Dormancy in stochastic interacting systems

Nandan, S.

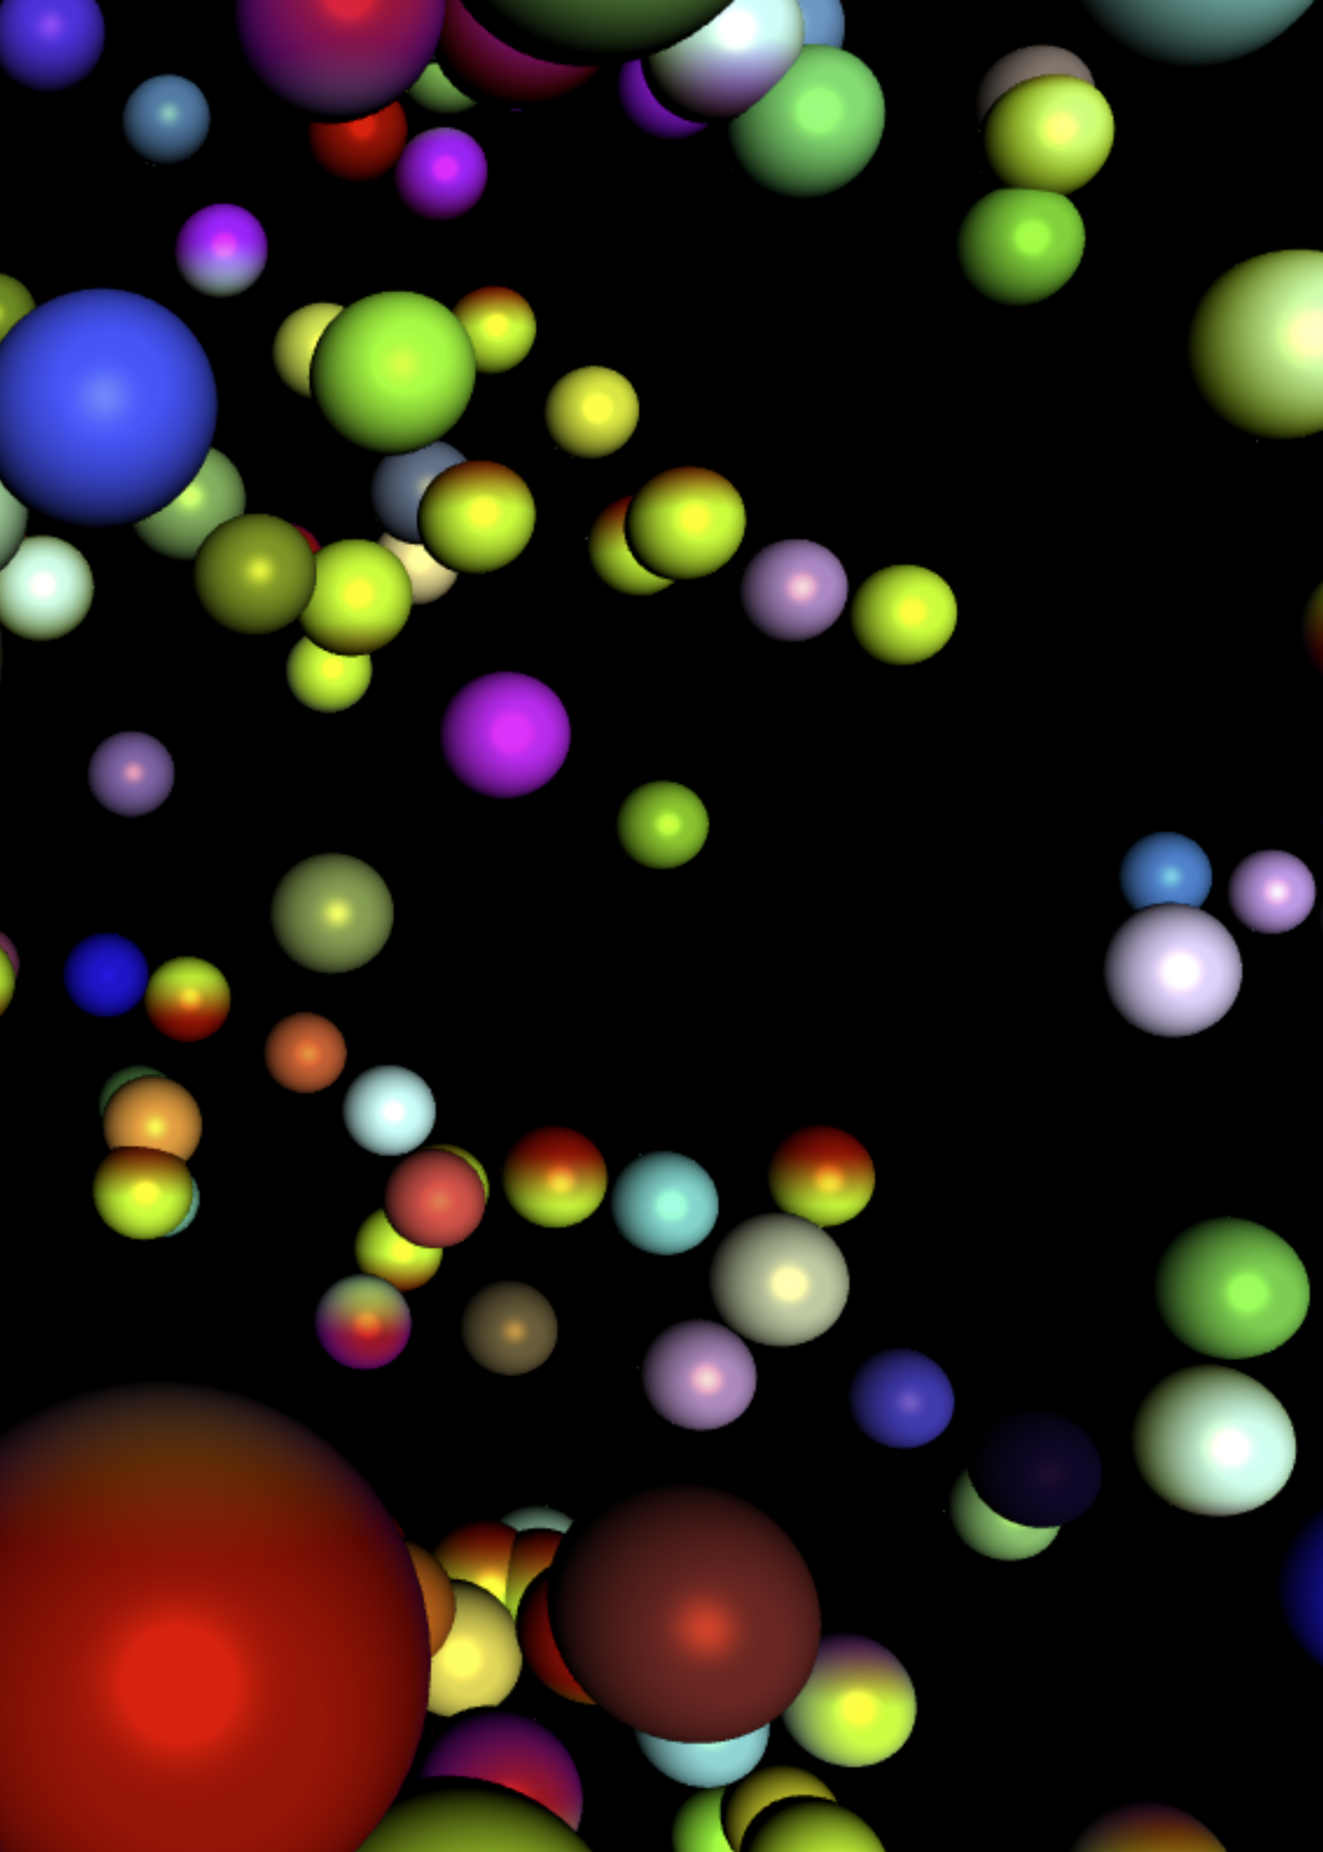

CHAPTER 4

Spatial populations with seed-banks in random environment

This chapter is based on the following paper:
S. Nandan. Spatial populations with seed-banks in random environment: III. Convergence towards mono-type equilibrium. *Electron. J. Probab.*, 28:1–36, 2023.

Abstract

We consider the spatially inhomogeneous Moran model with seed-banks introduced in [46]. Populations comprising *active* and *dormant* individuals are spatially structured in colonies labeled by $\mathbb{Z}^d$, $d \geq 1$. The population sizes are sampled from a *translation-invariant*, *ergodic*, *uniformly elliptic* field that constitutes a *static random environment*. Individuals carry one of two types: ♡ and ♠. Dormant individual resides in what is called a *seed-bank*. Active individuals *exchange* type from the seed-bank of their own colony, and *resample* type by choosing a parent uniformly at random from the distinct active populations according to a symmetric migration kernel. In [46] by exploiting a *dual* process given by an *interacting coalescing particle system*, we showed that the spatial system exhibits a dichotomy between *clustering* (mono-type equilibrium) and *coexistence* (multi-type equilibrium). In this paper, we identify the *domain of attraction* for each mono-type equilibrium in the clustering regime for an *arbitrary fixed* environment. Furthermore, we show that in dimensions $d \leq 2$, when the migration kernel is *recurrent*, for almost surely every realization of the environment, the system with an *initially consistent* type-distribution converges weakly to a mono-type equilibrium in which the probability of fixation to the all type-♡ configuration does not depend on the environment. An explicit formula for the fixation probability is given in terms of an annealed average of the type-♡ densities in the active and the dormant population, biased by the ratio of the two population sizes at the target colony.

Primary techniques employed in the proofs include stochastic duality and the environment process viewed from particle, introduced in [53] for random walk in random environment on a strip. A spectral analysis of Markov operator yields *quenched* weak convergence of the environment process associated with the *single-particle dual* process to a reversible ergodic distribution, which we transfer to the spatial system of populations by using duality.

§4.1 Introduction

In this chapter we study the spatial model with seed-banks introduced in Chapter 2 by treating the preassigned constant population sizes as an *environment* of the system. One of our main results in this chapter is that a full characterization of the domain of attraction for each mono-type equilibrium in the clustering regime is obtained for an *arbitrary fixed* environment (satisfying mild regularity conditions).

Recall that the constituent active and dormant populations in the spatial model maintain constant sizes over time. While this can be biologically explained by assuming that the system receives sufficient supply of environmental resources, a more natural extension would be to consider the model where population sizes come from a *random field* determined by environmental factors such as extreme temperatures, inadequate supply of food resources, etc. Research in this direction has started only recently (see e.g. [28, 17, 152]), although most results are available only for models that are scaled diffusively or are simulation based.

The novelty in the content of present chapter is that here we study the mono-type equilibrium behaviour of the spatial system with seed-banks introduced in Chapter 2 for the setting where the population sizes constitute a *static random* environment. In particular, the sizes are drawn from a *translation-invariant* and *ergodic* random field. Our contributions are two-fold:

(a) When the symmetric migration kernel is *recurrent* (which requires $d \leq 2$) and the random environment is *uniformly elliptic*, we show that the system started from an *initially consistent* type-distribution converges in law to a mono-type equilibrium for almost surely all realisation of the environment. In other words, we prove that the system undergoes *homogenisation* in the *quenched* setting.

(b) We show that, in the homogenised mono-type equilibrium, the *fixation probability* (in law) to the all type-$\heartsuit$ configuration is deterministic, i.e., does not depend on the realisation of the environment. We also provide an explicit formula for this probability.

The techniques used in the proof of the main theorems include stochastic duality, moment relations, semigroup expansion and the environment viewed from the particle recently introduced in [53] for random walk in random environment (RWRE) on a strip, and spectral analysis of Markov kernel operator.

Outline. The chapter is organised as follows. In Section 4.2 we recall some basic results from previous chapters, state our main theorems on the convergence of the system to a mono-type equilibrium, and explain the strategy of the proofs in detail. Section 4.3 is devoted to the analysis of dual process with a single lineage (or single particle) in random environment, where homogenisation results are derived for the associated *environment process*. In Section 4.4 we prove our main theorems using the results derived in Section 4.3. In Appendix B.1, we prove a result stated in Section 4.3 on the existence of a stationary distribution for the aforementioned environment process, and also give a proof of the *strong law of large numbers* for the single-particle dual, which is a result of independent interest. Finally, in Appendix B.2 we prove an auxiliary proposition relating weak convergence of Markov chain to the peripheral point-spectrum of a Markov operator, which is needed for the proof of our main theorems.

§4.2 Main theorems

In Section 4.2.1 we introduce some preliminary notations and set the stage to state our main results. In Section 4.2.2 we give our first main result on the convergence of the system in the clustering regime for an *arbitrary fixed environment* (Theorem 4.2.4). In Section 4.2.3 we consider the system in a *static random* environment that is drawn from a translation-invariant and ergodic field defined on a subset of *uniformly elliptic* environments, and present a *homogenisation* statement in the *quenched* setting on the convergence of the system to a mono-type equilibrium (Theorem 4.2.9–4.2.11). In Section 4.2.4 we discuss the results and shed light on the strategy of the proofs.

§4.2.1 Recollection of previous results and basic notations

Let us recall that under the resampling and exchange dynamics described in Section 2.2.1 of Chapter 2, the initial population sizes $(N_i, M_i)_{i\in\mathbb{Z}^d}$ remain constant over time. Thus, we can naturally think of the sizes of the populations as a *static environment* for the spatial process in (2.2). Throughout the sequel we denote by $\mathfrak{e} := (N_i, M_i)_{i\in\mathbb{Z}^d} \in (\mathbb{N}\times\mathbb{N})^{\mathbb{Z}^d}$ a typical choice for the sizes of the constituent populations and refer to it as the *environment*. From here onwards, we adopt the convention of adding a superscript (or subscript) with Fraktur font to emphasize the dependence of a variable on the realisation of the environment. Let us also recall that the Markov process associated to the spatial system is an interacting particle system denoted by

$$Z^{\mathfrak{e}} := (Z^{\mathfrak{e}}(t))_{t\geq 0}, \qquad Z^{\mathfrak{e}}(t) := (X_i^{\mathfrak{e}}(t), Y_i^{\mathfrak{e}}(t))_{i\in\mathbb{Z}^d}, \tag{4.1}$$

and lives on the inhomogeneous state space

$$\mathcal{X}^{\mathfrak{e}} := \prod_{i\in\mathbb{Z}^d} [N_i]\times[M_i]. \tag{4.2}$$

The superscript $\mathfrak{e}$ indicates the dependence of the process $Z^{\mathfrak{e}}$ on the environment $\mathfrak{e} = (N_i, M_i)_{i\in\mathbb{Z}^d}$, and the pair $(X_i^{\mathfrak{e}}(t), Y_i^{\mathfrak{e}}(t)) \in [N_i]\times[M_i]$ represents the number of active, respectively, dormant individuals of type $\heartsuit$ at time t at colony i. Let $\mathcal{P}^{\mathfrak{e}}$ be the set of probability distributions on $\mathcal{X}^{\mathfrak{e}}$ defined by

$$\mathcal{P}^{\mathfrak{e}} := \big\{\mathcal{P}^{\mathfrak{e}}_\theta\colon\ \theta\in[0,1]\big\}, \qquad \mathcal{P}^{\mathfrak{e}}_\theta := (1-\theta)\delta_{\spadesuit} + \theta\delta_{\heartsuit}, \tag{4.3}$$

where $\delta_\heartsuit$ (resp. $\delta_\spadesuit$) is the Dirac distribution concentrated at the all type-$\heartsuit$ configuration $\mathfrak{e} = (N_i, M_i)_{i\in\mathbb{Z}^d} \in \mathcal{X}^{\mathfrak{e}}$ (resp. the all type-$\spadesuit$ configuration $(0,0)_{i\in\mathbb{Z}^d} \in \mathcal{X}^{\mathfrak{e}}$). Recall that the process $Z^{\mathfrak{e}}$ is said to exhibit *clustering* if and only if the limiting distribution of $Z^{\mathfrak{e}}(t)$ (given that it exists) always falls in $\mathcal{P}^{\mathfrak{e}}$. Otherwise the process is said to be in the *coexistence* regime.

We throughout consider environments that are admissible in the following sense:

Definition 4.2.1 (Admissible environments). Consider the following three conditions for the environment $\mathfrak{e} = (N_i, M_i)_{i\in\mathbb{Z}^d} \in (\mathbb{N}\times\mathbb{N})^{\mathbb{Z}^d}$ and the migration kernel $a(\cdot\,,\cdot)$:

(a) $N_i \geq 2$ and $M_i \geq 2$ for all $i \in \mathbb{Z}^d$.

(b) $\sup_{i\in\mathbb{Z}^d\setminus\{0\}} \|i\|^{-\gamma} N_i < \infty$ and $\sum_{i\in\mathbb{Z}^d} \|i\|^{d+\gamma+\delta} a(0,i) < \infty$ for some $\gamma > 0$ and some $\delta > 0$.

(c) $\lim_{\|i\|\to\infty} \|i\|^{-1} \log N_i = 0$ and $\sum_{i\in\mathbb{Z}^d} \mathrm{e}^{\delta\|i\|} a(0,i) < \infty$ for some $\delta > 0$.

If (a) is satisfied, i.e., in each colony, both the active and the dormant population consist of at least two individuals, then we say that $\mathfrak{e}$ is *non-trivial.* Further, if either (b) or (c) is satisfied, then we say that $\mathfrak{e}$ is *compatible.* Non-trivial and compatible environments are referred to as *admissible* environments. The set of all admissible environments is denoted by $\mathcal{A}$. ■

Chapter 4

Remark 4.2.2. Observe from Theorem 2.2.2 in Chapter 2 that under Assumption 2.A, for any compatible environment, the Markov process $Z^{\mathfrak{e}}$ in (4.1) is well-defined. Condition (a) comes from Assumption 3.A which was made in Chapter 3 because of a technical requirement for determining the clustering regime of the process $Z^{\mathfrak{e}}$ and it can perhaps be removed with minor adaptations.

§4.2.2 Clustering in a fixed environment

In this chapter we refrain from reintroducing the dual process in full generality and only define a version of the dual consisting of a single particle in terms of a *coordinate process* $\Theta^{\mathfrak{e}}$. Informally, the process $\Theta^{\mathfrak{e}}$ keeps track of the location and the state of a single dual particle in time, while the general dual $Z^{\mathfrak{e}}_*$ describes the evolution of the particle via configurations in $\mathcal{X}^{\mathfrak{e}}_*$. The process $\Theta^{\mathfrak{e}}$ plays a key role in the proofs of all our main results, and will be our sole focus in Section 4.3. Later, in Section 4.4.1 we will explain via Lemma 4.4.2 how the single-particle process $\Theta^{\mathfrak{e}}$ is related to the general dual process $Z^{\mathfrak{e}}_*$. We refer the reader to Section 2.4.2 of Chapter 2 and Section 3.3.1 of Chapter 3 for further insight into the general dual process $Z^{\mathfrak{e}}_*$.

Definition 4.2.3 (Single-particle dual process). The single-particle dual process

$$\Theta^{\mathfrak{e}} := (\Theta^{\mathfrak{e}}(t))_{t\geq 0}, \qquad \Theta^{\mathfrak{e}}(t) = (x^{\mathfrak{e}}_t, \alpha^{\mathfrak{e}}_t), \tag{4.4}$$

in environment $\mathfrak{e} := (N_i, M_i)_{i\in\mathbb{Z}^d}$ is the continuous-time Markov chain on the state space

$$G := \mathbb{Z}^d \times \{0,1\} \tag{4.5}$$

with transition rates

$$\begin{aligned}(i,1) &\longrightarrow \begin{cases}(j,1) & \text{at rate } a(0,j-i), \quad j\in\mathbb{Z}^d, j\neq i\\ (i,0) & \text{at rate } \lambda,\end{cases}\\ (i,0) &\longrightarrow (i,1) \qquad \text{at rate } \lambda K_i,\end{aligned} \tag{4.6}$$

where $i\in\mathbb{Z}^d$ and the environment $\mathfrak{e}$ fixes K_i by (2.1). We define the time-t probability transition kernel $p_t^{\mathfrak{e}}(\cdot\,,\cdot)\,:G\times G\to[0,1]$ associated to $\Theta^{\mathfrak{e}}$ as

$$p_t^{\mathfrak{e}}(\eta,\xi) := P_\eta^{\mathfrak{e}}(\Theta^{\mathfrak{e}}(t)=\xi), \quad \eta,\xi\in G, \tag{4.7}$$

where $P_\eta^{\mathfrak{e}}$ is the law of the process $\Theta^{\mathfrak{e}}$ started at $\eta\in G$. ■

The coordinates $x_t^{\mathfrak{e}}$ and $\alpha_t^{\mathfrak{e}}$ in (4.4) represent, respectively, the location in $\mathbb{Z}^d$ and the state (active or dormant) of the particle at time t, where 0 stands for dormant and 1 stands for active. Note from (4.6) that only the wake-up rate of the particle depends on the environment $\mathfrak{e}$, and only via the ratios $(K_i)_{i\in\mathbb{Z}^d}$ defined in (2.1). Indeed, the average time spent in the dormant state by the particle at site i is proportional to K_i^{-1}, the relative strength of the seed-bank at colony i. The particle in the active state migrates according to the kernel $a(\cdot\,,\cdot)$, and so migration is not affected by the environment $\mathfrak{e}$, at least not in a direct manner. This makes the analysis of the single-particle process $\Theta^{\mathfrak{e}}$ in a typical *random* environment $\mathfrak{e}$ easier than the full dual process $Z_*^{\mathfrak{e}}$.

Let us now state the main result of this section.

Theorem 4.2.4 (Domain of attraction). *Suppose that the process* $Z^{\mathfrak{e}} := (Z^{\mathfrak{e}}(t))_{t\geq 0}$ *is in the clustering regime and* $Z^{\mathfrak{e}}(0) = (X_i^{\mathfrak{e}}(0), Y_i^{\mathfrak{e}}(0))_{i\in\mathbb{Z}^d}$ *has distribution* $\mu^{\mathfrak{e}}\in\mathcal{P}(\mathcal{X}^{\mathfrak{e}})$, *where* $\mathfrak{e} := (N_i,M_i)_{i\in\mathbb{Z}^d}\in\mathcal{A}$ *is an arbitrarily fixed environment. If* $\mu_t^{\mathfrak{e}}$ *denotes the time-t distribution of the process* $Z^{\mathfrak{e}}$, *then the following are equivalent:*

(a) $\mu_t^{\mathfrak{e}}$ *converges weakly as* $t\to\infty$.

(b) *For any* $(i,\alpha)\in G := \mathbb{Z}^d\times\{0,1\}$,

$$f^{\mathfrak{e}}(i,\alpha) := \lim_{t\to\infty}\sum_{(j,\beta)\in G} p_t^{\mathfrak{e}}((i,\alpha),(j,\beta))\,\mathbb{E}_{\mu^{\mathfrak{e}}}\left[\beta\,\frac{X_j^{\mathfrak{e}}(0)}{N_j} + (1-\beta)\,\frac{Y_j^{\mathfrak{e}}(0)}{M_j}\right] \text{ exists}, \tag{4.8}$$

where $p_t^{\mathfrak{e}}(\cdot\,,\cdot)$ *is as in Definition* 4.2.3.

Further, if any of the above two conditions is satisfied, then there exists $\theta_{\mathfrak{e}}\in[0,1]$ *such that* $f^{\mathfrak{e}}(\cdot)\equiv\theta_{\mathfrak{e}}$ *and*

$$\lim_{t\to\infty}\mu_t^{\mathfrak{e}} = (1-\theta_{\mathfrak{e}})\delta_{\spadesuit} + \theta_{\mathfrak{e}}\delta_{\heartsuit}. \tag{4.9}$$

The following corollary states that if the process $Z^{\mathfrak{e}}$ exhibits clustering and starts from an initial distribution that puts a constant density of type ♡ individuals at *infinity*, then with probability 1 the spatial process $Z^{\mathfrak{e}}$ converges towards a mono-type equilibrium. Further, the probability of fixation to the all type-♡ configuration in the attained equilibrium is given by the initial density of type ♡ in the populations at *infinity*.

Corollary 4.2.5. *Suppose that the process* $Z^{\mathfrak{e}}$ *is in the clustering regime and* $\mu_t^{\mathfrak{e}}$ *denotes the time-t distribution of the process, where* $\mathfrak{e} := (N_i, M_i)_{i\in\mathbb{Z}^d} \in \mathcal{A}$ *is fixed arbitrarily. If the initial distribution* $\mu^{\mathfrak{e}} := \mu_0^{\mathfrak{e}}$ *is such, that for some* $\theta_{\mathfrak{e}} \in [0,1]$,

$$\lim_{\|i\|\to\infty} \int_{\mathcal{X}^{\mathfrak{e}}} \tfrac{X_i}{N_i}\, \mathrm{d}\mu^{\mathfrak{e}}\{(X_k, Y_k)_{k\in\mathbb{Z}^d}\} = \lim_{\|i\|\to\infty} \int_{\mathcal{X}^{\mathfrak{e}}} \tfrac{Y_i}{M_i}\, \mathrm{d}\mu^{\mathfrak{e}}\{(X_k, Y_k)_{k\in\mathbb{Z}^d}\} = \theta_{\mathfrak{e}}, \tag{4.10}$$

then

$$\lim_{t\to\infty} \mu_t^{\mathfrak{e}} = (1-\theta_{\mathfrak{e}})\delta_{\spadesuit} + \theta_{\mathfrak{e}}\delta_{\heartsuit}. \tag{4.11}$$

Let us recall that in Chapter 3, the clustering criterion stated in Theorem 2.4.12 of Chapter 2 was further refined, and conditions on the environment $\mathfrak{e}$ and other parameters were obtained for which the process $Z^{\mathfrak{e}}$ exhibits clustering. In particular, it was shown (see Corollary 3.2.8 in Chapter 3) that clustering prevails under the following set of conditions:

Assumption 4.A (Clustering environment). The migration kernel $a(\cdot\,,\cdot)$ satisfying Assumption 2.A and the environment $\mathfrak{e} = (N_i, M_i)_{i\in\mathbb{Z}^d}$ are such that

(1) $a(\cdot\,,\cdot)$ is symmetric, i.e.,

$$a(0,i) = a(0,-i), \quad i \in \mathbb{Z}^d. \tag{4.12}$$

(2) $a(\cdot\,,\cdot)$ generates a recurrent random walk on $\mathbb{Z}^d$ that satisfies a local central limit theorem (LCLT). This requirement implicitly forces $d \le 2$ and requires the migration kernel $a(\cdot\,,\cdot)$ to have a finite second moment.

(3) The relative strength of the seed-banks determined by $\mathfrak{e}$ are spatially uniformly bounded, i.e.,

$$\sup_{i\in\mathbb{Z}^d} \tfrac{M_i}{N_i} < \infty. \tag{4.13}$$

(4) The sizes of the active populations determined by $\mathfrak{e}$ are *non-clumping*, i.e.,

$$\inf_{i\in\mathbb{Z}^d} \sum_{\|j-i\|\le R} \tfrac{1}{N_j} > 0 \quad \text{for some } R < \infty. \tag{4.14}$$

■

In view of the above, unless stated otherwise, we will throughout assume that Assumptions 2.A and 4.A are in force. We remark that the above conditions are sufficient but not necessary for the process $Z^{\mathfrak{e}}$ to remain in the clustering regime. The following corollary is immediate.

Corollary 4.2.6. *Suppose that Assumptions* 2.A *and* 4.A *are in force. Then the result in Theorem* 4.2.4 *holds.*

§4.2.3 Clustering in random environment

In this section we consider the process $Z^{\mathfrak{e}}$ in a static random environment $\mathfrak{e}$. Let us introduce the necessary notations before we present our main theorems. To simplify our analysis, we only consider *uniformly elliptic* environments.

Definition 4.2.7 (Uniformly elliptic environment). An environment $\mathfrak{e} \in (\mathbb{N}^2)^{\mathbb{Z}^d}$ with $\mathfrak{e} := (N_i, M_i)_{i \in \mathbb{Z}^d}$ is said to be *uniformly elliptic* if

$$(N_i, M_i) \in \{2, 3, \dots, \mathfrak{K}\}^2 \tag{4.15}$$

for all $i \in \mathbb{Z}^d$ and some natural number $\mathfrak{K} \geq 2$. The set of all environments satisfying (4.15) is denoted by $\mathcal{E}_{\mathfrak{K}}$. ■

From here onwards we fix a natural number $\mathfrak{K} \geq 2$, which we refer to as the ellipticity constant. We equip $\mathcal{E}_{\mathfrak{K}}$ with the product topology and the Borel σ-field Σ. The product topology is naturally induced by the metric $\mathcal{H} : \mathcal{E}_{\mathfrak{K}} \times \mathcal{E}_{\mathfrak{K}} \to [0, \infty)$,

$$\mathcal{H}((N_i, M_i)_{i \in \mathbb{Z}^d}, (\widehat{N}_i, \widehat{M}_i)_{i \in \mathbb{Z}^d}) := \sum_{i \in \mathbb{Z}^d} \frac{1}{2^{\|i\|}} \big[1 \wedge (|N_i - \widehat{N}_i| + |M_i - \widehat{M}_i|)\big]. \tag{4.16}$$

In this metric topology, $\mathcal{E}_{\mathfrak{K}}$ is a compact Polish space, and the Borel σ-field Σ becomes countably generated. Trivially, $\mathcal{E}_{\mathfrak{K}} \subset \mathcal{A}$ (see Definition 4.2.1) and so the process $Z^{\mathfrak{e}}$ is well-defined for any $\mathfrak{e} \in \mathcal{E}_{\mathfrak{K}}$. Note that any $\mathfrak{e} \in \mathcal{E}_{\mathfrak{K}}$ automatically satisfies conditions (3)–(4) in Assumption 4.A.

Definition 4.2.8 (Translation operators). For each $j \in \mathbb{Z}^d$, the shift operator $T_j : \mathcal{E}_{\mathfrak{K}} \to \mathcal{E}_{\mathfrak{K}}$ is defined by the map

$$\mathfrak{e} \mapsto T_j \mathfrak{e}, \quad T_j \mathfrak{e} := (N_{i+j}, M_{i+j})_{i \in \mathbb{Z}^d}, \tag{4.17}$$

where $\mathfrak{e} := (N_i, M_i)_{i \in \mathbb{Z}^d} \in \mathcal{E}_{\mathfrak{K}}$. The action of T_j on a set is interpreted pointwise, i.e., for $A \subset \mathcal{E}_{\mathfrak{K}}$, $T_j A := \{T_j \mathfrak{e} \,:\, \mathfrak{e} \in A\}$. ■

We impose the following assumption on the law of the random environment:

Assumption 4.B (Translation-invariant and ergodic field). The probability law $\bar{\mathbb{P}}$ of the random environment $\mathfrak{e}$ is defined on the measurable Polish space $(\mathcal{E}_{\mathfrak{K}}, \Sigma)$ and satisfies:

(1) For any $A \in \Sigma$ and $j \in \mathbb{Z}^d$, $\bar{\mathbb{P}}(T_j^{-1} A) = \bar{\mathbb{P}}(A)$.

(2) If $A \in \Sigma$ is such that $T_j^{-1} A = A$ for all $j \in \mathbb{Z}^d$, then $\bar{\mathbb{P}}(A) \in \{0, 1\}$.

We use $\bar{\mathbb{E}}$ to denote the expectation w.r.t. $\bar{\mathbb{P}}$. ■

We are now ready to state the main result of this section.

CHAPTER 4

Theorem 4.2.9 (Convergence in random environment). *Let $f_A, f_D : \mathcal{E}_{\mathfrak{K}} \to [0,1]$ be two Σ-measurable functions such that, for $\bar{\mathbb{P}}$-a.s. every realisation of $\mathfrak{e} := (N_i, M_i)_{i\in\mathbb{Z}^d}$, the initial law $\mu^{\mathfrak{e}} \in \mathcal{P}(\mathcal{X}^{\mathfrak{e}})$ of the process $Z^{\mathfrak{e}}$ satisfies the following for all $i \in \mathbb{Z}^d$:*

$$\int_{\mathcal{X}^{\mathfrak{e}}} \tfrac{X_i}{N_i}\, \mathrm{d}\mu^{\mathfrak{e}}\{(X_k, Y_k)_{k\in\mathbb{Z}^d}\} = f_A(T_i\mathfrak{e}), \qquad \int_{\mathcal{X}^{\mathfrak{e}}} \tfrac{Y_i}{M_i}\, \mathrm{d}\mu^{\mathfrak{e}}\{(X_k, Y_k)_{k\in\mathbb{Z}^d}\} = f_D(T_i\mathfrak{e}). \tag{4.18}$$

If Assumption 2.A *and conditions* (1)–(2) *in Assumption* 4.A *hold, then, for $\bar{\mathbb{P}}$-a.s. every realisation of the environment $\mathfrak{e}$, $Z^{\mathfrak{e}}(t)$ converges in law to $(1-\theta)\delta_{\spadesuit} + \theta\delta_{\heartsuit}$, where the fixation probability θ to the all type-$\heartsuit$ configuration $\mathfrak{e} \in \mathcal{X}^{\mathfrak{e}}$ does not depend on the realisation of the environment and is given by*

$$\theta = \frac{1}{1+\rho}\int_{\mathcal{E}_{\mathfrak{K}}} \left[f_A((N_k, M_k)_{k\in\mathbb{Z}^d}) + \tfrac{M_0}{N_0} f_D((N_k, M_k)_{k\in\mathbb{Z}^d})\right] \mathrm{d}\bar{\mathbb{P}}\{(N_k, M_k)_{k\in\mathbb{Z}^d}\}, \tag{4.19}$$

with $\rho := \bar{\mathbb{E}}\big[\frac{M_0}{N_0}\big] = \int_{\mathcal{E}_{\mathfrak{K}}} \frac{M_0}{N_0}\, \mathrm{d}\bar{\mathbb{P}}\{(N_k, M_k)_{k\in\mathbb{Z}^d}\}$, the average relative strength of the seed-bank in each colony.

Let us look at a simple example where the conditions in the above theorem are met.

Example 4.2.10 (Homogenised fixation probability). Fix $\kappa \in [0,1]$. Suppose that, for a typical environment $\mathfrak{e} := (N_i, M_i)_{i\in\mathbb{Z}^d}$ drawn from the law $\bar{\mathbb{P}}$, the process $Z^{\mathfrak{e}}$ starts with the initial law $\mu^{\mathfrak{e}} \in \mathcal{P}(\mathcal{X}^{\mathfrak{e}})$ given by

$$\mu^{\mathfrak{e}} := \bigotimes_{i\in\mathbb{Z}^d} \mathrm{Binomial}(N_i, \tfrac{\kappa}{N_i}) \otimes \mathrm{Uniform}([M_i]). \tag{4.20}$$

In other words, in the spatial system of populations with sizes $(N_i, M_i)_{i\in\mathbb{Z}^d}$, initially each active individual of colony i independently adopts type $\heartsuit$ with probability $\frac{\kappa}{N_i}$, and the number of type-$\heartsuit$ dormant individuals, which is given by $Y_i^{\mathfrak{e}}(0)$, is uniformly distributed over $[M_i] = \{0, 1, \ldots, M_i\}$. In this case, if we let $f_A : \mathcal{E}_{\mathfrak{K}} \to [0,1]$ to be the map $\mathfrak{e} \mapsto \frac{\kappa}{N_0}$ and $f_D : \mathcal{E}_{\mathfrak{K}} \to [0,1]$ to be the constant map $\mathfrak{e} \mapsto \frac{1}{2}$, then $\mu^{\mathfrak{e}}$ satisfies

$$\mathbb{E}_{\mu^{\mathfrak{e}}}\big[\tfrac{X_i^{\mathfrak{e}}(0)}{N_i}\big] = \tfrac{\kappa}{N_i} = f_A(T_i\mathfrak{e}), \quad \mathbb{E}_{\mu^{\mathfrak{e}}}\big[\tfrac{Y_i^{\mathfrak{e}}(0)}{M_i}\big] = \tfrac{1}{2} = f_D(T_i\mathfrak{e}), \tag{4.21}$$

for all $i \in \mathbb{Z}^d$. Thus, if the migration kernel $a(\cdot\,,\cdot)$ is symmetric, recurrent and satisfies a LCLT, then by Theorem 4.2.9 we have that, for $\bar{\mathbb{P}}$-a.s. every realisation of $\mathfrak{e}$, the process $Z^{\mathfrak{e}}$ converges in law to $(1-\theta)\delta_{\spadesuit} + \theta\delta_{\heartsuit}$, where θ is given by

$$\theta = \frac{1}{1+\bar{\mathbb{E}}[M_0/N_0]}\left[\bar{\mathbb{E}}\big[\tfrac{\kappa}{N_0}\big] + \tfrac{1}{2}\bar{\mathbb{E}}\big[\tfrac{M_0}{N_0}\big]\right]. \tag{4.22}$$

This tells that, in the long run, the probability of fixation of the spatial population to the all type-$\heartsuit$ configuration is θ and does not depend on the realisation of the environment $\mathfrak{e}$. Another interesting observation is that the fixation probability θ is an annealed average of the densities of type-$\heartsuit$ individuals. Therefore, θ is a function of the average type-$\heartsuit$ densities determined by the initial distribution $\mu^{\mathfrak{e}}$ and does *not* depend on any other parameters of the distribution. ■

The proof of Theorem 4.2.9 relies on the analysis of the single-particle process $\Theta^{\mathfrak{e}}$ in Definition 4.2.3 in a random environment $\mathfrak{e}$ drawn from the law $\bar{\mathbb{P}}$. In particular, at the heart of the proof lies an exploitation of the following homogenisation result, whose proof is deferred to Section 4.3.3.

Theorem 4.2.11 (Homogenisation of environment). *Let $f_A : \mathcal{E}_{\mathfrak{K}} \to \mathbb{R}$ and $f_D : \mathcal{E}_{\mathfrak{K}} \to \mathbb{R}$ be two bounded Σ-measurable functions. Then, under Assumption* 2.A *and conditions* (1)–(2) *in Assumption* 4.A*, for $\bar{\mathbb{P}}$-a.s. every realisation of $\mathfrak{e}$ and any $\alpha \in \{0,1\}$,*

$$\lim_{t\to\infty} \sum_{(j,\beta)\in G} p_t^{\mathfrak{e}}((0,\alpha),(j,\beta))\big[\beta f_A(T_j\mathfrak{e}) + (1-\beta) f_D(T_j\mathfrak{e})\big] = \theta, \tag{4.23}$$

where $p_t^{\mathfrak{e}}(\cdot\,,\cdot)$ is the time-t transition kernel of the single-particle dual process $\Theta^{\mathfrak{e}}$ given in Definition 4.2.3*, and*

$$\theta := \frac{1}{1+\rho}\int_{\mathcal{E}_{\mathfrak{K}}} \big[f_A((N_k,M_k)_{k\in\mathbb{Z}^d}) + \tfrac{M_0}{N_0} f_D((N_k,M_k)_{k\in\mathbb{Z}^d})\big]\, \mathrm{d}\bar{\mathbb{P}}\{(N_k,M_k)_{k\in\mathbb{Z}^d}\}, \tag{4.24}$$

with $\rho := \bar{\mathbb{E}}\big[\frac{M_0}{N_0}\big] = \int_{\mathcal{E}_{\mathfrak{K}}} \frac{M_0}{N_0}\, \mathrm{d}\bar{\mathbb{P}}\{(N_k,M_k)_{k\in\mathbb{Z}^d}\}$.

The interpretation of the above result is that, for $\bar{\mathbb{P}}$-a.s. every realisation of the environment $\mathfrak{e}$, the law of the *"environment viewed from the particle"* in the process $\Theta^{\mathfrak{e}}$ converges weakly to an invariant distribution. The precise meaning of the last statement will become clear in Section 4.3. Conditions (1)–(2) in Assumption 4.A play a crucial role in the proof. Theorem 4.2.11 combined with Theorem 4.2.4 enable us to prove Theorem 4.2.9.

Note that, in (4.23), the process $\Theta^{\mathfrak{e}}$ is assumed to start at $(0,\alpha) \in G$. However, this does not matter, because the law of the environment is translation-invariant. Indeed, we have the following corollary:

Corollary 4.2.12. *Suppose that Assumption* 2.A *and conditions* (1)–(2) *in Assumption* 4.A *hold. Let f_A, f_D and θ be as in Theorem* 4.2.11*. Then, for $\bar{\mathbb{P}}$-a.s. every realisation of $\mathfrak{e}$ and all $(i,\alpha) \in \mathbb{Z}^d \times \{0,1\}$,*

$$\lim_{t\to\infty} \sum_{(j,\beta)\in G} p_t^{\mathfrak{e}}((i,\alpha),(j,\beta))\big[\beta f_A(T_j\mathfrak{e}) + (1-\beta) f_D(T_j\mathfrak{e})\big] = \theta, \tag{4.25}$$

where $p_t^{\mathfrak{e}}(\cdot\,,\cdot)$ is as in Definition 4.2.3.

§4.2.4 Discussion

Clustering in fixed environment. In Theorem 2.4.9 of Chapter 2 we only showed convergence of the spatial process $Z^{\mathfrak{e}}$ to an equilibrium for a restricted class of initial distributions, namely, a product of binomials with parameters that are tuned to the environment $\mathfrak{e}$ and the density of type-♡ individuals in the populations. The main result of Section 4.2.2, namely, Theorem 4.2.4, fully characterises the set of initial distributions for which $Z^{\mathfrak{e}}$ admits convergence to equilibrium. The result is valid for any admissible environment $\mathfrak{e}$ in which $Z^{\mathfrak{e}}$ exhibits clustering. The proof follows from similar arguments used in the proof of the analogous results [112, Theorem 1.9(b)] and [140, Theorem 1.2] derived, respectively, in the context of the Voter model and the Stepping Stone model (see also e.g. [27, 12]). In Theorem 2.4.12 of Chapter 2 we showed that the process $Z^{\mathfrak{e}}$ clusters if and only if two dual particles in $Z_*^{\mathfrak{e}}$ coalesce into a single particle with probability 1. We also show in Theorem 4.4.4 in Section 4.4.1 that coalescence of two dual particles with probability 1 is equivalent to coalescence of any finite number of dual particles with probability 1. This consistency property of the dual process, which is purely a consequence of the duality relation between $Z^{\mathfrak{e}}$ and $Z_*^{\mathfrak{e}}$, is far from trivial, because the dual particles interact with each other.

To summarise, the process $Z^{\mathfrak{e}}$ admits *only* mono-type equilibria if and only if the evolution of the dual $Z_*^{\mathfrak{e}}$ is eventually governed by $p_t^{\mathfrak{e}}(\cdot\,,\cdot)$, the probability transition kernel of the single-particle dual $\Theta^{\mathfrak{e}}$ (recall Definition 4.2.3). Precisely because of this, we see in (4.8) that the domain of attraction for each mono-type equilibrium of the process $Z^{\mathfrak{e}}$ in the clustering regime is dictated by the limiting behaviour of $p_t^{\mathfrak{e}}(\cdot\,,\cdot)$ as $t \to \infty$. On the contrary, if the process $Z^{\mathfrak{e}}$ is in the coexistence regime (= existence of multi-type equilibria), then the evolution of the dual $Z_*^{\mathfrak{e}}$ is no longer described by $p_t^{\mathfrak{e}}(\cdot\,,\cdot)$ alone, and therefore providing an answer to similar questions in the case of coexistence is challenging. In particular, because of the presence of interactions in the dual and the lack of translation-invariance of the state space $\mathcal{X}^{\mathfrak{e}}$, the characterization of the domain of attraction for a multi-type equilibrium via Liggett-type conditions (see e.g. [112, Theorem 1.9(a)],[76]) is a highly non-trivial problem, and is closely related to the study of harmonic functions (see e.g. [141]) of the general dual process $Z_*^{\mathfrak{e}}$.

Clustering in random environment. Turning to the main result of Section 4.2.3, we see that Theorem 4.2.9 is a homogenisation statement on the convergence of the spatial system to a mono-type equilibrium in random environment. It states that if the population sizes are drawn from an ergodic and translation-invariant random field for which clustering prevails, and the initial average densities of type-♡ active and dormant individuals in each colony are modulated, respectively, by two global functions $f_A(\cdot)$ and $f_D(\cdot)$ of the population sizes, then the spatial system converges in law towards a mono-type equilibrium for almost all initial realisations of the sizes. In the attained equilibrium, the probability of fixation to the all type-♡ configuration is a weighted average of the two functions f_A and f_D, and is independent of the chosen initial population sizes. In other words, the spatial process $Z^{\mathfrak{e}}$ undergoes homogenisation, which, roughly speaking, can be viewed as a “weak law of large numbers”.

A closer look at the proof in Section 4.4.2 will reveal that the homogenisation comes, in essence, from the duality relation with the process $\Theta^{\mathfrak{e}}$ evolving in the same random environment. The homogenisation in the continuous-time process $\Theta^{\mathfrak{e}}$, in turn, is inherited from a discrete-time subordinate Markov chain $\widehat{\Theta}^{\mathfrak{e}}$ (see Definition 4.3.1 in Section 4.3.1). This $\widehat{\Theta}^{\mathfrak{e}}$ is embedded into the continuous-time process $\Theta^{\mathfrak{e}}$ and closely resembles a d-dimensional version of the random walk in random environment (RWRE) on a strip introduced in [21] (see also [54, 53, 62] for similar models and further references). However, results derived in that context do not immediately carry over to our setting, because $\widehat{\Theta}^{\mathfrak{e}}$ fails to meet some basic irreducibility hypotheses (see e.g. [21, Condition C]). Nonetheless, it turns out that $\widehat{\Theta}^{\mathfrak{e}}$ is easier to analyse than the RWRE on a strip, as some of its transition probabilities are controlled by deterministic parameters that do not depend on the environment $\mathfrak{e}$. To be precise, the step distribution of a particle evolving via $\widehat{\Theta}^{\mathfrak{e}}$ on the d-dimensional strip $\mathbb{Z}^d \times \{0,1\}$ is a preassigned probability distribution $\hat{p}(\cdot)$ on $\mathbb{Z}^d$ and, in fact, is defined in terms of the migration kernel $a(\cdot\,,\cdot)$ of the spatial process $Z^{\mathfrak{e}}$. This simplicity of the subordinate Markov chain, which is similar to a property found in for random walk in random scenery (see e.g., [44, 49]), allows us to answer some of the highly sought-after questions in the literature on RWRE. In particular, we are able to identify a stationary and ergodic distribution for the environment viewed from the particle, with an explicit expression for the density w.r.t. the initial law, and establish a strong law of large numbers for the location of the particle (see Section 4.3.2). Moreover, when $\hat{p}(\cdot)$ is symmetric and recurrent ($d \leq 2$), we show that the environment process converges weakly to the *reversible* stationary distribution in the *quenched* setting. The latter is a very powerful result, which ultimately causes the homogenisation found in the subordinate Markov chain $\widehat{\Theta}^{\mathfrak{e}}$, and later passes it on to the single-particle dual $\Theta^{\mathfrak{e}}$ as well.

As argued before, the spatial process $Z^{\mathfrak{e}}$ acquires the homogenisation via duality from $\Theta^{\mathfrak{e}}$. Indeed, a crucial observation will reveal that the homogenised fixation probability in (4.19) is nothing but the average of the two global functions f_A and f_D w.r.t. the invariant distribution of the environment process. The method employed in proving the quenched weak convergence of the environment process for $\widehat{\Theta}^{\mathfrak{e}}$ to the invariant distribution is not probabilistic and relies on ergodic theoretic tools. To be precise, we first show that the *peripheral point-spectrum* (i.e., the set of all eigenvalues of modulus 1) of the self-adjoint Markov kernel operator $\mathfrak{R}$ associated to the environment process is trivial (see Lemma 4.3.12 in Section 4.3.2) and afterwards invoke a generalised version of the fundamental theorem for Markov chains (see Proposition 4.3.10 in Section 4.3.2) to establish the convergence. This way of proving weak convergence of the environment process is non-standard in the literature on RWRE, where such convergences are often established by exploiting some form of regeneration structure, or results like a local central limit theorem for the relevant random walk (see e.g., [95, 106, 54, 9]). Admittedly, the analysis of the peripheral point-spectrum of a Markov kernel operator in the L_p ($p \geq 1$) space of its reversible distribution is non-trivial and requires knowledge of the explicit form of the distribution. However, in many random environment models, such as the random conductance model, the one-dimensional RWRE, etc., important results in the quenched setting are still incomplete, despite knowledge of the explicit reversible distribution. Perhaps such problems may be approached in a similar way.

§4.3 Single-particle dual in random environment

As indicated in the previous section, the single-particle dual process $\Theta^{\mathfrak{e}}$ (see Definition 4.2.3) serves as the main ingredient in proofs of all our main results. In this section we study $\Theta^{\mathfrak{e}}$ in a typical random environment $\mathfrak{e} \in \mathcal{E}_{\mathfrak{K}}$ drawn according to the law $\bar{\mathbb{P}}$

(see Assumption 4.B) and prove the homogenisation result stated in Theorem 4.2.11.

To avoid dealing with technicalities that arise in the context of continuous-time Markov processes, in Section 4.3.1 we transform the process $\Theta^{\mathfrak{e}}$ into a discrete-time Markov chain $\widehat{\Theta}^{\mathfrak{e}}$ using the well-known method of *uniformisation* by a Poisson clock. We also introduce an *auxiliary environment process* W associated to the Markov chain $\widehat{\Theta}^{\mathfrak{e}}$. In Section 4.3.2 we show that the environment process W converges weakly to an invariant distribution in the *quenched* setting. Finally, in Section 4.3.3 we prove Theorem 4.2.11 and Corollary 4.2.12 by transferring the convergence result on W to the continuous-time process $\Theta^{\mathfrak{e}}$.

§4.3.1 Subordinate Markov chain and auxiliary environment process

When a continuous-time Markov process on a countable state space retains *uniformly bounded* jump rates, it can be uniformised by a Poisson clock and a discrete-time subordinate Markov chain (see e.g., [113, Chapter 2]). The method of uniformisation essentially transforms a *variable-speed* continuous-time Markov process into a *constant-speed* continuous-time Markov process [11]. Observe from (4.6) that the jump rates of $\Theta^{\mathfrak{e}}$ (see Definition 4.2.3) are uniformly bounded when the chosen environment $\mathfrak{e}$ is uniformly elliptic, and therefore $\Theta^{\mathfrak{e}}$ is uniformisable for such an environment. We start by defining a subordinate Markov chain $\widehat{\Theta}^{\mathfrak{e}}$ corresponding to the process $\Theta^{\mathfrak{e}}$ in a uniformly elliptic environment $\mathfrak{e}$.

Definition 4.3.1 (Subordinate Markov chain). The subordinate Markov chain (see Fig. 4.1)

$$\widehat{\Theta}^{\mathfrak{e}} := (\widehat{\Theta}^{\mathfrak{e}}_n)_{n\in\mathbb{N}_0}, \qquad \widehat{\Theta}^{\mathfrak{e}}_n = (X^{\mathfrak{e}}_n, \alpha^{\mathfrak{e}}_n), \tag{4.26}$$

in a uniformly elliptic environment $\mathfrak{e} := (N_i, M_i)_{i\in\mathbb{Z}^d} \in \mathcal{E}_{\mathfrak{K}}$ is the discrete-time Markov chain on the state space $G = \mathbb{Z}^d \times \{0,1\}$ with transition probabilities

$$\begin{aligned}
(i,1) &\longrightarrow \begin{cases} (j,1) & \text{w.p. } (1-q_s)\hat{p}(j-i), \quad j\in\mathbb{Z}^d,\\ (i,0) & \text{w.p. } q_s, \end{cases}\\
(i,0) &\longrightarrow \begin{cases} (i,0) & \text{w.p. } 1-\omega(i),\\ (i,1) & \text{w.p. } \omega(i), \end{cases}
\end{aligned} \tag{4.27}$$

where $i \in \mathbb{Z}^d$, and the parameters q_s, $\omega := (\omega(k))_{k\in\mathbb{Z}^d}$ and $\hat{p} := (\hat{p}(k))_{k\in\mathbb{Z}^d}$ are determined by the exchange rate λ, the environment $\mathfrak{e}$, the migration kernel $a(\cdot\,,\cdot)$, and the ellipticity constant $\mathfrak{K} \geq 2$, as follows:

$$\begin{aligned}
q_s &:= \frac{\lambda}{c+\lambda+\lambda\mathfrak{K}}, \quad \omega(i) := \frac{\lambda K_i}{c+\lambda+\lambda\mathfrak{K}} = \frac{\lambda N_i}{M_i(c+\lambda+\lambda\mathfrak{K})},\\
\hat{p}(i) &:= \frac{\lambda\mathfrak{K}}{c+\lambda\mathfrak{K}}\mathbb{1}_{\{i=0\}} + \frac{a(0,i)}{c+\lambda\mathfrak{K}}\mathbb{1}_{\{i\neq 0\}},
\end{aligned} \qquad i\in\mathbb{Z}^d, \tag{4.28}$$

where c is the speed of migration defined in condition (3) of Assumption 2.A. We denote by $Q_{\mathfrak{e}}(\cdot\,,\cdot) : G\times G \to [0,1]$ the 1-step transition kernel of the chain $\widehat{\Theta}^{\mathfrak{e}}$, defined

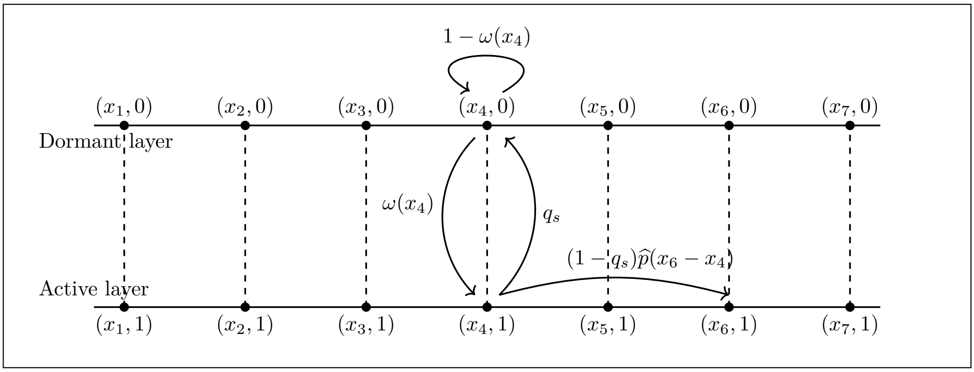

Figure 4.1: A schematic representation of the transition probabilities of a particle moving on the d-dimensional strip $\mathbb{Z}^d \times \{0,1\}$ according to $\widehat{\Theta}^{\mathfrak{e}}$. The particle is allowed to migrate in the bottom layer and while doing so remains in active *state. However, the particle becomes* dormant *by entering the top layer, and thus can not migrate.*

as

$$Q_{\mathfrak{e}}(\eta,\xi) := \widehat{P}^{\mathfrak{e}}_{\eta}(\widehat{\Theta}^{\mathfrak{e}}_1 = \xi), \quad \eta,\xi \in G, \tag{4.29}$$

where $\widehat{P}^{\mathfrak{e}}_{\eta}$ is the canonical law of $\widehat{\Theta}^{\mathfrak{e}}$ started at η. ■

Remark 4.3.2 (Well-posedness). Observe that $\hat{p}(\cdot)$ defines a probability distribution on $\mathbb{Z}^d$ and inherits the role of the migration kernel $a(0,\cdot)$. By the uniform ellipticity of the environment $\mathfrak{e} \in \mathcal{E}_{\mathfrak{K}}$, it follows that $\omega \in [\delta, 1-\delta]^{\mathbb{Z}^d}$ for some $\delta \in (0, \frac{1}{2})$ determined by c, λ and $\mathfrak{K}$. Thus, the transition probabilities in (4.27) are well-defined. From (4.28) we see that ω is the only parameter that depends on $\mathfrak{e}$ and plays the role of *random* environment for $\widehat{\Theta}^{\mathfrak{e}}$, while q_s takes over the role of λ, which is the rate of becoming dormant from the active state in the continuous-time process $\Theta^{\mathfrak{e}}$.

The subordinate Markov chain $\widehat{\Theta}^{\mathfrak{e}}$ describes the evolution of a particle moving on the d-dimensional strip $\mathbb{Z}^d \times \{0,1\}$ in discrete time. The coordinates $X^{\mathfrak{e}}_n$ and $\alpha^{\mathfrak{e}}_n$ give, respectively, the location in $\mathbb{Z}^{\mathrm{d}}$ and the state (active or dormant) at time $n \in \mathbb{N}_0$ of the particle evolving in the environment $\mathfrak{e}$ according to the transition probabilities given in (4.27). In each step, the particle in the active state, with probability $(1-q_s)$, performs random walk on $\mathbb{Z}^d$ according to the increment distribution $\hat{p}(\cdot)$, while, with probability q_s, it becomes dormant from the active state. The particle does not move in the dormant state and becomes active with a location-dependent probability determined by the environment $\mathfrak{e}$. The following property of the law of $\widehat{\Theta}^{\mathfrak{e}}$ is a consequence of the translation-invariance of $\mathbb{Z}^d$ and the migration kernel $a(\cdot\,,\cdot)$. The proof follows from an easy calculation of the transition probabilities of $\widehat{\Theta}^{\mathfrak{e}}$ given in (4.27), and is omitted for briefness.

Lemma 4.3.3 (Translation-invariance). *For any $(i,\alpha),(j,\beta) \in G$ and $n \in \mathbb{N}_0$,*

$$\widehat{P}^{\mathfrak{e}}_{(0,\alpha)}(\widehat{\Theta}^{\mathfrak{e}}_n = (j,\beta)) = \widehat{P}^{T_{-i}\mathfrak{e}}_{(i,\alpha)}(\widehat{\Theta}^{T_{-i}\mathfrak{e}}_n = (i+j,\beta)). \tag{4.30}$$

The connection between the discrete-time Markov chain $\widehat{\Theta}^{\mathfrak{e}}$ and the continuous-time Markov process $\Theta^{\mathfrak{e}}$ becomes apparent in the next lemma.

Lemma 4.3.4 (Uniformisation by Poisson clock). *Let $\mathfrak{e} \in \mathcal{E}_{\mathfrak{K}}$ be a uniformly elliptic environment and $(N_t)_{t\geq 0}$ be a Poisson process with rate $c+\lambda+\lambda\mathfrak{K}$ that is independent of the subordinate Markov chain $\widehat{\Theta}^{\mathfrak{e}}$. Then, under the assumption that the process $\Theta^{\mathfrak{e}}$ (see Definition* 4.2.3*) and the Markov chain $\widehat{\Theta}^{\mathfrak{e}}$ have the same initial distribution,*

$$(\Theta^{\mathfrak{e}}(t))_{t\geq 0} \overset{d}{=} (\widehat{\Theta}^{\mathfrak{e}}_{N_t})_{t\geq 0}. \tag{4.31}$$

In particular, for $\eta, \xi \in G$,

$$p^{\mathfrak{e}}_t(\eta,\xi) = \mathrm{e}^{-(c+\lambda+\lambda\mathfrak{K})t} \sum_{n=0}^{\infty} \tfrac{[(c+\lambda+\lambda\mathfrak{K})t]^n}{n!} Q^n_{\mathfrak{e}}(\eta,\xi), \tag{4.32}$$

where $p^{\mathfrak{e}}_t(\cdot\,,\cdot)$ and $Q_{\mathfrak{e}}(\cdot\,,\cdot)$ are as in Definition 4.2.3 *and Definition* 4.3.1*, respectively.*

Proof. Let $\mathcal{J}_{\mathfrak{e}}$ denote the infinitesimal generator of the process $\Theta^{\mathfrak{e}}$. The action of $\mathcal{J}_{\mathfrak{e}}$ on a bounded function $f \in \mathcal{F}_b(G)$ is given by

$$(\mathcal{J}_{\mathfrak{e}} f)(i,\alpha) = \begin{cases} \lambda[f(i,0)-f(i,1)] + \sum_{j\in\mathbb{Z}^d} a(i,j)[f(j,1)-f(i,1)], & \text{if } \alpha = 1,\\ \lambda K_i[f(i,1)-f(i,0)], & \text{if } \alpha = 0, \end{cases} \tag{4.33}$$

where $(i,\alpha) \in G$. Since $\mathfrak{e}$ is uniformly elliptic and the total speed of migration given by c is finite by virtue of Assumption 2.A, it is easily seen that $\mathcal{J}_{\mathfrak{e}}$ is a bounded operator. Thus $(\exp\{\mathcal{J}_{\mathfrak{e}} t\})_{t\geq 0}$ defines the semigroup of $\Theta^{\mathfrak{e}}$. In particular, the transition probability kernel $p^{\mathfrak{e}}_t(\cdot\,,\cdot)$ expands as

$$p^{\mathfrak{e}}_t(\cdot\,,\cdot) = \sum_{n=0}^{\infty} \mathcal{J}^n_{\mathfrak{e}}(\cdot\,,\cdot)\tfrac{t^n}{n!}, \tag{4.34}$$

where the generator $\mathcal{J}_{\mathfrak{e}}$ is viewed as a matrix. The claim follows from this expansion of $p^{\mathfrak{e}}_t(\cdot\,,\cdot)$ and the observation that

$$\mathcal{J}_{\mathfrak{e}} = (c+\lambda+\lambda\mathfrak{K})[Q_{\mathfrak{e}} - I], \tag{4.35}$$

where I is the identity operator (viewed as a matrix). Note that in (4.35) the translation-invariance of the migration kernel $a(\cdot\,,\cdot)$ is used. □

Below we define the "environment process" associated to the subordinate Markov chain $\widehat{\Theta}^{\mathfrak{e}}$. This process is defined in the same way as for RWRE on a strip (see e.g., [53, Definition 2.2]).

Definition 4.3.5 (Auxiliary environment process). Let $\widehat{\Theta}^{\mathfrak{e}} = (X^{\mathfrak{e}}_n, \alpha^{\mathfrak{e}}_n)_{n\in\mathbb{N}_0}$ with the canonical law $\widehat{P}^{\mathfrak{e}}_{(0,\alpha)}$ be the subordinate Markov chain (see Definition 4.3.1) started at $(0,\alpha) \in G$ in environment $\mathfrak{e} \in \mathcal{E}_{\mathfrak{K}}$. The auxiliary environment process W having initial distribution $\delta_{(\mathfrak{e},\alpha)}$ is the discrete-time process on $\Omega_{\mathfrak{K}} := \mathcal{E}_{\mathfrak{K}} \times \{0,1\}$ given by

$$W := (W_n)_{n\in\mathbb{N}_0}, \quad W_n := (\mathfrak{e}_n, \alpha_n) \text{ with } \mathfrak{e}_n := T_{X^{\mathfrak{e}}_n}\mathfrak{e},\ \alpha_n := \alpha^{\mathfrak{e}}_n, \tag{4.36}$$

and is defined on the same probability space of $\widehat{\Theta}^{\mathfrak{e}}$. ■

It is trivial to check that, for any $(\mathfrak{e},\alpha) \in \Omega_{\mathfrak{K}}$, W is a Markov chain on the state space $\Omega_{\mathfrak{K}}$ under the law $\widehat{P}^{\mathfrak{e}}_{(0,\alpha)}$, with initial distribution $\delta_{(\mathfrak{e},\alpha)}$ (by Lemma 4.3.3, also under the law $\widehat{P}^{\mathfrak{e}}_{(i,\alpha)}$, $i \in \mathbb{Z}^d$, with initial distribution $\delta_{(T_i\mathfrak{e},\alpha)}$).

The action of the Markov kernel operator $\mathfrak{R}$ associated to W on a bounded function $f \in \mathcal{F}_b(\Omega_{\mathfrak{K}})$ is given by

$$\mathfrak{R}f(\mathfrak{e},\alpha) := \widehat{E}^{\mathfrak{e}}_{(0,\alpha)}[f(W_1)] = \sum_{(j,\beta)\in G} Q_{\mathfrak{e}}((0,\alpha),(j,\beta))f(T_j\mathfrak{e},\beta), \tag{4.37}$$

where $(\mathfrak{e},\alpha) \in \Omega_{\mathfrak{K}}$ and $Q_{\mathfrak{e}}(\cdot\,,\cdot)$ is the 1-step transition kernel of $\widehat{\Theta}^{\mathfrak{e}}$ defined in (4.29). In particular,

$$\mathfrak{R}f(\mathfrak{e},\alpha) = \begin{cases} q_s\, f(\mathfrak{e},0) + (1-q_s)\displaystyle\sum_{j\in\mathbb{Z}^d} \hat{p}(j) f(T_j\mathfrak{e},1), & \text{if } \alpha = 1, \\ \omega(0)f(\mathfrak{e},1) + [1-\omega(0)]f(\mathfrak{e},0), & \text{if } \alpha = 0, \end{cases} \tag{4.38}$$

where q_s, $\hat{p}(\cdot)$ and $\omega := (\omega(k))_{k\in\mathbb{Z}^d}$ are defined in terms of $\mathfrak{e}$ and the other parameters in (4.28).

The Markov chain W describes the state of the environment from the point of view of a particle that moves on the d-dimensional strip $\mathbb{Z}^d \times \{0,1\}$ according to the chain $\widehat{\Theta}^{\mathfrak{e}}$. The definition of the process differs from the standard definition usually encountered in the literature on RWRE. This is because the particle moves on two copies of $\mathbb{Z}^d$ instead of one, and in order to preserve the Markov property we need an extra variable describing the layer on which the particle is present.

The state space $\Omega_{\mathfrak{K}}$ of the auxiliary environment process W, even though compact, is huge. Thus, at first glance, obtaining any useful information from W might seem to be an impossible task. In general, this difficulty is overcome by taking initial samples of the environment from an ergodic and translation-invariant law. In such settings, it often becomes possible to construct "by hand" an invariant distribution that is absolutely continuous w.r.t. the initial law. Invariant distributions having such characteristics, which guarantees its uniqueness as well (see e.g. [22, 100]), are an extremely powerful tool for deriving many interesting properties, such as laws of large numbers, central limit theorems etc., for the relevant process. In the next section we find an invariant distribution $\mathbb{Q}$ with such a property and prove weak convergence of W to the invariant distribution in the *quenched* setting.

§4.3.2 Stationary environment process and weak convergence

In this section we address the question of whether the auxiliary environment process W admits an invariant distribution that is "equivalent" to its initial distribution. The following result provides a positive answer:

Theorem 4.3.6 (Invariant distribution of environment process). *Let* $\mathbb{Q}$ *be the probability measure on* $(\Omega_{\mathfrak{K}}, \Sigma \otimes 2^{\{0,1\}})$ *defined by*

$$\mathrm{d}\mathbb{Q}\{(\mathfrak{e},\alpha)\} := \frac{u(\mathfrak{e},\alpha)}{1+\rho}\,\mathrm{d}\bar{\mathbb{P}}\{\mathfrak{e}\}, \tag{4.39}$$

where the law $\bar{\mathbb{P}}$ *defined on* $(\mathcal{E}_{\mathfrak{K}}, \Sigma)$ *is as in Assumption* 4.B, $\rho := \bar{\mathbb{E}}\big[\frac{M_0}{N_0}\big]$, *and the density* $u : \Omega_{\mathfrak{K}} \to (0, \mathfrak{K}]$ *is given by*

$$u((N_k, M_k)_{k\in\mathbb{Z}^d}, \alpha) = \begin{cases} 1 & \text{if } \alpha = 1, \\ \frac{M_0}{N_0} & \text{if } \alpha = 0. \end{cases} \tag{4.40}$$

The following hold:

(1) *The environment process* W *in Definition* 4.3.5 *is stationary and ergodic under the probability law* $\mathbb{Q}$.

(2) *Under condition* (1) *in Assumption* 4.A, $\mathbb{Q}$ *is reversible.*

Remark 4.3.7 (Validity in all dimensions). Part (1) of Theorem 4.3.6 holds without the imposition of condition (1) in Assumption 4.A. It essentially follows from the translation-invariance and ergodicity of the law $\bar{\mathbb{P}}$. Moreover, both part (1) and part (2) are valid in all dimensions $d \geq 1$. Assumption 2.A is crucial for the proof and can not be removed in a straightforward way.

The proof of Theorem 4.3.6 is mostly computational and is deferred to Appendix B.1. As an application of this result, in Appendix B.1 we also give a proof of strong law of large numbers for the subordinate Markov chain $\widehat{\Theta}^{\mathfrak{e}}$ (recall Definition 4.3.1), which is a result of independent interest.

Before we proceed further, let us explain what we mean by "equivalence" of the invariant distribution $\mathbb{Q}$ in the theorem and the initial law $\bar{\mathbb{P}}$ of the environment. In the literature on RWRE, this phenomenon is called "equivalence between the static and the dynamic points of view".

Lemma 4.3.8 (Equivalence of $\mathbb{Q}$ and $\bar{\mathbb{P}}$). *Let* $\mathbb{Q}, \bar{\mathbb{P}}$ *be as in Theorem* 4.3.6. *Then, for any measurable* $A \subseteq \Omega_{\mathfrak{K}} = \mathcal{E}_{\mathfrak{K}} \times \{0,1\}$, *the following are equivalent:*

(1) $\mathbb{Q}(A) = 1$.

(2) *There exists a* Σ*-measurable* $A' \subseteq \mathcal{E}_{\mathfrak{K}}$ *such that* $\bar{\mathbb{P}}(A') = 1$ *and* $A' \times \{0,1\} \subseteq A$.

Proof. Let $\theta := \frac{1}{1+\bar{\mathbb{E}}[M_0/N_0]} \in (0,1)$, and let μ be the probability measure on $(\mathcal{E}_{\mathfrak{K}}, \Sigma)$ defined by

$$\mu(E) = \tfrac{\theta}{1-\theta} \int_E \tfrac{M_0}{N_0}\,\mathrm{d}\bar{\mathbb{P}}\{(N_k, M_k)_{k\in\mathbb{Z}^d}\}, \quad E \in \Sigma. \tag{4.41}$$

Clearly, for any $E \in \Sigma$,

$$\mu(E) = 1 \text{ if and only if } \bar{\mathbb{P}}(E) = 1. \tag{4.42}$$

Suppose that (1) holds for some measurable $A \subseteq \Omega_{\mathfrak{K}}$. Note from (4.39) that

$$1 = \mathbb{Q}(A) = \theta\,\bar{\mathbb{P}}(A_1) + (1-\theta)\mu(A_0), \tag{4.43}$$

where

$$A_0 := \{\mathfrak{e} \,:\, (\mathfrak{e}, 0) \in A\}, \quad A_1 := \{\mathfrak{e} \,:\, (\mathfrak{e}, 1) \in A\}. \tag{4.44}$$

Since $\theta \in (0,1)$, this implies $\bar{\mathbb{P}}(A_1) = \mu(A_0) = 1$. Defining $A' = A_0 \cap A_1$, we see that (2) follows from (4.42).

Similarly, if (2) holds, then by (4.42), $\mathbb{Q}(A' \times \{0,1\}) = \theta\,\bar{\mathbb{P}}(A') + (1-\theta)\mu(A') = 1$. Thus, $\mathbb{Q}(A) \geq \mathbb{Q}(A' \times \{0,1\}) = 1$ and so (1) is proved. □

Our next goal is to prove weak convergence of the environment process W to the stationary law $\mathbb{Q}$ under the quenched law $\widehat{P}^{\mathfrak{e}}_{(0,\alpha)}$ for $\bar{\mathbb{P}}$-a.s. every realisation of the environment $\mathfrak{e} \in \mathcal{E}_{\mathfrak{K}}$. In particular, we have the following result:

Theorem 4.3.9 (Weak convergence of auxiliary environment). *Suppose that conditions* (1)–(2) *in Assumption* 4.A *hold. Let $f_A : \mathcal{E}_{\mathfrak{K}} \to \mathbb{R}$ and $f_D : \mathcal{E}_{\mathfrak{K}} \to \mathbb{R}$ be two bounded Σ-measurable functions. Then, for $\bar{\mathbb{P}}$-a.s. every realisation of $\mathfrak{e} \in \mathcal{E}_{\mathfrak{K}}$ and any $\alpha \in \{0,1\}$,*

$$\lim_{n\to\infty} \widehat{E}^{\mathfrak{e}}_{(0,\alpha)}[h(\mathfrak{e}_n, \alpha_n)] = \int_{\mathcal{E}_{\mathfrak{K}} \times \{0,1\}} h(\mathfrak{e}', \beta)\,\mathrm{d}\mathbb{Q}(\mathfrak{e}', \beta), \tag{4.45}$$

where h is the function $(\mathfrak{e}, \alpha) \mapsto \alpha f_A(\mathfrak{e}) + (1-\alpha) f_D(\mathfrak{e})$, $W = (\mathfrak{e}_n, \alpha_n)_{n\in\mathbb{N}_0}$ is the auxiliary environment process with law $\widehat{P}^{\mathfrak{e}}_{(0,\alpha)}$ defined in Definition 4.3.5*, and $\mathbb{Q}$ is the stationary law of W given in* (4.39)*.*

The proof of Theorem 4.3.9 is a consequence of the proposition stated below. This proposition is an analogue of the "fundamental theorem of Markov chains on countable state spaces" because it addresses Markov chains on general state spaces. We believe that this result is already known in the literature (see e.g., [114] or [23, 89, 35]) on ergodic theory on Markov chains, but we have been unable to find a reference with an explicit proof of the statement. For the sake of completeness, the proof is given in Appendix B.2.

Proposition 4.3.10 (Fundamental theorem of MC). *Let $(\Omega, \Sigma, \mathbb{Q})$ be a probability space, where the σ-field Σ is countably generated. Let $W := (W_n)_{n\in\mathbb{N}_0}$ be a Markov chain on the state space Ω, and assume that $\mathbb{Q}$ is a reversible and ergodic stationary distribution for W. If -1 is not an eigenvalue of the Markov kernel operator $\mathfrak{R} : L_\infty(\Omega, \mathbb{Q}) \to L_\infty(\Omega, \mathbb{Q})$ associated to W, then for every bounded measurable function $f \in \mathcal{F}_b(\Omega)$ and $\mathbb{Q}$-a.s. every $w \in \Omega$,*

$$\lim_{n\to\infty} \mathbb{E}_w[f(W_n)] = \int_\Omega f\,\mathrm{d}\mathbb{Q}, \tag{4.46}$$

where the expectation on the left is taken w.r.t. the law of W started at w.

Remark 4.3.11 (Convergence in total variation). The above proposition only establishes weak convergence and gives no information on the rate of convergence in (4.46). Under more stringent classical conditions on W, such as Harris recurrence or a

Doeblin criterion (see e.g., [121, 129] and [135, 104] for further references), uniqueness of the law $\mathbb{Q}$ holds and the chain converges in total variation norm from *all* initial starting points. The existence of a *spectral gap* of the operator $\mathfrak{R}$ results in *geometric ergodicity*, where the convergence takes place at an exponential rate (see e.g., [98]). However, under the assumption of only aperiodicity and ϕ-irreducibility of the Markov chain W, convergence in total variation holds only for $\mathbb{Q}$-a.s. all initial points.

Although in the above remark we discuss convergence of a Markov chain in total variation norm, the reader should not hope for such a strong convergence of the auxiliary environment process W given in Definition 4.3.5. Indeed, the process W is a highly "singular" Markov chain living on a huge state space $\Omega_{\mathfrak{K}}$ and admits infinitely many invariant distributions (e.g., take $\bar{\mathbb{P}} = \delta_{\mathfrak{e}}$, where $\mathfrak{e} = (N, M)_{i\in\mathbb{Z}^d}$ is a translation-invariant environment with $(N, M) \in \mathbb{N}^2$, and construct $\mathbb{Q}$ by (4.39)). Thus, it is very unlikely for W to be Harris recurrent, or to satisfy Doeblin-type conditions for that matter.

Proof of Theorem 4.3.9. By condition (1) of Assumption 4.A and Theorem 4.3.6, we see that $\mathbb{Q}$ is a reversible and ergodic distribution for the auxiliary environment process W. Observe from Proposition 4.3.10, if we are able to prove that -1 is not an eigenvalue of the Markov kernel operator $\mathfrak{R} : L_\infty(\Omega_{\mathfrak{K}}, \mathbb{Q}) \to L_\infty(\Omega_{\mathfrak{K}}, \mathbb{Q})$ given in (4.38), then we can find a measurable $E \subseteq \Omega_{\mathfrak{K}}$ such that $\mathbb{Q}(E) = 1$ and, for all $(\mathfrak{e}, \alpha) \in E$, (4.45) holds for the function h. In particular, using Lemma 4.3.8 we can find a measurable $E' \subset \mathcal{E}_{\mathfrak{K}}$ with $\bar{\mathbb{P}}(E') = 1$ and (4.45) holds for all $(\mathfrak{e}, \alpha) \in E' \times \{0, 1\}$. Thus, the proof is complete once we show that -1 is not an eigenvalue of $\mathfrak{R}$ when viewed as an operator on $L_\infty(\Omega_{\mathfrak{K}}, \mathbb{Q})$. We prove this in Lemma 4.3.12 stated below. □

Lemma 4.3.12 (Trivial peripheral point-spectrum). *Let $\mathfrak{R}$ be the Markov kernel operator (see* (4.38)*) of the auxiliary environment process W, and $\mathbb{Q}$ be the invariant distribution of W given in Theorem* 4.3.6*. If condition* (2) *in Assumption* 4.A *holds, then -1 is not an eigenvalue of the kernel operator $\mathfrak{R} : L_\infty(\Omega_{\mathfrak{K}}, \mathbb{Q}) \to L_\infty(\Omega_{\mathfrak{K}}, \mathbb{Q})$.*

Proof. Let $g \in L_\infty(\Omega_{\mathfrak{K}}, \mathbb{Q})$ be such that

$$\mathfrak{R}g = -g \quad \mathbb{Q}\text{-a.s.} \tag{4.47}$$

We show $g = 0$ a.s. As we will see below, this will follow from condition (2) in Assumption 4.A, which ensures that the increment distribution $\hat{p}(\cdot)$ defined in terms of $a(\cdot\,,\cdot)$ in (4.28) does not admit any non-constant and nonnegative bounded subharmonic function. With this aim, let $A \subseteq \Omega_{\mathfrak{K}}$ be measurable with $\mathbb{Q}(A) = 1$ and such that (4.47) holds for all $(\mathfrak{e}, \alpha) \in A$. Without loss of generality, we can also assume that

$$|g(\mathfrak{e}, \alpha)| \le \|g\|_\infty \quad \forall\ (\mathfrak{e}, \alpha) \in A. \tag{4.48}$$

By Lemma 4.3.8, there exists a measurable $A' \subseteq \mathcal{E}_{\mathfrak{K}}$ such that $\bar{\mathbb{P}}(A') = 1$ and (4.47) holds for all $(\mathfrak{e}, \alpha) \in A' \times \{0, 1\} \subseteq A$. Using (4.38), we compute $\mathfrak{R}g$ and obtain from (4.47) that

$$\begin{aligned} g(\mathfrak{e}, 0) &= -\big[\omega(0)g(\mathfrak{e}, 1) + (1 - \omega(0))g(\mathfrak{e}, 0)\big], \\ g(\mathfrak{e}, 1) &= -\Big[q_s\, g(\mathfrak{e}, 0) + (1 - q_s) \sum_{j\in\mathbb{Z}^d} \hat{p}(j)g(T_j\mathfrak{e}, 1)\Big], \quad \mathfrak{e} \in A', \end{aligned} \tag{4.49}$$

where, as before, $\omega, \hat{p}$ and q_s are defined by (4.28) in terms of $\mathfrak{e}$ and the other parameters. Now, using the translation invariance of $\bar{\mathbb{P}}$, we also have

$$\bar{\mathbb{P}}(B_{\text{inv}}) = 1, \quad B_{\text{inv}} := \bigcap_{j\in\mathbb{Z}^d} T_j^{-1}(A') \subseteq A', \tag{4.50}$$

where, trivially, B_{inv} is a translation-invariant set. We get from (4.49) that

$$\begin{aligned} g(\mathfrak{e},0) &= -\frac{\omega(0)}{2-\omega(0)} g(\mathfrak{e},1), \\ \sum_{j\in\mathbb{Z}^d} \hat{p}(j) g(T_j\mathfrak{e},1) &= -\Big[\frac{2-(1+q_s)\omega(0)}{(2-\omega(0))(1-q_s)}\Big] g(\mathfrak{e},1), \end{aligned} \tag{4.51}$$

for all $\mathfrak{e} \in B_{\text{inv}}$. By ellipticity (see Definition 4.2.7) of $\mathfrak{e} \in B_{\text{inv}}$, we can find a $\delta \in (0, \frac{1}{2})$ such that $\delta < \omega(0) < 1-\delta$ for all $\omega = (\omega(k))_{k\in\mathbb{Z}^d}$ determined by $\mathfrak{e} \in B_{\text{inv}}$. In particular, setting

$$C := \frac{1}{1-q_s}\Big[1 - \tfrac{1-\delta}{1+\delta}\, q_s\Big], \tag{4.52}$$

we see that

$$\frac{2-(1+q_s)\omega(0)}{(2-\omega(0))(1-q_s)} \geq C, \tag{4.53}$$

and also $C > 1$ as $\delta \in (0, \frac{1}{2})$. Combining the above with (4.51), we have

$$\Big|\sum_{j\in\mathbb{Z}^d} \hat{p}(j) g(T_j\mathfrak{e},1)\Big| = \Big|\tfrac{2-(1+q_s)\omega(0)}{(2-\omega(0))(1-q_s)}\Big| |g(\mathfrak{e},1)| \geq C|g(\mathfrak{e},1)|, \quad \mathfrak{e} \in B_{\text{inv}}. \tag{4.54}$$

Using the triangle inequality, we get

$$\sum_{j\in\mathbb{Z}^d} \hat{p}(j) |g(T_j\mathfrak{e},1)| \geq C|g(\mathfrak{e},1)|, \quad \mathfrak{e} \in B_{\text{inv}}. \tag{4.55}$$

Because B_{inv} is translation-invariant, the above implies that for any $\mathfrak{e} \in B_{\text{inv}}$ and all $i \in \mathbb{Z}^d$,

$$\sum_{j\in\mathbb{Z}^d} \hat{p}(j) |g(T_{i+j}\mathfrak{e},1)| \geq C|g(T_i\mathfrak{e},1)|. \tag{4.56}$$

Since $C > 1$, the above equation tells that, for a fixed $\mathfrak{e} \in B_{\text{inv}}$, the map $i \mapsto |g(T_i\mathfrak{e},1)|$ is a bounded (recall (4.48)) non-negative subharmonic function for $\hat{p}(\cdot)$. Now, by condition (2) in Assumption 4.A, a random walk on $\mathbb{Z}^d$ with increment distribution $\hat{p}(\cdot)$ defined as in (4.27) is irreducible and recurrent (see e.g., [107, Chapter 4]). Therefore, any bounded nonnegative subharmonic function of $\hat{p}(\cdot)$ on $\mathbb{Z}^d (d \leq 2)$ must be a constant (by an application of Doob's submartingale convergence theorem). In particular, for any $\mathfrak{e} \in B_{\text{inv}}$ and all $i \in \mathbb{Z}^d$,

$$|g(T_i\mathfrak{e},1)| = |g(\mathfrak{e},1)|. \tag{4.57}$$

Since $C > 1$, the only way in which (4.56) complies with (4.57), is when $|g(\mathfrak{e},1)| = 0$, so (4.51) implies that $g(\mathfrak{e},0) = 0$ as well. Thus, $g = 0$ on $B_{\text{inv}} \times \{0,1\}$ and, since $\bar{\mathbb{P}}(B_{\text{inv}}) = 1$, we see by Lemma 4.3.8 that $\mathbb{Q}(B_{\text{inv}} \times \{0,1\}) = 1$. □

Remark 4.3.13 (Peripheral point-spectrum in L_1). Using [89, Lemma 2], we can actually show that -1 is not an eigenvalue of $\mathfrak{R}$ in $L_1(\Omega_{\mathfrak{K}}, \mathbb{Q})$ as well. But convergence of $\mathfrak{R}^{2n} f$ may fail as $n \to \infty$, when it is merely assumed that $f \in L_1(\Omega_{\mathfrak{K}}, \mathbb{Q})$ (see e.g., [131]), and therefore Proposition 4.3.10 does not hold in general for such f.

§4.3.3 Transference of convergence: discrete to continuous

In this section we prove Theorem 4.2.11 and Corollary 4.2.12 by utilising the results derived in the Section 4.3.2.

Before we start with the proof of Theorem 4.2.11, let us briefly elaborate on its statement. In Section 4.3.1 we introduced in Definition 4.3.5 the discrete-time auxiliary environment process W associated to the subordinate Markov chain $\widehat{\Theta}^{\mathfrak{e}}$. We can also, in a similar fashion, extend the definition of W to construct a continuous-time environment process $w := (w_t)_{t\geq 0}$ for the single-particle dual $\Theta^{\mathfrak{e}}$ (recall Definition 4.2.3). Indeed, we obtain the process w by simply putting

$$w_t := (\mathfrak{e}_t, \alpha_t) \text{ with } \mathfrak{e}_t := T_{x_t^{\mathfrak{e}}}\,\mathfrak{e},\ \ \alpha_t := \alpha_t^{\mathfrak{e}}, \tag{4.58}$$

for each $t \geq 0$, where $\Theta^{\mathfrak{e}} = (x_t^{\mathfrak{e}}, \alpha_t^{\mathfrak{e}})_{t\geq 0}$ is as in Definition 4.2.3. Upon closer inspection of (4.10) and the definition of w, we see that Theorem 4.2.11 basically states that

$$\lim_{t\to\infty} E^{\mathfrak{e}}_{(0,\alpha)}[\alpha_t f_A(\mathfrak{e}_t) + (1-\alpha_t) f_D(\mathfrak{e}_t)] = \theta \tag{4.59}$$

for $\bar{\mathbb{P}}$-a.s. every realisation of the environment $\mathfrak{e}$, where f_A, f_D and θ are as in the theorem. In other words, (4.59) is equivalent to saying that the process w converges in distribution to the law $\mathbb{Q}$ given in (4.39) for $\bar{\mathbb{P}}$-a.s. every realisation of $\mathfrak{e} \in \mathcal{E}_{\mathfrak{K}}$ and any $\alpha \in \{0, 1\}$.

Proof of Theorem 4.2.11. From Lemma 4.3.4, we observe that

$$p_t^{\mathfrak{e}}((0,\alpha),(j,\beta)) = \sum_{n=0}^{\infty} \widehat{P}^{\mathfrak{e}}_{(0,\alpha)}(\widehat{\Theta}^{\mathfrak{e}}_n = (j,\beta))\,\mathbb{P}(N_t = n), \quad (j,\beta)\in G,\ \mathfrak{e}\in\mathcal{E}_{\mathfrak{K}},\ t\geq 0, \tag{4.60}$$

where $p_t^{\mathfrak{e}}(\cdot\,,\cdot)$ is as in Definition 4.2.3, $\widehat{\Theta}^{\mathfrak{e}} = (\widehat{\Theta}^{\mathfrak{e}}_n)_{n\in\mathbb{N}_0}$ is the subordinate Markov chain with law $\widehat{P}^{\mathfrak{e}}_{(0,\alpha)}$ (see Definition 4.3.1) and $(N_t)_{t\geq 0}$ is the Poisson process mentioned in the lemma, which is independent of $\widehat{\Theta}^{\mathfrak{e}}$. Thus, using the above, the left-hand side of (4.23), which we abbreviate by $l((\mathfrak{e},\alpha),t)$ for any $t \geq 0$, can be written as

$$\begin{aligned} l((\mathfrak{e},\alpha),t) &= \sum_{(j,\beta)\in G}\Big[\sum_{n\in\mathbb{N}_0} \widehat{P}^{\mathfrak{e}}_{(0,\alpha)}(\widehat{\Theta}^{\mathfrak{e}}_n = (j,\beta))\,\mathbb{P}(N_t = n)\Big]\big\{\beta f_A(T_j\mathfrak{e}) + (1-\beta) f_D(T_j\mathfrak{e})\big\} \\ &= \sum_{n\in\mathbb{N}_0}\Big[\sum_{(j,\beta)\in G} \widehat{P}^{\mathfrak{e}}_{(0,\alpha)}(W_n = (T_j\mathfrak{e},\beta))\,\big\{\beta f_A(T_j\mathfrak{e}) + (1-\beta) f_D(T_j\mathfrak{e})\big\}\Big]\mathbb{P}(N_t = n) \\ &= \sum_{n\in\mathbb{N}_0} \widehat{E}^{\mathfrak{e}}_{(0,\alpha)}\big[h(W_n)\big]\mathbb{P}(N_t = n), \end{aligned} \tag{4.61}$$

where the interchange of the order of summation in the second equality is justified by Fubini's theorem, $(W_n)_{n\in\mathbb{N}_0}$ is the auxiliary environment process (see Definition 4.3.5), and $h : \mathcal{E}_{\mathfrak{K}} \times \{0,1\} \to \mathbb{R}$ is the map $(\mathfrak{e},\alpha) \mapsto \alpha f_A(\mathfrak{e}) + (1-\alpha) f_D(\mathfrak{e})$. By virtue of Theorem 4.3.9, we can find a measurable $B \in \Sigma$ with $\bar{\mathbb{P}}(B) = 1$ such that, for all $\mathfrak{e} \in B$ and any $\alpha \in \{0,1\}$,

$$\lim_{n\to\infty} \widehat{E}^{\mathfrak{e}}_{(0,\alpha)}\big[h(W_n)\big] = \int_{\Omega_{\mathfrak{K}}} h(\mathfrak{b},\beta)\,\mathrm{d}\mathbb{Q}(\mathfrak{b},\beta) = \theta, \tag{4.62}$$

where θ is as in (4.24). Fix $\mathfrak{e} \in B$, $\alpha \in \{0,1\}$ and $\epsilon > 0$. By virtue of the above, we can find $N_{\mathfrak{e}} \in \mathbb{N}$ such that, for all $n \geq N_{\mathfrak{e}}$, $|\widehat{E}^{\mathfrak{e}}_{(0,\alpha)}\big[h(W_n)\big] - \theta| < \epsilon$. Finally, from (4.61), we get

$$\begin{aligned} |l((\mathfrak{e},\alpha),t) - \theta| &\leq \sum_{n=0}^{\infty} \big|\widehat{E}^{\mathfrak{e}}_{(0,\alpha)}\big[h(W_n)\big] - \theta\big|\,\mathbb{P}(N_t = n) \\ &\leq 2\|h\|_\infty\,\mathbb{P}(N_t < N_{\mathfrak{e}}) + \epsilon\,\mathbb{P}(N_t \geq N_{\mathfrak{e}}) \\ &\leq 2\|h\|_\infty\,\mathbb{P}(N_t < N_{\mathfrak{e}}) + \epsilon. \end{aligned} \tag{4.63}$$

Since $N_t \to \infty$ with probability 1 as $t \to \infty$, letting $t \to \infty$ in the above, we see

$$\limsup_{t\to\infty} |l((\mathfrak{e},\alpha),t) - \theta| \leq \epsilon. \tag{4.64}$$

As $\epsilon > 0$ is arbitrary, we get that

$$\lim_{t\to\infty} l((\mathfrak{e},\alpha),t) = \theta \tag{4.65}$$

for all $\mathfrak{e} \in B$ and $\alpha \in \{0,1\}$. This proves the claim in (4.23). □

Proof of Corollary 4.2.12. The proof basically follows from the translation-invariance of $\bar{\mathbb{P}}$ and Lemma 4.3.3. Indeed, using Theorem 4.2.11, we can find a measurable $B \in \Sigma$ such that $\bar{\mathbb{P}}(B) = 1$ and, for all $\mathfrak{e} \in B$, $\alpha \in \{0,1\}$,

$$\lim_{t\to\infty} \sum_{(j,\beta)\in G} p^{\mathfrak{e}}_t((0,\alpha),(j,\beta))\big[\beta f_A(T_j\mathfrak{e}) + (1-\beta) f_D(T_j\mathfrak{e})\big] = \theta, \tag{4.66}$$

where θ is as in (4.24). Letting $B_{\mathrm{inv}} := \cap_{j\in\mathbb{Z}^d} T_j^{-1}B$, we see that $B_{\mathrm{inv}} \in \Sigma$ is translation-invariant and $\bar{\mathbb{P}}(B_{\mathrm{inv}}) = 1$. In particular, for any $\mathfrak{e} \in B_{\mathrm{inv}}$ and all $(i,\alpha) \in \mathbb{Z}^d \times \{0,1\}$,

$$\lim_{t\to\infty} \sum_{(j,\beta)\in G} p^{T_i\mathfrak{e}}_t((0,\alpha),(j,\beta))\big[\beta f_A(T_j(T_i\mathfrak{e})) + (1-\beta) f_D(T_j(T_i\mathfrak{e}))\big] = \theta. \tag{4.67}$$

Also, using Lemma 4.3.3–4.3.4, we see that, for any $t \geq 0$ and $(j,\beta) \in \mathbb{Z}^d \times \{0,1\}$,

$$p^{T_i\mathfrak{e}}_t((0,\alpha),(j,\beta)) = p^{\mathfrak{e}}_t((i,\alpha),(i+j,\beta)), \quad \forall i \in \mathbb{Z}^d, \alpha \in \{0,1\}. \tag{4.68}$$

Combining the last two equations, for all $(i,\alpha) \in \mathbb{Z}^d \times \{0,1\}$, we get

$$\lim_{t\to\infty} \sum_{(j,\beta)\in G} p^{\mathfrak{e}}_t((i,\alpha),(i+j,\beta))\big[\beta f_A(T_{i+j}\mathfrak{e}) + (1-\beta) f_D(T_{i+j}\mathfrak{e})\big] = \theta, \tag{4.69}$$

which after a change of variable in the summation translates to

$$\lim_{t\to\infty} \sum_{(j,\beta)\in G} p_t^{\mathfrak{e}}((i,\alpha),(j,\beta))\big[\beta f_A(T_j\mathfrak{e}) + (1-\beta) f_D(T_j\mathfrak{e}))\big] = \theta. \tag{4.70}$$

The proof is complete by the observation that $\bar{\mathbb{P}}(B_{\text{inv}}) = 1$, and the above holds for any $\mathfrak{e} \in B_{\text{inv}}$. □

§4.4 Proof of main theorems

In this section we prove the two main results given in Section 4.2.2–4.2.3. In Section 4.4.1, we derive a consistency property of the general dual $Z_*^{\mathfrak{e}}$ of the process $Z^{\mathfrak{e}}$. Using this preliminary result on the dual, in Section 4.4.2 we prove Theorem 4.2.4, Corollary 4.2.5, and using Theorem 4.2.4 and the previous homogenisation result on the single-particle dual $\Theta^{\mathfrak{e}}$ (see Definition 4.3.1), we prove Theorem 4.2.9.

§4.4.1 Preliminaries: consistency of dual process

We start by recalling from Chapter 2 the duality relation between the spatial process $Z^{\mathfrak{e}}$ and the dual process $Z_*^{\mathfrak{e}}$ that will be needed for the proof of our main theorems.

Theorem 4.4.1 (Duality relation, [Corollary 2.4.6, Chapter 2]). *Suppose that Assumption* 2.A *is in force. Then, for every admissible environment $\mathfrak{e} = (N_i, M_i)_{i\in\mathbb{Z}^d} \in \mathcal{A}$, the following duality relation holds between the two processes $Z^{\mathfrak{e}}$ and $Z_*^{\mathfrak{e}}$:*

$$\mathbb{E}_U[D^{\mathfrak{e}}(Z^{\mathfrak{e}}(t), V)] = \mathbb{E}_*^V[D^{\mathfrak{e}}(U, Z_*^{\mathfrak{e}}(t))], \quad t \geq 0. \tag{4.71}$$

Here the expectation on the left (right) side is taken w.r.t. the law of $Z^{\mathfrak{e}}$ ($Z_*^{\mathfrak{e}}$) started at $U \in \mathcal{X}^{\mathfrak{e}}$ ($V \in \mathcal{X}_*^{\mathfrak{e}}$), and $D^{\mathfrak{e}} : \mathcal{X}^{\mathfrak{e}} \times \mathcal{X}_*^{\mathfrak{e}} \to [0,1]$ is the duality function defined by

$$D^{\mathfrak{e}}(U,V) = \prod_{i\in\mathbb{Z}^d} \frac{\binom{X_i}{n_i}}{\binom{N_i}{n_i}} \frac{\binom{Y_i}{m_i}}{\binom{M_i}{m_i}} \mathbb{1}_{n_i\le X_i, m_i\le Y_i}, \tag{4.72}$$

with $U = (X_i, Y_i)_{i\in\mathbb{Z}^d} \in \mathcal{X}^{\mathfrak{e}}$ and $V = (n_i, m_i)_{i\in\mathbb{Z}^d} \in \mathcal{X}_*^{\mathfrak{e}}$.

The next lemma establishes the relation between the process $\Theta^{\mathfrak{e}}$ and the general dual $Z_*^{\mathfrak{e}}$. We omit the proof for brevity, as this easily follows from the fact that any injective transformation preserves the Markov property and a unique such transformation exists that maps $\Theta^{\mathfrak{e}}$ to the dual process $Z_*^{\mathfrak{e}}$ started at a configuration consisting of only a single particle.

Lemma 4.4.2 (Relation between $\Theta^{\mathfrak{e}}$ and $Z_*^{\mathfrak{e}}$). *For $i \in \mathbb{Z}^d$, let $\vec{\delta}_{i,A}$ (resp. $\vec{\delta}_{i,D}$) $\in \mathcal{X}_*^{\mathfrak{e}}$ denote the configuration containing a single active (resp. dormant) particle at location i. Formally,*

$$\vec{\delta}_{i,A} := (\mathbb{1}_{\{n=i\}}, 0)_{n\in\mathbb{Z}^d}, \quad \vec{\delta}_{i,D} := (0, \mathbb{1}_{\{n=i\}})_{n\in\mathbb{Z}^d}, \tag{4.73}$$

and for $\eta = (i,\alpha) \in \mathbb{Z}^d \times \{0,1\}$, *let* $\vec{\delta}_\eta := \mathbb{1}_{\alpha=1}\,\vec{\delta}_{i,A} + \mathbb{1}_{\alpha=0}\,\vec{\delta}_{i,D}$. *If* $\mathbb{P}^{\varphi}_{\mathfrak{e}}$ *denotes the law of* $Z^{\mathfrak{e}}_*$ *started at* $\varphi \in \mathcal{X}^{\mathfrak{e}}_*$, *then, for all* $t \geq 0$,

$$p^{\mathfrak{e}}_t(\eta,\xi) = \mathbb{P}^{\vec{\delta}_\eta}_{\mathfrak{e}}(Z^{\mathfrak{e}}_*(t) = \vec{\delta}_\xi), \quad \eta,\xi \in \mathbb{Z}^d \times \{0,1\}, \tag{4.74}$$

where $p^{\mathfrak{e}}_t(\cdot\,,\cdot)$ *is as in Definition* 4.2.3.

The following lemma, which is essentially a consequence of Assumption 2.A, tells us that any bounded harmonic function of the single-particle dual process $\Theta^{\mathfrak{e}}$ is a constant.

Lemma 4.4.3 (Constant harmonics). *Let* $\Theta^{\mathfrak{e}} = (\Theta^{\mathfrak{e}}(t))_{t\geq 0}$ *be the process defined in Definition* 4.2.3 *started at* $\eta \in G$ *with law* $P^{\mathfrak{e}}_\eta$, *where* $G = \mathbb{Z}^d \times \{0,1\}$ *and* $\mathfrak{e} := (N_i, M_i)_{i\in\mathbb{Z}^d}$. *Let* $f : G \to \mathbb{R}$ *be a bounded harmonic function for* $P^{\mathfrak{e}}_\eta$, *i.e.,*

$$E^{\mathfrak{e}}_\eta[f(\Theta^{\mathfrak{e}}(t))] = f(\eta) \quad \text{for all } \eta \in G,\ t \geq 0. \tag{4.75}$$

Then f *is constant.*

Proof. Let $\mathcal{J}_{\mathfrak{e}}$ be the infinitesimal generator of the process $\Theta^{\mathfrak{e}}$. The action of $\mathcal{J}_{\mathfrak{e}}$ on f can be written in the following concise expression:

$$(\mathcal{J}_{\mathfrak{e}} f)(i,\alpha) := (\alpha\lambda + (1-\alpha)\lambda K_i)[f(i,1-\alpha) - f(i,\alpha)] + \alpha \sum_{j\in\mathbb{Z}^d} a(i,j)[f(j,\alpha) - f(i,\alpha)], \tag{4.76}$$

where $(i,\alpha) \in G$. Since f is harmonic, $(\mathcal{J}_{\mathfrak{e}} f) \equiv 0$ and, using the above, we have $f(i,\alpha) = f(i,1-\alpha)$ for all $(i,\alpha) \in G$, which in turn implies that the function $i \mapsto f(i,1)$ is harmonic for $a(\cdot\,,\cdot)$. Applying the Choquet-Deny theorem to the irreducible and translation-invariant kernel $a(\cdot\,,\cdot)$, we get the result. □

By using the duality relation stated in Theorem 4.4.1 and exploiting the clustering criterion given in Theorem 2.4.12 of Chapter 2, we obtain that coalescence of two dual particles with probability 1 is equivalent to coalescence of any number of dual particles with probability 1.

Theorem 4.4.4 (Lineage consistency). *Let* $\mathbb{P}^{\varphi}_{\mathfrak{e}}$ *denote the law of the dual process* $Z^{\mathfrak{e}}_*$ *started at* $\varphi := (n_i, m_i)_{i\in\mathbb{Z}^d} \in \mathcal{X}^{\mathfrak{e}}_*$ *and evolving in environment* $\mathfrak{e} := (N_i, M_i)_{i\in\mathbb{Z}^d}$. *Let* τ *be first time when all particles have coalesced into a single particle in the dual process, i.e.,*

$$\tau := \inf\{t \geq 0 \,:\, |Z^{\mathfrak{e}}_*(t)| = 1\}, \tag{4.77}$$

where $|\varphi| := \sum_{i\in\mathbb{Z}^d}(n_i + m_i)$ *is the total number of initial dual particles. Then the following are equivalent:*

(a) $\mathbb{P}^{\varphi}_{\mathfrak{e}}(\tau < \infty) = 1$ *for all* $\varphi \in \mathcal{X}^{\mathfrak{e}}_*$ *with* $|\varphi| = 2$.

(b) $\mathbb{P}^{\varsigma}_{\mathfrak{e}}(\tau < \infty) = 1$ *for all* $\varsigma \in \mathcal{X}^{\mathfrak{e}}_*$ *with* $|\varsigma| \geq 2$.

Proof. By irreducibility of the dual process $Z_*^{\mathfrak{e}}$, it suffices to prove the equivalence of the two statements for fixed $\varphi, \varsigma \in \mathcal{X}_*^{\mathfrak{e}}$ such that $|\varphi| = 2$ and $n := |\varsigma| \geq 2$. If $n = 2$, then there is nothing to prove. So assume that $n > 2$. It is straightforward to see from irreducibility and the Markov property of $Z_*^{\mathfrak{e}}$ that if $\mathbb{P}_{\mathfrak{e}}^{\varphi}(\tau = \infty) > 0$, then $\mathbb{P}_{\mathfrak{e}}^{\varsigma}(\tau = \infty) \geq \mathbb{P}_{\mathfrak{e}}^{\varsigma}(Z^*(t) = \varphi)\mathbb{P}_{\mathfrak{e}}^{\varphi}(\tau = \infty) > 0$. Hence (b) implies (a).

To prove that (a) implies (b), assume $\mathbb{P}_{\mathfrak{e}}^{\varphi}(\tau < \infty) = 1$ and, for $t \geq 0$, set $I_t := |Z_*^{\mathfrak{e}}(t)|$. Note that, since $Z_*^{\mathfrak{e}}$ is a coalescent process, I_t is an integer-valued bounded random variable that is decreasing in t a.s. Thus, $I := \lim_{t\to\infty} I_t$ exists a.s. and it is enough to prove that $I = 1$ a.s. To this purpose, let $\theta \in (0,1)$ be fixed arbitrarily, and let $Z^{\mathfrak{e}}$ be the spatial process started at the initial distribution $\mu_\theta^{\mathfrak{e}}$ given by

$$\mu_\theta^{\mathfrak{e}} := \bigotimes_{i\in\mathbb{Z}^d} \text{Binomial}(N_i, \theta) \otimes \text{Binomial}(M_i, \theta). \tag{4.78}$$

By Theorem 2.4.9 of Chapter 2, the process $Z^{\mathfrak{e}}$ converges to an equilibrium ν_θ. Also, by our assumption that $\mathbb{P}_{\mathfrak{e}}^{\varphi}(\tau < \infty) = 1$ and Theorem 2.4.12 of Chapter 2, we have

$$\nu_\theta = (1-\theta)\delta_{\spadesuit} + \theta\delta_{\heartsuit}, \tag{4.79}$$

where $\delta_{\heartsuit}$ (resp. $\delta_{\spadesuit}$) is the Dirac distribution concentrated at the all type-$\heartsuit$ configuration $\mathfrak{e} \in \mathcal{X}^{\mathfrak{e}}$ (resp. the all type-$\spadesuit$ configuration $(0,0)_{i\in\mathbb{Z}^d} \in \mathcal{X}^{\mathfrak{e}}$). Furthermore, if $D^{\mathfrak{e}}(\cdot\,,\cdot)$ is the duality function in (4.72), then combining Theorem 2.4.9 of Chapter 2 and the above we get

$$\theta = \mathbb{E}_{\nu_\theta}\big[D^{\mathfrak{e}}(Z^{\mathfrak{e}}(0), \varsigma)\big] = \lim_{t\to\infty} \mathbb{E}_{\mathfrak{e}}^{\varsigma}\big[\theta^{I_t}\big] = \mathbb{E}_{\mathfrak{e}}^{\varsigma}\big[\theta^{I}\big] \quad \text{(bounded convergence)}, \tag{4.80}$$

which implies that $\mathbb{E}_{\mathfrak{e}}^{\varsigma}\big[\theta(1-\theta^{I-1})\big] = 0$. Since $\theta \in (0,1)$, we have that $I = 1$ almost surely. □

§4.4.2 Proofs: clustering in fixed and random environment

We are now ready to prove the two main theorems.

Proof of Theorem 4.2.4. To show that (a) implies (b), suppose that $\mu_t^{\mathfrak{e}}$ converges weakly to $\nu \in \mathcal{P}(\mathcal{X}^{\mathfrak{e}})$ as $t \to \infty$. Let $\theta_{\mathfrak{e}} := \mathbb{E}_\nu\big[\frac{X_0^{\mathfrak{e}}(0)}{N_0}\big] \in [0,1]$ be fixed. Since the system is in the clustering regime by assumption, $\delta_{\spadesuit}$ and $\delta_{\heartsuit}$ are the only two extremal equilibria for the process $Z^{\mathfrak{e}}$. Hence, we must have that

$$\nu = (1-\theta_{\mathfrak{e}})\delta_{\spadesuit} + \theta_{\mathfrak{e}}\delta_{\heartsuit}, \tag{4.81}$$

where $\delta_{\heartsuit}$ (resp. $\delta_{\spadesuit}$) is the Dirac distribution concentrated at the all type-$\heartsuit$ configuration $\mathfrak{e} \in \mathcal{X}^{\mathfrak{e}}$ (resp. $(0,0)_{i\in\mathbb{Z}^d} \in \mathcal{X}^{\mathfrak{e}}$). We show that $f^{\mathfrak{e}} \equiv \theta_{\mathfrak{e}}$, which will settle (b) along with the last statement of the theorem. To this end, for each $t \geq 0$, let $f_t^{\mathfrak{e}} : G \to [0,1]$ be defined as

$$f_t^{\mathfrak{e}}(\eta) := \sum_{(j,\beta)\in G} p_t^{\mathfrak{e}}(\eta, (j,\beta)) \int_{\mathcal{X}^{\mathfrak{e}}} \big[\beta\tfrac{X_j}{N_j} + (1-\beta)\tfrac{Y_j}{M_j}\big]\, \mathrm{d}\mu^{\mathfrak{e}}\{(X_k, Y_k)_{k\in\mathbb{Z}^d}\}, \quad \eta \in G. \tag{4.82}$$

Let $\eta = (i, \alpha) \in G$ be arbitrary, and let $Z_*^{\mathfrak{e}} := (Z_*^{\mathfrak{e}}(t))_{t\geq 0}$ be the dual process started at $\vec{\delta}_\eta := \mathbb{1}_{\alpha=1}\,\vec{\delta}_{i,A} + \mathbb{1}_{\alpha=0}\,\vec{\delta}_{i,D}$, where for each $i \in \mathbb{Z}^d$ the configurations $\vec{\delta}_{i,A}, \vec{\delta}_{i,D} \in \mathcal{X}_*^{\mathfrak{e}}$ are defined as in (4.73). In other words, $\vec{\delta}_\eta$ is the configuration with a single dual particle located at $i \in \mathbb{Z}^d$ with state α. Recall from Definition 4.2.3 that the time-t transition kernel $p_t^{\mathfrak{e}}(\cdot\,,\cdot)$ of the single-particle dual process $\Theta^{\mathfrak{e}}$ is defined as

$$p_t^{\mathfrak{e}}(\eta, \zeta) := P_\eta^{\mathfrak{e}}(\Theta^{\mathfrak{e}}(t) = \zeta), \quad \eta, \zeta \in G. \tag{4.83}$$

Using Lemma 4.4.2 and appealing to the monotone convergence theorem, we get from (4.82) that

$$f_t^{\mathfrak{e}}(\eta) = \int_{\mathcal{X}^{\mathfrak{e}}} \mathbb{E}_{\mathfrak{e}}^{\vec{\delta}_\eta}\left[D^{\mathfrak{e}}(z, Z_*^{\mathfrak{e}}(t))\right] \mathrm{d}\mu^{\mathfrak{e}}\{z\}, \tag{4.84}$$

where the expectation is w.r.t. the law of the dual process $Z_*^{\mathfrak{e}}$, and $D^{\mathfrak{e}}(\cdot\,,\cdot)$ is the duality function in (4.72). Furthermore, applying the duality relation between $Z^{\mathfrak{e}}$ and $Z_*^{\mathfrak{e}}$ to the above identity, we get

$$f_t^{\mathfrak{e}}(\eta) = \mathbb{E}_{\mu^{\mathfrak{e}}}\left[D^{\mathfrak{e}}(Z^{\mathfrak{e}}(t), \vec{\delta}_\eta)\right] = \int_{\mathcal{X}^{\mathfrak{e}}} D^{\mathfrak{e}}(z, \vec{\delta}_\eta)\, \mathrm{d}\mu_t^{\mathfrak{e}}\{z\}. \tag{4.85}$$

However, since $\mu_t^{\mathfrak{e}} \overset{weak}{\longrightarrow} \nu$ as $t \to \infty$, combining the above with (4.81), we see that

$$f^{\mathfrak{e}}(\eta) = \lim_{t\to\infty} f_t^{\mathfrak{e}}(\eta) = \int_{\mathcal{X}^{\mathfrak{e}}} D^{\mathfrak{e}}(z, \vec{\delta}_\eta)\, \mathrm{d}\nu\{z\} = \theta_{\mathfrak{e}}, \tag{4.86}$$

and hence the claim is proved.

To prove the converse, for $t \geq 0$, let $f_t^{\mathfrak{e}} : G \to [0, 1]$ be as in (4.82). Applying Fubini's theorem to (4.84), for any $\eta \in G$ we have

$$f_t^{\mathfrak{e}}(\eta) = \mathbb{E}_{\mathfrak{e}}^{\vec{\delta}_\eta}\Big[\int_{\mathcal{X}^{\mathfrak{e}}} D^{\mathfrak{e}}(z, Z_*^{\mathfrak{e}}(t))\, \mathrm{d}\mu^{\mathfrak{e}}\{z\}\Big]. \tag{4.87}$$

Using the Markov property of $Z_*^{\mathfrak{e}}$, we note that, for $t, s \geq 0$ and $\eta \in G$,

$$f_{s+t}^{\mathfrak{e}}(\eta) = \sum_{\zeta\in G} p_s^{\mathfrak{e}}(\eta, \zeta) f_t^{\mathfrak{e}}(\zeta). \tag{4.88}$$

Since by assumption $f^{\mathfrak{e}}(\eta) = \lim_{t\to\infty} f_t^{\mathfrak{e}}(\eta)$ exists for any $\eta \in G$, letting $t \to \infty$ in the above identity, we obtain

$$\begin{aligned} f^{\mathfrak{e}}(\eta) &= \lim_{t\to\infty} \sum_{\zeta\in G} p_s^{\mathfrak{e}}(\eta, \zeta) f_t^{\mathfrak{e}}(\zeta) = \sum_{\zeta\in G} p_s^{\mathfrak{e}}(\eta, \zeta)\left[\lim_{t\to\infty} f_t^{\mathfrak{e}}(\zeta)\right] \quad \text{(dominated convergence)} \\ &= \sum_{\zeta\in G} p_s^{\mathfrak{e}}(\eta, \zeta) f^{\mathfrak{e}}(\zeta) = E_\eta^{\mathfrak{e}}\left[f^{\mathfrak{e}}(\Theta^{\mathfrak{e}}(s))\right]. \end{aligned} \tag{4.89}$$

Hence, in particular, $f^{\mathfrak{e}}$ is harmonic for the process $(\Theta^{\mathfrak{e}}(t))_{t\geq 0}$ and thus, by Lemma 4.4.3, $f^{\mathfrak{e}} \equiv \theta_{\mathfrak{e}}$ for some $\theta_{\mathfrak{e}} \in [0, 1]$. It only remains to show that $\mu_t^{\mathfrak{e}}$ converges weakly as $t \to \infty$. This is equivalent to showing that, for any $\varphi \in \mathcal{X}_*^{\mathfrak{e}}$, $\lim_{t\to\infty} \mathbb{E}_{\mu^{\mathfrak{e}}}\left[D^{\mathfrak{e}}(Z^{\mathfrak{e}}(t), \varphi)\right]$ exists. Because $\mathcal{P}(\mathcal{X}^{\mathfrak{e}})$ is compact (as $\mathcal{X}^{\mathfrak{e}}$ is) in the topology of weak convergence,

$(\mu_t^{\mathfrak{e}})_{t\geq 0}$ is tight. Finally, the existence of the limit ensures the convergence of the associated finite-dimensional distributions, because the family of functions $\{D^{\mathfrak{e}}(\,\cdot\,,\varphi) : \varphi \in \mathcal{X}_*^{\mathfrak{e}}\}$ fixes the mixed moments of the finite-dimensional distributions of $Z^{\mathfrak{e}}$ (see Proposition 2.6.4 in Chapter 2), and therefore is convergence determining. Let $\varphi = (n_i, m_i)_{i\in\mathbb{Z}^d} \in \mathcal{X}_*^{\mathfrak{e}}$ be fixed, and $Z_*^{\mathfrak{e}}$ be the dual process started at φ. First note that if $|\varphi| = \sum_{i\in\mathbb{Z}^d}(n_i + m_i) = 1$, then the limit exists and equals $\theta_{\mathfrak{e}}$ by our assumption. Indeed, if $|\varphi| = 1$, then $\varphi = \vec{\delta}_\zeta$ for some $\zeta \in G$. As a consequence of duality and (4.84), we see that $\mathbb{E}_{\mu^{\mathfrak{e}}}\big[D^{\mathfrak{e}}(Z^{\mathfrak{e}}(t),\varphi)\big] = f_t^{\mathfrak{e}}(\zeta)$ and hence

$$\lim_{t\to\infty} \mathbb{E}_{\mu^{\mathfrak{e}}}\big[D^{\mathfrak{e}}(Z^{\mathfrak{e}}(t),\varphi)\big] = \lim_{t\to\infty} \mathbb{E}_{\mathfrak{e}}^{\vec{\delta}_\zeta}\Big[\int_{\mathcal{X}^{\mathfrak{e}}} D^{\mathfrak{e}}(z, Z_*^{\mathfrak{e}}(t))\,\mathrm{d}\mu^{\mathfrak{e}}\{z\}\Big] = \lim_{t\to\infty} f_t^{\mathfrak{e}}(\zeta) = f^{\mathfrak{e}}(\zeta) = \theta_{\mathfrak{e}}. \tag{4.90}$$

Now, let us fix $\varphi \in \mathcal{X}_*^{\mathfrak{e}}$ such that $|\varphi| \geq 2$. Since the system is in the clustering regime, by virtue of Theorem 2.4.12 stated in Chapter 2, condition (a) in Theorem 4.4.4 is satisfied. Hence from part (b) of Theorem 4.4.4 it follows that $\tau < \infty$ a.s., where

$$\tau := \inf\{t \geq 0 \,:\, |Z_*^{\mathfrak{e}}(t)| = 1\}. \tag{4.91}$$

Using duality and the strong Markov property of the dual process, we see that

$$\begin{aligned}
\lim_{t\to\infty} \mathbb{E}_{\mu^{\mathfrak{e}}}\Big[D^{\mathfrak{e}}(Z^{\mathfrak{e}}(t),\varphi)\Big] &\overset{Fubini}{=} \lim_{t\to\infty} \mathbb{E}_{\mathfrak{e}}^{\varphi}\Big[\int_{\mathcal{X}^{\mathfrak{e}}} D^{\mathfrak{e}}(z, Z_*^{\mathfrak{e}}(t))\,\mathrm{d}\mu^{\mathfrak{e}}\{z\}\Big]\\
&= \lim_{t\to\infty} \mathbb{E}_{\mathfrak{e}}^{\varphi}\Big[\int_{\mathcal{X}^{\mathfrak{e}}} D^{\mathfrak{e}}(z, Z_*^{\mathfrak{e}}(t))\,\mathrm{d}\mu^{\mathfrak{e}}\{z\}; \tau \leq t\Big]\\
&\quad + \lim_{t\to\infty} \underbrace{\mathbb{E}_{\mathfrak{e}}^{\varphi}\Big[\int_{\mathcal{X}^{\mathfrak{e}}} D^{\mathfrak{e}}(z, Z_*^{\mathfrak{e}}(t))\,\mathrm{d}\mu^{\mathfrak{e}}\{z\} \mid \tau > t\Big]}_{\leq 1} \mathbb{P}_{\mathfrak{e}}^{\varphi}(\tau > t)\\
&= \lim_{t\to\infty} \mathbb{E}_{\mathfrak{e}}^{\varphi}\Big[\mathbb{E}_{\mathfrak{e}}^{Z_*^{\mathfrak{e}}(\tau)}\Big[\int_{\mathcal{X}^{\mathfrak{e}}} D^{\mathfrak{e}}(z, Z_*^{\mathfrak{e}}(t-\tau))\,\mathrm{d}\mu^{\mathfrak{e}}\{z\}\Big]; \tau \leq t\Big]\\
&= \lim_{t\to\infty} \mathbb{E}_{\mathfrak{e}}^{\varphi}\Big[\sum_{\zeta\in G} f_{t-\tau}^{\mathfrak{e}}(\zeta)\mathbb{1}_{\{Z_*^{\mathfrak{e}}(\tau)=\vec{\delta}_\zeta\}}; \tau \leq t\Big],
\end{aligned} \tag{4.92}$$

where we use that the second term after the first equality converges to 0 because $\tau < \infty$ a.s., and the last equality follows from (4.84) and the fact that $Z_*^{\mathfrak{e}}(\tau) = \vec{\delta}_\zeta$ for some $\zeta \in G$. Finally, by an application of the dominated convergence theorem, we get

$$\begin{aligned}
&\lim_{t\to\infty} \mathbb{E}_{\mu^{\mathfrak{e}}}\Big[D^{\mathfrak{e}}(Z^{\mathfrak{e}}(t),\varphi)\Big]\\
&\quad= \mathbb{E}_{\mathfrak{e}}^{\varphi}\Big[\sum_{\zeta\in G}\big(\lim_{t\to\infty} f_{t-\tau}^{\mathfrak{e}}(\zeta)\big)\mathbb{1}_{\{Z_*^{\mathfrak{e}}(\tau)=\vec{\delta}_\zeta\}}; \tau < \infty\Big]\\
&\quad= \mathbb{E}_{\mathfrak{e}}^{\varphi}\Big[\sum_{\zeta\in G} f^{\mathfrak{e}}(\zeta)\mathbb{1}_{\{Z_*^{\mathfrak{e}}(\tau)=\vec{\delta}_\zeta\}}; \tau < \infty\Big] = \theta_{\mathfrak{e}}\,\mathbb{P}_{\mathfrak{e}}^{\varphi}(\tau < \infty) \quad (\text{since } f^{\mathfrak{e}} \equiv \theta_{\mathfrak{e}})\\
&\quad= \theta_{\mathfrak{e}}.
\end{aligned} \tag{4.93}$$

This shows that there exists $\nu \in \mathcal{P}(\mathcal{X}^{\mathfrak{e}})$ such that $\mu_t^{\mathfrak{e}}$ converges weakly to ν as $t \to \infty$. Since the system clusters by assumption, we must have

$$\nu = (1-\theta_{\mathfrak{e}})\delta_{\spadesuit} + \theta_{\mathfrak{e}}\delta_{\heartsuit}, \tag{4.94}$$

where $\delta_\heartsuit$ (resp. $\delta_\spadesuit$) is the Dirac distribution concentrated at the all type-$\heartsuit$ configuration $\mathfrak{e} \in \mathcal{X}^\mathfrak{e}$ (resp. the all type-$\spadesuit$ configuration $(0,0)_{i\in\mathbb{Z}^d} \in \mathcal{X}^\mathfrak{e}$). $\square$

Proof of Corollary 4.2.5. The proof basically exploits Theorem 4.2.4 and the fact that the particle associated to the process $\Theta^\mathfrak{e}$ eventually leaves any finite region of the state space $G = \mathbb{Z}^d \times \{0,1\}$ with probability 1. It suffices to prove that condition (b) in Theorem 4.2.4 is satisfied. Let $f : \mathbb{Z}^d \times \{0,1\} \to [0,1]$ be the map

$$f(i,\alpha) := \alpha \mathbb{E}_{\mu^\mathfrak{e}}\big[\tfrac{X_i^\mathfrak{e}(0)}{N_i}\big] + (1-\alpha)\mathbb{E}_{\mu^\mathfrak{e}}\big[\tfrac{Y_i^\mathfrak{e}(0)}{M_i}\big], \quad (i,\alpha) \in \mathbb{Z}^d \times \{0,1\}, \tag{4.95}$$

and let $\epsilon > 0$ be arbitrary. By (4.10), there exists $N \in \mathbb{N}$ such that, for all $i \in \mathbb{Z}^d$, $\|i\| > N$ and $\alpha \in \{0,1\}$, $|f(i,\alpha) - \theta_\mathfrak{e}| < \epsilon$. Thus, if $p_t^\mathfrak{e}(\cdot\,,\cdot)$ is the time-t transition kernel of the process $(\Theta^\mathfrak{e}(t))_{t\geq 0}$ in Definition 4.2.3, then for any $\eta \in G$ and $t \geq 0$,

$$\begin{aligned}
&\Bigg|\sum_{(j,\beta)\in G} p_t^\mathfrak{e}(\eta,(j,\beta))\Big\{\beta\,\mathbb{E}_{\mu^\mathfrak{e}}\big[\tfrac{X_j^\mathfrak{e}(0)}{N_j}\big] + (1-\beta)\mathbb{E}_{\mu^\mathfrak{e}}\big[\tfrac{Y_j^\mathfrak{e}(0)}{M_j}\big]\Big\} - \theta_\mathfrak{e}\Bigg| \\
&\quad\leq \sum_{\substack{(j,\beta)\in G,\\ \|j\|\leq N}} p_t(\eta,(j,\beta))\underbrace{\big|f(j,\beta)-\theta_\mathfrak{e}\big|}_{\leq 2} + \sum_{\substack{(j,\beta)\in G,\\ \|j\|> N}} p_t^\mathfrak{e}(\eta,(j,\beta))\underbrace{\big|f(j,\beta)-\theta_\mathfrak{e}\big|}_{\leq \epsilon} \\
&\quad\leq 2\,P_\eta^\mathfrak{e}(\Theta^\mathfrak{e}(t) \in \Lambda_N \times \{0,1\}) + \epsilon\,P_\eta^\mathfrak{e}(\Theta^\mathfrak{e}(t) \notin \Lambda_N \times \{0,1\}),
\end{aligned} \tag{4.96}$$

where $\Lambda_N := \mathbb{Z}^d \cap [0,N]^d$, and $P_\eta^\mathfrak{e}$ denotes the law of $(\Theta^\mathfrak{e}(t))_{t\geq 0}$ started at η. Since Λ_N is finite, $\lim_{t\to\infty} P_\eta^\mathfrak{e}(\Theta^\mathfrak{e}(t) \in \Lambda_N \times \{0,1\}) = 0$, and so letting $t \to \infty$ in (4.96), we get

$$\limsup_{t\to\infty}\Bigg|\sum_{(j,\beta)\in G} p_t^\mathfrak{e}(\eta,(j,\beta))\Big\{\beta\,\mathbb{E}_{\mu^\mathfrak{e}}\big[\tfrac{X_j^\mathfrak{e}(0)}{N_j}\big] + (1-\beta)\mathbb{E}_{\mu^\mathfrak{e}}\big[\tfrac{Y_j^\mathfrak{e}(0)}{M_j}\big]\Big\} - \theta_\mathfrak{e}\Bigg| \leq \epsilon. \tag{4.97}$$

As ϵ is arbitrary, we see that

$$\lim_{t\to\infty} \sum_{(j,\beta)\in G} p_t^\mathfrak{e}(\eta,(j,\beta)) f(j,\beta) = \theta_\mathfrak{e} \tag{4.98}$$

and hence the claim follows from Theorem 4.2.4. $\square$

Proof of Theorem 4.2.9. We exploit Theorem 4.2.4 and the homogenisation result in Corollary 4.2.12. We see that, because of conditions (1)–(2) in Assumption 4.A and ellipticity of the environments $\mathfrak{e} \in \mathcal{E}_\mathfrak{K}$, the process $Z^\mathfrak{e}$ is in the clustering regime for every environment $\mathfrak{e} \in \mathcal{E}_\mathfrak{K}$. Also, by virtue of Corollary 4.2.12 and the assumption in (4.18) on initial distributions, there exists $B \in \Sigma$ such that $\bar{\mathbb{P}}(B) = 1$, and for all $\mathfrak{e} \in B$ condition (b) of Theorem 4.2.4 holds. Furthermore, we see from Corollary 4.2.12, that the limiting value in that condition is independent of the environment $\mathfrak{e}$, and is given by (4.19). Hence the result follows. $\square$