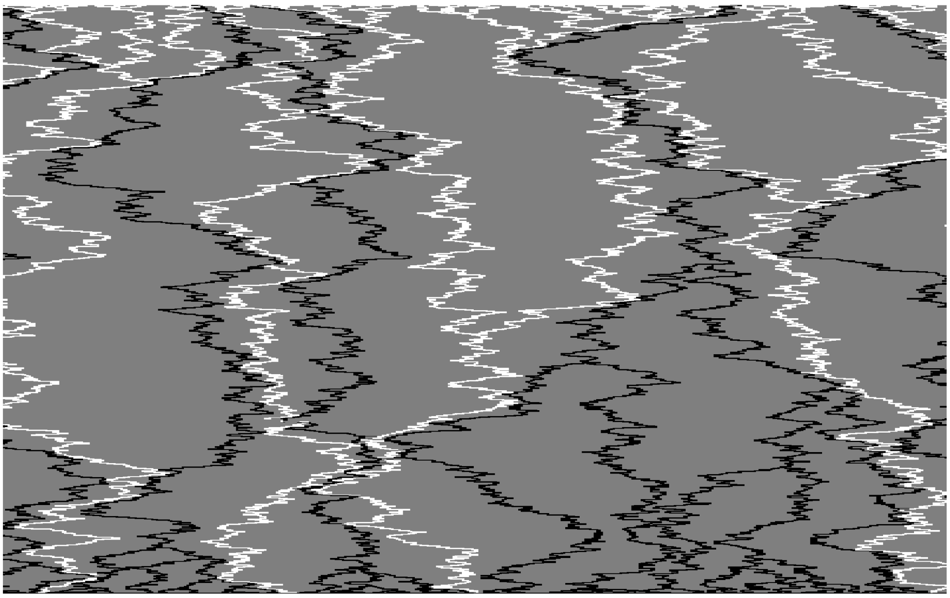

Figure 4.4: The sets of all Brownian web paths and dual Brownian web paths started at two fixed times $s < t$.

$m_{\rm out}(z) = 2$ and even $m_{\rm out}(z) = 3$. The key to understanding this is (again) duality.

We say that a path $\pi \in \Pi^\uparrow$ *enters* a point $z = (x, u) \in \mathbb{R}^2$ if $\sigma_\pi < u$ and $\pi(u) = x$. We denote the set of Brownian web paths entering z by

$$\mathcal{W}_{\rm in}(z) := \big\{\pi \in \mathcal{W} : \pi \text{ enters } z\big\}.$$

We define an equivalence relation $\sim$ on $\mathcal{W}_{\rm in}(z)$ by setting $\pi \sim \pi'$ if and only if there exists a time s with $\sigma_\pi \vee \sigma_{\pi'} \leq s < u$ such that $\pi(t) = \pi'(t)$ for all $t \in [s, u]$ and we let

$$\dot{\mathcal{W}}_{\rm in}(z) := \big\{\dot\pi \in \mathcal{W} : \pi \in \mathcal{W}_{\rm in}(z)\big\} \quad \text{with} \quad \dot\pi := \{\pi' \in \mathcal{W}_{\rm in}(z) : \pi' \sim \pi\}$$

denote the set of equivalence classes. With these conventions,

$$m_{\rm in}(z) := \big|\dot{\mathcal{W}}_{\rm in}(z)\big|$$

denotes the number of nonequivalent paths $\pi \in \mathcal{W}$ entering z. We call $(m_{\rm in}(z), m_{\rm out}(z))$ the *type* of a point $z \in \mathbb{R}^2$.

Theorem 4.17 (Special points of the Brownian web) *Let $\mathcal{W}$ be a Brownian web. Then almost surely, all points in $\mathbb{R}^2$ are of one of the following types:*

$$(0,1), \quad (0,2), \quad (0,3), \quad (1,1), \quad (1,2), \quad (2,1),$$

and all these types occur. For each deterministic $t \in \mathbb{R}$*, almost surely, all points in* $\mathbb{R} \times \{t\}$ *are of one of the following types:*

$$(0,1), \quad (0,2), \quad (1,1),$$

and all these types occur. A deterministic point $(x,t) \in \mathbb{R}^2$ *is almost surely of type* $(0,1)$.

Points of type $(1,2)$ are further distinguished into points of type $(1,2)_{\rm l}$ and $(1,2)_{\rm r}$, depending on whether the incoming path is the left or right of the two outgoing paths. The proof of Theorem 4.17 needs a few preparations.

Recall that in Section 2.7, we defined paths $\pi \in \Pi(\overline{\mathbb{R}})$ as subsets $\pi \subset \mathcal{R}(\overline{\mathbb{R}})$ with certain special properties. Our definitions excluded the case that $\pi = \emptyset$, but for some purposes it is convenient to allow the empty paths, so we define $\Pi_0(\overline{\mathbb{R}}) := \Pi(\overline{\mathbb{R}}) \cup \{\emptyset\}$. If $[s,u] \subset \mathbb{R}$ is a compact nonempty interval, then we define the restriction of a path $\pi \in \Pi(\overline{\mathbb{R}})$ to the interval $[s,u]$ as

$$\pi\big|_{[s,u]} := \big\{(x,t) \in \pi : t \in [s,u]\big\}.$$

Note that $\pi\big|_{[s,u]} \in \Pi_0(\overline{\mathbb{R}})$ for all $\pi \in \Pi(\overline{\mathbb{R}})$. If $\mathcal{A} \subset \Pi(\overline{\mathbb{R}})$ is a set of paths, then we define

$$\mathcal{A}\big|_{[s,u]} := \big\{\pi\big|_{[s,u]} : \pi \in \mathcal{A}\big\} \backslash \{\emptyset\},$$

where now we remove the empty path so that $\mathcal{A}\big|_{[s,u]} \subset \Pi(\overline{\mathbb{R}})$.

Lemma 4.18 (Independent increments) *Let* $-\infty < t_0 \leq \cdots \leq t_n < \infty$*. Then the restricted Brownian webs*

$$\mathcal{W}\big|_{[t_0,t_1]}, \ldots, \mathcal{W}\big|_{[t_{n-1},t_n]}$$

are independent.

Proof (sketch) We use discrete approximation. Let $-\infty < s < u < \infty$, choose positive constants ε_n, tending to zero, and let $s_n, u_n \in \mathbb{Z}$ satisfy $\varepsilon_n^2 s_n \to s$ and $\varepsilon_n^2 u_n \to u$. Let $\mathcal{U}$ be the set of paths in an arrow configuration. Then we claim that

$$\mathbb{P}\big[\theta_{\varepsilon_n}(\overline{\mathcal{U}}\big|_{[s_n,u_n]}) \in \cdot\,\big] \underset{n\to\infty}{\Longrightarrow} \mathbb{P}\big[\mathcal{W}\big|_{[s,u]} \in \cdot\,\big], \tag{4.16}$$

where $\Rightarrow$ denotes weak convergence of probability laws on $\Pi(\overline{\mathbb{R}})$. This would quite easily follow from Theorem 3.18 if the map $\mathcal{A} \mapsto \mathcal{A}\big|_{[s,u]}$ were continuous, but that is not the case (Exercise 4.19 below). Nevertheless, (4.16) can be proved by adapting the proof of Theorem 3.18 in a straightforward manner. For brevity, we skip the details.

Now let $-\infty < t_0 \leq \cdots \leq t_m < \infty$. Then we can find $t_1^n, \ldots, t_m^n \in \mathbb{Z}$ with $t_1^n \leq \cdots \leq t_m^n$, such that $\varepsilon_n^2 t_i^n \to t_i$ for each $0 \leq i \leq m$. Since the restricted discrete webs

$$\overline{\mathcal{U}}\big|_{[t_0^n, t_1^n]}, \ldots, \overline{\mathcal{U}}\big|_{[t_{m-1}^n, t_m^n]}$$

are independent, taking the limit, we see that the same is true for the restricted Brownian webs. ∎

Exercise 4.19 *Show that the map $\mathcal{A} \mapsto \mathcal{A}\big|_{[s,u]}$ is not continuous with respect to the topology on $\mathcal{K}(\Pi(\overline{\mathbb{R}}))$.*

The following lemma shows how the type of a point in the dual Brownian web can be derived from its type in the Brownian web. This lemma will also be key in understanding why certain types of points must exist, or on the other hand do not exist; see Figure 4.5.

Lemma 4.20 (Types of points in dual web) *Let $(\hat{m}_{\rm in}(z), \hat{m}_{\rm out}(z))$ denote the type of a point $z \in \mathbb{R}^2$ in the dual Brownian web $\hat{\mathcal{W}}$. Then for each $z \in \mathbb{R}^2$,*

$$m_{\rm out}(z) = \hat{m}_{\rm in}(z) + 1 \quad \textit{and} \quad \hat{m}_{\rm out}(z) = m_{\rm in}(z) + 1.$$

Proof By symmetry, it suffices to prove that $m_{\rm out}(z) = \hat{m}_{\rm in}(z) + 1$. If there is an incoming path in $\hat{\mathcal{W}}$ at z, then forward paths started on either side of such a dual path cannot coalesce until the starting time of the dual path, since otherwise the dual path would enter the wedge defined by these forward paths. As a result, since the incoming paths divide the area just above z into $\hat{m}_{\rm in}(z) + 1$ regions, approaching the point z from different directions, using the compactness of $\mathcal{W}$, we see that there are at least $\hat{m}_{\rm in}(z) + 1$ distinct paths in $\mathcal{W}$ starting at z. On the other hand, if there are two outgoing paths in $\mathcal{W}$ at z, then any dual path that is started between these paths must stay between these forward paths and pass through z. Therefore, $\hat{m}_{\rm in} \geq m_{\rm out} - 1$. Together with our earlier claim that $m_{\rm out}(z) \geq \hat{m}_{\rm in}(z) + 1$, this proves the claim. ∎

Proof of Theorem 4.17 (sketch)] Let $\mathcal{D} \subset \mathbb{R}^2$ be countable and dense and let $\mathcal{W}(\mathcal{D})$ be the associated skeleton of the Brownian web. By Lemma 4.7 and the remark below it, if $\pi \in \mathcal{W}$ enters a point z, then there exists a $\pi' \in \mathcal{W}(\mathcal{D})$ such that π and π' are two equivalent paths entering z. Thus, for each $z \in \mathbb{R}^2$, we have that $m_{\rm in}(z)$ is also the number of non-equivalent paths in the skeleton $\mathcal{W}(\mathcal{D})$ entering z. A completely analogue statement holds for paths in the dual web.

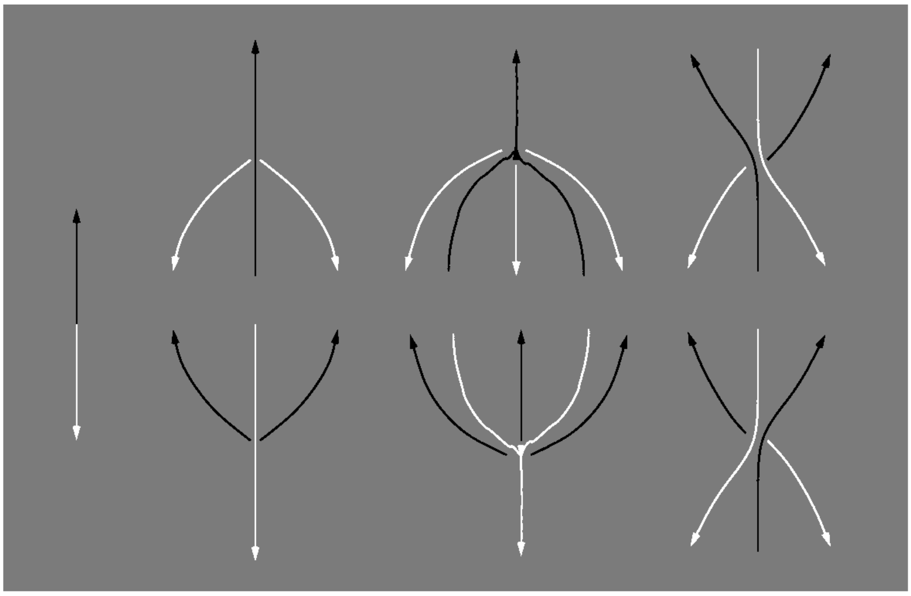

Figure 4.5: Possibe types of points in the Brownian web and its dual.

If $z \in \mathbb{R}^2$ is a deterministic point, then Theorem 3.7 tells us that $m_{\rm out}(z) = 1$ and $\hat{m}_{\rm out}(z) = 1$. By Lemma 4.20, this implies that z is almost surely of type $(0, 1)$.

If $t \in \mathbb{R}$ is a deterministic time, then clearly there exist (random) $x \in \mathbb{R}$ such that $m_{\rm in}(x, t) = 1$, while by our previous argument there also exist $x \in \mathbb{R}$ such that $m_{\rm in}(x, t) = 0$. We claim there exist no $x \in \mathbb{R}$ such that $m_{\rm in}(x, t) \geq 2$. Indeed, by the remarks at the beginning of our proof, for this to be true there would have to exist skeletal paths $\pi, \pi' \in \mathcal{W}(\mathcal{D})$ with $\pi(t) = \pi'(t)$ while $\pi(s) \neq \pi'(s)$ for all $\sigma_\pi \vee \sigma_{\pi'} \leq s < t$. For any two paths in $\{\pi_z : z \in \mathcal{D}\}$, this event clearly has probability zero. Since $\mathcal{D}$ is countable, we can conclude such paths π, π' do not exist. By a similar argument, we see that there exist no $x \in \mathbb{R}$ such that $m_{\rm in}(x, t) = 1$ and $\hat{m}_{\rm in}(x, t) = 1$. Indeed, for this to be true, a path in $\{\pi_z : z \in \mathcal{D}\}$ started below time t and a dual path in $\{\hat{\pi}_z : z \in \mathcal{D}\}$ started above time t would have to be at the same location at time t. By Lemma 4.18, what happens below time t is independent of what happens above time t, and therefore such an event has probability zero. We conclude that if $t \in \mathbb{R}$ is a deterministic time, then for each $x \in \mathbb{R}$, one of the following statements must be true: 1. $m_{\rm in}(x, t) = \hat{m}_{\rm in}(x, t) = 0$ (which is true for deterministic x), 2. $m_{\rm in}(x, t) = 0$ and $\hat{m}_{\rm in}(x, t) = 1$, and 3. $m_{\rm in}(x, t) = 1$ and $\hat{m}_{\rm in}(x, t) = 0$, and all these cases occur. By Lemma 4.20, it follows that all points on $\mathbb{R} \times \{t\}$ are of the types $(0, 1)$, $(0, 2)$, $(1, 1)$, and all these types occur.

To complete the proof, we must show that (at random times), points of type (0,3), (1,2), and (2,1) also occur, but no other types of points can occur. It is clear that there exist points $z \in \mathbb{R}^2$ with $m_{\rm in}(z) = 2$, but by the remarks at the beginning of our proof, there exist no points $z \in \mathbb{R}^2$ with $m_{\rm in}(z) \geq 3$, because for that to occur, three Brownian motions started from deterministic points would have to coalesce in one and the same point, which has probability zero. There are in fact only countably many points z with $m_{\rm in}(z) = 2$, since these are the coalescence points of the countable collection of coalescing Brownian motions $\{\pi_z : z \in \mathcal{D}\}$.

We claim that each point z with $m_{\rm in}(z) = 2$ satisfies $\hat{m}_{\rm in}(z) = 0$. To see this, we need to consider the joint law of two forward paths started from deterministic points and one dual path started from a deterministic point. We claim that if we forget about the trajectory of the coalescing forward Brownian motions after their coalescence time, then the conditional law of the dual path given the trajectories of the forward paths up to their coalescence time is described by a Brownian motion with Skorohod reflection off the two forward paths. The proof of this is similar to the proof of Theorem 4.16, so we skip the details. Now the conditional probability that the dual path hits the two forward paths exactly in the point where they coalesce is zero, which implies there are no points with $m_{\rm in}(z) = 2$ and $\hat{m}_{\rm in}(z) \geq 1$.

By Lemma 4.20, these arguments show that there exist points of type $(2, 1)$, and by duality also of type $(0, 3)$, but no points of type (n, m) with $n \geq 3$, or $n = 2$ and $m \geq 2$, or $n = 1$ and $m \geq 3$.

By Lemma 4.20, and our previous remarks, to complete the proof and show that there exist points of type $(1, 2)$, it suffices to prove that there exist point $z \in \mathbb{R}^2$ with $m_{\rm in}(z) = 1$ and $\hat{m}_{\rm in}(z) = 1$. This follows from the fact that forward paths reflect off dual paths, proved in Theorem 4.16. ∎

Exercise 4.21 *Let $\mathcal{W}$ be a Brownian web. Recall the definition of the maximal path $\pi^+_z \in \mathcal{W}(z)$ ($z \in \mathbb{R}^2$) from Lemma 4.6. Say that a point $z \in \mathbb{R}^2$ is a* point of continuity *of the map $z \mapsto \pi^+_z$ if*

$$\pi^+_{z_n} \to \pi^+_z \quad \textit{for all } z_n \in \mathbb{R}^2 \textit{ s.t. } z_n \to z.$$

Show that the set of points of continuity of the map $z \mapsto \pi^+_z$ is given by

$$\big\{z \in \mathbb{R}^2 : \hat{m}_{\rm in}(z) = 0\big\}.$$

Exercise 4.22 *Let $x \in \mathbb{R}$ be deterministic. Try to determine which types of points almost surely occur, or do not occur, in the set $\{x\} \times \mathbb{R}$. Note: there is one type of point for which this question is not so easy to answer rigorously (although you may guess the correct answer).*

4.5 The continuum voter model

In this section, we return to the processes $(A_t)_{t\geq 0}$ and $(B_t)_{t\geq 0}$ defined in (4.2). We will show that they are indeed Markov processes, which can be thought of as the continuum limits of the voter model and coalescing random walks, respectively.

Lemma 4.23 (Stochastic flow property) *Let $t \in \mathbb{R}$ be deterministic. Then almost surely,*

$$\mathcal{X}_{t,u} \circ \mathcal{X}_{s,t} = \mathcal{X}_{s,u} \quad \textit{and} \quad \mathcal{Y}_{t,u} \circ \mathcal{Y}_{s,t} = \mathcal{Y}_{s,u} \qquad (s \leq t \leq u).$$

Proof We first prove the inclusion $\mathcal{X}_{t,u} \circ \mathcal{X}_{s,t}(A) \supset \mathcal{X}_{s,u}(A)$ for each closed $A \subset \overline{\mathbb{R}}$. If $x \in \mathcal{X}_{s,u}(A)$, then there exists a $\hat{\pi} \in \hat{\mathcal{W}}(x,u)$ such that $\hat{\pi}(s) \in A$. Set $x' := \hat{\pi}(t)$. By Excercise 3.19, there exists a $\hat{\pi}' \in \hat{\mathcal{W}}(x',t)$ that coincides with $\hat{\pi}$ on $(-\infty,t]$. It follows that $x' \in \mathcal{X}_{s,t}(A)$, which by the fact that $\hat{\pi}(t) = x'$ implies that $x \in \mathcal{X}_{t,u} \circ \mathcal{X}_{s,t}(A)$.

If t is deterministic, then the opposite conclusion can also almost surely be draw. If $x \in \mathcal{X}_{t,u} \circ \mathcal{X}_{s,t}(A)$, then there exists an $x' \in \overline{\mathbb{R}}$, $\hat{\pi} \in \hat{\mathcal{W}}(x,t)$, and $\hat{\pi}' \in \hat{\mathcal{W}}(x',t)$ such that $\hat{\pi}(t) = x'$ and $\hat{\pi}'(s) \in A$. Since t is deterministic and $\hat{m}_{\rm in}(x',t) = 1$, Theorem 4.17 allows us to conclude that (x',t) is of type $(1,1)$. But then $\hat{\pi}'$ must coincide with $\hat{\pi}$ on $(-\infty,t]$, which implies that $x \in \mathcal{X}_{s,u}(A)$.

This concludes the proof that if $t \in \mathbb{R}$ is deterministic, then $\mathcal{X}_{t,u} \circ \mathcal{X}_{s,t} = \mathcal{X}_{s,u}$ a.s. for all $s \leq t$, $u \geq t$, and closed $A \subset \overline{\mathbb{R}}$. The proof that $\mathcal{Y}_{t,u} \circ \mathcal{Y}_{s,t} = \mathcal{Y}_{s,u}$ is basically the same. ∎

Proposition 4.24 (Continuum voter model) *Let $\mathcal{K}(\overline{\mathbb{R}})$ be the space of all compact subsets of $\overline{\mathbb{R}}$, equipped with the Hausdorff topology. Then setting*

$$\left.\begin{array}{l} P_t(A,\,\cdot\,) := \mathbb{P}\big[\mathcal{X}_{0,t}(A) \in \,\cdot\,\big], \\ Q_t(A,\,\cdot\,) := \mathbb{P}\big[\mathcal{Y}_{0,t}(A) \in \,\cdot\,\big] \end{array}\right\} \qquad \big(A \in \mathcal{K}(\overline{\mathbb{R}}),\ t \geq 0),$$

defines transition kernels on $\mathcal{K}(\overline{\mathbb{R}})$ such that for each $A, B \in \mathcal{K}(\overline{\mathbb{R}})$, the processes $(A_t)_{t\geq 0}$ and $(B_t)_{t\geq 0}$ defined in (4.2) are Markov processes with transition kernels $(P_t)_{t\geq 0}$ and $(Q_t)_{t\geq 0}$, respectively.

Proof We need to show that for each deterministic $0 \leq t \leq u$,

$$\mathbb{P}\big[A_u \in \,\cdot\, \big|\, (A_s)_{0\leq s\leq t}\big] = P_{u-t}(A_t,\,\cdot\,) \quad \text{a.s.}$$

Since A is deterministic and $A_s = \mathcal{X}_{0,s}(A)$ $(0 \leq s \leq t)$, where $\mathcal{X}_{0,s}$ is a function of $\mathcal{W}\big|_{[0,s]}$, we see that $(A_s)_{0\leq s\leq t}$ is a function of $\mathcal{W}\big|_{[0,t]}$. Lemma 4.23 tells us moreover that

$$A_u = \mathcal{X}_{0,u}(A) = \mathcal{X}_{t,u} \circ \mathcal{X}_{0,t}(A) = \mathcal{X}_{t,u}(A_t).$$

Since $\mathcal{X}_{t,u}$ is a function of $\mathcal{W}\big|_{[t,u]}$, by Lemma 4.18, it is independent of $\mathcal{W}\big|_{[0,t]}$ and hence also of $(A_s)_{0\leq s\leq t}$. This implies that

$$\mathbb{P}\big[\mathcal{X}_{t,u}(A_t)\in\,\cdot\,\big|\,(A_s)_{0\leq s\leq t}=(a_s)_{0\leq s\leq t}\big]=\mathbb{P}\big[\mathcal{X}_{t,u}(a_t)\in\,\cdot\,\big]=P_{u-t}(a_t,\,\cdot\,)$$

for almost every $(a_s)_{0\leq s\leq t}$ with respect to the law of $(A_s)_{0\leq s\leq t}$. It follows that

$$\mathbb{P}\big[A_u\in\,\cdot\,\big|\,(A_s)_{0\leq s\leq t}\big]=\mathbb{P}\big[\mathcal{X}_{t,u}(A_t)\in\,\cdot\,\big|\,(A_s)_{0\leq s\leq t}\big]=P_{u-t}(A_t,\,\cdot\,)\quad\text{a.s.}$$

The proof for $(B_t)_{t\geq 0}$ is completely the same. ∎

Remark An alternative construction of a continuum voter models has been given by Steve Evans in [Eva97]. His construction is based on duality with coalescing motions, which can be very general Markov processes (not necessarily Brownian motions, and with state space much more general than the real line), and he also allows voter models with infinitely many types. Because of the high level of abstraction, the paper is a bit hard to read.

Exercise 4.25 *Let $\mathcal{W}$ be a Brownian web. Show that the set $\mathcal{I}:=\big\{x\in\overline{\mathbb{R}}:\hat{m}_{\rm in}(x,0)=1\big\}$ is a.s. countable. Conditional on $\mathcal{W}$, let $\big(\chi(x)\big)_{x\in\mathcal{I}}$ be i.i.d. uniformly distributed $\{0,1\}$-valued random variables, and define $(A_t)_{t>0}$ by*

$$A_t:=\big\{x\in\overline{\mathbb{R}}:\exists\hat{\pi}\in\hat{\mathcal{W}}(x,t)\ \textit{s.t.}\ \chi(\hat{\pi}(0))=1\big\}.$$

Show that for deterministic $0<s<t$, one has

$$A_t=\mathcal{X}_{s,t}(A_s)\quad\text{a.s.}$$

Sketch a proof that the process $(A_t)_{t>0}$ is the scaling limit of voter models started in i.i.d. uniformly distributed initial laws (compare the picture in Section 1.2).

4.6 The Arratia flow

Let $(\mathcal{W},\hat{\mathcal{W}})$ be a double Brownian web. Recall the definition of the minimal and maximal paths $\pi^{\pm}_z$ starting at a point $z\in\mathbb{R}^2$ in Lemma 4.6. For $z=(y,t)\in\mathbb{R}^2$, we similarly let $\hat{\pi}^-_z$ and $\hat{\pi}^+_z$ denote the unique elements of $\mathcal{W}(z)$ such that $\hat{\pi}^-_z(s)\leq\hat{\pi}(s)\leq\hat{\pi}^+_z(s)$ for all $s\leq t$. For each $s,t\in\mathbb{R}$ with $s\leq t$, we define a map $\Phi_{s,t}:\overline{\mathbb{R}}\to\overline{\mathbb{R}}$ by

$$\Phi_{s,t}(x):=\pi^+_{(x,s)}(t)\qquad(s\leq t,\ x\in\mathbb{R}).\tag{4.17}$$

Similarly, we set

$$\hat{\Phi}_{t,s}(y) := \hat{\pi}^{+}_{(y,t)}(s) \qquad (s \leq t,\ y \in \mathbb{R}). \tag{4.18}$$

We say that a function $f : \overline{\mathbb{R}} \to \overline{\mathbb{R}}$ is *monotone* if it is nondecreasing, i.e., $x \leq y$ implies $f(x) \leq f(y)$. If $f : \mathbb{R} \to \mathbb{R}$ is monotone and right-continuous with $\lim_{x\to\pm\infty} f(x) = \pm\infty$, then we define its *generalised inverse* as

$$f^{-1}(y) := \sup\big\{x \in \mathbb{R} : f(x) \leq y\big\} \qquad (y \in \mathbb{R}). \tag{4.19}$$

Then f^{-1} is monotone and right-continuous with $\lim_{y\to\pm\infty} f^{-1}(y) = \pm\infty$, and its generalised inverse is the function f. See Figure 4.6 for an illustration.

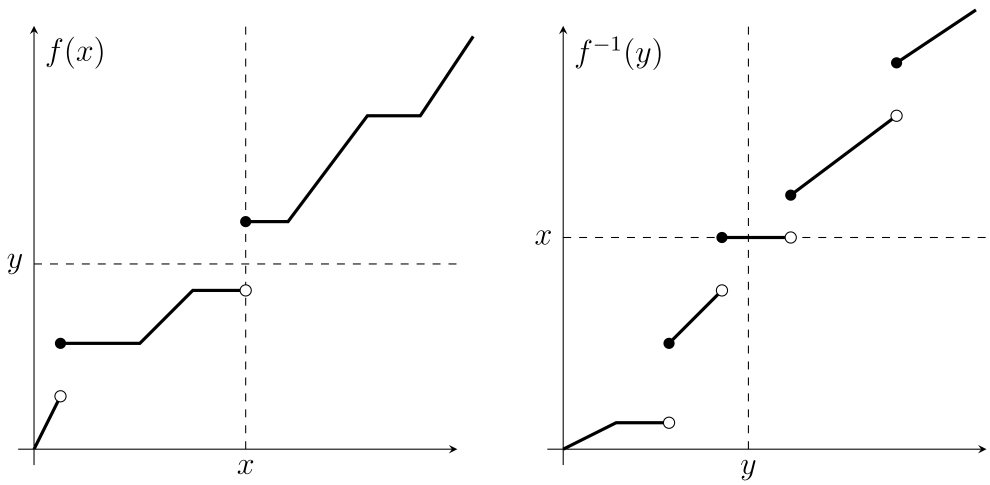

Figure 4.6: The generalised inverse $f^{-1}(y) := \sup\big\{x \in \mathbb{R} : f(x) \leq y\big\}$ of a monotone right-continuous function $f : \mathbb{R} \to \mathbb{R}$.

Proposition 4.26 (Arratia flow) *Almost surely, for all $s, t \in \mathbb{R}$ with $s \leq t$, the functions $\Phi_{s,t}$ are monotone and right-continuous with $\lim_{x\to\pm\infty} \Phi_{s,t}(x) = \pm\infty$, and $\Phi_{s,s}$ is the identity map for each $s \in \mathbb{R}$. Moreover, for each deterministic $s, t, u \in \mathbb{R}$ with $s \leq t \leq u$, one has*

$$\Phi_{t,u} \circ \Phi_{s,t} = \Phi_{s,u} \quad \text{a.s.} \tag{4.20}$$

For each deterministic $t_0 \leq \cdots \leq t_n$,

$$\textit{the maps } \Phi_{t_0,t_1}, \ldots, \Phi_{t_{n-1},t_n} \textit{ are independent.} \tag{4.21}$$

Finally, almost surely for all $s, t \in \mathbb{R}$ with $s \leq t$, the random map $\hat{\Phi}_{t,s}$ is the generalised inverse of $\Phi_{s,t}$.

Property (4.20), together with the fact that $\Phi_{s,s}$ is the identity map for each $s \in \mathbb{R}$, say that $(\Phi_{s,t})_{s\leq t}$ is a *stochastic flow*, and property (4.21) says that $(\Phi_{s,t})_{s\leq t}$ has *independent increments*. One has to be careful with the order of the "for all" and "almost sure" statements. We will see in the proof of Proposition 4.26 that while property (4.20) holds almost surely at deterministic times, it fails to hold simultaneously for all times, since there almost surely exist random times $s \leq t \leq u$ such that (4.20) does not hold. By contrast, all other statements of the proposition hold almost surely simultaneously for all times (deterministic or random).

Proof of Proposition 4.26 Since paths in the Brownian web coalesce as soon as they meet (Lemma 4.5), the maps $\Phi_{s,t}$ are clearly monotone for all $s \leq t$, and right-continuity follows immediately from Lemma 4.6. By monotonicity, the limit $\lim_{x\to\infty} \Phi_{s,t}(x)$ exist. If this limit were finite, then that would imply the existence of paths in the Brownian net coming in from infinity, which by Lemma 3.8 does not happen, so we conclude that $\lim_{x\to\infty} \Phi_{s,t}(x) = \infty$ and similarly $\lim_{x\to-\infty} \Phi_{s,t}(x) = -\infty$. These same arguments also show that the functions $\hat{\Phi}_{t,s}$ are monotone and right-continuous with $\lim_{x\to\pm\infty} \hat{\Phi}_{t,s}(x) = \pm\infty$.

It is clear from the definition that $\Phi_{s,s}$ is the identity map for each $s \in \mathbb{R}$. We next prove (4.20). Since $\Phi_{s,s}$ is the identity map for each $s \in \mathbb{R}$, it suffices to prove the statement for $s < t < u$. Thus, we need to show that for deterministic times $s < t < u$, one almost surely has

$$\pi^+_{(\pi^+_{(x,s)}(t),t)}(u) = \pi^+_{(x,s)}(u) \qquad (x \in \mathbb{R}). \tag{4.22}$$

Since t is deterministic, by Theorem 4.17, almost surely for all $x \in \mathbb{R}$, the point $(\pi^+_{(x,s)}(t), t)$ is of one of the types $(0,1)$, $(0,2)$, or $(1,1)$. Since the path $\pi^+_{(x,s)}$ enters this point, we must have $m_{\rm in}(\pi^+_{(x,s)}(t), t) \geq 1$, which allows us to conclude that the point $(\pi^+_{(x,s)}(t), t)$ must be of type $(1,1)$. It follows that $\pi^+_{(\pi^+_{(x,s)}(t),t)}$ must be the continuation of $\pi^+_{(x,s)}$ and (4.22) holds. Note, however, that this argument essentially uses that t is deterministic. There exist random times t such that for a suitable choice of $s < t$ and $x \in \mathbb{R}$, the point $(\pi^+_{(x,s)}(t), t)$ is of type $(1,2)_{\rm l}$. In such cases, the path $\pi^-_{(\pi^+_{(x,s)}(t),t)}$ is the continuation of $\pi^+_{(x,s)}$ and (4.22) fails.

The independent increment property (4.21) follows from the analogue proerty of the Brownian web, stated in Lemma 4.18. To complete the proof, we must show that for all $s, t \in \mathbb{R}$ with $s \leq t$, the random map $\hat{\Phi}_{t,s}$ defined in (4.18) is the generalised inverse of $\Phi_{s,t}$. We first make some general observations. Let $\mathcal{F}$ denote the space of all monotone and right-continuous functions with $\lim_{x\to\pm\infty} f(x) = \pm\infty$. We adopt the notation

$$f(x-) := \lim_{x'\uparrow x} f(x) \qquad (x \in \mathbb{R},\ f \in \mathcal{F}).$$

Then $x \mapsto f(x-)$ is the left-continuous modification of f. It is not hard to check that

$$f(x) < y \quad \Leftrightarrow \quad x < f^{-1}(y-) \qquad (x, y \in \mathbb{R},\ f \in \mathcal{F}), \tag{4.23}$$

where $y \mapsto f^{-1}(y-)$ is the left-continuous modification of the generalised inverse of f. Therefore, to show that $\hat\Phi_{t,s}$ is the generalised inverse of $\Phi_{s,t}$, it suffices to show that

$$\Phi_{s,t}(x) < y \quad \Leftrightarrow \quad x < \hat\Phi_{t,s}(y-) \qquad (s \leq t,\ x, y \in \mathbb{R}).$$

The left-continuous modification of $\hat\Phi_{t,s}$ is given by

$$\hat\Phi_{s,t}(y-) = \hat\pi^-_{(y,t)}(s) \qquad (s \leq t,\ y \in \mathbb{R}),$$

so we need to show that

$$\pi^+_{(x,s)}(t) < y \quad \Leftrightarrow \quad x < \hat\pi^-_{(y,t)}(t). \tag{4.24}$$

By the symmetry of the problem, it suffices to prove only the implication $\Rightarrow$. Since forward and dual paths do not cross, $\pi^+_{(x,s)}(t) < y$ clearly implies $x \leq \pi^-_{(y,t)}(t)$. To see that in fact we must have $x < \hat\pi^-_{(y,t)}(t)$ imagine that, conversely, $x = \hat\pi^-_{(y,t)}(t)$. Then there must exist a $\pi \in \mathcal{W}(x,s)$ such that $\hat\pi^-_{(y,t)} \leq \pi$ on $[s,t]$. Since forward and dual paths interact via Skorohod reflection, π must lie strictly on the right of $\hat\pi^-_{(y,t)}$ at most times in $[s,t]$, so $\pi^+_{(x,s)}$ is not the maximal element of $\mathcal{W}(x,s)$, contradicting our assumptions. ∎

Remark Our proof actually shows that for each deterministic $t \in \mathbb{R}$, almost surely, (4.20) holds for all $s, u \in \mathbb{R}$ such that $s \leq t \leq u$.

The following lemma, which is of interest in its own right, prepares for an alternative proof of the fact that $\hat\Phi_{t,s}$ is the generalised inverse of $\Phi_{s,t}$.

Lemma 4.27 (Evolution of halflines) *The maps $\mathcal{X}_{s,u}$ defined in (4.1) almost surely satisfy*

$$\mathcal{X}_{s,u}\big([-\infty, x]\big) = [-\infty, \pi^+_{(x,s)}(u)] \qquad (s \leq u,\ x \in \mathbb{R}).$$

Proof By definition, $y \in \mathcal{X}_{s,u}\big([-\infty, x]\big)$ if and only if there exists a $\hat{\pi} \in \hat{\mathcal{W}}(y,u)$ such that $\hat{\pi}(s) \leq x$. We need to show that $y \in \mathcal{X}_{s,u}\big([-\infty, x]\big)$ if and only if $y \leq \pi^+_{(x,s)}(u)$. If $y \leq \pi^+_{(x,s)}(u)$, then there exists a $\hat{\pi} \in \hat{\mathcal{W}}(y,u)$ such that $\hat{\pi}(t) \leq \pi^+_{(x,s)}(t)$ for all $t \in [s,u]$. In particular, $\hat{\pi}(s) \leq x$, so we conclude that $y \in \mathcal{X}_{s,u}\big([-\infty, x]\big)$. Conversely, if $y \in \mathcal{X}_{s,u}\big([-\infty, x]\big)$, then there exists a $\hat{\pi} \in \hat{\mathcal{W}}(y,u)$ such that $\hat{\pi}(s) \leq x$. It follows that there exists a $\pi \in \mathcal{W}(x,s)$ such that $\hat{\pi}(t) \leq \pi(t)$ for all $t \in [s,u]$. This implies $y = \hat{\pi}(u) \leq \pi(u) \leq \pi^+_{(x,s)}(u)$. ∎

Using Lemma 4.27, we can alternatively prove (4.24) by writing

$$\begin{aligned} \pi^+_{(x,s)}(t) < y \quad &\Leftrightarrow \quad \mathcal{X}_{s,t}\big([-\infty, x]\big) \cap [y, \infty] = \emptyset \\ &\Leftrightarrow \quad [-\infty, x] \cap \hat{\mathcal{Y}}_{t,s}\big([y,\infty]\big) = \emptyset \quad \Leftrightarrow \quad x < \hat{\pi}^-_{(y,t)}(t), \end{aligned}$$

where we have used the definition of the map $\hat{\mathcal{Y}}_{t,s}$ in (4.1) and the duality of Lemma 4.2.

We extend the maps $\Phi_{s,t}$ to $\overline{\mathbb{R}}$ by setting $\Phi_{s,t}(\pm\infty) := \pm\infty$, and similarly for $\hat{\Phi}_{t,s}$. For any set $A \subset \overline{\mathbb{R}}$, we let $\Phi_{s,t}(A) := \{\Phi_{s,t}(x) : x \in A\}$ denote the image of A under $\Phi_{s,t}$. We also let $\Phi^{-1}_{s,t}(A) := \{x \in \overline{\mathbb{R}} : \Phi_{s,t}(x) \in A\}$ denote the inverse image of a set $A \subset \overline{\mathbb{R}}$ under the map $\Phi_{s,t}$. Note that in view of Proposition 4.3, for $s < t$, the maps $\Phi_{s,t}$ are very much *not* one-to-one. Therefore, there is a big difference between the inverse image of a set under $\Phi_{s,t}$, and the image of the same set under the generalised inverse $\hat{\Phi}_{t,s}$. The following lemmas link these images and inverse images to the maps $\mathcal{X}_{s,t}, \mathcal{Y}_{s,t}, \hat{\mathcal{X}}_{t,s}, \hat{\mathcal{Y}}_{t,s}$ defined in (4.1). Below, for obvious notational reasons, we let

$$\mathrm{clos}(A) := \overline{A}$$

denote the closure of a set $A \subset \overline{\mathbb{R}}$.

Lemma 4.28 (Images) *For each deterministic $t \in \mathbb{R}$ and deterministic closed set $A \subset \overline{\mathbb{R}}$, one has almost surely*

$$\mathcal{Y}_{t,u}(A) = \Phi_{t,u}(A) \quad (t \leq u) \quad \textit{and} \quad \hat{\mathcal{Y}}_{t,s}(A) = \hat{\Phi}_{t,s}(A) \qquad (s \leq t).$$

Lemma 4.29 (Inverse images) *For each deterministic $t \in \mathbb{R}$, one has almost surely*

$$\mathcal{X}_{s,t}(A) = \mathrm{clos}\big(\hat{\Phi}^{-1}_{t,s}(A)\big) \quad \textit{and} \quad \hat{\mathcal{X}}_{u,t}(A) = \mathrm{clos}\big(\Phi^{-1}_{t,u}(A)\big)$$

for all $s \leq t$ and $u \geq t$ and for all closed sets $A \subset \overline{\mathbb{R}}$.

Proof of Lemma 4.28 By the symmetry between the Brownian web and the dual Brownian web, it suffices to prove only the statement for $\mathcal{Y}_{s,t}$. Filling in the definitions, we see that

$$\begin{aligned}\mathcal{Y}_{s,t}(A) &= \big\{\pi(t) : \pi \in \mathcal{W}(A \times \{s\})\big\},\\ \Phi_{s,t}(A) &= \big\{\pi^+_{(x,s)}(t) : x \in A\big\}.\end{aligned}$$

This immediately implies that the inclusion $\Phi_{s,t}(A) \subset \mathcal{Y}_{s,t}(A)$ holds almost surely for all $s \leq t$ and for all closed $A \subset \overline{\mathbb{R}}$. We will show that if s and A are deterministic, then the opposite inclusion also holds almost surely. Since s is deterministic, Theorem 4.17 tells us that almost surely $m_{\rm out}(x,s) \leq 2$ for all $x \in \mathbb{R}$, so

$$\mathcal{Y}_{s,t}(A) = \big\{\pi^\pm_{(x,s)}(t) : s \in A\big\} \quad \text{a.s.}$$

Thus, if the inclusion $\Phi_{s,t}(A) \subset \mathcal{Y}_{s,t}(A)$ is strict, then there exist $x \in A$ and $y \in \mathbb{R}$ such that $\pi^-_{(x,s)}(t) = y$ but $\pi^+_{(x',s)}(t) \neq y$ for all $x' \in A$. By Lemmas 4.6 and 4.7, this implies that there exists an $\varepsilon > 0$ such that $(x - \varepsilon, x) \cap A = \emptyset$. It is not hard to see that the set

$$\partial_- A := \big\{x \in A : \exists \varepsilon > 0 \text{ s.t. } (x-\varepsilon,x) \cap A = \emptyset\big\}$$

of "left boundary points" of A is countable. Since A is deterministic, Theorem 4.17 tells us that each point (x,s) with $x \in \partial_- A$ is almost surely of type $(0,1)$. This contradicts the fact that $\pi^-_{(x,s)}(t) = y$ but $\pi^+_{(x',s)}(t) \neq y$, so we conclude that the inclusion $\Phi_{s,t}(A) \subset \mathcal{Y}_{s,t}(A)$ is in fact an equality. ∎

Proof of Lemma 4.29 By the symmetry between the Brownian web and the dual Brownian web, it suffices to prove only the statement for $\mathcal{X}_{s,t}$. Filling in the definitions, we see that

$$\begin{aligned}\mathcal{X}_{s,t}(A) &= \big\{y \in \overline{\mathbb{R}} : \exists \hat\pi \in \hat{\mathcal{W}}(y,t) \text{ s.t. } \hat\pi(s) \in A\big\},\\ \hat\Phi^{-1}_{t,s}(A) &= \big\{y \in \overline{\mathbb{R}} : \hat\pi^+_{(y,t)}(s) \in A\big\}.\end{aligned}$$

It follows from Lemma 4.6 that

$$\hat\pi^-_{(y,s)} = \lim_{y' \uparrow y} \pi^+_{(y',s)} \quad \text{and} \quad \hat\pi^+_{(y,s)} = \lim_{y' \downarrow y} \pi^+_{(y',s)} \qquad \big((y,s) \in \mathbb{R}^2\big). \tag{4.25}$$

Since $y \in \mathrm{clos}\big(\hat\Phi^{-1}_{t,s}(A)\big)$ if and only if it is the limit of $y_n \in \hat\Phi^{-1}_{t,s}(A)$, using moreover Lemma 4.7, it follows that

$$\mathrm{clos}\big(\hat\Phi^{-1}_{t,s}(A)\big) = \big\{y \in \overline{\mathbb{R}} : \hat\pi^+_{(y,t)}(s) \in A\big\} \cup \big\{y \in \overline{\mathbb{R}} : \hat\pi^-_{(y,t)}(s) \in A\big\}. \tag{4.26}$$

It follows that $\mathrm{clos}\big(\hat\Phi^{-1}_{t,s}(A)\big) \subset \mathcal{X}_{s,t}(A)$, and this inclusion is strict if and only if there exists an $y \in \mathbb{R}$ such that $\hat\pi^-_{(y,t)}(s) \notin A$, $\hat\pi^+_{(y,t)}(s) \notin A$, but $\hat\pi(s) \in A$

for some $\pi \in \hat{\mathcal{W}}(y,t)$. This is possible only if (y,t) is of type $(2,1)$ (see Figure 4.5), but by Theorem 4.17 such points almost surely do not occur at deterministic times t. ∎

Exercise 4.30 *Recall the definition of points of types $(1,2)_{\rm l}$ and $(1,2)_{\rm r}$. For each $z \in \mathbb{R}$, define*

$$\pi_z^{\uparrow} := \begin{cases} \pi_z^- & \text{if } z \text{ is of type } (1,2)_{\rm l}, \\ \pi_z^+ & \text{otherwise.} \end{cases}$$

Modify the definition of the Arratia flow in (4.17) by replacing $\pi_{x,s}^+$ by $\pi_{x,s}^{\uparrow}$. Show that with this modified definition, $\Phi_{s,t}$ may fail to be right-continuous for some s,t, but on the other hand the stochastic flow property (4.20) now holds almost surely for all $s \leq t \leq u$ simultaneously.

Chapter 5

The Brownian net

5.1 Adding branching and deaths

As in Chapter 3, we let $\mathbb{Z}^2_{\rm even}$ and $\mathbb{Z}^2_{\rm odd}$ denote the even and odd sublattices of $\mathbb{Z}^2$. Generalising the set-up of Chapter 3, let $\omega = (\omega_z)_{z\in\mathbb{Z}^2_{\rm even}}$ be an i.i.d. collection of random variables that take values in the subsets of $\{-1,+1\}$. We can use ω to define a random directed graph with vertex set $\mathbb{Z}^2_{\rm even}$ and set of oriented edges

$$\vec{E} := \big\{\big((x,t),(x+y,t+1)\big) : (x,t)\in\mathbb{Z}^2_{\rm even},\ y\in\omega_{(x,t)}\big\}.$$

Generalising our earlier definition, we will call the random directed graph $(\mathbb{Z}^2_{\rm even},\vec{E})$ an *arrow configuration*. In particular, when ω_z takes the values $\{-1\}$ and $\{+1\}$ with equal probabilities, this is an arrow configuration as defined in Section 3.1. In the present chapter, we look at sequences ω^n of arrow configurations where $\omega^n = (\omega^n_z)_{z\in\mathbb{Z}^2_{\rm even}}$, for each $n\geq 1$, is a an i.i.d. collection with common law

$$\begin{aligned}&\mathbb{P}\big[\omega^n_z=\{-1\}\big]=l_n,\quad \mathbb{P}\big[\omega^n_z=\{+1\}\big]=r_n,\\ &\mathbb{P}\big[\omega^n_z=\{-1,+1\}\big]=b_n,\quad \mathbb{P}\big[\omega^n_z=\emptyset\big]=d_n.\end{aligned}\tag{5.1}$$

Here l_n is the probability that at a given point $z\in\mathbb{Z}^2_{\rm even}$, there starts (only) an arrow to the left, r_n is the probability of an arrow to the right, b_n is the branching probability, i.e., the probability that both arrows are present, and d_n is the death probability, i.e., the probability that no arrows are present.

Recall that σ_π and τ_π denote the starting time and final time of a path $\pi\in\Pi(\overline{\mathbb{R}})$. Generalising our definition from Section 3.1, we say that π is a *open path in the arrow configuration* ω^n if $\pi\in\Pi(\overline{\mathbb{R}})$ has following properties:

(i) $\big(\pi(t),t\big)\in\mathbb{Z}^2_{\rm even}$ $(t\in\mathbb{Z},\ t\geq\sigma_\pi)$,

(ii) $\pi(t+1)-\pi(t)\in\omega_{(\pi(t),t)}$ $(t\in\mathbb{Z},\ t\geq\sigma_\pi)$,

(iii) $\pi(t+s)=(1-s)\pi(t)+s\pi(t+1)$ $(0\leq s\leq 1,\ t\in\mathbb{Z},\ t\geq\sigma_\pi)$.

We let $\mathcal{V}_n$ denote the set of all open paths in ω^n. Note that even in the special case when $l_n=r_n=\frac{1}{2}$ and $b_n=d_n=0$, this is not quite the same object as the set $\mathcal{U}$ defined in Section 3.1, since we allow open paths to end at some final time $\tau_\pi<\infty$. We let $\overline{\mathcal{V}}_n$ denote the closure of $\mathcal{V}_n$ in $\Pi(\overline{\mathbb{R}})$. In this chapter, we will sketch a proof of the following theorem. Recall that θ_ε denotes the diffusive scaling map defined in (3.2).

Theorem 5.1 (The Brownian net with killing) *Let ε_n be positive constants tending to zero and let $\alpha\in\mathbb{R}$ and $\beta,\delta\in[0,\infty)$. Let ω_n be arrow configurations with probabilities l_n,r_n,b_n,d_n satisfying*

$$\varepsilon_n^{-1}(r_n-l_n)\underset{n\to\infty}{\longrightarrow}\alpha,\quad \varepsilon_n^{-1}b_n\underset{n\to\infty}{\longrightarrow}\beta,\quad\text{and}\quad \varepsilon_n^{-2}d_n\underset{n\to\infty}{\longrightarrow}\delta. \tag{5.2}$$

Let $\mathcal{V}_n$ be the set of open paths in the arrow configuration ω_n. Then

$$\mathbb{P}\big[\theta_{\varepsilon_n}(\overline{\mathcal{V}}_n)\in\cdot\,\big]\underset{n\to\infty}{\Longrightarrow}\mathbb{P}\big[\mathcal{N}_*\in\cdot\,\big], \tag{5.3}$$

where $\Rightarrow$ denotes weak convergence of probability laws on the space $\mathcal{K}(\Pi(\overline{\mathbb{R}}))$ of compact sets of paths, equipped with the Hausdorff topology, and $\mathcal{N}_*$ is a random compact subset of $\Pi(\overline{\mathbb{R}})$, whose law only depends on the parameters α,β,δ.

Exercise 5.2 *Show that the conditions (5.2) are equivalent to*

$$\left.\begin{aligned} l_n&=\tfrac12-\tfrac12(\beta+\alpha)\varepsilon_n+o(\varepsilon_n),\\ r_n&=\tfrac12-\tfrac12(\beta-\alpha)\varepsilon_n+o(\varepsilon_n),\\ b_n&=\beta\varepsilon_n+o(\varepsilon_n),\\ d_n&=\delta\varepsilon_n^2+o(\varepsilon_n^2),\end{aligned}\right\}\quad\text{as } n\to\infty.$$

In Theorem 5.1, the Brownian net with killing is obtained as the limit of branching-coalescing random walks in discrete time. A similar result is expected to hold for the collections of open paths in the graphical representations of biased voter models, though the details have nowhere been written down. This is similar to what we saw in Section 3.7.

For most of the chapter, we will be concerned with the case that $d_n=0$ for all n, and hence also $\delta=0$. This will allow us to work with the space $\Pi^\uparrow$ of upward paths as we are used to from Chapter 5. In Section 5.8, we will briefly indicate how the arguments can be generalised to allow for a positive death probability. For simplicity, in what follows, we will moreover focus on the case that $\alpha=0$ and $\beta=1$. In this case, the limiting object in (5.3) is known as the *standard Brownian net*.

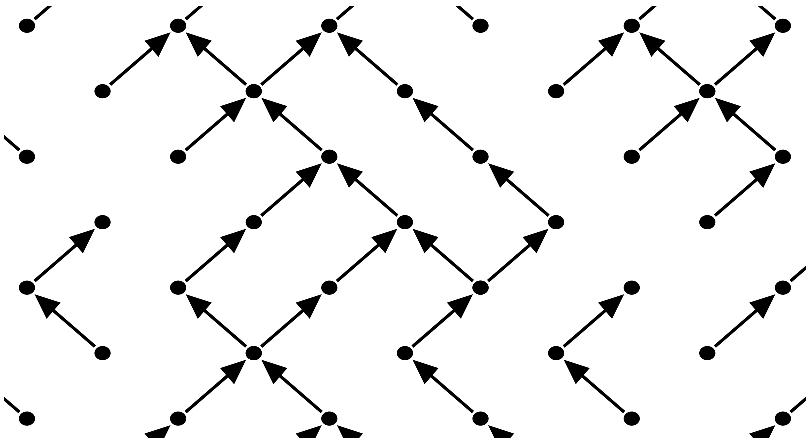

Figure 5.1: An arrow configuration with branching and deaths.

5.2 Left and right paths

We consider a sequence ω^n of arrow configurations as in the previous section with

$$d_n = 0, \quad \varepsilon_n^{-1}(r_n - l_n) \underset{n\to\infty}{\longrightarrow} 0 \quad \text{and} \quad \varepsilon_n^{-1} b_n \underset{n\to\infty}{\longrightarrow} 1. \tag{5.4}$$

We define $\mathcal{V}_n$ as in the previous section and set $\mathcal{U}_n := \mathcal{V}_n \cap \Pi^\uparrow$. Since the death probability is zero, $\mathcal{V}_n$ can simply be recovered from $\mathcal{U}_n$ by adding all shortened paths, that are cut off at an arbitrary time in $\mathbb{Z}$. Thus, all information is contained in the set $\mathcal{U}_n$ and we can continue to work with the space $\Pi^\uparrow$ that we are used to from the previous chapter.

We define collections of $\{-1,+1\}$-valued random variables

$$\omega^{\mathrm{l},n} = \left(\omega_z^{\mathrm{l},n}\right)_{z\in\mathbb{Z}^2_{\mathrm{even}}} \quad \text{and} \quad \omega^{\mathrm{r},n} = \left(\omega_z^{\mathrm{r},n}\right)_{z\in\mathbb{Z}^2_{\mathrm{even}}}$$

by

$$\omega_z^{\mathrm{l},n} := \begin{cases} -1 & \text{if } -1 \in \omega_z^n, \\ +1 & \text{otherwise,} \end{cases} \quad \text{and} \quad \omega_z^{\mathrm{r},n} := \begin{cases} +1 & \text{if } +1 \in \omega_z^n, \\ -1 & \text{otherwise.} \end{cases}$$

Then $\omega^{\mathrm{l},n}$ and $\omega^{\mathrm{r},n}$ correspond to arrow configurations of the type we used in Chapter 3 to approximate the Brownian web. They are constructed from ω^n by making a choice at each branching point z, in such a way that at each such point, $\omega_z^{\mathrm{l},n}$ only contains the left arrow and $\omega_z^{\mathrm{r},n}$ only contains the right arrow. We let $\mathcal{U}_n^{\mathrm{l}}$ and $\mathcal{U}_n^{\mathrm{r}}$ denote the sets of all paths in $\Pi^\uparrow$ that are open in the arrow configurations $\omega^{\mathrm{l},n}$ and $\omega^{\mathrm{r},n}$, respectively. We call $\mathcal{U}_n^{\mathrm{l}}$ and $\mathcal{U}_n^{\mathrm{r}}$ the collections of *left* and *right* open paths.

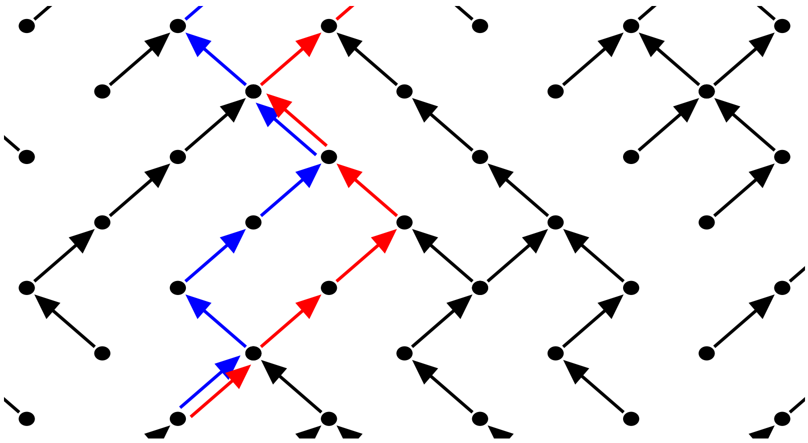

Figure 5.2: A left and a right open path in an arrow configuration with branching.

We also define dual arrow configurations

$$\hat{\omega}^{\mathrm{l},n} = \big(\hat{\omega}^{\mathrm{l},n}_z\big)_{z\in\mathbb{Z}^2_{\mathrm{odd}}} \quad \text{and} \quad \hat{\omega}^{\mathrm{r},n} = \big(\hat{\omega}^{\mathrm{r},n}_z\big)_{z\in\mathbb{Z}^2_{\mathrm{odd}}}$$

in terms of $\omega^{\mathrm{l},n}$ and $\omega^{\mathrm{r},n}$ as in (3.17). Then $\hat{\omega}^{\mathrm{l},n}$ is equally distributed with $\omega^{\mathrm{l},n}$ after a rotation over 180 degrees (but not after mirroring in the horizontal axis!) and a shift to the odd sublattice, and the same is true for $\hat{\omega}^{\mathrm{r},n}$ and $\omega^{\mathrm{r},n}$. We define $\hat{\omega}^n = \big(\hat{\omega}^n_z\big)_{z\in\mathbb{Z}^2_{\mathrm{odd}}}$ by

$$\hat{\omega}^n_z := \{\hat{\omega}^{\mathrm{l},n}_z, \hat{\omega}^{\mathrm{r},n}_z\}.$$

Finally, we let

$$\mathcal{U}^*_n, \quad \mathcal{U}^{\mathrm{l}*}_n, \quad \text{and} \quad \mathcal{U}^{\mathrm{r}*}_n$$

denote all paths in $\Pi^\downarrow$ that are open in the arrow configurations $\hat{\omega}^n$, $\hat{\omega}^{\mathrm{l},n}$, and $\hat{\omega}^{\mathrm{r},n}$, respectively. Figure 5.2 shows a left and right path in an arrow configuration with branching, and Figure 5.3 shows the corresponding left and right arrow configurations and their duals.

We define a *Brownian web with drift* α in the same way as the standard Brownian web, except that the coalescing Brownian motions now have drift α. If $\mathcal{W}$ is a standard Brownian web, then we can construct a Brownian web with drift α by setting $\mathcal{W}' := \{\pi' : \pi \in \mathcal{W}\}$ with $\pi' := \{(x+\alpha t, t) : (x,t) \in \pi\}$.

Theorem 5.3 (Scaling limit of left and right paths) *Let ε_n be positive constants, tending to zero, let $\mathcal{U}^{\mathrm{l}}_n, \mathcal{U}^{\mathrm{r}}_n$ be the collections of left and right paths*

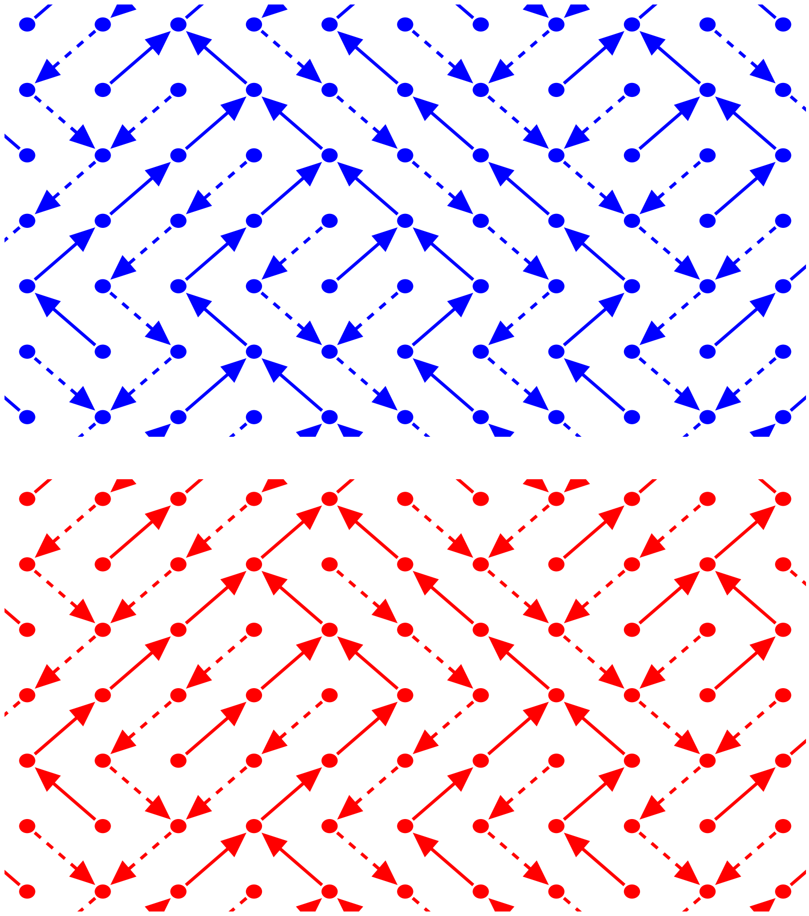

Figure 5.3: A left arrow configuration (blue) and a right arrow configuration (red), together with their dual arrow configurations.

in arrow configurations ω^n *satisfying (5.4), and* $\mathcal{U}_n^{\mathrm{l}*}, \mathcal{U}_n^{\mathrm{r}*}$ *be the associated dual left and right paths. Then*

$$\mathbb{P}\big[\theta_{\varepsilon_n}(\overline{\mathcal{U}}^{\mathrm{l}}_n, \overline{\mathcal{U}}^{\mathrm{l}*}_n) \in \cdot\,\big] \underset{n\to\infty}{\Longrightarrow} \mathbb{P}\big[(\mathcal{W}^{\mathrm{l}}, \hat{\mathcal{W}}^{\mathrm{l}}) \in \cdot\,\big],$$
$$\mathbb{P}\big[\theta_{\varepsilon_n}(\overline{\mathcal{U}}^{\mathrm{r}}_n, \overline{\mathcal{U}}^{\mathrm{r}*}_n) \in \cdot\,\big] \underset{n\to\infty}{\Longrightarrow} \mathbb{P}\big[(\mathcal{W}^{\mathrm{r}}, \hat{\mathcal{W}}^{\mathrm{r}}) \in \cdot\,\big],$$

where $\Rightarrow$ *denotes weak convergence of probability laws on the space* $\mathcal{K}(\Pi^{\uparrow}) \times \mathcal{K}(\Pi^{\downarrow})$*, and* $(\mathcal{W}^{\mathrm{l}}, \hat{\mathcal{W}}^{\mathrm{l}})$ *and* $(\mathcal{W}^{\mathrm{r}}, \hat{\mathcal{W}}^{\mathrm{r}})$ *are double Brownian webs with drift* -1 *and* $+1$*, respectively.*

Proof The proof of Theorem 3.18 carries over with only a very minor change: when we prove convergence of finite dimensional distributions as in Proposition 3.3, the limit is a system of coalescing Brownian motions with drift -1

or +1, respectively. Indeed, letting L^n_z and R^n_z denote the unique left and right paths in $\mathcal{U}_n$ starting from a point $z \in \mathbb{Z}^2_{\rm even}$, we observe that

$$\begin{aligned}\mathbb{E}\big[L^n_z(t+1) - L^n_z(t)\big] &= r_n - l_n - b_n \sim -\varepsilon_n,\\ \mathbb{E}\big[R^n_z(t+1) - L^n_z(t)\big] &= r_n - l_n + b_n \sim +\varepsilon_n\end{aligned}$$

as $n \to \infty$, which is easily seen to imply that L^n_z and R^n_z converge after diffusive rescaling to Brownian motions with drift −1 and +1, respectively. ∎

Theorem 5.3 does not tell us anything about the limit of the joint law of $\overline{\mathcal{U}}^{\rm l}_n$ and $\overline{\mathcal{U}}^{\rm r}_n$. It turns out that in the limit, the interacting between left and right paths is a form of sticky reflection. In view of this, in the next section, we will study sticky reflection. This will then be used to prove a result about the scaling limit of the joint law of $\overline{\mathcal{U}}^{\rm l}_n$ and $\overline{\mathcal{U}}^{\rm r}_n$, which in the end will be used to prove Theorem 5.1, first under the more restrictive assumptions (5.4), and then generally.

5.3 Sticky reflection

In this section, we study sticky reflection. This is similar to Skorohod reflection (Lemma 4.9), but a bit more complicated. Recall the definition of the function spaces $\mathcal{C}, \mathcal{C}_0, \mathcal{C}^+_0$, and $\mathcal{C}_{\rm pos}$ in Section 4.2. We also define $\mathcal{C}^1_0 \subset \mathcal{C}^+_0$ by

$$\mathcal{C}^1_0 := \big\{f \in \mathcal{C}_0 : 0 \leq f(t) - f(s) \leq t - s \ \forall 0 \leq s \leq t\big\}.$$

Let $g(0) \in [0,\infty)$, $f \in \mathcal{C}_0$, and $h \in \mathcal{C}^+_0$ be given. By definition, a solution to the *sticky reflection equation*[1]

$$\mathrm{d}g(t) = \mathrm{d}f(T_t) + \mathrm{d}h(S_t) \qquad (t \geq 0) \tag{5.5}$$

is a triple (g, S, T) of functions $g \in \mathcal{C}_{\rm pos}$ and $S, T \in \mathcal{C}^1_0$ such that

(i) $g(t) = g(0) + f(T_t) + h(S_t) \quad (t \geq 0)$,

(ii) $\displaystyle\int_0^\infty 1_{\{g(t)>0\}}\mathrm{d}h(S_t) = 0$,

(iii) $S_t + T_t = t \quad (t \geq 0)$.

[1]This is my terminology. I do not know if this precise definition has been invented before, although there is an extensive literature on diffusion processes with sticky reflection.

If we replace the condition $g \in \mathcal{C}_{\rm pos}$ by the weaker condition $g \in \mathcal{C}$, then we say that (g, S, T) is a *signed* solution to the sticky reflection equation (5.5). Our first result says that sticky reflection equations have solutions, and that under mild conditions, such solutions are unique. Note that if h is strictly increasing, then the set H defined below is empty and hence the condition $H \cap M = \emptyset$ is trivially fulfilled.

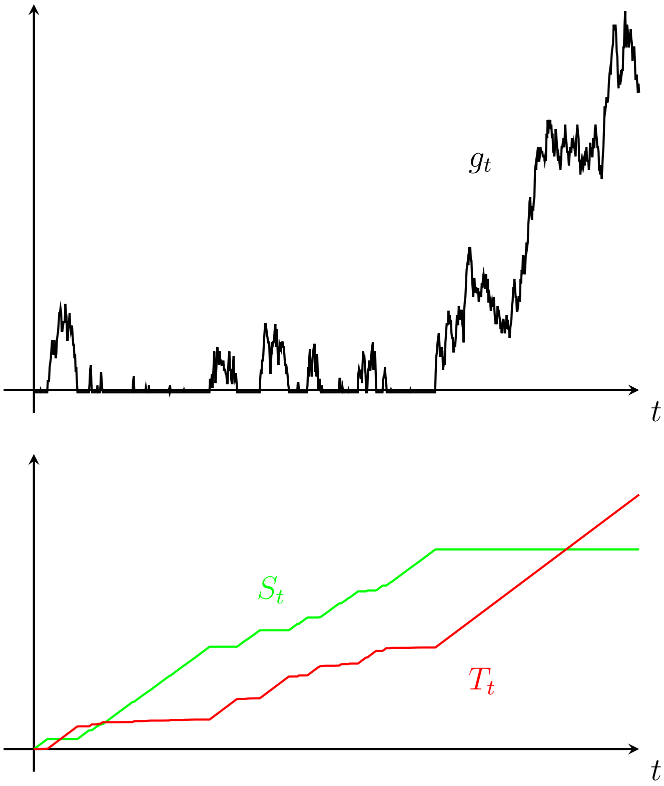

Figure 5.4: Sticky reflected Brownian motion: the solution (g, S, T) to the sticky reflection equation (5.5) in the case that f is a Brownian path and $h(t) := t$ $(t \geq 0)$ is the identity function.

Proposition 5.4 (Sticky reflection) *For each $g(0) \in [0, \infty)$, $f \in \mathcal{C}_0$, and $h \in \mathcal{C}_0^+$, there exists a solution (g, S, T) to the sticky reflection equation (5.5). Let $\tilde{f}(t) := g(0) + f(t)$ $(t \geq 0)$ and set*

$$H := \big\{r \geq 0 : \exists s < t \text{ s.t. } h(s) = r = h(t)\big\},$$
$$M := \big\{r \geq 0 : \exists s < t \text{ s.t. } m_s(\tilde{f}) = -r = m_t(\tilde{f})\big\},$$

where $m_t(\tilde f)$ is defined in (4.6). Then solutions to the sticky reflection equation (5.5) are unique if and only if $H \cap M = \emptyset$.

Proof (sketch) Let $\tilde f(t) := g(0) + f(t)$ $(t \geq 0)$. We claim that (g, S, T) with $g \in \mathcal{C}_{\rm pos}$ and $S, T \in \mathcal{C}^1_0$ solves the sticky reflection equation (5.5) if and only if

(i)' $g(t) = \tilde f(T_t) - m_{T_t}(\tilde f) \quad (t \geq 0)$,

(ii)' $h(S_t) = -m_{T_t}(\tilde f) \quad (t \geq 0)$,

(iii) $S_t + T_t = t \quad (t \geq 0)$,

To prove this, assume that (g, S, T) solves the sticky reflection equation (5.5). Set $F(t) := f(T_t)$ and $\psi(t) := h(S_t)$ $(t \geq 0)$. Conditions (i) and (ii) of the definition of a solution to the sticky reflection equation then say that (g, ψ) solves the Skorohod reflection equation

$$\mathrm{d}g(t) = \mathrm{d}F(t) + \mathrm{d}\psi(t) \qquad (t \geq 0). \tag{5.6}$$

Applying Lemma 4.9 to (5.6), we see that any solution (g, S, T) to the sticky reflection equation (5.5) satisfies (i)'. Combining (i) and (i)', we see that moreover (ii)' holds. Assume, conversely, that (g, S, T) with $g \in \mathcal{C}_{\rm pos}$ and $S, T \in \mathcal{C}^1_0$ satisfy conditions (i)', (ii)', and (iii). Set $F(t) := f(T_t)$ and $\psi(t) := h(S_t)$ $(t \geq 0)$. Then (ii)' implies $\psi(t) = -m_{T_t}(\tilde f)$ $(t \geq 0)$ and therefore (i)' and Lemma 4.9 imply that (g, ψ) solves the Skorohod reflection equation (5.6), which implies that (g, S, T) satisfies conditions (i) and (ii) of the definition of a solution to the sticky reflection equation.

We see immediately from (i)', (ii)', and (iii) that if $S \in \mathcal{C}^1_0$ satisfies

$$h(S_t) + m_{t-S_t}(\tilde f) = 0 \qquad (t \geq 0), \tag{5.7}$$

then setting

$$T_t := t - S_t \quad \text{and} \quad g(t) := \tilde f(T_t) - m_{T_t}(\tilde f) \qquad (t \geq 0)$$

yields a solution (g, S, T) to the sticky reflection equation (5.5), and each solution is of this form. This motivates us to define

$$\begin{aligned} S^-_t &:= \inf\big\{s \in [0, t] : h(s) + m_{t-s}(\tilde f) = 0\big\}, \\ S^+_t &:= \sup\big\{s \in [0, t] : h(s) + m_{t-s}(\tilde f) = 0\big\}. \end{aligned}$$

We claim that

(a) $S^-, S^+ \in \mathcal{C}^1_0$.

(b) A function $S \in \mathcal{C}_0^1$ satisfies (5.7) if and only if $S_t^- \leq S \leq S_t^+$ $(t \geq 0)$.

From this, it is not hard to see that solutions to the sticky reflection equation (5.5) always exist (since we can take $S = S^-$ or $= S^+$), and that they are unique if and only if $S_t^- = S_t^+$ for all $t \geq 0$, which is easily seen to be equivalent to the condition $H \cap M = \emptyset$. ∎

Exercise 5.5 *Prove the claims (a) and (b) in the proof of Proposition 5.4.*

Our next result says that under suitable conditions, the solution (g, S, T) of a sticky reflection equation of the form (5.5) depends continuously on the initial state $g(0)$ and the driving processes f and h. In what follows, we will see that in applications of this proposition, it will be important that we allow the approximating solutions to be only signed solutions, i.e., the g_n may take negative values.

Proposition 5.6 (Continuous parameter dependence) *Assume that $f_n, f \in \mathcal{C}_0$ and $h_n, h \in \mathcal{C}_0^+$ satisfy $f_n \to f$ and $h_n \to h$ locally uniformly. For each n, let (g_n, S^n, T^n) be a signed solution to the sticky reflection equation*

$$\mathrm{d}g_n(t) = \mathrm{d}f_n(T_t^n) + \mathrm{d}h_n(S_t^n) \qquad (t \geq 0).$$

Assume that

$$g_n(0) \underset{n\to\infty}{\longrightarrow} g(0) \quad \text{and} \quad \liminf_{n\to\infty} g_n(t) \geq 0 \quad (t \geq 0).$$

Assume moreover that the sticky reflection equation

$$\mathrm{d}g(t) = \mathrm{d}f(T_t) + \mathrm{d}h(S_t) \qquad (t \geq 0)$$

has a unique solution (g, S, T) with initial state $g(0)$. Then one has

$$g_n \to g, \quad S^n \to S, \quad \text{and} \quad T^n \to T$$

locally uniformly.

The proof of Proposition 5.6 depends on two lemmas. We first state the lemmas, then show how they imply Proposition 5.6, and finally prove the lemmas.

Lemma 5.7 (Precompactness of solutions) *Let $\mathcal{A} \subset [0,\infty) \times \mathcal{C}_0 \times \mathcal{C}_0^+$. Define $\mathcal{B} \subset \mathcal{C} \times \mathcal{C}_0^1 \times \mathcal{C}_0^1$ to be the set of all triples (g, S, T) that are a signed solution of a sticky reflection equation of the form*

$$\mathrm{d}g(t) = \mathrm{d}f(T_t) + \mathrm{d}h(S_t) \qquad (t \geq 0)$$

with $(g(0), f, h) \in \mathcal{A}$. If $\mathcal{A}$ is a precompact subset of $[0,\infty) \times \mathcal{C}^2$ (equipped with the product topology), then $\mathcal{B}$ is a precompact subset of $\mathcal{C}^3$.

Lemma 5.8 (Limits of solutions) *Assume that $f_n, f \in \mathcal{C}_0$ and $h_n, h \in \mathcal{C}_0^+$ satisfy $f_n \to f$ and $h_n \to h$ locally uniformly. For each n, let (g_n, S^n, T^n) be a signed solution to the sticky reflection equation*

$$\mathrm{d}g_n(t) = \mathrm{d}f_n(T^n_t) + \mathrm{d}h_n(S^n_t) \qquad (t \geq 0),$$

and assume that $g_n \to g$, $S^n \to S$, and $T^n \to T$ locally uniformly for some $g, S, T \in \mathcal{C}$. Then (g, S, T) is a signed solution to the sticky reflection equation

$$\mathrm{d}g(t) = \mathrm{d}f(T_t) + \mathrm{d}h(S_t) \qquad (t \geq 0).$$

Proof of Proposition 5.6 By Lemma 2.2, our assumptions imply that the set

$$\big\{(g_n(0), f_n, h_n) : n \in \mathbb{N}\big\}$$

is a precompact subset of $[0, \infty) \times \mathcal{C}^2$. By Lemma 5.7, this implies that the set

$$\big\{(g_n, S^n, T^n) : n \in \mathbb{N}\big\}$$

is a precompact subset of $\mathcal{C}^3$. Lemma 5.8 implies that each subsequential limit of the sequence $(g_n, S^n, T^n)_{n\in\mathbb{N}}$ is a signed solution to the limiting sticky reflection equation

$$\mathrm{d}g(t) = \mathrm{d}f(T_t) + \mathrm{d}h(S_t) \qquad (t \geq 0).$$

Our condition $\liminf_{n\to\infty} g_n(t) \geq 0$ $(t \geq 0)$ implies that $g \in \mathcal{C}_{\mathrm{pos}}$, so our signed solution is in fact a true, nonnegative solution. By assumption, the limiting sticky reflection equation has a unique solution (g, S, T). Therefore, we can apply Lemma 2.2 to conclude that $(g_n, S^n, T^n) \to (g, S, T)$. ∎

Recall from (2.16) that for each $T < \infty$, the *modulus of continuity* of a function $f \in \mathcal{C}_{[0,\infty)}(\mathbb{R})$ is given by

$$m_{K,\delta}(f) = \sup\big\{\big|f(s) - f(t)\big| : 0 \leq s \leq t \leq K,\ t - s \leq \delta\big\},$$

and that a set $\mathcal{D} \subset \mathcal{C}_{[0,\infty)}(\mathbb{R})$ is *equicontinuous* if

$$\lim_{\delta\to 0} \sup_{f\in\mathcal{D}} m_{K,\delta}(f) = 0 \qquad (K < \infty).$$

By the Arzela-Ascoli theorem (see Theorem 2.30 and Lemma 2.27), a subset $\mathcal{D} \subset \mathcal{C}_{[0,\infty)}(\mathbb{R})$ is precompact if and only if

(a) $\mathcal{D}$ is equicontinuous,

(b) For each $K < \infty$, there exists a $C < \infty$ such that $|f(t)| \leq C$ for all $t \in [0, K]$.

Proof of Lemma 5.7 It suffices to show that each of the sets

$$\{S : (g, S, T) \in \mathcal{B}\}, \quad \{T : (g, S, T) \in \mathcal{B}\}, \quad \text{and} \quad \{g : (g, S, T) \in \mathcal{B}\}$$

is equicontinuous and satisfies the compact containment condition (b) above.

For each $(g, S, T) \in \mathcal{B}$, since $S, T \in \mathcal{C}_0^1$ it is clear that $\{S : (g, S, T) \in \mathcal{B}\}$ and $\{T : (g, S, T) \in \mathcal{B}\}$ are equicontinuous and satisfy the compact containment condition (b).

Using the fact that $S, T \in \mathcal{C}_0^1$, we see that if $0 \leq s \leq t \leq K$ satisfy $t - s \leq \delta$, then $0 \leq S_s \leq S_t \leq K$ and $S_t - S_s \leq \delta$, and likewise with S replaced by T. By the definition of a signed solution, each $(g, S, T) \in \mathcal{B}$ solves an equation of the form

$$g(t) = g(0) + f(T_t) + h(S_t) \qquad (t \geq 0), \tag{5.8}$$

for some $(g(0), f, h) \in \mathcal{A}$. Using (5.8) and our previous observations about S and T, we see that

$$m_{K,\delta}(g) \leq m_{K,\delta}(f) + m_{K,\delta}(h) \qquad (K < \infty,\ \delta > 0). \tag{5.9}$$

In view of this, the equicontinuity of $\{g : (g, S, T) \in \mathcal{B}\}$ follows from the equicontinuity of $\{f : (x, f, g) \in \mathcal{A}\}$ and $\{h : (x, f, g) \in \mathcal{A}\}$, which is a result of the Arzela-Ascoli theorem and our assumption that $\mathcal{A}$ is precompact. Using (5.8) once more, we can estimate

$$\big|g(t)\big| \leq \big|g(0)\big| + \sup_{s\in[0,t]} \big|f(s)\big| + \big|h(t)\big| \qquad (t \geq 0).$$

The precompactness of $\mathcal{A}$ implies that $\{x : (x, f, h) \in \mathcal{A}\}$ is bounded. By the Arzela-Ascoli theorem, the precompactness of $\mathcal{A}$ also implies that the sets $\{f : (x, f, h) \in \mathcal{A}\}$ and $\{h : (x, f, h) \in \mathcal{A}\}$ satisfy the compact containment condition (b) above. Using this and our estimate, we see that $\{g : (g, S, T) \in \mathcal{B}\}$ satisfies the compact containment condition (b). ∎

Proof of Lemma 5.8 Since $S^n, T^n \in \mathcal{C}_0^1$ for each n, taking the limit, we see that $S, T \in \mathcal{C}_0^1$. Since

$$g_n(t) = g_n(0) + f_n(T^n(t)) + h_n(S^n(t)) \quad (t \geq 0)$$

for all n, taking the limit, we see that (g, S, T) satisfies condition (i) of the definition of a signed solution to the sticky reflection equation. Similarly, since $S^n(t) + T^n(t) = t \quad (t \geq 0)$ for all n, taking the limit, we see that condition (iii) is satisfied. It remains to prove that (g, S, T) satisfies condition (ii).

For each $\varepsilon > 0$, we can find a continuous function $\rho_\varepsilon : [0,\infty) \to [0,1]$ such that $\rho_\varepsilon(0) = 0$ and $\rho_\varepsilon(x) = 1$ for all $t \geq \varepsilon$. Then

$$\int_0^t \rho_\varepsilon(g_n(s))\mathrm{d}h_n(S^n_s) \leq \int_0^\infty 1_{\{g_n(s)>0\}}\mathrm{d}h_n(S^n_s) = 0$$

for all $t \in [0,\infty)$, $\varepsilon > 0$, and n. Using Lemma 5.9 below, it follows that

$$\int_0^t 1_{\{g(s)\geq\varepsilon\}}\mathrm{d}h(S_s) \leq \int_0^t \rho_\varepsilon(g(s))\mathrm{d}h(S_s) = \lim_{n\to\infty}\int_0^t \rho_\varepsilon(g_n(s))\mathrm{d}h_n(S^n_s) = 0$$

for all $t \in [0,\infty)$ and $\varepsilon > 0$. Letting first $\varepsilon \downarrow 0$ and then $t \uparrow \infty$, using dominated convergence and monotone convergence, we see that g and S satisfy condition (ii) of the definition of a solution to the sticky reflection equation. ∎

Lemma 5.9 (Convergence of integrals) *Let $t > 0$, let $F_n, G_n, F, G \in \mathcal{C}_{[0,t]}(\mathbb{R})$ satisfy $F_n \to F$ and $G_n \to G$ uniformly, and assume that F_n, F are nondecreasing. Then*

$$\int_0^t G_n(s)\,\mathrm{d}F_n(s) \underset{n\to\infty}{\longrightarrow} \int_0^t G(s)\,\mathrm{d}F(s).$$

Proof Let $\mu_n = \mathrm{d}F_n$ and $\mu = \mathrm{d}F$, i.e., μ is the unique finite measure on $[0,t]$ such that $\mu([0,s]) = F(s)$ $(s \in [0,t])$, and similarly F_n is the "distribution function" of μ_n. It is well-known that a sequence of finite measures μ_n on $[0,t]$ converge to a limit μ if and only if their distribution functions satisfy $F_n(t) \to F(t)$ for each continuity point t of F. In particular, our condition that $F_n \to F$ uniformly implies that $\mu_n \Rightarrow \mu$. It follows that

$$\int_0^t G(s)\,\mathrm{d}F_n(s) \underset{n\to\infty}{\longrightarrow} \int_0^t G(s)\,\mathrm{d}F(s).$$

Since

$$\begin{aligned}\int_0^t G(s)\,\mathrm{d}F_n(s) - \|G - G_n\|_\infty F_n(t) &\leq \int_0^t G_n(s)\,\mathrm{d}F_n(s),\\ \int_0^t G_n(s)\,\mathrm{d}F_n(s) &\leq \int_0^t G(s)\,\mathrm{d}F_n(s) + \|G - G_n\|_\infty F_n(t),\end{aligned}$$

we see that

$$\begin{aligned}\int_0^t G(s)\,\mathrm{d}F(s) &\leq \liminf_{n\to\infty}\int_0^t G(s)\,\mathrm{d}F_n(s),\\ \limsup_{n\to\infty}\int_0^t G(s)\,\mathrm{d}F_n(s) &\leq \int_0^t G(s)\,\mathrm{d}F(s).\end{aligned}$$

∎

5.4 Sticky reflected random walk

In the next section, we will show that pairs of random walks, consisting of one left path and one right path in arrow configurations satisfying (5.4), converge in the diffusive scaling limit to pairs of drifted Brownian motions with a form of sticky reflection. In the present section, we give a simpler application of Propositions 5.4 and 5.6. This will not be needed in what follows, but serves as a useful illustration of the main ideas and a warm-up for the next section.

Fix positive constants ε_n, tending to zero, and for each n, let $(X^n_k)_{k\geq 0}$ be a Markov chain on $\{-1,0,1,2,\ldots\}$ with transition kernel P_n given by

$$\begin{aligned} &P_n(-1,0) := \varepsilon_n, \quad P_n(-1,-1) := 1-\varepsilon_n, \\ &P_n(k,k-1) = P(k,k+1) = \tfrac{1}{2} \quad (k \geq 0). \end{aligned}$$

We extend X^n to all real times $t \geq 0$ by linear interpolation and define a diffusively rescaled process $X^{(n)}$ by

$$X^{(n)}(\varepsilon_n^2 t) := \varepsilon_n X^n(t) \qquad (t \geq 0).$$

We will prove the following theorem. Note that (5.10) is a special case of (5.5) where the function h from (5.5) is the identity function $h(t) = t$ $(t \geq 0)$.

Theorem 5.10 (Scaling limit of sticky reflected random walk) *Assume that $X^{(n)}_0 \geq 0$ is deterministic and $X^{(n)}_0 \to x_0$ as $n \to \infty$. Let $(X_t, S_t, T_t)_{t\geq 0}$ be the a.s. unique solution with initial condition $X_0 = x_0$ of the sticky reflection equation*

$$\mathrm{d}X_t = \mathrm{d}B_{T_t} + \mathrm{d}S_t \qquad (t \geq 0), \tag{5.10}$$

where $(B_t)_{t\geq 0}$ is a standard Brownian motion. Then one has

$$\mathbb{P}\big[\big(X^{(n)}(t)\big)_{t\geq 0} \in \cdot\,\big] \underset{n\to\infty}{\Longrightarrow} \mathbb{P}\big[\big(X_t\big)_{t\geq 0} \in \cdot\,\big],$$

where $\Rightarrow$ denotes weak convergence of probability measures on $\mathcal{C}_{[0,\infty)}(\mathbb{R})$, equipped with the topology of locally uniform convergence.

Our proof strategy will be to relate X^n to a solution of a sticky reflection equation driven by processes F and H^n that have as diffusive scaling limits Brownian motion and the identity function $h(t) = t$ $(t \geq 0)$. Theorem 5.10 will then follow as an application of Propositions 5.4 and 5.6.

Let $(\omega_i)_{i\geq 1}$ be i.i.d. uniformly distributed on $\{-1,+1\}$ and let $(F_k)_{k\geq 0}$ be the random walk defined by

$$F_k := \sum_{i=1}^{k} \omega_i \qquad (k \in \mathbb{N}).$$

Fix positive constants ε_n, tending to zero, and for each n, let $(H^n_k)_{k\geq 0}$ be a Markov chain with initial state $H^n_0 := 0$ and transition kernel Q^n given by

$$Q^n(k,k+1) := \varepsilon_n \quad \text{and} \quad Q^n(k,k) := 1-\varepsilon_n \qquad (k\in\mathbb{N}).$$

We inductively define a process $(X^n_k, S^n_k, T^n_k)_{k\geq 0}$ with $S^n_0 = 0 = T^n_0$ by

$$\begin{aligned} X^n_{k+1} &:= X^n_k + \big(F(T^n_{k+1}) - F(T^n_k)\big) + \big(H^n(S^n_{k+1}) - H^n(S^n_k)\big) \\ S^n_{k+1} &:= S^n_k + 1_{\{X^n_k=-1\}} \quad \text{and} \quad T^n_{k+1} := T^n_k + 1_{\{X^n_k\geq 0\}} \end{aligned}$$

$(k\geq 0)$. These definitions are illustrated in Figure 5.5. It is not hard to see that $(X^n_k)_{k\geq 0}$ is the Markov chain with transition kernel P_n defined above.

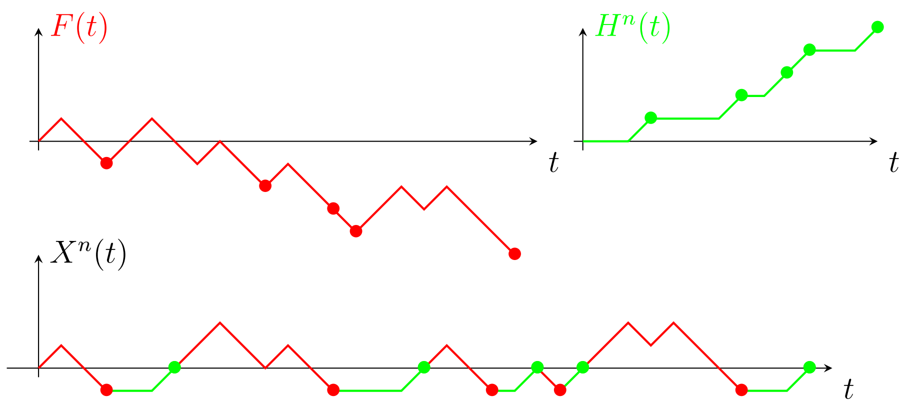

Figure 5.5: A sticky reflected random walk X^n constructed from a random walk F and a reflection process H^n.

We interpolate the processes F, H^n and X^n, S^n, T^n linearly between integer times. Then it is easy to see that

(i) $X^n(t) = X^n(0) + F(T^n_t) + H^n(S^n_t) \qquad (t\geq 0),$

(ii) $\displaystyle\int_0^\infty 1_{\{X^n(t)>0\}} \mathrm{d}H^n(S^n_t) = 0,$

(iii) $S^n_t + T^n_t = t \qquad (t\geq 0),$

i.e., (X^n, S^n, T^n) is a signed solution to the sticky reflection equation

$$\mathrm{d}X^n(t) = \mathrm{d}F(T^n_t) + \mathrm{d}H^n(S^n_t) \qquad (t\geq 0).$$

We define rescaled processes $X^{(n)}, S^{(n)}, T^{(n)}, F^{(n)}$, and $H^{(n)}$ by

$$X^{(n)}(\varepsilon_n^2 t) := \varepsilon_n X^n(t), \quad F^{(n)}(\varepsilon_n^2 t) := \varepsilon_n F^n(t), \quad H^{(n)}(\varepsilon_n^2 t) := \varepsilon_n H^n(t),$$
$$S^{(n)}_{\varepsilon_n^2 t} := \varepsilon_n^2 S^n_t, \quad T^{(n)}_{\varepsilon_n^2 t} := \varepsilon_n^2 T^n_t.$$

$(t \geq 0)$. Note that we rescale the processes diffusively: time is rescaled by ε_n^2 and space is rescaled by ε_n. Here $X^{(n)}, F^{(n)}$, and $H^{(n)}$ are functions from time to space, but $S^{(n)}, T^{(n)}$ are functions that map times into times (hence the at first sight different scaling). It is straightforward to check that the rescaled processes $(X^{(n)}, S^{(n)}, T^{(n)})$ is a signed solution to the sticky reflection equation

$$\mathrm{d}X^{(n)}(t) = \mathrm{d}F^{(n)}(T^{(n)}_t) + \mathrm{d}H^{(n)}(S^{(n)}_t) \qquad (t \geq 0).$$

Proof of Theorem 5.10 It follows from Donsker's invariance principle that

$$\mathbb{P}\big[\big(F^{(n)}(t)\big)_{t\geq 0} \in \cdot\,\big] \underset{n\to\infty}{\Longrightarrow} \mathbb{P}\big[\big(B_t\big)_{t\geq 0} \in \cdot\,\big],$$

where $\Rightarrow$ denotes weak convergence of probability measures on $\mathcal{C}_{[0,\infty)}(\mathbb{R})$, equipped with the topology of locally uniform convergence, and $(B_t)_{t\geq 0}$ is a standard Brownian motion. Using the weak law of large numbers, it is not hard to show that moreover

$$\mathbb{P}\big[\sup_{t\in[0,K]} \big|H^{(n)}(t) - t\big| \geq \delta\big] \underset{n\to\infty}{\longrightarrow} 0 \qquad \big(K < \infty,\ \delta > 0\big).$$

In other words, the process $(H^{(n)}(t))_{t\geq 0}$ converges to the identity function in probability with respect to the topology of locally uniform convergence. Equivalently, this says that the law of $(H^{(n)}(t))_{t\geq 0}$ converges weakly to the delta-measure on the identity function $I_t := t$ $(t \geq 0)$.

By Skorohod's representation theorem, we can couple our random variables such that

$$\big(F^{(n)}(t)\big)_{t\geq 0} \underset{n\to\infty}{\longrightarrow} \big(B_t\big)_{t\geq 0} \quad \text{and} \quad \big(H^{(n)}(t)\big)_{t\geq 0} \underset{n\to\infty}{\longrightarrow} \big(I_t\big)_{t\geq 0} \quad \text{a.s.}$$

Using moreover that $X^{(n)}_0 \to x_0$, we can apply Proposition 5.6 to conclude that almost surely

$$X^{(n)} \to B, \quad S^{(n)} \to S, \quad \text{and} \quad T^{(n)} \to T$$

locally uniformly, where (X, S, T) is the unique solution of (5.10). Since almost sure convergence implies convergence in law, the claim of the theorem follows. ∎

5.5 The interaction of left and right paths

In this section, we return to the left and right paths introduced in Section 5.2. Our aim is to describe the joint law of one left and one right path. We fix a sequence ε_n of positive constants, tending to zero, and let ω^n be a sequence of arrow configurations satisfying (5.4). We let $\mathcal{U}^{\rm l}_n$ and $\mathcal{U}^{\rm r}_n$ denote the collections of left and right open paths in ω^n, and for each $z \in \mathbb{Z}^2_{\rm even}$, we let L^n_z and R^n_z denote the unique elements of $\mathcal{U}^{\rm l}_n(z)$ and $\mathcal{U}^{\rm r}_n(z)$, respectively. We fix two sequences of even integers $x^{\rm l}_n$ and $x^{\rm r}_n$ with $x^{\rm l}_n \leq x^{\rm r}_n$ for each n such that

$$\varepsilon_n x^{\rm l}_n \underset{n\to\infty}{\longrightarrow} x^{\rm l} \quad \text{and} \quad \varepsilon_n x^{\rm r}_n \underset{n\to\infty}{\longrightarrow} x^{\rm r} \tag{5.11}$$

for some $x^{\rm l}, x^{\rm r} \in \mathbb{R}$. We write

$$\left.\begin{array}{ll} L^n(t) := L^n_{(x^{\rm l},0)}(t), & R^n(t) := R^n_{(x^{\rm r},0)}(t), \\ L^{(n)}(\varepsilon_n^2 t) := \varepsilon_n L^n(t), & R^{(n)}(\varepsilon_n^2 t) := \varepsilon_n R^n(t). \end{array}\right\} \quad (t \geq 0). \tag{5.12}$$

Our aim is to determine the limit as $n \to \infty$ of the joint law of the diffusively rescaled left and right paths $L^{(n)}$ and $R^{(n)}$.

To this aim, we define processes $(S^n_k)_{k\geq 0}$ and $(T^n_k)_{k\geq 0}$ by

$$S^n_k := \sum_{i=0}^{k-1} 1_{\{L^n_i = R^n_i\}} \quad \text{and} \quad T^n_k := \sum_{i=0}^{k-1} 1_{\{L^n_i < R^n_i\}} \qquad (k \geq 0),$$

and we inductively define processes $(V^n_k, \tilde V^n_k, \tilde W^n_k, W^n_k)_{k\geq 0}$ with initial states $V^n_0 = \tilde V^n_0 = \tilde W^n_0 = W^n_0 = 0$ by

$$\left.\begin{array}{l} V^n(T^n_{k+1}) := V^n(T^n_k) + 1_{\{L^n_k < R^n_k\}}\big(L^n_{k+1} - L^n_k\big), \\ \tilde V^n(S^n_{k+1}) := \tilde V^n(S^n_k) + 1_{\{L^n_k = R^n_k\}}\big(L^n_{k+1} - L^n_k\big), \\ \tilde W^n(S^n_{k+1}) := \tilde W^n(S^n_k) + 1_{\{L^n_k = R^n_k\}}\big(R^n_{k+1} - R^n_k\big), \\ W^n(T^n_{k+1}) := W^n(T^n_k) + 1_{\{L^n_k < R^n_k\}}\big(R^n_{k+1} - R^n_k\big), \end{array}\right\} \quad (k \geq 0).$$

These definitions are illustrated in Figure 5.6. Note that our definitions are consistent in the sense that in the first line we get $V^n(T^n_{k+1}) := V^n(T^n_k)$ if $T^n_{k+1} = T^n_k$, and likewise in the other three lines. We observe that

$$(V^n_k)_{k\geq 0}, \quad (\tilde V^n_k, \tilde W^n_k)_{k\geq 0}, \quad \text{and} \quad (W^n_k)_{k\geq 0} \text{ are independent.}$$

Recall from (5.1) that l_n, r_n, b_n are the probabilities that at a point $z \in \mathbb{Z}^2_{\rm even}$ there starts only left arrow, only a right arrow, or both types of arrows,

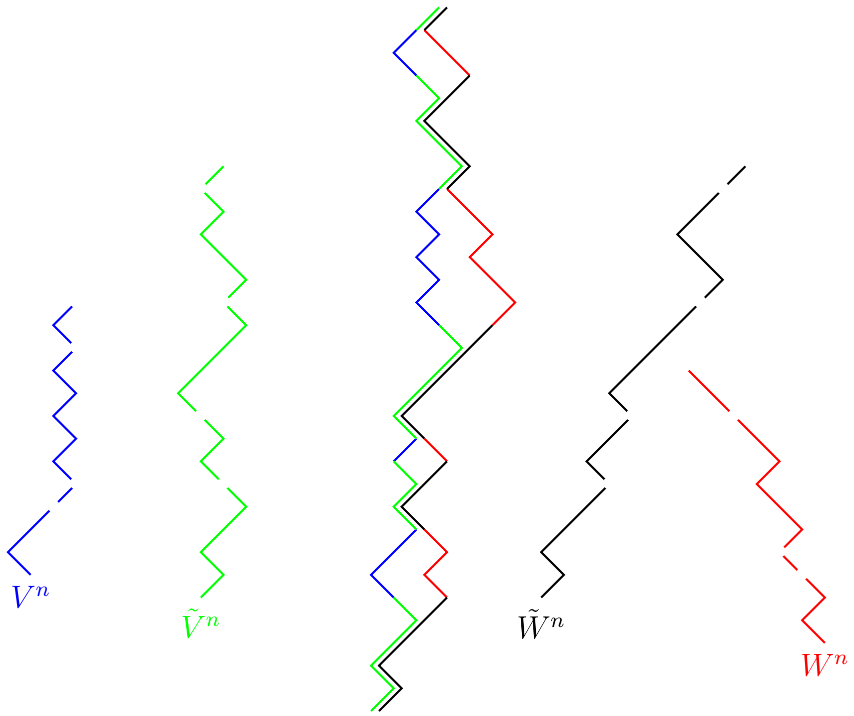

Figure 5.6: Decomposition of a left path and a right path (middle picture) into four random walks $V^{(n)}, \tilde{V}^{(n)}, \tilde{W}^{(n)}$ and $W^{(n)}$. Here $V^{(n)}$ and $\tilde{V}^{(n)}$ have a drift to the left while $\tilde{W}^{(n)}$ and $W^{(n)}$ have a drift to the right. The random walks $\tilde{V}^{(n)}$ and $\tilde{W}^{(n)}$ are highly correlated, but $V^{(n)}$ and $W^{(n)}$ are independent of each other and of the pair $(\tilde{V}^{(n)}, \tilde{W}^{(n)})$.

respectively. We observe that $(V^n_k)_{k\geq 0}$ and $(\tilde{V}^n_k)_{k\geq 0}$ are random walks with transition kernel

$$P^{\mathrm{l}}_n(x, x-1) = l_n + b_n \quad \text{and} \quad P^{\mathrm{l}}_n(x, x+1) = r_n,$$

while $(\tilde{W}^n_k)_{k\geq 0}$ and $(W^n_k)_{k\geq 0}$ are random walks with transition kernel

$$P^{\mathrm{r}}_n(x, x-1) = l_n \quad \text{and} \quad P^{\mathrm{l}}_n(x, x+1) = r_n + b_n.$$

We interpolate the processes $V^n, \tilde{V}^n, \tilde{W}^n, W^n$ and L^n, R^n, S^n, T^n linearly between integer times. Then it is easy to see that

(i) $L^n(t) = L^n(0) + V^n(T^n_t) + \tilde{V}^n(S^n_t) \qquad (t \geq 0),$

(ii) $R^n(t) = R^n(0) + W^n(T^n_t) + \tilde{W}^n(S^n_t) \qquad (t \geq 0),$

(iii) $\displaystyle\int_0^\infty 1_{\{L^n(t)+1<R^n(t)\}} \mathrm{d}S^n_t = 0,$

(iv) $S^n_t + T^n_t = t \qquad (t \geq 0).$

We define rescaled processes V^n, W^n and S^n, T^n by

$$V^{(n)}(\varepsilon_n^2 t) := \varepsilon_n V^n(t), \quad W^{(n)}(\varepsilon_n^2 t) := \varepsilon_n W^n(t),$$
$$S^{(n)}_{\varepsilon_n^2 t} := \varepsilon_n^2 S^n_t, \quad T^{(n)}_{\varepsilon_n^2 t} := \varepsilon_n^2 T^n_t.$$

$(t \geq 0)$, and we similarly define $\tilde V^{(n)}, \tilde W^{(n)}$ in terms of $\tilde V^n, \tilde W^n$. The rescaled processes satisfy conditions similar to (i)–(iv) above, except that in (iii) the indicator of the set $\{t \geq 0 : L^n(t) + 1 < R^n(t)\}$ should of course be replaced by the indicator of $\{t \geq 0 : L^{(n)}(t) + \varepsilon_n < R^{(n)}(t)\}$. The following proposition follows easily from (5.4) and Donsker's invariance principle, so we omit the proof.

Proposition 5.11 (Convergence of the driving noise) *Let $B^{\mathrm{l}}, B^{\mathrm{s}}$, and B^{r} be three independent Brownian motions. Then one has*

$$\mathbb{P}\big[\big(V^{(n)}(t), \tilde V^{(n)}(t), \tilde W^{(n)}(t), W^{(n)}(t)\big)_{t\geq 0} \cdot \big]$$
$$\underset{n\to\infty}{\Longrightarrow} \mathbb{P}\big[\big(B^{\mathrm{l}}_t - t, B^{\mathrm{s}}_t - t, B^{\mathrm{s}}_t + t, B^{\mathrm{r}}_t + t\big)_{t\geq 0}\big],$$

where $\Rightarrow$ denotes weak convergence of probability measures on $\mathcal{C}_{[0,\infty)}(\mathbb{R}^4)$, equipped with the topology of locally uniform convergence.

Our calculations and observations so far motivate the following definition. Let $l(0), r(0) \in \mathbb{R}$ with $l(0) \leq r(0)$ be given, together with $v, \tilde v, \tilde w, w \in \mathcal{C}_0$ which satisfy $\tilde w - \tilde v \in \mathcal{C}_0^+$. By definition, a solution to the *left-right equation*

$$\left.\begin{aligned} \mathrm{d}l(t) &= \mathrm{d}v(T_t) + \mathrm{d}\tilde v(S_t), \\ \mathrm{d}r(t) &= \mathrm{d}w(T_t) + \mathrm{d}\tilde w(S_t), \end{aligned}\right\} \quad (t \geq 0) \tag{5.13}$$

is a quadruple (l, r, S, T) where $l, r \in \mathcal{C}$ and $S, T \in \mathcal{C}_0^1$ satisfy $l \leq r$ and

(i) $l(t) = l(0) + v(T_t) + \tilde v(S_t) \qquad (t \geq 0),$

(ii) $r(t) = r(0) + w(T_t) + \tilde w(S_t) \qquad (t \geq 0),$

(iii) $\displaystyle\int_0^\infty 1_{\{l(t)<r(t)\}} \mathrm{d}h(S_t) = 0$ with $h := \tilde w - \tilde v,$

(iv) $S_t + T_t = t \qquad (t \geq 0).$

In analogy with our earlier terminology for sticky reflection equations, if we drop the condition that $l \leq r$, then we say that (l, r, S, T) is a *signed* solution to the left-right equation (5.13).

Proposition 5.12 (Left-right equation) *Assume that $l(0), r(0) \in \mathbb{R}$ and $v, \tilde{v}, \tilde{w}, w \in \mathcal{C}_0$ satisfy $l(0) \leq r(0)$ and $\tilde{w} - \tilde{v} \in \mathcal{C}_0^+$. Then the left-right equation (5.13) has a solution (l, r, S, T). If $\tilde{w} - \tilde{v}$ is strictly increasing, then this solution is unique.*

Proof Let

$$f := w - v \quad \text{and} \quad h := \tilde{w} - \tilde{v}.$$

We observe that if (g, S, T) solves the sticky reflection equation

$$\mathrm{d}g(t) = \mathrm{d}f(T_t) + \mathrm{d}h(S_t) \qquad (t \geq 0) \tag{5.14}$$

with initial state $g(0) = r(0) - l(0)$, then setting

$$\left.\begin{aligned} l(t) &:= l(0) + v(T_t) + \tilde{v}(S_t), \\ r(t) &:= r(0) + w(T_t) + \tilde{w}(S_t) \end{aligned}\right\} \quad (t \geq 0)$$

yields a solution to the left-right equation (5.13). In view of this, existence of solutions to the sticky reflection equation follows from Proposition 5.4.

We observe that if (l, r, S, T) solves the left-right equation (5.13), then setting $g := r - l$ yields a solution to the sticky reflection equation (5.14). Proposition 5.4 tells us that solutions to the latter are unique if $\tilde{w} - \tilde{v}$ is strictly increasing. In particular, in this case, S and T are uniquely determined, and hence, by conditions (i) and (ii) of the definition of a solution to left-right equation, so are l and r. ∎

In view of Proposition 5.11, we are interested in solutions (L, R, S, T) to the left-right equation (5.13)

$$\left.\begin{aligned} \mathrm{d}L(t) &= \mathrm{d}V(T_t) + \mathrm{d}\tilde{V}(S_t), \\ \mathrm{d}R(t) &= \mathrm{d}W(T_t) + \mathrm{d}\tilde{W}(S_t), \end{aligned}\right\} \quad (t \geq 0), \tag{5.15}$$

where

$$\left.\begin{aligned} V(t) &:= B^{\mathrm{l}}_t - t, & \tilde{V}(t) &:= B^{\mathrm{s}}_t - t, \\ W(t) &:= B^{\mathrm{r}}_t + t, & \tilde{W}(t) &:= B^{\mathrm{s}}_t + t, \end{aligned}\right\} \quad (t \geq 0),$$

and $B^{\mathrm{l}}, B^{\mathrm{s}}, B^{\mathrm{r}}$ are three independent Brownian motions. Using the fact that $S_t + T_t = t$, we can write

$$\mathrm{d}(B^{\mathrm{l}}_{T_t} - T_t) + \mathrm{d}(B^{\mathrm{s}}_{S_t} - S_t) = \mathrm{d}B^{\mathrm{l}}_{T_t} - \mathrm{d}T_t + \mathrm{d}B^{\mathrm{s}}_{S_t} - \mathrm{d}S_t, = \mathrm{d}B^{\mathrm{l}}_{T_t} + \mathrm{d}B^{\mathrm{s}}_{S_t} - \mathrm{d}t,$$

which motivates us to rewrite (5.15) in the simpler form (5.16) below. Below is the main result of this section.

Theorem 5.13 (Scaling limit of a left and right path) *Let ω^n be a sequence of arrow configurations satisfying (5.4). Assume (5.11) and let $(L^{(n)}(t))_{t\geq 0}$ and $(R^{(n)}(t))_{t\geq 0}$ be defined as in (5.12). Let $B^{\rm l}, B^{\rm s}, B^{\rm r}$ be three independent Brownian motions and let (L, R, S, T) be the a.s. unique solution to the left-right equation*

$$\left.\begin{aligned} \mathrm{d}L(t) &= \mathrm{d}B^{\rm l}_{T_t} + \mathrm{d}B^{\rm s}_{S_t} - \mathrm{d}t, \\ \mathrm{d}R(t) &= \mathrm{d}B^{\rm r}_{T_t} + \mathrm{d}B^{\rm s}_{S_t} - \mathrm{d}t, \end{aligned}\right\} \quad (t \geq 0). \tag{5.16}$$

Then one has

$$\mathbb{P}\big[\big(L^{(n)}(t), R^{(n)}(t)\big)_{t\geq 0} \cdot\ \big] \underset{n\to\infty}{\Longrightarrow} \mathbb{P}\big[\big(L(t), R(t)\big)_{t\geq 0}\big],$$

where $\Rightarrow$ denotes weak convergence of probability measures on $\mathcal{C}_{[0,\infty)}(\mathbb{R}^2)$, equipped with the topology of locally uniform convergence.

Proof (sketch) The proof is almost identical to the proof of Theorem 5.10, so we only sketch the main line of the argument. Lemmas 5.7 and 5.8, that were formulated for signed solutions to sticky reflection equations, generalise in a straightforward way to signed solutions to left-right equations, and hence so does Proposition 5.6. If we slightly change our definition of $L^n(t)$ in (5.12) by putting $L^n(t) := L^n_{(x^{\rm l},0)}(t) + 1$ instead of $:= L^n_{(x^{\rm l},0)}(t)$ $(t \geq 0)$, then for each n we have that $(L^{(n)}, R^{(n)}, S^{(n)}, S^{(n)})$ is a signed solution to the left-right equation

$$\left.\begin{aligned} \mathrm{d}L^{(n)}(t) &= \mathrm{d}V^{(n)}(T^{(n)}_t) + \mathrm{d}\tilde{V}^{(n)}(S^{(n)}_t), \\ \mathrm{d}R^{(n)}(t) &= \mathrm{d}W^{(n)}(T^{(n)}_t) + \mathrm{d}\tilde{W}^{(n)}(S^{(n)}_t), \end{aligned}\right\} \quad (t \geq 0).$$

By Proposition 5.11 and Skorohod's representation theorem, we can couple our random variables such that almost surely

$$\begin{aligned} \big(V^{(n)}(t)\big)_{t\geq 0} &\underset{n\to\infty}{\longrightarrow} \big(B^{\rm l}_t - t\big)_{t\geq 0}, \qquad & \big(\tilde{V}^{(n)}(t)\big)_{t\geq 0} &\underset{n\to\infty}{\longrightarrow} \big(B^{\rm s}_t - t\big)_{t\geq 0}, \\ \big(W^{(n)}(t)\big)_{t\geq 0} &\underset{n\to\infty}{\longrightarrow} \big(B^{\rm r}_t + t\big)_{t\geq 0}, \qquad & \big(\tilde{W}^{(n)}(t)\big)_{t\geq 0} &\underset{n\to\infty}{\longrightarrow} \big(B^{\rm s}_t + t\big)_{t\geq 0}, \end{aligned}$$

where $\to$ denotes locally uniform convergence. Using the analogon of Proposition 5.6 for left-right equations, it follows that for this coupling almost surely

$$\big(L^{(n)}(t)\big)_{t\geq 0} \underset{n\to\infty}{\longrightarrow} \big(L(t)\big)_{t\geq 0} \quad \text{and} \quad \big(R^{(n)}(t)\big)_{t\geq 0} \underset{n\to\infty}{\longrightarrow} \big(R(t)\big)_{t\geq 0}.$$

Of course, for this last statement it does not matter whether we have defined $L^t(t) := L^n_{(x^l,0)}(t) + 1$ or $:= L^n_{(x^l,0)}(t)$ $(t \geq 0)$. Since almost sure convergence implies weak convergence in law, this completes the proof of the theorem. ∎

Remark 1 Our proof of Theorem 5.13 is an adaptation of the proof of [SS08, Prop. 5.2]. Our description of the joint law of a left and right path is based on the left-right equation (5.13). Alternatively, it is shown in [SS08] that subject to the condition that $L(t) \leq R(t)$ for all $t \geq 0$, the following stochastic differential equation (SDE) has a unique weak solution, that also describes joint law of a left and right path:

$$\begin{aligned}
\mathrm{d}L(t) &= 1_{\{L(t)<R(t)\}}\mathrm{d}\tilde{B}^{\mathrm{l}}(t) + 1_{\{L(t)=R(t)\}}\mathrm{d}\tilde{B}^{\mathrm{s}}(t) - \mathrm{d}t,\\
\mathrm{d}R(t) &= 1_{\{L(t)<R(t)\}}\mathrm{d}\tilde{B}^{\mathrm{r}}(t) + 1_{\{L(t)=R(t)\}}\mathrm{d}\tilde{B}^{\mathrm{s}}(t) + \mathrm{d}t.
\end{aligned}$$

Here $\tilde{B}^{\mathrm{l}}, \tilde{B}^{\mathrm{r}}, \tilde{B}^{\mathrm{s}}$ are three independent Brownian motions. Yet another useful characterisation of the joint law of a left and right path, which is formulated in terms of the drift, quadratic variation, and cross-variation of L and R, can be found in [SS19, Prop. 3.2].

Remark 2 Our understanding of sticky reflection is not as good as for Skorohod reflection. Recall that in Section 4.3, we described the conditional law of a forward path in the Brownian web given a path in the dual web in terms of Skorohod reflection. Similarly, for two forward paths in the Brownian web, which are just coalescing Brownian motions, we have a (very easy) description of the conditional law of one path given the other one, which allowed us to give an easy inductive description of any finite number of coalescing Brownian motions. By contrast, even though we have a good description of their *joint* law, it does not seem easy to give a description of the *conditional* law of a right forward path given a left forward path.

Exercise 5.14 (Positive Lebesgue time) *Let (L, R, S, T) be a solution of the left-right equation (5.15) started in an initial state such that $L(0) = R(0)$. Prove that*

$$\mathbb{P}\Big[\int_0^\infty 1_{\{L(t)=R(t)\}}\mathrm{d}t > 0\Big] > 0.$$

Exercise 5.15 (Right-left pair) *Let ω^n be a sequence of arrow configurations satisfying (5.4), and let $(\omega^{\mathrm{l},n}_z)_{z\in\mathbb{Z}^2_{\mathrm{even}}}$ and $(\omega^{\mathrm{r},n}_z)_{z\in\mathbb{Z}^2_{\mathrm{even}}}$ be the collections of $\{-1,+1\}$-valued random variables defined in Section 5.2. For each n, define paths $(x^n_k, y^n_k)_{k\geq 0}$ starting in $x^n_0 = 0 = y^n_0$ by the inductive formulas*

$$x^n_{k+1} := \begin{cases} x^n_k + \omega^{\mathrm{l},n}_{(x^n_k,k)} & \text{if } x^n_k = y^n_k,\\ x^n_k + \omega^{\mathrm{r},n}_{(x^n_k,k)} & \text{if } x^n_k < y^n_k,\end{cases}$$

and

$$y^n_{k+1} := \begin{cases} y^n_k + \omega^{\mathrm{r},n}_{(y^n_k,k)} & \text{if } x^n_k = y^n_k, \\ y^n_k + \omega^{\mathrm{l},n}_{(y^n_k,k)} & \text{if } x^n_k < y^n_k. \end{cases}$$

Note that this says that when there is a choice, x^n takes the left arrow if x^n and y^n are at the same position, but the right arrow otherwise, and the other way round for the path y^n. As a consequence, $(x^n_k)_{k\geq 0}$ on its own is not a Markov chain and neither is $(y^n_k)_{k\geq 0}$. Nevertheless, the joint process $(x^n_k, y^n_k)_{k\geq 0}$ is a Markov chain. Describe the diffusive scaling limit of this Markov chain by means of a left-right equation.

5.6 The left-right Brownian web

Let ε_n be positive constants, tending to zero, let ω^n be a sequence of arrow configurations satisfying (5.4), and let $\mathcal{U}^{\mathrm{l}}_n$ and $\mathcal{U}^{\mathrm{r}}_n$ be the collections of left and right paths in ω^n, respectively. In Theorem 5.3, we have shown that the diffusivelt rescaled collection of paths $\theta_{\varepsilon_n}(\overline{\mathcal{U}}^{\mathrm{l}}_n)$ converges in law to a Brownian web $\mathcal{W}^{\mathrm{l}}$ with drift -1, and likewise $\theta_{\varepsilon_n}(\overline{\mathcal{U}}^{\mathrm{r}}_n)$ converges in law to a Brownian web $\mathcal{W}^{\mathrm{r}}$ with drift $+1$. In this section, we will show that also the joint law of $\theta_{\varepsilon_n}(\overline{\mathcal{U}}^{\mathrm{l}}_n, \overline{\mathcal{U}}^{\mathrm{r}}_n)$ converges, and characterise the joint law of the limit object $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$. We will call this limit object a *left-right Brownian web*, and we will call $\mathcal{W}^{\mathrm{l}}$ and $\mathcal{W}^{\mathrm{r}}$ the associated *left Brownian web* and *right Brownian web*, respectively.

Recall that $\Pi = \Pi(\overline{\mathbb{R}})$ denotes the space of all paths, which may have finite starting and final times, and $\Pi^\uparrow$ is the subspace of upward paths, which have infinite final times. Let $\mathcal{A} \subset \Pi^\uparrow$ be a collection of upward paths, let $A \subset \mathcal{R}(\overline{\mathbb{R}})$ be a closed set, and let $\mathrm{int}(A)$ denote its interior. Then we define the *restriction* of $\mathcal{A}$ to A as

$$\mathcal{A}\big|_A := \overline{\big\{\pi \in \Pi : \pi \subset \mathrm{int}(A),\ \exists \pi' \in \mathcal{A} \text{ s.t. } \pi \subset \pi'\big\}},$$

where the overbar means that we take the closure in the topology on Π. It is easy to see that if $\mathcal{A}$ is compact, then so is $\mathcal{A}\big|_A$. In particular, if $\mathcal{W}$ is a Brownian web, then $\mathcal{W}\big|_A$ is a random variable taking values in the Polish space $\mathcal{K}(\Pi)$. Below is a conjecture that is so far unproven.

Conjecture 5.16 (Left-right Brownian web) *There exists a random variable $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ with values in $\mathcal{K}(\Pi^\uparrow) \times \mathcal{K}(\Pi^\uparrow)$, whose law is uniquely characterised by the following properties.*

(i) *$\mathcal{W}^{\mathrm{l}}$ is distributed as a Brownian web with drift -1 and $\mathcal{W}^{\mathrm{r}}$ is distributed as a Brownian web with drift $+1$.*

(ii) *For each $z \in \mathbb{R}^2$, the joint law of a.s. unique paths $\pi^{\rm l}_z \in \mathcal{W}^{\rm l}(z)$ and $\pi^{\rm r}_z \in \mathcal{W}^{\rm r}(z)$ is described by the left-right equation (5.16).*

(iii) *If $A_1, \ldots, A_n \subset \mathbb{R}^2$ are disjoint closed sets, then the random variables $(\mathcal{W}^{\rm l}\big|_{A_1}, \mathcal{W}^{\rm r}\big|_{A_1}), \ldots, (\mathcal{W}^{\rm l}\big|_{A_n}, \mathcal{W}^{\rm r}\big|_{A_n})$ are independent.*

The difficult part of Conjecture 5.16 is the claim that properties (i)–(iii) uniquely characterise the joint law of $(\mathcal{W}^{\rm l}, \mathcal{W}^{\rm r})$. If Conjecture 5.16 were proved, then by combining Theorems 5.3 and 5.13, it would be possible to give a very short proof of the following result.

Theorem 5.17 (Convergence to the left-right Brownian web) *Let ω^n be a sequence of arrow configurations satisfying (5.4), and let $\mathcal{U}^{\rm l}_n$ and $\mathcal{U}^{\rm r}_n$ be the collections of left and right paths in ω^n, respectively. There exists a random variable $(\mathcal{W}^{\rm l}, \mathcal{W}^{\rm r})$ with values in $\mathcal{K}(\Pi^\uparrow) \times \mathcal{K}(\Pi^\uparrow)$, called a left-right Brownian web, such that*

$$\mathbb{P}\big[\theta_{\varepsilon_n}(\overline{\mathcal{U}}^{\rm l}_n, \overline{\mathcal{U}}^{\rm r}_n) \in \cdot\,\big] \underset{n\to\infty}{\Longrightarrow} \mathbb{P}\big[(\mathcal{W}^{\rm l}, \mathcal{W}^{\rm r}) \in \cdot\,\big],$$

where $\Rightarrow$ denotes weak convergence of probability laws on the space $\mathcal{K}(\Pi^\uparrow) \times \mathcal{K}(\Pi^\uparrow)$.

Theorem 5.17 has been proved in [SS08]. In the remainder of this section, we will sketch its proof. In the absence of Conjecture 5.16, we will need another characterisation of the limit object, the left-right Brownian web $(\mathcal{W}^{\rm l}, \mathcal{W}^{\rm r})$. The characterisation will be a bit more complicated than Conjecture 5.16, but nevertheless very similar in spirit. Since $\mathcal{W}^{\rm l}$ and $\mathcal{W}^{\rm r}$ are Brownian webs, for each deterministic $z \in \mathbb{R}^2$, there almost surely exist unique paths $\pi^{\rm l}_z$ and $\pi^{\rm r}_z$ such that $\pi^{\rm l}_z \in \mathcal{W}^{\rm l}(z)$ and $\pi^{\rm r}_z \in \mathcal{W}^{\rm r}(z)$. In view of the characterisation of the Brownian web (Theorem 3.7), in order to characterise the joint law of $\mathcal{W}^{\rm l}$ and $\mathcal{W}^{\rm r}$, it suffices to describe the joint law of

$$\pi^{{\rm q}_1}_{z_1}, \ldots, \pi^{{\rm q}_m}_{z_m} \tag{5.17}$$

for any deterministic finite collection of points $z_1, \ldots, z_m \in \mathbb{R}^2$ and sequence ${\rm q}_1, \ldots, {\rm q}_m$ with ${\rm q}_i \in \{{\rm l}, {\rm r}\}$ for all $1 \leq i \leq m$. Without loss of generality, we can assume that the time coordinates of $z_i = (x_i, t_i)$ $(1 \leq i \leq m)$ are ordered as $t_1 \leq \cdots \leq t_m$. Then it suffices to describe during each of the time intervals $[t_1, t_2], \ldots, [t_{n-1}, t_m]$ and $[t_m, \infty)$ the joint evolution (which is Markovian) of the paths whose starting time lies before the initial time of the interval. In view of this, we can without loss of generality assume that $t_1 = \ldots = t_m = 0$.

We assume now that $t_1 = \ldots = t_m = 0$. Without loss of generality, we also assume that $x_1 \leq \cdots \leq x_m$. We can also assume that $x_i < x_{i+1}$ whenever

$q_i = q_{i+1}$, since paths of the same type (left or right) coalesce as soon as they meet. We can then group left and right paths that immediately follow after each other (in this order), leaving the remaining paths as singletons. For example, if at time zero we have a collection of left and right paths that ordered from left to right, looks like this $LRLLRLRRRLR$, then we group them as follows:

$$\{LR\}\{L\}\{LR\}\{LR\}\{R\}\{R\}\{LR\}.$$

Let τ_1 denote the first time when a path from one group meets a path from another group. It may be that multiple paths meet at such a time. However, it is not hard to see that with probability one, at the time τ_1, there exist two consecutive groups so that all paths that meet at time τ_1 belong to these two groups.

If at the the time τ_1, two left paths meet, then they coalesce, so the total number of left paths that we have to follow decreases by one. Likewise, if two right paths meet, they also coalesce. We are now in a similar situation as at time zero and can again group the remaining paths into singletons and pairs consisting of one left path and one right path (in this order). We then let τ_2 denote the first time when a path of these newly created groups meets a path of another group. Inductively, we define $\tau_3, \tau_4, \ldots$ in the same fashion. Then during each of the random time intervals $[0, \tau_1], [\tau_1, \tau_2], \ldots$, we can specify the joint law of our paths by saying that the groups evolve independently in such a way that:

- each group consisting of a single left path evolves as a Brownian motion with drift -1,
- each group consisting of a single right path evolves as a Brownian motion with drift $+1$,
- each group consisting of a left and a right path evolves as a solution to the left-right equation (5.16).

Note that at each of the times $\tau_1, \tau_2, \ldots$, either two left paths coalesce, or two right paths coalesce, or a right and left path change their order, in the sense that the right path was on the left of the left path before, but has to stay on the right after. This means that there are only finitely many times $\tau_1, \ldots, \tau_N$, and we can specify the evolution of the left and right paths on each of the intervals $[0, \tau_1], \ldots, [\tau_{N-1}, \tau_N]$ and $[\tau_N, \infty)$ according to the rules above.

Filling in the technical details is a bit cumbersome, especially since we are working with stopping times, but the description above gives the main idea. Using this idea, one can give a rigorous definition of a collection of *left-right coalescing Brownian motions*. We cite the following result from [SS08, Prop. 5.2]. (This reference is, admittedly, also a bit sketchy on the technical details.)

Proposition 5.18 (Convergence of finite dimensional distributions) *Let ε_n be positive constants, tending to zero and let ω^n be a sequence of arrow configurations satisfying (5.4). Fix $z_1, \ldots, z_m \in \mathbb{R}^2$ and $z_i^n \in \mathbb{Z}^2_{\mathrm{even}}$ such that $\theta_{\varepsilon_n}(z_i^n) \to z_i$ as $n \to \infty$ $(1 \leq i \leq m)$. Fix $\mathrm{q}_1, \ldots, \mathrm{q}_m$ with $\mathrm{q}_i \in \{\mathrm{l}, \mathrm{r}\}$ and depending on whether $\mathrm{q}_i = \mathrm{l}$ or $= \mathrm{r}$, let $\pi_i^{\mathrm{q}_i, n}$ denote the unique left or right path in ω^n starting at z_i^n. Then*

$$\mathbb{P}\big[\theta_{\varepsilon_k}(\pi_1^{\mathrm{q}_1,n}, \ldots, \pi_m^{\mathrm{q}_m,n}) \in \cdot\,\big] \underset{k\to\infty}{\Longrightarrow} \mathbb{P}\big[(\pi_{z_1}^{\mathrm{q}_1}, \ldots, \pi_{z_m}^{\mathrm{q}_m}) \in \cdot\,\big],$$

where $\Rightarrow$ denotes weak convergence of probability measures on $(\Pi^{\uparrow})^m$, and $(\pi_{z_1}^{\mathrm{q}_1}, \ldots, \pi_{z_m}^{\mathrm{q}_m})$ is a collection of left-right coalescing Brownian motions starting from $z_1, \ldots, z_m$.

It is clear that left and right random paths in an arrow configuration are consistent in the sense of Kolmogorov's extension theorem, and hence by Proposition 5.18 the same must be true for left-right coalescing Brownian motions. In view of this, if $\mathcal{D} \subset \mathbb{R}^2$ is a deterministic countable dense set, then we can construct a collection $(\pi_z^{\mathrm{l}}, \pi_z^{\mathrm{r}})_{z\in\mathcal{D}}$ of left-right coalescing Brownian motions started from $\mathcal{D}$. By Theorem 3.7, setting

$$\mathcal{W}^{\mathrm{l}} := \overline{\{\pi_z^{\mathrm{l}} : z \in \mathcal{D}\}} \quad \text{and} \quad \mathcal{W}^{\mathrm{r}} := \overline{\{\pi_z^{\mathrm{l}} : z \in \mathcal{D}\}}$$

then defines two Brownian webs $\mathcal{W}^{\mathrm{l}}$ and $\mathcal{W}^{\mathrm{r}}$ with drift -1 and $+1$, respectively. By definition, we call $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ the *left-right Brownian web*. We let $\hat{\mathcal{W}}^{\mathrm{l}}, \hat{\mathcal{W}}^{\mathrm{r}}$ denote the dual Brownian webs associated with $\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}}$. Theorem 5.17 is implied by the following theorem, that gives a somewhat more complete picture.

Theorem 5.19 (Convergence to the left-right Brownian web) *Let ε_n be positive constants tending to zero and let ω^n be a sequence of arrow configurations satisfying (5.4). Let $\mathcal{U}_n^{\mathrm{l}}$ and $\mathcal{U}_n^{\mathrm{r}}$ be the collections of left and right paths in ω^n, respectively, and let $\mathcal{U}_n^{\mathrm{l}*}$ and $\mathcal{U}_n^{\mathrm{r}*}$ be the collections of dual left and right paths. Then one has*

$$\mathbb{P}\big[\theta_{\varepsilon_n}(\overline{\mathcal{U}}_n^{\mathrm{l}}, \overline{\mathcal{U}}_n^{\mathrm{r}}, \overline{\mathcal{U}}_n^{\mathrm{l}*}, \overline{\mathcal{U}}_n^{\mathrm{r}*}) \in \cdot\,\big] \underset{n\to\infty}{\Longrightarrow} \mathbb{P}\big[(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}}, \hat{\mathcal{W}}^{\mathrm{l}}, \hat{\mathcal{W}}^{\mathrm{r}}) \in \cdot\,\big],$$

where $\Rightarrow$ denotes weak convergence of probability laws on the space $\mathcal{K}(\Pi^{\uparrow})^2 \times \mathcal{K}(\Pi^{\downarrow})^2$, $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ is a left-right Brownian web, and $\hat{\mathcal{W}}^{\mathrm{l}}$ and $\hat{\mathcal{W}}^{\mathrm{r}}$ are the associated dual webs.

Proof Convergence of $(\overline{\mathcal{U}}^{\rm l}_n,\overline{\mathcal{U}}^{\rm l*}_n)$ to $(\mathcal{W}^{\rm l},\hat{\mathcal{W}}^{\rm l})$ and of $(\overline{\mathcal{U}}^{\rm r}_n,\overline{\mathcal{U}}^{\rm r*}_n)$ to $(\mathcal{W}^{\rm r},\hat{\mathcal{W}}^{\rm r})$ follows from Theorem 3.18. Using Lemma 3.14, it follows that the laws of the random variables $\theta_{\varepsilon_n}(\overline{\mathcal{U}}^{\rm l}_n,\overline{\mathcal{U}}^{\rm r}_n,\overline{\mathcal{U}}^{\rm l*}_n,\overline{\mathcal{U}}^{\rm r*}_n)$ are tight, so by going to a subsequence, we may assume that they converge in law to some random variable $(\mathcal{V}^{\rm l},\mathcal{V}^{\rm r},\hat{\mathcal{V}}^{\rm l},\hat{\mathcal{V}}^{\rm r})$. In view of Lemma 2.2, it suffices to show that $(\mathcal{V}^{\rm l},\mathcal{V}^{\rm r},\hat{\mathcal{V}}^{\rm l},\hat{\mathcal{V}}^{\rm r})$ is equal in law to $(\mathcal{W}^{\rm l},\mathcal{W}^{\rm r},\hat{\mathcal{W}}^{\rm l},\hat{\mathcal{W}}^{\rm r})$. Since a web is a.s. uniquely determined by its dual, it suffices to show that $(\mathcal{V}^{\rm l},\mathcal{V}^{\rm r})$ is equal in law to $(\mathcal{W}^{\rm l},\mathcal{W}^{\rm r})$. By Theorem 3.18, $\mathcal{V}^{\rm l}$ is a Brownian web with drift -1 and $\mathcal{V}^{\rm r}$ is a Brownian web with drift $+1$. Therefore, by Theorem 3.7, we know that at each deterministic $z\in\mathbb{R}^2$, the sets $\mathcal{V}^{\rm l}(z)$ and $\mathcal{V}^{\rm r}(z)$ almost surely contain a single path. To show that $(\mathcal{V}^{\rm l},\mathcal{V}^{\rm r})$ is equal in law to $(\mathcal{W}^{\rm l},\mathcal{W}^{\rm r})$, by Theorem 3.7, it suffices to show that $(\mathcal{V}^{\rm l},\mathcal{V}^{\rm r})$ has the right finite dimensional distributions, i.e., we must show that the left and right paths started from finitely many points are distributed as left-right coalescing Brownian motions. This follows from Proposition 5.18, so the proof is complete. ∎

Recall that $-\pi:=\{(-x,-t):(x,t)\in\pi\}$ is our notation for a path π, rotated over 180 degrees, and that $-\mathcal{W}:=\{-\pi:\pi\in\mathcal{W}\}$. It is not hard to see that $-\hat{\mathcal{W}}^{\rm l}$ is equally distributed with $\mathcal{W}^{\rm l}$ (both are Brownian webs with drift -1) and $-\hat{\mathcal{W}}^{\rm r}$ is equally distributed with $\mathcal{W}^{\rm r}$. In fact, a stronger statement holds.

Lemma 5.20 (Dual left-right Brownian web) *Let $(\mathcal{W}^{\rm l},\mathcal{W}^{\rm r})$ be a left-right Brownian web and let $\hat{\mathcal{W}}^{\rm l}$ and $\hat{\mathcal{W}}^{\rm r}$ be the dual Brownian webs associated with $\mathcal{W}^{\rm l}$ and $\mathcal{W}^{\rm r}$. Then $(-\hat{\mathcal{W}}^{\rm l},-\hat{\mathcal{W}}^{\rm r})$ is equally distributed with $(\mathcal{W}^{\rm l},\mathcal{W}^{\rm r})$.*

Proof It is straightforward to check that $(-\mathcal{U}^{\rm l*}_n,-\mathcal{U}^{\rm r*}_n)$ is equally distributed with $(\mathcal{U}^{\rm l}_n,\mathcal{U}^{\rm r}_n)$, so the claim follows from finite approximation, using Theorem 5.19). ∎

5.7 The hopping and wedge constructions

We continue to assume that ε_n are positive constants tending to zero and that ω^n is a sequence of arrow configurations satisfying (5.4). Ultimately, we are not interested in left and right paths only, but in the scaling limit of the set $\mathcal{U}_n$ of *all* open paths in the arrow configuration ω^n. In this section, we describe a random compact set of paths $\mathcal{N}$ that we will call the *Brownian net* and that will turn out to be the scaling limit of the sets $\mathcal{U}_n$. We will obtain $\mathcal{N}$ as a function of a left-right Brownian web $(\mathcal{W}^{\rm l},\mathcal{W}^{\rm r})$. In fact, we will describe two different ways to construct $\mathcal{N}$ from $(\mathcal{W}^{\rm l},\mathcal{W}^{\rm r})$. The fact that both constructions yield the same object will be important when we prove the convergence in law of the collection $\theta_{\varepsilon_n}(\overline{\mathcal{U}}_n)$ of rescaled discrete open paths to the Brownian net $\mathcal{N}$.

The two constructions we will use are called the *hopping construction* and the *wedge construction* of the Brownian net. They will yield two sets of paths $\mathcal{N}_-$ and $\mathcal{N}_+$ that will later be shown to be equal, similar to the statement of Theorem 3.15. In fact, the wedge construction, which yields $\mathcal{N}_+$, is extremely similar to the definition of $\mathcal{W}_+$ in Theorem 3.15. We start with the hopping construction, however, which is a bit more complicated than the construction of $\mathcal{W}_-$ in Theorem 3.15, since it requires a new concept: hopping.

Let $(\mathcal{W}^{\rm l}, \mathcal{W}^{\rm r})$ be a left-right Brownian web and let $\pi^{\rm l}_1, \pi^{\rm r}_2, \pi^{\rm l}_3, \ldots$ be a finite sequence of paths that are alternatively taken from $\mathcal{W}^{\rm l}$ and $\mathcal{W}^{\rm r}$, such that

$$\sigma_{\pi^{\rm l}_1} < \sigma_{\pi^{\rm r}_2} < \sigma_{\pi^{\rm l}_3} < \cdots$$

and

$$\pi^{\rm r}_2(\sigma_{\pi^{\rm r}_2}) < \pi^{\rm l}_1(\sigma_{\pi^{\rm r}_2}), \quad \pi^{\rm r}_2(\sigma_{\pi^{\rm l}_3}) < \pi^{\rm l}_3(\sigma_{\pi^{\rm l}_3}), \ldots$$

i.e., the second path, which is a right path, is started on the left of the first path, which is a left path, and then the third path, which is a left path, is started on the right of the second path and so on; see Figure 5.7.

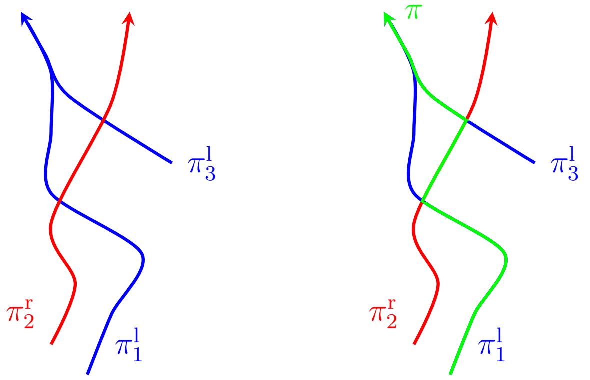

Figure 5.7: A path π constructed by hopping between left and right paths $\pi^{\rm l}_1, \pi^{\rm r}_2, \pi^{\rm l}_3$.

Recall that

$$\tau(\pi_1, \pi_2) := \inf\{t > \sigma_{\pi_1} \vee \sigma_{\pi_2} : \pi_1(t) = \pi_2(t)\} \qquad (\pi_1, \pi_2 \in \Pi^\uparrow)$$

denotes the first meeting time of two upward paths π_1, π_2. Let us assume that

$$\tau(\pi^{\rm l}_1, \pi^{\rm r}_2) < \sigma_{\pi^{\rm l}_3}, \quad \tau(\pi^{\rm r}_2, \pi^{\rm l}_3) < \sigma_{\pi^{\rm r}_4}, \ldots$$

i.e., we start the third path only after the first meeting time of the first two paths and so on. Then we can define a path π with starting time $\sigma_\pi := \sigma_{\pi^{\rm l}_1}$ by

$$\pi(t) := \begin{cases} \pi^{\rm l}_1(t) & \big(\sigma_{\pi^{\rm l}_1} \le t \le \tau(\pi^{\rm l}_1, \pi^{\rm r}_2)\big), \\ \pi^{\rm r}_2(t) & \big(\tau(\pi^{\rm l}_1, \pi^{\rm r}_2) \le t \le \tau(\pi^{\rm r}_2, \pi^{\rm l}_3)\big), \\ \pi^{\rm l}_3(t) & \big(\tau(\pi^{\rm r}_2, \pi^{\rm l}_3) \le t \le \tau(\pi^{\rm l}_3, \pi^{\rm r}_4)\big), \end{cases}$$

and so on, i.e., we start by following the path $\pi^{\rm l}_1$, then "hop" onto the path $\pi^{\rm r}_2$ at the first time when $\pi^{\rm l}_1$ meets $\pi^{\rm r}_2$, and so on, until we arrive at the last path in our finite sequence, which we follow till time $+\infty$. We fix a countable dense set $\mathcal{D} \subset \mathbb{R}^2$ and let

$$\begin{aligned}\mathcal{N}_- := \text{the closure of } \big\{\pi : \pi \text{ is obtained by hopping} \\ \text{between paths in } (\pi^{\rm l}_z)_{z\in\mathcal{D}} \text{ and } (\pi^{\rm r}_z)_{z\in\mathcal{D}}\big\}.\end{aligned}$$

This completes the description of the hopping construction of the Brownian net. We make one simple observation.

Lemma 5.21 (Compactness of the Brownian net) *Almost surely, $\mathcal{N}_-$ is a compact subset of $\Pi^\uparrow$.*

Proof Let us write $\mathcal{N}_- = \overline{\mathcal{N}'_-}$, where $\mathcal{N}'_-$ is the set of paths that can be constructed by hopping between paths in $(\pi^{\rm l}_z)_{z\in\mathcal{D}}$ and $(\pi^{\rm r}_z)_{z\in\mathcal{D}}$. We need to show that $\mathcal{N}'_-$ is almost surely precompact. We apply Proposition 2.32. We need to show that

$$\begin{aligned}\mathbb{P}\big[\,\big|\pi(u) - \pi(t)\big| \ge \varepsilon \text{ for some } \pi \in \mathcal{N}'_- \text{ and } \sigma_\pi \le t \le u \\ \text{s.t. } (\pi(t), t) \in [-T, T]^2,\ u - t \le \delta\big] \underset{\delta\to 0}{\longrightarrow} 0 \quad \forall T < \infty,\ \varepsilon > 0.\end{aligned} \tag{5.18}$$

Let $\mathcal{W}^{\rm l} := \overline{\{\pi^{\rm l}_z : z \in \mathcal{D}\}}$ and $\mathcal{W}^{\rm r} := \overline{\{\pi^{\rm r}_z : z \in \mathcal{D}\}}$ be the left and right Brownian webs constructed from our left and right paths. Paths in $\mathcal{N}'_-$ cannot cross paths in $\mathcal{W}^{\rm l}$ from right to left and they cannot cross paths in $\mathcal{W}^{\rm r}$ from left to right. Therefore, letting $\pi^{\rm l-}_z$ and $\pi^{\rm r+}_z$ as in Lemma 4.6 denote the minimal left and maximal right path starting at a point z, we have

$$\pi^{\rm l-}_{(\pi(t),t)}(u) \le \pi(u) \le \pi^{\rm r+}_{(\pi(t),t)}(u)$$

for all $\pi \in \mathcal{N}'_-$ and $\sigma_\pi \le t \le u$. In view of this, the fact that $\mathcal{N}'_-$ satisfies (5.18) follows from the fact that $\mathcal{W}^{\rm l}$ and $\mathcal{W}^2$ satisfy (5.18), which in turns follows from their almost sure compactness and Proposition 2.32. ∎

We next describe the wedge construction. Let $(\mathcal{W}^{\rm l}, \mathcal{W}^{\rm r})$ be a left-right Brownian web and let $\hat{\mathcal{W}}^{\rm l}, \hat{\mathcal{W}}^{\rm r}$ be the associated dual webs. Recall from

Section 3.5 that $W(\hat{\pi}_1, \hat{\pi}_2)$ denotes the wedge defined by two downward paths $\hat{\pi}_1$ and $\hat{\pi}_2$. In the same section, we also defined what it means for a forward path π to enter a wedge $W(\hat{\pi}_1, \hat{\pi}_2)$. We again fix a countable dense set $\mathcal{D} \subset \mathbb{R}^2$ and define

$$\begin{aligned}\mathcal{N}_+ := \big\{\pi \in \Pi^\uparrow : &\pi \text{ does not enter wedges} \\ &\text{of the form } W(\hat{\pi}^{\rm r}_{z_1}, \hat{\pi}^{\rm l}_{z_2}) \text{ with } z_1, z_2 \in \mathcal{D}\big\}.\end{aligned}$$

This construction is known as the wedge construction of the Brownian net. See Figure 5.8 for an illustration. Note that here the left boundary of the wedge is formed by a dual right path and the right boundary is a dual left path. Because of the drift, these paths may fail to meet so the wedge may be infinite in size. In particular, the fact that paths do not enter wedges of this form implies that paths in $\mathcal{N}_+$ do not cross dual left paths from right to left, or dual right paths from left to right.

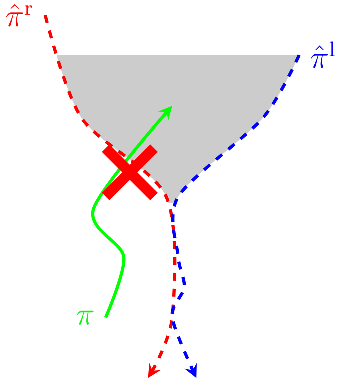

Figure 5.8: Illustration of the wedge construction of the Brownian net. Paths $\pi \in \mathcal{N}_+$ cannot enter wedges $W(\hat{\pi}^{\rm r}, \hat{\pi}^{\rm l})$ defined by a dual right and left path.

The following theorem, first proved in [SS08, Lemmas 4.5 and 4.7], is similar to Theorem 3.15 (and in fact historically predates it). We call the compact set $\mathcal{N} := \mathcal{N}_- = \mathcal{N}_+$ from the following theorem the *Brownian net*.

Theorem 5.22 (Characterisation of the Brownian net) *Let $\mathcal{D}$ be a countable dense subset of $\mathbb{R}^2$ and let $\mathcal{N}_-$ and $\mathcal{N}_+$ be defined in terms of a left-right Brownian web $(\mathcal{W}^{\rm l}, \mathcal{W}^{\rm r})$ and its dual as above. Then $\mathcal{N}_- = \mathcal{N}_+$.*

Proof (partial) Here we only prove the inclusion $\mathcal{N}_+ \subset \mathcal{N}_-$. The proof of the other inclusion will be combined with the proof of Theorem 5.23 below. The argument is similar to the proof of Theorem 3.15. We fix $\pi \in \mathcal{N}_+$,

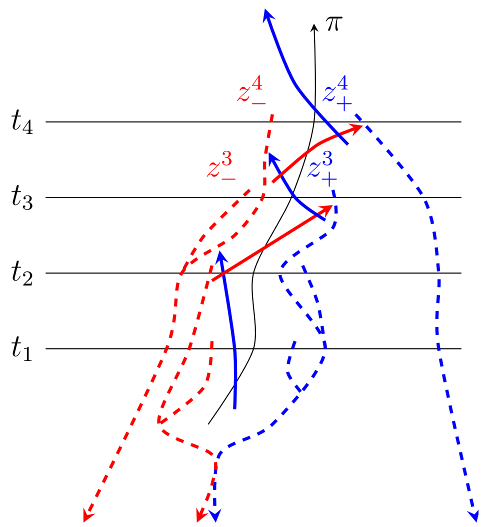

Figure 5.9: Construction showing that each path π that does not enter wedges of the dual left-right Brownian web can be approximated by hopping between paths in the forward left-right Brownian web. The structure of dual paths that capture the forward paths is reminiscent of a fish trap (Figure 5.10).

$\sigma_\pi < t_1 < \cdots < t_m$, and $\varepsilon > 0$. We claim that we can construct a path $\pi^{\rm hop}$ by hopping finitely often between paths in $(\pi^{\rm l}_z)_{z\in\mathcal{D}}$ and $(\pi^{\rm r}_z)_{z\in\mathcal{D}}$, such that $\sigma_\pi < \sigma_{\pi^{\rm hop}} < t_1$ and $|\pi^{\rm hop}(t_i) - \pi(t_i)| \leq \varepsilon$ for all $i = 1, \ldots, m$. To see this, for each $i = 1, \ldots, m$, we choose $z^i_\pm = (x^i_\pm, t^i_\pm) \in \mathcal{D}$ such that $t^i_\pm > t_i$ and

$$\pi(t_i) - \varepsilon < \hat\pi^{\rm r}_{z_-}(t_i) < \pi(t_i) < \hat\pi^{\rm l}_{z_+}(t_i) < \pi(t_i) + \varepsilon.$$

See Figure 5.9. Since π does not enter the wedge $W(\hat\pi^{\rm r}_{z^i_-}, \hat\pi^{\rm l}_{z^i_+})$, the meeting time of $\hat\pi^{\rm r}_{z^i_-}$ and $\hat\pi^{\rm l}_{z^i_+}$ must satisfy

$$\tau(\hat\pi^{\rm r}_{z^i_-}, \hat\pi^{\rm l}_{z^i_+}) \leq \sigma_\pi,$$

and we have $\hat\pi^{\rm r}_{z^i_-}(t) \leq \pi(t) \leq \hat\pi^{\rm l}_{z^i_+}(t)$ for all $t \in [\sigma_\pi, t_i]$. We can now choose $z = (x, s) \in \mathcal{D}$ such that $\sigma_\pi < s < t_1$ and

$$\sup_{1\leq i\leq m} \hat\pi^{\rm r}_{z^i_-}(t_1) < \pi^{\rm l}_z(t_1) < \inf_{1\leq i\leq m} \hat\pi^{\rm l}_{z^i_+}(t_1).$$

The forward left path $\pi^{\rm l}_z$ cannot cross any of the left downward paths $\hat\pi^{\rm l}_{z^i_+}$, but it can cross the right downward paths $\hat\pi^{\rm r}_{z^i_-}$. Just before it does so, however, we can hop onto a cleverly chosen forward right path and continue until it

Figure 5.10: A fish trap. Picture reused from:
https://commons.wikimedia.org/wiki/File:Stellnetzfischerei_(Reusen).jpg.

threatens to cross one of the left downward paths $\hat{\pi}^1_{z^i_+}$. Just before it does, we can again hop onto a left path, and so on.

We claim that in this way, we can construct a hopping path that after a finite number of steps arrives at the last time t_m. Indeed, since the dual left paths on the right of π cannot meet the dual right paths on the left of π (since otherwise π would enter a wedge created by two of these paths), there is some positive δ such that the distance between the closest dual left path on the right of π and the closest dual right path on the left of π is at least δ at any time between the time when we started our hopping path and the last time t_m. We can construct a hopping path so that the position where we hop from a left to a right path is always less than $\delta/3$ from the closest dual right path on the left, and similarly, the position where we hop from a right to a left path is always less than $\delta/3$ from the closest dual left path to the right. The times when we hop are increasing, so either we reach t_m in a finite number of steps, or the times when we hop increase to a limit that is $\leq t_m$. But then our hopping path comes infinitely often in the neighbourhood of two points that lie at least a distance $\delta/3$ apart. This clearly violates the equicontinuity of the left and right webs.

Let us write $\mathcal{N}_- = \overline{\mathcal{N}'_-}$, where $\mathcal{N}'_-$ is the set of paths that can be constructed by hopping between paths in $(\pi^{\rm l}_z)_{z\in\mathcal{D}}$ and $(\pi^{\rm r}_z)_{z\in\mathcal{D}}$. Since our hopping construction terminates after a finite number of steps, we have shown that for each $\pi \in \mathcal{N}_+$, $\varepsilon > 0$, and $t_1 < \cdots < t_n$, there exists a path $\pi' \in \mathcal{N}'_-$ such that $|\pi(t) - \pi'(t)| \leq \varepsilon$. Using the fact that $\mathcal{N}_-$ is compact (Lemma 5.21), we can now repeat the arguments at the end of the proof of Theorem 3.15 to show that $\mathcal{N}_+ \subset \overline{\mathcal{N}}'_-$. ∎

5.8 Convergence to the Brownian net

We are finally ready to state and prove the main result of this section.

Theorem 5.23 (Convergence to the Brownian net) *Let ε_n be positive constants tending to zero, let ω^n be a sequence of arrow configurations satisfying (5.4), and let $\mathcal{U}_n$ be the set of all open upward paths in ω^n. Then*

$$\mathbb{P}\big[\theta_{\varepsilon_n}(\overline{\mathcal{U}}_n) \in \cdot\,\big] \underset{n\to\infty}{\Longrightarrow} \mathbb{P}\big[\mathcal{N} \in \cdot\,\big],$$

where $\mathcal{N} := \mathcal{N}_- = \mathcal{N}_+$ is defined as in Theorem 5.22.

Proof This is very similar to the proof of Theorem 3.18. We start by showing that the laws

$$\big\{P\big[\theta_{\varepsilon_n}(\overline{\mathcal{U}}_n) \in \cdot\,\big] : n \in \mathbb{N}\big\}$$

are tight. We apply Proposition 2.33. For all $\pi \in \mathcal{U}_n$, we can estimate

$$\pi^{\rm l}_{(\pi(s),s)}(t) - \pi(s) \leq \pi(t) - \pi(s) \leq \pi^{\rm r}_{(\pi(s),s)}(t) - \pi(s),$$

where $\pi^{\rm l}_{(\pi(s),s)}$ and $\pi^{\rm r}_{(\pi(s),s)}$ denote the unique element of $\mathcal{U}^{\rm l}(\pi(s),s)$ and $\mathcal{U}^{\rm r}(\pi(s),s)$, respectively. In view of this, the tightness of the laws of $\theta_{\varepsilon_n}(\overline{\mathcal{U}}_n)$ follows easily from the tightness of the laws of $\theta_{\varepsilon_n}(\overline{\mathcal{U}}^{\rm l}_n)$ and $\theta_{\varepsilon_n}(\overline{\mathcal{U}}^{\rm r}_n)$ (Proposition 3.16).

With tightness proved, in view of Prohorov's theorem (Theorem 2.12) and Lemma 2.2, to prove the theorem, it suffices to prove that if a subsequence of the $\theta_{\varepsilon_n}(\overline{\mathcal{U}}_n)$ converges in law to a limit $\mathcal{N}$, then $\mathcal{N}$ is a Brownian net. We assume therefore, from now on, that we are given a subsequence such that the $\theta_{\varepsilon_n}(\overline{\mathcal{U}}_n)$ converges in law to a limit $\mathcal{N}$. By Skorohod's representation theorem, we can couple our random variables such that this convergence is almost sure:

$$\theta_{\varepsilon_n}(\overline{\mathcal{U}}_n) \underset{n\to\infty}{\longrightarrow} \mathcal{N}. \tag{5.19}$$

Similar to what we did in the proof of Theorem 3.18, we can extend this to include also the almost sure convergence of a number of other objects, that we already know converge in law.

We fix a deterministic countable dense set $\mathcal{D} \subset \mathbb{R}^2$. For each $z \in \mathcal{D}$, we choose $z_n \in \mathbb{Z}^2_{\rm even}$ and $z^n \in \mathbb{Z}^2_{\rm odd}$ such that $\theta_{\varepsilon_n}(z_n) \to z$ and $\theta_{\varepsilon_n}(z^n) \to z$. We let $\pi_n^{(n)\,\rm l}$ and $\pi_n^{(n)\,\rm r}$ denote the diffusively rescaled left and right paths started from z_n and we let $\hat\pi_n^{(n)\,\rm l}$ and $\hat\pi_n^{(n)\,\rm r}$ denote the diffusively rescaled dual left and right paths started from z^n. Then, for a suitable coupling, we have

$$\pi_z^{(n)\,\rm q} \underset{n\to\infty}{\longrightarrow} \pi_z^{\rm q} \text{ a.s.} \quad \text{and} \quad \tau(\pi_{z_1}^{(n)\,\rm q_1}, \pi_{z_2}^{(n)\,\rm q_2}) \underset{n\to\infty}{\longrightarrow} \tau(\pi_{z_1}^{\rm q_1}, \pi_{z_2}^{\rm q_2}) \text{ a.s.}$$

for all $z, z_1, z_2 \in \mathcal{D}$ and $\rm q, q_1, q_2 \in \{l, r\}$, and likewise for downward paths, where

$$\big(\pi_z^{\rm l}, \pi_z^{\rm r}, \hat\pi_z^{\rm l}, \hat\pi_z^{\rm r}\big)_{z\in\mathcal{D}}$$

are the forward and dual left and right paths started from $\mathcal{D}$ in a left-right Brownian web. With all this set up, we will show that the limit in (5.19) satisfies

$$\mathcal{N}_- \subset \mathcal{N} \subset \mathcal{N}_+, \tag{5.20}$$

where $\mathcal{N}_-$ and $\mathcal{N}_+$ are defined in terms of the forward and dual left and right paths started from $\mathcal{D}$. In particular, this then proves that $\mathcal{N}_- \subset \mathcal{N}_+$, which was the missing part of the proof of Theorem 5.22. Since the inclusion $\mathcal{N}_+ \subset \mathcal{N}_-$ has already been proved with the fish trap argument from Section 5.7, this then concludes the proofs of both Theorem 5.22 and 5.23.

It therefore remains to prove (5.20). We start by proving the inclusion $\mathcal{N}_- \subset \mathcal{N}$. Since $\mathcal{N}$ is closed, it suffices to prove $\mathcal{N}'_- \subset \mathcal{N}$ where $\mathcal{N}'_-$ is the set of paths that can be constructed by hopping between paths in $(\pi_z^{\rm l})_{z\in\mathcal{D}}$ and $(\pi_z^{\rm r})_{z\in\mathcal{D}}$. The statement now follows from the fact that $\tau(\pi_{z_1}^{(n)\,\rm l_1}, \pi_{z_2}^{(n)\,\rm r_2}) \underset{n\to\infty}{\longrightarrow} \tau(\pi_{z_1}^{\rm l}, \pi_{z_2}^{\rm r})$ for each $z_1, z_2 \in \mathcal{D}$, and the fact that in an arrow configuration, any path constructed by hopping between left and right paths is an open upward path.

The inclusion $\mathcal{N} \subset \mathcal{N}_+$ follows on the other hand from the fact that $\tau(\hat\pi_{z_1}^{(n)\,\rm l_1}, \hat\pi_{z_2}^{(n)\,\rm r_2}) \underset{n\to\infty}{\longrightarrow} \tau(\hat\pi_{z_1}^{\rm l}, \hat\pi_{z_2}^{\rm r})$ for each $z_1, z_2 \in \mathcal{D}$, the fact that in an arrow configuration, open paths cannot enter wedges, and Lemma 3.12. ∎

Exercise 5.24 *Show that almost surely, there exist no $\pi \in \mathcal{N}$, $\hat\pi^{\rm r} \in \hat{\mathcal{W}}^{\rm r}$ and $\hat\pi^{\rm l} \in \hat{\mathcal{W}}^{\rm l}$ such that π enters the wedge $W(\hat\pi^{\rm r}, \hat\pi^{\rm l})$. Note that we do not assume that the starting points of $\hat\pi^{\rm r}$ and $\hat\pi^{\rm l}$ lie in some fixed, deterministic, countable dense set. Hint: Lemma 4.7.*

5.9 The Brownian net with killing

We conclude this chapter with a very crude sketch of the proof of Theorem 5.1.

Proof of Theorem 5.1 (crude idea) We first consider the special case that $\alpha = 0$, $\beta = 1$, and $d_n = 0$ for all n. In this case, Theorem 5.23 tells us that the set $\theta_{\varepsilon_n}(\overline{\mathcal{V}}_n) \cap \Pi^\uparrow$ converges in law to the Brownian net $\mathcal{N}$. Since $d_n = 0$, each open path in an arrow configuration ω_n can be extended to an upward path, so

$$\overline{\mathcal{V}}_n = \big\{\pi \in \Pi : \exists \pi' \in \overline{\mathcal{V}}_n \cap \Pi^\uparrow \text{ s.t. } \pi \subset \pi'\big\}.$$

Using this, Theorem 5.23 is easily seen to imply (5.3), where the limit is

$$\mathcal{N}_* := \big\{\pi \in \Pi : \exists \pi' \in \mathcal{N} \text{ s.t. } \pi \subset \pi'\big\}.$$

It is not hard to generalise this to arbitrary $\alpha \in \mathbb{R}$ and $\beta > 0$. We just need to generalise our earlier definition of a left-right Brownian web in such a way that the left Brownian web $\mathcal{W}^{\rm l}$ has drift $\alpha - \beta$ and the right Brownian web $\mathcal{W}^{\rm r}$ has drift $\alpha + \beta$, and all arguments go through in a trivial way. We can even allow for the case $\beta = 0$, where now the limit is the Brownian web. Indeed, if $\beta = 0$, then Theorem 5.17 (convergence to the left-right Brownian web) remains true, where now the left and right Brownian webs are a.s. equal. This can be seen by adapting Theorem 5.13 (scaling limit of a left and right path), where now one does not need the left-right equation but simply uses that $L^{(n)}(t) \leq R^{(n)}(t)$ while $\mathbb{E}[R^{(n)}(t) - L^{(n)}(t)] \to 0$ as $n \to \infty$, for each $t \geq 0$.

In view of this, the real challenge is to prove Theorem 5.1 when the death probability d_n may be positive. Clearly, in this case we can no longer work with half-infinite paths. Nevertheless, in each arrow configuration ω_n, at each $z \in \mathbb{Z}^2_{\rm even}$, there start a unique *maximal* left path $\pi^{n\,\rm l}_z$ and right path $\pi^{n\,\rm r}_z$, which are defined by the fact that at branching points, they always choose the left or right arrow, respectively, and they only stop once they reach a death point, i.e., a point with no outgoing arrows. It is easy to check that under the conditions (5.2), diffusively rescaled left (resp. right) paths converge to Brownian motions with drift $\alpha - \beta$ (resp. $\alpha + \beta$) and an exponential life time with parameter β (i.e., with mean β^{-1}).

In view of this, one can still prove an analogon of Theorem 5.17 (convergence to the left-right Brownian web). Using the hopping construction, one can also still prove a lower bound $\mathcal{N}_-$ on the scaling limit of the set of all open paths. The most difficult part of the proof is to get a matching upper bound. We can naturally couple our arrow configurations to arrow configurations that have no deaths. These then give rise to a limiting left-right Brownian web that can be used to define wedges. Naturally, even with deaths, open paths cannot enter these wedges. To get a good upper bound, one needs to add one more condition that takes into account the deaths. It turns out that the right condition is that paths cannot pass through points where other paths have died. Proving this requires a better understanding of the Brownian net (such as density calculations and the concept of meshes). We cite [NRS15] for those who want to know more. ∎

Chapter 6

Properties of the Brownian net

6.1 The continuum biased voter model

We recall that in (4.1) we constructed collections of random maps $(\mathcal{X}_{s,t})_{s\leq t}$ and $(\mathcal{Y}_{s,t})_{s\leq t}$ that could be interpreted as scaling limits of the stochastic flows associated with a voter model and its associated dual system of coalescing random walks, constructed from their graphical representations. In the present section, we study the analogue objects for the scaling limits of biased voter models and their associated systems of branching and coalescing random walks. We start with a useful lemma, that is a direct consequence of the hopping and wedge constructions of the Brownian net. This lemma is illustrated in Figure 6.1.

Lemma 6.1 (Connections in the Brownian net) *Let $\mathcal{N}$ be a standard Brownian net. Then almost surely, for all $a, b, s, t \in \mathbb{R}$ with $a \leq b$ and $s < t$, if some $\pi \in \mathcal{N}(\mathbb{R} \times \{s\})$ satisfies $\pi(t) \in [a, b]$, then*

$$\tau\big(\hat{\pi}^{\mathrm{r}-}_{(a,t)}, \hat{\pi}^{\mathrm{l}+}_{(b,t)}\big) \leq s \quad \textit{and} \quad \hat{\pi}^{\mathrm{r}-}_{(a,t)} \leq \pi \leq \hat{\pi}^{\mathrm{l}+}_{(b,t)} \ \textit{on}\ [s,t]. \tag{6.1}$$

Conversely, almost surely, for all $a, b, s, t \in \mathbb{R}$ with $a \leq b$ and $s < t$, if there exist $\hat{\pi}^{\mathrm{r}} \in \hat{\mathcal{W}}^{\mathrm{r}}(a,t)$ and $\hat{\pi}^{\mathrm{l}} \in \hat{\mathcal{W}}^{\mathrm{l}}(b,t)$ such that $\hat{\pi}^{\mathrm{r}} < \hat{\pi}^{\mathrm{l}}$ on (s,t), then for each $x \in [\hat{\pi}^{\mathrm{r}}(s), \hat{\pi}^{\mathrm{l}}(s)]$, there exists a path $\pi \in \mathcal{N}(x,s)$ such that $\hat{\pi}^{\mathrm{r}} \leq \pi \leq \hat{\pi}^{\mathrm{l}}$ on $[s,t]$.

Proof By the wedge characterisation of the Brownian net, if $\pi \in \mathcal{N}(\mathbb{R} \times \{s\})$ satisfies $\pi(t) \in [a, b]$, then by the fact that π does not enter wedges (Exercise 5.24), for each $\varepsilon > 0$, paths $\hat{\pi}^{\mathrm{r}} \in \hat{\mathcal{W}}^{\mathrm{r}}(a-\varepsilon, t)$ and $\hat{\pi}^{\mathrm{l}} \in \hat{\mathcal{W}}^{\mathrm{l}}(b+\varepsilon, t)$ must satisfy

$$\tau\big(\hat{\pi}^{\mathrm{r}}, \hat{\pi}^{\mathrm{l}}\big) \leq s \quad \text{and} \quad \hat{\pi}^{\mathrm{r}} \leq \pi \leq \hat{\pi}^{\mathrm{l}} \text{ on } [s,t].$$

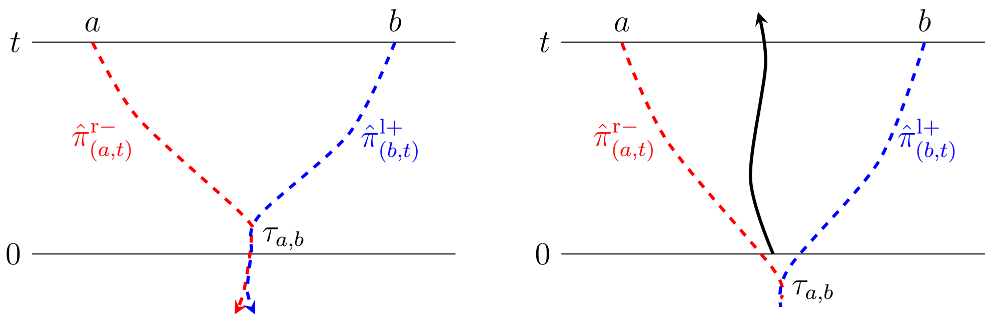

Figure 6.1: Illustration of Lemma 6.1. If $\tau_{a,b} > 0$, then no path in $\mathcal{N}$ starting at time 0 can pass through $[a, b]$ at time t. On the other hand, if $\tau_{a,b} \leq 0$, then at any point (x, s) with $\hat\pi^{\rm r-}_{(a,t)}(0) \leq x \leq \hat\pi^{\rm l+}_{(b,t)}(0)$ there starts a path $\pi \in \mathcal{N}$ with $\pi(t) \in [a, b]$.

Letting $\varepsilon \to 0$, we see that $\hat\pi^{\rm r-}_{(a,t)} \leq \pi \leq \hat\pi^{\rm l+}_{(b,t)}$ on $[s, t]$. Using Lemma 4.7, we moreover see that $\tau\big(\hat\pi^{\rm r-}_{(a,t)}, \hat\pi^{\rm l+}_{(b,t)}\big) \leq s$, proving (6.1).

Conversely, if for some $a, b, s, t \in \mathbb{R}$ with $a \leq b$ and $s < t$, there exist $\hat\pi^{\rm r} \in \hat{\mathcal{W}}^{\rm r}(a, t)$ and $\hat\pi^{\rm l} \in \hat{\mathcal{W}}^{\rm l}(b, t)$ such that $\hat\pi^{\rm r} < \hat\pi^{\rm l}$ on $[s, t]$, then by the fish-trap argument in the proof of Theorem 5.22, for each $x \in [\hat\pi^{\rm r}(s), \hat\pi^{\rm l}(s)]$, we can construct a path $\pi \in \mathcal{N}(x, s)$ such that $\hat\pi^{\rm r} \leq \pi \leq \hat\pi^{\rm l}$ on $[s, t]$. Using the compactness of $\mathcal{N}$, we can relax the condition that $\hat\pi^{\rm r} < \hat\pi^{\rm l}$ on $[s, t]$ to the weaker condition $\hat\pi^{\rm r} < \hat\pi^{\rm l}$ on (s, t). ∎

The following lemma is similar to Lemma 3.8. We leave its proof as an exercise to the reader.

Lemma 6.2 (Trivial paths) *Let $\mathcal{N}$ be a Brownian net. Then $\Pi^\uparrow_{\rm triv} \subset \mathcal{N}$ a.s. and each $\pi \in \mathcal{N} \backslash \Pi^\uparrow_{\rm triv}$ satisfies $\pi(t) \in \mathbb{R}$ for all $\sigma_\pi \leq t < \infty$.*

Exercise 6.3 *Prove Lemma 6.2.*

It is easy to see that a Brownian net $\mathcal{N}$ almost surely uniquely determines its associated left-right Brownian web $(\mathcal{W}^{\rm l}, \mathcal{W}^{\rm r})$, and hence also their assoctiated dual Brownian webs $\hat{\mathcal{W}}^{\rm l}$ and $\hat{\mathcal{W}}^{\rm r}$. In Lemma 5.20, we have seen that $(-\hat{\mathcal{W}}^{\rm l}, -\hat{\mathcal{W}}^{\rm r})$ is equally distributed with $(\mathcal{W}^{\rm l}, \mathcal{W}^{\rm r})$. It follows that we can construct a random collection of downward paths $\hat{\mathcal{N}}$ associated with $(\hat{\mathcal{W}}^{\rm l}, \hat{\mathcal{W}}^{\rm r})$, so that the triple $(\hat{\mathcal{W}}^{\rm l}, \hat{\mathcal{W}}^{\rm r}, \hat{\mathcal{N}})$ is equally distributed with $(-\mathcal{W}^{\rm l}, -\mathcal{W}^{\rm r}, -\mathcal{N})$. We call $\hat{\mathcal{N}}$ the *dual Brownian net* associated with $\mathcal{N}$.

In analogy with the definitions in (4.1), we define

$$\begin{aligned}\mathcal{X}_{s,t}(A) &:= \big\{x \in \overline{\mathbb{R}} : \exists \hat{\pi} \in \hat{\mathcal{N}}(x,t) \text{ s.t. } \hat{\pi}(s) \in A\big\},\\ \mathcal{Y}_{s,t}(A) &:= \big\{\pi(t) : \pi \in \mathcal{N}(A \times \{s\})\big\},\\ \hat{\mathcal{X}}_{t,s}(A) &:= \big\{x \in \overline{\mathbb{R}} : \exists \pi \in \mathcal{N}(x,s) \text{ s.t. } \pi(t) \in A\big\},\\ \hat{\mathcal{Y}}_{t,s}(A) &:= \big\{\hat{\pi}(s) : \hat{\pi} \in \hat{\mathcal{N}}(A \times \{t\})\big\}.\end{aligned} \tag{6.2}$$

We can think of the maps $(\mathcal{X}_{s,t})_{s\leq t}$ as a continuum analogue of the stochastic flow $(\mathbf{X}_{s,t})_{s\leq t}$ defined in Section 1.1. Let us fix closed sets $A, B \subset \overline{\mathbb{R}}$ and define, in analogy with (1.5),

$$A_t := \mathcal{X}_{0,t}(A) \quad \text{and} \quad B_t := \mathcal{Y}_{0,t}(B) \qquad (t \geq 0). \tag{6.3}$$

Then we can think of the process $(A_t)_{t\geq 0}$ as of some sort of continuum version of the biased voter model and similarly, we can think of $(B_t)_{t\geq 0}$ as a continuum version of branching and coalescing random walks. We call $(A_t)_{t\geq 0}$ the *continuum biased voter model* and $(B_t)_{t\geq 0}$ the *branching-coalescing point set.*

Informally, we can think of the branching-coalescing point set as branching and coalescing Brownian motions. This informal description is a bit too simplistic, however, since we cannot simply construct the process by letting coalescing Brownian motions branch with a finite rate. Indeed, such a description would not make sense, since whenever the branching would create two Brownian motions on the same positions, the two would coalesce immediately. We should think of the branching-coalescing point set as coalescing Brownian motions which in addition branch with "infinite" rate. However, since most of the particles created due to the branching disappear immediately due to the coalescence, on macroscopis scales, we see only finitely many successful branchings.

Lemma 6.4 (Basic properties) *One has $\mathcal{X}_{s,t}(A) \in \mathcal{K}(\overline{\mathbb{R}})$ for each $A \in \mathcal{K}(\overline{\mathbb{R}})$ and $s \leq t$. Moreover,*

$$\mathcal{X}_{s,t}(A \cup B) = \mathcal{X}_{s,t}(A) \cup \mathcal{X}_{s,t}(B) \qquad \big(A, B \in \mathcal{K}(\overline{\mathbb{R}}),\ s \leq t\big). \tag{6.4}$$

Analogue statements hold with $\mathcal{X}_{s,t}$ replaced by $\mathcal{Y}_{s,t}$. For each $A, B \in \mathcal{K}(\overline{\mathbb{R}})$ and $s, t \in \mathbb{R}$ with $s \leq t$, one has

$$1_{\{\mathcal{X}_{s,t}(A) \cap B \neq \emptyset\}} = 1_{\{A \cap \hat{\mathcal{Y}}_{t,s}(B) \neq \emptyset\}}. \tag{6.5}$$

Proof This follows from the same proofs as Lemmas 4.1 and 4.2. ∎

Lemma 6.1 implies in particular that (compare Lemma 4.27)

$$\mathcal{X}_{s,u}\big([x,y]\big) = \begin{cases} \big[\pi^{\mathrm{l}-}_{(x,s)}(t), \pi^{\mathrm{r}+}_{(y,s)}(t)\big] & \text{if } u \leq \tau\big(\pi^{\mathrm{l}-}_{(x,s)}, \pi^{\mathrm{r}+}_{(y,s)}\big), \\ \emptyset & \text{if } u > \tau\big(\pi^{\mathrm{l}-}_{(x,s)}, \pi^{\mathrm{r}+}_{(y,s)}\big). \end{cases} \tag{6.6}$$

Note that as a consequence, the continuum biased voter model $(A_t)_{t\geq 0}$ defined in (6.3), started in an initial state of the form $A_0 = [x, y]$, has left- but not right-continuous sample paths. When we combine (6.6) with (6.4), we see that the continuum biased voter model, started from a finite union of compact real intervals, has a rather simple description. At each time, it consists of a finite union of compact real intervals, whose boundaries evolve as drifted Brownian motions.

It is sometimes useful to view the continuum biased voter model and the branching-coalescing point set as taking values in the space of all closed subsets of the real line, instead of the space of all compact subsets of the extended real line. We observe that

$$\mathcal{K}' := \big\{A \in \mathcal{K}_+(\overline{\mathbb{R}}) : \{-\infty, \infty\} \subset A\big\}$$

is a closed subset of $\mathcal{K}_+(\overline{\mathbb{R}})$, and hence compact, by the compactness of the latter. We let $\mathrm{Cl}(\mathbb{R})$ denote the set of all closed subsets of the real line. We observe that the map $A \mapsto A \cap \mathbb{R}$ is a bijection from $\mathcal{K}'$ to $\mathrm{Cl}(\mathbb{R})$, which allows us to identify these two spaces. We equip $\mathrm{Cl}(\mathbb{R})$ with the topology that comes from its identification with $\mathcal{K}'$, making it into a compact metrisable space. It follows from Lemma 6.2 that $A_0 \in \mathcal{K}'$ implies $A_t \in \mathcal{K}'$ for all $t \geq 0$ and similarly for $(B_t)_{t\geq 0}$, so we can view these processes as processes with state space $\mathrm{Cl}(\mathbb{R})$ instead of $\mathcal{K}(\overline{\mathbb{R}})$.

As a consequence of (6.5), the continuum biased voter model and the branching-coalescing point set started in deterministic initial states satisfy a duality relation of the form

$$\mathbb{P}\big[A_0 \cap B_t \neq \emptyset\big] = \mathbb{P}\big[A_t \cap B_0 \neq \emptyset\big] \qquad \big(A_0, B_0 \in \mathrm{Cl}(\mathbb{R})\big). \tag{6.7}$$

One can prove that knowing this expression for all A_0 that are finite unions of compact intervals uniquely determines the law of B_t, viewed as a random variable with values in the space $\mathrm{Cl}(\mathbb{R})$. Thus, using duality, we can uniquely characterise the transition probabilities of the branching-coalescing point set in terms of the simpler continuum biased voter model.

It is possible to show that the operators $(\mathcal{X}_{s,t})_{s\leq t}$ and $(\mathcal{Y}_{s,t})_{s\leq t}$ have the stochastic flow property, similar to what we proved in the unbiased case (Lemma 4.23), and using this one can also show that the continuum biased voter model $(A_t)_{t\geq 0}$ and the branching-coalescing point set $(B_t)_{t\geq 0}$ are Markov processes, similar to our earlier Proposition 4.24. For brevity, we skip the details.

For the branching-coalescing point set, more is known. It has been proved in [SS08, Thm 1.11] that the branching-coalescing point set is a Feller process with continuous sample paths. Abstract Hille-Yosida theory tells us that each

Feller process is uniquely characterised by its generator, so in principle it should be possible to give a description of the branching-coalescing point set in terms of its generator. It is an open problem to give an explicit description of this generator. Perhaps the duality relation (6.7) is a good starting point to get an idea what sort of functions the domain of the generator should contain.

6.2 The branching-coalescing point set

The following proposition is similar to Proposition 4.3, but its proof is a bit more involved.

Proposition 6.5 (Density of the branching-coalescing point set) *The branching-coalescing point set satisfies*

$$\mathbb{E}\big[\big|\mathcal{Y}_{0,t}(\overline{\mathbb{R}}) \cap [a,b]\big|\big] = (b-a)\cdot\Big(\frac{e^{-t}}{\sqrt{\pi t}} + 2\Phi(\sqrt{2t})\Big) \tag{6.8}$$

$(a, b \in \mathbb{R},\ a < b,\ t > 0)$, *where* $\Phi(x) := \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x} e^{-y^2/2}\mathrm{d}y$ *is the distribution function of the normal distribution.*

Let

$$\Psi(t) := \frac{e^{-t}}{\sqrt{\pi t}} + 2\Phi(\sqrt{2t}) \qquad (t \geq 0) \tag{6.9}$$

denote the function on the right-hand side of (6.8). We observe that

$$\Psi(t) \sim \frac{1}{\sqrt{\pi t}} \quad \text{as } t \to 0 \quad \text{and} \quad \Psi(t) \underset{t\to\infty}{\longrightarrow} 2.$$

This shows that for small times, the density of the branching-coalescing point set is asymptotically the same as for the coalescing point set, but for large times, it is quite different since the density does not go to zero but tends to a positive limit.

Proof of Proposition 6.5 For $\varepsilon, t > 0$, set

$$F_\varepsilon(t) := \mathbb{P}\big[\tau\big(\hat{\pi}^{\mathrm{r}}_{(0,t)}, \hat{\pi}^{\mathrm{l}}_{(\varepsilon,t)}\big) > 0\big].$$

Then Lemma 6.1 tells us that

$$\mathbb{P}\big[\mathcal{Y}_{0,t}(\overline{\mathbb{R}}) \cap [0,\varepsilon] \neq \emptyset\big] = 1 - F_\varepsilon(t) \qquad (\varepsilon, t > 0),$$

and the claim of the proposition will follow from the argument used in the proof of Proposition 4.3, provided we show that

$$\lim_{\varepsilon\to 0} \varepsilon^{-1}\big(1 - F_\varepsilon(t)\big) = \Psi(t) \qquad (t > 0), \tag{6.10}$$

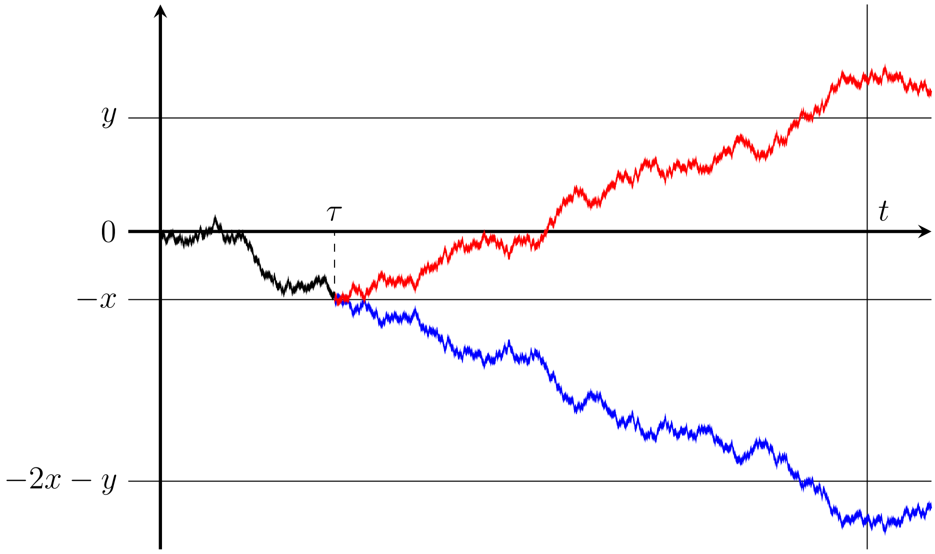

Figure 6.2: The reflection principle: $\mathbb{P}[\tau < t,\ B_t > y] = \mathbb{P}[B_t < -2x - y]$.

where Ψ is the function defined in (6.9). The difference of a right and left path is a Brownian motion with drift 2 and twice the quadratic variation of a standard Brownian motion. Therefore, we can express the probability we are interested in in terms of a standard Brownian motion $(B_t)_{t\geq 0}$ as

$$\begin{aligned} F_\varepsilon(t) &= \mathbb{P}\big[\inf_{0\leq s\leq t}\big(\sqrt{2}B_s + 2s\big) \leq -\varepsilon\big] \\ &= \mathbb{P}\big[\inf_{0\leq s\leq t}\big(B_s + \sqrt{2}s\big) \leq -\varepsilon/\sqrt{2}\big] \qquad (\varepsilon, t > 0). \end{aligned}$$

In line with notation introduced in (4.6), let us set

$$B'_t := B_t + \sqrt{2}t, \quad m_t(B) := \inf_{0\leq s\leq t} B_s, \quad \text{and} \quad m_t(B') := \inf_{0\leq s\leq t} B'_s \qquad (t \geq 0).$$

Using the reflection principle (see Figure 6.2), we see that

$$\mathbb{P}\big[-m_t(B) \geq x,\ B_t \geq y\big] = \mathbb{P}\big[B_t \geq 2x + y\big] \qquad (x \geq 0,\ y \geq -x,\ t \geq 0).$$

Differentiating the normal distribution, we see that the joint density of the law of $(-m_t(B), B_t)$ on the set $\{(x,y) \in \mathbb{R}^2 : x \geq 0,\ y \geq -x\}$ is given by

$$\begin{aligned} &\frac{1}{\sqrt{2\pi t}}\frac{\partial^2}{\partial x \partial y}\int_{2x+y}^{\infty} e^{-\frac{1}{2}z^2/t}\mathrm{d}z = \frac{1}{\sqrt{2\pi t}}\frac{\partial}{\partial x}e^{-\frac{1}{2}(2x+y)^2 t^{-1}} \\ &= \frac{1}{\sqrt{2\pi t}}2(2x+y)t^{-1}e^{-\frac{1}{2}(2x+y)^2 t^{-1}} =: H_t(x,y). \end{aligned}$$

By Girsanov's formula, the law of $(B'_s)_{0\le s\le t}$ has a density with respect to the law of $(B_s)_{0\le s\le t}$, as follows:

$$\mathbb{P}\big[(B'_s)_{0\le s\le t}\in \mathrm{d}f\big] = e^{\sqrt{2}f_t - t}\,\mathbb{P}\big[(B_s)_{0\le s\le t}\in \mathrm{d}f\big] \qquad \big(f\in \mathcal{C}_{[0,t]}(\mathbb{R})\big).$$

As a consequence, we obtain that

$$\begin{aligned}\mathbb{P}\big[-m_t(B'))\in \mathrm{d}x,\ B'_t\in \mathrm{d}y\big] &= e^{\sqrt{2}y - t}H_t(x,y)\,\mathrm{d}x\,\mathrm{d}y\\ &=: H'_t(x,y)\,\mathrm{d}x\,\mathrm{d}y.\end{aligned}$$

This allows us to express the probability we are interested in as

$$\begin{aligned}F_\varepsilon(t) = \mathbb{P}\big[-m_t(B')\ge \varepsilon/\sqrt{2}\big] &= \int_{\varepsilon/\sqrt{2}}^\infty \mathrm{d}x\int_{-x}^\infty \mathrm{d}y\, H'_t(x,y)\\ = \underbrace{\int_{-\varepsilon/\sqrt{2}}^\infty \mathrm{d}y\int_{\varepsilon/\sqrt{2}}^\infty \mathrm{d}x\, H'_t(x,y)}_{=:\,\mathrm{I}} &+ \underbrace{\int_{-\infty}^{-\varepsilon/\sqrt{2}} \mathrm{d}y\int_{-y}^\infty \mathrm{d}x\, H'_t(x,y)}_{=:\,\mathrm{II}}.\end{aligned}$$

We calculate

$$\begin{aligned}\mathrm{I} &= \frac{e^{-t}}{\sqrt{2\pi t}}\int_{-\varepsilon/\sqrt{2}}^\infty \mathrm{d}y\, e^{\sqrt{2}y}\int_{\varepsilon/\sqrt{2}}^\infty \mathrm{d}x\, 2(2x+y)t^{-1}e^{-\frac12(2x+y)^2t^{-1}}\\ &= \frac{e^{-t}}{\sqrt{2\pi t}}\int_{-\varepsilon/\sqrt{2}}^\infty \mathrm{d}y\, e^{\sqrt{2}y}e^{-\frac12(\sqrt{2}\varepsilon+y)^2t^{-1}}\\ &= e^{-2\varepsilon}\int_{-\varepsilon/\sqrt{2}}^\infty \mathrm{d}y\,\frac{1}{\sqrt{2\pi t}}e^{-\frac12(y+\sqrt{2}\varepsilon-\sqrt{2}\,t)^2t^{-1}}\\ &= e^{-2\varepsilon}\mathbb{P}\big[\sqrt{t}N-\sqrt{2}\varepsilon+\sqrt{2}\,t\ge -\varepsilon/\sqrt{2}\big]\\ &= e^{-2\varepsilon}\mathbb{P}\Big[N\ge \frac{\varepsilon}{\sqrt{2t}}-\sqrt{2t}\Big] = e^{-2\varepsilon}\Phi\Big(\sqrt{2t}-\frac{\varepsilon}{\sqrt{2t}}\Big),\end{aligned}$$

where N denotes a standard normally distributed random variable and in the second step we have used that

$$\begin{aligned}&-t+\sqrt{2}y-\tfrac12(\sqrt{2}\varepsilon+y)^2t^{-1} = -\tfrac12 t^{-1}\big[(y+\sqrt{2}\varepsilon)^2+2t^2-\sqrt{2}\,ty\big]\\ &= -\tfrac12 t^{-1}\big[(y+\sqrt{2}\varepsilon)^2+2t^2-\sqrt{2}\,t(y+\sqrt{2}\varepsilon)+2\varepsilon t\big]\\ &\quad= -\tfrac12 t^{-1}\big[(y+\sqrt{2}\varepsilon-\sqrt{2}\,t)^2+2\varepsilon t\big] = -2\varepsilon-\tfrac12(y+\sqrt{2}\varepsilon-\sqrt{2}\,t)^2t^{-1}.\end{aligned}$$

In a similar way, we calculate

$$\begin{aligned}
\mathrm{II} &= \frac{e^{-t}}{\sqrt{2\pi t}} \int_{-\infty}^{-\varepsilon/\sqrt{2}} \mathrm{d}y\, e^{\sqrt{2}y} \int_{-y}^{\infty} \mathrm{d}x\ 2(2x+y)t^{-1} e^{-\frac{1}{2}(2x+y)^2 t^{-1}} \\
&= \frac{e^{-t}}{\sqrt{2\pi t}} \int_{-\infty}^{-\varepsilon/\sqrt{2}} \mathrm{d}y\, e^{\sqrt{2}y} e^{-\frac{1}{2}(-2y+y)^2 t^{-1}} \\
&= \int_{-\infty}^{-\varepsilon/\sqrt{2}} \mathrm{d}y\, \frac{1}{\sqrt{2\pi t}} e^{-\frac{1}{2}(y-\sqrt{2}\,t)^2 t^{-1}} \\
&= \mathbb{P}\big[\sqrt{t}N + \sqrt{2}\,t \leq -\varepsilon/\sqrt{2}\big] = \Phi\Big(-\sqrt{2t} - \frac{\varepsilon}{\sqrt{2t}}\Big).
\end{aligned}$$

Putting everything together, we find that

$$F_\varepsilon(t) = e^{-2\varepsilon}\Phi\Big(\sqrt{2t} - \frac{\varepsilon}{\sqrt{2t}}\Big) + \Phi\Big(-\sqrt{2t} - \frac{\varepsilon}{\sqrt{2t}}\Big).$$

It follows that

$$\begin{aligned}
\tfrac{\partial}{\partial\varepsilon} F_\varepsilon(t) = &-2e^{-2\varepsilon}\,\Phi\Big(\sqrt{2t} - \frac{\varepsilon}{\sqrt{2t}}\Big) - e^{-2\varepsilon}\frac{1}{\sqrt{2t}}\Phi'\Big(\sqrt{2t} - \frac{\varepsilon}{\sqrt{2t}}\Big) \\
&- \frac{1}{\sqrt{2t}}\Phi'\Big(-\sqrt{2t} - \frac{\varepsilon}{\sqrt{2t}}\Big),
\end{aligned}$$

and

$$\begin{aligned}
\Psi(t) = -\tfrac{\partial}{\partial\varepsilon} F_\varepsilon(t)\big|_{\varepsilon=0} &= 2\Phi\big(\sqrt{2t}\big) + \frac{1}{\sqrt{2t}}\Phi'\big(\sqrt{2t}\big) + \frac{1}{\sqrt{2t}}\Phi'\big(-\sqrt{2t}\big) \\
&= 2\Phi\big(\sqrt{2t}\big) + \frac{2}{\sqrt{2t}}\frac{1}{\sqrt{2\pi}}e^{-t},
\end{aligned}$$

which agrees with (6.9). ∎

Although Proposition 6.5 shows that the branching-coalescing point set comes down from infinity, in is not true that $\mathcal{Y}_{0,t}(\overline{\mathbb{R}})$ is a locally finite point set for all $t > 0$. Indeed, it has been shown in [SSS09, Prop 3.14] that there exists a dense set of random times at which $\mathcal{Y}_{0,t}(\overline{\mathbb{R}})$ does not contain any isolated points. Similarly, the set of times t when $|\mathcal{Y}_{0,t}(\{0\})| = \infty$ is dense in $[0,\infty)$. This is a consequence of the "infinite" branching rate, which makes the number of particles explode immediately. To prove these results, one first needs to show that the set of times when a sticky reflected Brownian motion is at the origin is nowhere dense, which is quite easy to do. The consequence is that for the branching-coalescing point set started with a single particle, within each time interval of positive length, one can find a shorter time interval during which there are at least two particles. By the same principle, within such a shorter interval, one can find an even shorter time interval during which one of these particles has split and hence there are at least three particles, and so on, ad infinitum.

We conclude this section with two useful consequences of Proposition 6.5. Before we do so, we state the following lemma, that is similar to Lemma 4.18.

Lemma 6.6 (Independent increments) *Let $\mathcal{N}$ be a Brownian net and let $-\infty < t_0 \leq \cdots \leq t_n < \infty$. Then the restricted Brownian nets*

$$\mathcal{N}\big|_{[t_0,t_1]}, \ldots, \mathcal{N}\big|_{[t_{n-1},t_n]}$$

are independent.

Proof Same as the proof of Lemma 4.18, except that instead of Theorem 3.18 we now need to use (or rather slightly generalise) Theorem 5.23. ∎

We recall from Section 4.4 that a path $\pi \in \Pi^{\uparrow}$ enters a point $z = (x, u) \in \mathbb{R}^2$ if $\sigma_\pi < u$ and $\pi(u) = x$. We denote the set of Brownian net paths entering z by

$$\mathcal{N}_{\rm in}(z) := \big\{\pi \in \mathcal{N} : \pi \text{ enters } z\big\}.$$

Similar notation applies to the dual Brownian net $\hat{\mathcal{N}}$.

Lemma 6.7 (Forward and dual paths) *Let $\mathcal{N}$ be a Brownian net and let $\hat{\mathcal{N}}$ its associated dual Brownian net. Then for each deterministic $t \in \mathbb{R}$, there almost surely do not exist $x \in \mathbb{R}$ such that $\mathcal{N}_{\rm in}(x,t) \neq \emptyset$ and $\hat{\mathcal{N}}_{\rm in}(x,t) \neq \emptyset$. Moreover, for each $\pi \in \mathcal{N}$ and $\hat{\pi} \in \hat{\mathcal{N}}$, the set*

$$\big\{t \in [\sigma_\pi, \tau_{\hat{\pi}}] : \pi(t) = \hat{\pi}(t) \in \mathbb{R}\big\}$$

has Lebesgue measure zero.

Proof We write $\mathcal{Y}_{s,t}(\mathbb{R}) := \{\pi(t) : \pi \in \mathcal{N}(\mathbb{R} \times \{s\})\} = \mathcal{Y}_{s,t}(\overline{\mathbb{R}}) \cap \mathbb{R}$ and use similar notation for $\hat{\mathcal{Y}}_{u,t}$. By Proposition 6.5, for each deterministic $s < t < u$, the sets $\mathcal{Y}_{s,t}(\mathbb{R})$ and $\hat{\mathcal{Y}}_{u,t}(\mathbb{R})$ are locally finite. By Lemma 6.6, they are also independent. Using also that their laws are translation invariant, it follows that $\mathcal{Y}_{s,t}(\mathbb{R}) \cap \hat{\mathcal{Y}}_{u,t}(\mathbb{R}) = \emptyset$ a.s. In particular, this holds for all $s, u \in \mathbb{Q}$ with $s < t < u$, proving the first claim.

To prove also the second claim, fix deterministic $a, b, s, u \in \mathbb{R}$ with $a < b$ and $s < u$, and let C be the random subset of $[s, u]$ defined by

$$C := \big\{t \in [s,u] : \exists \pi \in \mathcal{N}(\overline{\mathbb{R}} \times \{s\}),\ \hat{\pi} \in \hat{\mathcal{N}}(\overline{\mathbb{R}} \times \{u\}) \text{ s.t. } \pi(t) = \hat{\pi}(t) \in [a,b]\big\}.$$

Using the fact that $\mathcal{N}$ and $\hat{\mathcal{N}}$ are compact, it is easy to see that C is closed, so it is certainly measurable. By Fubini and what we have just proved

$$\mathbb{E}\Big[\int_s^u 1_C(t)\,\mathrm{d}t\Big] = \int_s^u \mathbb{P}[t \in C]\,\mathrm{d}t = 0.$$

Now if $\pi \in \mathcal{N}$ and $\hat{\pi} \in \hat{\mathcal{N}}$ are arbitrary, then by what we have just proved, for each $a, b, s, u \in \mathbb{Q}$ such that $s \leq \sigma_\pi \leq \tau_{\hat{\pi}} \leq u$, the set

$$\big\{t \in [s,u] : \pi(t) = \hat{\pi}(t) \in [a,b]\big\}$$

has Lebesgue measure zero. Letting $s \downarrow \sigma_\pi$, $u \uparrow \tau_{\hat{\pi}}$, $a \to -\infty$, and $b \to \infty$, the claim follows. ∎

Lemma 6.8 (Bounding left and right paths) *Let $\mathcal{N}$ be a Brownian net and let $(\mathcal{W}^{\rm l}, \mathcal{W}^{\rm r})$ be its associated left-right Brownian web. Let $t \in \mathbb{R}$ be deterministic. Then almost surely, for each $x \in \mathbb{R}$ and for each $\pi \in \mathcal{N}_{\rm in}(x,t)$, there exist an $s \in \mathbb{R}$ with $\sigma_\pi \leq s < t$, as well as $\pi^{\rm l} \in \mathcal{W}^{\rm l}_{\rm in}(x,t)$ and $\pi^{\rm r} \in \mathcal{W}^{\rm r}_{\rm in}(x,t)$ with $\sigma_{\pi^{\rm l}} \vee \sigma_{\pi^{\rm r}} \leq s$, such that $\pi^{\rm l} \leq \pi \leq \pi^{\rm r}$ on $[s,\infty)$.*

Proof Let $S < t$ be deterministic, let $\pi \in \mathcal{N}$ satisfy $\sigma_\pi \leq S$, and let $b := \pi(t)$. By Proposition 6.5, the set $\mathcal{Y}_{S,t} \subset \mathbb{R}$ is locally finite, so there exists an $a \in \mathcal{Y}_{S,t}$ with $a < b$ such that $(a,b) \cap \mathcal{Y}_{S,t} = \emptyset$. By Lemma 6.1, for each $\varepsilon > 0$, we have

$$\tau\big(\hat{\pi}^{\rm r-}_{(a+\varepsilon,t)}, \hat{\pi}^{\rm l+}_{(b-\varepsilon,t)}\big) > S.$$

Taking the limit, we conclude that

$$\tau\big(\hat{\pi}^{\rm r+}_{(a,t)}, \hat{\pi}^{\rm l-}_{(b,t)}\big) \geq S.$$

Let $\mathcal{D} \subset \mathbb{R}^2$ be a deterministic countable dense set. By the remark below Lemma 4.7, if some $\pi^{\rm l} \in \mathcal{W}^{\rm l}$ and $\pi^{\rm r} \in \mathcal{W}^{\rm r}$ meet in a point (x,S), then there must be skeletal paths $\tilde{\pi}^{\rm l} \in \mathcal{W}^{\rm l}(\mathcal{D})$ and $\tilde{\pi}^{\rm r} \in \mathcal{W}^{\rm r}(\mathcal{D})$ that also meet in (x,S). Therefore, since S is deterministic and since left and right paths started from deterministic points do not meet at deterministic times, we conclude that

$$\tau := \tau\big(\hat{\pi}^{\rm r+}_{(a,t)}, \hat{\pi}^{\rm l-}_{(b,t)}\big) > S.$$

By Lemma 6.7, there exists a time $s \in \mathbb{R}$ with $\tau < s < t$ and $x \in \mathbb{R}$ such that $\hat{\pi}^{\rm l-}_{(b,t)}(s) < x < \pi(s)$. Now any $\pi^{\rm l} \in \mathcal{W}^{\rm l}(x,s)$ must satisfy $\hat{\pi}^{\rm l-}_{(b,t)} \leq \pi^{\rm l}$ on $[s,t]$ and $\pi^{\rm l} \leq \pi$ on $[s,\infty)$. It follows that $\pi^{\rm l} \in \mathcal{W}^{\rm l}_{\rm in}(x,t)$. By symmetry, we see that there must also exist an $s' < t$ and $\pi^{\rm r} \in \mathcal{W}^{\rm r}_{\rm in}(x,t)$ such that $\pi \leq \pi^{\rm r}$ on $[s',\infty)$. Since $S < t$ is arbitrary, the claim of the lemma follows. ∎

6.3 The backbone

The fact that the density $\Psi(t)$ tends to a positive limit as $t \to \infty$ suggests that the branching-coalescing point set should have an invariant law, and also, that the Brownian net should contain bi-infinite paths (contrary to the Brownian web, see Lemma 4.4). The following proposition confirms this. Note that according to our notation, $\mathcal{N}(*,-\infty)$ is the set of all paths $\pi \in \mathcal{N}$ with starting point $(*,-\infty)$, i.e., $\mathcal{N}(*,-\infty) = \mathcal{N} \cap \Pi^{\updownarrow}$. The set of paths $\mathcal{N}(*,-\infty)$ is called the *backbone* of the Brownian net. See Figure 6.3 for an illustration.

Proposition 6.9 (Backbone of the Brownian net) *The set*

$$\{\pi(0) : \pi \in \mathcal{N}(*,-\infty)\} \cap \mathbb{R}$$

is a Poisson point process with intensity 2. Moreover, $\mathcal{N}(*,-\infty)$ is equal in law to $-\mathcal{N}(*,-\infty)$.

Remark This proposition shows that the law of a Poisson point process with intensity 2 is an invariant law for the branching-coalescing point set. Moreover, since $\mathcal{N}(*,-\infty)$ is equal in law to $-\mathcal{N}(*,-\infty)$, this invariant law is reversible.

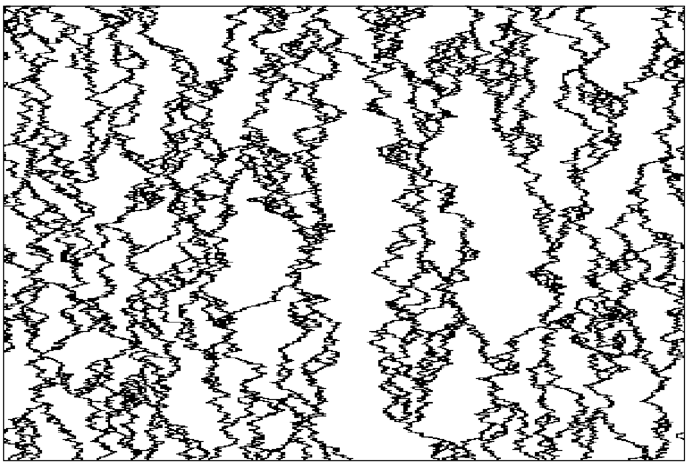

Figure 6.3: The backbone of the Brownian net.

We will first prove an analogue of Proposition 6.9 for the set of open paths in an arrow configuration, and then prove the statement about the Brownian net by finite approximation.

Proposition 6.10 (Backbone of an arrow configuration) *Let $\mathcal{U}$ be the set of open upward paths in an arrow configuration $\omega = (\omega_z)_{z\in\mathbb{Z}^2_{\rm even}}$ with*

$$\mathbb{P}\big[\omega_z = \{-1\}\big] = l, \quad \mathbb{P}\big[\omega_z = \{+1\}\big] = r,$$
$$\mathbb{P}\big[\omega_z = \{-1,+1\}\big] = b, \quad \mathbb{P}\big[\omega_z = \emptyset\big] = 0,$$

where $l + r + b = 1$ and $b > 0$. Set

$$X_t := \{\pi(t) : \pi \in \mathcal{U}(*,-\infty)\} \qquad (t \in \mathbb{Z}). \tag{6.11}$$

Then for each $t \in \mathbb{Z}_{\rm even}$ (resp. $t \in \mathbb{Z}_{\rm odd}$) the events $\{x \in X_t\}_{x\in\mathbb{Z}_{\rm even}}$ (resp. $\{x \in X_t\}_{x\in\mathbb{Z}_{\rm odd}}$) are independent with

$$\mathbb{P}[x \in X_t] = \frac{b}{(b+l)(b+r)} \qquad \big((x,t) \in \mathbb{Z}^2_{\rm even}\big). \tag{6.12}$$

Moreover, $\mathcal{U}(*,-\infty)$ is equal in law to $-\mathcal{U}(*,-\infty)$.

Proof We will write $\omega_t := (\omega_{(x,t)})_{x\in Z_{\rm even}}$ if t is even and $:= (\omega_{(x,t)})_{x\in Z_{\rm odd}}$ if t is odd. Fix $s \in \mathbb{Z}_{\rm even}$ and $p \in [0,1]$. Let X'_s be a random subset of $\mathbb{Z}_{\rm even}$ such that the events $\{x \in X'_s\}_{x\in\mathbb{Z}_{\rm even}}$ are independent and have probability p, and X'_s is independent of ω_s. Let X'_{s+1} be the random subset of $\mathbb{Z}_{\rm odd}$ defined by

$$X'_{s+1} := \big\{x-1 : x \in X'_s,\ -1 \in \omega_{(x,0)}\big\} \cup \big\{x+1 : x \in X'_s,\ +1 \in \omega_{(x,0)}\big\}.$$

We will show that it is possible to choose p such that the events

$$\begin{aligned} A_{x,y} &:= \{x \in X'_s \text{ and } y - x \in \omega_{(x,s)}\} \\ &\qquad \text{with} \quad x \in \mathbb{Z}_{\rm even},\ y \in \mathbb{Z}_{\rm odd},\ |x-y| = 1 \end{aligned} \tag{6.13}$$

are all independent. Clearly, $A_{x,y}$ is independent of $A_{x',y'}$ if $x \neq x'$, so it suffices to choose p such that $A_{x,x-1}$ is independent of $A_{x,x+1}$ for each x. We calculate

$$\mathbb{P}(A_{x,x-1}) = p(b+l), \quad \mathbb{P}(A_{x,x+1}) = p(b+r), \quad \text{and} \quad \mathbb{P}(A_{x,x-1} \cap A_{x,x+1}) = pb,$$

so we need

$$p^2(b+l)(b+r) = pb,$$

which is trivially satisfied for $p = 0$ and less trivially for

$$p = \frac{b}{(b+l)(b+r)}. \tag{6.14}$$

From now on, we fix p as in (6.14). We also fix $s \in \mathbb{Z}_{\rm even}$ and let X'_s be a random subset of $\mathbb{Z}_{\rm even}$ such that the events $\{x \in X'_s\}_{x\in\mathbb{Z}_{\rm even}}$ are i.i.d. with probability p and independent of $(\omega_t)_{t\geq s}$. For each $t \in \mathbb{Z}$ with $t \geq s$ define

$$X'_t = \big\{\pi(t) : \pi \in \mathcal{U}(X'_s \times \{s\})\big\} \qquad (t \geq s).$$

By our earlier remarks, for each $t \in \mathbb{Z}_{\rm even}$ with $t \geq s$, we have that the events $\{x \in X'_t\}_{x\in\mathbb{Z}_{\rm even}}$ are i.i.d. with probability p and independent of $(\omega_u)_{u\geq t}$. It follows that the joint laws of $(X'_t)_{t\geq s}$ and $(\omega_t)_{t\geq s}$ for different values of s are consistent, so by Kolmogorov's extension theorem we can couple ω to a stationary process $(X'_t)_{t\in\mathbb{Z}}$ such that

(i) For each $s \in \mathbb{Z}_{\rm even}$, the events $\{x \in X'_s\}_{x\in\mathbb{Z}_{\rm even}}$ are i.i.d. with probability p and independent of $(\omega_t)_{t\geq s}$.

(ii) $X'_t = \big\{\pi(t) : \pi \in \mathcal{U}(X'_s \times \{s\})\big\}$ for each $t \in \mathbb{Z}$ with $t \geq s$.

Recall from Section 3.1 that each arrow configuration ω defines a random directed graph $(\mathbb{Z}^2_{\rm even}, \vec{E})$ by

$$\vec{E} := \big\{\big((x,t),(x+\omega_{(x,t)}, t+1)\big) : (x,t) \in \mathbb{Z}^2_{\rm even}\big\}.$$

We let $\vec{E}'$ be the subset of arrows defined by

$$\vec{E}' := \big\{\big((x,t),(y,t+1)\big) \in \vec{E} : x \in X'_t\big\},$$

and we let $\mathcal{U}'$ be the subset of $\mathcal{U}$ consisting of all open upward paths π such that $\pi(t) \in X'_t$ for all $t \in I(\pi) \cap \mathbb{Z}$. Equivalently, this says that $\mathcal{U}'$ is the set of all upward paths π such that

$$\big((\pi(t),t),(\pi(t+1),t+1)\big) \in \vec{E}' \qquad (t \in I(\pi) \cap \mathbb{Z}),$$

with linear interpolation between integer times. It follows from the independence of the events in (6.13) that if we rotate the oriented graph $(\mathbb{Z}^2_{\rm even}, \vec{E}')$ over 180 degrees and reverse the direction of all arrows, then the new graph obtained in this way is equal in law to the original one. We call this the *rotational symmetry of* $\vec{E}'$.

Using this, we see immediately that each $(x,t) \in \mathbb{Z}^2_{\rm even}$ with $x \in X'_t$ is not only the starting point of a forward open path, but also the endpoint of an open path with starting time $-\infty$. As a consequence, for each $t \in \mathbb{Z}$, we have $X'_t \subset X_t$, where X_t is defined in (6.11). We claim that this is in fact an equality.

It suffices to prove this for $t = 0$. We start by noting that for each $y \in \mathbb{Z}_{\rm even}$, we can find $x, z \in X_0$ with $x < y < z$. The right open path starting from $(x,0)$ and the left open path starting from $(z,0)$ a.s. meet eventually, say in the point (y',t). These are paths in the oriented graph $(\mathbb{Z}^2_{\rm even},\vec{E}')$, so using rotational symmetry, we see that similarly, for each $y\in\mathbb{Z}_{\rm even}$, there a.s. exists a point $(y',-t)\in\mathbb{Z}^2_{\rm even}$ at which there start paths in the oriented graph $(\mathbb{Z}^2_{\rm even},\vec{E}')$ that pass at time zero on the left and right of y, respectively.

We now prove that $X_0\subset X'_0$. Let $y\in X_0$. By our previous argument, there a.s. exists a point $(y',-t)\in\mathbb{Z}^2_{\rm even}$ such that $y'\in X'_{-t}$ and such that there exist open paths that start at $(y',-t)$ and that pass at time zero on the left and right of y, respectively. By the definition of X_0, there also exists an open path $\pi\in\mathcal{U}(*,-\infty)$ with $y=\pi(0)$. Since this open path must cross the two open paths starting from $(y',-t)$, there must also exist an open path from $(y',-t)$ to $(y,0)$, proving that $y\in X'_0$.

From the fact that $X_t=X'_t$ and our construction of the latter, we see that for each $t\in\mathbb{Z}_{\rm even}$ (resp. $t\in\mathbb{Z}_{\rm odd}$) the events $\{x\in X_t\}_{x\in\mathbb{Z}_{\rm even}}$ (resp. $\{x\in X_t\}_{x\in\mathbb{Z}_{\rm odd}}$) are independent with probabilities given in (6.12). The fact that $\mathcal{U}(*,-\infty)$ is equal in law to $-\mathcal{U}(*,-\infty)$ follows from the rotational symmetry of $\vec{E}'$. ∎

Proof of Proposition 6.9 Let ε_n be positive constants tending to zero, let ω^n be a sequence of arrow configurations satisfying (5.4), and let $\mathcal{U}_n$ be the set of all open upward paths in ω^n. Then Theorem 5.23 tells us that

$$\mathbb{P}\big[\theta_{\varepsilon_n}(\overline{\mathcal{U}}_n)\in\cdot\,\big]\underset{n\to\infty}{\Longrightarrow}\mathbb{P}\big[\mathcal{N}\in\cdot\,\big],$$

where $\mathcal{N}$ is the Brownian net. Since $\overline{\mathcal{U}}_n(*,-\infty)$ is a closed subset of $\overline{\mathcal{U}}_n$, it is compact. In view of Lemma 2.17, the tightness of the laws $\mathbb{P}\big[\theta_{\varepsilon_n}(\overline{\mathcal{U}}_n)\in\cdot\,\big]$ means that for each $\eta>0$, there exists a compact set $\mathcal{C}\subset\Pi^\uparrow$ such that $\inf_n\mathbb{P}[\theta_{\varepsilon_n}(\overline{\mathcal{U}}_n)\subset\mathcal{C}]\geq 1-\eta$. Since $\overline{\mathcal{U}}_n(*,-\infty)\subset\overline{\mathcal{U}}_n$, it immediatey follows that the laws $\mathbb{P}\big[\theta_{\varepsilon_n}(\overline{\mathcal{U}}_n(*,-\infty))\in\cdot\,\big]$ are also tight. Thus, by going to a subsequence, we can assume that

$$\mathbb{P}\big[\theta_{\varepsilon_n}(\overline{\mathcal{U}}_n,\overline{\mathcal{U}}_n(*,-\infty))\in\cdot\,\big]\underset{n\to\infty}{\Longrightarrow}\mathbb{P}\big[(\mathcal{N},\mathcal{N}')\in\cdot\,\big],$$

where $\mathcal{N}$ is a Brownian net and $\mathcal{N}'$ is a random compact subset of $\Pi^\uparrow$. Since limits of bi-infinite paths are bi-infinite we have $\mathcal{N}'\subset\Pi^\updownarrow$, and since $\overline{\mathcal{U}}_n(*,-\infty)\subset\overline{\mathcal{U}}_n$ for all n we have $\mathcal{N}'\subset\mathcal{N}$. Proposition 6.10 tells us that for each $t\in\mathbb{Z}_{\rm even}$, the set X_t in (6.11) is an i.i.d. subset of $\mathbb{Z}_{\rm even}$ with intensity

$$p_n:=\frac{b_n}{(b_n+l_n)(b_n+r_n)}\sim\frac{\varepsilon_n}{(\frac{1}{2}+\varepsilon_n)^2}\sim 4\varepsilon_n.$$

Taking the scaling limit, using (5.4), we see that

$$\big\{\pi(t):\pi\in\mathcal{N}'\big\}\cap\mathbb{R}$$

is a Poisson point set with intensity 2. (Note that the limiting intensity is 2 and not 4 since X_t is an i.i.d. subset of $\mathbb{Z}_{\rm even}$ not $\mathbb{Z}$.) The fact that $\mathcal{N}'$ is equal in distribution with $-\mathcal{N}'$ follows from the rotational symmetry of $\mathcal{U}_n(*, -\infty)$. Therefore, to complete the proof, it suffices to prove that $\mathcal{N}' = \mathcal{N}(*, -\infty)$. The inclusion $\mathcal{N}' \subset \mathcal{N}(*, -\infty)$ is clear since $\mathcal{N}' \subset \Pi^{\updownarrow}$ and $\mathcal{N}' \subset \mathcal{N}$. The opposite inclusion can be proved in almost exactly the same way as the inclusion $X_t \subset X'_t$ in the proof of Proposition 6.10, so we omit the details and only sketch the main line of the argument.

First, one needs to prove that if a path $\pi \in \mathcal{N}$ crosses a path $\pi' \in \mathcal{N}'$, then the path constructed by first following π' and then hopping onto π at their first meeting time is also a path in $\mathcal{N}'$. The approximating backbones clearly have this property and since crossing of the limiting paths means that also all except finitely many of the approximating paths must cross, this property is preserved in the limit. (Compare Exercise 6.11 below.) Next, similar to what we did in the proof of Proposition 6.10, one can use rotational symmetry to show that each path $\pi \in \mathcal{N}(*, -\infty) \backslash \Pi^{\uparrow}_{\rm triv}$ must cross a path in $\mathcal{N}'$ at some time s and hence, by our earlier remark, there must exist a path $\pi' \in \mathcal{N}'$ such that $\pi(t) = \pi'(t)$ for all $t \geq s$. It is then not hard to see that the crossing time s can be chosen arbitrarily small, so taking the limit, using the compactness of $\mathcal{N}'$, one obtains that $\pi \in \mathcal{N}'$. ∎

Exercise 6.11 (Hopping in the net) *Let $\mathcal{N}$ be a Brownian net and let $\pi, \pi' \in \mathcal{N}$ satisfy $\pi(s) < \pi'(s)$ and $\pi'(u) < \pi(u)$ for some $\sigma_\pi \vee \sigma_{\pi'} \leq s < u$. Define $\tau := \inf\{t > \sigma_\pi \vee \sigma_{\pi'} : \pi(t) = \pi'(t)\}$ and let π'' be the path defined by $\sigma_{\pi''} := \sigma_\pi$, $\pi''(t) := \pi(t)$ for $\sigma_\pi \leq t \leq \tau$, and $\pi''(t) := \pi'(t)$ for $t \geq \tau$. Show that $\pi'' \in \mathcal{N}$. Hint: use finite approximation.*

6.4 Law of a forward and dual path

In Section 5.5, we described the scaling limit of the joint law of a left and right path in an arrow configuration, which was then used in Section 5.6 to describe the law of the left-right Brownian web. This approach closely follows the original introduction of the left-right Brownian web in [SS08]. In the present section, we will described the scaling limit of the joint law of a forward left and dual right path in an arrow configuration. Since a Brownian web is almost surely uniquely determined by its dual, this approach can be used to give an alternative characterisation of the left-right Brownian web which, as we will see, has certain advantages over the characterisation given in Section 5.6.

Let ε_n be positive constants tending to zero and let ω^n be a sequence of arrow configurations satisfying (5.4). As in Section 5.2, we let $\omega^{\mathrm{l},n}$ and $\omega^{\mathrm{r},n}$ denote the left and right arrow configurations associated with ω^n and we let $\hat\omega^{\mathrm{l},n}$ and $\hat\omega^{\mathrm{r},n}$ denote the corresponding dual arrow configurations. We denote the set of open upward paths in $\omega^{\mathrm{l},n}$ and $\omega^{\mathrm{r},n}$ by $\mathcal{U}^{\mathrm{l}}_n$ and $\mathcal{U}^{\mathrm{r}}_n$ and we denote the set of open downward paths in $\hat\omega^{\mathrm{l},n}$ and $\hat\omega^{\mathrm{r},n}$ by $\mathcal{U}^{\mathrm{l}*}_n$ and $\mathcal{U}^{\mathrm{r}*}_n$.

We now proceed very similarly to what we did in the proof of Theorem 4.16. We fix $(y_n, u_n) \in \mathbb{Z}^2_{\mathrm{odd}}$ and let $\hat R_n$ be the unique element of $\mathcal{U}^{\mathrm{r}*}(y_n, u_n)$. We also fix $(x_n, s_n) \in \mathbb{Z}^2_{\mathrm{even}}$ and let $(X^n_k)_{k\geq s_n+1}$ be i.i.d. $\{-1,+1\}$-valued random variables, independent of $\hat R_n$, such that $\mathbb{P}[X^n_k = -1] = l_n + b_n$ and $\mathbb{P}[X^n_k = +1] = l_n$. We let L_n be the random walk that is defined for integer times by

$$L_n(t) := x_n + \sum_{k=s_n+1}^{t} X^n_k \qquad (t \geq s_n),$$

and then for general $t \geq s_n$ by linear interpolation. We can then define a reflected random walk $L'_n = (L'_n(t))_{t\geq s_n}$ started at $L'_n(s_n) = x_n$ first for integer times by

$$L'_n(t+1) := \begin{cases} L'_n(t) + \omega^{\mathrm{l},n}_{(L'_n(t),t)} & \text{if } t < u_n \text{ and } \hat R_n(t+1) = L'_n(t), \\ L'_n(t) + X^n_{t+1} & \text{otherwise,} \end{cases} \tag{6.15}$$

and then for general $t \geq s_n$ by linear interpolation. Then it is easy to see that the conditional law of L'_n given $\hat R_n$ is precisely the conditional law of the the unique element of $\mathcal{U}^{\mathrm{l}}(x_n, s_n)$ given $\hat R_n$.

We write $L'_n(t) = L_n(t) + \Psi_n(t)$, where Ψ_n is a reflection function. We now distinguish two cases. If $x_n < \hat R_n(s_n)$, then we observe that precisely as in the proof of Theorem 4.16, L'_n is the path L_n reflected to the left off $(\hat R_n(t) - 1)_{t\leq u_n}$ in the sense of Lemma 4.14. In the opposite case, when $\hat R_n(s_n) < x_n$, the evolution of L'_n is initially equal to the evolution of the the path L_n reflected to the right off $(\hat R_n(t)+1)_{t\leq u_n}$ in the sense of Lemma 4.13, but L'_n starts to behave differently after the random time

$$\tau_n := \inf\big\{t \geq s_n :\ t < u_n,\ \hat R_n(t+1) = L'_n(t),\\ \omega^n_{(L'_n(t),t)} = \{-1,+1\},\ X^n_{t+1} = -1\big\},$$

because when $(L'_n(t), t)$ is a branching point and $\hat R_n(t+1) = L'_n(t)$, the path L'_n can cross $\hat R_n$. After crossing, L'_n starts to evolve as the path L_n reflected to the left off $(\hat R_n(t) - 1)_{t\leq u_n}$ in the sense of Lemma 4.14. See Figure 6.4 for an illustration.

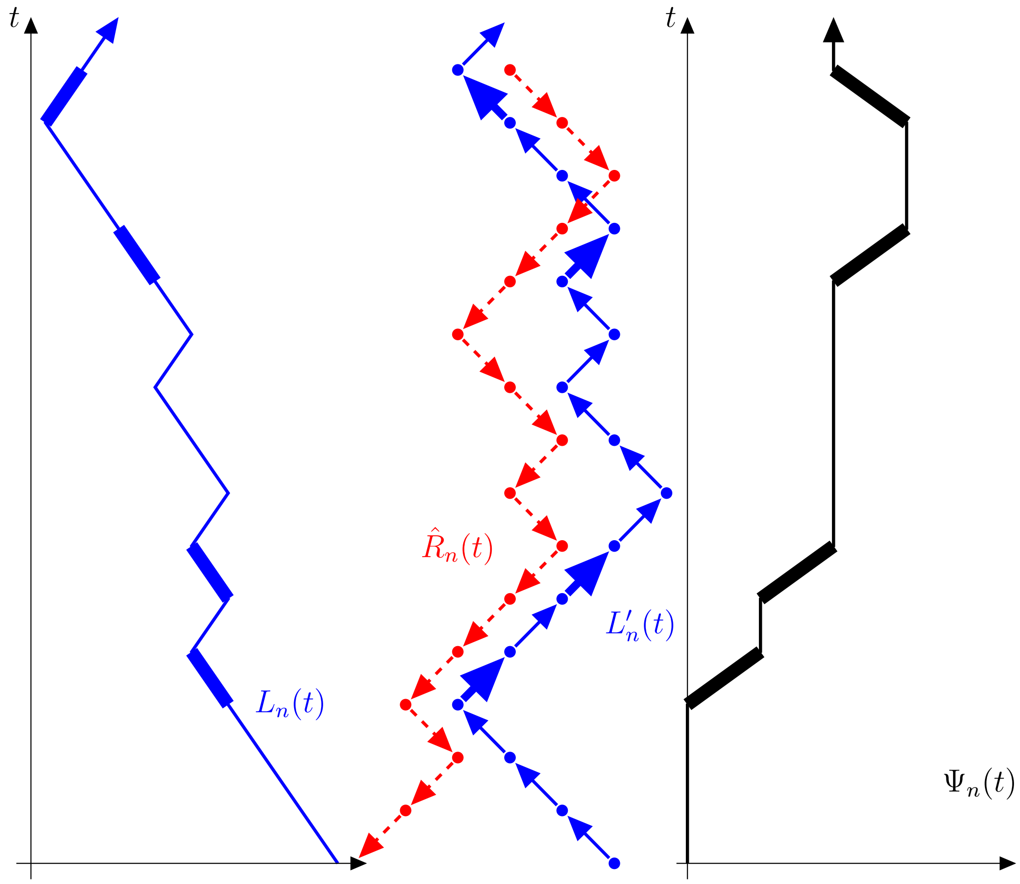

Figure 6.4: A forward left path reflected to the right off a dual right path $\hat{R}_n$. The reflected path is $L'_n(t) = L_n(t) + \Psi_n(t)$, where L_n is an independent random walk and Ψ_n is a reflection term. The steps where reflection takes place are indicated with a greater line thickness.

We observe that for all times $t \leq \tau_n$, we have that $\Psi_n(t)$ is twice the number of times $t' \in \{s_n, \ldots, t-1\}$ when $\hat{R}_n(t'+1) = L'_n(t')$ and $X^n_{t'+1} = -1$, i.e., the number of time when L'_n has attempted to cross $\hat{R}_n$. Each attempt is succesful with probability b_n, so the maximal height that Ψ_n will reach before it starts to decrease again is a random variable of the form $H_n = 2(G_n - 1)$, where G_n is geometrically disctributed with success probability b_n. In the diffusive scaling limit, we have $b_n \sim \varepsilon_n$ while we rescale Ψ_n by a factor ε_n, which means the maximal height the reflection term Ψ in the scaling limit will reach is exponentially distributed with mean 2.

Recall the reflection map Φ defined in (4.14), that takes as its input a pair

$(\pi,\hat\pi)\in\Pi^\uparrow\times\Pi^\downarrow$ and produces as its output a pair (π',ψ) where π' is the path π reflected off $\hat\pi$, and ψ is a reflection function. Our previous considerations motivate us to define a modified map Φ' of the form

$$\Pi^\uparrow\times\Pi^\downarrow\times[0,\infty)\ni(\pi,\hat\pi,T)\overset{\Phi'}{\longmapsto}(\pi'',\psi^-,\psi^+)\in\Pi^\uparrow\times\mathcal{C}(\mathbb{R})^2, \qquad (6.16)$$

with the following description: we first calculate $(\pi',\psi):=\Phi(\pi,\hat\pi)$, where Φ is the reflection map defined in (4.14). Let S be the random time defined by

$$S:=\begin{cases}\inf\{t\geq\sigma_\pi:\psi(t)=T\} & \text{if } \sigma_\pi<\tau_{\hat\pi} \text{ and } \hat\pi(\sigma_\pi)<\pi(\sigma_\pi),\\ 0 & \text{otherwise,}\end{cases} \qquad (6.17)$$

with the convention that $\inf\emptyset:=\infty$. In the case when $S<\infty$, we define a modified forward path $\tilde\pi$ by

$$\sigma_{\tilde\pi}:=S \quad\text{and}\quad \tilde\pi(t):=\pi(t)+\hat\pi(S)-\pi(S) \qquad (t\geq S),$$

and we set $(\pi^*,\psi^*):=\Phi(\tilde\pi,\hat\pi)$, where Φ is the reflection map in (4.14). We then define the map Φ' by setting

$$\big(\pi''(t),\psi^-(t),\psi^+(t)\big):=\begin{cases}\big(\pi'(t),0,\psi(t)\big) & (t\leq S),\\ \big(\pi^*(t),\psi^*(t),\psi(S)\big) & (S\leq t).\end{cases}$$

Note that here the precise way we have defined the reflection map in (4.14) when the forward path starts on the position of the dual path becomes important. In Section 4.3, we used the convention that in such a case, the reflected path is reflected to the left off the dual path, which is what we need here. Note also that as a result of our definitions

$$\pi''(t)=\pi(t)+\psi^+(t)-\psi^-(t) \qquad (t\geq\sigma_\pi), \qquad (6.18)$$

so that the difference $\psi^+-\psi^-$ of the two reflection functions corresponds to the reflection function Ψ_n that we introduced earlier in the context of reflected left paths in arrow configurations. In the light of our previous considerations, the following theorem should not come as a surprise.

Theorem 6.12 (Law of forward left and dual right path) *Assume that* $(\mathcal{W}^{\rm l},\mathcal{W}^{\rm r},\hat{\mathcal{W}}^{\rm l},\hat{\mathcal{W}}^{\rm r})$ *is a left-right Brownian web together with its associated dual Brownian webs. Let* $(x,s),(y,u)\in\mathbb{R}^2$, *and let* $\pi^{\rm l},\hat\pi^{\rm r}$ *be the almost surely unique paths such that* $\pi^{\rm l}\in\mathcal{W}^{\rm l}(x,s)$ *and* $\hat\pi^{\rm r}\in\hat{\mathcal{W}}^{\rm r}(y,u)$. *Let* $B=(B_t)_{t\geq0}$ *be a standard Brownian motion, independent of* $\hat\pi$, *and let* T *be an exponentially distributed random variable with mean 2, independent of* $\hat\pi$ *and* B. *Let* $L=(L_t)_{t\geq s}$ *be defined by* $L_{s+t}:=x+B_t-t$ $(t\geq0)$ *and let* $(\pi',\psi^-,\psi^+):=\Phi'(L,\hat\pi,T)$, *where* Φ' *is the map in (6.16). Then* $(\pi,\hat\pi)$ *is equal in law to* $(\pi',\hat\pi)$.

The proof of Theorem 6.12 uses finite approximation and is very similar to the proof of Theorem 4.16. We will use the following lemma, that plays the same role as Lemma 4.15. Recall that we equipped the space $\mathcal{C}(\mathbb{R})$ with the topology of uniform convergence. By definition, a *point of continuity* of Φ' is a triple $(\pi,\hat\pi,T)\in\Pi^\uparrow\times\Pi^\downarrow\times[0,\infty)$ with the property that if $(\pi_n,\hat\pi_n,T_n)\in\Pi^\uparrow\times\Pi^\downarrow\times[0,\infty)$ converge to $(\pi,\hat\pi,T)$, then $\Phi'(\pi_n,\hat\pi_n,T_n)$ converges to $\Phi'(\pi,\hat\pi,T)$ (in the topology on $\Pi^\uparrow\times\mathcal{C}(\mathbb{R})^2$).

Lemma 6.13 (Almost sure points of continuity) *Let $\pi\in\Pi^\uparrow$ and $\hat\pi\in\Pi^\downarrow$ be deterministic paths that satisfy either $\sigma_\pi\geq\tau_{\hat\pi}$ or $\sigma_\pi<\tau_{\hat\pi}$ and $\hat\pi(\sigma_\pi)\neq\pi(\sigma_\pi)$, and let T be a $[0,\infty)$-valued random variable with an atomless distribution. Then $(\pi,\hat\pi,T)$ is almost surely a point of continuity of Φ'.*

Proof Let $(\pi,\hat\pi,T)\in\Pi^\uparrow\times\Pi^\downarrow\times[0,\infty)$ be deterministic and assume that $(\pi_n,\hat\pi_n,T_n)\in\Pi^\uparrow\times\Pi^\downarrow\times[0,\infty)$ converge to $(\pi,\hat\pi,T)$. Assume also that either $\sigma_\pi\geq\tau_{\hat\pi}$ or $\sigma_\pi<\tau_{\hat\pi}$ and $\hat\pi(\sigma_\pi)\neq\pi(\sigma_\pi)$. Then Lemma 4.15 tells us that $\Phi(\pi_n,\hat\pi_n)\to\Phi(\pi,\hat\pi)$. Let us write $(\pi'_n,\psi_n):=\Phi(\pi_n,\hat\pi_n)$ and $(\pi',\psi):=\Phi(\pi,\hat\pi)$, and let us define S_n and S in terms of ψ_n and T_n, respectively ψ and T, as in (6.17). Provided that $S_n\to S$, we can then again apply Lemma 4.15 as well as the remark below it to conclude that $\Phi'(\pi_n,\hat\pi_n,T_n)$ converges to $\Phi'(\pi,\hat\pi,T)$.

What is complicating matters is that according to our definition, S does not depend continuously on ψ and T. Indeed, if $\psi_n\to\psi$ (uniformly on $\mathbb{R}$) and $T_n\to T$, then S_n may fail to converge to S precisely when $\inf\{t\geq\sigma_\pi:\psi(t)=T\}$ differs from $\sup\{t\geq\sigma_\pi:\psi(t)=T\}$, i.e., when the nondecreasing function ψ has a plateau precisely at the level T. However, each nondecreasing real function can have only countably many plateaus, so if T is random with an atomless distribution, then this problem almost surely does not occur. ∎

Proof of Theorem 6.12 Most of the work has already been done. We use finite approxiation. We fix $(x_n,s_n)\in\mathbb{Z}^2_{\rm even}$ and $(y_n,u_n)\in\mathbb{Z}^2_{\rm odd}$ such that

$$\theta_{\varepsilon_n}(x_n,s_n)\underset{n\to\infty}{\longrightarrow}(x,s)\quad\text{and}\quad\theta_{\varepsilon_n}(y_n,u_n)\underset{n\to\infty}{\longrightarrow}(y,u).$$

If $u\leq s$, then we can choose $u_n\leq s_n$ and the statement of the theorem follows easily from Theorem 5.19, so we assume from now on that $s<u$. We then choose also $s_n<u_n$ for all n.

We let ω^n be arrow configurations satisfying (5.4) and let $\hat R_n$ be the unique element of $\mathcal{U}^{\rm r*}(y_n,u_n)$. We let L_n be an independent drifted random walk started from (x_n,s_n) as before, and we define a reflected path L'_n as in (6.15). Then the conditional law of L'_n given $\hat R_n$ is precisely the conditional law of the the unique element of $\mathcal{U}^1(x_n,s_n)$ given $\hat R_n$.

For our present purposes, it will be convenient to construct L'_n in a slightly different way. We define a shifted random walk $\tilde L_n$ by $\tilde L_n(t) := L_n(t)+1$ $(t\geq s_n)$ if $L_n(s_n)<\hat R_n(s_n)$ and $\tilde L_n(t) := L_n(t)-1$ $(t\geq s_n)$ if $\hat R_n(s_n)<L_n(s_n)$ and let G_n be a geometrically distributed random variable, independent of everything else, such that

$$\mathbb{P}[G_n=k]=(1-\varepsilon_n)^{k-1}\varepsilon_n \qquad (k\geq 1).$$

We then set

$$\big(\tilde L'_n,\Psi^-_n,\Psi^+_n\big) := \Phi'(\tilde L_n,\hat R_n,2G_n), \tag{6.19}$$

where Φ' is the map in (6.16). Let

$$S_n := \inf\big\{t\geq s_n : \Psi^+(t)=2G_n\big\}.$$

Then it is not hard to check (see Figure 6.4) that the reflected random walk L'_n above is equal in law to the path defined at integer times $t\geq s_n$ by

$$L'_n(t) := \begin{cases} \tilde L'_n(t)+1 & (s_n\leq t\leq S_n-1),\\ \tilde L'_n(t)-1 & (S_n\leq t),\end{cases}$$

and then by linear interpolation for all $t\geq s_n$. We denote diffusively rescaled paths and functions by

$$(\hat R_{(n)},\tilde L_{(n)},\tilde L'_{(n)},L'_{(n)},\Psi^-_{(n)},\Psi^+_{(n)}) := \theta_{\varepsilon_n}(\hat R_n,\tilde L_n,\tilde L'_n,L'_n,\Psi^-_n,\Psi^+_n).$$

It is easy to see that as a consequence of (6.19), we have

$$\big(\tilde L'_{(n)},\Psi^-_{(n)},\Psi^+_{(n)}\big) := \Phi'(\tilde L_{(n)},\hat R_{(n)},2\varepsilon_n G_n).$$

Theorem 5.19 tells us that $(L'_{(n)},\hat R_{(n)})$ converges in law to $(\pi,\hat\pi)$, the forward and dual path from Theorem 6.12. Since L'_n and $\tilde L'_n$ differ at most by one, the last statement remains true if we replace $L'_{(n)}$ by $\tilde L'_{(n)}$. On the other hand, the triple $(\tilde L_{(n)},\hat R_{(n)},2\varepsilon_n G_n)$ converges in law to $(L,\hat\pi,T)$, where $\hat\pi$ is the same as before, L is an independent Brownian motion with drift -1, started from (x,s), and T is an independent exponentially distributed random variable with mean 2. Using Skorohod's representation theorem, we can couple our random variables such that the convergence is almost sure. The claim of the theorem then follows from Lemma 6.13. ∎

Remark Since a Brownian web is almost surely uniquely determined by its dual, instead of characterising the law of a left-right Brownian web using the joint law of a left forward path and a right forward path as we did in

Section 5.6, one could also try characterisation based on the joint law of a left forward path and a right dual path. This approach has the advantage that we also have a good understanding of the *conditional* law of a left forward path given a right dual path. As we mentioned in Section 5.5, we are lacking a good description of the conditional law of a left forward path given a right forward path.

6.5 Relevant separation points

Let $\mathcal{N}$ be a Brownian net and let $(\mathcal{W}^{\mathrm{l}}, \mathcal{W}^{\mathrm{r}})$ be its associated left-right Brownian web. We say that a left path π^{l} and a right path π^{r} *separate* in a point $z = (x,t)$ if $\pi^{\mathrm{l}} \in \mathcal{W}^{\mathrm{l}}_{\mathrm{in}}(z)$, $\pi^{\mathrm{r}} \in \mathcal{W}^{\mathrm{r}}_{\mathrm{in}}(z)$, and there exists a $U > t$ such that $\pi^{\mathrm{l}}(u) < \pi^{\mathrm{r}}(u)$ for all $u \in (t,U)$. A *separation point* of the Brownian net $\mathcal{N}$ is any point $z \in \mathbb{R}^2$ for which there exist left and right paths that separate in z. Separation points of the dual Brownian net $\hat{\mathcal{N}}$ are defined analogously.

Separation points can be viewed as the continuous analogue of branching points in an arrow configuration, in the sense that at a separation point (x,t), a path in the Brownian net with starting time $S < t$ can choose whether to turn left or right. Such a choice may or may not have a big influence on how the path in the Brownian net continues. If the left and right path that separate at (x,t) meet again very soon after t, then whether one turns left or right at (x,t) does not make a big difference. This is the idea behind the following definition.

Let $S,U \in \overline{\mathbb{R}}$ satisfy $S < U$. By definition, an *(S,U)-relevant separation point* of the Brownian net $\mathcal{N}$ is a point $z = (x,t) \in \mathbb{R}^2$ with $S < t < U$ such that:

(i) there exist $\pi^{\mathrm{l}} \in \mathcal{W}^{\mathrm{l}}_{\mathrm{in}}(z)$ and $\pi^{\mathrm{r}} \in \mathcal{W}^{\mathrm{r}}_{\mathrm{in}}(z)$ such that $\pi^{\mathrm{l}} < \pi^{\mathrm{r}}$ on (t,U),

(ii) there exists a path $\pi \in \mathcal{N}(\mathbb{R} \times \{S\})$ such that $\pi(t) = x$.

Clearly, each separation point is (S,U)-relevant for some S and U, since for the path π in point (ii) we can take either of the paths π^{l} and π^{r}.

By the remark below Lemma 4.7, if $\mathcal{D} \subset \mathbb{R}^2$ is a countable dense set, then for each separation point z, there exists skeletal paths $\pi^{\mathrm{l}} \in \mathcal{W}^{\mathrm{l}}(\mathcal{D})$ and $\pi^{\mathrm{r}} \in \mathcal{W}^{\mathrm{r}}(\mathcal{D})$ that separate in z. Since two skeletal paths separate in at most countably many points, it follows that the set of all separation points of a Brownian net is a.s. countable. In fact, a much stronger statement holds. Recall that a set $R \subset \mathbb{R}^2$ is called *locally finite* if $R \cap K$ is a finite set for all compact $K \subset \mathbb{R}^2$.

Proposition 6.14 (Set of relevant separation points) *For each $S, U \in \overline{\mathbb{R}}$ with $S < U$, the set of (S,U)-relevant separation points of the Brownian net is a locally finite subset of $\mathbb{R}^2$.*

The proof of Proposition 6.14 needs some preparations. We start with a simple observation.

Lemma 6.15 (Deterministic times) *For each deterministic $t \in \mathbb{R}$, the set $\mathbb{R} \times \{t\}$ almost surely does not contain any separation points of the Brownian net.*

Proof Let $z = (x,t) \in \mathbb{R}^2$ and assume that $\pi^{\rm l} \in \mathcal{W}^{\rm l}_{\rm in}(z)$ and $\pi^{\rm r} \in \mathcal{W}^{\rm r}_{\rm in}(z)$ satisfy $\pi^{\rm l} < \pi^{\rm r}$ on (t,U) for some $U > t$. Then by Lemma 6.1 applied to the dual Brownian net $\hat{\mathcal{N}}$, there exists a $\hat\pi \in \hat{\mathcal{N}}(\overline{\mathbb{R}} \times \{U\})$ such that $\hat\pi(t) = x$. Thus $\mathcal{W}^{\rm l}_{\rm in}(z) \neq \emptyset$ and $\hat{\mathcal{N}}_{\rm in}(z) \neq \emptyset$, which by Lemma 6.7 a.s. does not occur if t is deterministic. ∎

We now start preparing for the proof of Proposition 6.14 in earnest. The proof will be based on a density argument, that uses a sort of "approximate relevant separation points. Fix deterministic $S, U \in \overline{\mathbb{R}}$ with $S < U$, and let $(Y_{S,t})_{t \geq S}$ and $(\hat{Y}_{U,t})_{t \leq U}$ be defined by

$$\begin{aligned} Y_{S,t} &:= \big\{\pi(t) : \pi \in \mathcal{N},\ \sigma_\pi = S\big\} \qquad (t \geq S),\\ \hat{Y}_{U,t} &:= \big\{\hat\pi(t) : \hat\pi \in \hat{\mathcal{N}},\ \tau_{\hat\pi} = U\big\} \qquad (t \leq U). \end{aligned}$$

Let $s, u \in \mathbb{R}$ be deterministic times such that $S < s < u < U$. By Lemma 6.8, almost surely, for each $v \in Y_{S,s}$ there exist $\pi^{\rm l} \in \mathcal{W}^{\rm l}_{\rm in}(v,s)$ and $\pi^{\rm r} \in \mathcal{W}^{\rm r}_{\rm in}(v,s)$. Using notation first introduced in Exercise 4.30, we let

$$\pi^{\rm l\uparrow}_{(v,s)} \quad \text{and} \quad \pi^{\rm r\uparrow}_{(v,s)}$$

denote the unique paths in $\mathcal{W}^{\rm l}(v,s)$ and $\mathcal{W}^{\rm r}(v,s)$ that are continuations of each path in $\mathcal{W}^{\rm l}_{\rm in}(v,s)$ or in $\mathcal{W}^{\rm r}_{\rm in}(v,s)$, respectively. Similarly, for each $y \in \hat{Y}_{U,u}$, we let

$$\hat\pi^{\rm l\downarrow}_{(y,u)} \quad \text{and} \quad \hat\pi^{\rm r\downarrow}_{(y,u)}$$

denote the unique paths in $\hat{\mathcal{W}}^{\rm l}(y,u)$ and $\hat{\mathcal{W}}^{\rm r}(y,u)$ that are continuations of each path in $\hat{\mathcal{W}}^{\rm l}_{\rm in}(y,u)$ or in $\hat{\mathcal{W}}^{\rm r}_{\rm in}(y,u)$, respectively. We claim that almost surely, for all $v \in Y_{S,s}$ and $y \in \hat{Y}_{U,u}$, the following statements are equivalent:

$$\text{(i) } \pi^{\rm l\uparrow}_{(v,s)} < y < \pi^{\rm r\uparrow}_{(v,s)} \quad \text{and} \quad \text{(ii) } \hat\pi^{\rm r\downarrow}_{(y,u)} < v < \hat\pi^{\rm l\downarrow}_{(y,u)}.$$

By the symmetry between the forward and dual Brownian net, it suffices to prove the implication (i)⇒(ii), and by the symmetry between left and right it suffices to prove that $y < \pi^{\mathrm{r}\uparrow}_{(v,s)}$ implies $\hat\pi^{\mathrm{r}\downarrow}_{(y,u)} < v$. But this follows from the fact that forward and dual right paths do not cross, and (by Theorem 4.17), points with both an incoming forward and dual right path do not occur at deterministic times.

The considerations above motivate us to define, for each deterministic $S, U \in \overline{\mathbb{R}}$ and $s, u \in \mathbb{R}$ with $S < s < u < U$,

$$\begin{aligned} Q_{S,U}(s,u) &:= \big\{(v,y) : v \in Y_{S,s},\ y \in \hat Y_{U,u},\ \pi^{\mathrm{l}\uparrow}_{(v,s)} < y < \pi^{\mathrm{r}\uparrow}_{(v,s)}\big\}, \\ &= \big\{(v,y) : v \in Y_{S,s},\ y \in \hat Y_{U,u},\ \hat\pi^{\mathrm{r}\downarrow}_{(y,u)} < v < \hat\pi^{\mathrm{l}\downarrow}_{(y,u)}\big\}. \end{aligned}$$

When s and u are close together, we can view elements of $Q_{S,U}(s,u)$ as "approximate (S,U)-relevant separation points". This idea will be made more precise in the proof of Proposition 6.14. We first prove the following lemma.

Lemma 6.16 (Approximate separation points) *Let $S, U \in \overline{\mathbb{R}}$ and $s, u \in \mathbb{R}$ satisfy $S < s < u < U$, and let $a, b \in \mathbb{R}$ satisfy $a < b$. Then*

$$\begin{aligned} &\mathbb{E}\big[\big|\{(v,y) \in Q_{S,U}(s,u) : v \in [a,b]\}\big|\big] \\ &\qquad = 2(u-s)(b-a)\Psi(s-S)\Psi(U-u). \end{aligned} \tag{6.20}$$

where $\Psi(t)$ denotes the density of the branching-coalescing point set, defined in (6.9), and we use the convention $\Psi(\infty) := 2$.

Proof By Lemma 6.6, the sets $Y_{S,s}$ and $\hat Y_{U,u}$ are independent of each other and of the restriction of $\mathcal{N}$ to the time interval $[s,u]$. By Propositions 6.5 and 6.9,

$$\mathbb{E}\big[\big|Y_{S,s} \cap [a,b]\big|\big] = (b-a)\Psi(s-S).$$

Now if we condition on $Y_{S,s}$, then under the conditional law, for each $v \in Y_{S,s}$, the sets $\mathcal{W}^{\mathrm{l}}(v,s)$ and $\mathcal{W}^{\mathrm{r}}(v,s)$ almost surely contain unique paths $\pi^{\mathrm{l}}_{(v,s)}$ and $\pi^{\mathrm{r}}_{(v,s)}$, and these are distributed as the solution to left-right equation (5.16) started from (v,v). In particular, individually, they are just Brownian motions with drift -1 and $+1$, so

$$\mathbb{E}\big[\pi^{\mathrm{r}}_{(v,s)}(u) - \pi^{\mathrm{l}}_{(v,s)}(u)\,\big|\,Y_{S,s}\big] = 2(u-s).$$

Finally, if we condition both on $Y_{S,s}$ and the restriction of $\mathcal{N}$ to the time interval $[a,b]$, then for each $v \in Y_{S,s}$, under the conditional law, the random variable $\big|\hat Y_{U,u} \cap [\pi^{\mathrm{l}}_{(v,s)}(u), \pi^{\mathrm{r}}_{(v,s)}(u)]\big|$ has expectation

$$\big(\pi^{\mathrm{r}}_{(v,s)}(u) - \pi^{\mathrm{l}}_{(v,s)}(u)\big)\Psi(U-u).$$

Putting everything together, we arrive at (6.20). ∎

Proof of Proposition 6.14 We will prove the claim under the additional assumption that $S,U \in \mathbb{R}$. The more general claim then follows since for $S' \leq S < U \leq U'$ we have that $\{(x,t) \in R_{S',U'} : S < t < U\}$ is a subset of $R_{S,U}$. For each $n \geq 1$, we choose $S = s_0^n < \cdots < s_n^n = U$ in such a way that

$$\sup\big\{|s_{i+1}^n - s_i^n| : 0 \leq i \leq n-1\big\} \underset{n\to\infty}{\longrightarrow} 0, \tag{6.21}$$

and we define

$$R_{S,U}^n := \big\{(v,y,i) : 0 \leq i \leq n-1,\ (v,y) \in Q_{S,U}(s_i^n, s_{i+1}^n)\big\}.$$

We define a relation $\sim$ between $R_{S,U}$ and $R_{S,U}^n$ as follows. By definition, elements $(x,t) \in R_{S,U}$ and $(v,y,i) \in R_{S,U}^n$ satisfy $(x,t) \sim (v,y,i)$ if $s_i^n < t < s_{i+1}^n$ and there exist $\pi \in \mathcal{N}(\mathbb{R}\times\{S\})$ and $\hat\pi \in \hat{\mathcal{N}}(\mathbb{R}\times\{U\})$ such that $\pi(s_i^n) = v$, $\pi(t) = x = \hat\pi(t)$, and $y = \hat\pi(s_{i+1}^n)$. We claim that for each n,

$$\forall (x,t) \in R_{S,U}\ \exists (v,y,i) \in R_{S,U}^n \text{ s.t. } (x,t) \sim (v,y,i). \tag{6.22}$$

To prove this, fix $z = (x,t) \in R_{S,U}$ and choose $\pi^{\rm l}, \pi^{\rm r}$, and π are as in points (i) and (ii) of the definition of an (S,U)-relevant separation point. By Lemma 6.1 applied to the dual Brownian net $\hat{\mathcal{N}}$, there exists a dual path $\hat\pi \in \hat{\mathcal{N}}$ with $\tau_{\hat\pi} = U$ such that $\pi^{\rm l} \leq \hat\pi \leq \pi^{\rm r}$ on $[t,U]$. We fix such a dual path $\hat\pi$. By Lemma 6.15, at deterministic times there a.s. are no separation points. Therefore, for each n, there is a unique $0 \leq i \leq n-1$ such that $s_i^n < t < s_{i+1}^n$. Set $s := s_i^n$ and $u := s_{i+1}^n$. We claim that setting $v := \pi(s)$ and $y := \hat\pi(u)$, we have that $(v,y) \in Q_{S,U}(s,u)$ and hence $(v,y,i) \in R_{S,U}^n$. Clearly $v \in Y_{S,s}$ and $y \in \hat Y_{U,u}$, so it remains to prove that

$$\pi_{(v,s)}^{{\rm l}\uparrow}(u) < y < \pi_{(v,s)}^{{\rm r}\uparrow}(s).$$

To see this, we observe that $\pi_{(v,s)}^{{\rm l}\uparrow} \leq \pi$ on $[s,\infty)$ and hence $\pi_{(v,s)}^{{\rm l}\uparrow}(t) \leq \pi(t) = \pi^{\rm l}(t)$. Since left paths cannot cross each other and coalesce if they meet at any time after their starting times (Lemma 4.5), using the fact that $\pi^{\rm l} \in \mathcal{W}_{\rm in}^{\rm l}(x,t)$, we see that $\pi_{(v,s)}^{{\rm l}\uparrow}(u) \leq \pi^{\rm l}(u)$. We already saw that $\pi^{\rm l}(u) \leq \hat\pi(u) = y$, and since u is deterministic, this inequality must be strict by Lemma 6.7. This proves that $\pi_{(v,s)}^{{\rm l}\uparrow}(u) < y$. A similar argument applies to $\pi_{(v,s)}^{{\rm r}\uparrow}$. This completes the proof that $(v,y) \in Q_{S,U}(s_i^n, s_{i+1}^n)$. The fact that $(x,t) \sim (v,y,i)$ is immediate from the definition of (v,y,i), so (6.22) is proved.

Fix $a,b,s,u \in \mathbb{R}$ with $a < b$ and $S < s < u < U$. Let O denote the open set

$$O := (a,b) \times (s,u).$$

We claim that

$$\big|R_{S,U} \cap O\big| \le \liminf_{n\to\infty} \big|\big\{(v,y,i) \in R^n_{S,U} : (v, s^n_i) \in O\big\}\big|, \qquad (6.23)$$

where we allow for the case that both sides of the inequality are infinite. To see this, let Δ be any finite subset of $R_{S,U} \cap O$. By (6.22), for each $(x,t) \in \Delta$ and for each n, we can choose $(v_n, y_n, i_n) \in R^n_{S,U}$ such that $(x,t) \sim (v_n, y_n, i_n)$. Recall that the latter implies that $s^n_i < t < s^n_{i+1}$ and there exists a $\pi \in \mathcal{N}$ with $\pi(s^n_i) = v_n$ and $\pi(t) = x$. Using (6.21) and the equicontinuity of $\mathcal{N}$, this implies that $(v_n, i_n) \to (x,t)$ as $n \to \infty$. Since the set O is open, it follows that $(v_n, i_n) \in O$ for all n large enough. Thus,

$$\liminf_{n\to\infty} \big|\big\{(v,y,i) \in R^n_{S,U} : (v, s^n_i) \in O\big\}\big| \ge |\Delta|$$

for each finite subset Δ of $R_{S,U} \cap O$, which implies (6.23).

By Fatou's lemma, (6.23) implies that

$$\mathbb{E}\big[|R_{S,U} \cap O|\big] \le \liminf_{n\to\infty} \mathbb{E}\big[\big|\big\{(v,y,i) \in R^n_{S,U} : (v, s^n_i) \in O\big\}\big|\big].$$

Using Lemma 6.16 and Riemann sum approximation of the integral, we see that

$$\lim_{n\to\infty} \mathbb{E}\big[\big|\big\{(v,y,i) \in R^n_{S,U} : (v, s^n_i) \in O\big\}\big|\big] = 2(b-a)\int_s^u \Psi(t-S)\Psi(U-t)\mathrm{d}t.$$

This proves

$$\mathbb{E}\big[\big|R_{S,U} \cap \big((a,b)\times(s,u)\big)\big|\big] \le 2(b-a)\int_s^u \Psi(t-S)\Psi(U-t)\mathrm{d}t \qquad (6.24)$$

for all $a < b$ and $S < s < u < U$. By monotone convergence, we can relax the conditions $S < s < u < U$ to $S \le s < u \le U$. We observe that by (6.9) we have that $\Psi(t) \sim ct^{-1/2}$ as $t \to 0$ for some $c > 0$, so

$$\int_S^U \Psi(t-S)\Psi(U-t)\mathrm{d}t < \infty.$$

Inserting this into our previous formula, we see that $R_{S,U}$ is a locally finite subset of $\mathbb{R}^2$, as claimed. ∎

The following theorem says that the inequality in (6.24) is in fact an equality.

Theorem 6.17 (Density of relevant separation points) *Assume that $S,s,u,U\in\overline{\mathbb{R}}$ and $a,b\in\mathbb{R}$ satisfy $S\leq s<u\leq U$ and $a<b$. Let $\mathcal{N}$ be a Brownian net and let $R_{S,U}$ be its set of (S,U)-relevant separation points. Then*

$$\mathbb{E}\big[\big|R_{S,U}\cap\big([a,b]\times[s,u]\big)\big|\big]=2(b-a)\int_s^u\Psi(t-S)\Psi(U-t)\mathrm{d}t,\qquad(6.25)$$

where $\Psi(t)$ denotes the density of the branching-coalescing point set, defined in (6.9), and $\Psi(\infty):=2$.

Proof (crude sketch) The first step of the proof is to show that

$$\forall(v,y,i)\in R^n_{S,U}\ \exists(x,t)\in R_{S,U}\ \text{s.t.}\ (x,t)\sim(v,y,i).\qquad(6.26)$$

Note that formulas (6.22) and (6.26) together say that the relation $\sim$ is a correspondence between $R_{S,U}$ and $R^n_{S,U}$ in the sense defined in Section 2.5.

To prove (6.26), fix deterministic $S<s<u<U$ and let $(v,y)\in Q_{S,U}(s,u)$. By the definition of $Q_{S,U}(s,u)$, there exist $\pi\in\mathcal{N}(\mathbb{R}\times\{S\})$ and $\hat\pi\in\hat{\mathcal{N}}(\mathbb{R}\times\{U\})$ such that $\pi(s)=v$, $\hat\pi(u)=y$, and $\pi^{\mathrm{l}\uparrow}_{(v,s)}(u)<y<\pi^{\mathrm{r}\uparrow}_{(v,s)}(u)$. Since s is deterministic, Lemma 6.8 tells us that there exist $\pi^{\mathrm{r}}\in\mathcal{W}^{\mathrm{r}}_{\mathrm{in}}(v,s)$ such that $\pi^{\mathrm{l}}\leq\pi\leq\pi^{\mathrm{r}}$ on $[s,\infty)$. Since s is deterministic, Theorem 4.17 now tells us that (v,s) is of type $(1,1)$ both in $\mathcal{W}^{\mathrm{l}}$ and $\mathcal{W}^{\mathrm{r}}$, so we must have $\pi^{\mathrm{l}-}_{(v,s)}=\pi^{\mathrm{l}}$ and $\pi^{\mathrm{r}+}_{(v,s)}=\pi^{\mathrm{r}}$ on $[s,\infty)$. Since u is deterministic, Lemma 6.7 tells us that $\pi^{\mathrm{l}}(u)\neq\hat\pi(u)\neq\pi^{\mathrm{r}}(u)$ a.s., so we must have $\pi^{\mathrm{l}}(u)<\hat\pi(u)<\pi^{\mathrm{r}}(u)$. By Lemma 6.1, the existence of the path $\hat\pi$ implies that $\pi^{\mathrm{l}}<\pi^{\mathrm{r}}$ on $[u,U)$. Together with the fact that $\pi^{\mathrm{l}}(s)=\pi^{\mathrm{r}}(s)$, this allows us to define $t\in(s,u)$ by

$$t:=\sup\big\{t'\in(s,U):\pi^{\mathrm{l}}(t')=\pi^{\mathrm{r}}(t')\big\}.$$

Since $\pi^{\mathrm{l}}\leq\pi\leq\pi^{\mathrm{r}}$ on $[s,\infty)$ and $\pi^{\mathrm{l}}<\pi^{\mathrm{r}}$ on (t,U), we see that $z:=\big(\pi(t),t\big)$ is an (S,U)-relevant separation point. Applying this to $s=s^n_i$ and $u=s^n_{i+1}$, we obtain (6.26).

With formula (6.26) proved, we know that the relation $\sim$ is a correspondence between $R_{S,U}$ and $R^n_{S,U}$. The remainder of the proof is rather technical, so we only sketch the details. Let $O:=(a,b)\times(s,u)$ and let $R^n_{S,U}(O):=\{(v,y,i)\in R^n_{S,U}:(v,s^n_i)\in O\}$. The main idea is to show that for large enough n the relation $\sim$ is a bijection between $R_{S,U}\cap O$ and $R^n_{S,U}(O)$.

Since $\sim$ is a correspondence, for each $(x,t)\in R_{S,U}$, we can choose $(v_n,y_n,i_n)\in R^n_{S,U}$ such that $(x,t)\sim(v_n,y_n,i_n)$. Using the equicontinuity of the Brownian net, one can show that this implies that $(v_n,s^n_i)\to(x,t)$. Using this and the fact that O is open and $R_{S,U}$ is locally finite, it is not hard to show that for large enough n, the relation $\sim$ is a correspondence between

$R_{S,U} \cap O$ and $R^n_{S,U}(O)$, and each element of $R^n_{S,U}(O)$ corresponds to at most one element of $R_{S,U} \cap O$.

To see that conversely, for large enough n, each element of $R_{S,U} \cap O$ corresponds to at most one element of $R^n_{S,U}(O)$, one can use that the sets $Q_{S,U}(s,u)$ are unlikely to contain two points (v,y) and (v',y') for which v and v' lie very close to each other. In the proof of [SSS09, Prop. 2.9], this is done by showing that one can change the definition of $Q_{S,U}(s,u)$ so that it does not contain such points without changing the density in (6.20) by much. This part of the proof is a bit technical.

After showing that for large enough n the relation $\sim$ is a bijection between $R_{S,U} \cap O$ and $R^n_{S,U}(O)$, one obtains in particular that

$$\big|R_{S,U} \cap O\big| \geq \limsup_{n\to\infty} \big|R^n_{S,U}(O)\big|.$$

The idea is now to take expectations on both sides and apply Lemma 6.16 to obtain a matching upper bound for (6.24). To justify interchanging the limit superior and the expectation, one still needs to do some technical work. Since we will never actually use Theorem 6.17 but be satisfied with the weaker Proposition 6.14, we refer to [SSS09, Prop. 2.9] for the details. ∎

6.6 Structure of separation points

We recall from Section 4.4 that points $z \in \mathbb{R}^2$ are distinguished into different types according to the local structure of the Brownian web at these points. For a left-right Brownian web $(\mathcal{W}^{\rm l}, \mathcal{W}^{\rm r})$, one can wonder what combinations of types in $\mathcal{W}^{\rm l}$ and $\mathcal{W}^{\rm r}$ are possible. This question has completely been answered in [SSS09]. The main result of that paper is summarised in Figure 6.5.

In this section, we will only be interested in separation points. The following proposition identifies their position in the table Figure 6.5. Indeed, they are precisely the points that are these are of type $(1,2)_{\rm l}$ in $\mathcal{W}^{\rm l}$ and of type $(1,2)_{\rm r}$ in $\mathcal{W}^{\rm r}$.

Proposition 6.18 (Structure of separation points) *Let $\mathcal{N}$ be a Brownian net, let $(\mathcal{W}^{\rm l}, \mathcal{W}^{\rm r})$ be its associated left-right Brownian web, let $(\hat{\mathcal{W}}^{\rm l}, \hat{\mathcal{W}}^{\rm r})$ be their associated dual Brownian webs, and let $\hat{\mathcal{N}}$ be the associated dual Brownian net. Then almost surely, for all $(x,t) \in \mathbb{R}^2$, the following statements are equivalent.*

(i) *(x,t) is a separation point of the Brownian net $\mathcal{N}$,*

		$\mathcal{W}^{\rm r}$						
		$(0,1)$	$(1,1)$	$(0,2)$	$(1,2)_{\rm l}$	$(1,2)_{\rm r}$	$(2,1)$	$(0,3)$
	$(0,1)$	X	X	X	X			
	$(1,1)$	X	X		X			
	$(0,2)$	X		X	X			
$\mathcal{W}^{\rm l}$	$(1,2)_{\rm l}$				X	X		
	$(1,2)_{\rm r}$	X	X	X		X		
	$(2,1)$						X	
	$(0,3)$							X

Figure 6.5: Possible combinations of types of points in $\mathcal{W}^{\rm l}$ and $\mathcal{W}^{\rm r}$.

(ii) (x,t) *is a separation point of the dual Brownian net* $\hat{\mathcal{N}}$*,*

(iii) (x,t) *is of type* $(1,2)_{\rm l}$ *in* $\mathcal{W}^{\rm l}$ *and of type* $(1,2)_{\rm r}$ *in* $\mathcal{W}^{\rm r}$*.*

(iv) *There exist* $\pi^{\rm r}\in\mathcal{W}^{\rm r}$ *and* $\hat\pi^{\rm l}\in\hat{\mathcal{W}}^{\rm l}$ *with* $\sigma_{\pi^{\rm r}}<t<\tau_{\hat\pi^{\rm l}}$ *such that* $\pi^{\rm r}\leq\hat\pi^{\rm l}$ *on* $[\sigma_{\pi^{\rm r}},t]$ *and* $\hat\pi^{\rm l}\leq\pi^{\rm r}$ *on* $[t,\tau_{\hat\pi^{\rm l}}]$*.*

Proof We start by proving the implication (iv)⇒(i). Let (x,t) be as in point (iv). Choose $S,U\in\mathbb{Q}$ such that $\sigma_{\pi^{\rm r}}<S<t<U\tau_{\hat\pi^{\rm l}}]$. By Proposition 6.14, for each $S,U\in\mathbb{Q}$, the set of (S,U)-relevant separation points is locally finite, so there must exist an $s\in(S,t)$ such that the set $\{(\pi(s'),s'):s<s'<t\}$ contains no (S,U)-relevant separation points.

Choose $u\in(s,t)\cap\mathbb{Q}$. By Lemma 6.8, there exists a path $\pi^{\rm l}_n\in\mathcal{W}^{\rm l}_{\rm in}(\pi^{\rm r}(u),u)$ such that $\pi^{\rm l}\leq\pi^{\rm r}$ on $[u,\infty)$. Let

$$\tau:=\sup\big\{t'\in[u,U]:\pi^{\rm l}(t)=\pi^{\rm r}(t)\big\}.$$

Since left paths and dual left paths do not cross, we must have $\pi^{\rm l}\leq\hat\pi^{\rm l}$. By Lemma 6.7, at deterministic times, there are no points with both and incoming forward and dual path, so we must have $\pi^{\rm l}(U)<\hat\pi^{\rm l}(U)$. It follows that $\pi^{\rm l}<\pi^{\rm r}$ on $(t,U]$, since otherwise they would create a wedge that is entered by $\hat\pi^{\rm l}$. This shows that $\tau\leq t$. On the other hand, we cannot have $\tau<t$ since that would contradict our assumption that $\{(\pi(s'),s'):s<s'<t\}$ contains no (S,U)-relevant separation points. We conclude that $\pi^{\rm l}$ and $\pi^{\rm r}$ separate in (x,t).

We next prove the implication (i)⇒(iii). Assume that $z=(x,t)$ is a separation point of the Brownian net $\mathcal{N}$. Then, by the definition of a separation point, there exist $\pi^{\rm l}\in\mathcal{W}^{\rm l}_{\rm in}(z)$ and $\pi^{\rm r}\in\mathcal{W}^{\rm r}_{\rm in}(z)$ such that $\pi^{\rm l}<\pi^{\rm r}$ on $(s,U]$

for some $U > s$. There must exist $S \in \mathbb{Q}$ satisfying $\sigma_{\pi^{\rm l}} < S < s$. By Proposition 6.14, for each such S, the set of (S, U)-relevant separation points is locally finite, so there must exist a $u > s$ such that the set $\{(\pi(t), t) : s < t < u\}$ contains no (S, U)-relevant separation points.

Choose $t_n \in \mathbb{Q} \cap (t, u)$ with $t_n \to t$. By Lemma 6.8, for each n, there exists a path $\pi^{\rm r}_n \in \mathcal{W}^{\rm r}_{\rm in}(\pi^{\rm l}(t_n), t_n)$ such that $\pi^{\rm l} \leq \pi^{\rm r}_n$ on $[t_n, \infty)$. For each n, there must be some time $u_n \in [u, U]$ such that $\pi^{\rm l}(u_n) = \pi^{\rm r}_n(u_n)$, for if that would not be the case, then setting

$$\tau := \sup\big\{t \in [s, u) : \pi^{\rm l}(t) = \pi^{\rm r}_n(t)\big\},$$

we would have that $(\pi(\tau), \tau)$ is an (S, U)-relevant separation point with $s < \tau < u$, which contradicts our choice of u. Letting $n \to \infty$, using the compactness of $\mathcal{W}^{\rm r}$ to select a convergent subsequence, we see that there must exist a path $\tilde\pi^{\rm r} \in \mathcal{W}^{\rm r}_{\rm in}(z)$ such that $\pi^{\rm l} \leq \tilde\pi^{\rm r}$ on $[s, \infty)$ and $\pi^{\rm l}(t) = \tilde\pi^{\rm r}(t)$ for some $t \in [u, U]$. In particular, this implies that $\tilde\pi^{\rm r} < \pi^{\rm r}$ on $(s, u]$, since otherwise $\tilde\pi^{\rm r}$ would by Lemma 4.5 have to coalesce with $\pi^{\rm r}$ which would contradict the fact that $\pi^{\rm l} < \pi^{\rm r}$ on $(s, U]$.

This shows that $\mathcal{W}^{\rm r}(z)$ contains apart from the path $\pi^{\rm r} \in \mathcal{W}^{\rm r}_{\rm in}(z)$ another path $\tilde\pi^{\rm r}$ that lies on the left of $\pi^{\rm r}$. By Theorem 4.17, it follows that z is of type $(1, 2)_{\rm r}$ in $\mathcal{W}^{\rm r}$. By symmetry, the same argument shows that z is of type $(1, 2)_{\rm l}$ in $\mathcal{W}^{\rm l}$. By Lemma 4.20 and the fact that forward paths in a Brownian web cannot cross dual paths, it follows that z is of type $(1, 2)_{\rm l}$ in $\hat{\mathcal{W}}^{\rm l}$ and of type $(1, 2)_{\rm r}$ in $\hat{\mathcal{W}}^{\rm r}$.

We next prove the implication (iii)⇒(iv). Assume that $z = (x, t)$ is of type $(1, 2)_{\rm l}$ in $\mathcal{W}^{\rm l}$ and of type $(1, 2)_{\rm r}$ in $\mathcal{W}^{\rm r}$. Let $\pi^{\rm r} \in \mathcal{W}^{\rm r}_{\rm in}(z)$ and $\hat\pi^{\rm l} \in \hat{\mathcal{W}}^{\rm l}_{\rm in}(z)$. Since right forward paths cannot cross left dual paths from right to left, one of the following three statements must be true.

I There exists a $u \in (t, \tau_{\hat\pi^{\rm l}}]$ such that $\pi^{\rm r} \leq \hat\pi^{\rm l}$ on $[\sigma_{\pi^{\rm r}}, u]$.

II There exists a $s \in [\sigma_{\pi^{\rm r}}, t)$ such that $\hat\pi^{\rm l} \leq \pi^{\rm r}$ on $[s, \tau_{\hat\pi^{\rm l}}]$.

III $\pi^{\rm r} \leq \hat\pi^{\rm l}$ on $[\sigma_{\pi^{\rm r}}, t]$ and $\hat\pi^{\rm l} \leq \pi^{\rm r}$ on $[t, \tau_{\hat\pi^{\rm l}}]$.

We will rule out I and II, which leaves III as the only remaining possibility. By Lemma 6.7, forward and dual paths in the Brownian net spend zero Lebesgue time together, so if I holds, then there must exist a dense set of times $t' \in [\sigma_{\pi^{\rm r}}, u]$ such that $\pi^{\rm r}(t') < \hat\pi^{\rm l}(t')$. This allows us to choose $(x', t') \in \mathbb{R}^2$ with $t < t' < u$ and $\pi^{\rm r}(t') < x < \hat\pi^{\rm l}(t')$. Let $\hat\pi^{\rm r} \in \hat{\mathcal{W}}^{\rm r}(x', t')$. Then $\hat\pi^{\rm r}$ is contained between $\pi^{\rm r}(t')$ and $\hat\pi^{\rm l}(t')$ and hence must pass through (x, t). However, this implies that (x, t) is of type $(1, 2)_{\rm l}$ in $\mathcal{W}^{\rm r}$, contradicting our assumption.

To also rule out II, we use arguments similar to those we have already seen. Since dual right paths cannot cross forward left paths from left to right, using the compactness of $\hat{\mathcal{W}}^{\mathrm{r}}$, we see that there must exist a path $\hat{\pi}^{\mathrm{r}}_1 \in \hat{\mathcal{W}}^{\mathrm{r}}(x,t)$ such that $\hat{\pi}^{\mathrm{r}}_1 \leq \pi^{\mathrm{l}}$ on $[\sigma_{\pi^{\mathrm{l}}}, t]$. Similarly, there must exist a path $\hat{\pi}^{\mathrm{r}}_2 \in \hat{\mathcal{W}}^{\mathrm{r}}(x,t)$ such that $\pi^{\mathrm{l}} \leq \hat{\pi}^{\mathrm{r}}_2$ on $[\sigma_{\pi^{\mathrm{r}}}, t]$. Now if II holds, then using the local finiteness of relevant separation points, we see that there must exist a third path $\hat{\pi}^{\mathrm{r}}_3 \in \hat{\mathcal{W}}^{\mathrm{r}}(x,t)$ that does not separate from the path $\hat{\pi}^{\mathrm{l}} \in \hat{\mathcal{W}}^{\mathrm{l}}_{\mathrm{in}}(x,t)$ on some interval $[s,t]$ of positive length. By Theorem 4.17, the existence of such a third path contradicts the existence of an incoming right path at (x,t).

We have now proved (iv)⇒(i)⇒(iii)⇒(iv), showing that all these conditions are equivalent. By Theorem 4.17, condition (iii) is equivalent to (x,t) being of type $(1,2)_{\mathrm{l}}$ in $\hat{\mathcal{W}}^{\mathrm{l}}$ and of type $(1,2)_{\mathrm{r}}$ in $\hat{\mathcal{W}}^{\mathrm{r}}$, which by what we have already proved is equivalent to (ii). ∎

Bibliography

[Arr79] R. Arratia. Coalescing Brownian motions on the line. Ph.D. Thesis, University of Wisconsin, Madison, 1979.

[Arr81] R. Arratia. Coalescing Brownian motions and the voter model on $\mathbb{Z}$. Unpublished partial manuscript. Available from rarratia@math.usc.edu.

[BBI01] D. Burago, Y. Burago, and S. Ivanov. *A Course in Metric Geometry.* Graduate Studies in Mathematics 33. AMS, Providence, 2001.

[Bil99] P. Billingsley. *Convergence of Probability Measures. 2nd ed.* Wiley, New York, 1999.

[BMSV06] S. Belhaouari, T. Mountford, R. Sun, and G. Valle. Convergence results and sharp estimates for the voter model interfaces. *Electron. J. Prob.* 11 (2006), 768–801.

[Bou58] N. Bourbaki. *Éléments de Mathématique. VIII. Part. 1: Les Structures Fondamentales de l'Analyse. Livre III: Topologie Générale. Chap. 9: Utilisation des Nombres Réels en Topologie Générale. 2iéme éd.* Actualités Scientifiques et Industrielles 1045. Hermann & Cie, Paris, 1958.

[Cho69] G. Choquet. *Lectures on Analysis. Volume I. Integration and Topological Vector Spaces.* Benjamin, London, 1969.

[Dud02] R.M. Dudley. *Real Analysis and Probability.* Reprint of the 1989 edition. Camebridge University Press, Camebridge, 2002.

[EFS17] A. Etheridge, N. Freeman, and D. Straulino. The Brownian net and selection in the spatial Λ-Fleming-Viot process. *Electron. J. Probab.* 22. (2017), Paper No. 39, 1–36.

[EK86] S.N. Ethier and T.G. Kurtz. *Markov Processes; Characterization and Convergence.* John Wiley & Sons, New York, 1986.

[Eng89] R. Engelking. *General Topology.* Heldermann, Berlin, 1989.

[Eva97] S.N. Evans. Coalescing Markov labelled partitions and a continuous sites genetics model with infinitely many types. *Ann. Inst. H. Poincaré Probab. Statist.* 33(3) (1997), 339–358.

[FINS01] L.R.G. Fontes, M. Isopi, C.M. Newman, and D. Stein. Aging in 1D discrete spin models and equivalent systems. *Phys. Rev. Lett.* 87 (2001), 110201.

[FINR04] L.R.G. Fontes, M. Isopi, C.M. Newman and K. Ravishankar. The Brownian web: characterization and convergence. *Ann. Probab.* 32(4) (2004), 2857–2883.

[Gri79] D. Griffeath. *Additive and Cancellative Interacting Particle Systems.* Lecture Notes in Math. 724, Springer, Berlin, 1979.

[KS91] I. Karatzas and S.E. Shreve. *Brownian motion and stochastic calculus. 2nd ed.* Graduate Texts in Mathematics 113. Springer, New York, 1991.

[Lig99] T.M. Liggett. *Stochastic Interacting Systems: Contact, Voter and Exclusion Process.* Springer, Berlin, 1999.

[McK69] H.P. McKean. *Stochastic integrals.* Academic Press, New York, 1969.

[Mey76] P.A. Meyer. Un cours sur lest intégrales stochastiques. Pages 245–400 in: *Sém. de Probabilités X*, Lecture Notes in Math. 511. Springer, Berlin, 1976.

[NRS05] C.M. Newman, K. Ravishankar, and R. Sun. Convergence of coalescing nonsimple random walks to the Brownian web. *Electron. J. Prob.* 10 (2005), 21–60.

[NRS10] C.M. Newman, K. Ravishankar, and E. Schertzer. Marking $(1,2)$ points of the Brownian web and applications. *Ann. Inst. Henri Poincaré Probab. Statist.* 46 (2010), 537–574.

[NRS15] C.M. Newman, K. Ravishankar, and E. Schertzer. The Brownian net with killing. *Stochastic Processes Appl.* 125 (2015), 1148–1194.

[NRS17] C.M. Newman, K. Ravishankar, and E. Schertzer. Perturbations of the voter model in one-dimension. *Electron. J. Probab.* 22 (2017), paper no. 34, 1–42.

[NS22] N. Freeman and J.M. Swart. Skorohod's topologies on path space. Preprint, 2022, arXiv:2301.05637.

[RW87] L.C.G. Rogers and D. Williams. *Diffusions, Markoc processes, and martingales*, Vol. 2. Wiley, Chichester, 1987.

[SS08] R. Sun and J.M. Swart. The Brownian net. *Ann. Probab.* 36(3), 1153-1208, 2008.

[SS19] E. Schertzer and R. Sun. Perturbations of supercritical oriented percolation and sticky Brownian webs. Pages 241–261 in: V. Sidoravicius (ed.), *Sojourns in probability theory and statistical physics. II. Brownian web and percolation, a festschrift for Charles M. Newman,* Springer Proc. Math. Stat. 299. Springer, Singapore, 2019.

[SSS09] E. Schertzer, R. Sun, and J.M. Swart. Special points of the Brownian net. *Electron. J. Probab.* 14 (2009), paper No. 30, 805–864.

[SSS14] E. Schertzer, R. Sun, and J.M. Swart. Stochastic flows in the Brownian web and net. *Mem. Am. Math. Soc.* Vol. 227 (2014), Nr. 1065.

[SSS16] E. Schertzer, R. Sun, and J.M. Swart. The Brownian Web, the Brownian Net, and their Universality. Pages 270-368 in: P. Contucci and C. Giardinà (Eds.) *Advances in Disordered Systems, Random Processes and Some Applications*, Cambridge University Press, 2016.

[STW00] F. Soucaliuc, B. Tóth, W. Werner. Reflection and coalescence between one-dimensional Brownian paths. *Ann. Inst. Henri Poincaré Probab. Statist.* 36 (2000), 509–536.

[Swa22] J.M. Swart. *A Course in Interacting Particle Systems.* Lecture notes (2022), arXiv:1703.10007v4.

[TW98] B. Tóth and W. Werner. The true self-repelling motion. *Probab. Theory Related Fields* 111 (1998), 375–452.

[Oxt80] J.C. Oxtoby. *Measure and Category. Second Edition.* Springer, New York, 1980.

Index