

The Brownian web is a two-dimensional black noise

Tom Ellis* Ohad N. Feldheim†

March 2012

Abstract

The Brownian web is a random variable consisting of a Brownian motion starting from each space-time point on the plane. These are independent until they hit each other, at which point they coalesce. Tsirelson mentions this model in [10], along with planar percolation, in suggesting the existence of a two-dimensional black noise. A two-dimensional noise is, roughly speaking, a random object on the plane whose distribution is translation invariant and whose behavior on disjoint subsets is independent. Black means sensitive to the resampling of sets of arbitrarily small total area.

Tsirelson implicitly asks: “Is the Brownian web a two-dimensional black noise?”. We give a positive answer to this question, providing the second known example of such after the scaling limit of critical planar percolation.

1 Introduction

In this paper we study a stochastic object called the *Brownian web*. We research this object in the context of the theory of classical and non-classical noises, developed by Boris Tsirelson (see [11] for a survey). Our main result is that, in the terminology of this theory, the Brownian web is a two-dimensional *black* noise. Roughly speaking, the Brownian web is a random variable which

*School of Mathematics, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel. This work was supported in part by a Wyndham Deedes Memorial Travel Scholarship from The Anglo-Israel Association.

†School of Mathematics, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel. E-mail: ohad_f@netvision.net.il. Research supported by an ERC advanced grant.

assigns to every space-time point in $\mathbb{R} \times \mathbb{R}$ a standard Brownian motion starting at that point. The motions in each finite subcollection are independent until the first time that one hits another and from thereon those two coalesce, continuing together. This object was originally studied more than twenty-five years ago by Arratia [1], motivated by a study of the asymptotics of one-dimensional voter models, and later by Tóth and Werner [7], motivated by the problem of constructing continuum “self-repelling motions”, by Fontes, Isopi, Newman and Ravishankar [3], motivated by its relevance to “aging” in statistical physics of one-dimensional coarsening, and by Norris and Turner [5] regarding a scaling limit of a two-dimensional aggregation process. A rigorous notion of the Brownian web in our context can be found in [9] for the case of coalescing Brownian motions on a circle. The above also provide their own constructions of the Brownian web.

The Brownian web functions as an important example in the theory of classical and non-classical noises. In this theory a noise is a probability space equipped with a collection of sub- σ -fields indexed by the open rectangles (possibly infinite) of \mathbb{R}^d . The sub- σ -field associated to a rectangle is intended to represent the behavior of a stochastic object within that rectangle. The σ -fields must satisfy the following three properties:

- the σ -fields associated to disjoint rectangles of \mathbb{R}^d are independent,
- translations on \mathbb{R}^d act in a way that preserves the probability measure,
- the σ -field associated to a rectangle is generated by the two σ -fields associated to two smaller rectangles which partition it.

Two natural examples of noises are the Gaussian white noise and the Poisson noise. These noises are called classical, or white, meaning that resampling of a small portion of \mathbb{R}^d doesn’t change the observables of the process very much.

The foundational result of Tsirelson and Vershik [8] showed that there exist non-classical noises. Indeed they showed the existence of non-classical noises that are as far from white as could be, and these are called black. The defining property of a black noise is that all its observables are sensitive, i.e. for any particular observable, resampling a small scattered portion of the noise renders that observable nearly independent of its original value. (For a thorough discussion of black and white, classical and non-classical noises see [11]). Tsirelson showed in [10] (Theorem 7c2) that the Brownian web, when considered as a time-indexed random process, is one-dimensional black noise. We extend this result by showing:

Theorem 1.1. The Brownian web is a two-dimensional black noise.

This makes the Brownian web only the second known two-dimensional black noise after Schramm and Smirnov's recent result on the scaling limit of critical planar percolation [6].

Of the three properties required for a process to be a noise, the first two, i.e. translation invariance and independence on disjoint domains, hold trivially for the Brownian web. Furthermore, once we have shown that the Brownian web is a two-dimensional noise it will follow that it is a two-dimensional *black* noise, through a general argument.

The main difficulty in proving Theorem 1.1 is to show that the σ -field associated to any rectangle is generated by the two σ -fields associated to any two rectangles that partition it. The major milestone towards this result is to show the special case when the large rectangle is the whole plane, and the smaller rectangles that partition it are the upper and lower half-planes.

Theorem 1.2. In the Brownian web, the σ -field associated to the whole plane is generated by that associated to the upper half-plane and that associated to the lower half-plane.

The rest of the paper goes as follows: in Section 2 we define the Brownian web formally; in Section 3 we restate Theorem 1.2 as Theorem 3.1; we then reduce this theorem to a convergence result for an auxiliary process which we prove in Section 4. In Section 5 we extend Theorem 1.2 to hold for the σ -fields associated to horizontal strips as well; in Section 6 we extend further to all rectangles. In addition we define noises properly and conclude by proving Theorem 1.1. Section 7 is devoted to remarks, open problems and acknowledgements.

2 Definition of the Brownian web

The Brownian web is the continuum scaling limit of a system of independent-coalescing random walks (see [9]). Constructing the continuum version raises several technical difficulties addressed in [3],[5],[7]. Nonetheless, all the constructions share the following property, which we use as a definition:

Denote $\mathcal{S} = \{(s, t) \in \mathbb{R}^2 : s \leq t\}$. A Brownian web on a probability space Ω is a (jointly) measurable mapping $\phi : \Omega \times \mathcal{S} \times \mathbb{R} \rightarrow \mathbb{R}$, $(\omega, (s, t), x) \mapsto \phi_{st}(x)$ (suppressing ω in the notation) such that for every finite collection of starting points $(s_1, x_1), (s_2, x_2), \dots, (s_n, x_n)$, the collection of processes $\phi_{s_1 \cdot}(x_1), \phi_{s_2 \cdot}(x_2), \dots, \phi_{s_n \cdot}(x_n)$ forms a system of n independent-coalescing Brownian motions.

A system of n independent-coalescing Brownian motions is a finite collection of stochastic processes (X^1, X^2, \dots, X^n) such that each X^i starts at




Figure 1: Some trajectories of the Brownian web. A particular trajectory is marked.

some point x_i at some time s_i , and (X^1, X^2, \dots, X^n) are independent until the first time T at which $X^i(T) = X^j(T)$ for some $i \neq j$. From this time onwards $X^i(T)$ and $X^j(T)$ coalesce and continue with the rest of the X^k (for $k \neq i, j$) as a system of $n - 1$ independent-coalescing Brownian motions. Several trajectories of a Brownian web can be seen in Figure 1.

3 Recovering the web from half-planes

We introduce three σ -fields generated by the Brownian web, and use them to restate Theorem 1.2 as Theorem 3.1.

Write \mathcal{X} for the collection of trajectories comprising the whole web, that is $\{t \mapsto \phi_{st}(x) : (s, x) \in \mathbb{R}^2\}$, and \mathcal{F} for the σ -field generated by the whole web, i.e. generated by \mathcal{X} . We further introduce \mathcal{F}_+ and \mathcal{F}_- , the σ -fields generated by the web on the upper and lower half-planes respectively. Formally,

Definition. For any path f we write $\mathcal{R}_+(f)$ for f stopped at the first time it is outside the upper half-plane. $\mathcal{R}_+(t \mapsto \phi_{st}(x))$ is therefore the trajectory of the web ϕ started at the point (s, x) and stopped at the first time it is outside the upper half-plane (if (s, x) is outside the upper half-plane, f is stopped immediately). We define \mathcal{F}_+ to be the σ -field generated by the collection of paths $\{\mathcal{R}_+(X) : X \in \mathcal{X}\}$, and define \mathcal{F}_- analogously.

Note that \mathcal{F}_+ and \mathcal{F}_- are independent by the definition of a system of n independent-coalescing Brownian motions.

For σ -fields \mathcal{F}_a , \mathcal{F}_b , \mathcal{F}_c we write $\mathcal{F}_a = \mathcal{F}_b \otimes \mathcal{F}_c$ when \mathcal{F}_a is generated by \mathcal{F}_b and \mathcal{F}_c (up to sets of measure 0), and \mathcal{F}_b and \mathcal{F}_c are independent. We are now ready to restate Theorem 1.2.

Theorem 3.1. $\mathcal{F} = \mathcal{F}_- \otimes \mathcal{F}_+$.

The difficulty is to show that $\mathcal{F} \subseteq \mathcal{F}_- \otimes \mathcal{F}_+$, because the reverse inclusion follows directly from the definitions. To avoid having to state results as “for all $X \in \mathcal{X} \dots$ ”, here and for the rest of the paper we write X for an arbitrary element of \mathcal{X} . This we do purely for the sake of notational simplicity.

\mathcal{F} is generated by such processes, so our theorem reduces to the following lemma.

Lemma 3.2. X is $\mathcal{F}_- \otimes \mathcal{F}_+$ -measurable.

To prove this we must show that X can be recovered using trajectories starting in the upper half-plane which stop when they hit 0, and trajectories starting on the lower half-plane which stop when they hit 0.

To do so, we might have liked to recover X (starting on, say, the upper half-plane) by following it until it hits 0, then by following its continuation within the lower half-plane. However, this is impossible since when trajectories of Brownian motion cross 0 they do so infinitely often within a finite period of time.

To overcome this problem we introduce in Section 3.1 a process X^ϵ , for each $\epsilon > 0$, which is $\mathcal{F}_- \otimes \mathcal{F}_+$ -measurable. In Section 4 we show that X^ϵ is an approximation of X in the sense that X^ϵ converges to X (in probability) as $\epsilon \rightarrow 0$. This will imply that X itself is $\mathcal{F}_- \otimes \mathcal{F}_+$ -measurable.

3.1 The perturbed process

From our arbitrary X we now construct X^ϵ , a “perturbed” version of X , which depends also on ψ_t , some Brownian motion independent of ϕ (measurable with respect to \mathcal{G} , say, where \mathcal{G} is independent of \mathcal{F}).

In the definition of X^ϵ we give the word “follows” two distinct meanings. We say X^ϵ follows ϕ on $[s, u]$ if $X_t^\epsilon = \phi_{st}(X_s^\epsilon)$ for $t \in [s, u]$. That is if the trajectory of X^ϵ follows the trajectory of the web starting from point X_s^ϵ at time s and up to time u . We say X^ϵ follows ψ on $[s, u]$ if $X_t^\epsilon = X_s^\epsilon + \psi_t - \psi_s$ for $t \in [s, u]$.

Definition 3.3. The perturbed process X^ϵ starts at the same time and position as X and alternates between two states. In state S_ϕ it follows ϕ while, in state S_ψ it follows ψ . The process starts in state S_ϕ and the transition from state S_ϕ to state S_ψ occurs when X^ϵ hits 0, while the transition from




Figure 2: A sample of a web trajectory and the corresponding perturbed processes. The top left image depicts the perturbed process, with state S_ψ in bold. The top right image depicts the original web trajectory. The center image illustrates both processes together, showing the segments where they coalesce.

state S_ψ to state S_ϕ occurs when X^ϵ hits $\pm\epsilon$. See Figure 2 for an illustration of sample paths of X and X^ϵ .

The following lemma specifies in what sense the perturbed process is an approximation of the trajectory of the web. Here and in the rest of the paper the convergence is uniform on compacts in probability.

Lemma 3.4. $X^\epsilon \xrightarrow{\mathbb{P}} X$ as $\epsilon \rightarrow 0$.

Lemma 3.2 reduces to Lemma 3.4.

Proof of reduction. X^ϵ is $\mathcal{F}_- \otimes \mathcal{F}_+ \otimes \mathcal{G}$ -measurable, so we use Lemma 3.4 to conclude that X is also $\mathcal{F}_- \otimes \mathcal{F}_+ \otimes \mathcal{G}$ -measurable. However, since X is actually independent of \mathcal{G} we can use a basic result on tensor products of

Hilbert spaces (for example Equation (2c8) of [12]) to conclude that X is in fact $\mathcal{F}_- \otimes \mathcal{F}_+$ -measurable. \square

We devote the following section to the proof of Lemma 3.4.

4 Convergence of the perturbed process

In this section we prove that $X^\epsilon \xrightarrow{\mathbb{P}} X$ as $\epsilon \rightarrow 0$. This statement depends only on the joint distribution of X and X^ϵ . We therefore define $Y(\epsilon) = Y = (X^\epsilon, X)$ (generally suppressing the ϵ dependence in the notation). Let us describe the distribution of Y as a two-dimensional random process.

We classify the behavior of the process into three states according to the behavior of X^ϵ with respect to X .

- If X^ϵ is in state S_ϕ and is coalesced with X we say Y is in state S_ϕ^{1D} .
- If X^ϵ is in state S_ϕ and is *not* coalesced with X we say Y is in state S_ϕ^{2D} .
- If X^ϵ is in state S_ψ we say Y is in state S_ψ^{2D} .

Y starts in S_ϕ^{1D} . From S_ϕ^{1D} , Y can only transition to S_ψ^{2D} . This transition occurs when Y hits the origin, as the coalesced X and X^ϵ will continue together until they leave their current half-plane. From S_ψ^{2D} , Y can only transition to S_ϕ^{2D} . This transition occurs when X^ϵ leaves the $(-\epsilon, \epsilon)$ interval (i.e. Y hits either of the $x = \pm\epsilon$ lines). From S_ϕ^{2D} , Y can either transition to S_ϕ^{1D} if X and X^ϵ coalesce (i.e. Y hits the line $x = y$) or transition to S_ψ^{2D} if X^ϵ hits 0 (i.e. Y hits the $x = 0$ line). States and possible transitions of Y are summarized in Figure 3.

The form of the labels given to the states is justified by the following.

Observation. In S_ϕ^{1D} , Y follows the law of a (time scaled) one-dimensional Brownian motion on the line $x = y$. In S_ϕ^{2D} and S_ψ^{2D} , Y follows the law of a two-dimensional Brownian motion.

Additionally observe that by the scale-invariance of Brownian motion, the distribution of the sample paths of $Y(\epsilon)/\epsilon$ is independent of ϵ .

Define $A_\delta = \{(x, y) : |x - y| = \delta\}$. To prove Lemma 3.4 we use the following property of Y :

Lemma 4.1. For given $P > 0$, $\delta > 0$ the probability that Y hits A_δ before it hits (P, P) is $o(1)$ as $\epsilon \rightarrow 0$.







	State	Illustration	Law	Next	Trans. Cond.
X^ϵ	S_ϕ^{1D}		equal	S_ψ^{2D}	hits 0
	S_ϕ^{2D}		indep.	S_ψ^{2D}	hits 0
	S_ψ^{2D}		indep.	S_ϕ^{1D}	hits +/- epsilon
$Y = (x, y)$ $x = X^\epsilon$ $y = X$	S_ϕ^{1D}		equal	S_ψ^{2D}	$x = y = 0$
	S_ϕ^{2D}		indep.	S_ψ^{2D}	$x = 0$
	S_ψ^{2D}		indep.	S_ϕ^{1D}	$x = y$

Figure 3: Illustrated states and transitions of Y

We delay the proof of Lemma 4.1 to Section 4.1.

Proof of Lemma 3.4. Write s for the time at which Y starts. Lemma 3.4 is equivalent to: for all $\delta > 0$, $u > s$

$$\mathbb{P}(Y_t \in \{(x, y) : |x - y| < \delta\} \text{ for all } t \in [s, u]) \rightarrow 1 \text{ as } \epsilon \rightarrow 0.$$

The above statement can be rephrased as

For all $\delta > 0$, $u > s$, the probability that before time u , Y has hit A_δ is $o(1)$ as $\epsilon \rightarrow 0$.

We prove this as follows. For any $\eta > 0$, choose P so that the probability that standard Brownian motion travels from 0 to P in time less than $u - s$ is less than η . Apply Lemma 4.1 to choose ϵ_0 such that, for all $\epsilon < \epsilon_0$, the probability that $Y(\epsilon)$ hits A_δ before (P, P) is less than η . Then the probability that $Y(\epsilon)$ hits A_δ before (P, P) or takes less time than $u - s$ to reach (P, P) is less than 2η . Thus the probability that $Y(\epsilon)$ hits A_δ before time u is less than 2η . \square

4.1 Excursions of Y

In this section we prove the following, which is slightly stronger than Lemma 4.1.

Lemma 4.2. For given $P > 0$, $\delta > 0$ the probability that Y hits A_δ before it hits (P, P) is $O(\frac{1}{\log 1/\epsilon})$.

We begin by introducing the notion of an excursion of Y . Almost surely, the times at which $Y = (0, 0)$ (which are stopping times) form an infinite discrete collection $T_0 < T_1 < \dots$. We say “the probability that an excursion hits a set U is p ” if $\mathbb{P}(Y_t \in U \text{ for some } t \in [T_0, T_1]) = p$. Observe that by equidistribution this probability is the same when t ranges over $[T_i, T_{i+1}]$, and note that the hitting events in question are independent.

Our approach to proving Lemma 4.2 is to demonstrate that

$$\mathbb{P}(\text{an excursion hits } (P, P)) \gg \mathbb{P}(\text{an excursion hits } A_\delta) \text{ as } \epsilon \rightarrow 0.$$

This is realized through the next pair of lemmas whose proofs we delay until Section 4.2.

Lemma 4.3. For given $\delta > 0$, $\mathbb{P}(\text{an excursion hits } A_\delta)$ is $O(\epsilon)$.

Lemma 4.4. For given $P > 0$, $\mathbb{P}(\text{an excursion hits } (P, P))$ is $\Omega(\epsilon \log 1/\epsilon)$.

Proof of Lemma 4.2. Y consists of a sequence of excursions, each of which satisfies exactly one of the following conditions

- the excursion hits A_δ (with probability $O(\epsilon)$),
- the excursion does not hit A_δ but does hit (P, P) (with probability $\Omega(\epsilon \log 1/\epsilon) - O(\epsilon)$, which is itself $\Omega(\epsilon \log 1/\epsilon)$),
- the excursion does not hit A_δ or (P, P) .

The excursions are independent, so the probability that Y hits A_δ before (P, P) is

$$\frac{O(\epsilon)}{\Omega(\epsilon \log 1/\epsilon) + O(\epsilon)} = O\left(\frac{1}{\log 1/\epsilon}\right).$$

□

4.2 Proofs of the excursion lemmas

In this section we give the proof of Lemmas 4.3 and 4.4. For convenience we rotate (and scale) $Y = (X^\epsilon, X)$, defining

$$Z(\epsilon) = Z = (Z^1, Z^2) = \frac{1}{2}(X^\epsilon + X, X^\epsilon - X).$$

As for Y , Z has the following “scale invariance” property: the distribution of sample paths of $Z(\epsilon)/\epsilon$ is independent of ϵ .

Proof of Lemma 4.3. Consider the process Y between times T_0 and T_1 . Once Y arrives at S_ϕ^{1D} it can never hit A_δ before hitting $(0, 0)$. Thus our goal is to show that with probability at least $1 - O(\epsilon)$, Y arrives in S_ϕ^{1D} before hitting A_δ . Next, we observe the following two auxiliary claims:

Claim 1. Whenever $Z^2 = 0$ the probability that subsequently Y arrives at S_ϕ^{1D} before Z^2 hits $\pm\epsilon/2$ is a constant (independent of ϵ).

Claim 2. Whenever $Z^2 = \pm\epsilon/2$ then there is probability equal to ϵ/δ of Z^2 hitting $\pm\delta/2$ before it hits 0.

Claim 1 follows from scale invariance, while Claim 2 is a standard martingale result on Brownian motion (observing that on the relevant time interval Z^2 is a standard Brownian motion).

The reduction of Lemma 4.3 to those two claims is similar to the proof of Lemma 4.2. Claims 1 and 2 imply that between a time when $Z^2 = 0$ and the next time that $Z^2 = 0$ after having hit $\pm\epsilon/2$,

$$\frac{\mathbb{P}(Y \text{ hits } A_\delta)}{\mathbb{P}(Y \text{ arrives at } S_\phi^{1D})} \leq \frac{(1 - C)(\epsilon/\delta)}{C} = O(\epsilon)$$

where C is the constant of Claim 1. As the behavior of Y is independent on those intervals, we deduce Lemma 4.3. \square

Proof of Lemma 4.4. We bound below the probability that an excursion hits (P, P) , i.e. Z hits $(P, 0)$ before returning to $(0, 0)$. We do this by considering the probability that the excursion takes the following form: Z travels from $(0, 0)$ to the line segment $Q = [0, \epsilon] \times \{\epsilon\}$, then hits the horizontal axis for the first time in $[\epsilon, 1] \times \{0\}$, then travels to $(P, 0)$, before returning to $(0, 0)$.

After a stopping time at which $Y = Z = (0, 0)$ there is a positive probability K that Z hits Q before returning to $(0, 0)$. By scale invariance K is independent of ϵ .

Consider Z after hitting some point in Q . We now bound the hitting density of this process on the horizontal axis. Regardless of the point in Q , this density for points on $[\epsilon, 1] \times \{0\}$ is at least

$$\frac{1}{\pi\epsilon} \frac{1}{1 + (x/\epsilon)^2} dx.$$

This follows directly from the classical result that the hitting density on a line of the two-dimensional Brownian motion is a Cauchy distribution (see for example Theorem 2.37 of [4]).

On hitting a point $(x, 0)$ for $x \in [\epsilon, 1]$ the process transitions from state S_ϕ^{2D} to state S_ϕ^{1D} . When in state S_ϕ^{1D} , Z behaves as a one-dimensional Brownian motion on the horizontal axis until it hits $(0, 0)$. By the same martingale argument which justifies Claim 2, the probability of subsequently hitting $(P, 0)$ before $(0, 0)$ is x/P . Integrating this against the hitting density we get that the probability that Z started from some point in Q hits the horizontal axis in $[\epsilon, 1] \times \{0\}$ and then travels to $(P, 0)$ before returning to $(0, 0)$ is at least

$$\frac{1}{\pi P} \int_\epsilon^1 \frac{x/\epsilon}{1 + (x/\epsilon)^2} dx = \frac{\epsilon}{2\pi P} \log \left(\frac{1 + (1/\epsilon)^2}{2} \right), \text{ which is } \Omega(\epsilon \log 1/\epsilon).$$

\square

5 Recovering the web from strips

The next step towards proving Theorem 1.1 is to show that the σ -field associated to a horizontal strip is generated by those σ -fields associated to any two substrips which partition the larger strip. To do so, we follow closely the structure of Section 3.

In the same way we defined \mathcal{F}_+ and \mathcal{F}_- , we introduce a σ -field \mathcal{F}_{yz} for each $y < z \in [-\infty, \infty]$, generated by the web in the horizontal strip $(-\infty, \infty) \times (y, z)$. Formally,

Definition 5.1. For any path f we write $\mathcal{R}_{yz}(f)$ for f stopped at the first time it is outside $(-\infty, \infty) \times (y, z)$, i.e. $\mathcal{R}_{yz}(f)(t) = f(t \wedge T_f)$ where $T_f = \inf\{t : f(t) \notin (y, z)\}$ (as in Section 3, if f starts outside $(-\infty, \infty) \times (y, z)$, it is stopped immediately). We define \mathcal{F}_{yz} to be the σ -field generated by the collection of paths $\{\mathcal{R}_{yz}(X) : X \in \mathcal{X}\}$. The association of strips $(-\infty, \infty) \times (y, z)$ to σ -fields \mathcal{F}_{yz} we call the *horizontal factorization of the Brownian web*.

With these definitions, \mathcal{F} of Section 3 is $\mathcal{F}_{-\infty, \infty}$, \mathcal{F}_+ is $\mathcal{F}_{0, \infty}$ and \mathcal{F}_- is $\mathcal{F}_{-\infty, 0}$. Note that \mathcal{F}_{wx} and \mathcal{F}_{yz} are independent if the intervals (w, x) and (y, z) are disjoint.

Theorem 5.2. $\mathcal{F}_{xz} = \mathcal{F}_{xy} \otimes \mathcal{F}_{yz}$ for all $x < y < z$.

The Brownian web is translation invariant. Thus without loss of generality we can limit ourselves to $x = a$, $y = 0$, $z = b$ in the above theorem, for arbitrary fixed a and b . This reduces the theorem to

$$\mathcal{F}_{ab} = \mathcal{F}_{a0} \otimes \mathcal{F}_{0b}.$$

In the rest of this section we therefore write $\mathcal{R}(\cdot)$ for $\mathcal{R}_{ab}(\cdot)$.

Given the definition of \mathcal{F}_{ab} and since X is arbitrary, Theorem 5.2 is a consequence of the following.

Lemma 5.3. $\mathcal{R}(X)$ is $\mathcal{F}_{a0} \otimes \mathcal{F}_{0b}$ -measurable.

Similarly to Section 3, $\mathcal{R}(X^\epsilon)$ is constructed from trajectories that are \mathcal{F}_{a0} , \mathcal{F}_{0b} and \mathcal{G} -measurable only. Formally, $\mathcal{R}(X^\epsilon)$ is $\mathcal{F}_{a0} \otimes \mathcal{F}_{0b} \otimes \mathcal{G}$ -measurable. Thus as in the reduction of Lemma 3.2 to Lemma 3.4, Lemma 5.3 follows from

Lemma 5.4. $\mathcal{R}(X^\epsilon) \xrightarrow{\mathbb{P}} \mathcal{R}(X)$ as $\epsilon \rightarrow 0$.

We could show this convergence result directly by an extension of the argument we used for the half-planes in Section 3. However, knowing that $X^\epsilon \xrightarrow{\mathbb{P}} X$ is nearly enough, and all that is required in addition is that this convergence is preserved by $\mathcal{R}(\cdot)$. For this we use the following straightforward result in classical analysis.

Lemma 5.5. If T_f is not a turning point of the path f , then the map $f \mapsto \mathcal{R}(f)$ is continuous at f in the topology of uniform convergence on compacts.

Proof of Lemma 5.4. We know from Lemma 3.4 that $X^\epsilon \xrightarrow{\mathbb{P}} X$. In addition X is a Brownian motion so almost surely satisfies the condition of Lemma 5.5, i.e. “ T_X is not a turning point of the path X ”. We conclude that $\mathcal{R}(X^\epsilon) \xrightarrow{\mathbb{P}} \mathcal{R}(X)$ by the continuous mapping theorem (see for example [2] Theorem 2.7). \square

6 Conclusions about the noise

We conclude by supplying a formal framework for the statement of Theorem 1.1 followed by its proof. The following definition of noise is a straightforward extension of that due to Tsirelson (Definition 3d1 of [11]) to multiple dimensions.

A *d-dimensional noise* consists of a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, sub- σ -fields $\mathcal{F}_R \subset \mathcal{F}$ given for all open d -dimensional rectangles $R \subset \mathbb{R}^d$, and a measurable action $(T_h)_h$ of the additive group of \mathbb{R}^d on Ω , having the following properties:

- (a) $\mathcal{F}_R \otimes \mathcal{F}_{R'} = \mathcal{F}_{R''}$ whenever R and R' partition R'' , in the sense that $R \cap R' = \emptyset$ and the closure of $R \cup R'$ is the closure of R'' ,
- (b) T_h sends \mathcal{F}_R to \mathcal{F}_{R+h} for each $h \in \mathbb{R}^d$,
- (c) \mathcal{F} is generated by the union of all \mathcal{F}_R .

When $d = 1$ our definition coincides with that of Tsirelson. In that case, R ranges over all open intervals and condition (a) translates to $\mathcal{F}_{(s,t)} \otimes \mathcal{F}_{(t,u)} = \mathcal{F}_{(s,u)}$ whenever $s < t < u$.

As conditions (b) and (c) are immediate for the horizontal factorization of the Brownian web, Theorem 5.2 immediately implies the following:

Proposition. The horizontal factorization of the Brownian web is a (one-dimensional) noise.

Recall that the horizontal factorization of the Brownian web is an association of a σ -field to any horizontal strip (see Definition 5.1). Observe that the association arises from considering trajectories of the Brownian web stopped at the first time they are outside a particular strip. Similarly, we can associate a σ -field to any vertical strip, or indeed to any rectangle. The former association is the *vertical factorization of the Brownian web* and the latter the *two-dimensional factorization*.

We can extend existing results to derive the following:

Proposition. The Brownian web factorized on two-dimensional rectangles is a two-dimensional noise.

That is, when a rectangle is partitioned horizontally or vertically into two smaller rectangles, the σ -field of the larger is generated by the σ -fields of the two smaller. To see this holds for a rectangle partitioned by a horizontal split observe that this is a consequence of our result restricted to a finite time interval. To see it holds for a vertical split, observe that this is a straightforward modification of the earlier result that the vertical factorization of the Brownian web is a noise (see, for example, [10]).

Furthermore, by a general result of Tsirelson (see [12] Theorem 1e2), a two-dimensional noise is black when one of its one-dimensional factorizations is black. As the vertical factorization of the Brownian web is black (see [11]), we deduce Theorem 1.1.

7 Remarks and open problems

In this paper we present the second known example of a two-dimensional black noise, after the scaling limit of critical planar percolation, as proved by Schramm and Smirnov [6]. They also remark that σ -fields can be associated to a larger class of domains than just rectangles in a way that still allows the σ -field of a larger domain to be recovered from the σ -fields of two smaller domains that partition it. In particular, in Remark 1.8 they claim that this can be done for the scaling limit of site percolation on the triangular lattice, as long as border between those domains has Hausdorff dimension less than $5/4$, and cannot be done if the border has Hausdorff dimension greater than $5/4$. This raises the following question:

Open Problem 1. To what class of two-dimensional domains can the Brownian web noise be extended?

We expect the answer to be more sophisticated than for percolation, since the Brownian web is not rotationally invariant. This suggests that Hausdorff dimension is not a sufficient measurement to determine from which subdomains the Brownian web can be reconstructed. In some sense it is easier to reconstruct the Brownian web from vertical strips than it is from horizontal strips.

We may obtain a better understanding of Open Problem 1 if we can answer

Open Problem 2. Give an explicit example of domains to which the noise cannot be extended.

Analogy with general results in the one-dimensional case suggests that such domains should exist (see [11] Theorem 11a2 and Section 11b).

Moreover, having seen that the Brownian web is a two-dimensional black noise, further examples in two dimensions (and indeed in higher dimensions) are called for. Their discovery would hopefully shed light on the nature of black noises.

Open Problem 3. Find more examples of two-dimensional black noises. Show an example of a black noise in three dimensions or higher.

Readers may wish to note that in [12] Tsirelson extends the concept of a noise to a much more abstract and general setting. This allows results on noises to be formulated and proved without having an explicit underlying geometric base. However, our methods here which are concrete and geometric in nature are more conveniently described in terms of the earlier formulation.

7.1 Acknowledgements

The authors would like to thank Boris Tsirelson for his guidance in the field of noises, for introducing the question to us, and for helpful discussions. In addition we would like to thank Ron Peled and Naomi Feldheim for simplifying some of our arguments and their useful comments on preliminary drafts.

References

- [1] Arratia, R.A. (1979). Coalescing Brownian motions on the line. PhD thesis, University of Wisconsin, Madison.
- [2] Billingsley, P. (1999). Convergence of probability measures. Wiley Series in Probability and Statistics.
- [3] Fontes, L. R. G., Isopi, M., Newman, C. M. and Ravishankar, K. (2004). The Brownian web: characterization and convergence. *Ann. Probab.* 32(4):2857–2883.
- [4] Mörters, P. and Peres, Y. (2010). Brownian Motion. Cambridge University Press.
- [5] Norris J. and Turner, A. (2008). Planar aggregation and the coalescing Brownian flow. arXiv:0810.0211v1.

- [6] Schramm, O. and Smirnov, S. (2011). On the scaling limits of planar percolation. *Ann. Probab.* 39(5):1768–1814.
- [7] Toth, B. and Werner, W. (1998). The true self-repelling motion. *Probab. Theory Related Fields* 111(3):375–452.
- [8] Tsirelson, B. and Vershik, A. (1998). Examples of nonlinear continuous tensor products of measure spaces and non-Fock factorizations. *Reviews in Mathematical Physics* 10(1):81–145.
- [9] Tsirelson, B. (2002) White noises, black noises and other scaling limits. Lecture course, Tel Aviv University. <http://www.tau.ac.il/~tsirel/Courses/Noises/syllabus.html>
- [10] Tsirelson, B. (2004). Scaling Limit, Noise, Stability. *Lecture Notes in Mathematics* 1840, Springer: 1–106.
- [11] Tsirelson, B. (2004). Nonclassical stochastic flows and continuous products. *Probability Surveys* 1: 173–298.
- [12] Tsirelson, B. (2011). Noise as a Boolean algebra of sigma-fields. arXiv:1111.7270v1.