Random walks in dynamic random environments

Daniel Kious

Abstract

We review the recent progress made on random walk in dynamic random environments, and provide (almost) proofs for some of the results, insisting on the heuristic aspect. We start by recalling the situation when the environment is static, where atypical behaviours have been proved for the motion of the walk. We then study a first model of random walk on top of a dynamic environment: a random walk on independent spin flips. We will see that one can obtain satisfying results in that case, using a regeneration structure, due to the fact that the environment mixes uniformly and fast. Finally, we provide the results and conjectures around random walks on top of the exclusion process, which turns out to be much more challenging to study. Along the way, we will bring intuitive explanations on why usual techniques, such as the powerful Kipnis-Varadhan martingale approximation, cannot be used in this context (or at least not in an obvious manner). We will also try and justify, via some heuristic non-rigorous argument, some of the conjectures from the physics literature about a possible super-diffusive regime for the motion of the walk.
The author is grateful for two invitations to give a mini-course, which provided the opportunity to prepare these notes: the CIRM research school *Random walks: applications and interactions* in Marseille, and CLAPEM XVII in Montevideo, Uruguay, both in 2026.

University of Bath
Department of Mathematical Sciences
Bath BA2 7AY
United Kingdom
d.kious@bath.ac.uk

1 Introduction

The goal of these lectures notes is to present questions, as well as a few results, on random walks on the exclusion process, mostly in dimension 1, where interesting behaviour can possibly be observed. This process models the motion advected by a turbulent fluid and has been studied, under various forms, in the physics literature and in the mathematical literature. Later in these notes, we will present some puzzling questions raised in physics papers, see [3, 8], about a possible super-diffusive regime for the motion of the particle. We will even attempt to provide a (very) heuristic self-consistent argument, which may convince someone that, if super-diffusive, the scaling should be of the order of $t^{2/3}$, which is actually not the unanimous prediction.

When studying some models of random walks in dynamic random environment, many of the traditional techniques fail to apply and results become very challenging to obtain. For instance, one of the powerful techniques is the martingale approximation theory developed by Kipnis and Varadhan in [11], and generalised by many others. This technique requires to first understand the environment as seen from the walker which, for static environments, or reversible models, may be doable. Here, we will see that describing the environment seen from the walker is far from being easy, if possible at all. Hence, in order to study models in dynamic environment, where the walk is not reversible and the environment mixed slowly and non-uniformly, hence needs to develop new techniques in order to handle the long-range correlations created along the trajectory of the walk.

The general questions we will be interested in are around the recurrence/transience of the walk, whether it obeys a law of large numbers, i.e. if X_n/n converges almost surely to some limiting *speed* v, and whether a central limit theorem holds (a standard diffusive behaviour), or if one could expect some atypical fluctuations.

We will first review a static version of the model, where the environment is random but fixed through time. We will be interested in some of the atypical behaviours that have been proved for this model in [13]. In the second section, we will see our first example of dynamic environment: the independent spin flips. Unlike the exclusion process, this model mixes fast and uniformly, which allows one to develop regeneration methods to obtain limit theorems. In the third section, we will finally present the model of the random walk on the exclusion process, and discuss the results and proof of [9, 6].

2 The static model on $\mathbb{Z}$

We will here define a very particular case of random walk in random environment, far from being the most general model one can be interested in, or has been understood. Our goal is simply to introduce a static version of a model with dynamic environment where the environment is induced by some interacting particle system. We will also use a simple case for the law of the walk, again with the idea that we simply want to showcase the results available and some of the atypical behaviour that can be observed.

Generally speaking, random walks in random environments model the motion of a particle in a medium with impurities.

The environment. Let us start by defining our environment. Colour each vertex of $\mathbb{Z}$ black with probability $\rho \in (0,1)$, or white with probability $1-\rho$, independently over all vertices. This defines a random collection $\eta = (\eta(x))_{x\in\mathbb{Z}}$ in $\{0,1\}^{\mathbb{Z}}$ of independent Bernoulli random variable with parameter ρ, where we interpret a "1" (or black vertex) as an environment particle, and a "0" as a hole (i.e. a location without a particle). We call the parameter ρ the *density* of the environment, since it corresponds to the average density of particles in it. We will denote $\mathbf{P}^\rho$ the law of η with density $\rho \in (0,1)$, and $\mathbf{E}^\rho$ the corresponding expectation.

The walk. Fix a parameter $p_\bullet \in (1/2, 1)$. Fix a realisation η of the environment under $\mathbf{P}^\rho$ and let us define a nearest-neighbour discrete-time walk $(X_n)_{n\geqslant 0}$ conditionally on the realisation η. Let $X_0 = 0$ and, for $n \geqslant 0$, $X_{n+1} \in \{X_n - 1, X_n + 1\}$ and the evolution of the walk is given by the conditional transition probability

$$P^{\eta}\left(X_{n+1} = X_n + 1 \middle| \mathcal{F}_n\right) = \begin{cases} p_\bullet & \text{if } \eta(X_n) = 1, \\ 1 - p_\bullet & \text{if } \eta(X_n) = 0, \end{cases}$$

where $\mathcal{F}_n = \sigma(X_0, \ldots, X_n)$. This process is pictured in Figure 1 below.

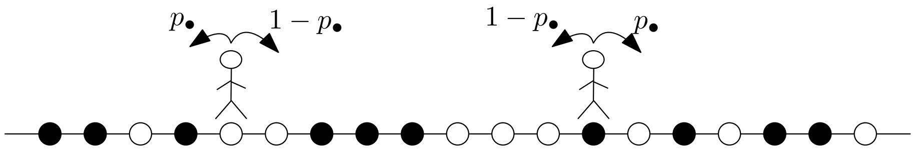

Figure 1: A random walk in a random environment η, where a vertex x is coloured black if $\eta(x) = 1$ and white if $\eta(x) = 0$.

This provides the law of (X_n) when the realisation of the environment has been fixed, i.e. under P^η which is referred to as the *quenched law* of the walk. One can also look at the law of walk after averaging over the environment by defining $\mathbb{P}^\rho[\cdot] = \mathbf{E}^\rho[P^\eta(\cdot)]$. The measure $\mathbb{P}^\rho$ then provides the *annealed law* of the walk, and we will denote $\mathbb{E}^\rho$ the corresponding expectation.
Under the quenched law, (X_n) is nothing else than a Markov chain, but in a totally inhomogeneous medium. Under the annealed law, the medium is average, which should smooth things out, but one drawback is that (X_n) is not a Markov chain anymore but becomes a self-interacting random walk: indeed, knowing the steps the walk took at x in the past provides information on the environment there and thus condition/influence the future evolution of the walk.

The results. A more general version of this model has been studied by Solomon [13] in 1975. As we will see below, this result is pretty much complete as far as recurrence/transience and law of large numbers go. The fluctuations of the walk has been studied as well, as we will see with results of Sinai [12] and of Kesten [10] which show a heavily sub-diffusive behaviour.

Theorem 2.1 (Corollary of [13])**.** *Let $\rho \in (0, 1)$ and $p_\bullet \in (1/2, 1)$.*

A. *(i) If $\rho > 1/2$ then $X_n \to +\infty$ $\mathbb{P}^\rho$-almost surely;*
(ii) If $\rho < 1/2$ then $X_n \to -\infty$ $\mathbb{P}^\rho$-almost surely;
(iii) If $\rho = 1/2$ then X_n is recurrent, i.e. $\limsup_n X_n = -\liminf_n X_n = +\infty$, $\mathbb{P}^\rho$-almost surely.

B. We have that $\lim_{n\to\infty} \frac{X_n}{n} = v(\rho)$, $\mathbb{P}^\rho$-almost surely, where

(i) If $\rho > p_\bullet$ then

$$v(\rho) = \frac{(2p_\bullet - 1)(\rho - p_\bullet)}{p_\bullet - \rho(2p_\bullet - 1)} > 0;$$

(ii) If $\rho < 1 - p_\bullet$ then

$$v(\rho) = -\frac{(2p_\bullet - 1)(\rho - 1 + p_\bullet)}{1 - p_\bullet + \rho(2p_\bullet - 1)} < 0;$$

(iii) If $\rho \in [1 - p_\bullet, p_\bullet]$, then $v(\rho) = 0$.

Remark 2.2. This result is quite complete in dimension one. Much less can be said in dimension two and more. A reason for that seems to be that this random walk is reversible only in dimension one, which allows one to study it via martingale methods and using the environment seen from the walker, as we will see in the proof below. In higher dimensions, reversibility is lost and these techniques collapse.

Remark 2.3. As can be seen in the result above, there exists, in dimension one, a regime where the walk is transient, say to $+\infty$, but has a limiting speed equal to 0, which is often called a transient sub-ballistic regime. Here, this happens when $\rho \in (1/2, p_\bullet]$. We can see a graph of the function $v(\cdot)$ in Figure 2, showing this extended zero-speed regime. Below, we will see that the question of a transient zero-speed

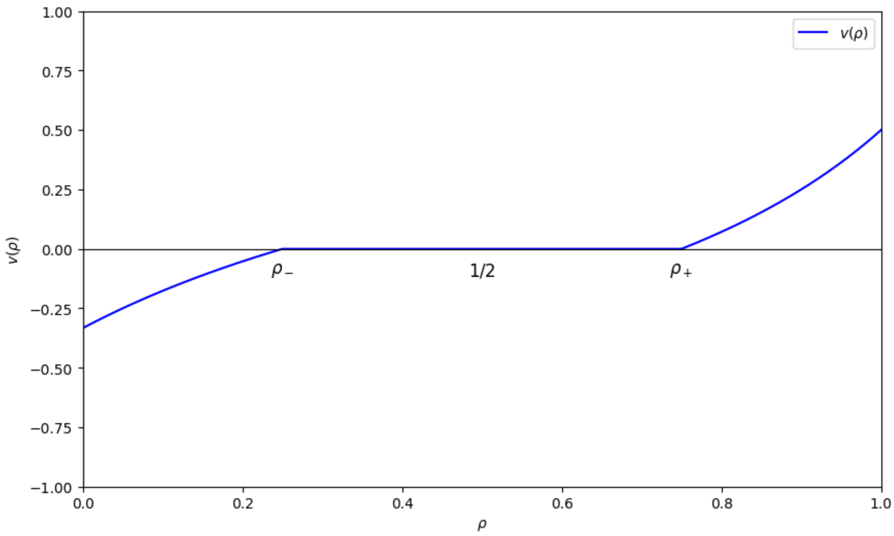

Figure 2: Graph of $v(\rho)$ for $p_\bullet = 3/4$.

regime for random walks on the exclusion process remained open until recently.

It should be noted that it is conjectured that this phenomenon is purely one-dimensional, in the sense that, for dimension two and more, if the random walk is transient in a given direction ℓ (and assuming it is uniformly elliptic), then it has to be ballistic, i.e. X_n/n converges a.s. to some non-zero vector (which has a positive scalar product with ℓ). This is the content of Sznitman's conjecture, who proved it under a stronger transience assumption in [14], known as Condition (T). A lot of work has been done towards solving this conjecture, in particular weakening Condition (T), see e.g. [2] among many other and more recent progress. Nevertheless, proving the conjecture, or proving that Condition (T) holds, still remains open and challenging today and is one of the central questions on the topic.

Proof. Here, we give a very partial proof of Theorem 2.1 due to Solomon. Throughout these notes, we will be more concerned about communicating the general ideas and methods of proofs rather than providing a complete proof, and we trust the reader to go and read the references given if they wish to understand the details of the proofs. The complete proof of a more general version of this result can be found in Solomon's original paper [13] or the excellent notes by Sznitman [15].

Proof of A(i). Here, we want to that when $p_\bullet > 1/2$, then the walk is transient to the right. The argument is purely based on a standard martingale method, hence it is pretty straightforward.... one you found the martingale! Here, viewed from today's modern techniques and in particular electrical networks, it is easy to see that this martingale is actually the signed resistance from 0 to the position of walk, which is now a usual suspect, hence should not be surprising anymore.
For an environment $\eta \in \{0,1\}^{\mathbb{Z}}$, define the function $f^\eta : \mathbb{Z} \to \mathbb{R}$ by

$$f^\eta(x) = \begin{cases} -\sum_{z=0}^{x-1} \left(\frac{1-p_\bullet}{p_\bullet}\right)^{\sum_{k=1}^{z}(2\eta(k)-1)} & \text{if } x \geqslant 0, \\ \sum_{z=x}^{-1} \left(\frac{1-p_\bullet}{p_\bullet}\right)^{-\sum_{k=z+1}^{0}(2\eta(k)-1)} & \text{if } x < 0, \end{cases}$$

noting that $f^\eta(0) = 0$ and $f^\eta(1) = -1$.

Exercise 1. Prove that f^η is harmonic, that is, prove that, for all $x \in \mathbb{Z}$,

$$p_{\eta(x)} f^\eta(x+1) + (1 - p_{\eta(x)}) f^\eta(x-1) = f^\eta(x),$$

where we use the notation $p_{\eta(x)} = p_\bullet \mathbb{1}_{\{\eta(x)=1\}} + (1-p_\bullet)\mathbb{1}_{\{\eta(x)=0\}}$.

Once you have proved the exercise above (because I know all of you will do it, obviously), it is a standard fact that this implies that

the process $\left(f^\eta\left(X_n\right)\right)_n$ the process is a martingale.

Now, let us rewrite slightly f^η as

$$f^\eta(x) = \begin{cases} -\sum_{z=0}^{x-1} \exp\left(-z\ln\left(\frac{p_\bullet}{1-p_\bullet}\right)\frac{\sum_{k=1}^{z}(2\eta(k)-1)}{z}\right) & \text{if } x \geqslant 0, \\ \sum_{z=x}^{-1} \exp\left(|z|\ln\left(\frac{p_\bullet}{1-p_\bullet}\right)\frac{\sum_{k=z+1}^{0}(2\eta(k)-1)}{|z|}\right) & \text{if } x < 0. \end{cases} \tag{2.1}$$

Since we are proving (i), we are in the case where $p_\bullet > 1/2$, and thus, by the usual law of large numbers for i.i.d. random variables, we have that

$$\lim_{z\to+\infty} \frac{\sum_{k=1}^{z}(2\eta(k)-1)}{z} = \lim_{z\to-\infty} \frac{\sum_{k=z+1}^{0}(2\eta(k)-1)}{|z|} = 2\rho - 1 > 0, \quad \mathbf{P}^\rho\text{-a.s.},$$

which thus implies that

$\lim_{x\to+\infty} f^\eta(x)$ exists and is finite $\mathbf{P}^\rho$-a.s., while $\lim_{x\to-\infty} f^\eta(x) = +\infty$.

By the martingale convergence theorem and because, $\mathbf{P}^\rho$-a.s., $\left(f^\eta\left(X_n\right)\right)_n$ is a lower-bounded martingale and therefore converges $f^\eta\left(X_n\right)$ P^η-a.s. to a finite limit, which necessarily implies that $X_n \to +\infty$ (otherwise, either the process would not converge, or would diverge to infinity). The proof of A(ii) is identical.

Proof of A(iii). Now, we assume $\rho = 1/2$. So, under $\mathbf{P}^\rho$, the processes $z \mapsto \sum_{k=1}^{z}(2\eta(k)-1)$ and $z \mapsto \sum_{k=-z+1}^{0}(2\eta(k)-1)$, for $z \geqslant 0$ are simply symmetric simple random walks, and thus, one has that, $\mathbf{P}^\rho$ almost surely,

$$\liminf_{z\to+\infty} \sum_{k=1}^{z}(2\eta(k)-1) = -\infty \quad \text{and} \quad \limsup_{z\to-\infty} \sum_{k=z+1}^{0}(2\eta(k)-1) = +\infty.$$

Therefore, both sums appearing on the right-hand side of (2.1) have infinitely many terms greater than 1 as $|x|$ goes to infinity. Hence, one has that, $\mathbf{P}^\rho$-a.s., $\lim_{x\to+\infty} f^\eta(x) = -\infty$ and $\lim_{x\to-\infty} f^\eta(x) = +\infty$. Now, let $A \in \mathbb{Z}$, let $T_A = \inf\{k \geqslant 0 : X_k = A\}$ and consider the stopped martingale $\left(f^\eta\left(X_{n\wedge T_A}\right)\right)_n$. This martingale is lower-bounded if $A > 0$, resp. upper bounded if $A < 0$, hence, by the martingale convergence theorem, converges to a finite value P^η almost surely. Therefore, the only possibility is that it converges to A, that is $T_A < \infty$ $\mathbf{P}^\rho$-a.s. (or P^η-a.s., for $\mathbf{P}^\rho$-almost all η). This implies that $(X_n)_n$ is recurrent.

Proof of B (i): The environment seen from the walker. Here, we use this occasion to introduce the method of the environment seen from the walker in a relatively simple context. This technique is very powerful, when it can be applied. The basic idea is to observe, at any given time, the environment from the point of view of the random walker: consequently, the picture we will see around us will evolve with time and we obtain a process living in the space of environments. If one can understand the evolution of this process, for instance by proving that it is a reversible Markov chain with some invariant measure (hopefully absolutely continuous with the a priori measure), then this will give us access to the environment surrounding the walker at large time and understand the local drift of the walk, eventually enabling us to prove a law of large numbers and identify the speed of the walk, typically via some martingale approximation.

Let us define the shift operator $\theta_.$ as follows: for all $y \in \mathbb{Z}$, $\theta_y \eta$ is an environment on $\{0,1\}^{\mathbb{Z}}$ such that $\theta_y \eta(x) = \eta(x+y)$. Then, the environment seen from the walker is the process $(\overline{\eta}_n)_n$ on $\{0,1\}^{\mathbb{Z}}$, defined, for all $n \in \mathbb{N}$, by

$$\overline{\eta}_n = \theta_{X_n} \eta.$$

In particular, we have that $\overline{\eta}_n(0) = \eta(X_n)$ is the state of the environment at the location of the walk at time n. Let us also define the *local drift* of the walk at time $n \in \mathbb{N}$ by

$$d(X_n, \eta) = (2p_\bullet - 1)\overline{\eta}_n(0) - (2p_\bullet - 1)(1 - \overline{\eta}_n(0)).$$

For all $n \in \mathbb{N}$, let us define

$$M_n = X_n - \sum_{k=0}^{n-1} d(X_k, \eta). \tag{2.2}$$

It is easy to check that $(M_n)_n$ is a martingale. Moreover, since $|M_{n+1} - M_n| \leqslant 2$ almost surely, we obtain by Azuma-Hoeffding inequality that

$$P^\eta \left(M_n > n^{3/4} \right) \leqslant \exp\left(-\frac{\sqrt{n}}{4} \right).$$

The above right-hand side is summable, hence by the first Borel-Cantelli lemma, we obtain

$$\boxed{\lim_{n\to+\infty} \frac{M_n}{n} \longrightarrow 0 \quad \mathbb{P}^\rho\text{-almost surely.}}$$

Hence, in view of the above and (2.2), proving a law of large numbers for X_n boils down to proving a law of large numbers for the integrated local drift on the right-hand side (2.2). For this purpose, we will study $(\overline{\eta}_n)_n$ the environment seen from the walker, as a Markov chain with kernel $\mathcal{R}$ given by

$$\mathcal{R}h(\eta) = p_{\eta(0)} h \circ \theta_1(\eta) + (1 - p_{\eta(0)}) h \circ \theta_{-1}(\eta),$$

where we again use the notation $p_{\eta(0)} = p_\bullet \eta(0) + (1 - p_\bullet)(1 - \eta(0))$. Now, we would like to find an invariant measure for this process: finding one may be easy (e.g. if the original measure is invariant) but here the formula is not exactly straightforward even if it becomes less surprising when interpreted as a multiple the effective resistance on the right of the walker.

Exercise 2. Define the function $g : \{0,1\}^{\mathbb{Z}} \to \mathbb{R}^+$ by

$$g(\eta) = v(\rho) \times \frac{1}{p_{\eta(0)}} \times \sum_{y=0}^{+\infty} \left(\frac{1 - p_\bullet}{p_\bullet} \right)^{\sum_{k=1}^{y} (2\eta(k) - 1)}, \tag{2.3}$$

where $v(\rho)$ is the expression given in the statement of B (i). Prove that $\mathbb{Q} = g\mathbf{P}^\rho$ is an invariant probability measure for the kernel $\mathcal{R}$. In other words, prove that, for all bounded measure function h, $\int \mathcal{R}h \mathrm{d}\mathbb{Q} = \int h \mathrm{d}\mathbb{Q}$, and that the total mass is equal to 1.

Since $g > 0$, we have that $\mathbb{Q} \sim \mathbf{P}^\rho$, that is, they are mutually *absolutely continuous*. Moreover, from quite general arguments, see [4], one can prove that $\mathbb{Q}$ is the unique ergodic invariant probability measure that is absolutely continuous with respect to $\mathbf{P}^\rho$.
Finally, by Birkhoff's ergodic theorem, we have that, $\mathbb{Q}$-a.s., and thus $\mathbb{P}^\rho$-a.s.,

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} d(X_k, \eta) = \lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} d(0, \overline{\eta}_k) = \mathbb{Q}\left[d(0, \overline{\eta}) \right].$$

Exercise 3. Calculate from (2.3) that

$$\mathbb{Q}\left[d(0, \overline{\eta}) \right] = v(\rho).$$

The above exercise allows us to conclude the law of large numbers for X_n/n with limiting speed $v(\rho)$.

Proof of B (ii)**:** This one is identical as above, with symmetric arguments.

Proof of B (iii)**:** We will describe only briefly how to prove this point, which is done via coupling. Consider $\rho \in [1-p_\bullet, p_\bullet]$, which is equivalent to $p_\bullet \geqslant \rho$ and $p_\bullet \geqslant 1-\rho$. We will define three environments: the original η, an environment η_- with lesser density, and an environment η_+ with greater density.
For $\varepsilon \in (0, 1-p_\bullet)$, define $\rho_+ = \rho + (p_\bullet - \rho) + \varepsilon > \rho$ and $\rho_- = \rho + (1 - p_\bullet - \rho) - \varepsilon < \rho$. We will define the triplet (η_-, η, η_+) such that, under some coupling measure $\mathbf{P}^{\rho,\varepsilon}$, $(\eta_-(z), \eta(z), \eta_+(z))$ are i.i.d. for $z \in \mathbb{Z}$ and

$$(\eta_-(z), \eta(z), \eta_+(z)) = \begin{cases} (0,0,0) & \text{with probability } 1-\rho_+, \\ (0,0,1) & \text{with probability } \rho_+ - \rho, \\ (0,1,1) & \text{with probability } \rho - \rho_-, \\ (1,1,1) & \text{with probability } \rho_-. \end{cases}$$

In particular, note that, $\mathbf{P}^{\rho,\varepsilon}$-almost surely, $\eta_-(z) \leqslant \eta(z) \leqslant \eta_+(z)$ for all $z \in \mathbb{Z}$. Now, one can naturally define three coupled walks on these environments $(X_n)_n$ on η, $(X_n^{(+)})_n$ on η_+ and $(X_n^{(-)})_n$ on η_-, so that they all start at 0 and $X_n^{(-)} \leqslant X_n \leqslant X_n^{(+)}$, almost surely for all $n \geqslant 0$ (let them walk independently in their respective environments when they do not share their location, and couple them monotonically when they are on the same site, using that their environments are nicely ordered).
Now, note that $\rho_+ > p_\bullet$ and $\rho_- < 1 - p_\bullet$, then both $X_n^{(+)}$ and $X_n^{(-)}$ satisfy a law of large numbers with the velocities given in B(i) and B(ii), hence we have that

$$-\frac{\varepsilon}{1-\varepsilon} \leqslant \liminf \frac{X_n}{n} \leqslant \limsup \frac{X_n}{n} \leqslant \varepsilon \frac{2p_\bullet - 1}{p_\bullet - (p_\bullet + \varepsilon)(2p_\bullet - 1)}.$$

□

Intuitive explanation for the transient zero-speed regime. Here, we want to provide a more intuitive justification for the zero-speed regime in B(iii), which will explain, qualitatively, why this is happening. According to the statement, we have a zero-speed transient regime if $1/2 < \rho < p_\bullet$, but here we will restrict ourselves to $p_\bullet > 2/3$ and $1/2 < \rho < (2p_\bullet - 1)/p_\bullet$, since it will make the exposition simpler (but more precise computations can be done to obtain the full result).
Start by considering an interval of length roughly $c\ln(n)$ which is empty for the environment, for some $c > 0$ to be chosen later. That is, we work on the event, for some $z \in \mathbb{Z}$,

$$E_n(z) = \bigcap_{x=0}^{\lfloor c\ln(n)\rfloor - 1} \{\eta(x+z) = 0\}.$$

On this event, let's look at how the time T_z it takes for the random walk to cross from the left-end to the right-end of the interval $[z, z + \lfloor c\ln(n)\rfloor - 1]$, reflecting at z: by standard Gambler's ruin estimates, the expectation of this crossing time is equal to

$$E[T_z] \geqslant c_0 \exp\left(c\ln(n)\ln\left(\frac{p_\bullet}{1-p_\bullet}\right)\right) = c_0 n^{c\ln\left(\frac{p_\bullet}{1-p_\bullet}\right)}. \tag{2.4}$$

Now, observe the probability for the environment to create such an interval which is completely empty:

$$\mathbf{P}^\rho(E_n(z)) \geqslant (1-\rho)^{c\ln(n)} = n^{c\ln(1-\rho)}.$$

Then, we can estimate the probability to see such an empty interval somewhere between 0 and n, by simply considering the non-overlapping intervals of length $c\ln(n)$ (i.e. choosing z multiple of $c\ln(n)$), so that we can use independence. Define the event:

$$\mathcal{E}_n = \bigcup_{k=0}^{\lfloor (n+1)/\lfloor c\ln(n)\rfloor\rfloor} E_n\left(k\lfloor c\ln(n)\rfloor\right)$$

$$\begin{aligned}\mathbf{P}^\rho\left(\mathcal{E}_n\right) &\geqslant 1-\left(1-n^{c\ln(1-\rho)}\right)^{\frac{n}{c\ln(n)}-1}\\ &\geqslant 1-\exp\left(-\left(\frac{n}{c\ln(n)}-1\right)n^{c\ln(1-\rho)}\right).\end{aligned}$$

Choosing $c\in(0,-1/\ln(1-\rho))$, we have that $\mathbf{P}^\rho\left(\mathcal{E}_n\right)\geqslant 1/2$ as soon as n is large enough.
Now, on the event $\mathcal{E}_n$, the time $\mathcal{T}_n$ it takes to go from 0 to n is lower bounded by the time just to cross the one empty interval once, hence its expectation is lower-bounded by the expectation calculated in (2.4), hence we have

$$\mathbb{E}^\rho\left[\mathcal{T}_n\right]\geqslant\mathbb{E}^\rho\left[\mathcal{T}_n\middle|\,\mathcal{E}_n\right]\mathbf{P}^\rho\left(\mathcal{E}_n\right)\geqslant n^{c\ln\left(\frac{p_\bullet}{1-p_\bullet}\right)}\times\frac{1}{2}.$$

If we can choose c so that the probability above is larger than n by an order, than the time it takes to go from 0 to n will be larger than n^a for some $a>1$, which indicates that X_n/n can only converge to 0. So, all in all, we need to be able to choose c such that

$$\frac{1}{\ln\left(\frac{p_\bullet}{1-p_\bullet}\right)}<c<-\frac{1}{\ln\left(1-\rho\right)}.$$

This is possible as soon as the interval $(1/2,(2p_\bullet-1)/p_\bullet)$ is non-empty and ρ belongs to it. This interval is non-empty for $p_\bullet>2/3$.

Fluctuations in the recurrent case, at $\rho=1/2$. Let us say a word about the fluctuations in the recurrent regime, known as Sinai random walk. Here, even if the martingale in the proof above could satisfy a CLT, the position of the walk itself does not satisfy a CLT. Indeed, it was proved by Sinai [12] that the position of the walk is actually heavily sub-diffusive and is typically at distance $\ln^2(n)$ at time n. The limit in law of $X_n/\ln^2(n)$ was identified by Kesten [10] and as a nice description. We will describe Kesten's result at the intuitive level. The description below is not very accurate, we ignore some subtilities and technicalities in the hope to convey the idea of what is happening more clearly.

If $\rho=1/2$, the environment can be described, as mentioned above, as a potential which in this case

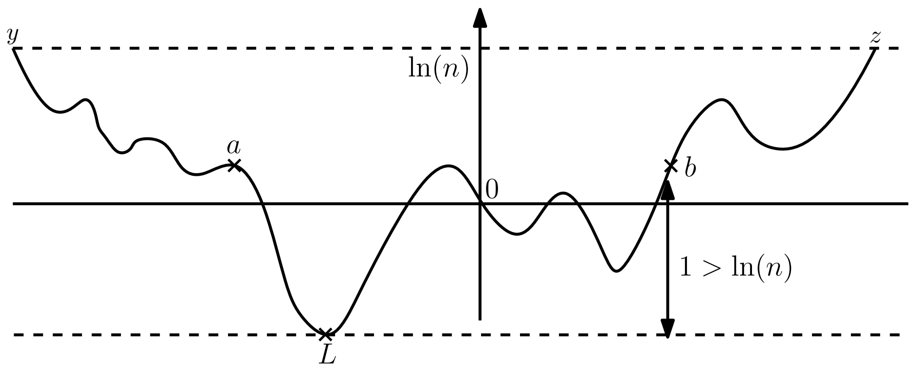

Figure 3: Graph of $v(\rho)$ for $p_\bullet=3/4$.

is nothing else than a symmetric simple random walk trajectory, going up at sites where $\eta(x)=0$ and down if $\eta(x)=1$. The random walk on this environment is then driven by the negative gradient of this potential. The potential will therefore have peaks and valleys (since it looks like a random walk trajectory), and there exists a valley of depth larger than $\ln(n)$ which is the closest to 0, i.e. there is a closest point $y<0$ and a closest point $z>0$ where the values of the potential exceed $\ln(n)$. Let x be the minimum of the potential between y and z (this is not exactly true, but let's think of x as such). This valley may itself contain other deep valleys. Then, Kesten defines a refinement procedure that allows to find the valley of depth larger than $\ln(n)$ for which the walk is the most likely to go and hit the bottom of this valley, say x, quickly. It will take a very long time to escape the depths of this valley, exponential in the depth of the valley and thus of an order longer than time n. Since the potential behaves like a random walk, x is typically at depth of order $\ln(n)$ for the potential and thus x is of order $\ln^2(n)$. Hence, at time n, X_n will be concentrated around this local minimiser, at distance of order $\ln^2(n)$.

3 Dynamic environments

In this section, we start discussing random walk on dynamic environment and point out to the difficulties encountered when trying to use the environment seen from the walker. The main object of these notes is the random walk on the exclusion process, which presents several difficulties: among those, 1) the random walk is not reversible, nor is the environment seen from the walker, 2) the environment is non-uniformly mixing, and 3) the environment mixes slowly. As we will see below, the lack of reversibility is not a definite barrier towards using methods à la Kipnis-Varadhan to study the environment seen from the walker but makes it certainly harder. We will argue below that, even being optimistic in thinking that one could in principle find an explicit invariant measure for the environment seen from the walker, there are other barriers that seem to indicate that this route is not likely to be successful, due to the strong and persisting interactions of the model.
We start by discussing a first model, which shares some of the features above, but where the environment mixes fast and uniformly, which enables us to use probabilistic methods in order to prove limit theorems on the trajectory of the walk.

3.1. A nice example: the random walk on independent spin flips. Here, we start considering the case where the environment is dynamical, described by a Markov process $(\eta_t(x), x \in \mathbb{Z})_{t\geqslant 0}$ on $\{0,1\}^{\mathbb{Z}}$. We initiate this process by choosing $(\eta_0(x), x \in \mathbb{Z})$ to be a collection of i.i.d. Bernoulli random variables with parameter $\rho \in (0,1)$, where ρ will again be the density of the environment. The dynamic of this Markov process are then given by the generator

$$\mathcal{L}f(\eta) = \sum_{x\in\mathbb{Z}} \nu\rho \left(f(\eta^{x,1}) - f(\eta)\right) + \sum_{x\in\mathbb{Z}} \nu(1-\rho) \left(f(\eta^{x,0}) - f(\eta)\right),$$

where, for $i \in \{0,1\}$, $\eta^{x,i}(x) = i$ and $\eta^{x,i}(z) = \eta(z)$ for all $z \in \mathbb{Z}\backslash\{x\}$, and where $\nu > 0$ is the rate of the environment. In words, at rate ν, each site x updates independently its configuration, taking value 1 with probability ρ and 0 with probability $1-\rho$. Let us denote $\mathbf{P}$ the law of the environment (we do not emphasise the parameters since they will be fixed throughout).
Similarly as before, we now define the random walk $(X_n)_n$ by first fixing a realisation of $(\eta_t)_t$ and defining $X_0 = 0$ and, for all $n \geqslant 0$,

$$P^\eta\left(X_{n+1} = X_n + 1 \,|\, \mathcal{F}_n\right) = \begin{cases} p_\bullet & \text{if } \eta_n(X_n) = 1, \\ 1 - p_\bullet & \text{if } \eta_n(X_n) = 0, \end{cases} \tag{3.1}$$

where $\mathcal{F}_n = \sigma(X_0, \ldots, X_n)$. Here again P^η is the quenched law of the walk and we can define the annealed law $\mathbb{P}$ of (X_n) by integrating this measure against the law $\mathbf{P}$ of the environment.

What are the traditional methods one may want to use? A natural approach is to see the (centred) position of the walk as a martingale plus some additive process, which is nothing else that the integrated local drift of the walk, similarly to what we established in (2.2). This integrated local drift can then itself be seen as a martingale plus an error term: two points one may want to prove would then be that the error term is small and that this second martingale does not 'cancel out' with the first martingale, i.e. the sum of these martingales has a non-trivial variance.
1. The classical Kipnis-Varadhan approach [11]. This is an analytical/algebraic L^2 method to prove

central limit theorems. A key starting point to apply this strategy is to work under the invariant measure for the environment seen from the walker: in the better cases, this invariant measure is explicit and it is known that the initial law of the environment is absolutely continuous with respect to it (or even better if they are the same measures). Generally, while it is often possible to know the existence of an invariant measure seen from the walker (for instance by considering the Cesàro mean), the absolute continuity is very far from trivial. We will describe the method in the continuous-time setting, since some quantities are nicer to write down in that case. We consider some environment $(\eta_t)_{t\geqslant 0}$ wit. As above, we define the environment seen from the walker $\overline{\eta} = (\overline{\eta}_t)_{t\geqslant 0}$ by setting, for all $t \geqslant 0$,

$$\overline{\eta}_t = \theta_{X_t} \circ \eta_t,$$

that is, $\overline{\eta}_t(\cot) = \eta_t(X_t + \cdot)$. This process has a generator which can be written as the sum of the generator coming from the dynamics of the environment and the generator coming from the jumps of the random walk, but we can also write it $\mathcal{L} = \mathcal{S} + \mathcal{A}$, where $\mathcal{S}$ is the symmetric part and $\mathcal{A}$ the anti-symmetric part of the generator in $L^2(\mu)$. If we assume that we start the process with measure μ which is reversible, and invariant, for $\overline{\eta}$, then the anti-symmetric part vanishes. Let's work in that case. Then, the Kipnis-Varadhan approach is a martingale approximation, where we write

$$X_t = M_t^{(1)} + \int_0^t d(\eta_s)\mathrm{d}s, \tag{3.2}$$

where $d(\cdot)$ is the instantaneous drift and $M^{(1)}$ is a martingale. Now the idea, is that $M^{(1)}$ should easily satisfy a CLT and then the goal is to prove that the integrated drift also satisfies a CLT. We will work in the case where the drift has mean 0 under μ, for simplicity (but things apply anyway if not). To prove that the integrated drift satisfies a CLT, we write it as a sum of another martingale and an error term, often called the corrector:

$$\int_0^t d(\eta_s)\mathrm{d}s = M_t^{(2)} + \chi_t. \tag{3.3}$$

Again, the martingale $M^{(2)}$ should easily satisfy a CLT. Then, then conclusion follows if we then prove two things: 1) the error term behaves sub-diffusively (i.e. $\chi_t/\sqrt{t}$ goes to 0 in $L^2(\mu)$), and 2) the variance of $M^{(1)} + M^{(2)}$ is non-trivial in the sense that it grows linearly, i.e. the two martingales do not cancel out each other.
To handle the error term, one can try and prove that the Green-Kubo condition is satisfied, which can be written as

$$\int_0^{+\infty} \mathbb{E}_\mu\left[d(\eta_0)d(\eta_t)\right]dt < +\infty. \tag{3.4}$$

One can see that the above integral says something about the correlations along the trajectory of the walk. This condition can be dealt with either if the correlations are clearly decaying fast enough or if some symmetry considerations, such as reversibility, allow to control the integral (we will comment more about this below). Note that, in the general non-reversible case (i.e. when the anti-symmetric part $\mathcal{A}$ of the generator does not vanish), then the integral above is taken over the dynamics for $\mathcal{S}$ only.
Proving that the sum $M^{(1)} + M^{(2)}$ has a linearly growing variance seems to be more model dependent and indicates that, along the walk, there is always a non-trivial local noise: the noise coming from the randomness of the random walk cannot be completely compensated by the noise coming from the environment.
Let us note that reversibility is not strictly necessary, and can be replaced by other conditions such as the Sector condition (which control the influence of the anti-symmetric part in terms of the Dirichlet form of the symmetric part), or by properties of the observable considered (the drift) for instance if it is divergence-free, i.e. the spatial gradient of some function, which allows to prove the Green-Kubo condition.
A famous example where this technique was applied, by Kipnis and Varadhan themselves [11], is to

prove a CLT for the motion of a tagged particle of a simple symmetric exclusion process, in dimension two and more. In that case, the environment seen from the walker has the product measure as reversible (invariant) measure, with an extra particle at 0 (this is the tagged particle). The Green-Kubo condition is satisfied due to the reversibility, which can be interpreted by the fact that any initial local drift will quickly be washed away: this is true in all dimension, including dimension 1. The fact that the martingale part has a linearly growing variance is true only in dimension two and more, not in dimension 1: in higher dimension, the tagged particle can easily go around any local trap, which is not the case in dimension 1 (hence the noise of the walk is compensated by the noise of the environment), where the motion of the tagged particle is actually sub-diffusive.

2. L^2-perturbations approach. In the example we want to consider in this section, the random walk on independent spin flips, it is already much less clear what would be the invariant measure for the environment seen from walker and whether absolute continuity holds (even though we know it exists, as mentioned above). Hence, it is hard to even get a start with the Kipnis-Varadhan method. For this reason, other methods have been developed: here we will briefly discuss the strategy adopted by Avena, Blondel and Faggionato in [1]. The idea for this paper is to consider the case where the random walk is a perturbation of an "unperturbed" random walk, where this reference walk has an invariant measure μ for the environment seen from the walker. Moreover, it is assumed that the environment satisfies a Poincaré inequality, which is an algebraic statement (asking for a spectral gap) which implies that the environment is exponentially mixing, in the sense that there exist a constant $C_P \in (0, \infty)$ such that, for any functions f and g in $L^2(\mu)$, we have

$$|\mathrm{Cov}_\mu\left(f(\overline{\eta_0}), g(\overline{\eta}_t)\right)| \leqslant e^{-\frac{t}{C_P}} \cdot \sqrt{\mathrm{Var}_\mu\left(f(\bar(\eta)\right) \cdot \mathrm{Var}_\mu\left(g(\bar(\eta)\right)}.$$

The method is quite similar to Kipnis-Varadhan in the sense that the authors use a martingale approximation as in (3.2) and (3.3) above and prove that the corrector χ_t behaves sub-diffusively. The Poincaré inequality allows to control the correlations along the trajectory of walk, similar to (3.4), and also to compare the random walk to the unperturbed random walk. Their results apply to the random walk on the independent spin flips as long as the rate of the environment belongs to some explicit interval, or if the bias of the walk is close enough to zero (with an explicit bound in terms of the Poincarè constant). Below, we provide an alternative proof.

3. Regeneration methods. Here, we will consider the random walk defined at the start of Section 3.1 and prove a law of large numbers and a central limit theorem by constructing regeneration times, very much in the spirit of the cone-mixing of Comets and Zeitouni [5]. For this purpose, let us first define a

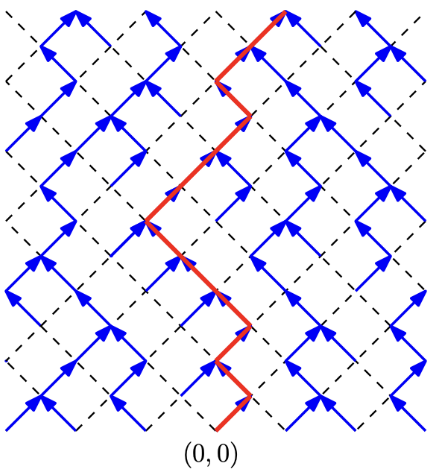

Figure 4: Graphical construction: the arrows are functions of the uniform random variables and the environment at each space-time point. Figure taken from [6].

graphical construction for the random walk, depicted in Figure 4.

To each space-time point $(x, n) \in \mathbb{Z} \times \mathbb{N}$, we attach a uniform random variable $U_{x,n}$, such that these random variables are i.i.d. over x and n. Then, given the environment $(\eta_t)_t$, we generate the walk by fixing $X_0 = 0$ and, for $n \in \mathbb{N}$, if $X_n = x$, then we let

$$X_{n+1} = \begin{cases} X_n - 1 & \text{if } U_{x,n} < p_\bullet - (2p_\bullet - 1)\eta_n(x), \\ X_n + 1 & \text{if } U_{x,n} \geqslant p_\bullet - (2p_\bullet - 1)\eta_n(x). \end{cases}$$

It is straightforward to check that the step has the same distribution as in (3.6). We extend the quenched and annealed measures to include those extra random variables. Stated as such, the random walk is then a function of the environment and of the uniform random variables.
Below, it will be convenient to consider the environment from time -1 rather than 0, hence we naturally extend η_t to all times $t \geqslant -1$.
For each space-time point $(x, n) \in \mathbb{Z} \times \mathbb{N}$, let us define the space-time cone rooted at (x, n) by

$$\mathcal{C}_{x,n} = \{(y, t) : t \geqslant 0, x - t \leqslant y \leqslant x + t\},$$

which is depicted in Figure 5. When started at (x, n), since its trajectory is Lipschitz, the walk will stay inside the this cone $\mathcal{C}_{x,n}$ from time n onwards.

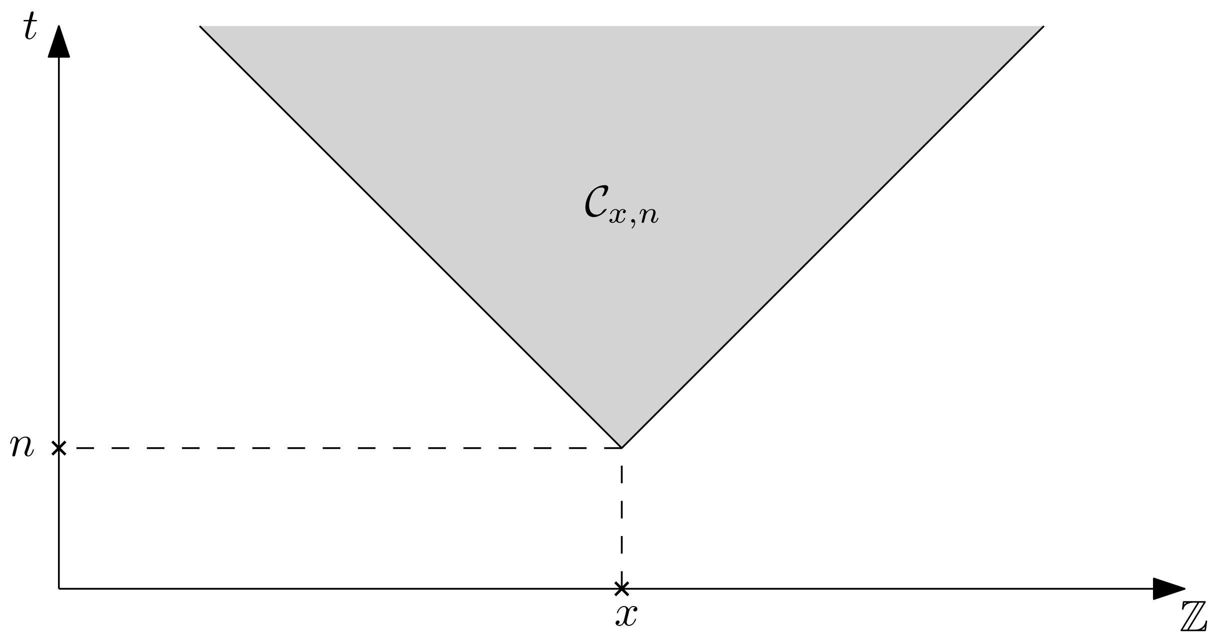

Figure 5: The cone $\mathcal{C}_{x,n}$.

Regeneration will occur at (x, n) when the following *good event* occurs:

$$B_{x,n} = \bigcap_{y \in \mathbb{Z}} B_{x,n}(y) \text{ where } B_{x,n}(y) = \{(\eta_{n-1+s}(y))_{s>0} \text{ updates before time } |y - x|\}.$$

On this good event $B_{x,n}$, we have that $((\eta_k(y), U_{y,k}), (y, k) \in \mathcal{C}_{x,n}))$ is independent of the collection $((\eta_t)_{0 \leqslant t \leqslant n-1}, (U_{\cdot,k})_{0 \leqslant k \leqslant n-1})$, since all points in the space-time cone have updated their configuration for the environment since time $n - 1$. Observing (3.6), one can see that, under the annealed law $\mathbb{P}$, $(X_{n+k} - X_n)_{k \geqslant 0}$ is independent of $(X_k)_{0 \leqslant k \leqslant n}$, since the former a function of $((\eta_k(y), U_{y,k}), (y, k) \in \mathcal{C}_{x,n}))$ and the latter is a function of $((\eta_t)_{0 \leqslant t \leqslant n-1}, (U_{\cdot,k})_{0 \leqslant k \leqslant n-1})$.
In words, on if $X_n = x$ and on $B_{x,n}$, the future increments of the walk are independent of its past trajectory: this is what we call a regeneration event.
Now let us prove that a regeneration happens with constant probability, regardless of the past. Indeed,

using that $B_{x,n}$ is measurable with respect to $(\eta_t, t > n-1)$, which is independent of $\mathcal{F}_n$, we have that

$$\begin{aligned}\mathbb{P}\left(B_{x,n} \middle| \mathcal{F}_n\right) &= \mathbb{P}\left(B_{0,0}\right) \\ &= \prod_{y\in\mathbb{Z}}\left(1-e^{-\nu\cdot(1+|y-x|)}\right) \\ &= \left(1-e^{-\nu}\right)\exp\left(2\sum_{k=1}^{\infty}\ln\left(1-e^{-\nu\cdot k}\right)\right) \\ &= C(\nu),\end{aligned}$$

where $C(\nu)$ is a positive constant, since $\sum_{k=1}^{\infty} e^{-\nu\cdot k} < \infty$. This tells us that, for any time n, as long as we have not revealed information about the environment after time $n-1$, we have a constant chance of regenerating. Intuitively, we can already guess that we will achieve regeneration after a geometric number of attempts. But, when we observe that regeneration fails at a given point, say $(0,0)$, i.e. we observe $B_{0,0}^c$, this conditions the future of the environment (and thus of the walk), hence we need to control the *failure time* R_0, which is a time after which we can safely attempt regeneration, without suffering from conditioning. Let us define

$$R_0 = \sup\left\{1+|y| : y\in\mathbb{Z} \text{ and } B_{0,0}(y) \text{ fails}\right\},$$

see Figure 6 for a depiction. Indeed, if R_0 is finite, then after time R_0, for each $y \in \mathbb{Z}$, we have either

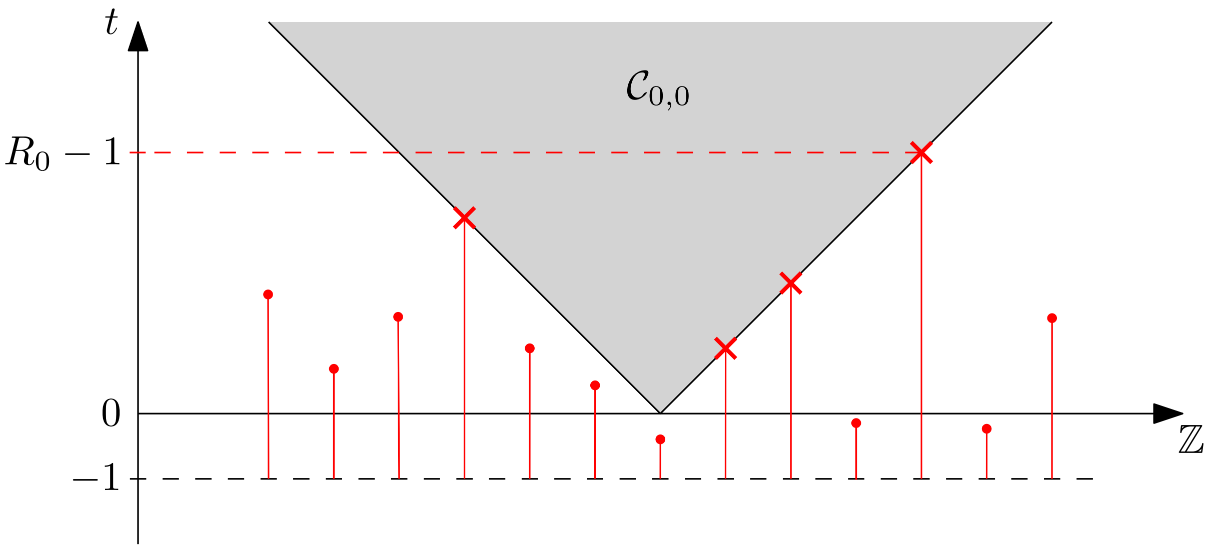

Figure 6: The red disks are places where the environment locally updates its configurations. The red lines hitting the cone $\mathcal{C}_{0,0}$ with a cross depicts locations where the environment does not update its configuration on time; all those locations are within distance $R_0 - 1$ from the origin.

observe 1) that $\eta_t(y)$ has refreshed on time, thus its configuration is unconditioned in the cone $\mathcal{C}_{0,0}$, or 2) that $\eta_t(y)$ has not refreshed by time $|R_0|$, but its refresh rate after that random time is again ν, by memoryless property.

We now need to control how big the failure time R_0; if it cannot be too big than one can see that the random walk will try and fails regeneration a few times (a geometric number of times), each time loosing a time distributed like R_0, and then it will achieve regeneration. This should allow us to control the time to the first regeneration event. For all $t > 0$, we have that, for some positive constant c whose value changes from line to line, and as long as t is large enough (only depending on ν), we have by a computation similar as for $B_{x,n}$,

$$\begin{aligned}
\mathbb{P}\left(R_0 \geqslant t+1 \middle| B_{0,0}^c\right) &= 1 - \frac{\mathbb{P}\left(\bigcup_{y:|y|<t} B_{0,0}^c(y)\right)}{\mathbb{P}\left(\bigcup_{y\in\mathbb{Z}} B_{0,0}^c(y)\right)} \mathbb{P}\left(\bigcap_{y:|y|\geqslant t} B_{0,0}(y)\right) \\
&\leqslant 1 - \frac{1}{2} \prod_{y:|y|\geqslant t} \mathbb{P}\left(B_{0,0}(y)\right) \\
&= 1 - \frac{1}{2}\exp\left(2\sum_{k=\lceil t\rceil}^{\infty} \ln\left(1-e^{-\nu\cdot k}\right)\right) \qquad (3.5)\\
&\leqslant 1 - \frac{1}{2}\exp\left(-4\sum_{k=\lceil t\rceil}^{\infty} e^{-\nu\cdot k}\right) \\
&\leqslant 1 - \frac{1}{2}\exp(-ce^{-\nu t}) \\
&\leqslant ce^{-\nu t}.
\end{aligned}$$

Hence, on the event $B_{0,0}^c$, we have a very good control on the tail of R_0. Note that, on the event $B_{0,0}$, we have that $R_0 = +\infty$.
It is now time to define our *first regeneration time* τ_1. For this purpose, let us define $S_0 = 0$ and, for all $k \in \mathbb{N}$,

$$S_{k+1} = S_k + \theta_{S_k} \circ R_0,$$

where $\theta_\cdot$ is the canonical shift along trajectories. Define the following random index:

$$K = \inf\left\{k \geqslant 0 : S_k < \infty \text{ and } S_{k+1} = \infty\right\}.$$

Then, we define the first regeneration time to be

$$\tau_1 = S_K.$$

Note that, one can write

$$\tau_1 = \sum_{i=1}^{G} R^{(i)},$$

where G is a geometric random variable on $\mathbb{N}$ (so, possibly 0), with success parameter $C(\nu)$, and $R^{(i)}$, $i \geqslant 1$, is an i.i.d. family of random variables, independent of G, and distributed like R_0 under $\mathbb{P}\left(\cdot \middle| B_{0,0}^c\right)$. Hence, since these random variables have exponential tails by (3.5) and G is geometric, we have that

$$\boxed{\mathbb{E}\left[\tau_1^m\right] < \infty \text{ for all } m \in \mathbb{N}.}$$

This regeneration time $\tau_1 = \tau_1(X,\eta)$ is not a stopping time. But τ_1 has the renewal property that, starting from τ_1, since the walk will stay inside $\mathcal{C}_{X_{\tau_1},\tau_1}$, and the environment inside that cone has a fresh configuration (from the starting, invariant distribution), the future increments of the walk are independent of its past and distributed as the law of the walk starting from 0. In particular, we can repeat this construction in order the second regeneration time, the third one, etc. We define, for all $k \geqslant 1$,

$$\tau_{k+1} = \tau_1 + \tau_k\left(X_{\tau_1+\cdot} - X_{\tau_1}, \eta_{\tau_1+\cdot}\left(X_{\tau_1}+\cdot\right)\right).$$

The key property is the following: setting $\tau_0 = 0$, the collection

$$\left(\left(X_{(n+\tau_k)\wedge\tau_{k+1}} - X_{\tau_k}\right)_{n\geqslant 0}, \left(\tau_{k+1}-\tau_k\right), k \geqslant 0\right)$$

is independent and identically distributed, with the same law as (X_{τ_1}, τ_1) under $\mathbb{P}^\rho$. In words, the increments of the walk between two regeneration times are i.id. and the times between two regeneration times are i.i.d. as well. From there, it becomes quite easy to obtain a law of large numbers and a (functional) Central limit theorem. Indeed, note that, using the usual law of large numbers for i.i.d. random variables, we obtain

$$\frac{X_{\tau_k}}{k} = \frac{1}{k}\sum_{i=0}^{k-1}\left(X_{\tau_{i+1}} - X_{\tau_i}\right) \longrightarrow \mathbb{E}\left[X_{\tau_1}\right] \text{ a.s. as } k \to +\infty$$
$$\frac{\tau_k}{k} = \frac{1}{k}\sum_{i=0}^{k-1}(\tau_{i+1} - \tau_i) \longrightarrow \mathbb{E}\left[\tau_1\right] \text{ a.s. as } k \to +\infty.$$

This implies that, almost surely,

$$\lim_{k\to+\infty}\frac{X_{\tau_k}}{\tau_k} = \frac{\mathbb{E}\left[X_{\tau_1}\right]}{\mathbb{E}\left[\tau_1\right]} =: v(\rho).$$

A simple interpolation argument, justifying that, at time n, we are never too far from the previous or next regeneration time, one obtains that X_n/n converges almost surely to $v(\rho)$ as well. The proof for the CLT goes along the same lines, we'll give a rough argument that the reader should be able to make rigorous without too much difficulty. For $n \in \mathbb{N}$, denote k_n the index such that τ_{k_n} is the closest regeneration time around n (for a rigorous argument, one should consider the last and next regeneration times), and note that one can write

$$\frac{X_n - nv(\rho)}{\sqrt{n}} \approx \frac{X_{\tau_{k_n}} - \tau_{k_n}v(\rho)}{\sqrt{k_n}} \cdot \sqrt{\frac{k_n}{\tau_{k_n}}}.$$

Similarly as before, using the traditional central limit theorem for sums of i.i.d. random variables, the first term on the right-hand side converges in distribution to a centred Gaussian random variable, while the second term converges in probability to the square root of $1/\mathbb{E}[\tau_1]$. Therefore, the product converges in distribution to a centred Gaussian random variable. Moreover, it can be shown that the variance is non-trivial and can be expressed in term of the variance of the walk at τ_1 and the variance of τ_1. With a bit of work, but using standard arguments for the convergence of stochastic processes, one can turn this to a functional central limit theorem and obtain

$$\left(\frac{X_{\lfloor nt\rfloor} - ntv(\rho)}{\sqrt{n}}\right)_{t\geqslant 0} \Rightarrow (B_t)_{t\geqslant 0},$$

where the convergence holds in distribution, for the annealed measure $\mathbb{P}^{\rho}$, in the space of càdlàg functions equipped with the Skorohod topology.

Remark 3.1. It is not too hard to see that the argument above can be carried on very similarly for the same model define on $\mathbb{Z}^d$, $d \geqslant 2$.

Remarks on the mixing properties of the environment. What makes the regeneration argument above work particularly well is the fact that the environment we consider has excellent mixing properties. First, the mixing is exponentially fast, meaning that it does not take long for the environment to become random again after observing it. Second, this mixing is *uniform*, meaning that the mixing properties stay excellent, regardless of the observation we make: however bad is the configuration, the environment will become *random and typical* very quickly. For instance, looking back at the Green-Kubo condition in (3.4), one can argue that, for instance at $\rho = 1/2$,

$$\mathbb{E}\left[d(\overline{\eta}_0)d(\overline{\eta}_t)\right] = (2p_\bullet - 1)^2\mathbb{P}\left((\eta_s(0)_{s\geqslant 0} \text{ does not update by time t}\right) \leqslant cd^{-\nu t},$$

which is clearly integrable. In the next model we consider, the environment will **not** be uniformly mixing and the mixing will be slow.

Remarks on the properties of $v(\cdot)$**.** We proved that a law of large numbers holds for all densities $\rho \in (0,1)$, with diffusive fluctuations. It can be easily proved, by symmetry arguments, that we necessarily have that $v(1/2) = 0$. Moreover, at $\rho = 1/2$, using for instance a result from [7], one has that the walk is recurrent. Now, we can even prove continuity and strict monotonicity of $v(\cdot)$ on $(0, 1)$. This can be done by coupling monotonically two environments, one with density ρ and one with density $\rho + \varepsilon$; then coupling two random walks, on of each of those environments, such that the first walk always stays on the left of the second walk; finally, one has to consider *joint regeneration times*, meaning times one regeneration happens on some sort of double cone, with one root at each of the positions of the random walks. This construction is a bit more technical, but will enjoy the same nice properties as before, in particular the first regeneration time will behave very nicely and have all moments. On one hand, taking ε small enough, the probability to see any difference between the two environments can be made arbitrarily small, providing continuity. On the other hand, for $\varepsilon > 0$, there is always a positive probability that the two walks separate right at their last step before regeneration, providing strict monotonicity. This argument will work well because the regeneration structure we defined works for *any* density $\rho \in (0, 1)$, and actually the regeneration event does not depend on the density, only on the update times. So, it should be not too hard for the reader to believe that the speed for this model looks like pictured in Figure 7. This should be put in contrast with Figure 1 and what happens for the random walk on static random

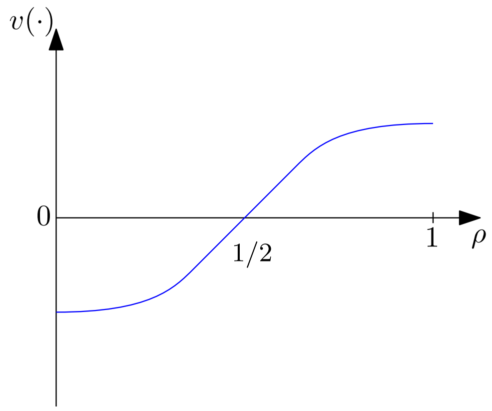

Figure 7: Speed for the random walk on independent spin flips. The speed is strictly monotone and continuous, with value 0 at $1/2$. The derivative is likely not correct, this is just for illustration purposes.

environment, where an open interval of densities have speed 0. Below, we will explain that this was an open question to know if this could happen for the random walk on the exclusion process, where the environment is dynamic, but much slower than the independent splin flips, and conservative, with long-range correlations.

3.2. The random walk on the simple symmetric exclusion process. This model is defined very similarly to what we did in the section above, except that the environment now is given by the states of a symmetric simple exclusion process. The exclusion process describes the motions of particles each performing a continuous-time symmetric simple random walk, at the exception that they obey the *exclusion rule*, meaning that they are not allowed to jump on top of each other, thus any jump leading to two particles being at the same location at the same time is suppressed. Now because all that matters for the random walk on top of that environment is the occupation field of that process, i.e. whether there is a particle at a given location at a given time, it is equivalent to consider the interchange process instead. For the interchange process, all location have a particle or a whole, and each edge has an exponential clock; when the clock rings, we swap the state of the two neighbouring vertices (possible changing nothing to the occupation field). It is not too hard to see that the occupation field of these two environments are the same. On the other hand, if we decide to tag a particle and follow its trajectory, this two models are very different: for the interchange process, a tagged particle will simple behave like a simple random walk (for the exclusion, the tagged particle behaves very differently and is actually sub-diffusive in dimension one).

So, let us now define the model, using the interchange model. Consider a Markov process $(\eta_t(x), x \in \mathbb{Z})_{t \geqslant 0}$ on $\{0, 1\}^{\mathbb{Z}}$. We initiate this process by choosing $(\eta_0(x), x \in \mathbb{Z})$ to be a collection of i.i.d. Bernoulli random variables with parameter $\rho \in (0,1)$, where ρ will again be the density of the environment. The dynamic of this Markov process are then given by the generator

$$\mathcal{L}f(\eta) = \sum_{x\in\mathbb{Z}} \nu \left(f(\eta^{x,x+1}) - f(\eta) \right),$$

where $\eta^{x,x+1}(x) = x+1$, $\eta^{x,x+1}(x+1) = x$ and $\eta^{x,x+1}(z) = \eta(z)$ for all $z \in \mathbb{Z}\backslash\{x, x+1\}$> As before, $\nu > 0$ is the rate of the environment. Let us denote $\mathbf{P}^\rho$ the law of the environment.
Again, we now define the random walk $(X_n)_n$ by first fixing a realisation of $(\eta_t)_t$ and defining $X_0 = 0$ and, for all $n \geqslant 0$, $X_{n+1} - X_n \in \{-1, 1\}$, with

$$P^\eta \left(X_{n+1} = X_n + 1 | \, \mathcal{F}_n \right) = \begin{cases} p_\bullet & \text{if } \eta_n(X_n) = 1, \\ p_\circ & \text{if } \eta_n(X_n) = 0, \end{cases} \tag{3.6}$$

where $\mathcal{F}_n = \sigma(X_0, \ldots, X_n)$. Here, we define two parameters such that $0 < p_\circ < p_\bullet < 1$, which are the two biases that the walk will use depending on whether it is on top of a whole or a particle at the time and location of the jump.
Once again, P^η is the quenched law of the walk and we can define the annealed law $\mathbb{P}^\rho$ of (X_n) by integrating this measure against the law $\mathbf{P}^\rho$ of the environment.

Why is this model hard to study? The environment we consider, the exclusion process, has massive differences compared to the environment of the previous section. The exclusion is a conservative system, meaning that the particles do not ever disappear or appear in the system, and its mixing is both non-uniform and slow. The fact that the mixing is non-uniform can for instance be seen as follows: let's observe the environment over a finite interval at time 0, and let's ask what is the probability to see a particle at location 0 at time t, conditional on this observation. If the mixing was uniform, this probability would be close to ρ, regardless of the event we observe at time 0. Here, on one hand, we can consider a large interval at time 0, around 0, and consider the event where this whole interval is empty of particles: then, the probability of seeing a particle at 0 at time t can be made arbitrarily small by taking the size of the interval large enough (a particle would need to travel for outside this interval to 0 by time t, which is not likely). On the other hand, if we consider the event where this interval is full of particles, then the probability to see a particle at 0 at time t can be made arbitrarily close to 1.
Next, the mixing is quadratic: if we observe a box of size n, it takes a time at least n^2 for it to look random again, which is a very long time! Alternatively, consider the interchange process and the event that there is a particle at 0 at time 0, then the probability to see the same particle back at 0 at time t is of the order of $1/\sqrt{t}$, since this particle performs a simple random walk. This means that the correlations in space-time can decay very slowly for some events: here $1/\sqrt{t}$ is considered slow, for instance because it is not integrable.
Could we apply the Kipnis-Varadhan method? Well, we can certainly define the environment seen from the walker. General arguments would give the existence of an invariant measure for it, but proving that our initial product measure is absolutely continuous with respect to it is a whole different story, and a priori not trivial at all. Hence, to apply Kipnis-Varadhan, we would need to have a good idea of what is the invariant measure for the environment seen from the walker, and then prove some absolute continuity, possibly uniqueness, or something that allows us to conclude that convergence towards this measure happens. It turns out that the product measure is not invariant for the environment seen from the walker (no product measure is). Someone more optimistic than the author may object that this only means that we were not able to find this convenient measure, which is fair enough. But, let us provide a better argument why we believe that Kipnis-Varadhan is not likely to work in that context.
Consider the case where $p_\circ = 1 - p_\bullet$ and $\rho = 1/2$, for simplicity, and let's assume we believe that the walk is diffusive in that case (which is not clear, as we will explain below). Assume that a much more clever friend provide us with a very convenient invariant measure seen from the walker, and let's even make it reversible, which is the best case scenario. But now, remember that a key element is the convergence of the integral of the two-point function in the Green-Kubo condition given in (3.4). Remember that this condition involves the product of the instantaneous drift at 0 and at time t. Considering the

interchange process as environment, each particle or hole perform a simple random walk. At time 0, the walker observe the configuration at 0 seeing, say, a particle (the same holds with a hole), then at time t, either this particle will be at the same location where X_t stands providing a positive correlation, or it will be somewhere else, and in this case, the correlation will be 0. Note that there should be any negative correlation created, since seeing a particle at time 0 should not increase the chances to see a hole at time t. Now, we know that this environment particle is diffusive and if we believe that the walk is diffusive as well, they will both be distributed at time t over a window of order $\sqrt{t}$, hence the probability that they land at the same position should be of the order $\sqrt{t}$ as well (we do not claim that this part is rigorous, due to the interactions along the trajectories, but we believe that this is reasonable). We give a depiction of this event in Figure 8. Hence, one can write

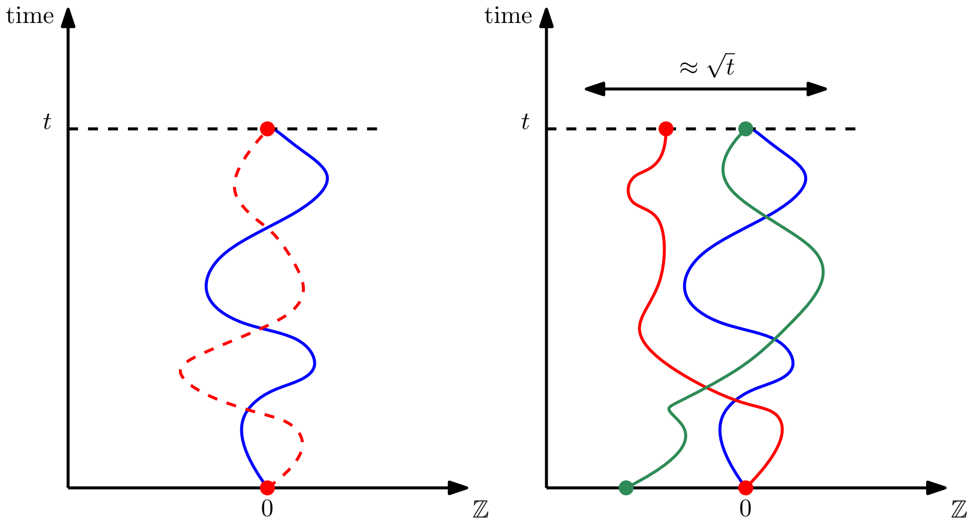

Figure 8: On the left-hand side, the event where the random walk meets again the same environment particle at time t as the one it saw at time 0. Since, for the interchange process, holes also behave like simple random walk, this red dot could be an actual particle for the environment, or a whole. On the right-hand side, we have the event where the particle/hole from time 0 lands at a different location than X_t: in that case, X_t stands on top of this green dot which represents either a hole or a particle and each case should happen with probability $1/2$, by symmetry.

$$\int_0^{+\infty} \mathbb{E}_\mu \left[d(\eta_0)d(\eta_t)\right] dt \geqslant \int_0^{+\infty} \frac{c}{\sqrt{t}} dt = +\infty. \tag{3.7}$$

We do not claim at all that this argument is rigorous, or even actually true, but we believe it to be a convincing hint that, even being optimistic about the invariant measure seen from the walker, the strong correlations along the trajectory of the walk for this model makes it very unlikely to be able to apply a method along the lines of Kipnis-Varadhan, at least not without having a major innovative approach that would allow to deal with those long-range correlations.

References

[1] L. Avena, O. Blondel, and A. Faggionato. Analysis of random walks in dynamic random environments via l^2-perturbations. *Stochastic Processes and their Applications*, 128(10):3490–3530, 2018.

[2] N. Berger, A. Drewitz, and A. F. Ramírez. Effective polynomial ballisticity conditions for random walk in random environment. *Communications on Pure and Applied Mathematics*, 67(12):1947–1973, 2014.

[3] A. Bohr, T. and Pikovsky. Anomalous diffusion in the kuramoto-sivashinsky equation. *Physical Review Letters*, 70, 1993.

[4] E. Bolthausen and A.-S. Sznitman. *Ten Lectures on Random Media*. Oberwolfach Seminars. Birkhäuser Basel, 2012.

[5] F. Comets and O. Zeitouni. A law of large numbers for random walks in random mixing environments. *The Annals of Probability*, 32(1B), 2004.

[6] G. Conchon-Kerjan, D. Kious, and Pierre-François. Sharp threshold for the ballisticity of the random walk on the exclusion process. *arXiv:2409.02096*, 2025.

[7] R. dos Santos and T. Orenshtein. Zero-one law for directional transience of one-dimensional random walks in dynamic random environments. *Electronic Communications in Probability*, 21 (15), 2016.

[8] M. Gopalakrishnan. Dynamics of a passive sliding particle on a randomly fluctuating surface. *Physical Review E*, 69, 2004.

[9] M. R. Hilário, D. Kious, and A. Teixeira. Random walk on the simple symmetric exclusion process. *Commun. Math. Phys.*, 379(1):61–101, 2020.

[10] H. Kesten. The limit distribution of sinai's random walk in random environment. *Physica A: Statistical Mechanics and Its Applications*, 138(1–2):299–309, 1986.

[11] C. Kipnis and S. R. S. Varadhan. Central limit theorem for additive functionals of reversible Markov processes and applications to simple exclusions. *Commun. Math. Phys.*, 104:1–19, 1986.

[12] Y. G. Sinai. The limiting behavior of a one-dimensional random walk in a random medium. *Theory of Probability & Its Applications*, 27(2):256–268, 1983.

[13] F. Solomon. Random Walks in a Random Environment. *The Annals of Probability*, 3(1):1 – 31, 1975.

[14] A.-S. Sznitman. On a class of transient random walks in random environments. *The Annals of Probability*, 29.

[15] A.-S. Sznitman. Topics in random walks in random environment. *School and Conference on Probability Theory, ICTP Lecture Notes*, 17:203–266, 2004.