Cubatures on Wiener Space, Spatial Coalescents, and Repulsion-Diffusion Equations

Peter Koepernik
St. Catherine's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

August 31st, 2025

Contents

Introduction **5**

Chapter I. Cubatures on Wiener Space 5

Chapter II. Genealogies of Spatial Population Models 9

Chapter III. Repulsion-Diffusion Equations 13

I Cubatures on Wiener Space **17**

1 Introduction 17

2 Background 20

2.1 Rough Path Theory 21

2.2 Cubature Formulae 26

3 The ARCANE Algorithm 29

3.1 Explanation of the Algorithm 30

Step 1: Construction of Approximate Cubature 30

Step 2: Recombination 33

Step 3: Sharpening 34

Step 4: Dyadic Construction 35

3.2 Theoretical Results 35

3.2.1 Asymptotic Consistency 35

3.2.2 Proofs of Lemmas 3.1 and 3.3 37

4 Empirical Results 38

5 Limitations 40

Acknowledgements 41

Bibliography 47

Appendices **48**

I.A Orthogonal Arrays 48

I.A.1 Linear Orthogonal Arrays 50

I.A.1.1 Direct linear constructions 51

I.A.1.2 Constructions from cap sets 53

I.A.2 Constructions from non-linear codes 54

I.B Theoretical Results 56

I.B.1 (Re-)statement of results . . . 56
I.B.2 Proofs . . . 56

II Genealogies of Spatial Population Models **65**
1 Introduction . . . 65
1.1 Coalescents . . . 65
1.2 Spatial Coalescents . . . 67
1.3 The Brownian Spatial Coalescent . . . 68
1.4 Sampling Consistency . . . 73
1.5 Associated Population Models . . . 77
1.5.1 Ξ-Fleming-Viot Process . . . 77
1.5.2 Spatial Cannings Models . . . 79
1.6 Related Literature . . . 81
1.7 Conclusion and Outlook . . . 83
1.7.1 Spatially Interactive Branching . . . 83
1.7.2 General Spatial Motions . . . 83
1.7.3 The Brownian Spatial Coalescent on $\mathbb{R}^d$ . . . 84
1.8 Structure . . . 85
2 Setup . . . 85
2.1 Non-spatial Coalescents . . . 85
2.2 Spatial Coalescents . . . 86
2.3 Brownian Spatial Coalescents . . . 88
3 Proofs . . . 94
3.1 Characterisation of Brownian Spatial Coalescents . . . 94
3.1.1 "If" Direction of Theorem 2.14 . . . 94
3.1.2 "Only if" Direction of Theorem 2.14 . . . 97
3.2 Characterisation of Label Invariance . . . 109
3.3 Characterisation of Sampling Consistency . . . 110
3.3.1 Proof of Lemma 3.14 . . . 113
3.3.2 Consequences of Sampling Consistency . . . 115
3.4 Drift Representation . . . 122
3.5 Associated Population Models . . . 123
3.5.1 Ξ-Fleming-Viot Process . . . 123
3.5.2 Spatial Cannings Models . . . 128
Acknowledgements . . . 132
Bibliography . . . 135

Appendices **136**
II.A Maps and Spaces . . . 136
II.B Disintegration of Brownian Motion (Lemma 3.22) . . . 137
II.C Characterisation of Label Invariance (Lemma 3.13) . . . 140
II.D Conditional Kernel of the Brownian Spatial Coalescent (Lemma 2.9) . . . 141
II.E N^ν is finite and continuous (Lemma 2.12) . . . 143
II.F Minor Technical Lemmas . . . 151

III Repulsion-Diffusion Equations **153**
1 Introduction . . . 153
1.1 Motivation—The Pain in the Torus . . . 153
1.2 Summary of results . . . 156
1.3 Related work . . . 158
2 Dynamical Properties . . . 160
2.1 Results . . . 160
2.2 Proofs . . . 163
2.2.1 Existence and regularity of solutions . . . 164
2.2.2 Maximum principle and global boundedness . . . 170
3 Stationary Distributions . . . 180
3.1 Results . . . 180
3.2 Proofs . . . 182
3.2.1 Proof of Lemma 3.8 . . . 186
Acknowledgements . . . 202
Bibliography . . . 205

Appendices **206**
III.A Fractional Sobolev spaces . . . 206

Introduction

Stochastic models are indispensable in the sciences, but simplifying assumptions are often necessary to retain mathematical or computational tractability. In many cases, this leads to models that are only selectively applicable or exhibit unrealistic behaviour. The guiding philosophy of this thesis is to make an attempt at closing this gap, in two ways: Firstly, by making generic simulation algorithms more efficient, which enables the use of more realistic models that would otherwise be computationally infeasible, across different application areas (Chapter I). And secondly, by building and analysing more realistic generalisations of two common models in the particular field of population biology (Chapters II and III).

As an integrated thesis, all three chapters are based on largely independent projects that have either been published or submitted for publication. They are self-contained with separate introductions, bibliographies, and appendices, and can be read independently of each other. In the remainder of this introduction, I give a brief, high-level motivation and summary of each of the three projects.

Chapter I. Cubatures on Wiener Space

Motivation Stochastic differential equations (SDEs) are ubiquitously used in sciences to model dynamical systems evolving in continuous time under the influence of random noise. Example application areas include chemistry [13], physics [24], population genetics [7], systems biology [29], social sciences [5, chapter 2], epidemiology [4, 20], mathematical finance [6], wind speed modelling [1], solar radiation forecasting [2, 16], research engineering [30], and modelling of aircraft dynamics [17, 18].

Mathematically, an SDE can be written as

$$\mathrm{d}y_t = \mu(y_t)\,\mathrm{d}t + \sum_{i=1}^{d} \sigma_i(y_t)\,\mathrm{d}W_t^i. \tag{1}$$

Here $y = \{y_t\}_{t\in[0,T]}$ evolves in $\mathbb{R}^e$ for some $e \in \mathbb{N}$ and describes the stochastic system in question (e.g. y_t could describe the concentrations at time t of some number e of reactants in a chemical reaction). $\mu\colon \mathbb{R}^e \to \mathbb{R}^e$ is called the *drift* and describes the deterministic behaviour of the system (i.e. it's behaviour in the absence of noise). Finally, the equation is influenced by some number $d \in \mathbb{N}$ of independent sources of noise, represented by

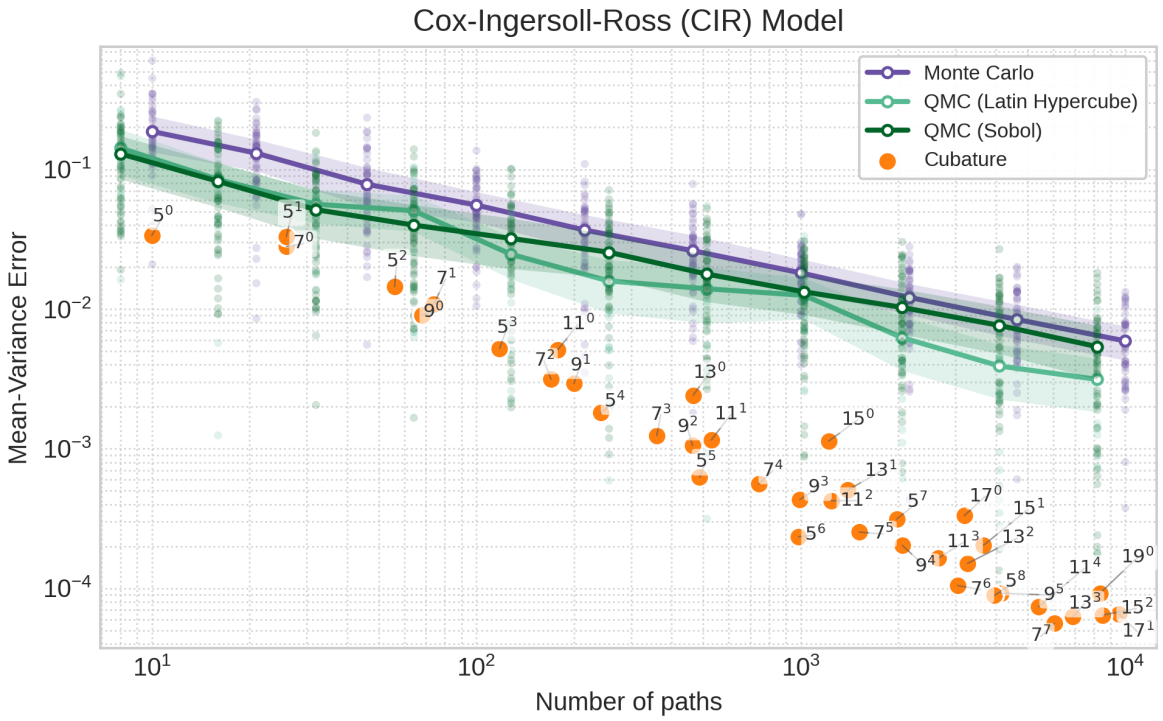

Figure 1: The SDE simulation error (a combination of the errors of the empirical mean and variance of the endpoint of the solution path, see (4.1)) of our cubature paths (orange) against an equivalent number of Monte Carlo paths (purple) and Quasi-Monte Carlo paths (green), evaluated on the CIR model, see (4.4).

a d-dimensional standard Brownian motion $W = (W^1, \dots, W^d) = \{W_t\}_{t\in[0,T]}$. Their influence on the system is described by the *diffusion* coefficients $(\sigma_i \colon \mathbb{R}^e \to \mathbb{R}^e)_{i=1}^d$.

In practice, one is often primarily interested in estimating statistics of the solution to (1), such as its mean or variance, or more generally the expectation of $f(y_T)$ for some functional $f \colon \mathbb{R}^e \to \mathbb{R}$. (The main alternative are applications where one is primarily interested in obtaining individual samples of the solution, such as in generative diffusion models.) If a closed-form analytical solution is not available, as is usually the case, then this is traditionally achieved using Monte Carlo estimation, which takes the average over some number M of independently sampled solutions of (1):

$$\mathbb{E}[f(y_T)] \approx \frac{1}{M} \sum_{i=1}^{M} f(y_T(\omega_i)), \tag{2}$$

where $\omega_i \colon [0,T] \to \mathbb{R}^d$ are independent Brownian sample paths, and $y_T(\omega_i) \in \mathbb{R}^e$ denotes the solution at time T of the system (1) driven by the path ω_i. However, by the Central Limit Theorem, the error of the Monte Carlo estimate only decays proportionally to $1/\sqrt{M}$, where M is the number of samples. An accurate estimation can therefore require an enormous number of sample paths and be computationally very expensive—especially for complex SDEs.

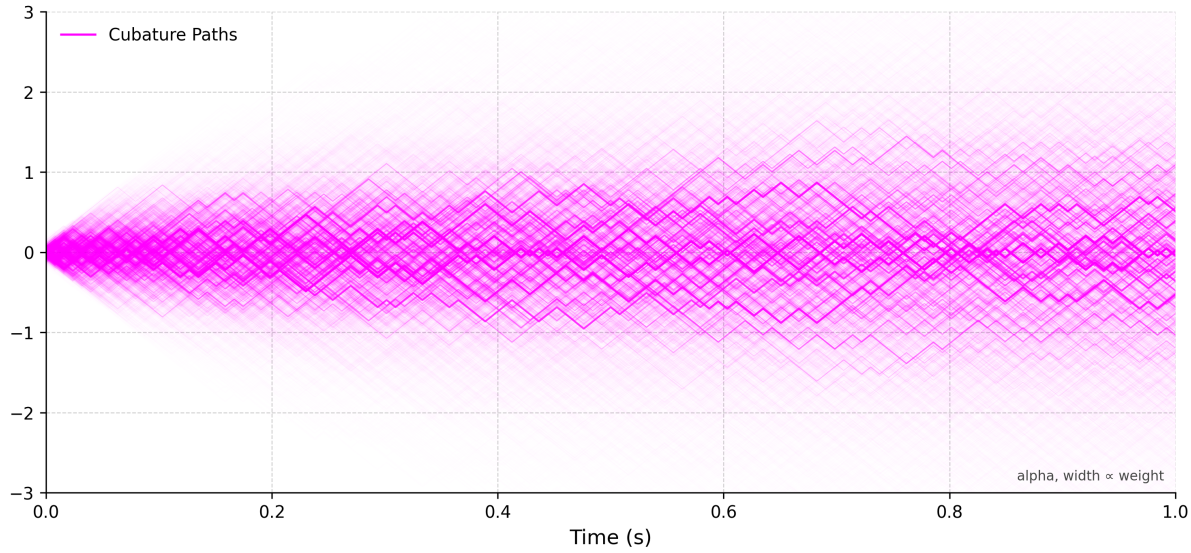

Figure 2: A visualisation of one of our largest cubature formulae, which has degree 17 and consists of 9574 paths.

Cubatures in Numerical Integration To get an idea how one might try to do better, consider for a moment the task of integrating some function $f\colon [0,1] \to \mathbb{R}$ over the unit interval. This integral can also be thought of as the expectation of $f(U)$ for a uniform random variable $U \sim \mathcal{U}[0,1]$, and therefore Monte-Carlo estimation is once again applicable:

$$\int_0^1 f(u)\,\mathrm{d}u = \mathbb{E}\left[f(U)\right] \approx \frac{1}{M}\sum_{i=1}^{M} f(u_i), \tag{3}$$

where $u_i \sim \mathcal{U}[0,1]$ are independent. However, in this case it is well-known that one can do significantly better. For example, by using uniformly spaced points $u_i = i/M$ instead of random ones, which already improves the order of convergence from $1/\sqrt{M}$ to $1/M$ (under reasonable assumptions on f, see the Euler-MacLaurin formula). One can do even better with so-called *quadrature formulae*: collections of points with weights, $(\lambda_i, u_i)_{i=1}^M$, where $u_i \in [0,1]$ and $\lambda_i \geq 0$ with $\sum_{i=1}^M \lambda_i$, that are chosen in such a way that the quadrature approximation

$$\int_0^1 f(u)\,\mathrm{d}u \approx \sum_{i=1}^{M} \lambda_i f(u_i) \tag{4}$$

is exact for polynomials f up to some maximum degree, which is also called the *degree* D of the quadrature formula. Equivalently, a collection (λ_i, u_i) is a quadrature of degree D if, thought of as an (atomic) probability measure, it matches the moments of a uniform random variable up to order D. In more than one dimension, this is called a *cubature formula* (of degree D). Simple, well-known examples of quadrature formulae include the Trapezoidal rule (degree 1) or Simpson's rule (degree 3). The error of the approximation

(4) depends on the regularity of f: if it has, say, k well-behaved derivatives, then the error is at most of order D^{-k} [3, eqs. (1) and (2)].

Cubatures on Wiener Space This analogy between estimating statistics of SDE's (2) and numerical integration (3) naturally motivates the question whether the idea of cubature formulae can be extended to the setting of Brownian motion. That is, can one find deterministic collections of weighted paths, $(\lambda_i, \omega_i)_{i=1}^M$, where $\lambda_i \geq 0$ with $\sum_{i=1}^M \lambda_i = 1$ and $\omega_i \colon [0,T] \to \mathbb{R}^d$, that match, in some appropriate sense, the "moments" of Brownian motion up to some degree D. Then, the Monte Carlo estimate (2) could be replaced by

$$\mathbb{E}[f(y_T)] \approx \sum_{i=1}^{M} \lambda_i f(y_T(\omega_i)). \tag{5}$$

This idea was first formalised in the seminal 2004 paper by Lyons and Victoir [19]. They suggest that, within the context of cubatures, the appropriate generalised notion of a (random) path's "moment" is (the expectation of) its *signature*: the signature $\mathrm{S}(\omega)$ of a path $\omega = (\omega^{(1)}, \dots, \omega^{(d)}) \colon [0,T] \to \mathbb{R}^d$ is the collection of iterated integrals of the form

$$\mathrm{S}_{0,T}^{(i_1 \dots i_k)}(\omega) = \int_{0<u_1<\dots<u_k<T} \mathrm{d}\omega_{u_1}^{(i_1)} \dots \mathrm{d}\omega_{u_k}^{(i_k)}.$$

Here $i_j \in \{0, \dots, d\}$, where by convention $\omega^{(0)}(u) = u$. The *degree* of the signature term is $k + |\{j \colon i_j = 0\}|$, that is zero-indices (which correspond to the time coordinate) are counted double. They define a *Cubature on Wiener Space* of degree D to be such a collection $(\lambda_i, \omega_i)_{i=1}^M$ of weighted continuous paths $\omega_i \colon [0,T] \to \mathbb{R}^d$ of finite variation such that they match, thought of as an atomic probability measure on $C([0,T], \mathbb{R}^d)$, the expectation of the Brownian signature for all terms with degree at most D. Here the Brownian signature is computed in the sense of the Stratonovich integral, and *Wiener space* refers to the space $C([0,T], \mathbb{R}^d)$ of continuous paths $[0,T] \to \mathbb{R}^d$ equipped with the Wiener measure—the law of Brownian motion. They proceed to show that the error of the cubature approximation (5), under suitable smoothness assumptions on the drift and diffusion coefficients of the SDE, is of order $T^{(D+1)/2}$ for small time intervals T. They give explicit constructions of cubature formulae of degree $D \leq 5$, and prove (non-constructively) that one exists for every d and D. Further constructions have been found in the subsequent literature [14, 28, 15, 21, 12], however none have exceeded $D = 7$.[1]

An idea based on the local error estimate of Lyons and Victoir [19] is to subdivide the interval $[0,T]$ into, say, N equally sized subintervals, and apply a cubature $(\lambda_i, \omega_i)_{i=1}^M$

[1] In the one-dimensional setting, higher order cubature formulae have technically been constructed [14], but they consist of high order "rough paths"; solving the SDE (1) driven by such a path involves the computation of hundreds to thousands of high order derivatives (up to eleventh order) of the drift and diffusion vector fields, as well as their evaluation at every step of the numerical solver. Even if the derivatives were efficiently computed by a symbolic program, solving (1) for one of these paths would take up to thousands of times longer than for an ordinary (e.g. piecewise linear) path.

(scaled to an interval of length T/N) on each of them. The local error estimate then gives a global error bound of order $N^{-(D-1)/2}$ (for fixed T). This is equivalent to using (5) directly on the interval $[0,T]$ with the *product cubature* $(\lambda_{i_1} \cdot \ldots \cdot \lambda_{i_N}, \omega_{i_1} \odot \ldots \odot \omega_{i_N})_{i_1,\ldots,i_N=1}^{M}$, where $\odot$ denotes path concatenation. The problem with this idea is that the size of this product cubature grows exponentially in the number of sub-divisions N, while the error only decays polynomially in N. This amounts to a poly-logarithmic decay of the error as a function of the number of summands in the approximation (5), much worse than Monte Carlo. Our primary contribution is two-fold:

(i) We prove that the error of the cubature approximation (5) goes to zero *globally*, for arbitrary $T > 0$ and without product constructions, as a function of the degree of the cubature, under suitable assumptions on the SDE coefficients and the cubature (see Theorems I.3.4 and I.3.6).

(ii) We introduce a scalable, parallelisable, and GPU-optimised algorithm called ARCANE that automatically constructs cubature formulae of arbitrary dimension and degree. It recovers the previous state of the art of $D = 7$ in seconds, and constructs cubatures of degree up to $D = 17$ within a less than two hours on low-end compute.

Based on our theoretical result, we apply our high-degree cubature formulae *globally*, that is we use (5) directly without subdivision, and demonstrate empirically that the resulting approximation error is orders of magnitudes smaller than the Monte Carlo error obtained from an equivalent number of sample paths, consistently across several different example SDEs. See Fig. 1 for the CIR model, and Fig. 2 for an illustration of our degree 17 cubature formula.

Chapter II. Genealogies of Spatial Population Models

The *genealogy* of a set of individuals is their joint ancestral tree. This is best illustrated by example: Figure 3 (left) shows a simple population of constant size five; each row represents one generation, and arrows indicate parenthood. (Note that every individual in this population has exactly one parent. This is true for haploid organisms, but it also holds in diploid individuals like humans on the level of individual genes.) Then the genealogy of the three individuals sampled from the most recent generation (the three marked individuals in the bottom row) can be obtained by following the arrows in reverse. This process of *ancestral lineages* evolving backwards in time and merging whenever a common ancestor is reached is also called a *coalescent process*, and can be more abstractly represented as a tree such as the one on the right of Figure 3.

Describing such processes formally requires some mathematical setup. First of all, each of the lineages will require some kind of label. There are many valid choices, but a common and convenient one is to use disjoint set-valued labels, because there is a natural

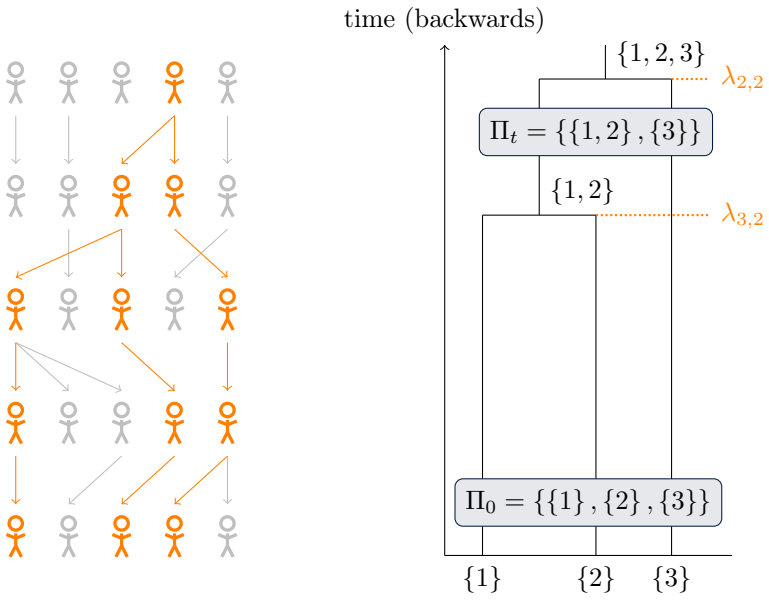

Figure 3: Illustration of a simple population and the genealogy of a sample from it (left), and the corresponding abstract coalescent process (right).

way of joining them whenever lineages coalesce (by taking their union). For example, if we label the three initial lineages $\{1\}$, $\{2\}$, and $\{3\}$, then the lineage resulting from the merge of the first two has the label $\{1,2\}$, and the common ancestor of all three is $\{1,2,3\}$. Then the *state* of the process Π_t at any given time is the set of active lineages, see Figure 3 (right). This motivates the following definition. Let $\mathcal{P}$ denote the set of partitions of finite subsets of $\mathbb{N}$.

Definition. A *coalescent* is a $\mathcal{P}$-valued homogeneous Markov process $(\Pi_t)_{t\geq 0}$ whose transitions can be obtained by merging partition elements, and which is independent of particle labels (equivalently, $(|\Pi_t|)_{t\geq 0}$ is Markov).

As a time-homogeneous Markov process in a discrete state space, the law of a coalescent is uniquely determined by its transition rates. The most general possible transition a coalescent could undergo would be the simultaneous merging of $m \in \mathbb{N}$ disjoint groups of lineages of sizes k_1 to k_m into one each. See Figure 4 for an illustration. We call such an event a $\vec{k} = (k_1, \ldots, k_m)$-merger. The rate at which each of the $\binom{n}{k_1 \ldots k_m (n-\sum_i k_i)}$ possible $\vec{k}$-mergers happens while there are currently n lineages is denoted $\lambda_{n,\vec{k}}$. For example, the transition rates of the two merge events in Figure 3 (right) are $\lambda_{3,2}$ and $\lambda_{2,2}$.

We now have a clean definition of a coalescent process, but nothing in that definition actually stipulates that this process must actually describe the genealogies of some population (model). A natural necessary condition is sampling consistency:

Definition. A coalescent is called *sampling consistent* if the restriction of the process started from $n+1$ particles to a subsample of size n has the same law as the process started directly from the subsample.

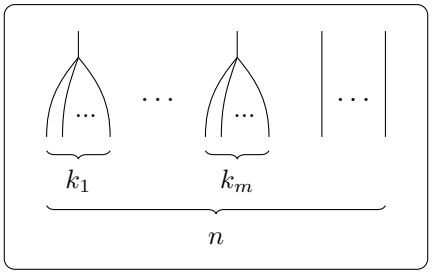

Figure 4: Illustration of the most general kind of merge event, a $\vec{k} = (k_1, \ldots, k_m)$-merger. The rate of each such merge event while there are currently n lineages is denoted $\lambda_{n,\vec{k}}$.

A full characterisation of sampling consistency for coalescents is known, discovered first in the case of no simultaneous mergers [8, 23, 26], and then in the general case [22, 25]. If there are no simultaneous mergers, then a coalescent is sampling consistent if and only if there exists a finite measure Λ on $[0, 1]$ such that

$$\lambda_{n,k} = \int_0^1 x^{k-2}(1-x)^{n-k}\Lambda(\mathrm{d}x), \qquad 2 \leq k \leq n.$$

It is called the Λ-*coalescent*. If simultaneous mergers are not excluded, then the coalescent is characterised instead through a more complicated formula by a finite measure Ξ on the infinite simplex. Furthermore, it turns out that sampling consistency is not only necessary but also sufficient. Indeed, every Λ- and Ξ-coalescent describes the genealogies of some (stationary, constant-size) population model [27, 22].

With this background in mind, the motivation behind this project is this: the Λ- and Ξ-coalescents are widely successful, and have been shown to arise ubiquitously across many different types of population models. However, their usefulness is ultimately limited by the fact that they treat all individuals as *indistinguishable*. Real populations evolve in (geographical and type) space, and knowing about the spatial locations of sampled individuals crucially shapes their genealogies.

In this project, we introduce the *Brownian spatial coalescent*, a coalescent process in which lineages evolve in continuous space of any dimension, and in the most general case of simultaneous multiple mergers. One of its key novelties is that it enjoys an axiomatic characterisation in terms of (a suitable spatial variant of) sampling consistency in complete analogy to the non-spatial Ξ-coalescents. In fact, there turns out to be a one-to-one correspondence, leading to what we will call the *Brownian spatial* Ξ*-coalescent*. Furthermore, each Brownian spatial coalescent describes the genealogies of a suitable population model in which individuals evolve as Brownian motions; for example, the genealogies of the well-known Fleming-Viot superprocess are described by the Brownian spatial Kingman coalescent (that is, the Brownian spatial Λ-coalescent where Λ is a point mass at zero). Another interesting result of this project is a representation of the joint evolution of lineages in the Brownian spatial coalescent as Brownian motions with a

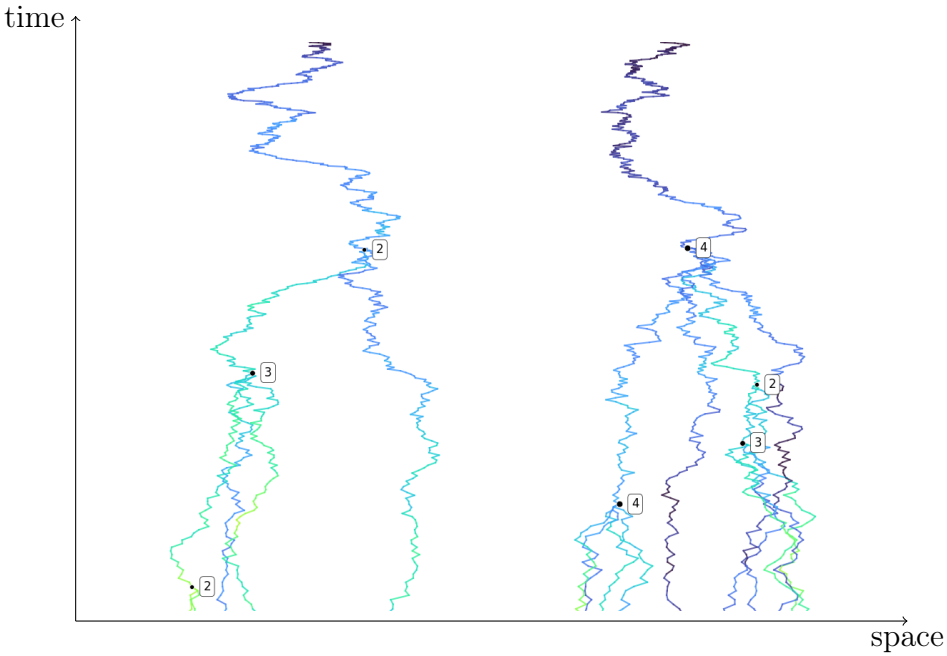

Figure 5: Two simulations of the Brownian spatial Λ-coalescent with $\Lambda = [0, 1]$, one started from five lineages (left), and one from ten (right).

coupled drift. It is well-known that Brownian motions never meet in dimensions $d \geq 2$, hence the coupled drift must be singular. Indeed, we find the following:

(i) For a binary merger, the drift pulls each particle in the direction of the other, with strength of order $\frac{d-1}{r}$, where r is the Euclidean separation of the particles.

(ii) For a ternary merger, the drift pulls each particle in the direction of the *midpoint* of the other two, and the strength of the drift is singular as a function of the largest pairwise distance amongst the three particles.

(iii) For two simultaneous binary mergers, i.e. an event where particles 1 and 2 merge at the same time as 3 and 4, if the locations of the lineages are x_i for $i = 1, 2, 3, 4$, then the drift pulls the $2d$-dimensional vector (x_1, x_3) in the direction of (x_2, x_4) and vice versa, with a strength proportional to $\frac{2d-1}{r}$, where r is their $2d$-dimensional Euclidean separation.

The Brownian spatial coalescent is, to our knowledge, the first spatial coalescent model in which lineages move as Brownian motions and can meet in multiple and simultaneous merge events in arbitrary dimensions, in fact meet at all in dimensions $d \geq 2$. Two simulations of the Brownian spatial Λ-coalescent with $\Lambda = [0, 1]$ are in Figure 5. Coalescence events are labelled by the number of merging lineages.

Chapter III. Repulsion-Diffusion Equations

The phenomenon that has been infamously dubbed the "pain in the torus" by Felsenstein in 1985 [11] can be summarised as follows: *in one and two dimensions, spatial population models with independent critical branching and independent diffusive spatial motion concentrate in increasingly few well-separated clumps, until they eventually die out.* The latter fact is simply due to the criticality of the branching, and can be overcome by conditioning on survival or by starting with an infinite population. The "clumping" issue is more inherently connected to the interplay of the branching mechanism and the diffusive motion in low dimensions. A crude heuristic goes as follows: suppose we start with some number N of particles, which we can think of as being the ancestors of N separate families. Then at some later time t, the critical, independent branching will have eliminated all but order $O(N/t)$ of the families, but those that survive will consist of order $O(t)$ offspring. The diffusive spatial motion will have spread out each of the families over an area with radius of order $O(\sqrt{t})$, resulting in a population density of order $O(t^{1-d/2})$. This heuristic gives the correct asymptotic order except if $d = 2$, where the correct result is $O(\log t)$. Hence, if $d \geq 2$ then the population density of surviving families will grow without bound, while space between these "clumps" grows increasingly empty. Until, at some time $t \gtrsim N$, all families will have gone extinct.

As a model for a spatially evolving population (for which $d = 2$ is the most natural dimension), this is very unrealistic behaviour, and a better model should reflect stable long-term population dynamics. One of the most obvious reasons this clumping phenomenon does not occur in nature is that real individuals do not behave independently from surrounding individuals. Indeed, a high population density leads to resource scarcity and decreases the average number of offspring, and causes migration away from the overcrowded area. The former effect has already been studied and was shown to lead to stable long-term dynamics [9]. The latter, however, has not yet been studied in this context. A natural way to implement this is to introduce a pairwise repulsion between individuals: we let an individual at position x experience a drift $-\nabla W(x-y)$ if there is another individual at y. Here $W\colon \mathbb{R}^d \to \mathbb{R}$ is the *interaction potential*, which must be decreasing in order for the interaction to be repulsive. Motivated by the heuristic above, we then aim to study the effect of the repulsion on the asymptotic population density of "surviving families". We motivate how this question can be simplified to studying the long-term asymptotics of a partial differential equation (PDE) of the form

$$\partial_t \rho = \frac{1}{2}\Delta\rho + \nabla \cdot (\rho \nabla W \star \rho) + f, \tag{6}$$

where $\rho\colon \mathbb{R}^d \to [0,\infty)$ represents a population density, the term $\frac{1}{2}\Delta\rho$ represents the diffusive motion of the particles, the term involving W represents the pairwise repulsion, and finally the constant term $f\colon \mathbb{R}^d \to [0,\infty)$ represents a constant immigration of mass.

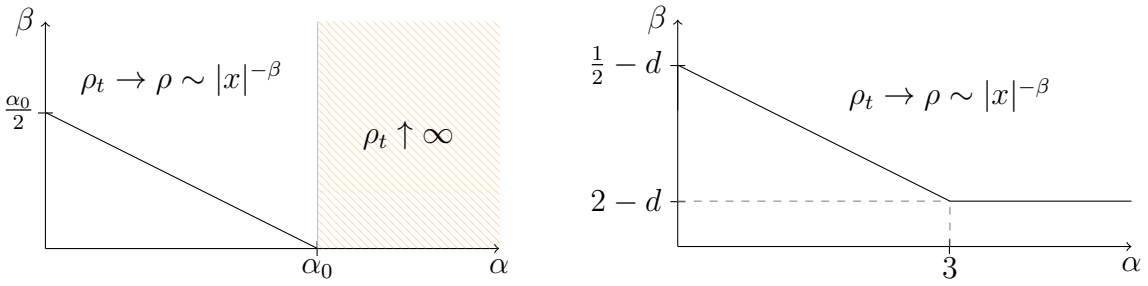

Figure 6: Illustration of the regimes of our conjecture for $d = 1, 2$ (left) and $d \geq 3$ (right). Here, $\alpha_0 = 1$ if $d = 1$ and $\alpha_0 = 3$ if $d = 2$.

This arises from the fact that we study a critically branching process *conditioned on survival*. Otherwise it would be guaranteed to die out, and what we are interested in is the asymptotic behaviour of the population's population density in the unlikely event that it *does* survive for a long time. It turns out that conditioning a population process with critical, independent branching on survival equates to a constant immigration of new living individuals (see e.g. Evan's immortal particle [10]). This immigration is represented in the PDE by f.

With this motivation in mind, we seek to find conditions on the repulsion W under which the solution to (6) stays globally bounded. We show that a *sufficient* condition is if the strength of the repulsion is of order $|\nabla W(x)| \sim |x|^{1-d}$, but we do not believe this to be sharp. In fact, some of our results give rise to the following conjecture.

Conjecture. Suppose that $|\nabla W(x)| \sim |x|^{-\alpha}$ for some $\alpha \geq 0$.

(i) $d = 1$.

(i.1) If $\alpha \in [0, 1)$, then ρ_t converges to a stationary distribution with $\rho(x) \sim |x|^{-\frac{1-\alpha}{2}}$.

(i.2) If $\alpha > 1$, then $\rho_t \uparrow \infty$ locally uniformly.

(ii) $d = 2$.

(ii.1) If $\alpha \in [0, 3)$, then ρ_t converges to a stationary distribution with $\rho(x) \sim |x|^{-\frac{3-\alpha}{2}}$.

(ii.2) If $\alpha > 3$, then $\rho_t \uparrow \infty$ locally uniformly.

(iii) If $d \geq 3$, then ρ_t converges to a stationary distribution with

$$\rho(x) \sim \begin{cases} |x|^{\frac{1+\alpha}{2}-d}, & \alpha \in [0, 3), \\ |x|^{2-d}, & \alpha \geq 3. \end{cases}$$

See Figure 6 for an illustration of the different regimes.

Bibliography

[1] J. Pablo Arenas-López and Mohamed Badaoui. "A Fokker–Planck equation based approach for modelling wind speed and its power output". In: *Energy Conversion and Management* 222 (2020).

[2] Jordi Badosa et al. "Day-Ahead Probabilistic Forecast of Solar Irradiance: A Stochastic Differential Equation Approach". In: *Renewable Energy: Forecasting and Risk Management*. 2018, pp. 73–93.

[3] H. Brass. "Error Bounds Based on Approximation Theory". In: *Numerical Integration: Recent Developments, Software and Applications*. Ed. by Terje O. Espelid and Alan Genz. Dordrecht: Springer Netherlands, 1992, pp. 147–163. ISBN: 978-94-011-2646-5. DOI: 10.1007/978-94-011-2646-5_12.

[4] Yongli Cai et al. "A stochastic SIRS epidemic model with infectious force under intervention strategies". In: *Journal of Differential Equations* 259.12 (2015), pp. 7463–7502.

[5] Loren Cobb. *Mathematical Frontiers Of The Social And Policy Sciences*. Routledge, New York, 1981.

[6] John C. Cox, Jonathan E. Ingersoll, and Stephen A. Ross. "A Theory of the Term Structure of Interest Rates". In: *Econometrica* 53.2 (1985), pp. 385–407.

[7] Peter Czuppon and Arne Traulsen. "Understanding evolutionary and ecological dynamics using a continuum limit". In: *Ecology and Evolution* 11.11 (2021), pp. 5857–5873.

[8] Peter Donnelly and Thomas G Kurtz. "Particle representations for measure-valued population models". In: *The Annals of Probability* 27.1 (1999), pp. 166–205.

[9] Alison M. Etheridge. "Survival and extinction in a locally regulated population". In: *Ann. Appl. Probab.* 14.1 (2004), pp. 188–214. DOI: 10.1214/aoap/1075828051.

[10] Steven N. Evans. "Two representations of a conditioned superprocess". In: *Proc. R. Soc. Edinb., Sect. A, Math.* 123.5 (1993), pp. 959–971. DOI: 10.1017/S0308210500029619.

[11] Joseph Felsenstein. "A Pain in the Torus: Some Difficulties with Models of Isolation by Distance". In: *The American Naturalist* 109.967 (1975), pp. 359–368. ISSN: 00030147, 15375323.

[12] Emilio Ferrucci et al. *High-degree cubature on Wiener space through unshuffle expansions*. 2024. arXiv: 2411.13707 [math.PR].

[13] Daniel T Gillespie. "The chemical Langevin equation". In: *The Journal of Chemical Physics* 113.1 (2000), pp. 297–306.

[14] Lajos Gergely Gyurkó and Terry Lyons. "Efficient and Practical Implementations of Cubature on Wiener Space". In: *Stochastic Analysis 2010*. Springer Berlin Heidelberg, 2011, pp. 73–111.

[15] Satoshi Hayakawa and Ken'ichiro Tanaka. "Monte Carlo construction of cubature on Wiener space". In: *Japan Journal of Industrial and Applied Mathematics* 39.2 (2022), pp. 543–571.

[16] Emil B. Iversen et al. "Probabilistic forecasts of solar irradiance using stochastic differential equations". In: *Environmetrics* 25.3 (2014), pp. 152–164.

[17] Haiquan Li et al. "Identification of a Stochastic Dynamic Model for Aircraft Flight Attitude Based on Measured Data". In: *International Journal of Aerospace Engineering* 2023.1 (2023).

[18] Ying Liu et al. "A stochastic approximation method for probability prediction of docking success for aerial refueling". In: *Applied Soft Computing* 103 (2021).

[19] Terry Lyons and Nicolas Victoir. "Cubature on Wiener Space". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 460.2041 (2004), pp. 169–198.

[20] Yoshihiro Maki and Hideo Hirose. "Infectious Disease Spread Analysis Using Stochastic Differential Equations for SIR Model". In: *2013 4th International Conference on Intelligent Systems, Modelling and Simulation*. 2013, pp. 152–156.

[21] Anatoliy Malyarenko and Hossein Nohrouzian. "Testing Cubature Formulae on Wiener Space Versus Explicit Pricing Formulae". In: *Stochastic Processes, Statistical Methods, and Engineering Mathematics*. Cham: Springer International Publishing, 2022, pp. 223–248.

[22] Martin Möhle and Serik Sagitov. "A Classification of Coalescent Processes for Haploid Exchangeable Population Models". In: *The Annals of Probability* 29.4 (2001), pp. 1547–1562. DOI: 10.1214/aop/1015345761.

[23] Jim Pitman. "Coalescents With Multiple Collisions". In: *The Annals of Probability* 27.4 (1999), pp. 1870–1902. DOI: 10.1214/aop/1022874819.

[24] Hamood Ur Rehman et al. "Analysis of Brownian motion in stochastic Schrödinger wave equation using Sardar sub-equation method". In: *Optik* 289 (2023).

[25] Serik Sagitov. "Convergence to the coalescent with simultaneous multiple mergers". In: *Journal of Applied Probability* 40.4 (2003), pp. 839–854. DOI: 10.1239/jap/1067436085.

[26] Serik Sagitov. "The General Coalescent with Asynchronous Mergers of Ancestral Lines". In: *Journal of Applied Probability* 36.4 (1999), pp. 1116–1125. ISSN: 00219002.

[27] Jason R Schweinsberg. *Coalescents with simultaneous multiple collisions*. University of California, Berkeley, 2001.

[28] Yuji Shinozaki. "Construction of a Third-Order K-Scheme and Its Application to Financial Models". In: *SIAM Journal on Financial Mathematics* 8.1 (2017), pp. 901–932.

[29] Marc Vaisband et al. "Loss formulations for assumption-free neural inference of SDE coefficient functions". In: *Nature Partner Journal: Systems Biology and Applications* 11.22 (2025).

[30] Yasen Wang et al. "Data-Driven Discovery of Stochastic Differential Equations". In: *Engineering* 17 (2022), pp. 244–252.

Chapter I

Cubatures on Wiener Space

1 Introduction

Stochastic differential equations (SDEs) are commonly used to model continuous-time phenomena evolving under the influence of random noise. For example, SDEs with *scalar noise* (that is, driven by a single Brownian motion) have been applied across a range of disciplines—such as wind speed modelling [1], solar radiation forecasting [2, 23], epidemiology [7, 30], social sciences [10, chapter 2], mathematical finance [11], physics [35], population genetics [14], research engineering [45] and systems biology [42]. In addition, SDEs with *low-dimensional noise* (e.g. driven by d independent Brownian motions with $d \leq 3$) are natural models for physical objects whose motion exhibits random fluctuations. For example, SDEs have been used to model aircraft dynamics [24, 26].

We consider low-dimensional SDEs, defined by Stratonovich integration[1], of the form:

$$dy_t = \mu(y_t)\,dt + \sum_{i=1}^{d} \sigma_i(y_t) \circ dW_t^i, \tag{1.1}$$

where the initial condition $y_0 = \xi$ and solution $y = \{y_t\}_{t\in[0,T]}$ take values in $\mathbb{R}^e$, $W = (W^1, \cdots, W^d) = \{W_t\}_{t\in[0,T]}$ is a standard d-dimensional Brownian motion, and $\mu, \sigma_i : \mathbb{R}^e \to \mathbb{R}^e$ are the vector fields (referred to as the drift and diffusion respectively).

In practice, one is often interested in estimating key statistics of the solution to the SDE (1.1), such as its mean or variance, or more generally the expectation of $f(y_T)$ for some functional $f\colon \mathbb{R}^e \to \mathbb{R}$. This is traditionally achieved using Monte Carlo estimation, which takes the average over some number M of independently sampled solutions of (1.1), see Table 1.1 (left). However, by the Central Limit Theorem, the error of the Monte Carlo estimate is inverse proportional to the square root of the number M of samples. An accurate estimation can therefore require an enormous number of sample paths and be computationally very expensive—especially for complex SDEs.

[1] This is just a matter of choice; SDEs in the Stratonovich sense can be transformed into equivalent SDEs in the Itô sense and vice versa.

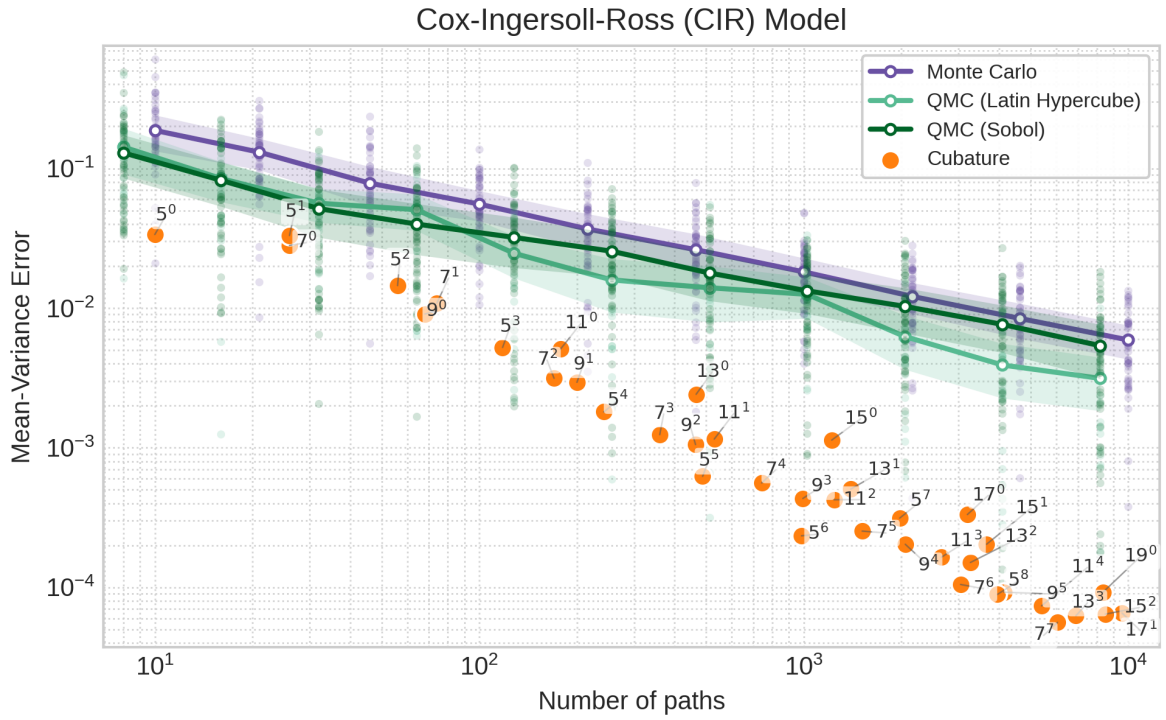

Figure 1.1: The SDE simulation error (a combination of the errors of the empirical mean and variance of the endpoint of the solution path, see (4.1)) of our cubature paths (orange) against an equivalent number of Monte Carlo paths (purple) and Quasi-Monte Carlo paths (green), evaluated on the CIR model, see (4.4).

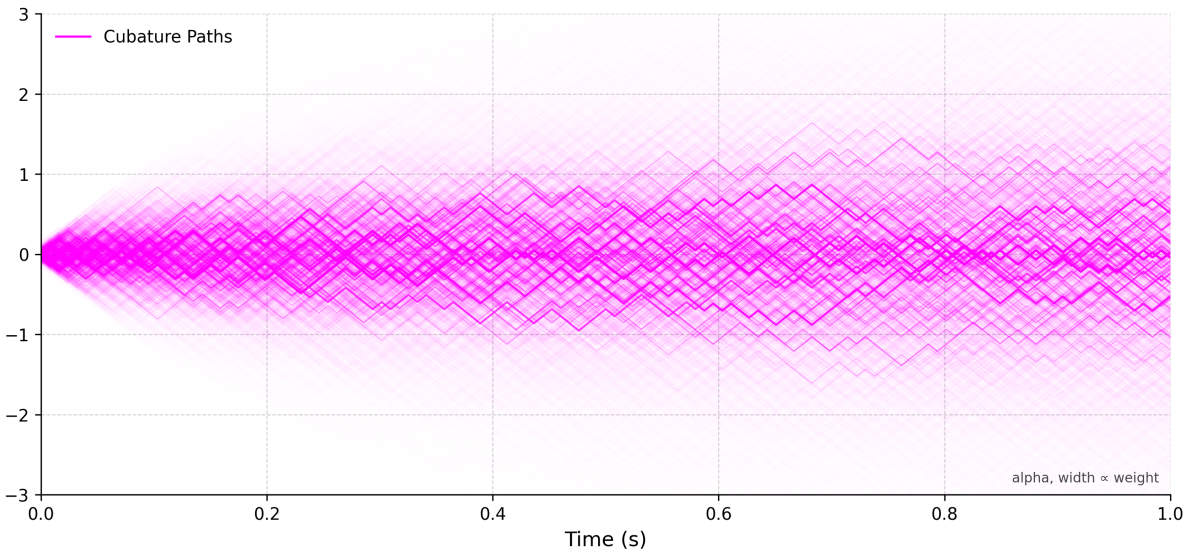

Figure 1.2: A visualisation of one of our largest cubature formulae, which has degree 17 and consists of 9574 paths.

	Monte Carlo	ARCANE (ours)
Estimate	$\mathbb{E}\big[f(y_T)\big] \approx \sum_{i=1}^{M} \lambda_i f\big(y_T(\omega_i)\big)$	(1.2)
Paths (ω_i)	randomly sampled Brownian paths	deterministic Brownian-like paths
Weights (λ_i)	uniform, $\lambda_i = \frac{1}{M}$	not necessarily uniform

Table 1.1: Comparison of the Monte Carlo and ARCANE estimators. Here, $y_T(\omega) \in \mathbb{R}^e$ denotes the solution at time T of the system (1.1) driven by a path $\omega\colon [0,T] \to \mathbb{R}^d$.

In this paper, we aim to reduce this estimation error by replacing the uniformly weighted, independently sampled Brownian paths in (1.2) with a carefully designed *deterministic* set of Brownian-like paths with (not necessarily uniform) associated probability weights, see Table 1.1 (right). These paths are piecewise linear so that (1.1) simply reduces to an ODE along each piece, which can be solved by standard methods. Experimentally, the estimation error using our paths is consistently several orders of magnitude lower than the error obtained by a similar number of Monte Carlo samples. An example of such a collection of Brownian-like paths produced by our methodology is illustrated in Figure 1.2, and the results of an experiment detailed in Section 4 is in Figure 1.1.

More precisely, we will propose a computational methodology called ARCANE[2], an efficient and scalable algorithm that produces what is known as *cubature formulae*: sets of weighted paths that match the "signature moments" [9] of Brownian motion up to some *degree* D. It was first observed by Lyons and Victoir in their seminal 2004 paper [28] that cubature formulae could be used in place of Monte Carlo estimation for simulating SDEs in the way we outlined in Table 1.1. They proved that cubature formulae exist for arbitrary Brownian dimension d and arbitrary degree D, but were only able to find explicit constructions for $D \leq 5$. Further constructions have been found in the subsequent literature [17, 37, 19, 31, 15], however none have exceeded $D = 7$. In the one-dimensional setting, higher order cubature formulae have technically been constructed [17], but they consist of high order "rough paths"; solving the SDE (1.1) driven by such a path involves the computation of hundreds to thousands of high order derivatives (up to eleventh order) of the drift and diffusion vector fields, as well as their evaluation at every step of the numerical solver. Even if the derivatives were efficiently computed by a symbolic program, solving (1.1) for one of these paths would take up to thousands of times longer than for an ordinary (e.g. piecewise linear) path.

Our ARCANE algorithm can automatically construct cubature formulae of arbitrary dimension and depth. They consist of piecewise linear paths and can therefore be used

[2] **A**lgorithm for **R**ecombination of **C**ubature formulae from Orthogonal **A**rrays that match **N**ested **E**xpected signatures

plug-and-play in place of random Brownian sample paths. Our algorithm is optimised for GPU and highly parallelisable; using ARCANE, we were able to construct a cubature with degree $D = 19$ in only a few GPU hours, and reproduce the previous state of the art at $D = 7$ in seconds. Using several real-world SDEs, we demonstrate that our ARCANE cubature formulae can achieve orders of magnitude more accuracy when estimating statistics for SDEs compared to traditional Monte Carlo simulation.

The core novelty behind ARCANE is a series of innovations that drastically reduce the run-time of a known and conceptually simple algorithm (from exponential to low-order polynomial in the relevant parameter, see Section 3 for details), combined with a highly optimized implementation of the core components of the algorithm in JAX [5]. Code for our ARCANE algorithm along with datasets containing the cubature formulae used in our experiments can be found at:

github.com/tttc3/FlipSampling

We summarise our key contributions below:

(1) We introduce ARCANE, a scalable, highly parallelisable and GPU-optimised algorithm for generating high-degree cubature formulae for SDE simulation.

(2) We publish both the cubature formulae that we constructed, as well as the code for the ARCANE algorithm, which can be used with more compute to construct even higher degrere cubatures.

(3) We demonstrate that cubature formulae obtained with this algorithm on low-end compute (a few GPU hours) already improve drastically over the previously best known cubature formulae, and lead to an SDE estimation error several orders of magnitude smaller compared with traditional Monte Carlo methods.

(4) We prove several theoretical results on the correctness of our approach, some of which may be of independent interest to members of the stochastic analysis and rough path theory communities.

The rest of this chapter is structured as follows. In Section 2, we introduce some useful background from rough path theory and the theory of cubatures. In Section 3, we explain our ARCANE algorithm in more detail and prove our theoretical results, and in Section 4, we present and discuss our empirical findings.

2 Background

This section contains relevant background on rough path theory (Section 2.1) and cubature formulae (Section 2.2), and introduces associated notation. None of the content is novel, so the reader familiar with the topics should feel free to skim or skip ahead.

2.1 Rough Path Theory

We give a brief motivation of and introduction to rough path theory. A more comprehensive introduction is [29].

Motivation

The goal of rough path theory is to study controlled differential equations (CDEs) of the form

$$\mathrm{d}Y_t = f(Y_t)\,\mathrm{d}X_t, \qquad Y_0 = \xi, \tag{2.1}$$

where $X = (X^{(0)}, X^{(1)}, \dots, X^{(d)})$ is the *control* or *signal* (living in $V \coloneqq \mathbb{R}^{d+1}$ for some $d \in \mathbb{N}$, where $X^{(0)}(t) = t$), and Y is the *response* (living in $W \coloneqq \mathbb{R}^m$ for some $m \in \mathbb{N}$). The solution, if it exists, is denoted $Y = I_f(X, \xi)$, and I_f is called the *Itô map* associated with f.

Recall that if $p \geq 1$ and $X\colon [0,1] \to V$ is continuous, then the *p-variation* of X is

$$\|X\|_p = \sup_{\mathcal{D}} \Big[\sum |X_{t_{i+1}} - X_{t_i}|^p\Big]^{1/p}, \tag{2.2}$$

where the supremum is over all subdivisions of $[0,1]$ (not to be confused with the L^p norm which we would denote by $\|\cdot\|_{L^p}$). If X is of bounded (1-)variation and f is sufficiently smooth (say Lipschitz), then classical theory ensures existence and uniqueness of the solution to (2.1), as well as continuity of the Itô map. More precisely, if X^n, $n \in \mathbb{N}$, and X are bounded variation paths, and $X^n \to X$ in 1-variation, then $I_f(X_n, \xi) \to I_f(X, \xi)$ in 1-variation for any initial condition ξ. However, in reality signals are often much more rough; a famous example is Brownian motion, whose paths almost-surely have finite p-variation if and only if $p > 2$, but the classical theory of controlled differential equation based on the Young integral fails as soon as $p \geq 2$. Equations controlled specifically by Brownian motion can still be made sense of in an inherently stochastic way using Itô calculus, but the aim of rough path theory is to develop a flexible framework in which to study (2.1) for a broad class of rough signals in a *pathwise* (i.e. deterministic) manner.

To summarise, what we would like to do is to extend the map I_f, say for a fixed, smooth f, to a space of functions that includes rough signals such as continuous paths that have finite p-variation only for $p > 2$, in such a way that I_f is continuous w.r.t. some complete metric. A naïve approach, however, is doomed to fail: there exist many examples of sequences of smooth paths (x_n^1) and (x_n^2) that converge to the same path x uniformly and in p-variation say for some $p > 2$, but such that $I_f(x_n^1)$ and $I_f(x_n^2)$ have different limits.

The problem, it turns out, is not due to a wrong choice of metric, or an inherent impossibility of (deterministically) integrating against rough signals, but rather of conceptual nature: the mistake lies in thinking of a rough signal as "just" a continuous path

$[0,1] \to V$; to characterise its behaviour as a driving signal in (2.1), more information is required. It turns out that, in many of the examples alluded to at the beginning of the paragraph, (x_n^1) and (x_n^2) do in fact both have well-defined but *distinct* limits in an appropriate space of so-called rough paths that both "live over" the same classical path $x\colon [0,1] \to V$.

To motivate what additional information may be required to characterise such a rough path, consider for a moment the case of CDEs driven by paths of bounded variation. If f is sufficiently regular, they can be solved through Picard iteration—i.e. letting $Y_t^0 = Y_0$ and iteratively solving $\mathrm{d}Y_t^{n+1} = f(Y_t^n)\,\mathrm{d}X_t$—which reveals that Y_t is actually a function of the collection of iterated integrals of the signal X on $[0,t]$, that is its *signature*

$$\mathrm{S}_{0,t}(X) \coloneqq \big(1, \mathrm{S}_{0,t}^1(X), \mathrm{S}_{0,t}^2(X), \ldots\big) \in T((V)) = \bigoplus_{k=0}^{\infty} V^{\otimes k},$$

where $V^{\otimes k}$ is the k-fold tensor product of V with itself, and

$$X_{0,t}^k \coloneqq \mathrm{S}_{0,t}^k(X) \coloneqq \int_{0<u_1<\ldots<u_k<t} \mathrm{d}X_{u_1} \otimes \ldots \otimes \mathrm{d}X_{u_k} \in V^{\otimes k}$$

can be thought of as compact notation for the collection of k-fold iterated integrals of $X = (X^{(0)}, X^{(1)}, \ldots, X^{(d)})\colon [0,1] \to V = \mathbb{R}^{d+1}$, i.e.

$$X_{0,t}^k = \left(X_{0,t}^{k,I}\right)_{I=(i_1\ldots i_k)\in[d]^k} \quad \text{where} \quad X_{0,t}^{k,I} = \int_{0<u_1<\ldots<u_k<t} \mathrm{d}X_{u_1}^{(i_1)} \ldots \mathrm{d}X_{u_k}^{(i_k)}.$$

Note that the first level of the signature are just the increments $X_{0,t}^1 = X_t - X_0$, which determine the path up to a translation that doesn't affect its behaviour as a signal.

From this perspective, it seems natural to think of a signal not as a path in V, but as a path in $T((V))$. If the underlying path has bounded variation, then this distinction is insubstantial because all levels of the signature are determined by the first level (the increments) through classical integration. But now say that we have a continuous path $X\colon [0,1] \to V$ with finite p-variation only for $p \in (2,3)$, such as a realisation of Brownian motion. Then we can define $X_{s,t}^1 = X_t - X_s$, but for higher levels we run into the old problem that iterated integrals of X cannot be defined classically. However, it turns out that if we *assign* a second level X^2 in such a way that

$$X_{s,t}^2 = X_{s,u}^2 + X_{u,t}^2 + X_{s,u}^1 \otimes X_{u,t}^1, \qquad \forall s < u < t, \tag{2.3}$$

an identity that holds automatically for the signature of a bounded variation path by standard properties of the integral, and such that X^2 has bounded p-variation in an appropriate sense, then there is a *unique* way of extending the signature to all higher levels in such a way that the natural generalisation of (2.3) holds at all levels (and that the entire signature has bounded p-variation in an appropriate sense). Such an object is

then called a (p-)*rough path* over the original path $X\colon [0,1] \to V$, and, crucially, there can be many rough paths over the same classical path X. The central result of rough path theory is that there is a natural way to define integration against and CDEs driven by rough paths, in such a way that the Itô map is continuous with respect to a suitably chosen metric.

Rough paths and the universal limit theorem

Let $\triangle = \{(s,t)\colon 0 \le s \le t \le 1\}$, and $X\colon \triangle \to T((V))$ be a continuous mapping. To avoid confusion, we will use lower case letters such as x to denote classical paths $[0,1] \to V$ in this section. The generalisation of (2.3) to higher levels is obtained by writing (the components of) $X_{s,t}$ in terms of $X_{s,u}$ and $X_{u,t}$ for some $s < u < t$, in the case where X is the signature of a bounded variation path. This can be done using standard properties of integrals, and the result can be written compactly as

$$X_{s,t} = X_{s,u} \otimes X_{u,t}, \qquad \forall 0 \le s \le u \le t \le 1, \tag{2.4}$$

also known as *Chen's identity*. The zero'th level of (2.4) is $1 = 1$, the first level is $X^1_{s,t} = X^1_{s,u} + X^1_{u,t}$, and the second level is (2.3).

Definition 2.1. Let $p \ge 1$. A *p-rough path* is a continuous mapping $X\colon \triangle \to T((V))$ that satisfies Chen's identity (2.4) and has finite p-variation in the sense that

$$\sup_{\mathcal{D}} \sum \|X^k_{s_i,s_{i+1}}\|^{p/k} < \infty \tag{2.5}$$

for all $k \in \mathbb{N}$, where the supremum is over finite subdivisions of $[0,1]$.

Note that (2.4) implies that the first level of a p-rough path is always the increments of a continuous path $[0,1] \to V$, which by (2.5) has finite p-variation. Note further that if $x\colon [0,1] \to V$ is a continuous path with finite p-variation for some $p \in [1,2)$, then its signature $\mathrm{S}(X) = (1, X^1, X^2, \ldots)$ is classically defined in terms of iterated (Young) integrals, and is a p-rough path according to Definition 2.1. Denote by $T^{(n)}(V) = \bigoplus_{k=0}^n V^{\otimes k}$ the truncated tensor algebra over V of degree n, equipped with the product $\boldsymbol{ab} = \boldsymbol{c}$ where $c_k = a_0 \otimes b_k + \ldots + a_k \otimes b_0$.

Theorem 2.2 (Extension Theorem)**.** *Let $p \ge 1$, and let $X\colon \triangle \to T^{(\lfloor p \rfloor)}(V)$ be a continuous mapping that satisfies Chen's identity and has finite p-variation in the sense that* (2.5) *holds for all $k \le \lfloor p \rfloor$. Then X has a unique extension to a p-rough path $X\colon \triangle \to T((V))$.*

See Theorem 3.7 in [29]. In particular, a p-rough path is uniquely determined by its truncation at level $\lfloor p \rfloor$, and we sometimes identify it with this truncation. Denote by

$\Omega_p(V)$ the space of p-rough paths, which is a complete metric space when equipped with the *p-variation metric* [27, Sect. 3.3.1]

$$d_p(X,Y) = \max_{i=1,\dots,\lfloor p \rfloor} \sup_{\mathcal{D}} \left(\sum \| X^k_{s_i,s_{i+1}} - Y^k_{s_i,s_{i+1}} \|^{\frac{p}{k}} \right)^{\frac{1}{p}}.$$

Note that $\Omega_1(V)$ can be identified with the space of (classical) continuous bounded variation paths $[0,1] \to V$ (modulo constant translations) with the 1-variation metric.

Definition 2.3. The space $G\Omega_p(V)$ of *geometric rough paths* is the closure of $\Omega_1(V)$ in $\Omega_p(V)$ w.r.t. the p-variation metric. It is in particular complete.

The inclusion $G\Omega_p(V) \subset \Omega_p(V)$ is strict except if $p < 2$. A simple example of a 2-rough path that is not geometric is $X_{s,t} = (1, 0, (t-s)\boldsymbol{a})$ for any $\boldsymbol{a} \in V^{\otimes 2} \setminus \{0\}$. It is easy to check that Chen's identity holds and X has finite p-variation for any $p \geq 2$, so X is a 2-rough path. But for any bounded variation path, and therefore any geometric p-rough path by approximation in d_p, it must hold that the symmetric part of $X^2_{s,t}$ is $\frac{1}{2}(X^1_{s,t})^{\otimes 2}$, but this is not true except if $\boldsymbol{a} = 0$.

Remark 2.4. An alternative way to introduce (geometric) rough paths directly, without any reference to tensor algebras and Chen's relation, would have been to say that the space of (geometric) p-rough paths is the completion of the space of continuous paths $[0,1] \to V$ of bounded variation w.r.t. a metric in which two paths are close if they and their first $\lfloor p \rfloor$ iterated integrals are close in p-variation.

The following theorem is the main result and justification of rough path theory. A function $f\colon A \to B$ between two finite-dimensional Banach spaces A, B is said to be $\mathrm{Lip}(\gamma)$ for some $\gamma > 0$ if it is bounded, $\lfloor \gamma \rfloor$ times continuously differentiable and such that the $\lfloor \gamma \rfloor$'th derivative is Hölder continuous with exponent $\gamma - \lfloor \gamma \rfloor$. Recall that $V = \mathbb{R}^{d+1}$ and $W = \mathbb{R}^m$.

Theorem 2.5 (Universal Limit Theorem)**.** *Let $p \geq 1$ and $\gamma > p$, and $f\colon W \to \mathcal{L}(V,W)$ be a Lip(γ) function. Then for all $X \in G\Omega_p(V)$ and $\xi \in W$, the equation*

$$\mathrm{d}Y_t = f(Y_t)\,\mathrm{d}X_t, \quad Y_0 = \xi,$$

has a unique solution in $G\Omega_p(W)$ (in an appropriate sense), and the map

$$I_f\colon G\Omega_p(V) \times W \to G\Omega_p(W)$$

that sends (X,ξ) to Y is continuous and equal to the unique continuous extension of the classical Itô map $I_f\colon \Omega_1(V) \times W \to \Omega_1(W)$ (which is continuous in p-variation).

A more precise statement and a proof are in Section 5.3 of [29].

Brownian motion as a rough path

Since the main application of rough path theory is to differential equations driven by Brownian motion, it will be useful to consider for a moment the special case of a geometric p-rough path when $p \in [2,3)$. As mentioned before, signatures of bounded variation paths and therefore geometric p-rough paths have the property that the symmetric part $S_{s,t} \in V^{\otimes 2}$ defined by

$$S_{s,t}^{(ij)} = \frac{1}{2}\left(X_{s,t}^{2,(ij)} + X_{s,t}^{2,(ji)}\right), \qquad (s,t) \in \triangle,$$

satisfies $S_{s,t} = \frac{1}{2}(X_{s,t}^1)^{\otimes 2}$. If we denote the anti-symmetric part by $A_{s,t} \in V^{\otimes 2}$, i.e.

$$A_{s,t}^{(ij)} = \frac{1}{2}\left(X_{s,t}^{2,(ij)} - X_{s,t}^{2,(ji)}\right), \qquad (s,t) \in \triangle,$$

then $X^2 = S + A$, which means that X as a whole is fully determined by its increments X^1 and the anti-symmetric part A of its second level,

$$X = \left(1, X^1, \frac{1}{2}(X^1)^{\otimes 2} + A\right).$$

The tensor A can be interpreted as the *area* of X, in the following sense: if X is the signature of a path $x \in \Omega_1(V)$, then

$$A_{s,t}^{ij} = \frac{1}{2} \iint\limits_{s \le u_1 \le u_2 \le t} \mathrm{d}x_{u_1}^i \, \mathrm{d}x_{u_2}^j - \mathrm{d}x_{u_1}^j \, \mathrm{d}x_{u_2}^i,$$

which is the area enclosed between the bounded variation path (x^i, x^j) and its secant on the time interval $[s,t]$. The area is "signed", e.g. the area enclosed by $t \mapsto (t, \sin(t))$ on $[0,\pi]$, $[\pi, 2\pi]$, and $[0, 2\pi]$ is, respectively, 2, -2, and 0. For genuine p-rough paths with $p \ge 2$, it is still common to refer to A, the asymmetric part of X^2, as its *area*.

Brownian motion is a random continuous path which almost-surely has finite p-variation for $p > 2$. By the considerations in the previous section, a (geometric) rough path over a fixed realisation of the Brownian increments is determined by the choice of its area tensor A, and this choice is not unique! If we choose A to be (a continuous version of) the *Lévy area* of Brownian motion, that is

$$A_{s,t}^{ij} = \iint\limits_{s \le u_1 \le u_2 \le t} \left(\mathrm{d}B_{u_1}^i \, \mathrm{d}B_{u_2}^j - \mathrm{d}B_{u_2}^i \, \mathrm{d}B_{u_1}^j\right), \qquad i,j \in [d],\ (s,t) \in \triangle,$$

where the integral is in the Itó or Stratonovich sense (which doesn't make a difference in this case), then we obtain what is called the *canonical Brownian rough path*. It can be obtained as the almost-sure limit of its piecewise linear approximation on an increasingly fine set of support points, so it is indeed geometric. It can further be thought of as the signature of Brownian motion where iterated integrals are calculated in the Stratonovich

sense. Indeed, if $X\colon \triangle \to T((V))$ is a canonical Brownian rough path over a Brownian motion B, then for every $k \in \mathbb{N}$ and $J = (j_1 \dots j_k) \in [d]^k$, and every $0 \le s \le t \le 1$, almost-surely

$$\mathrm{S}^J_{s,t}(B) \coloneqq X^J_{s,t} = \int_{s<u_1<\dots<u_k<t} \circ\, \mathrm{d}B^{(j_1)}_{u_1} \dots \circ\, \mathrm{d}B^{(j_k)}_{u_k}, \tag{2.6}$$

where $\circ\, \mathrm{d}B$ denotes integration in the Stratonovich sense. A consequence of this is that the solution of a differential equation driven by Brownian motion when thought of as a rough path has the same law as the solution of the same equation when interpreted in the classical Stratonovich sense; hence why we use the Stratonovich interpretation of (1.1) (recall, however, that Stratonovich- and Itô SDEs can easily be transformed into one antoher). The advantage of considering Brownian motion as a rough path is that, once a continuous version of the Lévy area has been fixed, we can solve all differential equations driven by Brownian motion in a pathwise sense on the same set of probability one, and without requiring predictability or other such inherently probabilistic conditions common in stochastic integration.

2.2 Cubature Formulae

Recall that our goal is to estimate statistics of the solution to SDEs of the form

$$\mathrm{d}y_t = \mu(y_t)\,\mathrm{d}t + \sum_{i=1}^{d} \sigma_i(y_t)\,\mathrm{d}W^i_t. \tag{2.7}$$

That is, we want to estimate $\mathbb{E}\left[f(y_T)\right]$ for some test function $f\colon \mathbb{R}^e \to \mathbb{R}$. If a closed-form analytical solution is not available, then this is traditionally achieved using Monte Carlo estimation, which takes the average over some number M of independently sampled solutions of (2.7):

$$\mathbb{E}[f(y_T)] \approx \frac{1}{M} \sum_{i=1}^{M} f(y_T(\omega_i)), \tag{2.8}$$

where $\omega_i\colon [0,T] \to \mathbb{R}^d$ are independent Brownian sample paths, and $y_T(\omega_i) \in \mathbb{R}^e$ denotes the solution at time T of the system (2.7) driven by the path ω_i. By the Central Limit Theorem, the error of the Monte Carlo estimate decays proportionally to $1/\sqrt{M}$, where M is the number of samples. To get an idea how one might try to do better, an analogy to classical numerical integration is useful.

Numerical Integration

Consider the task of integrating some function $f\colon [0,1]^d \to \mathbb{R}$ over the unit cube. This integral can also be thought of as the expectation of $f(U)$ for a uniform random variable $U \sim \mathcal{U}[0,1]$, and therefore Monte Carlo estimation is once again applicable:

$$\int_0^1 f(u)\,\mathrm{d}u = \mathbb{E}\left[f(U)\right] \approx \frac{1}{M} \sum_{i=1}^{M} f(u_i), \tag{2.9}$$

where $u_i \sim \mathcal{U}[0,1]$ are independent. The error is again of order $O(1/\sqrt{M})$. There are many approximation schemes that are known to be better (under suitable assumptions on f); amongst the most well-known and popular are *Quasi–Monte Carlo* methods and *Cubature Formulae*.

Quasi–Monte Carlo (QMC) methods replace the i.i.d. sequence of samples $u_i \sim \mathcal{U}[0,1]$ by a deterministic *low-discrepancy* sequence, which is designed to fill the unit cube more uniformly than randomly sampled points. QMC can achieve an error of order $O(M^{-1}(\log M)^d)$ if f is sufficiently smooth (has finite Hardy-Krause variation), see e.g. [44, 33]. Specific examples of low-discrepancy sequences include Sobol sequences [38], Latin hypercube sequences [32], and Halton sequences [18].

Cubature formulae are finite collections of points with weights, $(\lambda_i, u_i)_{i=1}^M$, where $u_i \in [0,1]^d$ and $\lambda_i \geq 0$ with $\sum_{i=1}^M \lambda_i$, that are chosen in such a way that the cubature approximation

$$\int_{[0,1]^d} f(u)\,\mathrm{d}u \approx \sum_{i=1}^{M} \lambda_i f(u_i) \tag{2.10}$$

is exact for polynomials f up to some maximum degree, which is called the *degree* D of the cubature formula. Equivalently, a collection (λ_i, u_i) is a cubature of degree D if, thought of as an (atomic) probability measure, it matches the moments of a uniform random variable up to order D. If f is at least $k+1$ times continuously differentiable, then cubature formulae can achieve an error of $O(M^{-k/d})$, see e.g. [21]. Simple, well-known examples of cubature formulae in one dimension include the trapezoidal rule (degree 1) or Simpson's rule (degree 3).

Quasi Monte Carlo on Wiener Space

To apply QMC to SDE simulation, one generates a low-discrepancy sequence, maps each point through the inverse normal CDF to obtain a deterministic sequence of Gaussian samples, and uses them, properly rescaled, as increments of discretized Brownian sample paths (as used, for example, in the Euler-Maruyama scheme). In our experiments, we include Sobol sequences [38] and Latin hypercube sequences [32] as QMC baselines; we also tested Halton sequences [18], but they were consistently inferior (typically close to ordinary Monte Carlo), so we omit them from the plots. See Section 4.

Cubature Formulae on Wiener Space

Extending the idea of a cubature formula to the setting of SDE simulation is less straightforward. It would involve finding a deterministic collection of weighted paths, $(\lambda_i, \omega_i)_{i=1}^M$, where $\lambda_i \geq 0$ with $\sum_{i=1}^M \lambda_i = 1$ and $\omega_i \colon [0,T] \to \mathbb{R}^d$, that matches, in some appropriate sense, the "moments" of Brownian motion up to some degree D. Then, the Monte Carlo

estimate (2.8) could be replaced by

$$\mathbb{E}[f(y_T)] \approx \sum_{i=1}^{M} \lambda_i f(y_T(\omega_i)). \tag{2.11}$$

This idea was first formalised in the seminal 2004 paper by Lyons and Victoir [28]. They suggest that, within the context of cubatures, the appropriate generalised notion of a random path's "moment" is its expected signature.

Definition 2.6 (Def. 2.2 in [28])**.** A *cubature on Wiener space* of degree $D \in \mathbb{N}$ at time $T > 0$ is a collection of weighted paths $(\lambda_i, \omega_i)_{i=1}^M$, where $\omega_i \colon [0,T] \to \mathbb{R}^d$ are continuous paths of finite variation, and $\lambda_i \geq 0$ with $\sum_{i=1}^M \lambda_i = 1$, such that

$$\sum_{i=1}^{M} \lambda_i \, \mathrm{S}^I_{0,T}(\omega_i) = \mathbb{E}\left[\mathrm{S}^I_{0,T}(B)\right],$$

for all multi-indices $I = (i_1, \ldots, i_k)$ with *degree* $\|I\| \coloneqq k + |\{j \colon i_j = 0\}| \leq D$.

That is, a weighted collection of paths $(\lambda_i, \omega_i)_{i=1}^M$ is a cubature on Wiener space of degree D if its expected signature matches that of Brownian motion up to degree D. Here the definition of the degree $\|I\|$ of a multi-index $I = (i_1, \ldots, i_k)$ is rooted in the fact that $\mathrm{S}^I_{0,t}(B)$ is equal in law to $(\sqrt{t})^{\|I\|} \, \mathrm{S}^I_{0,1}(B)$. *Wiener space* refers to the space $C([0,T], \mathbb{R}^d)$ of continuous paths $[0,T] \to \mathbb{R}^d$ equipped with the Wiener measure—the law of Brownian motion.

Theorem 2.7 (Thm. 2.4 in [28])**.** *For every $d, D \in \mathbb{N}$ and $T > 0$, there exists a cubature on Wiener space of degree D at time $T > 0$, with cardinarlity at most*

$$M(d, D) \coloneqq |\{I \in \{0, \ldots, d\}^* \colon \|I\| \leq D\}|, \tag{2.12}$$

where $\{0, \ldots, d\}^* = \bigcup_{k=0}^{\infty} \{0, \ldots, d\}^k$ is the set of multi-indices over the alphabet $\{0, \ldots, d\}$.

See Table 3.1 for a few example values of $M(d, D)$. Lyons and Victoir [28] proceed to show that the error of the cubature approximation (2.11) is of order $T^{(D+1)/2}$ for small time intervals T, under suitable smoothness assumptions.

Theorem 2.8 (Prop. 3.2 in [28])**.** *Let $(\lambda_i, \omega_i)_{i=1}^M$ be a cubature on Wiener space of degree D at time T. If $f \colon \mathbb{R}^e \to \mathbb{R}$ and the vector fields of the SDE are sufficiently smooth, then*

$$\left| \mathbb{E}\left[f(y_T)\right] - \sum_{i=1}^{M} \lambda_i f(y_T(\omega_i)) \right| \leq C T^{(D+1)/2},$$

where $C > 0$ depends on f and the vector fields, but not on T, the cubature, or the initial condition $y_0 \in \mathbb{R}^e$ of the SDE.

See Proposition 3.2 in [28] for precise conditions on f and the vector fields. Based on this local error estimate, they suggest partitioning $[0,T]$, if T is not small (e.g. $T=1$), into some number N of disjoint subintervals of length T/N, and applying a given cubature on each of them. The local error estimate then translates to a global error estimate of order $O(N^{-(D-1)/2})$, see the following theorem.

Theorem 2.9 (Thm. 3.3 in [28]). *Let $(\lambda_i, \omega_i)_{i=1}^M$ be a cubature on Wiener space of degree D at time T. Let $N \in \mathbb{N}$ and, for $j \in [N]$ and $i \in [M]$ define $\omega_{i,j}\colon [(j-1)T/N, jT/N] \to \mathbb{R}^{d+1}$ by $\omega_{i,j}^{(0)}(t) = t$ and $\omega_{i,j}^{(l)}(t) = N^{-1/2}\omega_i^{(l)}(N(t-(j-1)T/N))$ for $l = 1,\dots,d$. Then, if f and the vector fields of the SDE are sufficiently smooth,*

$$\left| \mathbb{E}\left[f(y_T)\right] - \sum_{i_1=1}^{M} \dots \sum_{i_N=1}^{M} \lambda_{i_1} \cdot \ldots \cdot \lambda_{i_N} f\big(y_T(\omega_{i_1,1} \odot \ldots \odot \omega_{i_N,N})\big) \right| \leq C N^{-(D-1)/2}, \quad (2.13)$$

where $C > 0$ is a constant that depends on f, the vector fields, and T, and where $\odot$ denotes path concatenation.

See Theorem 3.3 in [28] for precise conditions on f and the vector fields.

3 The ARCANE Algorithm

Using the terminology introduced in the background section, we start by reiterating the main idea of the ARCANE algorithm, then explain all of its steps in detail.

Motivation The main downside of the "product method" based on Theorem 2.9 is that the number of terms in the estimation in (2.13) grows exponentially in the number of sub-divisions N, while the error only decays polynomially in N. Specifically, as a function of the number $m = M^N$ of summands, the error is of order $O((\log m)^{-(D-1)/2})$, asymptotically much larger than the Monte Carlo error $O(m^{-1/2})$.

The main idea behind our approach is to use (2.11) directly, with no subdivision, and instead scale the *degree* D of the cubature. The local error estimate Theorem 2.8 does not apply in this setting (in fact, if $T > 1$ then the error bound diverges as $D \to \infty$), so we prove a new theoretical result that ensures that the error of the approximation (2.11) goes to zero as $D \to \infty$ regardless of $T > 0$ (see Theorems 3.4 and 3.6 below). Our proof method does not give quantitative error bounds, but we find empirically that the approximation error of the cubature formulae constructed by our method is orders of magnitudes smaller than the Monte Carlo error obtained from an equivalent number of sample paths, consistently across several different example SDEs (see Section 4).

The main reason this idea has not been tried before is the difficulty of finding cubature formulae of high degree. Lyons and Victoir proved their existence (recall Theorem 2.7), but they've proven incredibly difficult to construct. They themselves provided

constructions for $D = 5$, and over the past twenty years this has only been improved to $D = 7$. Our primary contribution is the ARCANE[3] algorithm, which automatically and efficiently constructs cubature formulae of arbitrary degree. Even on low-end compute (a single A100 GPU), we were able to reconstruct the state of the art $D = 7$ in seconds, and up to $D = 19$ within hours.

3.1 Explanation of the Algorithm

As an input to the ARCANE algorithm, we fix a Brownian dimension $d \geq 1$ and a cubature degree $D \geq 1$. The ARCANE algorithm consists of the following core steps, each of which we explain in more detail below.

- **Step 1:** We start by constructing an *approximate* cubature from a binary orthogonal array (or, equivalently, from an error-correcting code).
- **Step 2:** Then we compress that approximate cubature to size $M(d, D)$ (recall (2.12)) through an algorithm called "Recombination" [25, 40], implemented efficiently on GPU.
- **Step 3:** Next, we use a linear program to modify the weights of the compressed, approximate cubature in such a way that it becomes an exact cubature.
- **Step 4:** In an optional final step, we take the product of the cubature with itself and recombine it again, yielding a cubature that matches the Brownian signature moments to the given depth D not only on $[0, 1]$, but also on $[0, \frac{1}{2}]$ and $[\frac{1}{2}, 1]$. This step can be repeated for any number of times, yielding cubatures of arbitrary *dyadic depth*.

We now describe each step in more detail.

Step 1: Construction of Approximate Cubature

Outcome: Approximate cubature of size $(Nd)^{\lfloor D/2 \rfloor}$ (with uniform weights)

A simple way of generating approximate cubatures are product constructions: in the simplest case, we could divide $[0, 1]$ into N equally sized subintervals, and consider the (uniformly weighted) set of 2^{Nd} paths that take a step of size $\pm 1/\sqrt{N}$ in each of the d dimensions on each of the N intervals, with linear interpolations in between. This is nothing but the simple symmetric random walk on $[0, 1]$ with N steps and Brownian scaling, which is well-known to converge to Brownian motion as $N \to \infty$. Its expected signature can be shown to be within a relative error of $O(1/N)$ of that of Brownian

[3] **A**lgorithm for **R**ecombination of **C**ubature formulae from Orthogonal **A**rrays that match **N**ested **E**xpected signatures

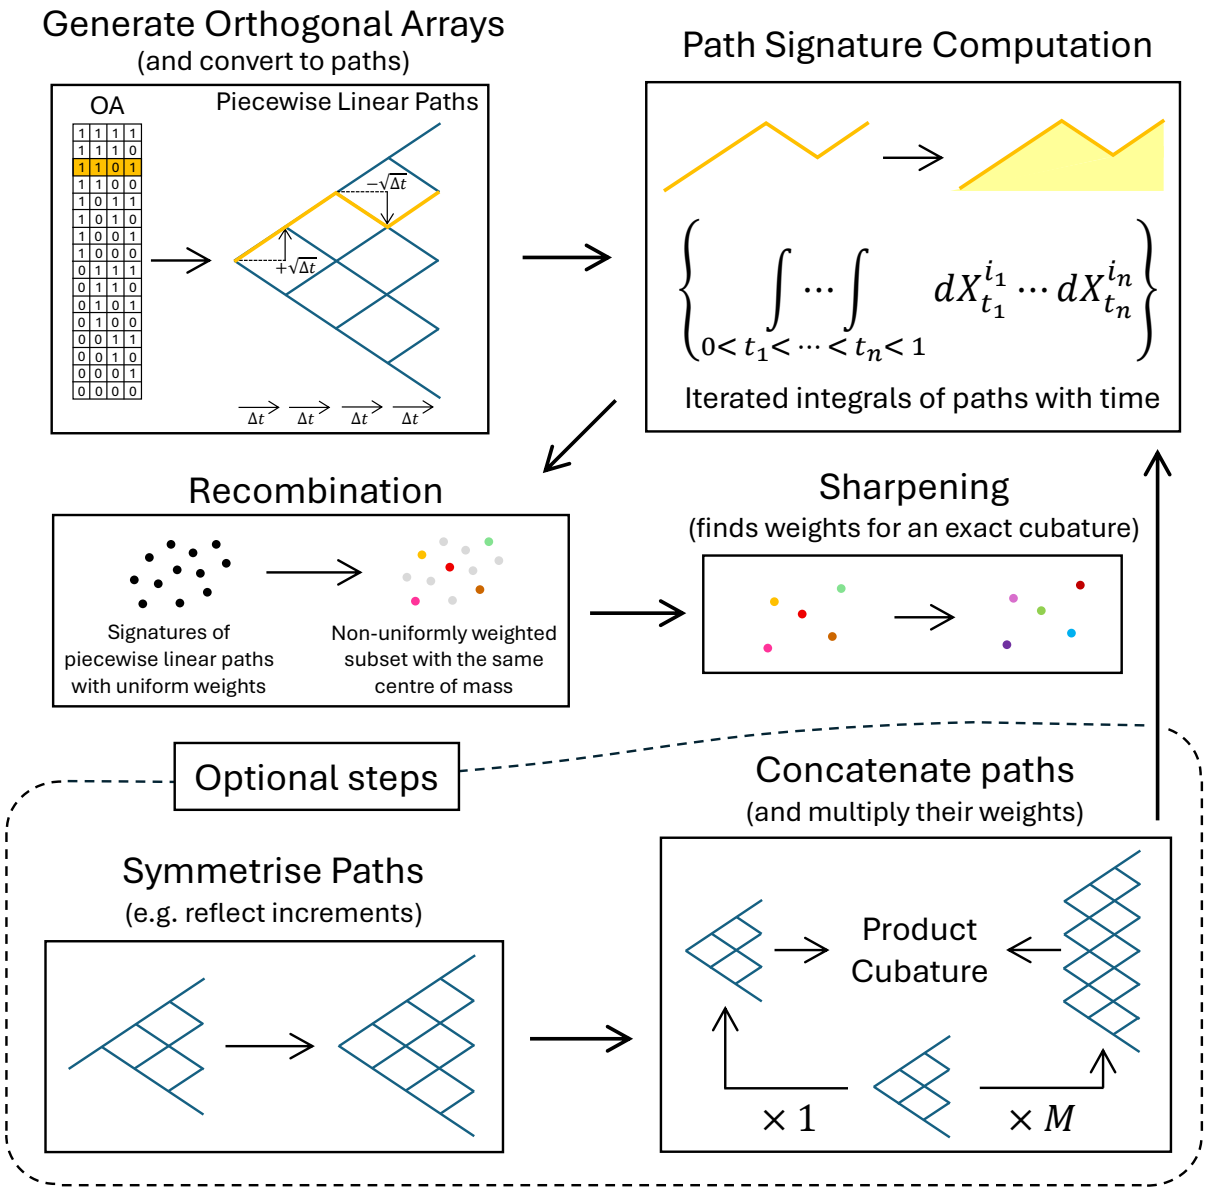

Figure 3.1: High-level overview of the ARCANE algorithm.

motion up to any fixed depth (see Lemma 3.9 below), making it a valid yet impractically large approximate cubature.

For convenience of notation, given a word $w \in \{0,1\}^{Nd}$, we denote by $\omega_w \colon [0,1] \to \mathbb{R}^{d+1}$ the realisation of the symmetric random walk whose increments are chosen according to the entries of w. That is, $\omega_w^{(0)}(t) = t$, $\omega_w^{(l)}(0) = 0$, and

$$\omega_w^{(l)}\left(\frac{j}{N}\right) = \omega_w^{(l)}\left(\frac{j-1}{N}\right) + \frac{w_{jd+l}}{\sqrt{N}}, \qquad j \in [N], l \in [d], \tag{3.1}$$

with linear interpolation in between.

Lemma 3.1. *The weighted collection of paths* $(\lambda_w, \omega_w)_{w\in\{0,1\}^{Nd}}$ *where* $\lambda_w = 2^{-Nd}$ *for all*

w satisfies

$$\left| \mathbb{E}\left[\mathrm{S}^I_{0,T}(B)\right] - \sum_{w \in \{0,1\}^{Nd}} \lambda_w \, \mathrm{S}^I_{0,T}(\omega_w) \right| \leq \frac{C}{N}$$

where $C = C(d, I) > 0$.

To explain how we improve on this basic idea, we need the following definition.

Definition 3.2. Let $s, R, C \in \mathbb{N}$ with $2 \leq s \leq C$. A matrix $A \in \{0,1\}^{R \times C}$ is called a *binary orthogonal array of strength* s if the $R \times s$ submatrix obtained by selecting any s columns has the property that all 2^s possible rows appear the same number of times.

The probabilistic interpretation of a binary orthogonal array with strength s and C columns is that of a coupling of C coin flips that are s-wise independent.

Our key observation is this: if $A \subset \{0,1\}^{Nd}$ is the set of rows of a binary orthogonal array of strength at least D, or equivalently A is an error-correcting code with "dual distance" at least $D+1$ (see Definition I.A.15 and Theorem I.A.16, or [20, Thm. 4.9]), then $(\lambda_w, \omega_w)_{w \in C}$ with uniform weights $\lambda_w = 1/|C|$ has the *same* expected signature as the full symmetric random walk.

Lemma 3.3. *If* $A \subset \{0,1\}^{Nd}$ *are the rows of an orthogonal array with strength at least* D*, then* $(\omega_w)_{w \in A}$ *and* $(\omega_w)_{w \in \{0,1\}^{Nd}}$*, each with uniform weights, have the same expected signature up to degree* D*. In particular,*

$$\left| \mathbb{E}\left[\mathrm{S}^I_{0,T}(B)\right] - \sum_{w \in A} \lambda_w \, \mathrm{S}^I_{0,T}(\omega_w) \right| \leq \frac{C}{N},$$

for a constant $C = C(d, I) > 0$.

If R is a power of two, then an orthogonal array of strength D with R columns can be constructed with $R^{\lfloor D/2 \rfloor}$ rows, based on a code of Bose and Ray [4], see Theorem I.A.4 (3a) for details. If Nd is a power of two, this leads to an approximate cubature of size $|A| = (Nd)^{\lfloor D/2 \rfloor}$, otherwise at most $(2Nd)^{\lfloor D/2 \rfloor}$.

This reduction in size compared with the naive full product construction with 2^{Nd} paths is transformative. For example, for dimension $d = 1$ and degree $D = 9$, if $N = 64$ (which we found to be minimal for the success of Step 3 below), the naive construction would consist of $2^N \approx 1.8 \times 10^{19}$ paths, which our construction reduces to $N^4 \approx 1.6 \times 10^7$ paths. More examples are in Table 3.1.

We note that orthogonal arrays have previously been used to construct cubature formulae for measures on $\mathbb{R}^n$ [43].

Table 3.1: For different degrees D, and $d = 1$ fixed, this table lists the minimal N (restricted to powers of 2) for which the sharpening in Step 3 succeeds, the output size of Step 1 (naive and actual), and an upper bound on the output size of Steps 2 and 3 (with and without symmetrisation). The value for N if $D = 19$ is a guess.

Dimension	Degree	Minimum N	Step 1 Output Size		Steps 2 & 3 Output Size	
			naive	actual	default	symmetrised
d	D	N	2^{Nd}	$(Nd)^{\lfloor D/2 \rfloor}$	$M(d,D)$	$2^d M_{\text{even}}(d,D)$
1	5	32	4.3×10^9	1.0×10^3	20	16
1	7	32	4.3×10^9	3.3×10^4	54	42
1	9	64	1.8×10^{19}	1.7×10^7	143	110
1	11	64	1.8×10^{19}	1.1×10^9	376	288
1	13	64	1.8×10^{19}	6.9×10^{10}	986	754
1	15	64	1.8×10^{19}	4.4×10^{12}	2583	1974
1	17	64	1.8×10^{19}	2.8×10^{14}	6764	5168
1	19	128	3.4×10^{38}	9.2×10^{18}	17710	13530

Step 2: Recombination

Outcome: Approximate cubature of size at most $M(d,D)$ (with non-uniform weights)

Recombination [25, 40] is an algorithm that takes a set of points $(x_i)_{i=1}^M$ in $\mathbb{R}^m$ for some $m \in \mathbb{N}$, with probability weights $(\lambda_i)_{i=1}^M$, and returns a subset $I \subset \{1, \dots, m\}$ of size $|I| \leq m+1$ together with new probability weights $(\tilde{\lambda}_i)_{i \in I}$ such that the weighted mean is preserved:

$$\sum_{i=1}^{M} \lambda_i x_i = \sum_{i \in I} \tilde{\lambda}_i x_i.$$

Note that Carathéodory's convex hull theorem guarantees that such a subset always exists, and that the size $m+1$ of the subset is optimal in general (in the sense that there exist examples where no smaller subsets with the desired property exist; for example, consider the $d+1$ corners of the standard simplex in $\mathbb{R}^d$ with e.g. uniform weights). The run-time of recombination is $O(Mm + \log(M/m)m^3)$ [39, page 61].

Since signatures can be viewed simply as vectors of real numbers, we can apply recombination to the approximate cubature obtained in Step 1 to obtain a new approximate cubature with the same expected signature, but size at most $M(d,D)$ (recall (2.12)), and in general with non-uniform weights. For example, if $d = 1$ and $D = 9$, then the size of the compressed cubature is at most 143, down from 1.6×10^7. Examples for more values of D and N are in Table 3.1.

We use the `baryx` library [13], which contains an efficient JAX [5] implementation of recombination that was implemented primarily by the first author of this paper.

Symmetrisation An optional speed-up in this step can be achieved by exploiting symmetries of Brownian motion: specifically, if $B = (B^{(1)}, \dots, B^{(d)})$ is a d-dimensional Brownian motion, then, for any $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_d) \in \{+1, -1\}^d$, B is equal in distribution to $\boldsymbol{\varepsilon} B \coloneqq (\varepsilon_1 B^{(1)}, \dots, \varepsilon_d B^{(d)})$. This implies that

$$\mathbb{E}\left[\mathrm{S}^I_{0,T}(B)\right] = \mathbb{E}\left[\mathrm{S}^I_{0,T}(\boldsymbol{\varepsilon} B)\right] = \left[\prod_{\substack{i \in I \\ i \neq 0}} \varepsilon_i\right] \mathbb{E}\left[\mathrm{S}^I_{0,T}(B)\right], \quad \forall \boldsymbol{\varepsilon} \in \{+1, -1\}^d, I \in \{0, \dots, d\}^*,$$

and therefore $\mathbb{E}\left[\mathrm{S}^I_{0,T}(B)\right] = 0$ if I is *odd* in the sense that at least some non-zero index $i \in \{1, \dots, d\}$ appears an odd number of times in I. Similarly, any (weighted) set of paths with the same symmetry automatically matches the expected Brownian signature in all odd terms. We can hence apply recombination to the output of Step 1, but only maintaining their expected signature in *even* (i.e. not odd) terms of degree up to D, which is significantly cheaper since recombination scales cubically in the number of terms whose weighted mean is preserved. The result $(\lambda_i, \omega_i)_{i=1}^M$ has size at most

$$M_{\text{even}}(d, D) = |\{I \in \{0, \dots, d\}^* \colon \|I\| \leq D \text{ and } I \text{ even}\}|$$

(cmp. (2.12)) and can then be "symmetrised" to obtain $(2^{-d}\lambda_i, \boldsymbol{\varepsilon}\omega_i)_{i \in [M], \boldsymbol{\varepsilon} \in \{+1,-1\}^d}$. The symmetrisation does not change the expected signature at even terms, and ensures that odd terms are zero, matching the expected Brownian signature exactly. The final size of this cubature is at most $2^d M_{\text{even}}(d, D)$, see Table 3.1 for a few example values.

Step 3: Sharpening

Outcome: Exact cubature of size at most $M(d, D)$.

In this step, we apply a linear programming solver (we use Gurobi [16] and HiGHS [22]) to the approximate cubature obtained in Step 2 to check if the weights can be adjusted in such a way that the resulting reweighted cubature is exact. We have no theoretical guarantees for this step beyond the fact that, for any given tolerance of the solver, it will succeed for sufficiently large N simply because the output of Step 2 will eventually be within that tolerance without any sharpening, due to Lemma 3.3. We record the minimum value of N (amongst powers of two) required for the sharpening to succeed as a function of the degree D in Table 3.1 for $d = 1$. We also found that it can be helpful to stop the recombination in Step 2 early, e.g. stop when the approximate cubature has been compressed to three or four times its final size $M(d, D)$, in order to give the linear program more degrees of freedom.

Step 4: Dyadic Construction

Outcome: Exact *dyadic* cubature.

Given an exact cubature formula $(\lambda_i, \omega_i)_{i=1}^M$, we can construct the product of the cubature with itself: first we rescale each of the paths $\omega_i \colon [0,1] \to \mathbb{R}^d$ to paths $\tilde{\omega}_{i,1} \colon [0,1/2] \to \mathbb{R}^{d+1}$ and $\tilde{\omega}_{i,2} \colon [1/2,1] \to \mathbb{R}^{d+1}$ as follows:

$$\tilde{\omega}_{i,1}^{(l)}(t) = \frac{\omega_i^{(l)}(2t)}{\sqrt{2}}, \quad t \in [0,1/2], \qquad \tilde{\omega}_{i,2}^{(l)}(t) = \frac{\omega_i^{(l)}(2t-1)}{\sqrt{2}}, \quad t \in [1/2,1],$$

for $l \in [d]$, and $\tilde{\omega}_{i,k}^{(0)}(t) = t$ $(k = 1, 2)$. Then the product cubature is

$$(\lambda_{i,j}^{\text{prod}}, \omega_{i,j}^{\text{prod}})_{i,j=1}^M, \quad \text{where} \quad \lambda_{i,j}^{\text{prod}} = \lambda_i \lambda_j, \qquad \omega_{i,j}^{\text{prod}} = \tilde{\omega}_{i,1} \odot \tilde{\omega}_{j,2}.$$

This cubature has the property that it matches the expected Brownian signature on $[0,1/2]$ and $[1/2,1]$ in addition to $[0,1]$. We then apply recombination to compress this cubature; the same symmetrisation trick used in Step 2 can be used here as well.

This process can be repeated iteratively, to obtain cubatures that match the expected Brownian signature on all intervals of the form $[k2^{-n_d}, (k+1)2^{-n_d}]$ for some $n_d \geq 0$, which we call the *dyadic depth* of the cubature (an ordinary cubature as produced by Step 3 has dyadic depth zero). We have a theoretical result (Theorem 3.4 below) that scaling the dyadic depth of a cubature alongside its degree is guaranteed to yield an asymptotically consistent SDE estimator.

3.2 Theoretical Results

3.2.1 Asymptotic Consistency

In this section, we want to establish conditions under which our cubatures are guaranteed to exhibit asymptotic consistency in the following sense:

> *If y^n is the solution to the SDE (1.1) driven by an ARCANE cubature $(\lambda_i^n, \omega_i^n)$, and y is the solution to the same SDE driven by Brownian motion, then $y^n \to y$ in distribution in the space of continuous paths.* (C)

We conjecture that (C) holds as long as $D_n \to \infty$, with no assumptions on the dyadic depth. We cannot show this (for the reader familiar with the language, we can show uniqueness of the subsequential limit, but not tightness), but we can prove that (C) holds if, in addition to $D_n \to \infty$, either the dyadic depth goes to infinity along with the degree (Theorem 3.4), or we skip the recombination step (Theorem 3.6). Additionally, we can show that $D_n \to \infty$ is *necessary* in a certain sense, regardless of the dyadic depth, see Lemma 3.8 below.

Theorem 3.4. *Suppose* $(\lambda_i^n, \omega_i^n)$ *is an ARCANE cubature with degree* $D_n \to \infty$*, and dyadic depth* $m_n \in \mathbb{N}$ *such that*

$$m_n \geq (\frac{3}{4} + \varepsilon) \log_2 N_n - C$$

for some $\varepsilon, C > 0$*. Then (*C*) holds.*

Remark 3.5. For the assertion of Theorem 3.4 to hold, it is in fact sufficient if the cubature matches the expected Brownian signature up to degree $D_n \to \infty$ only on $[0, 1]$, and to degree $D = 7$ on all of the finer dyadic intervals.

Theorem 3.6. *Let* $(\lambda_i^n, \omega_i^n)$ *be the approximate cubature resulting from Step 1 of the ARCANE algorithm, prior to recombination, sharpening, and any dyadic constructions, with* $N_n \to \infty$ *and degree* $D_n \to \infty$*. Then (*C*) holds.*

For the reader familiar with rough path theory, we remark that the essence in the proofs is showing that the cubatures converge in distribution to canonical Brownian motion *in the space of rough paths*, in which case (C) follows from the universal limit theorem [29, Section 3.5] (convergence in the space of ordinary paths would not imply (C)).

Remark 3.7. Work closely related to these types of theorems has been done on *Donsker-type theorems on rough path space*. Specifically, Breuillard, Friz, and Huesmann [6] give conditions under which a random walk with i.i.d. increments converges to Brownian motion in rough path space, and Bayer and Friz [3] prove convergence of product cubature formulae to Brownian motion in rough path space. The latter would imply (a stronger version of) our Theorem 3.4 *if* we didn't do the (essential) recombination step after the product construction. Nevertheless, the types of arguments used in the proof are very similar.

Some of the facts established in the proofs may be of independent interest, such as Lemma I.B.5, which gives conditions under which pointwise convergence of expected signatures of random rough paths implies convergence in distribution and vice versa.

Lemma 3.8. *For every* $D \in \mathbb{N}$ *there exists an example of a stochastic process that matches the expected Brownian signature up to degree* D *on* all *subintervals of* $[0, 1]$*, but is not a Brownian motion.*

The proofs of Theorems 3.4 and 3.6 are relatively involved and in Appendix I.B. The proof of Lemma 3.8 is simple and we present it here.

Proof of Lemma 3.8. Let B be a Brownian motion, and Z a real-valued random variable independent of B with $\mathbb{E}[Z^k] = 1$ for all even $k \leq D$. Then the process $X(t) = ZB(t)$ matches the expected Brownian signature up to degree D on any subinterval $[s, t] \subset [0, 1]$. It is easy to construct such a Z that also has $\mathbb{P}(Z = 0) > 0$, in which case X is not a Brownian motion because $\mathbb{P}(X \equiv 0) > 0$. □

3.2.2 Proofs of Lemmas 3.1 and 3.3

In this section, we validate Step 1 by proving Lemmas 3.1 and 3.3. We first prove a generalisation of Lemma 3.1, and then show that Lemma 3.1 implies Lemma 3.3.

To motivate how Lemma 3.1 could be generalised, notice that we can think of the set of 2^{Nd} realisations of the simple symmetric random walk on $[0,1]$ with N steps as an N-fold product cubature in the sense of Theorem 2.9. The "base cubature" represents a single step of the random walk: it consists of the 2^d paths of the form $\omega(t) = tw$ for $w \in \{1,-1\}^d$. It is easy to see that this is a cubature of degree 3.

More generally, suppose that $(\lambda_i, \omega_i)_{i=1}^M$ is a cubature of degree D, and consider its N-fold product cubature: for $j \in [N]$ and $i \in [M]$ define

$$\omega_{i,j} \colon \left[\frac{(j-1)T}{N}, \frac{jT}{N}\right] \to \mathbb{R}^{d+1}; \qquad \omega_{i,j}^{(0)}(t) = t, \quad \omega_{i,j}^{(l)}(t) = \frac{1}{\sqrt{N}} \omega_i^{(l)}\left(N\left(t - \frac{(j-1)T}{N}\right)\right),$$

for $l = 1, \dots, d$. This is the path ω_i scaled onto the interval $[(j-1)T/N, jT/N]$. Then the *N-fold product cubature of* $(\lambda_i, \omega_i)_{i=1}^M$ is

$$\left(\lambda_{i_1} \cdot \ldots \cdot \lambda_{i_N}, \omega_{i_1,1} \odot \ldots \odot \omega_{i_N,N}\right)_{i_1,\ldots,i_N=1}^M, \tag{3.2}$$

where $\odot$ denotes concatenation of paths.

Lemma 3.9. *If $(\lambda_i, \omega_i)_{i=1}^M$ is a cubature of degree D, then its N-fold product cubature is also of degree D, and for any multi-index $I \in \{0, \dots, d\}^*$ with degree $\|I\| > D$ there exists a constant $C = C(T, I) > 0$ such that*

$$\left| \mathbb{E}\left[\mathrm{S}_{0,T}^I(B)\right] - \sum_{i_1=1}^M \cdots \sum_{i_N=1}^M \lambda_{i_1} \cdot \ldots \cdot \lambda_{i_N} \, \mathrm{S}_{0,T}^I\left(\omega_{i_1,1} \otimes \ldots \otimes \omega_{i_N,N}\right) \right| \leq C N^{-(D-1)/2}.$$

Proof. Denote by $S_i \coloneqq \mathrm{S}_{0,T}(\omega_i)$ the signature of path ω_i. Then the signature of $\omega_{i,j}$, for any $j \in [N]$, is equal to the Brownian scaling $\tilde{S}_i$ of S_i onto an interval of length T/N, which is given by

$$\tilde{S}_i^I = N^{-\|I\|/2} S_i^I, \quad I \in \{0, \dots, d\}^*.$$

Denote by $\tilde{S}^{(1)}, \dots, \tilde{S}^{(N)}$ independent random variables distributed according to the scaled base cubature, that is

$$\mathbb{P}\left(\tilde{S}^{(1)} = \tilde{S}_{i_1}, \dots, \tilde{S}^{(N)} = \tilde{S}_{i_N}\right) = \prod_{j=1}^N \lambda_{i_j}, \quad i_1, \dots, i_N \in [M].$$

Then the expected signature of the product cubature (3.2) is

$$\mathbb{E}\left[\tilde{S}^{(1)} \otimes \ldots \otimes \tilde{S}^{(N)}\right] = \mathbb{E}\left[\tilde{S}^{(1)}\right]^{\otimes N},$$

due to the independence. Now, since $(\lambda_i, \omega_i)_{i=1}^M$ is a degree D cubature, $\mathbb{E}[\tilde{S}^{(1)}]$ coincides up to degree D with the expected Brownian signature on an interval of length T/N, which

is $\exp\left(\frac{T}{N}e_0 + \frac{T}{2N}\sum_{i=1}^{w} e_i \otimes e_i\right)$ [29]. Due to the Brownian scaling, all terms of degree higher than D are of order $N^{-(D+1)/2}$, which we collectively denote $O_{(>D)}(N^{-(D+1)/2})$. Thus,

$$\begin{aligned}\mathbb{E}\left[\tilde{S}^{(1)}\right]^{\otimes N} &= \left(\exp\left(\frac{T}{N}e_0 + \frac{T}{2N}\sum_{i=1}^{N} e_i \otimes e_i\right) + O_{(>D)}(N^{-(D+1)/2})\right)^{\otimes N} \\ &= \exp\left(Te_0 + \frac{T}{2}\sum_{i=1}^{N} e_i \otimes e_i\right) + O_{(>D)}\left(N^{-(D-1)/2}\right).\end{aligned}$$

That is, the expected signature is exact up to order D and has an absolute error of order $N^{-(D-1)/2}$ in higher degree terms, which was our claim. □

Lemma 3.1 follows by applying Lemma 3.9 to the "one-step" cubature $(\lambda_x, \omega_x)_{x\in\{+1,-1\}^d}$ defined by $\lambda_x = 2^{-d}$ and $\omega_x(t) = tx$.

Proof of Lemma 3.3. Let $A \subset \{0,1\}^{Nd}$ be the rows of an orthogonal array of strength at least D (recall Definition 3.2, and $(\lambda_w, \omega_w)_{w\in A}$ the associated cubature (recall (3.1)). Thinking of this cubature as a random path $[0,1] \to \mathbb{R}^{d+1}$ with N linearly interpolated increments of size $\pm 1/\sqrt{N}$, the definition of an orthogonal array is equivalent to the statement that any D of these increments are independent, and every individual increment is (marginally) equal to $+1/\sqrt{N}$ and $-1/\sqrt{N}$ with probability 1/2 each. Finally, the expected signature up to degree D is a function of expectations of products of up to D of the increments, and hence equal to that of the simple symmetric random walk. Then the claim follows from Lemma 3.1 □

4 Empirical Results

In order to be able to evaluate the accuracy of both the cubature and the Monte Carlo SDE solvers, we need a ground truth to compare against. Hence, we consider SDEs where the mean and variance of the solution,

$$m(t) = \mathbb{E}\left[Y_t\right], \qquad v(t) = \mathbb{V}(Y_t)$$

are available in closed form (as a function of the initial condition y_0 and the parameters of the SDE). Then, for a given set of Monte Carlo paths, or a given cubature $(\lambda_i, \omega_i)_{i=1}^{M}$, we calculate the empirical mean and variance,

$$\hat{m}(t) = \sum_{i=1}^{M} \lambda_i y_t(\omega_i), \qquad \hat{v}(t) = \sum_{i=1}^{M} \lambda_i y_t(\omega_i)^2 - \hat{m}(t)^2.$$

and combine them in the *Mean-Variance error*

$$\text{error} = \sqrt{(\hat{m}(T) - m(T))^2 + (\sqrt{\hat{v}(T)} - \sqrt{v(T)})^2}, \tag{4.1}$$

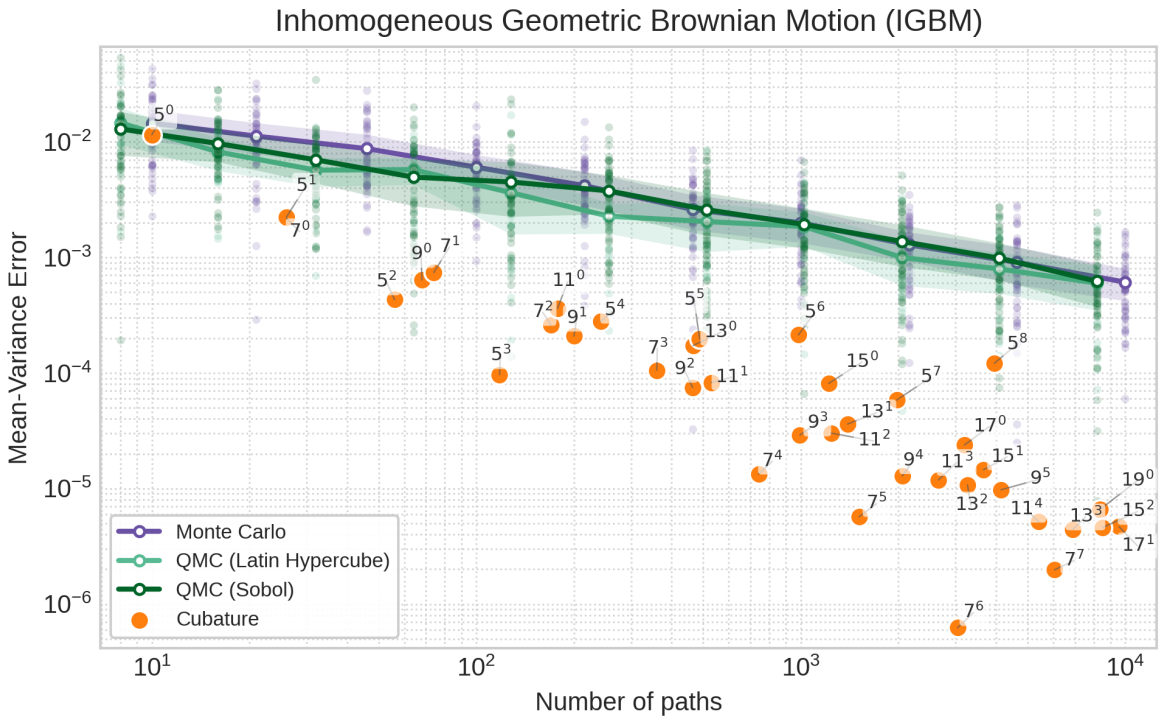

Figure 4.1: The (Mean-Variance) SDE simulation error of our cubature paths (orange) against an equivalent number of Monte Carlo paths (purple), evaluated on the IGBM model, see (4.2).

for a fixed time $T = 1.0$. This is the Wasserstein W_2 distance between two one-dimensional Gaussians $\mathcal{N}(\hat{m}(T), \hat{v}(T))$ and $\mathcal{N}(m(T), v(T))$. We now describe the three SDE's that our experiments are based on.

Inhomogeneous Geometric Brownian Motion (IGBM):

$$\mathrm{d}Y_t = a(b - Y_t)\,\mathrm{d}t + \sigma Y_t\,\mathrm{d}W_t, \tag{4.2}$$

where $a > 0$ is the mean reversion rate, $b \in \mathbb{R}$ is the long-term mean, and $\sigma > 0$ is the volatility. See eqs.(3) and (4) in [41] for the mean and variance. We use parameters $y_0 = 0.06$, $a = 0.1$, $b = 0.04$, $\sigma = 0.6$, as suggested in [8]. Results are in Figure 4.1. For this model, we also have an illustration of the actual solution paths for one of our smaller cubatures, compared with an equivalent number of Monte Carlo paths, see Figure 4.5.

Neutral Wright-Fisher Diffusion:

$$\mathrm{d}Y_t = \sqrt{\gamma Y_t(1 - Y_t)}\,\mathrm{d}W_t, \tag{4.3}$$

with variance parameter $\gamma > 0$. The mean is $m(t) = y_0$, and the variance is $v(t) = y_0(1 - y_0)(1 - \mathrm{e}^{-\gamma t})$. We use parameters $y_0 = 0.5$ and $\gamma = 0.1$. Results are in Figure 4.2.

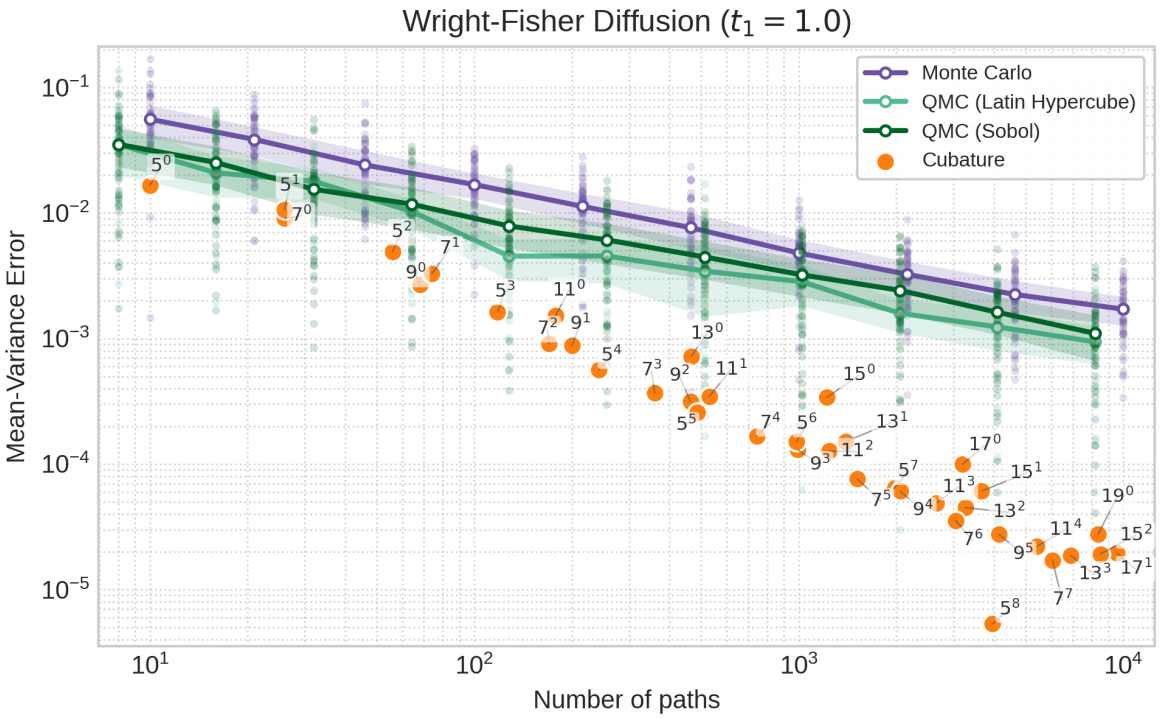

Figure 4.2: The (Mean-Variance) SDE simulation error of our cubature paths (orange) against an equivalent number of Monte Carlo paths (purple) and Quasi-Monte Carlo paths (green), evaluated on the Wright-Fisher model, see (4.3).

Cox-Ingersoll-Ross Model (CIR):

$$\mathrm{d}Y_t = a(b - Y_t)\,\mathrm{d}t + \sigma\sqrt{Y_t}\,\mathrm{d}W_t, \tag{4.4}$$

where $a > 0$ is the mean reversion rate, $b \in \mathbb{R}$ is the long-term mean, and $\sigma > 0$ is the volatility [12]. We use the parameters $y_0 = 2.0$, $a = 1.0$, $b = 3.0$, $\sigma = 0.5$, as suggested in [36]. The same source contains formulas for mean and variance.

For the CIR model, there is also a closed form analytical formula for the expected bond price,

$$\mathbb{E}\left[\exp\left(-\int_0^T Y_s\,\mathrm{d}s\right)\middle|\, Y_0\right],$$

see e.g. [34]. This is a particularly useful metric because, unlike the mean and variance of the endpoint, it depends on the entire solution path $(Y_s)_{0\le s\le T}$ rather than just its endpoint Y_T. Results are in Figures 4.3 and 4.4.

5 Limitations

Just as standard cubature formulae can lose accuracy when the test functions are non-smooth, the accuracy of our cubature formulae will depend heavily on the SDE. For example, if the SDE is highly oscillatory or simulated over too long a time horizon, then

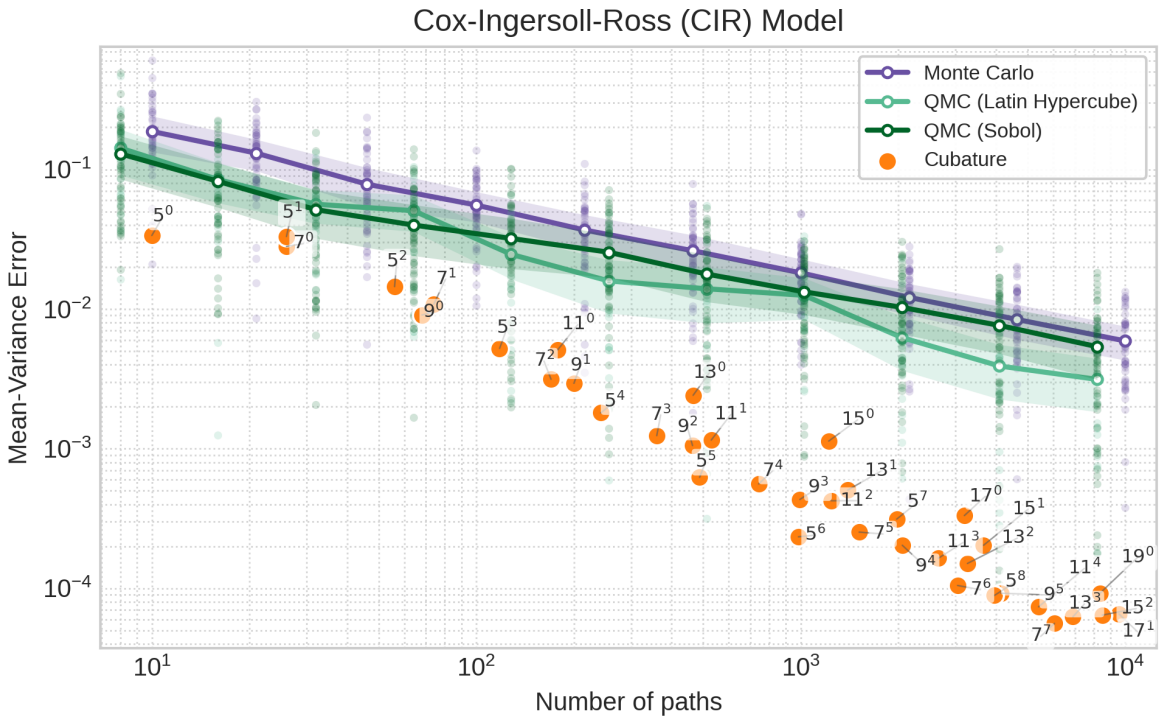

Figure 4.3: The (relative bond price) SDE simulation error of our cubature paths (orange) against an equivalent number of Monte Carlo paths (purple) and Quasi-Monte Carlo paths (green), evaluated on the CIR model, see (4.4).

it is possible for our formulae to become inaccurate. This phenomena can easily be seen for the Wright-Fisher diffusion simply by increasing γ, or equivalently by increasing the time horizon. Figure 5.1 shows the loss in accuracy for our ARCANE cubatures on the Wright-Fisher example when the time horizon is increased from $T = 1.0$ to $T = 10.0$.

Acknowledgments

We would like to thank Terry Lyons for many helpful discussions and advice. PK gratefully acknowledges support by the EPSRC grant EP/W523781/1.

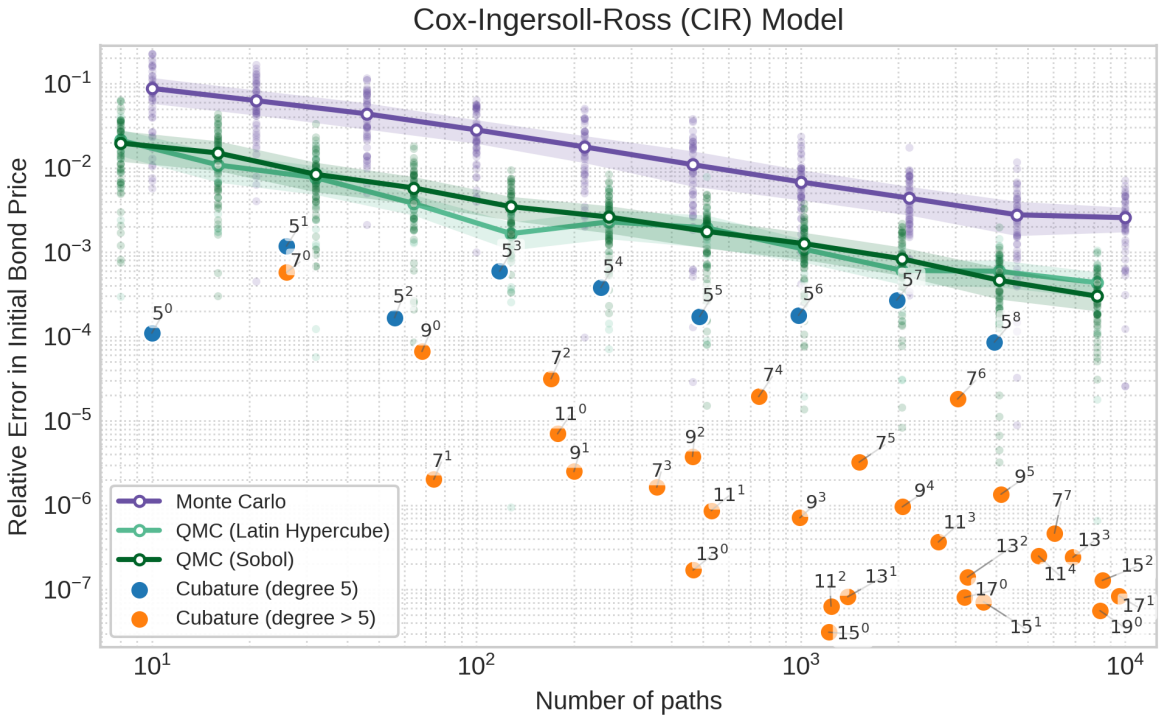

Figure 4.4: The SDE simulation error of our cubature paths (orange) against an equivalent number of Monte Carlo paths (purple) and Quasi-Monte Carlo paths (green), evaluated on the CIR model, see (4.4). Interestingly, the cubatures with degree 5 (highlighted in blue) seem to perform categorically worse.

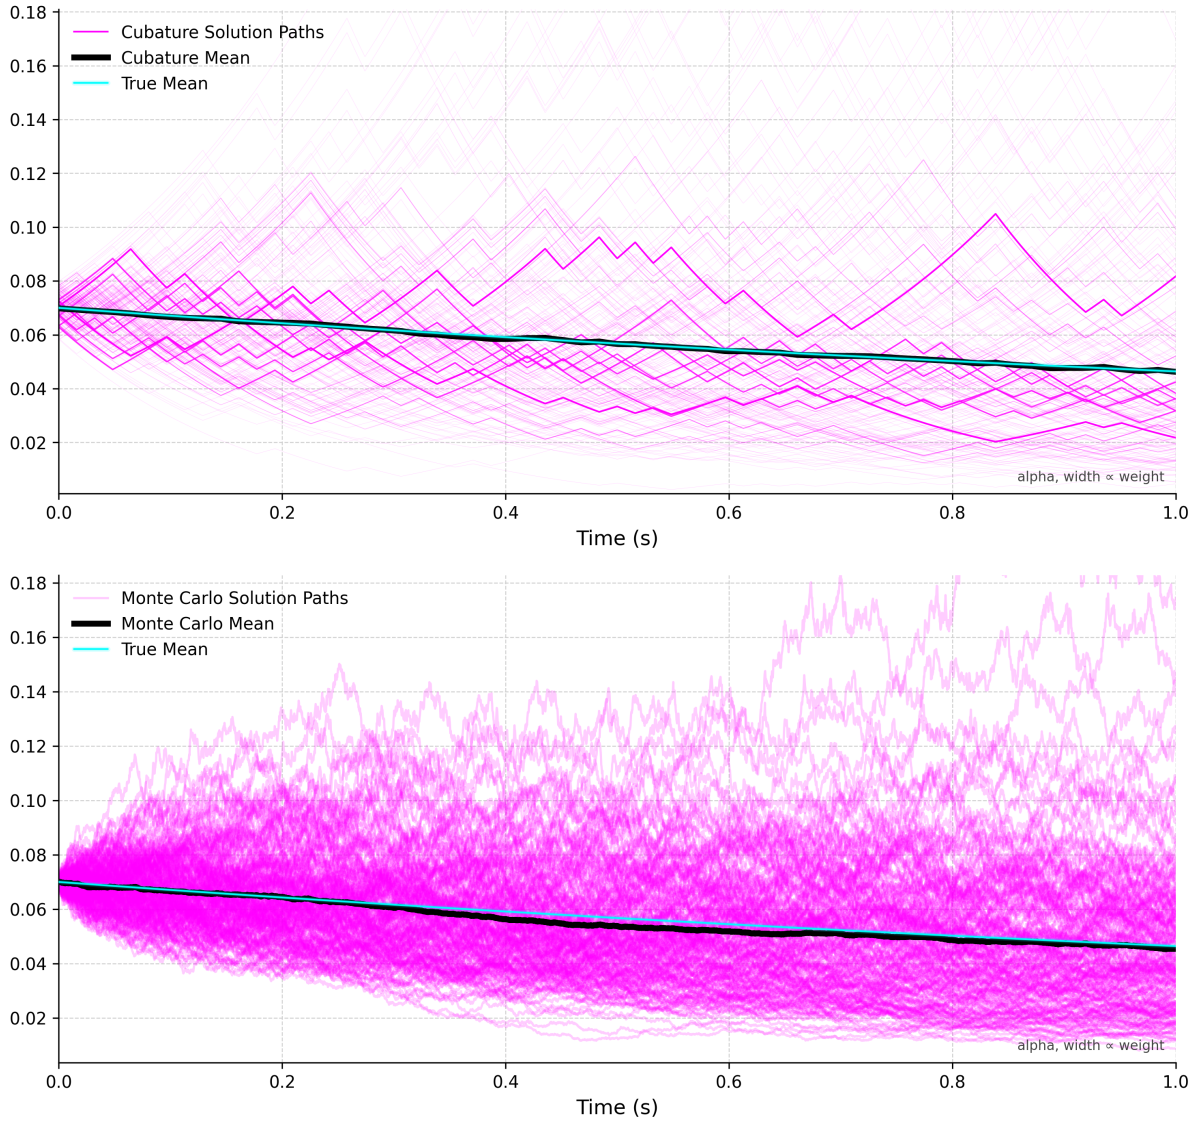

Figure 4.5: Illustration of the solution paths of the IGBM model, see (4.2), when driven by our 9^1 cubature (degree 9, dyadic depth 1), which consists of 200 paths (top), and the solution paths to the same SDE when driven by 200 randomly sampled (Monte Carlo) paths (bottom). Despite consisting of the same number of paths, and hence having the exact same cost when used to solve the SDE, the mean of the Cubature solution paths stays much closer to the true mean than that of the Monte Carlo solution paths. Note that this is one of our smallest cubature formulae, which was picked to make the errors large enough to be visually discernable.

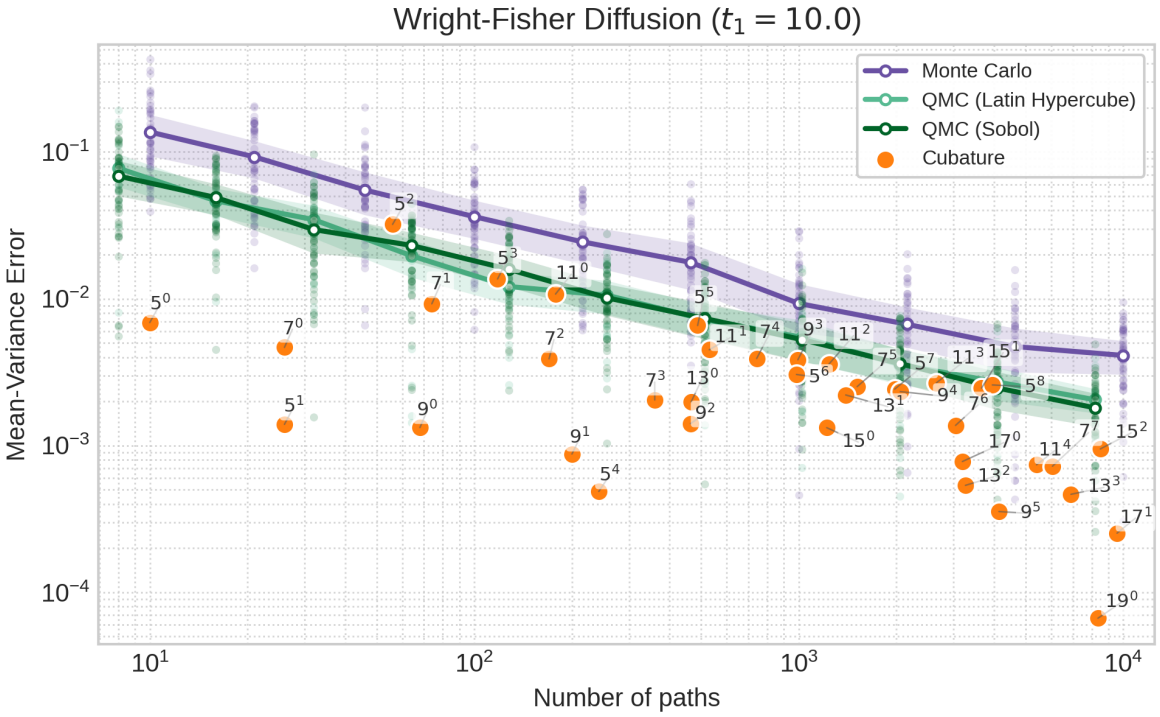

Figure 5.1: The (Mean-Variance) SDE simulation error of our cubature paths (orange) against an equivalent number of Monte Carlo paths (purple) and Quasi-Monte Carlo paths (green), evaluated on the Wright-Fisher model with time horizon $T = 1.0$.

Bibliography

[1] J. Pablo Arenas-López and Mohamed Badaoui. "A Fokker–Planck equation based approach for modelling wind speed and its power output". In: *Energy Conversion and Management* 222 (2020).

[2] Jordi Badosa et al. "Day-Ahead Probabilistic Forecast of Solar Irradiance: A Stochastic Differential Equation Approach". In: *Renewable Energy: Forecasting and Risk Management*. 2018, pp. 73–93.

[3] Christian Bayer and Peter K Friz. "Cubature on Wiener space: pathwise convergence". In: *Applied Mathematics & Optimization* 67.2 (2013), pp. 261–278.

[4] R.C. Bose and D.K. Ray-Chaudhuri. "On a class of error correcting binary group codes". In: *Information and Control* 3.1 (1960), pp. 68–79. ISSN: 0019-9958. DOI: https://doi.org/10.1016/S0019-9958(60)90287-4.

[5] James Bradbury et al. *JAX: composable transformations of Python+NumPy programs*. Version 0.3.13. 2018.

[6] Emmanuel Breuillard, Peter Friz, and Martin Huesmann. "From random walks to rough paths". In: *Proceedings of the American Mathematical Society* 137.10 (2009), pp. 3487–3496.

[7] Yongli Cai et al. "A stochastic SIRS epidemic model with infectious force under intervention strategies". In: *Journal of Differential Equations* 259.12 (2015), pp. 7463–7502.

[8] Luca Capriotti, Yupeng Jiang, and Gaukhar Shaimerdenova. "Approximation methods for inhomogeneous geometric Brownian motion". In: *International Journal of Theoretical and Applied Finance* 22.02 (2019), p. 1850055. DOI: 10.1142/S0219024918500553. eprint: https://doi.org/10.1142/S0219024918500553.

[9] Ilya Chevyrev and Harald Oberhauser. "Signature Moments to Characterize Laws of Stochastic Processes". In: *Journal of Machine Learning Research* 23.176 (2022), pp. 1–42.

[10] Loren Cobb. *Mathematical Frontiers Of The Social And Policy Sciences*. Routledge, New York, 1981.

[11] John C. Cox, Jonathan E. Ingersoll, and Stephen A. Ross. "A Theory of the Term Structure of Interest Rates". In: *Econometrica* 53.2 (1985), pp. 385–407.

[12] John C. Cox, Jonathan E. Ingersoll, and Stephen A. Ross. "A Theory of the Term Structure of Interest Rates". In: *Econometrica* 53.2 (1985), pp. 385–407. ISSN: 00129682, 14680262.

[13] Thomas Coxon. *baryx*. Version 0.1.0. 2025.

[14] Peter Czuppon and Arne Traulsen. "Understanding evolutionary and ecological dynamics using a continuum limit". In: *Ecology and Evolution* 11.11 (2021), pp. 5857–5873.

[15] Emilio Ferrucci et al. *High-degree cubature on Wiener space through unshuffle expansions*. 2024. arXiv: 2411.13707 [math.PR].

[16] Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*. 2025.

[17] Lajos Gergely Gyurkó and Terry Lyons. "Efficient and Practical Implementations of Cubature on Wiener Space". In: *Stochastic Analysis 2010*. Springer Berlin Heidelberg, 2011, pp. 73–111.

[18] J. H. Halton. "On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals". In: *Numerische Mathematik* 2 (1960), pp. 84–90.

[19] Satoshi Hayakawa and Ken'ichiro Tanaka. "Monte Carlo construction of cubature on Wiener space". In: *Japan Journal of Industrial and Applied Mathematics* 39.2 (2022), pp. 543–571.

[20] A Samad Hedayat, Neil James Alexander Sloane, and John Stufken. *Orthogonal arrays: theory and applications*. Springer Science & Business Media, 2012.

[21] Aicke Hinrichs et al. "Product rules are optimal for numerical integration in classical smoothness spaces". In: *Journal of Complexity* 38 (2017), pp. 39–49.

[22] Qi Huangfu and Julian A. J. Hall. "Parallelizing the dual revised simplex method". In: *Mathematical Programming Computation* 10.1 (2018), pp. 119–142. DOI: 10.1007/s12532-017-0130-5.

[23] Emil B. Iversen et al. "Probabilistic forecasts of solar irradiance using stochastic differential equations". In: *Environmetrics* 25.3 (2014), pp. 152–164.

[24] Haiquan Li et al. "Identification of a Stochastic Dynamic Model for Aircraft Flight Attitude Based on Measured Data". In: *International Journal of Aerospace Engineering* 2023.1 (2023).

[25] Christian Litterer and Terry Lyons. "High order recombination and an application to cubature on Wiener space". In: *Annals of Applied Probability* 22.4 (2012), pp. 1301–1327.

[26] Ying Liu et al. "A stochastic approximation method for probability prediction of docking success for aerial refueling". In: *Applied Soft Computing* 103 (2021).

[27] Terry Lyons and Zhongmin Qian. *System control and rough paths*. Oxford University Press, 2002.

[28] Terry Lyons and Nicolas Victoir. "Cubature on Wiener Space". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 460.2041 (2004), pp. 169–198.

[29] Terry J Lyons, Michael Caruana, and Thierry Lévy. *Differential equations driven by rough paths*. Springer, 2007.

[30] Yoshihiro Maki and Hideo Hirose. "Infectious Disease Spread Analysis Using Stochastic Differential Equations for SIR Model". In: *2013 4th International Conference on Intelligent Systems, Modelling and Simulation*. 2013, pp. 152–156.

[31] Anatoliy Malyarenko and Hossein Nohrouzian. "Testing Cubature Formulae on Wiener Space Versus Explicit Pricing Formulae". In: *Stochastic Processes, Statistical Methods, and Engineering Mathematics*. Cham: Springer International Publishing, 2022, pp. 223–248.

[32] M. D. McKay, R. J. Beckman, and W. J. Conover. "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code". In: *Technometrics* 21.2 (1979), pp. 239–245.

[33] H. Niederreiter. *Random Number Generation and Quasi-Monte Carlo Methods*. Vol. 63. CBMS–NSF Regional Conference Series in Applied Mathematics. SIAM, 1992.

[34] QuantPie. *CIR Price PDE*. `https://www.quantpie.co.uk/srm/cir_price_pde.php`. QuantPie Website. 2025.

[35] Hamood Ur Rehman et al. "Analysis of Brownian motion in stochastic Schrödinger wave equation using Sardar sub-equation method". In: *Optik* 289 (2023).

[36] Dialid Santiago. *5. Cox–Ingersoll–Ross Process. Understanding Quantitative Finance*. Quant Girl blog. 2024.

[37] Yuji Shinozaki. "Construction of a Third-Order K-Scheme and Its Application to Financial Models". In: *SIAM Journal on Financial Mathematics* 8.1 (2017), pp. 901–932.

[38] I. M. Sobol'. "On the distribution of points in a cube and the approximate evaluation of integrals". In: *USSR Computational Mathematics and Mathematical Physics* 7.4 (1967), pp. 86–112.

[39] Maria Tchernychova. "Caratheodory cubature measures". PhD thesis. University of Oxford, 2016.

[40] Maria Tchernychova. "Carathéodory cubature measures". PhD thesis. University of Oxford, 2015.

[41] Irene Tubikanec et al. *Qualitative properties of numerical methods for the inhomogeneous geometric Brownian motion*. 2021. arXiv: `2003.10193 [math.NA]`.

[42] Marc Vaisband et al. "Loss formulations for assumption-free neural inference of SDE coefficient functions". In: *Nature Partner Journal: Systems Biology and Applications* 11.22 (2025).

[43] Nicolas Victoir. "Asymmetric Cubature Formulae with Few Points in High Dimension for Symmetric Measures". In: *SIAM Journal on Numerical Analysis* 42.1 (2004), pp. 209–227. DOI: `10.1137/S0036142902407952`. eprint: `https://doi.org/10.1137/S0036142902407952`.

[44] X. Wang. "Low discrepancy sequences in high dimensions: How well do they integrate?" In: *Journal of Computational and Applied Mathematics* 215.2 (2008), pp. 301–314.

[45] Yasen Wang et al. "Data-Driven Discovery of Stochastic Differential Equations". In: *Engineering* 17 (2022), pp. 244–252.

Appendix

I.A Orthogonal Arrays

The first step of the ARCANE algorithm crucially relies on access to large, high-strength orthogonal arrays. For this reason, as a part of this project a considerable amount of effort went into finding (and efficiently implementing) orthogonal array constructions. All of the constructions presented in this section are implemented efficiently and publicly available in the following GitHub repository:

github.com/peter-koepernik/orray

Most of the constructions presented here are based on existing work and are attributed as such, but some are novel to the best of our knowledge.

Definition I.A.1. Let $q, s, R, C \in \mathbb{N}$ with $q \geq 2$ and $2 \leq s \leq C$. A matrix $A \in \{0, \dots, q-1\}^{R \times C}$ is called a *q-ary orthogonal array of strength s*, an $OA(C, R, q, s)$ for short, if the $R \times s$ submatrix obtained by selecting any s columns has the property that all q^s possible rows appear the same number of times.

The probabilistic interpretation of a q-ary orthogonal array with strength s and C columns is that of a coupling of C uniform random variables on $\{0, \dots, q-1\}$ that are s-wise independent.

In the context of the ARCANE algorithm, we are interested in coupling a set of $C = Nd$ coin flips in an s-wise independent way, where s is the desired degree of the approximate cubature. This is realised by an $OA(R, C, 2, s)$, and then each of the R rows corresponds to one path in the resulting cubature. Therefore, the parameter regime for orthogonal arrays that we are most interested in is that where $q = 2$ and s are fixed, $C = Nd$ is variable and large, and R is as small as possible as a function of C (given q and s).

Remark I.A.2. Recall from Section 3.2.2 that Step 1 of the ARCANE algorithm can be thought of as taking the N-fold product of a base cubature that represents a single step of a random walk. This "base cubature" could in principle be replaced by a different cubature (such as a d-dimensional bespoke Gaussian cubature), in which case an s-wise independent coupling would require a non-binary orthogonal array of strength s. This is

the main reason why we consider general q-ary orthogonal arrays in this section, rather than only binary ones.

A well-known *lower* bound on the number of rows of an orthogonal array with given s, q, C is given by *Rao's bound*.

Theorem I.A.3 (Rao [16])**.** *Any $OA(R, C, q, s)$ satisfies*

$$R \geq \sum_{i=0}^{\lfloor s/2 \rfloor} \binom{C}{i} (q-1)^i = \Theta\left(C^{\lfloor s/2 \rfloor}\right). \tag{I.A.1}$$

The Rao bound can be asymptotically achieved when $q = 2$—the case of most interest to us—as well as if $s = 2$, but for most parameter combinations with $s, q \geq 3$ the smallest possible asymptotic of $R = R(C)$ is unknown, and in particular whether or not they achieve the asymptotic growth $\Theta(C^{\lfloor s/2 \rfloor})$ suggested by (I.A.1).

In the following theorem, we summarise the best construction methods we are aware of (in the asymptotic regime we are interested in, i.e. s, q fixed and C large), together with their asymptotic relationship between R and C. Most of these are taken or adapted from existing literature, and we go into detail on both the constructions as well as the correct attributions in the subsequent sections.

Theorem I.A.4. *The following orthogonal arrays can be constructed.*

(1) ***Strength 2****: For a prime power q and $n \in \mathbb{N}$,*

$$OA\left(q^n, \frac{q^n - 1}{q - 1}, q, 2\right), \quad \textit{having} \quad R = 1 + (q-1)C,$$

which attains (I.A.1) *exactly.*

(2) ***Strength 3:*** *For a prime power q and $n \in \mathbb{N}$,*

$$OA\left(q^{3n+1}, q^{2n}, q, 3\right), \quad \textit{having} \quad R = qC^{3/2},$$

which does not attain (I.A.1) *asymptotically.*

(3) ***Binary Arrays ($q = 2$):*** *All of the following arrays attain* (I.A.1) *asymptotically.*

(i) (Bose-Ray) For $n \in \mathbb{N}$ and even $s \in \mathbb{N}$,

$$OA\left(2^{n\lfloor s/2 \rfloor}, 2^n - 1, 2, s\right), \quad \textit{having} \quad R = (C+1)^{\lfloor s/2 \rfloor}.$$

(ii) (Bose-Ray) For $n \in \mathbb{N}$ and odd $s \in \mathbb{N}$,

$$OA\left(2^{1+n\lfloor s/2 \rfloor}, 2^n, 2, s\right), \quad \textit{having} \quad R = 2C^{\lfloor s/2 \rfloor}.$$

(iii) (Kerdock) For $s = 5$ and an even integer $n \geq 4$,

$$OA\left(2^{2n}, 2^n, 2, 5\right), \quad \textit{having} \quad R = C^{\lfloor s/2 \rfloor}.$$

(iv) (Delsarte-Goethals) For $s = 7$ and an even integer $n \geq 4$,

$$OA\left(2^{3n-1}, 2^n, 2, 7\right), \quad \textit{having} \quad R = \frac{1}{2}C^{\lfloor s/2 \rfloor}.$$

(4) ***Ternary Arrays*** *($q = 3$):*

(i) For $s = 3$, and $n \in \mathbb{N}$,

$$\begin{aligned} &OA(3^{4n+1},\ 20^n, 3, 3), \quad \textit{having} \quad R = 3C^{4\log 3/\log 20} \approx 3C^{1.467}, \\ &OA(3^{5n+1},\ 45^n, 3, 3), \quad \textit{having} \quad R = 3C^{5\log 3/\log 45} \approx 3C^{1.443}, \\ &OA(3^{6n+1},112^n, 3, 3), \quad \textit{having} \quad R = 3C^{6\log 3/\log 112} \approx 3C^{1.397}, \end{aligned}$$

all three of which do not attain (I.A.1) *asymptotically.*

(ii) For $s = 4$, and $n \in \mathbb{N}$,

$$OA\left(3^{2n+1}, 3^n, 3, 4\right), \quad \textit{having} \quad R = 3C^2,$$

which attain (I.A.1) *asymptotically.*

(5) ***General:*** *For a prime number $q \in \mathbb{N}$, $s \in \mathbb{N}$, and $n \in \mathbb{N}$,*

$$OA\left(q^{ns}, q^n, q, s\right), \quad \textit{having} \quad R = C^s,$$

which does not attain (I.A.1) *asymptotically.*

I.A.1 Linear Orthogonal Arrays

With the exception of (3c) and (3d), all of the orthogonal arrays in Theorem I.A.4 belong to the special class of *linear orthogonal arrays*, which are orthogonal arrays $OA(R, C, q, s)$ whose rows form a vector space over a finite field with q elements. Such a field exists if and only if q is a prime power—in which case the field is unique up to isomorphism and denoted $\mathbb{F}_q$—and in that case $R = q^n$ where n is the dimension of the vector space spanned by the rows. This is why all of the constructions above require q to be a prime power.

To construct a linear array, say $OA(q^n, C, q, s)$, it suffices to specify the basis of the vector space spanned by the rows, which is a set of n vectors in $\mathbb{F}_q^C$. If we arrange them as the rows of an $n \times C$ matrix M, then it turns out that they give rise to (i.e. are the basis of the row space of) a linear orthogonal array of strength s if and only if the *columns* of M are s-wise linearly independent.

Lemma I.A.5. *Let q be a prime power, and M an $n \times C$ matrix with elements in $\mathbb{F}_q$. Then the $q^n \times C$ matrix whose rows are all possible $\mathbb{F}_q$–linear combinations of the rows of M form an orthogonal array of strength s if and only if any s columns of M are $\mathbb{F}_q$–linearly independent. In that case, we call M a* generator *of the orthogonal array.*

See e.g. Theorem 10.4 in [17]. They only state the "if" direction, but the "only if" statement follows from the same proof. Using Lemma I.A.5, the task reduces to finding a set of s-wise linearly independent vectors in $\mathbb{F}_q^n$ that is as large as possible. The next section contains constructions based directly on this idea.

I.A.1.1 Direct linear constructions

An instructive example is the case of strength $s = 2$: a set of vectors in $\mathbb{F}_q^n$ for some $n \in \mathbb{N}$ is 2-wise linearly independent if and only if no vector is a multiple of another. The $q^n - 1$ non-zero elements of $\mathbb{F}_q^n$ can be partitioned into disjoint sets of $q - 1$ vectors that are multiples of each other, and if we pick an arbitrary representative out of each of the $C \coloneqq (q^n - 1)/(q - 1)$ classes (for example the set of vectors whose first non-zero entry is equal to 1), then we obtain an $n \times C$ matrix over $\mathbb{F}_q$ whose columns are 2-wise linearly independent. By Lemma I.A.5, it generates an $OA(q^n, (q^n-1)/(q-1), 2, q)$, and this is exactly construction (1) from Theorem I.A.4. A reference for this construction is Corollary 10.5 in [17].

Construction I.A.6 (Theorem I.A.4 (1))**.** For a prime power q and $n \in \mathbb{N}$, there exists a linear

$$OA\left(q^n, \frac{q^n - 1}{q - 1}, 2, q\right).$$

It is generated in the sense of Lemma I.A.5 by an $n \times C$ matrix whose columns are all non-zero elements of $\mathbb{F}_q^n$ whose first non-zero coordinate is equal to 1.

Another relatively simple construction is Theorem I.A.4 (5).

Construction I.A.7 (Theorem I.A.4 (5))**.** For a prime number q, strength $s \in \mathbb{N}$, and $n \in \mathbb{N}$, there exists a linear

$$OA\left(q^{ns}, q^n, q, s\right).$$

It is generated by an $ns \times q^n$ matrix M that can be obtained as follows. First, create an $s \times q^n$ matrix M' over $\mathbb{F}_{q^n}$ whose columns are of the form $(1, x, x^2, \ldots, x^{s-1})$, with x ranging through all q^n elements of $\mathbb{F}_{q^n}$. Then, replace each entry in M' with a column vector of length n over $\mathbb{F}_q$ through an(y) identification of $\mathbb{F}_{q^n}$ with $\mathbb{F}_q^n$ as vector spaces over $\mathbb{F}_q$. The resulting $ns \times q^n$ matrix generates the array.

Proof. It suffices to show that any s columns of M' are $\mathbb{F}_{q^n}$–linearly independent. Indeed, then it follows that they are also $\mathbb{F}_q$–linearly independent, and therefore any s columns of M are $\mathbb{F}_q$–linearly independent.

Let $x_1, \ldots, x_s \in \mathbb{F}_{q^n}$ be distinct. Then we have to show that the $s \times s$ submatrix of M' corresponding to those s columns is non-singular, and indeed its determinant is given through the Vandermonde–formula by

$$\det \begin{pmatrix} 1 & \ldots & 1 \\ x_1 & \ldots & x_s \\ \vdots & \ddots & \vdots \\ x_1^{s-1} & \ldots & x_s^{s-1} \end{pmatrix} = \prod_{i<j} (x_j - x_i) \neq 0,$$

where multiplication etc. are performed in $\mathbb{F}_{q^n}$. □

If $n = 1$, then this is an $OA(q^s, q, s, q)$; by definition, *any* q-ary orthogonal array of strength s must have at least q^s rows, so the number of rows is minimal, and if $q > s$ then the number of columns is larger than that of the trivial array with the same N, q, and s (which has s columns). The largest possible number of columns given s, $q > s$, and $N = q^s$ is an open problem outside of a few known special cases; even if $s = 2$ the answer is unknown for numerous values of q (see e.g. Section 2 in [8], particularly the discussion following Corollary 2.22).

If $n > 1$, then this array stops being optimal in the same sense, and in fact the scaling of its number of rows, $R = C^s$, is worse than that of all the other arrays listed in Theorem I.A.4 in overlapping parameter regimes. But it is the best construction that we are aware of that works for arbitrary strength and $q \geq 3$. For $n = 1$, the construction is well-known, see e.g. [17, Corollary 10.7]. We haven't directly found the generalisation to $n > 1$ in the literature, but due to its simplicity assume that it is likewise well-known.

Construction I.A.8 (Theorem I.A.4 (3a))**.** For $n \in \mathbb{N}$ and even $s = 2u \in \mathbb{N}$, there exists a linear

$$OA(2^{nu}, 2^n - 1, 2, s).$$

It is generated by an $(nu) \times (2^n - 1)$ matrix M over $\mathbb{F}_2$ obtained in the following way. For every non-zero $x \in \mathbb{F}_{2^n}$, create the vector $(x, x^3, \ldots, x^{2u-1}) \in \mathbb{F}_{2^n}^u$. Replace each entry with a binary vector of length n through a (vector space) identification of $\mathbb{F}_{2^n}$ with $\mathbb{F}_2^n$ to yield a vector in $\mathbb{F}_2^{nu}$. The $2^n - 1$ vectors obtained in this way form the columns of M.

Proof. It suffices to prove that the columns of M constructed in this way are s-wise linearly independent, which is in Section 4 of a paper by Bose and Ray-Chaudhuri [2]. □

We haven't found this particular construction in the literature, but since the work of Bose and Ray-Chaudhuri [2] is well-known, we assume that we have just not been able to find it.

Theorem I.A.4 (3b) is obtained from (3a) through the following general construction for binary arrays, see e.g. Theorem 2.24 in [8]. In words, any binary array of even strength $s = 2u$ can be "upgraded" to strength $s = 2u + 1$ at the cost of doubling the number of rows and the additional benefit of adding one column.

Lemma I.A.9. *If A is an $OA(R, C, 2, s)$ for even $s \in \mathbb{N}$, then the $(2R) \times (C+1)$ matrix*

$$\left[\begin{array}{c|c} A & 0 \\ \overline{A} & 1 \end{array}\right]$$

is an $OA(2R, C+1, 2, s+1)$, where $\overline{A}$ denotes the array obtained from A by changing 1's for 0's and vice versa.

I.A.1.2 Constructions from cap sets

Let q be a prime power, and $d \in \mathbb{N}$. A *cap set* in $\mathbb{F}_q^d$ is a set of points such that no three of them are collinear, or, in other words, a 3-wise affinely independent set of points. Despite the simplicity of the following lemma, we have not been able to find it, or any of the OA constructions based on it in the literature, and therefore assume that they are novel.

Lemma I.A.10. *If the columns of an $n \times C$ matrix M over $\mathbb{F}_q$ are s-wise affinely independent, then the $(n+1) \times C$ matrix obtained by adding a row that only consists of 1's generates an $OA(q^{n+1}, C, q, s)$ in the sense of Lemma I.A.5.*

Proof. A set of vectors that is affinely independent becomes linearly independent when an additional coordinate equal to 1 is added to each vector. Then the statement follows from Lemma I.A.5. □

Furthermore, it is easy to see that products of cap sets are cap sets. In particular, if $S \subset \mathbb{F}_q^d$ is a cap set say of size k, then $S^n \subset \mathbb{F}_q^{dn}$ is a cap set of size k^n. Combining with Lemma I.A.10 gives the following.

Corollary I.A.11. *If $S \subset \mathbb{F}_q^d$ is a cap set of size C, then there exists an $OA(q^{nd+1}, C^d, q, 3)$.*

Using Corollary I.A.11, we can turn known cap-set constructions into orthogonal arrays. For instance, Theorem I.A.4 (2) is based on a known construction of a cap set in $\mathbb{F}_q^3$ for prime q.

Construction I.A.12 (Theorem I.A.4 (2))**.** For a prime number q and $n \in \mathbb{N}$, there exists a linear

$$OA(q^{3n+1}, q^{2n}, q, 3).$$

To construct it, let $aX^2 + X + b$ be an irreducible degree two polynomial over $\mathbb{F}_q$, and put

$$S = \{(x, y, ax^2 + xy + by^2) \colon x, y \in \mathbb{F}_q\}.$$

Then S is a cap set of size q^2 in $\mathbb{F}_q^3$, which generates the orthogonal array in the sense of Corollary I.A.11.

Proof. The cap set above can be derived from the cap set (or *ovoid*) in the projective space $PG(3, q)$ constructed in Example 1.4(1) in [14]. □

The arrays in Theorem I.A.4 (4a) are based on known cap sets in $\mathbb{F}_3^d$ for $d = 4, 5, 6$.

Construction I.A.13 (Theorem I.A.4 (4a))**.** The three classes of orthogonal arrays in Theorem I.A.4 (4a) are generated in the sense of Lemma I.A.10 by the following cap sets:

- A cap set of size 20 in $\mathbb{F}_3^4$,
- A cap set of size 45 in $\mathbb{F}_3^5$,
- A cap set of size 112 in $\mathbb{F}_3^6$,

which are explicitly given, constructed, or can be derived from the constructions given in [9, Fig. 1], [10, Section 2], and [15], respectively.

The three cap sets above are all of maximal size [9, 10, 15], see also sequence A090245 in the OEIS.

Sets of points in $\mathbb{F}_q^d$ that are s-wise affinely independent for $s > 3$ are sometimes called *generalised cap sets*, and a lot less well-studied. They also lack the convenient product property of (ordinary) cap sets: in fact, if $S \subset \mathbb{F}_q^d$ contains at least two points and is 4-wise affinely independent, then $S^2 \subset \mathbb{F}_q^{2d}$ is not. Indeed, let $u, v \in S$ be distinct, then (u, u), (u, v), (v, u), and (v, v) are all in S^2 and lie on a plane (i.e. are affinely dependent).

The only result in this area that we are aware of is a construction by Huang [11] of 4-wise affinely independent sets in $\mathbb{F}_3^n$.

Construction I.A.14 (Theorem I.A.4 (4b))**.** For every integer $n \in \mathbb{N}$, there exists a linear

$$OA(3^{2n+1}, 3^n, 3, 4).$$

For $x \in \mathbb{F}_{3^n}$, map the vector $(x, x^2) \in \mathbb{F}_{3^n}^2$ to a ternary vector of length $2n$ using an identification of $\mathbb{F}_{3^n}$ with $\mathbb{F}_3^n$ as vector spaces over $\mathbb{F}_3$. Looping through all 3^n possible values for x, this gives 3^n points in $\mathbb{F}_3^{2n}$. They are 4-wise affinely independent [11] and generate the array in the sense of Corollary I.A.11.

I.A.2 Constructions from non-linear codes

There is a connection between orthogonal arrays and codes, which we outline briefly.

Definition I.A.15. Let $N, k \in \mathbb{N}$, and $q \in \mathbb{N}$ be a prime power.

(i) A *code of size* $N \in \mathbb{N}$ *and length* $k \in \mathbb{N}$ *over* $\mathbb{F}_q$ is a matrix $C \in \mathbb{F}_q^{N \times k}$, whose rows are called *code words*, and we slightly abuse notation by writing $u \in C$ if u is a row in C.

(ii) The *distance* d of C is the minimal Hamming distance between any two distinct codewords, that is,

$$d = \min_{\substack{u,v \in C \\ u \neq v}} d_H(u,v), \quad \text{where} \quad d_H(u,v) = \#\{i \in [k] \colon u_i \neq v_i\}.$$

(iii) A code C is called *linear* if its code words are a vector space over $\mathbb{F}_q$. In that case, the *dual code* of C is its orthogonal complement,

$$C^\perp \coloneqq \{u \in \mathbb{F}_q^k \colon \forall v \in C \colon \langle u, v \rangle = 0\}.$$

The distance $d^\perp$ of $C^\perp$ is called the *dual distance* of C.

It is possible to extend the notion of dual distance to non-linear codes, see e.g. section 4.4 of [8]. This is important for the statement of the following theorem, but it is not necessary in any way to understand the definition for the purpose of reading the remainder of this section.

Theorem I.A.16 (Delsarte 1973 [6])**.** *A matrix with elements in* $\mathbb{F}_q$ *for some prime power* $q \in \mathbb{N}$ *is a code with dual distance at least* $t+1$ *if and only if it is an orthogonal array with strength at least* t*.*

See also Theorem 4.9 in [8] for a more recent exposition. The two remaining constructions, Theorem I.A.4 (3c) and (3d), are obtained from two non-linear codes called the *Kerdock code* [12] and the *Delsarte-Goethals code* [5], binary codes with dual distance 6 and 8, respectively.

Construction I.A.17 (Theorem I.A.4 (3c))**.** For $n \geq 4$ even, there exists a non-linear

$$OA(2^{2n}, 2^n, 2, 5).$$

The rows of the array are the codewords of the Kerdock code, whose construction can be derived from sections 3 and 4 of [7].

Construction I.A.18 (Theorem I.A.4 (3d))**.** For $n \geq 4$ even, there exists a non-linear

$$OA(2^{3n-1}, 2^n, 2, 7).$$

The rows of the array are the codewords of the Delsarte-Goethals code, whose construction can be derived from sections 3, 4, and 6 of [7].

I.B Theoretical Results

I.B.1 (Re-)statement of results

In this appendix, we state and then prove sligthly stronger and more precise versions of Theorems 3.4 and 3.6. We assume that the reader is familiar with some standard notions in probability theory, such as weak convergence of probability measures (i.e. convergence in distribution), which we denote $\mu_n \implies \mu$, and tightness.

Theorem 3.4*. *Suppose $(\lambda_i^n, \omega_i^n)$ is an ARCANE cubature with degree $D_n \to \infty$, and dyadic depth $m_n \in \mathbb{N}$ such that*

$$m_n \geq (\frac{3}{4} + \varepsilon) \log_2 N_n - C$$

for some $\varepsilon, C > 0$. Then there exists a $p \in (2,3)$ such that the cubature converges in distribution to Brownian motion as a p-rough path. In particular, if y^n denotes the solution to an SDE with smooth vector fields driven by the cubature, and y denotes the solution to the same SDE driven by Brownian motion, then $y^n \implies y$ in distribution as rough (and in particular as ordinary) paths.

Theorem 3.6*. *Let $(\lambda_i^n, \omega_i^n)$ be the approximate cubature resulting from the first step of the ARCANE algorithm, prior to recombination, sharpening, and any dyadic constructions, with $N_n \to \infty$ and degree $D_n \to \infty$. Then the assertion of Theorem 3.4 holds for* all $p \in (2,3)$.

Remark I.B.3. (i) The fact that the convergence of the cubature to Brownian motion holds on the level of *rough* paths is a much stronger assertion than just on the level of ordinary paths. In particular, the main assertion that solutions driven by the cubature paths converge to the real solutions only holds thanks to the universal limit theorem for rough paths (Theorem 2.5), and would not follow from convergence on the level of ordinary paths.

(ii) As mentioned in the main paper, for the assertion of Theorem 3.4* to hold, it is in fact sufficient if the cubature matches the expected Brownian signature up to degree $D_n \to \infty$ just on $[0,1]$, and only to degree $D = 7$ on all of the finer dyadic intervals.

I.B.2 Proofs

In the remainder of this appendix, we prove Theorems 3.4* and 3.6*. We start by establishing conditions under which weak convergence and pointwise convergence of expected signatures imply each other for rough paths.

Denote $V = \mathbb{R}^{d+1}$, and recall that $G\Omega_p(V)$ is the space of geometric p-rough paths $x = (x^0, x^1, \dots, x^d)$ in V defined on $[0, T]$ for some fixed and arbitrary $T > 0$ with $x^0(t) = t$. This is a complete and separable metric space [4, Proposition 39]. Recall that the signature is a map $\mathrm{S}\colon G\Omega_p \to T((V))$. Let

$$\mathcal{M} \coloneqq \{\mu \in \mathcal{M}_1(G\Omega_p)\colon \mu\left(\| \mathrm{S}^m \|_{V^{\otimes m}}\right) < \infty \forall m \in \mathbb{N}\}$$

where we used the shorthand $\mu(f)$ for $\int f \,\mathrm{d}\mu$ if f is a function and μ a probability measure such that the integral is well-defined, and $\mathcal{M}_1(X)$ for a measurable space X denotes the set of probability measures on X. In particular, if $\mu \in \mathcal{M}$, then $\mu(\mathrm{S}) \in T((V))$ exists and is the expected signature of the random rough path defined by μ.

Lemma I.B.4. *Let $\mu \in \mathcal{M}$ such that its expected signature has an infinite radius of convergence, that is*

$$\forall \lambda > 0\colon \sum_{m=1}^{\infty} \lambda^m \mu\left(\| \mathrm{S}^m \|\right) < \infty. \tag{I.B.1}$$

Then the expected signature of μ characterises it, that is

$$\forall \nu \in \mathcal{M}\colon \mu(\mathrm{S}) = \nu(\mathrm{S}) \implies \mu = \nu. \tag{I.B.2}$$

Furthermore, (I.B.1) *holds if μ is the law of canonical Brownian motion.*

Proof. Suppose μ satisfies the assumptions of the lemma above, and $\nu \in \mathcal{M}$ with $\mu(\mathrm{S}) = \nu(\mathrm{S})$. By [3, Prop. 6.1], this implies that the pushforward laws $\mathrm{S}\#\mu$ and $\mathrm{S}\#\nu$ coincide. But by [1, Theorem 1.1], a rough path $\boldsymbol{x} \in G\Omega_p$ is determined uniquely by its signature $\mathrm{S}(\boldsymbol{x})$ (we use here that $x^0(t) = t$), and therefore $\mu = \nu$.

Finally, (I.B.1) follows from standard exponential tail bounds on Brownian motion and e.g. [3, Theorem 6.3]. □

We say that a sequence (μ_n) in $\mathcal{M}$ is *uniformly integrable* if

$$\forall m \in \mathbb{N}\colon \sup_{n\in\mathbb{N}} \mu_n\left(\|S^m\| \mathbb{1}_{\{\|S^m\| \geq K\}}\right) \longrightarrow 0, \quad K \to \infty. \tag{I.B.3}$$

Lemma I.B.5. *Let (μ_n) be a uniformly integrable sequence in $\mathcal{M}$, and $\mu \in \mathcal{M}_1(G\Omega_p)$. Then the following hold.*

(i) If $\mu_n \implies \mu$ weakly, then $\mu \in \mathcal{M}$ and $\mu_n(\mathrm{S}) \to \mu(\mathrm{S})$ pointwise.

(ii) If (μ_n) is tight, and $\mu \in \mathcal{M}$ satisfies (I.B.2)*, then $\mu_n(\mathrm{S}) \to \mu(\mathrm{S})$ pointwise implies $\mu_n \implies \mu$ weakly.*

Proof. For $K > 0$ and $x \in \mathbb{R}$, write $\mathrm{clamp}_K(x) \coloneqq \min(K, \max(-K, x))$. To prove (i), assume that $\mu_n \implies \mu$. We first show that $\mu \in \mathcal{M}$. Let J be a multi-index, then

$$\mu(|S^J|) \leq \liminf_{n\to\infty} \mu_n(|S^J|) \leq \liminf_{n\to\infty} \mu_n(\mathrm{clamp}_K(|S^J|)) + \sup_{n\in\mathbb{N}} \mu_n(|S^J| \mathbb{1}_{\{|S^J| \geq K\}})$$

Since $\mathrm{clamp}_K(|S^J|)$ is continuous and bounded, the first term on the right-hand side converges to $\mu(\mathrm{clamp}_K(|S^J|)) < \infty$, and by uniform integrability the second term is finite for some $K > 0$ (in fact goes to zero as $K \to \infty$). We showed that $\mu \in \mathcal{M}$, in particular $\mu(S)$ exists. Next, again for any multi-index J and any $K > 0$,

$$\begin{aligned}\big|\mu_n(S^J) - \mu(S^J)\big| \le \big|\mu_n(\mathrm{clamp}_K(S^J)) - \mu(\mathrm{clamp}_K(S^J))\big| \\ + \mu_n(|S^J|\mathbb{1}_{\{|S^J|\ge K\}}) + \mu(|S^J|\mathbb{1}_{\{|S^J|\ge K\}}).\end{aligned}$$

Since $\mu_n \implies \mu$, the first term on the right-hand side goes to zero as $n \to \infty$ for any fixed $K > 0$, so

$$\limsup_{n\to\infty} \big|\mu_n(S^J) - \mu(S^J)\big| \le \limsup_{n\to\infty} \mu_n(|S^J|\mathbb{1}_{\{|S^J|\ge K\}}) + \mu(|S^J|\mathbb{1}_{\{|S^J|\ge K\}}).$$

By (I.B.3), the right-hand side goes to zero as $K \to \infty$, so the left-hand side is zero. Since J was arbitrary, this implies $\mu_n(\mathrm{S}) \to \mu(\mathrm{S})$ pointwise.

Now assume that $\mu \in \mathcal{M}$ satisfies (I.B.2), $\mu_n(\mathrm{S}) \to \mu(\mathrm{S})$ pointwise, and that (μ_n) is tight. It suffices to prove uniqueness of subsequential limits, so take some convergent subsequence $\mu_{k(n)} \implies \nu$. By (i) which we've already proved, $\nu \in \mathcal{M}$ and $\mu_{k(n)}(\mathrm{S}) \to \nu(\mathrm{S})$, but also $\mu_{k(n)}(\mathrm{S}) \to \mu(\mathrm{S})$, so $\mu(\mathrm{S}) = \nu(\mathrm{S})$, which implies $\nu = \mu$ by (I.B.2). □

The next lemma shows that if we want to apply Lemma I.B.5 (ii) in the case where μ is the distribution of (canonical rough) Brownian motion, then we get uniform integrability of (μ_n) for free.

Lemma I.B.6. *If $\mu_n, \mu \in \mathcal{M}$, $n \in \mathbb{N}$ such that $\mu_n(\mathrm{S}) \to \mu(\mathrm{S})$, and $\mu(\|\,\mathrm{S}^m\,\|^2) < \infty$ for all $m \in \mathbb{N}$, then $\sup_{n\in\mathbb{N}} \mu_n(\|\,\mathrm{S}^m\,\|^2) < \infty$, and in particular (μ_n) is uniformly integrable.*

Proof. For any fixed $m \in \mathbb{N}$ and multi-index J of length m, the shuffle product identity [13, Theorem 2.15] implies that there exists a set $\mathcal{J}$ of multi-indices of length $2m$ (the "shuffles" of J with itself) such that

$$|\,\mathrm{S}^J\,|^2 = (\mathrm{S}^J)^2 = \sum_{I\in\mathcal{J}} \mathrm{S}^I\,.$$

Therefore $\mu_n(\mathrm{S}) \to \mu(\mathrm{S})$ implies $\mu_n(|\,\mathrm{S}^J\,|^2) \to \mu(|\,\mathrm{S}^J\,|^2) < \infty$, in particular $\sup_{n\in\mathbb{N}} \mu_n(|\,\mathrm{S}^J\,|^2) < \infty$. □

To be able to apply Lemma I.B.5, we need to understand how to prove tightness for sequences of probability measures on $G\Omega_p$, which is the content of the next lemma.

For $\alpha \in (0, 1]$ and a function $f\colon [0, T] \to X$ for some normed space X, denote by

$$[f]_\alpha := \sup_{0\le s<t\le T} \frac{\|f(t) - f(s)\|}{|t - s|^\alpha}, \qquad \|f\|_\alpha := [f]_\alpha + \sup_{t\in[0,T]} \|f(t)\|$$

the α-Hölder seminorm and norm of f, respectively. We also write $[f]_{\mathrm{Lip}}$ for $[f]_\alpha$ if $\alpha = 1$, and similarly for $\|f\|_{\mathrm{Lip}}$. Sometimes it will be convenient to work with a version of the α-Hölder (semi-)norm that only evaluates f on subsequent dyadic points:

$$[f]_{\alpha,\mathcal{D}} \coloneqq \sup\left\{\frac{|f(t)-f(s)|}{|t-s|^\alpha} \colon 0 \le s < t \le 1, \exists k, n \in \mathbb{N}_0 \colon s = k2^{-n}, t = s + 2^{-n}\right\},$$

and similarly for $\|f\|_{\alpha,\mathcal{D}}$.

Lemma I.B.7. *Let (μ_n) be a sequence in $\mathcal{M}$, and suppose that, for some $\alpha > 0$,*

$$\limsup_{n\to\infty} \mu_n(x \colon \|x\|_{\alpha,\mathcal{D}} \ge K) \longrightarrow 0, \quad K \to \infty. \tag{I.B.4}$$

and

$$\limsup_{n\to\infty} \mu_n\left(x \colon \sup_{0<s<t<T} \frac{\|\,\mathrm{S}^2_{s,t}(x)\|}{|t-s|^{2\alpha}} \ge K\right) \longrightarrow 0, \quad K \to \infty. \tag{I.B.5}$$

Then (μ_n) is tight in $G\Omega_p$ for every $p > 1/\alpha$.

Proof. First of all, by Lemma I.B.8 below, (I.B.4) implies

$$\limsup_{n\to\infty} \mu_n(x \colon \|x\|_\alpha \ge K) \longrightarrow 0, \quad K \to \infty.$$

Now let

$$A_K \coloneqq \{x \colon \|x\|_\alpha \le K\} \cap \left\{x \colon \sup_{0<s<t<T} \frac{\|\,\mathrm{S}^2_{s,t}(x)\|}{|t-s|^{2\alpha}} \le K\right\}$$

By assumption, $\limsup_{n\to\infty} \mu_n(A_K^c) \to 0$ as $K \to \infty$, so it suffices to show that A_K for any $K > 0$ is relatively compact in $G\Omega_p$.

Let $1/\alpha < q < \min(3, p)$. Then by Proposition 1.15 in [13], if A_K is a bounded subset of $G\Omega_q$ and equicontinuous, then A_K is relatively compact in $G\Omega_p$. Equicontinuity follows from a uniform bound on the α-Hölder seminorm, and furthermore, for any $x \in A_K$, and any subdivision (t_i) of $[0,T]$,

$$\sum_i |x(t_i) - x(t_{i-1})|^q \le [x]_\alpha \sum_i |t_i - t_{i-1}|^{\alpha q} \le K \sum_i |t_i - t_{i-1}| = KT.$$

To prove boundedness of A_K as a subset of $G\Omega_q$, it remains to show that the q-variation of the second level of the signature is bounded in the sense of (2.5). Indeed, for any subdivision (t_i) of $[0,T]$,

$$\sum_i \|\,\mathrm{S}^2_{s,t}(x)\|^{q/2} \le K \sum_i |t_i - t_{i-1}|^{\alpha q} \le KT.$$

□

Lemma I.B.8. *For $f \colon [0,T] \to X$ for some normed space X, and $\alpha \in (0,1]$,*

$$[f]_\alpha \le \frac{2}{1-2^{-\alpha}} [f]_{\alpha,\mathcal{D}}.$$

Proof. Let $M \coloneqq [f]_{\alpha,\mathcal{D}}$, and let $0 \le s < t \le 1$. Suppose that neither s nor t are dyadic, otherwise the proof simplifies. Let $[v_0, w_0] \subsetneq [s,t]$ be the largest dyadic interval of the form $[k2^{-n}, (k+1)2^{-n}]$ that is fully contained in $[s,t]$. Then,

$$\frac{|f(t)-f(s)|}{|t-s|^\alpha} \le \frac{|f(t)-f(w_0)|}{|t-s|^\alpha} + \underbrace{\frac{|f(w_1)-f(w_0)|}{|w_1-w_0|^\alpha}}_{\le M} + \frac{|f(v_0)-f(s)|}{|t-s|^\alpha}$$

Now let $n_1 \ge n+1$ be the smallest integer such that $s < v_1 \coloneqq v_0 - 2^{-m}$. Then $|t-s| \ge 2|v_1 - v_0|$, so

$$\frac{|f(v_0)-f(s)|}{|t-s|^\alpha} \le \underbrace{\frac{|f(v_0)-f(v_1)|}{|t-s|^\alpha}}_{\le 2^{-\alpha}M} + \frac{|f(v_1)-f(s)|}{|t-s|^\alpha}.$$

Continuing in this way, we get

$$\frac{|f(v_0)-f(s)|}{|t-s|^\alpha} \le M\sum_{i=1}^{m} 2^{-\alpha i} + \frac{|f(v_m)-f(s)|}{|t-s|^\alpha}$$

for all m. By construction, $s < v_m < s + 2^{-n_m}$ where $n_m \ge n+m \to \infty$, so $f(v_m) \to f(s)$ and therefore letting $m \to \infty$ gives

$$\frac{|f(v_0)-f(s)|}{|t-s|^\alpha} \le M\sum_{i=1}^{\infty} 2^{-\alpha i} = M\frac{2^{-\alpha}}{1-2^{-\alpha}}.$$

Proceeding similarly with $\frac{|f(t)-f(w_1)|}{|t-s|^\alpha}$, we obtain

$$\frac{|f(t)-f(s)|}{|t-s|^\alpha} \le M\left(1 + 2\frac{2^{-\alpha}}{1-2^{-\alpha}}\right) \le \frac{2M}{1-2^{-\alpha}}.$$

Since s,t were chosen arbitrarily, this finishes the proof. □

The primary remaining ingredient in proving Theorems 3.4* and 3.6* is to confirm the conditions of Lemma I.B.7. This is the content of the next two lemmas.

Lemma I.B.9. *Let $\alpha \in (1/3, 1/2)$. Suppose we have probability measures μ_n on $C^1([0,1],\mathbb{R})$ that match the Brownian expected signature up to some degree $D > \frac{1}{1/2-\alpha}$ on dyadic intervals up to order $m_n \in \mathbb{N}$ (that is, on intervals of the form $[k2^{-m}, (k+1)2^{-m}]$ for $m \le m_n$), and*

$$\limsup_{n\to\infty} \mu_n\left(x\colon \|x\|_{Lip} \ge 2^{m_n(1-\alpha)}C\right) \longrightarrow 0, \quad C \to \infty. \tag{I.B.6}$$

Then, (μ_n) is tight in $G\Omega_p$ for all $p \in (1/\alpha, 3)$.

If μ_n is supported on piecewise linear paths with maximum slope $\sqrt{N_n}$ where $N_n = 2^{k_n}$, then (I.B.6) *is implied by $k_n \le 2(1-\alpha)m_n + C$ for some $C > 0$.*

Lemma I.B.10. *Let $N_n, D_n \in \mathbb{N}$ such that $N_n, D_n \to \infty$, and suppose that μ_n for $n \in \mathbb{N}$ is an atomic probability measure on $C^1([0,1], \mathbb{R})$ consisting of piecewise linear paths of step length $1/N$ and step size $\pm 1/\sqrt{N}$ such that every set of D_n increments are independent. Then (μ_n) is tight in $G\Omega_p$ for all $p \in (2,3)$.*

Before proving Lemmas I.B.9 and I.B.10, we show how they imply Theorems 3.4* and 3.6*.

Proof of Theorems 3.4 and 3.6*.* Let $(\lambda_i^n, \omega_i^n)$ be the ARCANE cubature specified in Theorem 3.4 (resp. Theorem 3.6), and let μ_n be the associated atomic probability measure on $G\Omega_p$:

$$\mu_n = \sum_i \lambda_i^n \delta_{\omega_i^n}.$$

Let μ denote the probability distribution of canonical Brownian motion on $G\Omega_p$, then by the universal limit theorem for rough paths (Theorem 2.5) it suffices to show that $\mu_n \implies \mu$. To do so, we verify the assumptions of Lemma I.B.5 (ii).

By Lemma I.B.9 (resp. Lemma I.B.10), (μ_n) is tight. Since μ_n is atomic it is a member of $\mathcal{M}$, and because $D_n \to \infty$ (and $N_n \to \infty$ for Theorem 3.6) we have $\mu_n(\mathrm{S}) \to \mu(\mathrm{S})$ pointwise, and therefore by Lemma I.B.6 also uniformly integrability. Finally, by Lemma I.B.4 we have that $\mu \in \mathcal{M}$ and that it satisfies (I.B.2). □

Proof of Lemma I.B.9. Consider a particular dyadic interval $[v, w]$ of length 2^{-k} with $k \geq m$, then

$$\frac{|x(w) - x(v)|}{|w - v|^\alpha} \leq \|x\|_{\mathrm{Lip}} |v - w|^{1-\alpha} \leq \|x\|_{\mathrm{Lip}} 2^{-k(1-\alpha)} \leq \|x\|_{\mathrm{Lip}} 2^{-m(1-\alpha)}.$$

Therefore,

$$\mu_n \left(x \colon \sup_{|s-t| > 2^{-m_n}} \frac{|x(t) - x(s)|}{|t-s|^\alpha} \geq C \right) \leq \mu_n \left(x \colon \|x\|_{\mathrm{Lip}} \geq 2^{m_n(1-\alpha)} C \right) \tag{I.B.7}$$

If μ_n is supported on paths with Lipschitz constant $\sqrt{N_n} = 2^{k_n/2}$, then it is sufficient if $k_n \leq 2(1-\alpha)m_n + C$ for some $C > 0$. By (I.B.6), the $\limsup_{n\to\infty}$ of the left-hand side above goes to zero as $C \to \infty$.

Now if $[v, w]$ is a dyadic interval of length 2^{-k} for some $k \leq m$, then the expected signature on that interval matches that of Brownian motion up to order D; in particular, the moments of the increment $|x(w) - x(v)|$ matches up to order D, so

$$\begin{aligned}
\mu_n \left(\frac{|x(w) - x(v)|}{|w - v|^\alpha} \geq C \right) &\leq C^{-D} 2^{k\alpha D} \mathbb{E}\left[|x(w) - x(v)|^D \right] \\
&\leq C^{-D} 2^{k\alpha D} \mathbb{E}\left[|\mathcal{N}(0, 2^{-k})|^D \right] \\
&= C^{-D} 2^{k\alpha D} 2^{-kD/2} \mathbb{E}\left[|\mathcal{N}(0,1)|^D \right] \\
&\leq c(D) C^{-D} 2^{-kD(1/2-\alpha)}
\end{aligned}$$

Given that $D > 1/(1/2-\alpha)$, we can proceed just as in the proof of Lemma I.B.10 to obtain

$$\limsup_{n\to\infty} \mu_n \left(x\colon \sup_{|s-t|\le 2^{-m_n}} \frac{|x(t)-x(s)|}{|t-s|^\alpha} \ge C \right) \longrightarrow 0, \quad C\to\infty,$$

complementing (I.B.7).

For the second level, just like in the proof of Lemma I.B.10 it suffices to study $S^{(10)}$. If $[v,w]$ is a dyadic interval of length 2^{-k} for some $k \le m_n$, then

$$\begin{aligned}\mu_n \left(\frac{|\,\mathrm{S}^{(10)}_{v,w}\,|}{|v-w|^{2\alpha}} \ge C \right) &\le C^{-D}|v-w|^{-2\alpha D} \underbrace{\mathbb{E}_n\left[|\,\mathrm{S}^{(10)}_{v,w}\,|^D\right]}_{\le c(D)|v-w|^{3D/2}} \\ &\le c(D) C^{-D} 2^{-2kD(3/4-\alpha)}.\end{aligned}$$

Again we proceed with a union bound over all dyadic intervals of length 2^{-k} and over all $k \le m_n$, which succeeds as long as

$$2D(3/4-\alpha) > 1 \iff 2D > \frac{1}{3/4-\alpha} \iff D > \frac{1}{3/4-\alpha},$$

which follows from our assumption that $D > 1/(1/2-\alpha)$.

If $v - w < 2^{-m}$, then

$$|\,\mathrm{S}^{(10)}_{v,w}\,| \le \|x\|_{\mathrm{Lip}} |\,\mathrm{S}^{(00)}_{v,w}\,| = \frac{1}{2}\|x\|_{\mathrm{Lip}} |v-w|^2,$$

so

$$\limsup_{n\to\infty} \mu_n \left(\frac{|\,\mathrm{S}^{(10)}_{v,w}\,|}{|v-w|^{2\alpha}} \ge C \right) \le \limsup_{n\to\infty} \mu_n \left(\|x\|_{\mathrm{Lip}} \ge 2C 2^{2m_n(1-\alpha)} \right) \longrightarrow 0, \quad C\to\infty.$$

□

Proof of Lemma I.B.10. Since μ_n is atomic it is automatically in $\mathcal{M}$. Thus by Lemma I.B.7 it suffices to prove (I.B.4) and (I.B.5) for arbitrary $\alpha \in (0, 1/2)$.

We assume for simplicity that $d = 1$; the proof for $d > 1$ involves more notation but is otherwise the same. We further assume for simplicity that $N_n = 2^{m_n}$ for some $m_n \in \mathbb{N}$, $m_n \to \infty$. Fix $n \in \mathbb{N}$, and abbreviate, if no confusion is possible $m = m_n$, $N = N_n$, etc. Consider a particular dyadic interval $[v,w]$ of length 2^{-k}. Note that if $k \ge m$, then, for any x in the support of μ_n,

$$\frac{|x(w)-x(v)|}{|w-v|^\alpha} \le \frac{N^{-1/2}}{2^{-k\alpha}} = 2^{k\alpha - m/2} \le 2^{-(1/2-\alpha)m} \to 0, \quad n\to\infty. \tag{I.B.8}$$

Hence let $k < m$. Then the increment $x(w)-x(v)$, under the probability measure μ_n, is a sum S of 2^{m-k} Rademacher random variables, any D of which are independent, multiplied with $1/\sqrt{N}$. In particular, moments of S up to order D coincide with the moments of

the same sum where *all* Rademacher variables are independent, which has the same distribution as $\frac{1}{\sqrt{N}}\left(2\mathrm{Bin}\left(2^{m-k},\frac{1}{2}\right)-2^{m-k}\right)$. Therefore, using Markov's inequality, for any $l\in\mathbb{N}$ with $2l\le D$,

$$\begin{aligned}\mu_n\left(\frac{|x(w)-x(v)|}{|w-v|^\alpha}\ge C\right)&=\mu_n\left(|S|\ge C2^{-k\alpha}\right)\\&\le C^{-2l}2^{2\alpha lk}\mathbb{E}\left[S^{2l}\right]\\&=C^{-2l}2^{2\alpha lk}(2/\sqrt{N})^{2l}\mathbb{E}\left[\left(\mathrm{Bin}\left(2^{m-k},\frac{1}{2}\right)-2^{m-k-1}\right)^{2l}\right]\\&\le c(l)C^{-2l}2^{2\alpha lk}2^{-ml}(2^{m-k-1})^l\\&\le c(l)C^{-2l}2^{-(1-2\alpha)lk}\end{aligned}$$

Here we used the well-known bound $\mathbb{E}[(\mathrm{Bin}(n,p)-np)^{2l}]\le c(l)(np)^l$. Using union bound over all 2^k such dyadic intervals, as well as over $k\in\{0,\dots,m-1\}$, we get, as long as $2^{-(1/2-\alpha)m}\le C$ (recall (I.B.8)),

$$\begin{aligned}\mu_n\left(x\colon [x]_{\alpha,\mathcal{D}}\ge C\right)&\le c(l)\sum_{k=0}^{m_n}2^kC^{-2l}2^{-(1-2\alpha)lk}\\&=c(l)C^{-2l}\sum_{k=0}^{m_n}2^{-k[(1-2\alpha)l-1]}.\end{aligned}$$

Now if $l>1/(1-2\alpha)$, then we get an upper bound of the form

$$\mu_n\left(x\colon [x]_{\alpha,\mathcal{D}}\ge C\right)\le c(l,\alpha)C^{-2l}$$

for some constant $c=c(l,\alpha)>0$. Since α is fixed, we can choose $l=l(\alpha)=\lceil 1/(1-2\alpha)\rceil$ and then for all $n\ge n_0(\alpha)$ we have $D_n\ge l$ and therefore

$$\limsup_{n\to\infty}\mu_n(x\colon [x]_{\alpha,\mathcal{D}}\ge C)\le c(\alpha)C^{-2l}\longrightarrow 0,\quad n\to\infty.$$

It remains to show (I.B.5). Note first that $\mathrm{S}^{(00)}_{s,t}(x)=(t-s)^2/2$ and $\mathrm{S}^{(11)}_{s,t}(x)=(x(t)-x(s))^2/2$, so it suffices to study $\mathrm{S}^{(10)}$ and $\mathrm{S}^{(01)}$. Since further $\mathrm{S}^{(10)}+\mathrm{S}^{(01)}=\mathrm{S}^{(1)}\,\mathrm{S}^{(0)}$, it in fact suffices to study only $\mathrm{S}^{(10)}$. Now, if $(s,t)=(v,w)$ is dyadic as above with $w-v=2^{-k}$, and the $N'\coloneqq 2^{m-k}$, D-wise independent Rademacher variables determining the increments in $[v,w]$ are denoted ε_i, then

$$\mathrm{S}^{(10)}_{v,w}(x)=\sum_{i=1}^{2^{m-k}}\frac{\varepsilon_i}{2N^{3/2}}+\sum_{i<j}\frac{\varepsilon_j}{N^{3/2}}$$

The first term is just $\frac{1}{2N}(x(w)-x(v))\le(w-v)(x(w)-x(v))$, the absolute value of which is at most $C|w-v|^{1+\alpha}\le C|w-v|^{2\alpha}$ with high probability by what we already proved.

The second term is equal to $N^{-3/2}\sum_j (j-1)\varepsilon_j$; note that

$$\begin{aligned}
\mathbb{E}\left[\left(\sum_j (j-1)\varepsilon_j\right)^{2l}\right] &= \sum_{j_1,\dots,j_{2l}=1}^{N'} \prod_{i=1}^{2l} \underbrace{(j_i-1)}_{\le N'} \underbrace{\mathbb{E}\left[\varepsilon_{j_1}\dots\varepsilon_{j_{2l}}\right]}_{\in\{0,1\}} \\
&\le \sum_{j_1,\dots,j_{2l}=1}^{N'} N'^{2l}\mathbb{E}\left[\varepsilon_{j_1}\dots\varepsilon_{j_{2l}}\right] \\
&= \mathbb{E}\left[\left(\sum_j N'\varepsilon_j\right)^{2l}\right] \\
&= N'^{2l}\mathbb{E}\left[\left(2\mathrm{Bin}\left(N',\frac{1}{2}\right)-N'\right)^{2l}\right] \\
&\le c(l)N'^{3l}.
\end{aligned}$$

All together, this gives

$$\begin{aligned}
\mu_n\left(\left|\sum_{i<j}\frac{\varepsilon_j}{N^{3/2}}\right| \ge C|v-w|^{2\alpha}\right) &\le c(l)C^{-2l}2^{4lk\alpha}\left(\frac{N'}{N}\right)^3 \\
&\le c(l)C^{-2l}2^{lk(4\alpha-3)},
\end{aligned}$$

which, with the usual union bound over all 2^{-k}-based dyadic intervals and $k=0,\dots,m$, gives an upper bound of

$$c(l)C^{-2l}\sum_{k=0}^{n} 2^{-k[l(4\alpha-3)-1]},$$

which again is summable for sufficiently large $l = l(\alpha)$. □

Chapter II

Genealogies of Spatial Population Models

1 Introduction

Coalescent processes arise when describing the genealogies, that is family relations, of a biological population by tracing the ancestries of a sample from the population backwards in time. If individuals in the population have exactly one parent, such as in haploid asexual populations, then every sampled individual has a single ancestral lineage going backwards in time, and the lineages associated with a set of individuals coalesce when their common ancestor is reached. We will not consider diploid populations, but remark that they can be treated as haploid populations of twice the size if one is only interested in a single gene, see e.g. [18], Remark 2.1. The most basic coalescent process is Kingman's coalescent, introduced in 1982 [26], in which every pair of lineages coalesces at constant rate, independently of all other pairs.

In this project, we are ultimately interested in the genealogies of *structured* populations, where every individual has an evolving type or location, but we begin with a brief overview of some classical theory in the unstructured case. The reader familiar with the field should feel free to skim Section 1.1.

1.1 Coalescents

Mathematically, a coalescent is typically described by giving disjoint, set-valued labels to the lineages, which merge when the lineages coalesce. Denote by $\mathcal{P}$ the set of partitions of finite subsets of $\mathbb{N}$.

Definition 1.1. A *coalescent* is a $\mathcal{P}$-valued right-continuous Markov process $(\Pi_t)_{t\geq 0}$ whose transitions can be obtained by merging partition elements. It is called *label invariant* if its law is invariant under changing the labels of the initial set of lineages.[1]

[1] In such a way that does not necessarily preserve their size. For example, the substitution $\{1,3\} \to \{3,4,9\}$ is allowed.

A precise definition is Definition 2.1. The label invariance condition is equivalent to requiring the counting process $(|\Pi_t|)_{t\geq 0}$ to be Markovian. Another equivalent phrasing is that for any $n, m \in \mathbb{N}$ and $k_1 \geq \ldots \geq k_m \geq 2$ with $\sum_{i=1}^m k_i \leq n$, every coalescence event in which m disjoint sets of lineages of sizes k_1 to k_m merge simultaneously while there are currently n lineages—called an $(n, \vec{k})$-merger—happens at the same rate, which we denote by $\lambda_{n,\vec{k}} = \lambda_{n,k_1,\ldots,k_m}$.

Not every coalescent in the sense of Definition 1.1 arises from the genealogies of a population model, and a natural necessary condition is *sampling consistency*: if we construct the genealogical tree corresponding to a sample of size $n+1$, then the induced genealogical tree for a subsample of size n will have the same distribution as the tree obtained by constructing the coalescent directly from the subsample. An equivalent definition is that there exists a coalescent started from all of $\mathbb{N}$ such that the coalescent started from a finite subset $A \subset \mathbb{N}$ is its restriction to A.

Theorem 1.2 ([13, 36, 39])**.** *A label invariant coalescent with no simultaneous mergers, that is $\lambda_{n,k_1,\ldots,k_m} = 0$ whenever $m > 1$, is sampling consistent if and only if there exists a finite measure Λ on $[0,1]$ such that*

$$\lambda_{n,k} = \int_0^1 p^{k-2}(1-p)^{n-k}\Lambda(\mathrm{d}p), \qquad 2 \leq k \leq n.$$

It is called the Λ-coalescent.

The Kingman coalescent arises as a special case if Λ is the unit mass at zero. If Λ is the uniform measure on $[0,1]$, then the Λ-coalescent corresponds to the clustering process constructed by Bolthausen and Sznitman [7]. A characterisation in the most general case of simultaneous multiple mergers is also known. Denote by $\triangle = \{\boldsymbol{\xi} = (\xi_1, \xi_2, \ldots)\colon \xi_1 \geq \xi_2 \geq \ldots \geq 0, |\boldsymbol{\xi}| \leq 1\}$ the infinite simplex, where $|\boldsymbol{\xi}| = \sum_{i=1}^\infty \xi_i$. Write $(u, v) = \sum_{i=1}^\infty u_i v_i$ for $u, v \in \triangle$.

Theorem 1.3 ([42, 33])**.** *A label invariant coalescent is sampling consistent if and only if there exists a finite measure $\Xi = a\delta_0 + \Xi_0$ on $\triangle$, where δ_0 is the unit mass at zero and Ξ_0 has no atom at zero, such that $\lambda_{n,2} = a$ for all $n \geq 2$, and for any $(n, \vec{k}) \neq (n, 2)$,*

$$\lambda_{n,\vec{k}} = \int_\triangle \sum_{l=0}^s \binom{s}{l} (1-|\boldsymbol{\xi}|)^{s-l} \sum_{i_1 \neq \ldots \neq i_{m+l}} \xi_{i_1}^{k_1} \ldots \xi_{i_m}^{k_m} \xi_{i_m+1} \ldots \xi_{i_{m+l}} \frac{\Xi_0(\mathrm{d}\boldsymbol{\xi})}{(\boldsymbol{\xi}, \boldsymbol{\xi})}, \tag{1.1}$$

where $s = n - \sum_i k_i$. It is called the Ξ-coalescent.

The Λ-coalescent arises as a special case if Ξ is the pushforward of Λ under the map $p \mapsto (p, 0, \ldots)$ from $[0,1]$ to $\triangle$. We introduced sampling consistency as a necessary condition for a coalescent to describe the genealogies of some population, but it is also sufficient. Möhle and Sagitov [33, 38] give conditions for the genealogies of a haploid population

evolving according to Cannings model—that is, a constant size population evolving in discrete generations with an exchangeable offspring mechanism—to converge as the population size tends to infinity when properly rescaled. The class of resulting coalescents is exactly the class of Ξ-coalescents. Conditions for the genealogies to converge to Kingman's coalescent were previously discovered by Kingman [27] with later improvements by Möhle [31, 32]. Analogous results for the Λ-coalescent were independently discovered by Donnelly and Kurtz [13] and Sagitov [39].

1.2 Spatial Coalescents

The Λ- and Ξ-coalescents appear as a universal limiting object for the genealogies of a wide range of population models [3, 14, 41, 40, 8, 10], including diploid populations [34, 5, 11] and populations where offspring sizes are not exchangeable [44]. However, they cannot fully describe the genealogies of a *structured* population, where every individual has a *type* or *location* evolving in some space E. Typically, type is the appropriate interpretation in a genetic model, and location is appropriate in a model of a geographically dispersing population. We usually refer to spatial locations from now on, but the genetic interpretation remains valid throughout. On top of the set of lineages Π_t, an associated coalescent process needs to keep track of the spatial location $\boldsymbol{X}_t(u)$ of every lineage $u \in \Pi_t$, which we summarise in a map $\boldsymbol{X}_t \colon \Pi_t \to E$.

Definition 1.4. A *spatial coalescent* is a right-continuous Feller Markov process $(\Pi_t, \boldsymbol{X}_t)_{t\geq 0}$ such that $\Pi_t \in \mathcal{P}$ and $\boldsymbol{X}_t \colon \Pi_t \to E$ for all $t \geq 0$, and all transitions in (Π_t) can be obtained by merging partition elements. It is called *label invariant* if its law is invariant under changing the labels of the initial set of lineages.

A precise definition is in Definitions 2.3 and 2.4. The Feller property—which is automatic in the non-spatial setting—is a form of continuity in the initial condition that, together with right-continuity of sample paths, implies the strong Markov property [37, Thm. 8.3]. The label invariance assumption is equivalent to the requirement that the empirical measure $\sum_{u\in\Pi_t} \delta_{\boldsymbol{X}_t(u)}$ is a Markov process, which was pointed out to us by Félix Foutel-Rodier.

The picture in the spatial setting is much less complete than in the non-spatial setting. Consider, for instance, a population model in which individuals move as independent Brownian motions on the d-dimensional (flat) torus $E = \mathbb{T}^d = \mathbb{R}^d/\mathbb{Z}^d$ forwards in time, while undergoing branching according to some unspecified (think general) mechanism. The well-known Fleming-Viot process and its Λ- and Ξ-generalisations [12, 4, 6], as well as the class of models considered by Donnelly and Kurtz in their lookdown papers [12, 13], when specialised to Brownian movement, fall into this category. Describing the genealogies of such a population model essentially requires partially time reversing

the process: given a finite sample from the population at present together with their spatial locations (with no knowledge of their history, or the spatial distribution of the remaining population), how do their lineages move through space backwards in time? Since Brownian motion is reversible, they will presumably follow Brownian motions, but there would have to be a coupled drift; if they were independent, lineages would never meet in dimensions $d \geq 2$, and even in $d = 1$ only in pairs. But what kind of drift leads to binary mergers in $d \geq 2$, or multiple, even simultaneous mergers in any dimension? These are interesting and basic questions, yet generally unanswered to the best of our knowledge. Our results will provide some answers (the curious reader may take a look at Example 1.17).

1.3 The Brownian Spatial Coalescent

The genealogies of a spatial model of the kind considered in the previous paragraph will be very complicated to describe in general; they depend on the initial condition of the process, and wouldn't be Markovian without conditioning at least on the total population size process (as is the case in the non-spatial setting, where genealogies are Markovian after a time change that depends on the total population size process, see e.g. [13], Theorem 5.1, for the Kingman coalescent). But if the population is at stationarity forwards in time, and the total population size is constant, then the genealogies are time homogeneous, and there is hope for them to be Markovian. If that is the case, they would be described by a spatial coalescent in the sense of Definition 1.4.

The goal of this project is to identify a universal class of spatial coalescents that describe the genealogies of constant size spatial population models at stationarity in which individuals follow independent Brownian motions. We will do so by axiomatically defining a class of spatial coalescents through a property that we would expect the genealogies of all such populations to satisfy. The idea is as follows: if we condition a realisation of the forwards in time population model on the initial and final locations of all individuals, as well as the times and locations of branch events, then, regardless of the branching mechanism, all individuals follow independent Brownian bridges along the branches of the "genealogical" tree (or forest in general) forwards in time. Interpreting this backwards in time, it means for the associated coalescent that, conditional on the genealogical tree and the times and locations of what are now the coalescence events, lineages follow independent Brownian bridges along the branches of the tree *backwards* in time.

Before turning this into a definition, we point out a subtle difficulty that arises when starting such a coalescent process from an initial configuration with more than one particle in the same location. If $d \geq 2$, then forwards in time, almost-surely the only occasion where any pair of particles would ever be in the same location is when they were just born in the same branching event. The same is true in $d = 1$ for three or more colliding

particles. Thus the coalescent started from such a configuration would have to merge those particles instantaneously, which it cannot if we require right-continuous paths. One way out would be to work with left-continuous processes, but that would leave many technical special cases in dealing with instantaneous mergers at time zero. The simpler solution is to exclude those points from the state space. For $\Pi \in \mathcal{P}$ denote by E^Π the set of maps $\Pi \to E$, and let

$$E_\circ^\Pi \coloneqq \left\{\boldsymbol{x} \in E^\Pi \colon \#\{\{u,v\} \subset \Pi \colon u \neq v, \boldsymbol{x}_u = \boldsymbol{x}_v\} \leq \mathbb{1}_{\{d=1\}}\right\} \tag{1.2}$$

denote the set of spatial decorations of Π in which no (if $d \geq 2$) or at most one (if $d = 1$) pair of particles reside in the same location. We occasionally identify $E_\circ^\Pi$ with an open subset of $E^{|\Pi|}$ by ordering the elements of Π in ascending order of their minima (and similarly in other contexts). Let

$$\mathcal{X} \coloneqq \left\{(\Pi, \boldsymbol{x}) \colon \Pi \in \mathcal{P}, \boldsymbol{x} \in E_\circ^\Pi\right\}.$$

Note that Π is the domain of $\boldsymbol{x}$ for $(\Pi, \boldsymbol{x}) \in \mathcal{X}$, so we may occasionally just write $\boldsymbol{x} \in \mathcal{X}$.

Definition 1.5. A spatial coalescent with state space $\mathcal{X}$ is called a *Brownian spatial coalescent* if, conditional on the genealogical forest, and times and locations of all coalescence events, lineages follow independent Brownian bridges along the branches of the forest, and Brownian motions after the last coalescence event in which they are involved.

A precise definition is Definition 2.10. We remark here that the only real obstacle in replacing Brownian motion with a much more general motion lies in finding the points that have to be excluded from the state space. We go into more detail in Section 1.7.2.

Remark 1.6. The defining property of a Brownian spatial coalescent should be satisfied by the genealogies of any population model in which individuals follow independent Brownian motions forwards in time, regardless of the branching mechanism. This includes spatial interactions such as competition or local regulation. However, it will turn out that assuming the Markov property of the genealogies eliminates *all* branching mechanism that non-trivially depend on the spatial distribution of the population. See Section 1.7.1 below.

We do not yet require any kind of sampling consistency. If a parallel is drawn with the non-spatial setting, then the Brownian spatial coalescent should be compared with the non-spatial coalescent in the sense of Definition 1.1; both contain many examples that do not describe the genealogies of any population. Our first result is a full characterisation of the class of Brownian spatial coalescents, which shows that this comparison is more apt than one might expect. Write M for the set of tuples $(n, \vec{k}) = (n, (k_1, \ldots, k_m))$ with $n \geq 2$, $1 \leq m \leq n$, $k_1 \geq \ldots \geq k_m \geq 1$ and $\sum_{i=1}^m k_i \leq n$. Write $\mathcal{M}_F(A)$ for the set of finite measures on a measurable space A.

Theorem 1.7. *Every label invariant Brownian spatial coalescent is uniquely characterised by a family of measures* $\boldsymbol{\nu} = (\nu_{n,k_1,\dots,k_n} \in \mathcal{M}_F(E^m))_{(n,\vec{k})\in M}$. *We call it the* Brownian spatial coalescent with transition measures $\boldsymbol{\nu}$.

The transition measures are reminiscent of the transition rates of a non-spatial coalescent, and in fact we will see in Remark 1.10 that if we formally replace E with a singleton, then the Brownian spatial coalescent with transition measures $(\nu_{n,\vec{k}})$ is just the non-spatial coalescent with transition rates $(\lambda_{n,\vec{k}} = |\nu_{n,\vec{k}}|)$. Here $|\nu|$ is the total mass of a measure ν.

We refer to Theorem 2.14 for a formal version of Theorem 1.7 that applies without the assumption of label invariance, and gives a precise description of the law of a Brownian spatial coalescent with given transition measures $\boldsymbol{\nu}$ building on the notation introduced in Section 2. Here, we give an informal description in the case where there are no simultaneous mergers, so $\boldsymbol{\nu} = (\nu_{n,k}\colon 2 \le k \le n)$, and we further assume that for every $n \ge 2$ there is some $k \in \{2,\dots,n\}$ with $\nu_{n,k} \neq 0$. This ensures that almost-surely all lineages will merge eventually, so the genealogical forest is always a tree. Under these assumptions, the law of the Brownian spatial coalescent with transition measures $\boldsymbol{\nu}$, started from initial locations x_1 to x_n is as follows. For every possible shape (topology) T of the genealogical tree on n leaves, a measure ν_T is defined on the space of time and spatial coordinates of coalescence events in the tree, which is most easily described by picture (Fig. 1.1): for every branch in the tree there is a spatial factor, illustrated in cyan in Fig. 1.1, where $(p_t\colon E \to (0,\infty))_{t>0}$ denotes the transition densities of a Brownian motion on E with periodic boundary conditions; and for every merge event there is a factor of the form $\mathrm{e}^{-|\nu_n|(\tau-\tau')}$, where τ and τ' are the times of the merge event and the previous merge event (or zero), respectively, and

$$|\nu_n| := \sum_{2\le k\le n} \binom{n}{k} |\nu_{n,k}|, \quad n \ge 2,$$

can be thought of as the analogue of the total event rate λ_n while at n lineages in the non-spatial setting. The density in the time coordinates is with respect to Lebesgue measure, and in the spatial coordinate z associated with an (n,k)-merger it is with respect to the transition measure $\nu_{n,k}(\mathrm{d}z)$. The measure ν_T is not finite for arbitrary choice of $\boldsymbol{\nu}$, but it will be for the transition measures of a Brownian spatial coalescent. Now,

(i) Sample a tree shape, where the probability of sampling T is proportional to $|\nu_T|$.

(ii) Given T, sample time and space coordinates of all merge events from $\nu_T/|\nu_T|$.

(iii) Given T and time and space coordinates of all merge events, run independent Brownian bridges along the branches of the tree, starting from x_1 to x_n at the leaves, and a Brownian motion after the final merge event.

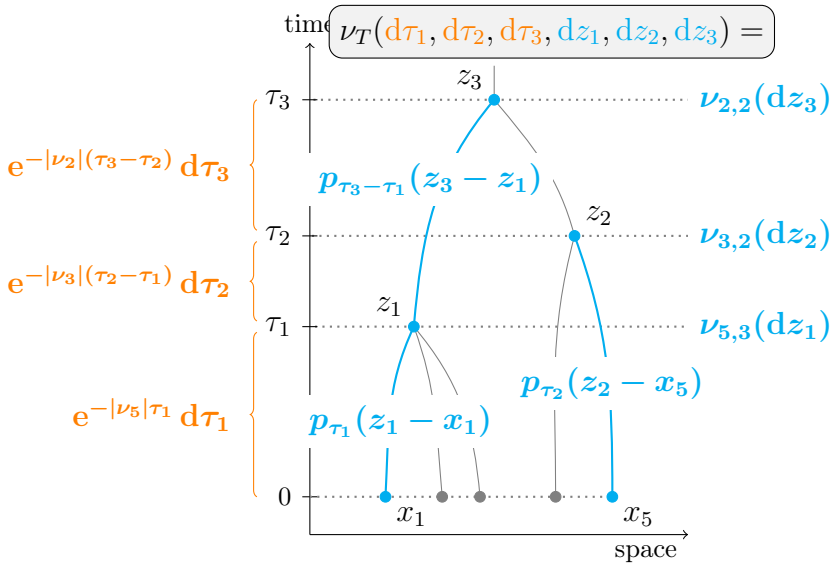

Figure 1.1: Illustration of the finite measure on the space and time decorations of a particular genealogical tree shape. Spatial factors associated with branches of the tree are only shown for three of the seven branches.

Equivalently to (i) and (ii), we could sample a tree together with its time and space decorations directly from the normalised formal sum of all the measures ν_T. The normalisation

$$N^{\boldsymbol{\nu}}(\boldsymbol{x}) = \sum_T |\nu_T| \tag{1.3}$$

as a function of the initial condition $\boldsymbol{x} = (x_1, \dots, x_n) \in \mathcal{X}$ plays an important role. For example, it decides for a given family $\boldsymbol{\nu}$ of transition measures whether an associated Brownian spatial coalescent exists.

Theorem 1.8. *For a given family $\boldsymbol{\nu} = (\nu_{n,k_1,\dots,k_m} \in \mathcal{M}_F(E^m))_{(n,\vec{k})\in M}$, a Brownian spatial coalescent with transition measures $\boldsymbol{\nu}$ exists if and only if $N^{\boldsymbol{\nu}} \colon \mathcal{X} \to [0,\infty]$ is finite and continuous.*

See Theorem 2.14. The main task in the proof of Theorem 1.8 is to show that the process defined through the transition measures $\boldsymbol{\nu}$ is a Markov process, which is not obvious from its description.

Continuity of $N^{\boldsymbol{\nu}}$ corresponds to the Feller property of the associated coalescent; for the latter to be well-defined, finiteness of $N^{\boldsymbol{\nu}}$ would suffice. We do not have a nice characterisation of the measures $\boldsymbol{\nu}$ for which $N^{\boldsymbol{\nu}}$ is finite and continuous, but a sufficient condition is that all transition measures are constant multiples of Lebesgue measure, see Lemma 2.12. Since $N^{\boldsymbol{\nu}}(\boldsymbol{x})$ is an integral over functions that are continuous in $\boldsymbol{x}$, dominated convergence gives a sufficient integrability condition. The function $N^{\boldsymbol{\nu}}$ also reveals a technical reason why it was necessary to exclude some points from the state space $\mathcal{X}$: If $\boldsymbol{x}$ approaches a point in the boundary of $\mathcal{X}$, then $N^{\boldsymbol{\nu}}(\boldsymbol{x})$ diverges if the transition

measure associated with the coalescent event that would have to happen instantaneously is non-zero, see also Remark 2.13.

An example and some remarks are in order.

Example 1.9. Starting from two lineages at locations x_1 and x_2, the probability that they coalesce at time $t > 0$ at location $z \in E$ is proportional to

$$\mathrm{e}^{-|\nu_2|t} p_t(x_1 - z) p_t(x_2 - z) \nu_{2,2}(\mathrm{d}z)\, \mathrm{d}t. \tag{1.4}$$

In particular, if $\nu_{2,2}(\mathrm{d}z) = |\nu_{2,2}|\, \mathrm{d}z$ (which will turn out to be a consequence of sampling consistency, see Theorem 1.12), then the probability that they reach a common ancestor at time $t > 0$ is proportional to

$$|\nu_2| \mathrm{e}^{-|\nu_2|t} p_{2t}(x_1 - x_2)\, \mathrm{d}t, \tag{1.5}$$

which is a form of the Wright-Malécot formula. It follows from (1.4) that conditional on t, the coalescence location z is equal in law to the midpoint of a Brownian bridge running from x_1 to x_2 in time $2t$.[2]

The normalisation in (1.5) as a function of x_1 and x_2 is

$$N^{\boldsymbol{\nu}}(x_1, x_2) = \int_0^\infty |\nu_2| \mathrm{e}^{-|\nu_2|t} p_{2t}(x_1 - x_2)\, \mathrm{d}t.$$

If $r = |x_1 - x_2|$ is small, then $N^{\boldsymbol{\nu}}(x_1, x_2)$ is of order r^{2-d} if $d \geq 3$, of order $\log \frac{1}{r}$ if $d = 2$, and of order 1 if $d = 1$. If the locations of more than two individuals are known, then similar calculations lead to higher-order versions of the Wright-Malécot formula.

Remark 1.10. (i) If we formally replace E with a singleton, and Brownian motion with the trivial Markov process on E, then the spatial factors in Fig. 1.1 vanish, the transition measures reduce to multiplicative constants, and by a simple argument about competing exponential clocks, the remaining exponential structure exactly describes the law of the non-spatial coalescent with transition rates $\lambda_{n,k} = |\nu_{n,k}|$. The same is true if simultaneous mergers are allowed.

(ii) This is a very "non-Markovian" description of a Markov process. The Markov property in terms of this description states that, if we sample the decorated coalescence tree at time zero and start running the particles along the corresponding Brownian bridges, and at some time $t > 0$ we resample the decorated coalescence tree for the remaining lineages using their current locations as initial conditions, and run them along the Brownian bridges associated with the new tree, then this does not affect the overall law of the process.

[2]This observation was made by Andraž Jelinčič after listening to a talk about the project.

The main work in the proof of Theorem 1.7 is to construct the measures $\nu_{n,\vec{k}}$ using only the defining property (Definition 1.5) and the Markov property of a given Brownian spatial coalescent. The idea is to apply the Markov property at the (random) time of the first merge event, and proceed inductively over the number of initial particles, starting with two.

1.4 Sampling Consistency

Our main result is a characterisation of sampling consistency within the class of Brownian spatial coalescents, analogous to the characterisations of the Λ- and Ξ-coalescents in the non-spatial setting (Theorems 1.2 and 1.3). A characterisation of this form is, to our knowledge, without precedent in the literature on spatial coalescents.

Sampling consistency for spatial coalescents is classically taken to mean that the distribution of the coalescent started from some set of n initial locations is the same as that of the coalescent started with an additional $n+1$'st particle at any fixed location and projected back onto the first n particles. We would expect the genealogies of a forward in time population model to have this property, but only if knowing the location of an additional individual does not reveal any information about the population. This is the case, for example, if the population is spread homogeneously throughout space and time, as is the case for a number of spatial models whose genealogies have been studied successfully, which includes the spatial Λ-Fleming-Viot process, or populations living on a discrete lattice with positive population density at every deme. We review related literature in detail in Section 1.6. But in our case, where the spatial distribution of the population fluctuates in time and space driven by its forward dynamics, the location of the additional particle will reveal information about the entire population. This will, in general, influence the genealogy of the first n particles. Instead, the additional particle has to be genuinely *sampled* from the (stationary) spatial distribution of the population, *conditional* on the n lineages whose locations are already known.

Definition 1.11. A spatial coalescent is *sampling consistent* if, for every choice of initial locations x_1 to x_n, there exists a probability measure $\mu_{x_1,\dots,x_n}$ on E such that the coalescent started from x_1 to x_n has the same distribution as the coalescent induced on the first n particles when starting with an additional particle x_{n+1} sampled from $\mu_{x_1,\dots,x_n}$.

See Definition 2.5 in the next section for a precise definition. We will prove that, as the motivation suggests, $\mu_{x_1,\dots,x_n}$ is the law of a sample from the stationary distribution of an associated forwards in time population model, conditional on having sampled x_1 to x_n in n previous samples. The familiar reader may also note that Definition 1.11 can be cast in the language of the Markov mapping theorem: $\mu_{x_1,\dots,x_n}$ is the kernel that recovers the state of the "larger" Markov process (the coalescent started from $n+1$ particles) from the state of the projected one (the coalescent started from n particles).

Theorem 1.12. *A label invariant Brownian spatial coalescent is sampling consistent if and only if there exists a finite measure Ξ on $\triangle$ such that the transition measures satisfy*

$$\nu_{n,\vec{k}}(\mathrm{d}\boldsymbol{z}) = \lambda_{n,\vec{k}}\,\mathrm{d}\boldsymbol{z}, \qquad (n,\vec{k}) \in M,$$

where $(\lambda_{n,\vec{k}})$ are the rates of the non-spatial Ξ-coalescent.

This gives rise to what we call the *Brownian spatial Ξ-coalescent*. If simultaneous mergers are not allowed, we get an analogous statement for what we call the *Brownian spatial Λ-coalescent*. If $\boldsymbol{\nu}$ are the transition measures of the Brownian spatial Ξ-coalescent, write $N^{\Xi} = N^{\boldsymbol{\nu}}$.

Proposition 1.13. *Let $\Xi \neq 0$ be a finite measure on $\triangle$. Then the family of probability measures associated with the Brownian spatial Ξ-coalescent through Definition 1.11, which we denote by $(\mu^{\Xi}_{x_1,\dots,x_n})$, is unique and*

$$\mu^{\Xi}_{x_1,\dots,x_n}(\mathrm{d}y) = \frac{N^{\Xi}(x_1,\dots,x_n,y)}{N^{\Xi}(x_1,\dots,x_n)}\,\mathrm{d}y. \tag{1.6}$$

Furthermore, there exists a unique random probability measure μ^{Ξ} on E such that

$$\mathbb{E}\left[\mu^{\Xi}(\mathrm{d}x_1)\dots\mu^{\Xi}(\mathrm{d}x_n)\right] = N^{\Xi}(x_1,\dots,x_n)\,\mathrm{d}x_1\dots\mathrm{d}x_n \tag{1.7}$$

as measures on E^n for every $n \in \mathbb{N}$.

Note that Proposition 1.13 reveals another meaning of the normalisation function N^{Ξ}: it describes the joint density of samples from a random realisation of μ^{Ξ}.

Remark 1.14. A spatial coalescent process is sampling consistent in the classical sense if and only if it is sampling consistent in the sense of Definition 1.11 with respect to (w.r.t.) *any* choice of probability measures $(\mu_{x_1,\dots,x_n})$. But since the probability measures in Proposition 1.13 are unique, there exists no Brownian spatial coalescent, except the trivial one, which is sampling consistent in the classical sense. This may explain why no non-trivial, sampling consistent spatial coalescent in which an individual lineage follows a Brownian motion has been found to this day in the setting of constant population density (where we would expect associated coalescents to be sampling consistent in the classical sense), except in one dimension where independent Brownian motions meet.

Eqs. (1.6) and (1.7) imply

$$\mathbb{E}\left[\mu^{\Xi}(\mathrm{d}x_1)\dots\mu^{\Xi}(\mathrm{d}x_n)\right] = \mathrm{d}x_1\mu^{\Xi}_{x_1}(\mathrm{d}x_2)\dots\mu^{\Xi}_{x_1,\dots,x_{n-1}}(\mathrm{d}x_n),$$

that is, $\mu^{\Xi}_{x_1,\dots,x_n}$ is the distribution of a sample from μ^{Ξ} conditional on having already sampled x_1 to x_n in n previous, independent samples. This confirms the motivation behind Definition 1.11, provided we can find a forwards in time population model with

stationary distribution μ^{Ξ} whose genealogies are described by the Brownian spatial Ξ-coalescent. This will be the Ξ-Fleming-Viot process.

Before showing that, we present a drift representation for the evolution of the Brownian spatial coalescent. The suggestion to look for a representation of this kind was made by Michal Bassan. For $\boldsymbol{\nu} \in R$ and $\boldsymbol{\lambda} \in R^\circ$ we write $\boldsymbol{\nu}(\mathrm{d}\boldsymbol{z}) = \boldsymbol{\lambda}\,\mathrm{d}\boldsymbol{z}$ as a short hand for $\nu_{n,\vec{k}}(\mathrm{d}\boldsymbol{z}) = \lambda_{n,\vec{k}}\,\mathrm{d}\boldsymbol{z}$ for all $(n,\vec{k}) \in M$.

Theorem 1.15. *If a Brownian spatial coalescent with transition measures $\boldsymbol{\nu}$ of the form $\boldsymbol{\nu}(\mathrm{d}\boldsymbol{z}) = \boldsymbol{\lambda}\,\mathrm{d}\boldsymbol{z}$ for some $\boldsymbol{\lambda} \in R^\circ$ is started from distinct initial locations x_1 to x_n, then the paths $\boldsymbol{Z}_t = (Z_t^1, \ldots, Z_t^n)$ of the lineages until just before the first merge event (or until just before the first time that any pair of lineages meet if $d = 1$) are described by the following stochastic differential equation:*

$$\mathrm{d}\boldsymbol{Z}_t = \mathrm{d}\boldsymbol{B}_t + \nabla \log N^{\boldsymbol{\nu}}(\boldsymbol{Z}_t)\,\mathrm{d}t, \tag{1.8}$$

where $(\boldsymbol{B}_t)$ is an nd-dimensional standard Brownian motion on E with periodic boundary conditions.

In particular, Theorem 1.15 applies to the Brownian spatial Ξ-coalescent. It also holds for a general Brownian spatial coalescent provided $N^{\boldsymbol{\nu}}$ is differentiable and the order of the differentiation and the integrals defining $N^{\boldsymbol{\nu}}$ can be exchanged, in the way that is needed in the proof; see Section 3.4. The following remark and example apply only to Brownian spatial coalescents of the form assumed in Theorem 1.15.

Remark 1.16. (i) If $d \geq 2$, then the drift representation determines the law of a Brownian spatial coalescent completely: run (1.8) until its maximal existence time, where the drift diverges and some set of lineages collide. Coalesce lineages that are in the same location and restart (1.8). Repeat until only one lineage is left (or a number k such that $|\nu_k| = 0$).

(ii) If $d = 1$, then lineages may meet in pairs without coalescing. At such points $N^{\boldsymbol{\nu}}$ is (continuous and) not differentiable, but it is possible to extend $\nabla \log N^{\boldsymbol{\nu}}$ to all of $\mathcal{X}$ in such a way that the drift representation (1.8) holds up until just before the first merge event. But unlike in $d \geq 2$, binary mergers always happen before the maximal existence time of (1.8), making the construction in (i) invalid, except if binary mergers are impossible because $\nu_{k,2} = 0$ for the relevant values of k. We leave it open to find a description of the law with which binary mergers occur conditional on the solution of (1.8); calculations suggest that they do *not* happen at a rate proportional to the collision local time of pairs of lineages, as is common in models of coalescing Brownian motions.

The function $N^{\boldsymbol{\nu}}$ is a sum over forests, whose number explodes combinatorially as the number of initial locations increases, but we can gain some interesting insights in simple special cases.

Example 1.17. Let $\boldsymbol{\nu} \in R$ with $\boldsymbol{\nu}(\mathrm{d}\boldsymbol{z}) = \boldsymbol{\lambda}\,\mathrm{d}\boldsymbol{z}$ for some $\boldsymbol{\lambda} \in R^\circ$, and consider the Brownian spatial coalescent with transition measures $\boldsymbol{\nu}$.

(i) If $\lambda_{2,2} > 0$ and the Brownian spatial coalescent is started from two lineages, then the drift symmetrically pulls them towards each other,[3] and the strength of the drift as their separation r tends to zero is of order 1 if $d = 1$, otherwise of order $1/r$. This is not surprising in light of the fact that the Brownian movement in $d \geq 2$ effectuates a repulsive drift of the same order $1/r$.

(ii) If $\lambda_{3,3} > 0$ and the Brownian spatial coalescent is started from three individuals, and conditioned on a triple merger, then the drift pulls each particle in the direction of the midpoint of the other two.[4] Without the conditioning, the drift is mixed in a complicated way with pairwise binary attractions.

(iii) If $\lambda_{4,2,2} > 0$ and $\lambda_{2,2} = 0$, and the Brownian spatial coalescent is started from four individuals at locations x_1 to x_4, and conditioned on a simultaneous binary merger of x_1, x_2, and x_3, x_4, then the drift symmetrically pulls (x_1, x_3) and (x_2, x_4) towards each other, and the strength of the drift is of order $1/r$ where r is their $2d$-dimensional Euclidean separation. We put $\lambda_{2,2} = 0$ because the drift is otherwise influenced in a complex way by the inevitable second merge event.

Remark 1.18. It is known in a variety of settings that the movement of a single lineage Z_t in a spatially evolving population is driven by an SDE of the form $\mathrm{d}Z_t = \mathrm{d}B_t + \nabla \log N(Z_t)\,\mathrm{d}t$, where N is something akin to a population density (for example in [9] the population density N is a centred one-dimensional Gaussian and the drift is $\nabla \log N(x) \propto -x$). In our case, we have more than one lineage, and the spatial distribution of the population is random, fluctuates, and does not admit a density except in one dimension (see e.g. [16], Chapter 2.9). In that light it seems surprising that the same representation still holds, by replacing N with the joint density of an n-sample from the population's stationary distribution.

[3] More precisely, the drift acting on a lineage located at x when there is another lineage at location y points in the same direction as $\int_0^\infty \mathrm{e}^{-|\nu_2|t} \nabla_x p_t(x-y)\,\mathrm{d}t$. On $\mathbb{R}^d$, this is the same direction as $y - x$, but on the torus it is more complex because of the periodicity. The difference becomes negligible for small separation.

[4] As for the binary merger, this is technically only true if we were on $\mathbb{R}^d$, but the difference is negligible for small separations of the three lineages.

1.5 Associated Population Models

1.5.1 Ξ-Fleming-Viot Process

Let us briefly define what has been called the Ξ-*Fleming-Viot process* [6], a probability measure valued Markov process on E (or on all of $\mathbb{R}^d$) that generalises the "generalised (or Λ-)Fleming-Viot process" coined in [4]. The Ξ-Fleming-Viot process is dual to the non-spatial Ξ-coalescent [6], in the same sense that the classical Fleming-Viot process is dual to Kingman's coalescent [12].

Let $\Xi = a\delta_0 + \Xi_0$ be a finite measure on $\triangle$, where δ_0 is the unit mass at zero and Ξ_0 has no atom at zero. The Ξ-Fleming-Viot process is most easily defined using a lookdown construction [12, 13, 6]. Let $\mathfrak{N}^{ij}$ for $i, j \in \mathbb{N}, i < j$ be a family of independent, rate a Poisson processes on $\mathbb{R}$, and $\mathfrak{M}$ a Poisson point process on $\mathbb{R} \times \triangle$, independent of the $\mathfrak{N}^{ij}$, with intensity

$$\mathrm{d}t \otimes \frac{\Xi_0(\mathrm{d}\boldsymbol{\xi})}{\langle \boldsymbol{\xi}, \boldsymbol{\xi} \rangle}.$$

We define a dynamic on E^∞, started from an initial configuration $\boldsymbol{Y}(0) = (Y_1(0), Y_2(0), \ldots)$ as follows.

(i) At a point t of the process $\mathfrak{N}^{ij}$, the j'th level particle "looks down" to the i'th level particle and copies its location.

(ii) At a point $(t, \boldsymbol{\xi})$ of the process $\mathfrak{M}$, we independently assign every level to a basket with label $i \in \mathbb{N}$ with probability ξ_i, and to no basket with probability $1 - \sum_i \xi_i$. Then every particle that has been assigned to a basket copies the location of the particle with the smallest label in the same basket.

(iii) All levels follow independent Brownian motions between reproductive events.

There may be an infinite number of lookdown events in finite time, but any fixed level is only hit by non-trivial reproductive events at a finite rate.

Remark 1.19. If $\xi_2 = \xi_3 = \ldots = 0$ for Ξ-a.e. $\boldsymbol{\xi} \in \triangle$, then Ξ reduces to a measure on $[0, 1]$ which is commonly denoted $\Lambda = a\delta_0 + \Lambda_0$, where Λ_0 has no atom at zero. In terms of Λ, step (ii) above can be summarised as follows: At a point (t, p) of a Poisson point process with intensity $\mathrm{d}t \otimes \frac{\Lambda_0(\mathrm{d}p)}{p^2}$ on $\mathbb{R} \times [0, 1]$, select every level independently with probability p, and each selected particle copies the location of the selected particle with the lowest level. The resulting particle process is called the Λ-*Fleming-Viot process*.

The point in this construction lies in the fact that it preserves exchangeability: If $\boldsymbol{Y}(0) = (Y_1(0), Y_2(0), \ldots)$ is exchangeable, then so is $\boldsymbol{Y}(t) = (Y_1(t), Y_2(t), \ldots)$ for every $t > 0$ [12, 6]. We denote the associated de Finetti measure by

$$\mathcal{Y}_t = \lim_{n \to \infty} \frac{1}{n} \sum_{i=1}^{n} \delta_{Y_i(t)}, \quad t \geq 0.$$

The Ξ-*Fleming-Viot process* started at $(\boldsymbol{Y}(0), \mathcal{Y}_0)$ is a càdlàg modification of the process $(\boldsymbol{Y}(t), \mathcal{Y}_t)_{t\geq 0}$ (only the second component has to be modified). For details of the construction see [6]. (The term "Ξ-Fleming-Viot process" usually just refers to the process $(\mathcal{Y}_t)$, but we find it convenient to include the particle representation explicitly.) If $\Xi = \delta_0$, then $(\mathcal{Y}_t)$ is the well-known Fleming-Viot process (this was the main subject of [12]), which can, for example, be obtained from a Dawson-Watanabe superBrownian motion by conditioning it on having constant, unit population size.

Proposition 1.20. *The stationary distribution of the Ξ-Fleming-Viot process (exists and) is equal to μ^Ξ. It can be constructed to run bi-infinitely at stationarity, which we denote $(\boldsymbol{Y}(t), \mathcal{Y}_t)_{t\in(-\infty,\infty)}$.*

More precisely, at stationarity the distribution of $\mathcal{Y}_t$ is μ^Ξ, and conditional on $\mathcal{Y}_t$, the vector $\boldsymbol{Y}(t)$ is an i.i.d. sequence of samples from $\mathcal{Y}_t$. Note, in particular, that the joint law of the first n levels $(Y_1(t), \dots, Y_n(t))$ at stationarity is given by (1.7) for any $n \in \mathbb{N}$.

Remark 1.21. The stationary distribution μ^Ξ is not reversible in the sense that the process $(\boldsymbol{Y}(t), \mathcal{Y}_t)_{t\in(-\infty,\infty)}$ running at stationarity has the same law as $(\boldsymbol{Y}(-t), \mathcal{Y}_{-t})_{t\in(-\infty,\infty)}$. The Fleming-Viot process is only reversible in this sense if the underlying spatial motion (or, in the genetics setting, mutation process) is a pure jump process in which the destination of each jump is independent of the location just before the jump. See [19, 23, 28, 35, 43] for literature on this topic.

Remark 1.22. A byproduct of this result is explicit formulas for samples from the stationary distribution of a Ξ-Fleming-Viot process through (1.7), where N^Ξ can be calculated from (1.3). For instance, recalling Example 1.9, the density of a two-sample from the stationary distribution μ^Ξ of the Ξ-Fleming-Viot process is

$$\mathbb{E}\left[\mu^\Xi(\mathrm{d}x_1)\mu^\Xi(\mathrm{d}x_2)\right] \propto \left(\int_0^\infty \mathrm{e}^{-|\Xi|t} p_{2t}(x_1 - x_2)\,\mathrm{d}t\right)\mathrm{d}x_1\,\mathrm{d}x_2,$$

normalised such that the integral over x_1 and x_2 is one. Here we used that $\lambda_2 = \lambda_{2,2}$, which equals $|\Xi|$ by (1.1). In particular, the distribution of a two-sample from a stationary Ξ-Fleming-Viot process only depends on $|\Xi|$. Formulas for higher order samples involve sums over an increasing number of trees, and depend on higher order rates of the Ξ-coalescent.

The particle representation $(\boldsymbol{Y}(t))$ can be used to define the genealogies of the Ξ-Fleming-Viot process. Denote by $L_k(t)$ the level of the ancestor at time $-t$ of the particle at level k at time 0 (assigned at jump times in such a way that L_k is right-continuous). Let Π_t be the partition of $\mathbb{N}$ induced by the equivalence relation

$$i \sim_t j \iff L_i(t) = L_j(t).$$

Note that $L_i(t) = \min\{j\colon i \sim_t j\}$ by the definition of the lookdown construction, so (Π_t) contains all information about the processes $(L_i(t)\colon i \in [n])$, where $[n] = \{1, \dots, n\}$. Denote by

$$\boldsymbol{X}_t(u) = \lim_{s \uparrow -t} Y_{\min u}(s), \qquad t \geq 0, u \in \Pi_t,$$

the position at time $-t$ of the common ancestor of the particles with levels in u at time zero (assigned again in such a way that $\boldsymbol{X}$ is right-continuous).

For $n \in \mathbb{N}$ and $t \geq 0$, write Π_t^n for the partition induced by Π_t on $[n]$, and $\boldsymbol{X}_t^n$ for the restriction of $\boldsymbol{X}_t$ to Π_t^n. Then $(\Pi_t^n, \boldsymbol{X}_t^n)_{t \geq 0}$ is the coalescent process that describes the genealogies of a sample of n individuals from the Ξ-Fleming-Viot process at stationarity. The initial condition is random: $\Pi_0 = \{\{1\}, \dots, \{n\}\}$ by construction, and $\boldsymbol{X}_0(\{i\}) = x_i$ is the location of the i'th level particle at time zero, so $(x_1, \dots, x_n)$ is an i.i.d. sample from $\mathcal{Y}_0 \sim \mu^{\Xi}$.

Theorem 1.23. *The law of $(\Pi_t^n, \boldsymbol{X}_t^n)_{t \geq 0}$ is that of a Brownian spatial Ξ-coalescent started from the random initial condition $(\Pi_0^n, \boldsymbol{X}_0^n)$.*

In fact, the Brownian spatial Ξ-coalescent can be used to describe the full time-reversal of the Ξ-Fleming-Viot process.

Theorem 1.24. *For any $n \in \mathbb{N}$, the law of the time reversal $(Y_1(-t), \dots, Y_n(-t))_{t \geq 0}$ at stationarity can be described as follows:*

(i) Sample initial points $y_1, \dots, y_n$ from a random realisation of μ^{Ξ}.

(ii) Evolve according to a Brownian spatial Ξ-coalescent.

(iii) When a coalescence event occurs, resample points using the measures $(\mu^{\Xi}_{x_1,\dots,x_n})$ so that the total number of lineages stays constant.

A precise statement, including how levels are reassigned at coalescence–resampling events, can be found in Section 3.5.1. As an example for rule (iii), if a ternary merger occurred and the remaining lineages are at locations $y_1, \dots, y_{n-2}$, we sample y_{n-1} from $\mu^{\Xi}_{y_1,\dots,y_{n-2}}$, and y_n from $\mu^{\Xi}_{y_1,\dots,y_{n-2},y_{n-1}}$, then resume the Brownian spatial Ξ-coalescent from $y_1, \dots, y_n$.

1.5.2 Spatial Cannings Models

Next, we present a class of particle models whose forward evolution and genealogies converge in the limit of large population size, properly rescaled, to the Ξ-Fleming-Viot process and Brownian spatial Ξ-coalescent, respectively. We stress that the convergence of the forward evolution isn't novel, and is essentially contained in [13]. We consider a spatial version of the Cannings model. For $N \in \mathbb{N}$, let $(O_1, \dots, O_N)$ be an exchangeable $\mathbb{N}_0^N$ valued random variable with $\sum_{i=1}^N O_i = N$, and let $T_N > 0$. We define an E^N-valued process $\boldsymbol{Y}^N(t) = (Y_1^N(t), \dots, Y_N^N(t))$ as follows.

(i) At rate T_N, draw a random realisation $(o_1, \ldots, o_N)$ of $(O_1, \ldots, O_N)$, independent of previous realisations and current spatial locations of all particles. Partition $[N]$ into sets of sizes $o_1, \ldots, o_N$ uniformly at random. Then every individual copies the location of the individual with the smallest level in the same set.

(ii) Individuals follow independent Brownian motions between reproduction events.

Step (i) may look more complicated than it is. It is equivalent to saying "individual i has o_i offspring", followed by a particular way of reassigning levels in such a way that preserves exchangeability. The version of this process with no spatial motions (and generations at discrete time steps rather than the points of a Poisson process) was studied in the coalescent context by many authors, including Möhle and Sagitov [33]. Following their line of reasoning, the probability that a particular $(n, \vec{k})$-merger occurs in a single generation is given by

$$p^N_{n,\vec{k}} = \frac{(N)_{n'}}{(N)_n} \mathbb{E}\left[(O_1)_{k_1} \ldots (O_m)_{k_m} O_{m+1} \ldots O_{n'}\right],$$

where $n' = n - \sum_{i=1}^n k_i + m$ is the number of lineages after the merge event, and $(k)_a = k(k-1) \cdot \ldots \cdot (k-a+1)$ for $a, k \in \mathbb{N}$, $a \leq k$ (see [33], Eq. (3)). They give necessary and sufficient conditions under which $T_N p^N_{n,\vec{k}}$ converges for every $(n, \vec{k}) \in M$ as $N \to \infty$, in which case the limits are the rates $\lambda_{n,\vec{k}}$ of a non-spatial Ξ-coalescent, to which the genealogies converge. We essentially prove a spatial analogue. Denote by

$$\mathcal{Y}^N_t = \frac{1}{N} \sum_{i=1}^N \delta_{Y^N_i(t)}, \quad t \geq 0,$$

the empirical measure of $\boldsymbol{Y}^N(t)$.

Proposition 1.25. *The process $(\boldsymbol{Y}^N, \mathcal{Y}^N)$ has a stationary distribution, and can be constructed to run bi-infinitely at stationarity, which we denote $(\boldsymbol{Y}^N(t), \mathcal{Y}^N_t)_{t \in (-\infty, \infty)}$.*

We can define the genealogies of $\mathcal{Y}^N$ in the same way as we did previously: $L^N_k(t)$ for $t \geq 0$ and $1 \leq k \leq N$ denotes the level of the ancestor at time $-t$ of the particle at level k at time 0. Then Π^N_t denotes the partition on $[N]$ induced by the equivalence relation "$i \sim j \iff L^N_i(t) = L^N_j(t)$". Finally $\boldsymbol{X}^N_t(u) = \lim_{s \uparrow -t} Y^N_{\min u}(s)$ for $t \geq 0$ and $u \in \Pi^N_t$. Write $\Pi^{N,n}_t$ for the partition induced by Π^N_t on $[n]$, and $\boldsymbol{X}^{N,n}_t$ for the restriction of $\boldsymbol{X}^N_t$ to $\Pi^{N,n}_t$.

Theorem 1.26. *For every $n \in \{1, \ldots, N\}$, the distribution of $(\Pi^{N,n}_t, \boldsymbol{X}^{N,n}_t)_{t \geq 0}$, conditional on initial locations, is the Brownian spatial coalescent with transition measures*

$$\nu^N_{n,\vec{k}}(\mathrm{d}\boldsymbol{z}) = T_N p^N_{n,\vec{k}} \, \mathrm{d}\boldsymbol{z}.$$

Furthermore, the following are equivalent.

(i) The process $((Y_1^N(t), \dots Y_n^N(t)))_{t\geq 0}$ *converges weakly as* $N \to \infty$ *for every* $n \in \mathbb{N}$.

(ii) The process $(\Pi_t^{N,n}, \boldsymbol{X}_t^{N,n})_{t\geq 0}$ *converges weakly as* $N \to \infty$ *for every* $n \in \mathbb{N}$.

(iii) The process $(\Pi_t^{N,n})_{t\geq 0}$ *converges weakly as* $N \to \infty$ *for every* $n \in \mathbb{N}$,

(iv) The limit $T_N p_{n,\vec{k}}^N$ *as* $N \to \infty$ *exists for all* $(n, \vec{k}) \in M$.

In that case, there exists a unique finite measure Ξ *on* $\triangle$ *such that, for every* $n \in \mathbb{N}$, $(Y_1^N, \dots, Y_n^N)$ *converges to the first* n *levels* $(Y_1, \dots, Y_n)$ *of the* Ξ*-Fleming-Viot process, and* $(\Pi^{N,n}, \boldsymbol{X}^{N,n})$ *converges to the Brownian spatial* Ξ*-coalescent.*

Here by weak convergence in (i), (ii), and (iii) we mean convergence in law in the Skorokhod topology. We repeat that the equivalence of (iii) and (iv) is proved in [33].

Remark 1.27. (i) By arguments in [13], in the case of no simultaneous mergers where Ξ reduces to a measure Λ on $[0, 1]$, if (i) to (iv) hold, then also the weak convergence $(\boldsymbol{Y}^N, \mathcal{Y}^N) \implies (\boldsymbol{Y}, \mathcal{Y})$ to the Λ-Fleming-Viot process holds. We expect the corresponding result to be true in the general case of simultaneous mergers.

(ii) The proof that the genealogies of the spatial Cannings model are a Brownian spatial coalescent (the first statement in Theorem 1.26) is constructive, see Lemma 3.32. Studying this class of spatial population models could thus have been a way to define the Brownian spatial coalescent in the first place, but it would not have revealed its axiomatic universality.

(iii) Similar to the Ξ-Fleming-Viot process, the full time reversal at stationarity of the spatial Cannings model can be described by the Brownian spatial coalescent "with resampling", with the transition measures defined in Theorem 1.26. See Lemma 3.33.

1.6 Related Literature

The inclusion of space (or type) in population and coalescent models is the subject of a vast amount of literature. The classical approach is to consider populations that are subdivided into demes of large constant size, each situated at a vertex of a graph. Lineages can merge with other lineages that are currently in the same deme, and migrate between neighbouring demes. We refer to [1] for a short review of classical results in this setting, and point out [29, 45] and references therein for recent results. If the population is thought to be (discretely) structured by type rather than space, then associated coalescents are commonly called *multi-type coalescents*. Recent results using this terminology include [24, 30, 25].

In reality, populations are often not subdivided but spread across a spatial continuum. One might be tempted to approximate this through deme-based models with small

granularity, but since local population size can be small this would break the crucial assumption of large population size in every deme; see also [18], Section 5. A common approach that incorporates continuous space directly is to assume that genealogical trees can be constructed from independent Brownian motions (or more general spatial motions [20]) which coalesce when they meet, either instantaneously or at a rate proportional to their collision local time. This is inherently limited to one dimension, and a suggested extension to higher dimensions has lineages coalesce at rates depending on their separation. The position of the common ancestor is typically taken to be a Gaussian centred on the midpoint between the two lineages immediately before the coalescence event. Aside from the fact that this behaviour is biologically unnatural, there is no corresponding forwards in time model for the evolution of the population [1].

A major breakthrough in the theory of spatial population models was the introduction of the Spatial Λ-Fleming-Viot process (SLFV) by Barton, Etheridge and Véber in [18, 1], see also [2] for a review. It models the genetic composition of a spatially structured population by a measure on $\mathbb{R}^d \times [0,1]$, where $\mathbb{R}^d$ is the geographical space and $[0,1]$ is the space of genetic types. The population density is modelled to be constant so that the spatial marginal is always Lebesgue measure. The population evolves through a sequence of events, each of which replaces a certain proportion of the population in a randomly chosen ball by offspring of an individual chosen at random within the same ball. If the radius is small, the event is interpreted as an ordinary reproduction event subject to local regulation, and if it is large one may think of it as an extinction–recolonisation event. The genealogies of the SLFV (also called its dual) have a simple description: backwards in time, whenever a lineage is hit by an event, with a certain probability (equal to the proportion of the population that was hit forwards in time) it jumps to a randomly chosen location within the same ball, otherwise it stays put. If a number of lineages are affected by the same event, they jump to the same location and coalesce. The dual of the SLFV was the first tractable model for the ancestry of a population evolving in a two-dimensional spatial continuum that was sampling consistent, had a natural corresponding forward in time population model, and in which independently evolving lineages fail to meet. A variant of the SLFV allows for more than one parent per event, leading to the possibility of simultaneous mergers in the dual.

Despite its success, the SLFV has weaknesses. It only models populations with perfectly homogeneous spatial distribution, which in real populations fluctuates strongly in time and both geographical and type space [18], Chapter 7. Secondly, individual lineages in the SLFV evolve through a sequence of correlated jumps, but in many applications it would be more natural if they followed Brownian motions, or the paths of other continuous Markov processes. The Ξ-Fleming-Viot processes are a class of spatial population model that showcase both characteristics.

A different approach to spatial coalescents is to replace geographical space with a tree-like structure, such as the hierarchical group which is often used to mimic higher-dimensional spaces. A spatial variant of the Λ-coalescent on the hierarchical group is the subject of [22]. In [21], Freeman also works on a space with a hierarchical structure, on which he considers a coalescent model called the segregated Λ-coalescent.

1.7 Conclusion and Outlook

We discuss some consequences of our results, and possible further research directions.

1.7.1 Spatially Interactive Branching

The Brownian spatial coalescent was axiomatically defined through a property that the genealogies of any spatial population model in which individuals follow independent Brownian motions should satisfy. This includes models with spatially interactive branching mechanisms such as competition or local regulation. Yet Section 1.5.2 shows that the full set of sampling consistent Brownian spatial coalescents is exhausted by spatial models in which the branching mechanism is *oblivious* to spatial locations. The only other assumptions made in the definition of Brownian spatial coalescents are time-homogeneity and the Markov property. If the spatial model has a stationary distribution, which will certainly be the case for some interesting spatially interactive branching mechanisms, then their genealogies will be time homogeneous, which means they *cannot be Markovian.* Thus, describing the genealogies of populations whose branching mechanisms have interesting spatial interactions, at least in the present setting, requires the study of non-Markovian coalescent processes. Since we expect most of our results to remain true for much more general spatial motions (see next paragraph), this observation is not limited to population models with Brownian movement.

1.7.2 General Spatial Motions

The definition of the Brownian spatial coalescent and most theorems and proofs (with the exception of the drift representation Theorem 1.15) work with little modification when Brownian motion on the torus is replaced by another well-behaved Markov process (Y_t) on some compact Polish space E with a stationary distribution λ and continuous transition densities (q_t) w.r.t. λ, such that its time reversal w.r.t. λ is also a well-behaved Markov process. Then (q_t) takes the place of (p_t), and λ takes the place of Lebesgue measure in the main theorems.

Example 1.28. Examples include random walks, α-stable motion or diffusions on fractals. One could also take a cartesian product of a geographical and a type space, for example $E = \mathbb{T}^d \times \{1, \dots, k\}$, where the torus is interpreted as geographical space, and

$\{1, \dots, k\}$ is type space. The motion could be a random walk on type space, and Brownian motion in geographical space, with a speed that may depend on the type.

The only major difficulty is understanding which points have to be excluded from the state space. One approach could be to exclude those points where the normalisation $N^{\boldsymbol{\nu}}$ diverges, but then it remains to show that these points will almost-surely never be reached by the evolution of the coalescent. This comes down to characterising whether or not k independent copies of the spatial motion have a positive probability of meeting simultaneously, and if so in which points. In fact, by studying the normalisation of the "(Y_t)-spatial δ_1-coalescent", in which only (n, n)-mergers for $n \in \mathbb{N}$ are possible, this gives rise to the following conjecture.

Conjecture 1.29. *Under some regularity conditions on (Y_t), the following are equivalent for every $k \in \mathbb{N}$, $k \geq 2$, and $x_1, \dots, x_k \in E$.*

(i) If k independent copies of Y are started from initial locations x_1 to x_k, then with positive probability there exists a time at which all k copies are in the same location.

(ii)
$$\int_0^1 \int_E \prod_{i=1}^{k} q_s(x_i, z)^k \lambda(\mathrm{d}z)\, \mathrm{d}s < \infty.$$

A simple case is if the motion is a random walk on some finite graph E. Then no points have to be excluded, and all of our results (except Theorem 1.15) apply with obvious modifications.

1.7.3 The Brownian Spatial Coalescent on $\mathbb{R}^d$

We have some unpublished results about Brownian spatial coalescents evolving on all of $\mathbb{R}^d$, which can be defined analogously to Definition 1.5. Corresponding to each is again a family of transition measures $\boldsymbol{\nu} = (\nu_{n,\vec{k}})$, where $\nu_{n,k_1,\dots,k_m}$ is a locally finite measure on $(\mathbb{R}^d)^m$. If $\boldsymbol{\lambda} = (\lambda_{n,\vec{k}})$ are the rates of a Ξ-coalescent, then $\nu_{n,\vec{k}}(\mathrm{d}\boldsymbol{z}) = \lambda_{n,\vec{k}}\, \mathrm{d}\boldsymbol{z}$ is again a valid choice, and the corresponding Brownian spatial coalescent describes the genealogies of a Ξ-Fleming-Viot process evolving on $\mathbb{R}^d$, at "stationarity modulo translation" (cmp. [12], Theorem 2.9).

If $d \geq 3$, then the class of Brownian spatial coalescents that describe the genealogies of some population is larger than that: for each Brownian spatial coalescent on $\mathbb{R}^d$ that we described in the previous paragraph, there exists another whose definition is the same except that all exponential factors in the densities on tree decorations (of the form $\mathrm{e}^{-|\nu_n|(\tau-\tau')}$) are removed. If $\nu_{n,2}(\mathrm{d}z) = \mathrm{d}z$ for all $n \in \mathbb{N}$, and all other transition measures are equal to zero, then this coalescent describes the genealogies of the infinite-mass superBrownian motion on $\mathbb{R}^d$ at stationarity (which exists iff $d \geq 3$ [16], Chapter

2.7). It has the property that, from any initial configuration with at least two particles, the probability that no further merge event happens is strictly between zero and one. The existence of coalescents with this property on $\mathbb{R}^d$ is directly linked to the existence of infinite-mass stationary distributions for associated Dawson-Watanabe superprocesses on $\mathbb{R}^d$, which is known to correspond to transience of the underlying spatial motion [17, 15]. This connection can also be seen from the perspective of the Brownian spatial coalescent: Brownian motion is transient if and only if $\int_C p_t(x)\,\mathrm{d}x$ is integrable over $t \in (0,\infty)$ for every compact $C \subset \mathbb{R}^d$ (this characterisation holds mutatis mutandis for more general spatial motions), which is certainly necessary for the normalisation $N^{\boldsymbol{\nu}}(\boldsymbol{x})$ to be finite, since integrability at large τ is no longer guaranteed by the exponential factors.

Finally, it is not clear how to define sampling consistency on $\mathbb{R}^d$ in such a way that it captures the coalescents associated with infinite-mass populations, because their stationary distributions are random measures of infinite mass that cannot be normalised to be probability measures.

1.8 Structure

In Section 2, we set up the notation necessary to formalise some of the definitions and statements from the introduction, and introduce the framework used for the proofs, which are in Section 3. Some technical proofs are postponed to the Appendix, most notably the proof of Lemma 2.12, which states that the function $N^{\boldsymbol{\nu}}$ is continuous under certain conditions; it is conceptually straight-forward but long and technical.

2 Setup

2.1 Non-spatial Coalescents

Recall that $\mathcal{P}$ denotes the (countable) set of partitions of finite subsets of $\mathbb{N}$, which we equip with the discrete topology. For $\Pi, \Pi' \in \mathcal{P}$ write $\Pi \leq \Pi'$ if Π' is a refinement of Π, and define

$$\Omega_0 \coloneqq \{\Pi \in R([0,\infty), \mathcal{P})\colon \Pi_s \geq \Pi_t \forall 0 \leq s \leq t\}\,,$$

where $R([0,\infty),\mathcal{P})$ is the space of right-continuous $\mathcal{P}$-valued paths. (Due to the monotonicity condition, every path in Ω_0 must in fact also have left limits.) We equip Ω_0 with its Borel σ-algebra $\mathcal{F}^{\Pi} \coloneqq \sigma(\Pi_t\colon t \geq 0)$, and filtration $\mathcal{F}^{\Pi}_t \coloneqq \bigcap_{s>t} \sigma(\Pi_r\colon r \leq s)$ for $t \geq 0$, where $(\Pi_t)_{t\geq 0} \coloneqq \mathrm{id}_{\Omega_0}$ denotes the canonical process in Ω_0. We further equip Ω_0 with the product topology, which is the smallest topology on Ω with respect to which $\boldsymbol{X}_t\colon \Omega \to \mathcal{X}$ is continuous for all $t \geq 0$.

Definition 2.1. A *(non-spatial) coalescent* is a time-homogeneous Ω_0-valued Markov process $(\mathbb{P}^{\Pi})_{\Pi\in\mathcal{P}}$. It is called *label invariant* if it is independent of particle labels in the

sense that any transition that is the result of merging m disjoint sets of lineages of sizes $k_1 \geq \ldots \geq k_m$, while at n particles, occurs at the same rate $\lambda_{n,k_1,\ldots,k_m}$.

In general, a coalescent is characterised by its transition rates $\lambda_{\Pi,\Pi'} \geq 0$ for $\Pi > \Pi'$, which induces a bijection between the set of coalescents and the set R° of non-negative arrays $\boldsymbol{\lambda} = (\lambda_{\Pi,\Pi'} \colon \Pi > \Pi')$. For $\Pi \in \mathcal{P}$ and $u \subset \mathbb{N}$, write $\Pi \setminus\setminus u \coloneqq \{v \setminus u \colon v \in \Pi, v \neq u\}$ for the partition induced by Π on $\mathbb{N} \setminus u$.

Definition 2.2. A coalescent is called *sampling consistent* if, for any $\Pi \in \mathcal{P}$ with $|\Pi| \geq 2$ and $u \in \Pi$,

$$\mathbb{P}^{\Pi}((\Pi_t \setminus\setminus u)_{t\geq 0} \in \cdot) = \mathbb{P}^{\Pi \setminus\setminus u}((\Pi_t)_{t\geq 0} \in \cdot).$$

2.2 Spatial Coalescents

Recall that $E = \mathbb{R}^d/\mathbb{Z}^d$ denotes the d-dimensional square flat torus with periodic boundary conditions. It is a Polish space with the Euclidean metric

$$\rho(x + \mathbb{Z}^d, y + \mathbb{Z}^d) = \inf\{|x - y + \vec{k}| \colon \vec{k} \in \mathbb{Z}^d\}, \quad x, y \in \mathbb{R}^d. \tag{2.1}$$

Recall the definitions of E^Π and $E^\Pi_\circ$ for $\Pi \in \mathcal{P}$ (see (1.2)), and $\mathcal{X}$. The identification of $E^\Pi_\circ$ with an open subset of $E^{|\Pi|}$ defines on it a topology and Lebesgue measure, and then $\mathcal{X}$ is a Polish space with the topology inherited from $E^\Pi_\circ$, $\Pi \in \mathcal{P}$, in which $(\Pi_n, \boldsymbol{x}_n) \to (\Pi, \boldsymbol{x})$ iff $\Pi_n = \Pi$ for all but finitely many n, and $\boldsymbol{x}_n \to \boldsymbol{x}$. Since $\Pi = \mathrm{dom}(\boldsymbol{x})$ is the domain of $\boldsymbol{x}$ for $(\Pi, \boldsymbol{x}) \in \mathcal{X}$, we are justified to occasionally just write $\boldsymbol{x} \in \mathcal{X}$. Denote

$$\Omega \coloneqq \{(\Pi_t, \boldsymbol{X}_t) \in R([0,\infty), \mathcal{X}) \colon \Pi_t \leq \Pi_s \forall t \geq s \geq 0\},$$

where $R([0,\infty), \mathcal{X})$ denotes the space of right-continuous functions $[0,\infty) \to \mathcal{X}$. We equip Ω with the product topology as well as its Borel σ-algebra $\mathcal{F}^{\boldsymbol{X}} \coloneqq \sigma(\boldsymbol{X}_t \colon t \geq 0)$ and the filtration $\mathcal{F}^{\boldsymbol{X}}_t \coloneqq \bigcap_{s>t} \sigma(\boldsymbol{X}_r \colon r \leq s)$ for $t \geq 0$, where $(\boldsymbol{X}_t)_{t\geq 0} \coloneqq (\Pi_t, \boldsymbol{X}_t)_{t\geq 0} \coloneqq \mathrm{id}_\Omega$ denotes the canonical process in Ω. Write $C_b(\mathcal{X})$ for the space of bounded and continuous functions $\mathcal{X} \to \mathbb{R}$.

Definition 2.3. A *spatial coalescent* is a time-homogeneous Ω-valued Markov process $(\mathbb{P}^{\boldsymbol{x}})_{\boldsymbol{x}\in\mathcal{X}}$ whose *semigroup* $(S_t)_{t\geq 0}$, defined by

$$S_t(\boldsymbol{x}, A) \coloneqq \mathbb{P}^{\boldsymbol{x}}(\boldsymbol{X}_t \in A), \tag{2.2}$$

for $\boldsymbol{x} \in \mathcal{X}$, $t \geq 0$, and $A \subset \mathcal{X}$ measurable, has the property that $S_t f \in C_b(\mathcal{X})$ for all $f \in C_b(\mathcal{X})$, which we refer to as the *Feller property*.

The Feller property—which is automatic in the non-spatial setting—is a form of continuity in the initial condition that, together with right-continuity of sample paths, implies the strong Markov property [37, Thm. 8.3]: if T is an $(\mathcal{F}_t^{\boldsymbol{X}})$-stopping time, then

$$\mathbb{P}^{\boldsymbol{x}}\left(T<\infty, (\boldsymbol{X}_{T+t})_{t\geq 0}\in\cdot\,\middle|\,\mathcal{F}_T^{\boldsymbol{X}}\right)=\mathbb{1}_{\{T<\infty\}}\mathbb{P}^{\boldsymbol{X}_T}((\boldsymbol{X}_t)_{t\geq 0}\in\cdot),\qquad \mathbb{P}^{\boldsymbol{x}}\text{-a.s.}$$

for every $\boldsymbol{x}\in\mathcal{X}$.

The formal definition of label invariance in the spatial setting is a bit more technical than in the non-spatial setting, because a convenient formulation in terms of transition rates is not available. (Within the class of Brownian spatial coalescents, Lemma 3.13 provides an analogous characterisation.) Suppose that $\Pi_0,\Pi_1\in\mathcal{P}$ with $|\Pi_0|=|\Pi_1|$. Then the law of a label invariant spatial coalescent should not be affected if we change the labels of the initial set of lineages with a bijection $\iota\colon\Pi_0\to\Pi_1$. To apply this change of labels to the entire process $(\Pi_t)_{t\geq 0}$, we extend ι to a map $\bigcup_{\Pi\leq\Pi_0}\Pi\to\bigcup_{\Pi\leq\Pi_1}\Pi$ with the property that $\iota(u\cup v)=\iota(u)\cup\iota(v)$ whenever $u,v,u\cup v$ are in the extended domain of ι (this extension is unique). To perform the change of labels in the spatial variables, we extend ι for every $\Pi\subset\bigcup_{\Pi'\leq\Pi_0}\Pi'$ to a map $E_\circ^{\Pi}\to E_\circ^{\iota(\Pi)}$ through $\iota(\boldsymbol{x})=\boldsymbol{x}\circ\iota^{-1}$. Denote by $\mathcal{I}(\Pi_0,\Pi_1)$ the set of bijections $\iota\colon\Pi_0\to\Pi_1$ that are extended in the way described above.

Definition 2.4. A spatial coalescent $(\mathbb{P}^{\boldsymbol{x}})_{\boldsymbol{x}\in\mathcal{X}}$ is called *label invariant* if for any two sets $\Pi_0,\Pi_1\in\mathcal{P}$ of equal size, $\iota\in\mathcal{I}(\Pi_0,\Pi_1)$, and any $\boldsymbol{x}\in E_\circ^{\Pi_0}$,

$$\mathbb{P}^{\iota(\boldsymbol{x})}((\boldsymbol{X}_t)_{t\geq 0}\in\cdot)=\mathbb{P}^{\boldsymbol{x}}((\iota(\boldsymbol{X}_t))_{t\geq 0}\in\cdot).\tag{2.3}$$

We now make precise the notion of sampling consistency for spatial coalescents introduced informally in Definition 1.11. If $\boldsymbol{x}\in E^A$ and $\boldsymbol{y}\in E^B$ for disjoint sets A and B, we can join the two maps and denote the result by $\boldsymbol{xy}\in E^{A\cup B}$, so

$$\boldsymbol{xy}\colon A\cup B\to E;\, u\mapsto\begin{cases}\boldsymbol{x}(u), & u\in A,\\ \boldsymbol{y}(u), & u\in B.\end{cases}\tag{2.4}$$

For $(\Pi,\boldsymbol{x})\in\mathcal{P}$, $u\subset\mathbb{N}$ with $u\cap\Pi=\emptyset$ and $y\in E$, write $\boldsymbol{x}y\colon\Pi\cup\{u\}\to E$ for the extension of $\boldsymbol{x}$ to $\Pi\cup\{u\}$ with $u\mapsto y$, where u is suppressed in the notation and will be clear from context. Then for $(\Pi,\boldsymbol{x})\in\mathcal{X}$ and any such $u\subset\mathbb{N}$, and $\Pi'\in\mathcal{P}$ disjoint from Π, write

$$\begin{aligned}E_{\boldsymbol{x}}&\coloneqq\{y\in E\colon \boldsymbol{x}y\in E_\circ^{\Pi\cup\{u\}}\},\\ E_{\boldsymbol{x}}^{\Pi'}&\coloneqq\{\boldsymbol{y}\in E_\circ^{\Pi'}\colon \boldsymbol{xy}\in E_\circ^{\Pi\cup\Pi'}\}.\end{aligned}\tag{2.5}$$

(The first definition does not depend on the choice of u.) For $(\Pi,\boldsymbol{x})\in\mathcal{X}$ and any $u\in\Pi$, define $(\boldsymbol{x}\setminus\!\setminus u)\colon(\Pi\setminus\!\setminus u)\to E$ by $v\setminus u\mapsto\boldsymbol{x}(v)$. Write $\mathcal{M}_1(S)$ for the set of probability measures on a measurable space S.

Definition 2.5. A spatial coalescent $(\mathbb{P}^{\boldsymbol{x}})_{\boldsymbol{x}\in\mathcal{X}}$ is called *sampling consistent* if there exists a family of probability measures $(\mu_{\boldsymbol{x}} \in \mathcal{M}_1(E_{\boldsymbol{x}})\colon \boldsymbol{x} \in \mathcal{X})$ such that for any $\Pi \in \mathcal{P}$, $|\Pi| \geq 2$, and $u \in \Pi$, $\boldsymbol{x} \in E_\circ^{\Pi\setminus\setminus u}$,

$$\int_{E_{\boldsymbol{x}}} \mathbb{P}^{\boldsymbol{x}y}((\boldsymbol{X}_t \setminus\setminus u)_{t\geq 0} \in \cdot)\mu_{\boldsymbol{x}}(\mathrm{d}y) = \mathbb{P}^{\boldsymbol{x}}((\boldsymbol{X}_t)_{t\geq 0} \in \cdot). \tag{2.6}$$

With these definitions, the projection of a (label invariant or sampling consistent) spatial coalescent onto $\mathcal{P}$ is a (label invariant or sampling consistent) coalescent.

2.3 Brownian Spatial Coalescents

The law of a Brownian spatial coalescent will be defined using densities over time and space decorations of all possible shapes of the genealogical forest, as outlined in the introduction. We introduce the necessary notation.

Definition 2.6. A *forest* is a set whose members form a strictly decreasing (necessarily finite) sequence in $\mathcal{P}$.

The notation for a generic forest is $F = \{\Pi_0^F, \ldots, \Pi_m^F\}$, where $\mathrm{lf}(F) \coloneqq \Pi_0^F$ is the set of leaves, Π_i^F for $i \in [m]$ is the set of nodes (or vertices) immediately after the ith merge event, and $\mathrm{rt}(F) \coloneqq \Pi_m^F$ is the set of roots. Note that this notion of a forest encodes its topology *and* the order and possibly simultaneousness of merge events, but no branch lengths or spatial information. This is important because already in the non-spatial setting, forests with the same topology but different order of merge events may have different probability. Write $\mathbb{F}$ for the set of all forests, and for $\Pi, \Pi' \in \mathcal{P}$ with $\Pi' \leq \Pi$,

$$\begin{aligned}\mathbb{F}(\Pi) &= \{F \in \mathbb{F}\colon \mathrm{lf}(F) = \Pi\}\,,\\ \mathbb{F}(\Pi, \Pi') &= \{F \in \mathbb{F}\colon \mathrm{lf}(F) = \Pi, \mathrm{rt}(F) = \Pi'\}\,.\end{aligned}$$

We call $F \in \mathbb{F}$ a *tree* if $|\mathrm{rt}(F)| = 1$, and denote by $\mathbb{T}$ and $\mathbb{T}(\Pi)$ for $\Pi \in \mathcal{P}$ the set of all trees, and the set of trees with leaves Π, respectively. Write $\mathrm{nd}(F) \coloneqq \bigcup_{i=0}^m \Pi_i^F$ for the set of all nodes, and $\mathrm{nd}^\circ(F) \coloneqq \mathrm{nd}(F) \setminus \mathrm{lf}(F)$ for the non-leaf nodes of F. Call F *trivial* if $m = 0$, that is if $\mathrm{rt}(F) = \mathrm{lf}(F)$. The maps

$$\mathrm{ch}_F\colon \mathrm{nd}(F) \to 2^{\mathrm{nd}(F)}, \qquad \mathrm{pr}_F\colon \mathrm{nd}(F) \to \mathrm{nd}(F) \cup \{\emptyset\}$$

assign to a node its children and its parent, respectively, where $\mathrm{ch}_F(u) \coloneqq \emptyset$ for $u \in \mathrm{lf}(F)$ and $\mathrm{pr}_F(u) \coloneqq \emptyset$ for $u \in \mathrm{rt}(F)$. Figure 2.1 illustrates this notation.

Definition 2.7. A *time decoration* of a forest F is a map $\tau\colon F \to (0, \infty)$ with $\tau(\mathrm{lf}(F)) = 0$ and $\tau(\Pi) < \tau(\Pi')$ for $\Pi > \Pi'$. Write $\mathrm{tm}(F)$ for the set of time decorations of F, which

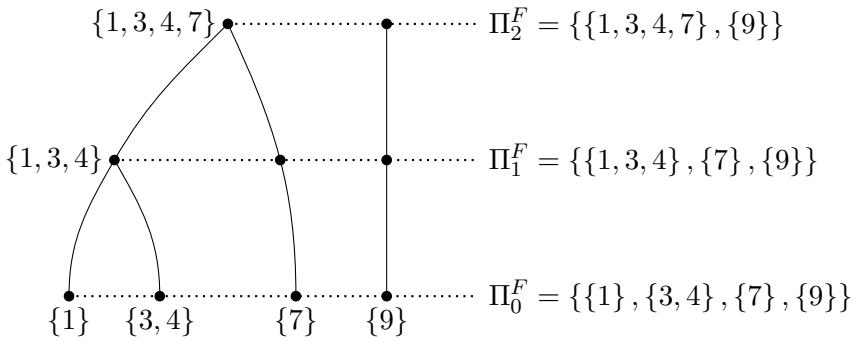

Figure 2.1: Illustration of the notation used for forests. In this example, $\mathrm{ch}_F(\{1,3,4\}) = \{\{1\},\{3,4\}\}$, and $\mathrm{pr}_F(\{7\}) = \{1,3,4,7\}$, and $\mathrm{ch}_F(\{9\}) = \mathrm{pr}_F(\{9\}) = \emptyset$.

we equip with the *upper limit topology*, that is $\tau^n \to \tau$ in $\mathrm{Tm}(F)$ if for every $\Pi \in F$, $\tau^n(\Pi) \to \tau(F)$ *and* there is $n_0 \in \mathbb{N}$ such that $\tau^n(\Pi) \leq \tau(\Pi)$ for all $n \geq n_0$. Write

$$\mathrm{tm}(\mathbb{F}) = \{(F,\tau)\colon F \in \mathbb{F}, \tau \in \mathrm{tm}(F)\}$$

for the set of time decorated forests, and define $\mathrm{tm}(S)$ for $S \subset \mathbb{F}$ analogously.

Since $F = \mathrm{dom}(\tau)$ for $(F,\tau) \in \mathrm{tm}(\mathbb{F})$ we occasionally just write $\tau \in \mathrm{tm}(\mathbb{F})$. If $F = \{\Pi_0^F, \ldots, \Pi_m^F\} \in \mathbb{F}$ and $\tau \in \mathrm{tm}(F)$, then for $i \in [m-1]$ we write $\tau(\Pi_i^F)^+ \coloneqq \tau(\Pi_{i+1}^F)$, and $\tau(\Pi_m^F)^+ \coloneqq \infty$. We extend τ to $\mathrm{nd}(F)$ by writing

$$\tau_u \coloneqq \min\{\tau(\Pi)\colon u \in \Pi \in F\}, \quad u \in \mathrm{nd}(F),$$

for the time of a nodes "birth", and we set $\tau_\emptyset = \infty$ so that $\tau_{\mathrm{pr}_F(u)} = \infty$ if u has no parent. See Fig. 2.2 for an illustration.

If F is trivial, then $\mathrm{tm}(F)$ is a singleton whose only element is defined by $\tau_{\mathrm{lf}(F)} = 0$, and we equip $\mathrm{tm}(F)$ with the measure that assigns mass one to the single member. Otherwise, since the Borel σ-algebra of the upper limit topology coincides with that of the usual Euclidean topology, we can define Lebesgue measure on $\mathrm{tm}(F)$ using the obvious bijection with an open subset of $\mathbb{R}^m$. In both cases we denote the integral of a function f against the measure on $\mathrm{tm}(F)$ by $\int f(\tau)\,\mathrm{d}\tau$. The set $\mathrm{tm}(\mathbb{F})$ inherits its topology from $\mathrm{tm}(F)$, noting it can be identified with a discrete union.

Definition 2.8. A *spatial decoration* of a forest F is a map $\xi\colon \mathrm{nd}^\circ(F) \to E$ such that $\xi\big|_\Pi \in E_\circ^\Pi$ for all $\Pi \in F$. Write $\mathrm{sp}(F)$ for the set of spatial decorations of F.

If F is non-trivial, we define a topology and Lebesgue measure on $\mathrm{sp}(F)$ using the obvious bijection with an open subset of $E^{|\mathrm{nd}^\circ(F)|}$. If F is trivial then $\mathrm{sp}(F)$ is a singleton comprised of the unique map $\xi\colon \emptyset \to E$, and we equip $\mathrm{sp}(F)$ with the measure that assigns mass one to the single member. In both cases write $\int f(\xi)\,\mathrm{d}\xi$ for the integral of a function f against the measure on $\mathrm{sp}(F)$.

The following maps assign to a path in Ω_0 its associated undecorated and time decorated coalescence forest, respectively.

$$\begin{aligned} \mathrm{Fr}\colon \Omega_0 \to \mathbb{F}; &\qquad (\Pi_t) \mapsto \{\Pi_t \colon t \geq 0\} \quad \text{(the range of the path } (\Pi_t)\text{)} \\ \mathrm{Tm}\colon \Omega_0 \to \mathrm{tm}(\mathbb{F}); &\quad \omega = (\Pi_t) \mapsto \Big(\mathrm{Fr}(\omega), \big[\Pi \mapsto \inf\{t > 0 \colon \Pi_t = \Pi\}\big] \Big) \end{aligned}$$

Both maps are measurable with respect to the product σ-algebra on Ω_0, see Lemmas II.A.1 and II.A.3. We can think of a (non-spatial) coalescent as a random, time decorated forest, and indeed Tm is a bijective and in fact bimeasurable map. In particular, the laws $\mathbb{P}^\Pi$ of a coalescent are determined by the pushforward laws $\mathrm{Tm}\,\#\mathbb{P}^\Pi$ on the set of time-decorated forests, which can be calculated explicitly using simple arguments of competing exponential clocks, see Lemma 3.16 below. We identify Fr and Tm with their extensions to Ω (obtained by composing with the projection $\Omega \to \Omega_0$).

In the spatial setting, we can assign to a path in Ω a forest decorated with space in addition to time coordinates. For $\omega = (\boldsymbol{x}_t) \in \Omega$ with $F = \mathrm{Fr}(\omega)$, we write

$$\mathrm{Sp}(\omega) \coloneqq \big[u \mapsto \boldsymbol{x}_{\mathrm{Tm}(\omega)_u}(u)\big] \in \mathrm{sp}(F)$$

for the associated space decoration of F. We prove that this map is measurable with respect to the product σ-algebra on Ω in Lemma II.A.4. Let $\mathrm{dec}(F) = \mathrm{tm}(F) \times \mathrm{sp}(F)$ for $F \in \mathbb{F}$, and write

$$\mathrm{dec}(\mathbb{F}) = \{(F, \tau, \xi) \colon F \in \mathbb{F}, \tau \in \mathrm{tm}(F), \xi \in \mathrm{sp}(F)\}$$

for the set of time and space decorated forests, and define $\mathrm{dec}(S)$ for $S \subset \mathbb{F}$ analogously. We write $F^\star = (F, \tau, \xi)$ for a generic element of $\mathrm{dec}(\mathbb{F})$. The topology on $\mathrm{dec}(\mathbb{F})$ is inherited from that of $\mathrm{dec}(F)$, $F \in \mathbb{F}$, noting it can be identified with a discrete union.

Since $F = \mathrm{dom}(\tau)$ for $(F, \tau, \xi) \in \mathrm{dec}(\mathbb{F})$ we occasionally just write $(\tau, \xi) \in \mathrm{dec}(\mathbb{F})$. The following map assigns to a path in Ω its associated decorated forest, see Fig. 2.2 for an illustration.

$$\mathrm{Dec}\colon \Omega \to \mathrm{dec}(\mathbb{F}); \quad \omega = (\boldsymbol{x}_t) \mapsto \Big(\mathrm{Tm}(\omega), \big[u \mapsto \boldsymbol{x}_{\mathrm{Tm}(\omega)_u}(u)\big] \Big). \tag{2.7}$$

It encodes all information about the path ω except the spatial motion of lineages in between merge events. Brownian spatial coalescents are exactly those for which this motion is Brownian conditional on the decorated forest, which means that its laws $\mathbb{P}^{\boldsymbol{x}}$ are determined completely by the laws $\mathrm{Dec}\,\#\mathbb{P}^{\boldsymbol{x}}$ on the space of decorated coalescence forests. To make this precise, we introduce a stochastic kernel $K_{\boldsymbol{x}}(F^\star, \cdot) \in \mathcal{M}_1(\Omega)$ describing the law of a Brownian spatial coalescent conditional on its decorated forest $F^\star \in \mathrm{dec}(\mathbb{F})$ and initial condition $\boldsymbol{x} \in E_\circ^{\mathrm{lf}(F)}$ (see Lemma 2.9). Then a spatial coalescent is a Brownian spatial coalescent if and only if $K_{\boldsymbol{x}}$ is a conditional probability of $\mathbb{P}^{\boldsymbol{x}}$ given Dec, that is if $\mathbb{P}^{\boldsymbol{x}} = (\mathrm{Dec}\,\#\mathbb{P}^{\boldsymbol{x}}) \otimes K_{\boldsymbol{x}}$ for all $\boldsymbol{x} \in \mathcal{X}$. See Definition 2.10 below. Measurability of Dec is proved in Lemma II.A.5.

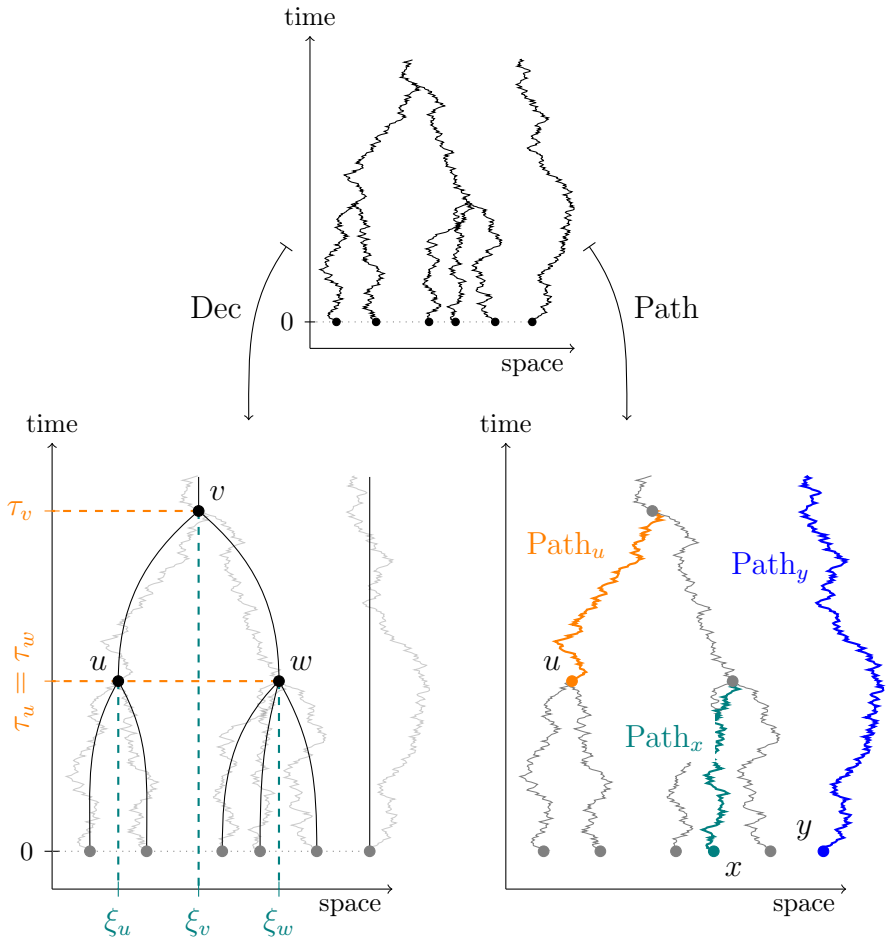

Figure 2.2: At the top is an illustration of an element $\omega \in \Omega$. The map Dec (bottom left) extracts the abstract coalescence forest, and times and spatial locations of the merge events. The Path maps (bottom right) extract the motion of particles along branches of the coalescence forest.

The following maps capture the motion of particles along each branch of the coalescence forest:

$$\begin{aligned} \text{Path}_u^{F^\star} \colon \{\text{Dec} = F^\star\} &\to R([\tau_u, \tau_{\text{pr}_F(u)}), E), \\ (\boldsymbol{x}_t) &\mapsto (\boldsymbol{x}_t(u))_{\tau_u \le t < \tau_{\text{pr}_F(u)}}, \end{aligned}$$

for $F^\star = (F, \tau, \xi) \in \text{Dec}(\mathbb{F})$ and $u \in \text{nd}(F)$, where $\{\text{Dec} = F^\star\} = \{\omega \in \Omega \colon \text{Dec}(\omega) = F^\star\}$. See again Fig. 2.2. These maps are continuous (if the codomain is equipped with the product topology, like the domain). Denote the law of a standard Brownian motion in E (with periodic boundary conditions) started at some x at time $t \ge 0$ by $B^{(t,x)+}$, the law of a Brownian bridge started at time $s \ge 0$ at $x \in E$, ending at time $t > s$ at $y \in E$ by $B^{(s,x)\to(t,y)}$, the law of the same bridge followed by a Brownian motion started at time t at y by $B^{(s,x)\to(t,y)+}$ etc. Recall our notation for joining maps: if $\boldsymbol{x} \in E_\circ^{\text{lf}(F)}$ and $\xi \in \text{sp}(F)$ then $\boldsymbol{x}\xi$ is the map $\text{nd}(F) \to E$ that extends ξ to $\text{lf}(F)$ using $\boldsymbol{x}$.

Lemma 2.9. *Given $F^\star = (F, \tau, \xi) \in \mathrm{Dec}(\mathbb{F})$ and $\boldsymbol{x} \in E_\circ^{\mathrm{lf}(F)}$, there is a unique law $K_{\boldsymbol{x}}(F^\star, \cdot) \in \mathcal{M}_1(\Omega)$ under which $\boldsymbol{X}_0 = \boldsymbol{x}$ and $\mathrm{Dec} = F^\star$ a.s., and*

$$\big(\mathrm{Path}_u^{F^\star}(\boldsymbol{X})\colon u \in \mathrm{nd}(F)\big)$$

is a family of independent random variables such that the law of $\mathrm{Path}_u^{F^\star}(\boldsymbol{X})$ is $B^{(\tau_u, (\boldsymbol{x}\xi)_u)+}$ if $\mathrm{pr}_F(u) = \emptyset$, and $B^{(\tau_u, (\boldsymbol{x}\xi)_u) \to (\tau_{\mathrm{pr}_F(u)}, \xi_{\mathrm{pr}_F(u)})}$ if $\mathrm{pr}_F(u) \neq \emptyset$. There exists an extension of these laws to a family $(K_{\boldsymbol{x}}(F^\star, \cdot)\colon \boldsymbol{x} \in \mathcal{X}, F^\star \in \mathrm{dec}(\mathbb{F}))$ such that

$$\mathcal{X} \times \mathrm{dec}(\mathbb{F}) \to \mathcal{M}_1(\Omega); \quad (\boldsymbol{x}, F^\star) \mapsto K_{\boldsymbol{x}}(F^\star, \cdot)$$

is continuous. In particular, $(\boldsymbol{x}, F^\star) \mapsto K_{\boldsymbol{x}}(F^\star, A)$ is measurable for any measurable $A \subset \Omega$.

The proof of Lemma 2.9 is in Appendix II.D. The exact definition of $K_{\boldsymbol{x}}(F^\star, \cdot)$ when $\boldsymbol{x} \notin E_\circ^{\mathrm{lf}(F)}$ is irrelevant, we only need this extension to ensure $K_{\boldsymbol{x}}(F^\star, \cdot)$ can be integrated over $F^\star \in \mathrm{dec}(\mathbb{F})$ for fixed $\boldsymbol{x} \in \mathcal{X}$, like in (2.8) below.

Definition 2.10. A spatial coalescent $(\mathbb{P}^{\boldsymbol{x}})_{\boldsymbol{x} \in \mathcal{X}}$ is called a *Brownian spatial coalescent* if $K_{\boldsymbol{x}}$ is a conditional probability of $\mathbb{P}^{\boldsymbol{x}}$ given Dec, that is if

$$\mathbb{P}^{\boldsymbol{x}}(\cdot) = P^{\boldsymbol{x}} \otimes K_{\boldsymbol{x}}(\cdot) = \int_{\mathrm{dec}(\mathbb{F})} K_{\boldsymbol{x}}(F^\star, \cdot) P^{\boldsymbol{x}}(\mathrm{d}F^\star), \tag{2.8}$$

for every $\boldsymbol{x} \in \mathcal{X}$, where $P^{\boldsymbol{x}} \coloneqq \mathrm{Dec}\,\#\mathbb{P}^{\boldsymbol{x}} \in \mathcal{M}_1(\mathrm{dec}(\mathbb{F}))$.

In particular, a Brownian spatial coalescent is fully determined by the laws $P^{\boldsymbol{x}}$, so we can essentially think of it as a random, time and space decorated forest. This is a useful simplification, and the main work in the characterisation of Brownian spatial coalescents Theorem 2.14 is to understand which families of laws $(P^{\boldsymbol{x}})_{\boldsymbol{x} \in \mathcal{X}}$ give rise to a Markov process on Ω through (2.8).

We close by formalising the characterisation of a Brownian spatial coalescent in terms of its transition measures. Recall the heuristic description we provided in Section 1.3: for every tree (forest in general) we define a density on its spatial and time decorations, comprised of exponential factors for every merge event, and spatial factors associated with each branch of the forest. The density is w.r.t. Lebesgue measure in the time coordinates, and w.r.t. the transition measures in the spatial coordinates.

Definition 2.11. A family of *transition measures* is a collection $\boldsymbol{\nu} = (\nu_{\Pi,\Pi'}\colon \Pi' < \Pi)$, where $\nu_{\Pi,\Pi'}$ is a finite measure on $E_\circ^{\Pi' \setminus \Pi}$. Write R for the set of all families of transition measures.

The spatial factors (recall Fig. 1.1) do not depend on the transition measures, and are given for a non-trivial forest $F \in \mathbb{F}$, $\tau \in \mathrm{tm}(F)$ and $\boldsymbol{x} \in E_\circ^{\mathrm{lf}(F)}$, by

$$f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x}) \coloneqq \prod_{u \in \mathrm{nd}(F) \setminus \mathrm{rt}(F)} p(\tau_{\mathrm{pr}_F(u)} - \tau_u, (\boldsymbol{x}\xi)_u - \xi_{\mathrm{pr}_F(u)}), \qquad \xi \in \mathrm{sp}(F). \tag{2.9}$$

Here and in the following, we sometimes write $p(t, x)$ instead of $p_t(x)$ if it benefits legibility. If F is trivial we put $f_{\mathrm{sp}}|_{\mathrm{sp}(F)} \equiv 1$. The dependence on F is implicit in ξ and τ, but we will occasionally write f_{sp}^F to make it explicit. Given transition measures $\boldsymbol{\nu} \in R$ and $\Pi \in \mathcal{P}$, write

$$|\nu_\Pi| \coloneqq \sum_{\Pi' < \Pi} |\nu_{\Pi,\Pi'}|,$$

which is zero if $|\Pi| = 1$. If the process is label invariant, then $|\nu_\Pi|$ only depends on $n = |\Pi|$, and we write $|\nu_n|$ like we did in the introduction. It turns out that, as in the non-spatial setting, if $|\nu_\Pi| = 0$ then Π is an absorbing state (in the sense that almost-surely no further coalescence events occur), and if $|\nu_\Pi| > 0$ then almost-surely there is at least one more coalescence event. In particular, given an initial state $(\Pi, \boldsymbol{x}) \in \mathcal{X}$, only forests $F \in \mathbb{F}(\Pi)$ with $|\nu_{\mathrm{rt}(F)}| = 0$ are possible. Given such a forest, the full (unnormalised) density on $\mathrm{dec}(F)$ is

$$f_{\boldsymbol{\nu}}(\tau, \xi \,|\, \boldsymbol{x}) \coloneqq f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x}) \prod_{(\Pi,\Pi') \in F} \mathrm{e}^{-|\nu_\Pi|(\tau_{\Pi'} - \tau_\Pi)}, \qquad (\tau, \xi) \in \mathrm{dec}(F). \tag{2.10}$$

Here we identified F with the set $\{(\Pi_0^F, \Pi_1^F), \ldots, (\Pi_{m-1}^F, \Pi_m^F)\}$, which we will do in similar situations in the future if it is convenient and unambiguous. if $|\nu_{\mathrm{rt}(F)}| > 0$, then we put $f_{\boldsymbol{\nu}}(\cdot \,|\, \boldsymbol{x})|_{\mathrm{dec}(F)} \equiv 0$, which defines $f_{\boldsymbol{\nu}}(\cdot \,|\, \boldsymbol{x})$ on all of $\mathrm{dec}(\mathbb{F})$. Note if F is trivial and $|\nu_{\mathrm{rt}(F)}| = 0$ then $f_{\boldsymbol{\nu}}(\cdot \,|\, \boldsymbol{x})|_{\mathrm{dec}(F)} \equiv 1$. The density $f_{\boldsymbol{\nu}}$ is w.r.t. Lebesgue measure in the time coordinates, and w.r.t. the relevant transition measures in the spatial coordinates: if $F \in \mathbb{F}$ is non-trivial, write

$$\boldsymbol{\nu}_F(\mathrm{d}\xi) \coloneqq \prod_{(\Pi,\Pi') \in F} \nu_{\Pi,\Pi'}(\mathrm{d}\xi_{\Pi,\Pi'}), \tag{2.11}$$

where $\xi_{\Pi,\Pi'} \coloneqq \xi|_{\Pi' \setminus \Pi}$ for $\xi \in \mathrm{sp}(F)$ and $(\Pi, \Pi') \in F$, so that $\boldsymbol{\nu}_F$ is a finite measure on $\mathrm{sp}(F)$. If F is trivial then $\boldsymbol{\nu}_F$ denotes the unique probability measure on the singleton $\mathrm{sp}(F)$. It remains to define the normalisation. For $(\Pi, \boldsymbol{x}) \in \mathcal{X}$ and $F \in \mathbb{F}(\Pi)$ write

$$N_F^{\boldsymbol{\nu}}(\boldsymbol{x}) \coloneqq \int_{\mathrm{dec}(F)} f_{\boldsymbol{\nu}}(\tau, \xi \,|\, \boldsymbol{x}) \boldsymbol{\nu}_F(\mathrm{d}\xi)\, \mathrm{d}\tau, \qquad N^{\boldsymbol{\nu}}(\boldsymbol{x}) \coloneqq \sum_{F \in \mathbb{F}(\Pi)} N_{\boldsymbol{\nu}}^F(\boldsymbol{x}),$$

defining functions respectively on $E_\circ^{\mathrm{lf}(F)}$ and $\mathcal{X}$. These quantities are not obviously, and will not generally, be finite. The following lemma gives a sufficient condition. If $\boldsymbol{\lambda} \in R^\circ$ then we formally write $\boldsymbol{\nu}(\mathrm{d}\xi) = \boldsymbol{\lambda}\, \mathrm{d}\xi$ if $\nu_{\Pi,\Pi'}(\mathrm{d}\xi) = \lambda_{\Pi,\Pi'}\, \mathrm{d}\xi$ for all $\Pi' < \Pi$. We write $\boldsymbol{\nu}(\mathrm{d}\xi) \sim \mathrm{d}\xi$ if there exists $\boldsymbol{\lambda} \in R^\circ$ with $\boldsymbol{\nu}(\mathrm{d}\xi) = \boldsymbol{\lambda}\, \mathrm{d}\xi$.

Lemma 2.12. *If $\boldsymbol{\nu}(\mathrm{d}\xi) \sim \mathrm{d}\xi$ then $N^{\boldsymbol{\nu}}$ is continuous (in particular finite).*

Remark 2.13. Under the assumption of Lemma 2.12, $N^{\boldsymbol{\nu}}(\boldsymbol{x}) \to \infty$ as $\boldsymbol{x}$ approaches $E^{\Pi} \setminus E^{\Pi}_{\circ}$ for some $\Pi \in \mathcal{P}$.

Given a measure P on $\mathrm{dec}(\mathbb{F})$, it will be a convenient notation to write $P(F, \mathrm{d}\tau, \mathrm{d}\xi)$ for the measure on $\mathrm{dec}(F)$ defined by $\int_A P(F, \mathrm{d}\tau, \mathrm{d}\xi) \coloneqq P(\{F\} \times A)$ for $A \subset \mathrm{dec}(F)$. Analogously in similar contexts. With this notation, a formal statement of Theorems 1.7 and 1.8 is as follows.

Theorem 2.14. *A spatial coalescent is a Brownian spatial coalescent if and only if there exists $\boldsymbol{\nu} \in R$ such that $N^{\boldsymbol{\nu}}$ is continuous and*

$$P^{\boldsymbol{x}}(F, \mathrm{d}\tau, \mathrm{d}\xi) = \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f_{\boldsymbol{\nu}}(\tau, \xi \,|\, \boldsymbol{x}) \boldsymbol{\nu}_F(\mathrm{d}\xi)\, \mathrm{d}\tau. \tag{2.12}$$

We call it the Brownian spatial coalescent with transition measures $\boldsymbol{\nu}$. *In that case, $\boldsymbol{x} \mapsto \mathbb{P}^{\boldsymbol{x}}$ is continuous w.r.t. the topology of weak convergence.*

If a Brownian spatial coalescent is label invariant, then the family of transition measures reduces to a family $(\nu_{n,\vec{k}} \colon (n, \vec{k}) \in M)$ as we wrote in the introduction, see Lemma 3.13.

3 Proofs

3.1 Characterisation of Brownian Spatial Coalescents

In this section we prove Theorem 2.14.

3.1.1 "If" Direction of Theorem 2.14

Fix $\boldsymbol{\nu} \in R$ throughout, denote by $(P^{\boldsymbol{x}})_{\boldsymbol{x} \in \mathcal{X}}$ the laws defined in (2.12), and $\mathbb{P}^{\boldsymbol{x}} = P^{\boldsymbol{x}} \otimes K_{\boldsymbol{x}}$ for $\boldsymbol{x} \in \mathcal{X}$. We need to show that the coalescent process defined by $(\mathbb{P}^{\boldsymbol{x}})_{\boldsymbol{x} \in \mathcal{X}}$ is Markov. Recall that the associated semigroup is denoted $(S_t)_{t \geq 0}$, see (2.2). For fixed $\Pi' \leq \Pi$ and $A \subset E^{\Pi'}_{\circ}$ we write

$$S_t((\Pi, \boldsymbol{x}), (\Pi', A)) \coloneqq S_t((\Pi, \boldsymbol{x}), \{(\Pi', \boldsymbol{y}) \colon \boldsymbol{y} \in A\}),$$

which defines a sub-probability measure on $E^{\Pi'}_{\circ}$ with total mass $\mathbb{P}^{(\Pi, \boldsymbol{x})}(\Pi_t = \Pi')$. We introduce some additional notation. Let $F \in \mathbb{F}$ be a (possibly trivial) forest, $\boldsymbol{x} \in E^{\mathrm{lf}(F)}_{\circ}$, $\tau \in \mathrm{tm}(F)$, $\xi \in \mathrm{sp}(F)$, and $t > \tau_{\mathrm{rt}(F)}$, $\boldsymbol{y} \in E^{\mathrm{rt}(F)}_{\circ}$. Then

$$f_{\boldsymbol{\nu}}(\tau, \xi, t, \boldsymbol{y} \,|\, \boldsymbol{x}) \coloneqq f_{\boldsymbol{\nu}}(\tau, \xi \,|\, \boldsymbol{x}) \mathrm{e}^{-|\nu_{\mathrm{rt}(F)}|(t - \tau_{\mathrm{rt}(F)})} \prod_{u \in \mathrm{rt}(F)} p_{t - \tau_u}(\boldsymbol{y}_u - (\boldsymbol{x}\xi)_u)$$

If $t \leq \tau_{\mathrm{rt}(F)}$ we put $f_{\boldsymbol{\nu}}(\tau, \xi, t, \boldsymbol{y} \,|\, \boldsymbol{x}) \coloneqq 0$. Think of this as a version of $f_{\boldsymbol{\nu}}(\tau, \xi \,|\, \boldsymbol{x})$ with "terminal condition $\boldsymbol{y}$ at time t".

Lemma 3.1. *If* $(\Pi, \boldsymbol{x}) \in \mathcal{X}$ *and* $\Pi' \leq \Pi$*, then*

$$S_t((\Pi, \boldsymbol{x}), (\Pi', \mathrm{d}\boldsymbol{y})) = \frac{N^{\boldsymbol{\nu}}(\boldsymbol{y})}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} \sum_{F \in \mathbb{F}(\Pi, \Pi')} \left(\int_{\mathrm{dec}(F)} f_{\boldsymbol{\nu}}(\tau, \xi, t, \boldsymbol{y} \,|\, \boldsymbol{x}) \boldsymbol{\nu}_F(\mathrm{d}\xi)\, \mathrm{d}\tau \right) \mathrm{d}\boldsymbol{y}.$$

We introduce notation for the proof. If $\Pi_1 \geq \Pi_2 \geq \Pi_3$ and $F \in \mathbb{F}(\Pi_1, \Pi_2), F' \in \mathbb{F}(\Pi_2, \Pi_3)$, then we write $FF' \coloneqq F \cup F' \in \mathbb{F}(\Pi_1, \Pi_3)$ for the concatenation of F and F'. If further $\tau \in \mathrm{tm}(F)$ and $\tau' \in \mathrm{tm}(F')$, and $t > \tau_{\mathrm{rt}(F)}$, write

$$(\tau/t/\tau') \colon FF' \to (0, \infty); \quad \Pi \mapsto \begin{cases} \tau(\Pi), & \Pi \in F, \\ t + \tau'(\Pi), & \Pi \in F' \setminus \{\mathrm{lf}(F')\}, \end{cases} \tag{3.1}$$

so that $(\tau/t/\tau') \in \mathrm{tm}(FF')$ and $\xi\xi' \in \mathrm{sp}(FF')$ (the latter is defined in the sense of (2.4)). If A is a statement, and x is an expression that evaluates to a real number if A is true (and may otherwise be ill-defined), then we use $[x]_A$ as short-hand for x if A is true, and 1 otherwise.

Proof of Lemma 3.1. The possible values of Fr of a path in $\mathcal{P}$ that starts at Π and passes through Π' are exactly FF' for $F \in \mathbb{F}(\Pi, \Pi')$ and $F' \in \mathbb{F}(\Pi')$. If $\mathrm{Fr} = FF'$, then Tm has to be of the form $(\tau/t/\tau')$ for $\tau \in \mathrm{tm}(F)$ with $\tau_{\Pi'} < t$ and $\tau' \in \mathrm{tm}(F')$, and Sp has to be of the form $\xi\xi'$ for $\xi \in \mathrm{sp}(F)$ and $\xi' \in \mathrm{sp}(F')$. Conditional on $\mathrm{Fr} = FF'$, $\mathrm{Tm} = (\tau/t/\tau')$, and $\mathrm{Sp} = \xi\xi'$, which happens with probability

$$\frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f_{\boldsymbol{\nu}}((\tau/t/\tau'), \xi\xi' \,|\, \boldsymbol{x}) \boldsymbol{\nu}_F(\mathrm{d}\xi) \boldsymbol{\nu}_{F'}(\mathrm{d}\xi')\, \mathrm{d}\tau\, \mathrm{d}\tau', \tag{3.2}$$

the probability density of $\boldsymbol{X}_t \in (\Pi', \mathrm{d}\boldsymbol{y})$ is, by (2.12) and definition of $K_{\boldsymbol{x}}$,

$$\begin{aligned} &K_{\boldsymbol{x}}\Big((FF', (\tau/t/\tau'), \xi\xi'), \{\boldsymbol{X}_t \in (\Pi', \mathrm{d}\boldsymbol{y})\}\Big) \\ &\qquad = \prod_{u \in \Pi'} p_{t-\tau_u}(y_u - (\boldsymbol{x}\xi)_u) \left[\frac{p_{\tau'_{\mathrm{pr}'(u)}}(y_u - \xi'_{\mathrm{pr}'(u)})}{p_{t+\tau'_{\mathrm{pr}'(u)}-\tau_u}((\boldsymbol{x}\xi)_u - \xi'_{\mathrm{pr}'(u)})} \right]_{\mathrm{pr}'(u) \neq \emptyset} \mathrm{d}\boldsymbol{y}, \end{aligned} \tag{3.3}$$

where $\mathrm{pr} \coloneqq \mathrm{pr}_F$ and $\mathrm{pr}' \coloneqq \mathrm{pr}_{F'}$. Multiplying (3.2) and (3.3) gives, after some careful cancellations,

$$\begin{aligned} &\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = FF', \mathrm{Tm} \in \mathrm{d}(\tau/t/\tau'), \mathrm{Sp} \in \mathrm{d}(\xi\xi'), \boldsymbol{X}_t \in (\Pi', \mathrm{d}\boldsymbol{y})) \\ &\qquad = \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f_{\boldsymbol{\nu}}(\tau, \xi, t, \boldsymbol{y} \,|\, \boldsymbol{x}) f_{\boldsymbol{\nu}}(\tau', \xi' \,|\, \boldsymbol{y}) \boldsymbol{\nu}_F(\mathrm{d}\xi) \boldsymbol{\nu}_{F'}(\mathrm{d}\xi')\, \mathrm{d}\tau\, \mathrm{d}\tau'\, \mathrm{d}\boldsymbol{y}. \end{aligned}$$

If F' is trivial and $\lambda_{\Pi'} > 0$, then the equality holds because both $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = FF')$ and thus the left-hand side (LHS), and $f_{\boldsymbol{\nu}}|_{\mathrm{dec}(F')}$ and thus the right-hand side (RHS) are zero. Integrating over τ, ξ, τ', ξ' gives

$$\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = FF', \boldsymbol{X}_t \in (\Pi', \mathrm{d}\boldsymbol{y})) = \frac{N^{F'}_{\boldsymbol{\nu}}(\boldsymbol{y})}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} \left(\int_{\mathrm{dec}(F)} f_{\boldsymbol{\nu}}(\tau, \xi, t, \boldsymbol{y} \,|\, \boldsymbol{x}) \boldsymbol{\nu}_F(\mathrm{d}\xi)\, \mathrm{d}\tau \right) \mathrm{d}\boldsymbol{y},$$

and then summing over F and F' gives the claim. □

Lemma 3.2. *If $t, s > 0$, then $S_t S_s = S_{t+s}$.*

A crucial ingredient in the proof is the following identity, which is an immediate consequence of the definitions. If $F, F' \in \mathbb{F}$ with $\mathrm{rt}(F) = \mathrm{lf}(F')$, $(\tau, \xi) \in \mathrm{dec}(F)$, $(\tau', \xi') \in \mathrm{dec}(F')$, $t, t' > 0$ with $\tau_{\mathrm{rt}(F)} < t$, and $\boldsymbol{y}' \in E_\circ^{\mathrm{lf}(F')}$, then

$$\int_{E_\circ^{\mathrm{rt}(F)}} f_{\boldsymbol{\nu}}(\tau, \xi, t, \boldsymbol{y} \,|\, \boldsymbol{x}) f_{\boldsymbol{\nu}}(\tau', \xi', t', \boldsymbol{y}' \,|\, \boldsymbol{y}) \,\mathrm{d}\boldsymbol{y} = f_{\boldsymbol{\nu}}((\tau/t/\tau'), \xi\xi', t + t', \boldsymbol{y}' \,|\, \boldsymbol{x}). \tag{3.4}$$

Proof of Lemma 3.2. Fix $(\Pi_1, \boldsymbol{x})$, $t, s > 0$, and $\Pi_2 \le \Pi_1$. Then by Lemma 3.1,

$$\begin{aligned}
&(S_t S_s)((\Pi_1, \boldsymbol{x}), (\Pi_2, \mathrm{d}\boldsymbol{y}))/\,\mathrm{d}\boldsymbol{y} \\
&\quad= \sum_{\Pi_2 \le \Pi' \le \Pi_1} \int_{\boldsymbol{z} \in E_\circ^{\Pi'}} S_s((\Pi', \boldsymbol{z}), (\Pi_2, \mathrm{d}\boldsymbol{y})) S_t((\Pi_1, \boldsymbol{x}), (\Pi', \mathrm{d}\boldsymbol{z}))/\,\mathrm{d}\boldsymbol{y} \\
&\quad= \sum_{\Pi_2 \le \Pi' \le \Pi_1} \int \mathrm{d}\boldsymbol{z} \left(\frac{N^{\boldsymbol{\nu}}(\boldsymbol{y})}{N^{\boldsymbol{\nu}}(\boldsymbol{z})} \sum_{F_2 \in \mathbb{F}(\Pi', \Pi_2)} \int_{\mathrm{dec}(F_2)} f_{\boldsymbol{\nu}}(\tau_2, \xi_2, s, \boldsymbol{y} \,|\, \boldsymbol{z}) \boldsymbol{\nu}_{F_2}(\mathrm{d}\xi_2) \,\mathrm{d}\tau_2 \right. \\
&\qquad\qquad \left. \times \frac{N^{\boldsymbol{\nu}}(\boldsymbol{z})}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} \sum_{F_1 \in \mathbb{F}(\Pi_1, \Pi')} \int_{\mathrm{dec}(F_1)} f_{\boldsymbol{\nu}}(\tau_1, \xi_1, t, \boldsymbol{z} \,|\, \boldsymbol{x}) \boldsymbol{\nu}_{F_1}(\mathrm{d}\xi_1) \,\mathrm{d}\tau_1 \right) \\
&\quad\overset{(3.4)}{=} \frac{N^{\boldsymbol{\nu}}(\boldsymbol{y})}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} \sum_{\substack{\Pi_2 \le \Pi' \le \Pi_1 \\ F_1 \in \mathbb{F}(\Pi_1, \Pi') \\ F_2 \in \mathbb{F}(\Pi', \Pi_2)}} \int_{\mathrm{dec}(F_1)} \int_{\mathrm{dec}(F_2)} f_{\boldsymbol{\nu}}((\tau_1/t/\tau_2), \xi_1\xi_2, t + s, \boldsymbol{y} \,|\, \boldsymbol{x}) \boldsymbol{\nu}_{F_1}(\mathrm{d}\xi_1) \,\mathrm{d}\tau_1 \\
&\qquad\qquad \boldsymbol{\nu}_{F_2}(\mathrm{d}\xi_2) \,\mathrm{d}\tau_2.
\end{aligned} \tag{3.5}$$

If $\Pi_2 \le \Pi' \le \Pi_1$ and $F_1 \in \mathbb{F}(\Pi_1, \Pi')$, $F_2 \in \mathbb{F}(\Pi', \Pi_2)$, then $\xi \in \mathrm{sp}(F_1 F_2)$ is always of the form $\xi = \xi_1\xi_2$ for $\xi_{1,2} \in \mathrm{sp}(F_{1,2})$, and $\tau \in \mathrm{tm}(F_1 F_2)$ is of the form $\tau = (\tau_1/t/\tau_2)$ for $\tau_{1,2} \in \mathrm{tm}(F_{1,2})$ if and only if $\tau_{\Pi'} < t < \tau_{\Pi'}^+$. Thus the final integral on the RHS of (3.5) is equal to

$$\int_{\mathrm{dec}(F_1 F_2)} \mathbb{1}_{\{\tau_{\Pi'} < t < \tau_{\Pi'}^+\}} f_{\boldsymbol{\nu}}(\tau, \xi, t + s, \boldsymbol{y} \,|\, \boldsymbol{x}) \,\mathrm{d}\tau \,\mathrm{d}\xi,$$

so

$$\begin{aligned}
&(S_t S_s)((\Pi_1, \boldsymbol{x}), (\Pi_2, \mathrm{d}\boldsymbol{y}))/\,\mathrm{d}\boldsymbol{y} \\
&\quad= \frac{N^{\boldsymbol{\nu}}(\boldsymbol{y})}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} \sum_{F \in \mathbb{F}(\Pi_1, \Pi_2)} \sum_{\Pi' \in F} \int_{\mathrm{dec}(F)} \mathbb{1}_{\{\tau_{\Pi'} < t < \tau_{\Pi'}^+\}} f_{\boldsymbol{\nu}}(\tau, \xi, t + s, \boldsymbol{y} \,|\, \boldsymbol{x}) \boldsymbol{\nu}_F(\mathrm{d}\xi) \,\mathrm{d}\tau \\
&\quad= \frac{N^{\boldsymbol{\nu}}(\boldsymbol{y})}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} \sum_{F \in \mathbb{F}(\Pi_1, \Pi_2)} \int_{\mathrm{dec}(F)} f_{\boldsymbol{\nu}}(\tau, \xi, t + s, \boldsymbol{y} \,|\, \boldsymbol{x}) \boldsymbol{\nu}_F(\mathrm{d}\xi) \,\mathrm{d}\tau \\
&\quad= S_{t+s}((\Pi_1, \boldsymbol{x}), (\Pi_2, \mathrm{d}\boldsymbol{y}))/\,\mathrm{d}\boldsymbol{y}.
\end{aligned}$$

□

This proves that the Brownian spatial coalescent with transition measures $\boldsymbol{\nu}$ is a Markov process.

Lemma 3.3. *The map $\boldsymbol{x} \mapsto \mathbb{P}^{\boldsymbol{x}}$ is continuous w.r.t. the topology of weak convergence of probability measures. In particular, $(\mathbb{P}^{\boldsymbol{x}})_{\boldsymbol{x}\in\mathcal{X}}$ has the Feller property.*

Proof. Suppose $O \subset \Omega$ is open, and $(\Pi_n, \boldsymbol{x}_n) \to (\Pi, \boldsymbol{x})$ in $\mathcal{X}$, without loss of generality $\Pi_n = \Pi$ for all $n \in \mathbb{N}$. Then,

$$\begin{aligned}\mathbb{P}^{\boldsymbol{x}_n}(O) &= \int_{\mathrm{dec}(\mathbb{F})} K_{\boldsymbol{x}_n}(F^\star, O) P^{\boldsymbol{x}_n}(\mathrm{d}F^\star) \\ &= \sum_{F\in\mathbb{F}(\Pi)} \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x}_n)} \int_{\mathrm{dec}(F)} K_{\boldsymbol{x}_n}((F,\tau,\xi), O) f_{\boldsymbol{\nu}}(\tau, \xi \,|\, \boldsymbol{x}_n) \boldsymbol{\nu}_F(\mathrm{d}\xi)\, \mathrm{d}\tau.\end{aligned}$$

By Lemmas 2.9 and 2.12, and definition of $f_{\boldsymbol{\nu}}$, we have $N^{\boldsymbol{\nu}}(\boldsymbol{x}_n) \to N^{\boldsymbol{\nu}}(\boldsymbol{x})$, $f_{\boldsymbol{\nu}}(\tau, \xi \,|\, \boldsymbol{x}_n) \to f_{\boldsymbol{\nu}}(\tau, \xi \,|\, \boldsymbol{x})$, and $\underline{\lim}_{n\to\infty} K_{\boldsymbol{x}_n}(F^\star, O) \geq K_{\boldsymbol{x}}(F^\star, O)$ for all $F^\star = (F, \tau, \xi) \in \mathrm{dec}(\mathbb{F})$. Thus by Fatou's lemma,

$$\underline{\lim_{n\to\infty}}\, \mathbb{P}^{\boldsymbol{x}_n}(O) \geq \sum_{F\in\mathbb{F}(\Pi)} \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} \int_{\mathrm{dec}(F)} K_{\boldsymbol{x}}((F,\tau,\xi), O) f_{\boldsymbol{\nu}}(\tau, \xi \,|\, \boldsymbol{x}_n) \boldsymbol{\nu}_F(\mathrm{d}\xi)\, \mathrm{d}\tau = \mathbb{P}^{\boldsymbol{x}}(O).$$

This proves that $\boldsymbol{x} \mapsto \mathbb{P}^{\boldsymbol{x}}$ is continuous by the Portmanteau theorem. □

3.1.2 "Only if" Direction of Theorem 2.14

Fix a Brownian spatial coalescent process, that is a spatial coalescent whose laws can be written as $\mathbb{P}^{\boldsymbol{x}} = P^{\boldsymbol{x}} \otimes K_{\boldsymbol{x}}$ where $P^{\boldsymbol{x}} = \mathrm{Dec}\,\#\mathbb{P}^{\boldsymbol{x}}$ for $\boldsymbol{x} \in \mathcal{X}$ (recall Definition 2.10). We have to prove that the laws $(P^{\boldsymbol{x}})$ are of the form (2.12).

The first step revolves around the following idea. Suppose $(\Pi, \boldsymbol{x}) \in \mathcal{X}$, $s > 0$, and $F_0 \in \mathbb{F}(\Pi)$ is the trivial forest. Then we can evaluate $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0, \boldsymbol{X}_s \in \mathrm{d}\boldsymbol{y})$, a sub-probability measure on $E_\circ^\Pi$, in two ways: by applying the Markov property at time s it equals $S_s(\boldsymbol{x}, \mathrm{d}\boldsymbol{y})\mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F_0)$, and by (2.12) it equals $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0) K_{\boldsymbol{x}}(F_0, \boldsymbol{X}_s \in \mathrm{d}\boldsymbol{y})$, giving

$$S_s(\boldsymbol{x}, \mathrm{d}\boldsymbol{y})\mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F_0) = \mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0) \prod_{u\in\Pi} p_s(\boldsymbol{x}_u - \boldsymbol{y}_u)\, \mathrm{d}\boldsymbol{y}$$

as measures on $E_\circ^\Pi$, where we expanded $K_{\boldsymbol{x}}(F_0, \boldsymbol{X}_s \in \mathrm{d}\boldsymbol{y})$. This already implies the main claim of the following proposition in the case where $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0) > 0$ for all $\boldsymbol{x} \in E_\circ^\Pi$. For $F \in \mathbb{F}$, $\tau \in \mathrm{tm}(F)$ and $s \geq 0$ we write $\tau + s$ for the map $u \mapsto \tau_u + s$, and $\tau \geq s$ if $\tau = \tau' + s$ for some $\tau' \in \mathrm{tm}(F)$.

Proposition 3.4. *For every $\Pi \in \mathcal{P}$, there is a unique constant $\beta_\Pi \geq 0$, and a continuous function $\widetilde{N}_\Pi \colon E_\circ^\Pi \to (0, \infty)$ unique up to a constant positive multiple such that for all $s > 0$,*

$$S_s((\Pi, \boldsymbol{x}), (\Pi, \mathrm{d}\boldsymbol{y})) = \mathrm{e}^{-\beta_\Pi s} \frac{\widetilde{N}_\Pi(\boldsymbol{y})}{\widetilde{N}_\Pi(\boldsymbol{x})} \prod_{u\in\Pi} p_s(\boldsymbol{x}_u - \boldsymbol{y}_u)\, \mathrm{d}\boldsymbol{y}. \tag{3.6}$$

Furthermore, for any $F \in \mathbb{F}(\Pi)$, exactly one of the following hold

(i) $\mathbb{P}^{(\Pi,\boldsymbol{x})}(\mathrm{Fr} = F) > 0$ *for all* $\boldsymbol{x} \in E_\circ^\Pi$*, in which case we call* F possible*,*

(ii) $\mathbb{P}^{(\Pi,\boldsymbol{x})}(\mathrm{Fr} = F) = 0$ *for all* $\boldsymbol{x} \in E_\circ^\Pi$*, in which case we call* F impossible*.*

If $\beta_\Pi > 0$ *then the trivial forest is impossible, and if* $\beta_\Pi = 0$ *then the trivial forest is the only possible forest and* $\widetilde{N}_\Pi$ *is constant.*

It will turn out that $\beta_\Pi = |\nu_\Pi|$, and that $\widetilde{N}_\Pi$ is (a constant multiple of) the normalisation $N^{\boldsymbol{\nu}}$ for the transition measures $\boldsymbol{\nu} \in R$ that define this Brownian spatial coalescent, see Theorem 3.9. Proposition 3.4 is proved in a sequence of lemmas.

Lemma 3.5. *If* $(\Pi, \boldsymbol{x}) \in \mathcal{X}$*,* $t > 0$*, then* $S_t((\Pi, \boldsymbol{x}), (\Pi, \mathrm{d}\boldsymbol{y}))$ *is not zero as a measure on* $E_\circ^\Pi$*.*

Proof. Fix $(\Pi, \boldsymbol{x}) \in \mathcal{X}$. Since $(\Pi_t, \boldsymbol{X}_t)_{t\geq 0}$ is right-continuous, we have $\Pi_t \to \Pi_0 = \Pi$ a.s. as $t \to 0$, in particular in distribution, so $S_t(\boldsymbol{x}, E_\circ^\Pi) = \mathbb{P}^{\boldsymbol{x}}(\Pi_t = \Pi) \to 1$ as $t \to 0$. This implies the claim for sufficiently small $s > 0$, and repeated applications of $S_{2s} = S_s S_s$ finish the proof. □

Lemma 3.6. *For every* $\Pi \in \mathcal{P}$ *there are a unique* $\beta_\Pi \geq 0$ *and a function* $\widetilde{N}_\Pi \colon E_\circ^\Pi \to (0, \infty)$ *unique up to a constant positive multiple such that* (3.6) *holds. Every* $F \in \mathbb{F}(\Pi)$ *is either possible or impossible, and if the trivial forest* $F_0 \in \mathbb{F}(\Pi)$ *is possible, then* $\widetilde{N}_\Pi(\boldsymbol{x})\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0)$ *is constant in* $\boldsymbol{x} \in E_\circ^\Pi$ *and positive. Furthermore, for every possible* $F \in \mathbb{F}(\Pi)$*, every* $\boldsymbol{x}, \boldsymbol{y} \in E_\circ^\Pi$*, and* $s \geq 0$*,*

$$\mathrm{e}^{\beta_\Pi s} \frac{\widetilde{N}_\Pi(\boldsymbol{x}) P^{\boldsymbol{x}}(F, \mathrm{d}(\tau + s), \mathrm{d}\xi)}{\prod\limits_{\substack{u \in \Pi \\ \mathrm{pr}_F(u) \neq \emptyset}} p(\tau_{\mathrm{pr}_F(u)} + s, \xi_{\mathrm{pr}_f(u)} - \boldsymbol{x}_u)} = \frac{\widetilde{N}_\Pi(\boldsymbol{y}) P^{\boldsymbol{y}}(F, \mathrm{d}\tau, \mathrm{d}\xi)}{\prod\limits_{\substack{u \in \Pi \\ \mathrm{pr}_F(u) \neq \emptyset}} p(\tau_{\mathrm{pr}_F(u)}, \xi_{\mathrm{pr}_F(u)} - \boldsymbol{y}_u)} \tag{3.7}$$

as measures on $\mathrm{dec}(F)$*.*

Proof. Fix $\Pi \in \mathcal{P}$ throughout the proof. Uniqueness is straightforward: if (3.6) holds with β and $\widetilde{N}$, but also with β' and $\widetilde{N}'$, then for every $s > 0$, $\boldsymbol{x} \in E_\circ^\Pi$, and Lebesgue-almost all $\boldsymbol{y} \in E_\circ^\Pi$ (with the exceptional null-set depending on $\boldsymbol{x}$ and s),

$$\mathrm{e}^{-\beta s} \frac{\widetilde{N}(\boldsymbol{y})}{\widetilde{N}(\boldsymbol{x})} = \mathrm{e}^{-\beta' s} \frac{\widetilde{N}'(\boldsymbol{y})}{\widetilde{N}'(\boldsymbol{x})}. \tag{3.8}$$

Letting $s \to 0$ in $\mathbb{Q}$, it follows that for every $\boldsymbol{x}$, and all $\boldsymbol{y} \notin N_{\boldsymbol{x}}$ for a null-set $N_{\boldsymbol{x}} \subset E_\circ^\Pi$, we have $\widetilde{N}(\boldsymbol{y})/\widetilde{N}(\boldsymbol{x}) = \widetilde{N}'(\boldsymbol{y})/\widetilde{N}'(\boldsymbol{x})$. For a fixed choice $\boldsymbol{x}_0$ and $a \coloneqq \widetilde{N}(\boldsymbol{x}_0)/\widetilde{N}'(\boldsymbol{x}_0) > 0$, this implies $\widetilde{N}(\boldsymbol{y}) = a\widetilde{N}'(\boldsymbol{y})$ for all $\boldsymbol{y} \in E_\circ^\Pi \setminus N_{\boldsymbol{x}_0}$. Then for any $\boldsymbol{x} \in E_\circ^\Pi$ we can choose $\boldsymbol{y} \in E_\circ^\Pi \setminus (N_{\boldsymbol{x}_0} \cup N_{\boldsymbol{x}})$ to obtain $\widetilde{N}(\boldsymbol{x}) = \frac{\widetilde{N}(\boldsymbol{y})}{\widetilde{N}'(\boldsymbol{y})}\widetilde{N}'(\boldsymbol{x}) = a\widetilde{N}'(\boldsymbol{x})$. Back into (3.8) this also implies $\beta = \beta'$. Conversely, if (3.7) holds for some function $\widetilde{N}_\Pi$, then it also holds for $c\widetilde{N}_\Pi$ for any $c > 0$.

If $|\Pi| = 1$ then (3.6) holds for $\beta_\Pi = 0$ and $\widetilde{N}_\Pi \equiv 1$ (because a single lineage follows a Brownian motion), and the only element of $\mathbb{F}(\Pi)$ is the trivial forest, which makes the remaining statements trivial. Assume $|\Pi| > 1$ for the rest of the proof. Let $F_0 \in \mathbb{F}(\Pi)$ be the trivial forest. For $s > 0$ and $\boldsymbol{x} \in E_\circ^\Pi$, we evaluate $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0, \boldsymbol{X}_s \in \mathrm{d}\boldsymbol{y})$, which is a sub-probability measure on $E_\circ^\Pi$, in two ways. On the one hand, by applying the Markov property at time s it equals $S_s(\boldsymbol{x}, \mathrm{d}\boldsymbol{y})\mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F_0)$. On the other hand, by (2.12) it equals $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0)K_{\boldsymbol{x}}(F_0, \boldsymbol{X}_s \in \mathrm{d}\boldsymbol{y})$. This gives

$$S_s(\boldsymbol{x}, \mathrm{d}\boldsymbol{y})\mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F_0) = \mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0) \prod_{u \in \Pi} p_s(\boldsymbol{x}_u - \boldsymbol{y}_u)\, \mathrm{d}\boldsymbol{y} \tag{3.9}$$

as measures on $E_\circ^\Pi$, where we expanded $K_{\boldsymbol{x}}(F_0, \boldsymbol{X}_s \in \mathrm{d}\boldsymbol{y})$. We now explain why F_0 is either possible or impossible. If $\mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F_0) = 0$ on a set of positive Lebesgue measure, then integrating (3.9) over $\boldsymbol{y}$ in this set gives zero on the LHS, and the integral on the RHS is zero only if $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0) = 0$, so we must have $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0) = 0$ for all $\boldsymbol{x} \in E_\circ^\Pi$ and F_0 is impossible. Now assume that $\mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F_0) > 0$ Lebesgue-almost everywhere. Then integrating (3.9) over $\boldsymbol{y} \in E_\circ^\Pi$ gives a positive number on the LHS by Lemma 3.5, and $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0)$ on the RHS, so $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0) > 0$ for all $\boldsymbol{x} \in E_\circ^\Pi$ and F_0 is possible. In that case, (3.9) implies (3.6) with $\beta_\Pi = 0$ and $\widetilde{N}_\Pi(\boldsymbol{x}) = 1/\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0)$. Eq. (3.7) for the trivial forest F_0 is just $\widetilde{N}_\Pi(\boldsymbol{x})\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0) = \widetilde{N}_\Pi(\boldsymbol{y})\mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F_0)$, which holds by definition of $\widetilde{N}_\Pi$.

Let now $F \in \mathbb{F}(\Pi)$ be non-trivial. For $s > 0$ and $\boldsymbol{x} \in E_\circ^\Pi$, we evaluate

$$\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F, \mathrm{Tm} \in \mathrm{d}(\tau + s), \mathrm{Sp} \in \mathrm{d}\xi, \boldsymbol{X}_s \in \mathrm{d}\boldsymbol{y}), \tag{3.10}$$

which is a sub-probability measure on $\mathrm{dec}(F) \times E_\circ^\Pi$, in two ways. On the one hand, by applying the Markov property at time s it equals $S_s(\boldsymbol{x}, \mathrm{d}\boldsymbol{y})P^{\boldsymbol{y}}(F, \mathrm{d}\tau, \mathrm{d}\xi)$. On the other hand, by conditioning on $\mathrm{Dec} = (F, \tau, \xi)$ and (2.12), the expression (3.10) equals

$$P^{\boldsymbol{x}}(F, \mathrm{d}(\tau + s), \mathrm{d}\xi)K_{\boldsymbol{x}}((F, \tau + s, \xi), \{\boldsymbol{X}_s \in \mathrm{d}\boldsymbol{y}\}).$$

Thus we obtain

$$\begin{aligned} &S_s(\boldsymbol{x}, \mathrm{d}\boldsymbol{y})P^{\boldsymbol{y}}(F, \mathrm{d}\tau, \mathrm{d}\xi) \\ &\qquad = P^{\boldsymbol{x}}(F, \mathrm{d}(\tau + s), \mathrm{d}\xi) \prod_{u \in \Pi} p_s(\boldsymbol{x}_u - \boldsymbol{y}_u) \left[\frac{p(\tau_{\mathrm{pr}(u)}, \xi_{\mathrm{pr}(u)} - \boldsymbol{y}_u)}{p(\tau_{\mathrm{pr}(u)} + s, \xi_{\mathrm{pr}(u)} - \boldsymbol{x}_u)}\right]_{\mathrm{pr}(u) \neq \emptyset} \mathrm{d}\boldsymbol{y}, \end{aligned} \tag{3.11}$$

as an equality of measures on $\mathrm{tm}(F) \times \mathrm{sp}(F) \times E_\circ^\Pi$, where we expanded $K_{\boldsymbol{x}}((F, \tau + s, \xi), \boldsymbol{X}_s \in \mathrm{d}\boldsymbol{y})$ on the RHS, and abbreviated $\mathrm{pr} = \mathrm{pr}_F$.

We now explain why F is either possible or impossible. Suppose that $\mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F) = 0$ for $\boldsymbol{y}$ in a set of positive Lebesgue measure. Integrating over this set and $(\tau, \xi) \in \mathrm{Dec}(F)$ gives zero on the LHS, and on the RHS a quantity that is zero if and only

if $P^{\boldsymbol{x}}(F, \{\tau \geq s\}) = 0$, which converges to $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F)$ as $s \to 0$ by continuity from below, so $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F) = 0$ for all $\boldsymbol{x} \in E_\circ^\Pi$ and F is impossible. Now suppose that $\mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F) > 0$ for Lebesgue almost-all $\boldsymbol{y} \in E_\circ^\Pi$. Then $S_s(\boldsymbol{x}, \mathrm{d}\boldsymbol{y})$ has a Lebesgue density. Indeed, integrating (3.11) over $\boldsymbol{y} \in N$ for a Lebesgue null set N, and $(\tau, \xi) \in \mathrm{dec}(F)$, gives zero on the RHS, and on the LHS it gives $\int_N \mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F) S_s(\boldsymbol{x}, \mathrm{d}\boldsymbol{y})$, which was positive if $S_s(\boldsymbol{x}, N) > 0$. Then, for an arbitrary $s > 0$,

$$\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F) \geq \mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F, \tau \geq s) = \int_{E_\circ^\Pi} \mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F) S_s(\boldsymbol{x}, \mathrm{d}\boldsymbol{y}) > 0, \tag{3.12}$$

so F is possible. In the second step, we applied the Markov property at time s, and the last expression is positive because $\mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F) > 0$ for Lebesgue-a.e. $\boldsymbol{y} \in E_\circ^\Pi$, and $S_s(\boldsymbol{x}, \cdot)$ is non-zero (Lemma 3.5) and has a Lebesgue density. In fact, (3.12) shows that if F is possible, then $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F, \tau \geq s) > 0$ for all $\boldsymbol{x} \in E_\circ^\Pi$ and $s \geq 0$.

We now prove (3.6) in case the trivial forest is impossible, as well as (3.7) for non-trivial forests. Thus we may assume that there exists a possible non-trivial forest $F \in \mathbb{F}(\Pi)$, for otherwise we are already done. Let such an F be given, and define for every $\boldsymbol{x} \in E_\circ^\Pi$,

$$G(\boldsymbol{x}, F, \mathrm{d}\tau, \mathrm{d}\xi) \coloneqq \frac{P^{\boldsymbol{x}}(F, \mathrm{d}\tau, \mathrm{d}\xi)}{\prod\limits_{\substack{u \in \Pi \\ \mathrm{pr}_F(u) \neq \emptyset}} p(\tau_{\mathrm{pr}_F(u)}, \xi_{\mathrm{pr}_F(u)} - \boldsymbol{x}_u)}, \tag{3.13}$$

which is a non-zero measure on $\mathrm{dec}(F)$ that is finite on $\{\tau \geq a\} \times \mathrm{sp}(F) \subset \mathrm{dec}(F)$ for every $a > 0$. Recall from the paragraph following (3.12) that $P^{\boldsymbol{x}}(F, \{\tau \geq a\}) > 0$ for every $a \geq 0$ and $\boldsymbol{x} \in E_\circ^\Pi$, so

$$G_F(\boldsymbol{x}, s) \coloneqq G(\boldsymbol{x}, F, \{\tau > 1 + s\} \times \mathrm{sp}(F)) = \int\limits_{\mathrm{tm}(F)} \mathbb{1}_{\{\tau > 1+s\}} \int\limits_{\mathrm{sp}(F)} G(\boldsymbol{x}, F, \mathrm{d}\tau, \mathrm{d}\xi) \in (0, \infty)$$

for all $\boldsymbol{x} \in E_\circ^\Pi$ and $s \geq 0$. By continuity from below, $G_F(\boldsymbol{x}, \cdot)$ is right-continuous on $[0, \infty)$. Rearranging (3.11) gives

$$G(\boldsymbol{x}, F, \mathrm{d}(\tau + s), \mathrm{d}\xi)\, \mathrm{d}\boldsymbol{y} = \frac{S_s(\boldsymbol{x}, \mathrm{d}\boldsymbol{y})}{\prod_{u \in \Pi} p_s(\boldsymbol{x}_u - \boldsymbol{y}_u)} G(\boldsymbol{y}, F, \mathrm{d}\tau, \mathrm{d}\xi).$$

If we integrate this over $\{\tau > 1 + t\} \times \mathrm{sp}(F) \subset \mathrm{dec}(F)$ for some $t \geq 0$, we obtain

$$\frac{S_s(\boldsymbol{x}, \mathrm{d}\boldsymbol{y})}{\prod_{u \in \Pi} p_s(\boldsymbol{x}_u - \boldsymbol{y}_u)} = \frac{G_F(\boldsymbol{x}, t + s)}{G_F(\boldsymbol{y}, t)}\, \mathrm{d}\boldsymbol{y} \tag{3.14}$$

for $s > 0$. In particular, the RHS does not depend on F, and since the LHS is independent of t, we can equate the RHS for $t = 0$ and $t \geq 0$ to obtain that for every $s > 0, t \geq 0$, $\boldsymbol{x} \in E_\circ^\Pi$,

$$g(s, \boldsymbol{x}, \boldsymbol{y}) \coloneqq \frac{G_F(\boldsymbol{x}, s)}{G_F(\boldsymbol{y}, 0)} = \frac{G_F(\boldsymbol{x}, t + s)}{G_F(\boldsymbol{y}, t)}, \tag{3.15}$$

for Lebesgue-almost all $\boldsymbol{y} \in E_\circ^\Pi$, and g does not depend on F. By right-continuity of $G_F(\boldsymbol{x}, \cdot)$, the same is true for $s = 0$ (and any $t \geq 0$). Thus for $t, s \geq 0$ and $\boldsymbol{x} \in E_\circ^\Pi$, there is a Lebesgue-null set $N_{\boldsymbol{x},s,t} \subset E_\circ^\Pi$ such that for all $\boldsymbol{z} \notin N_{\boldsymbol{x},s,t}$, and all $\boldsymbol{y} \in E_\circ^\Pi$,

$$g(s, \boldsymbol{x}, \boldsymbol{z}) g(t, \boldsymbol{z}, \boldsymbol{y}) = \frac{G_F(\boldsymbol{x}, s)}{G_F(\boldsymbol{z}, 0)} \frac{G_F(\boldsymbol{z}, t)}{G_F(\boldsymbol{y}, 0)} = \frac{G_F(\boldsymbol{x}, s+t)}{G_F(\boldsymbol{z}, t)} \frac{G_F(\boldsymbol{z}, t)}{G_F(\boldsymbol{y}, 0)} = g(s+t, \boldsymbol{x}, \boldsymbol{y}).$$

This implies[5] that there exists $\beta \in \mathbb{R}$ such that for $\boldsymbol{x} \in E_\circ^\Pi$, $s \in \mathbb{Q} \cap [0, \infty)$, and Lebesgue-almost all $\boldsymbol{y} \in E_\circ^\Pi$,

$$g(s, \boldsymbol{x}, \boldsymbol{y}) = \mathrm{e}^{-\beta s} \frac{G_F(\boldsymbol{x}, 0)}{G_F(\boldsymbol{y}, 0)} = \mathrm{e}^{-\beta s} \frac{\widetilde{N}_F(\boldsymbol{y})}{\widetilde{N}_F(\boldsymbol{x})}, \tag{3.16}$$

where $\widetilde{N}_F(\boldsymbol{x}) \coloneqq G_F(\boldsymbol{x}, 0)^{-1}$. If the trivial forest is impossible, in which case we have not yet found β_Π and $\widetilde{N}_\Pi$, (3.14) to (3.16) imply (3.6) for $s \in \mathbb{Q} \cap (0, \infty)$ with $\beta_\Pi \coloneqq \beta$ and $\widetilde{N}_\Pi \coloneqq \widetilde{N}_F$ (the fact that this choice for β_Π and $\widetilde{N}_\Pi$ works implies by uniqueness that the construction above would have yielded the same $\widetilde{N}_F$ for any F, up to a constant multiple). In that case, if we assume $\beta_\Pi < 0$, then integrating (3.6) over $\boldsymbol{y} \in E_\circ^\Pi$ gives

$$1 \geq \int S_s(\boldsymbol{x}, \mathrm{d}\boldsymbol{y}) \geq \mathrm{e}^{|\beta_\Pi| s} \frac{1}{\widetilde{N}_\Pi(\boldsymbol{x})} (\min_E p_s)^{|\Pi|} \int \widetilde{N}_\Pi(\boldsymbol{y}) \, \mathrm{d}\boldsymbol{y} \overset{s \to \infty}{\longrightarrow} \infty,$$

a contradiction. We now explain why it sufficed to show (3.6) for $s \in \mathbb{Q}_+ \coloneqq \mathbb{Q} \cap (0, \infty)$. Denote the RHS in (3.6) for fixed $\boldsymbol{x} \in E_\circ^\Pi$ temporarily by $h(s, \boldsymbol{y}) \, \mathrm{d}\boldsymbol{y}$, which has now been shown to be a finite measure on $E_\circ^\Pi$ for all $s \in \mathbb{Q}_+$ (because the LHS in (3.6) is a finite measure). Then $\widetilde{N}_\Pi$ is integrable because $\inf p_s > 0$ for any fixed $s \in \mathbb{Q}_+$, which implies that $h(s, \boldsymbol{y}) \, \mathrm{d}\boldsymbol{y}$ is a finite measure for all $s > 0$, and enables an application of dominated convergence to show that $s \mapsto h(s, \boldsymbol{y}) \, \mathrm{d}\boldsymbol{y}$ is continuous in the weak topology on $\mathcal{M}_F(E_\circ^\Pi)$. The LHS of (3.6) is right-continuous in s w.r.t. the same topology because $(\Pi_t, \boldsymbol{X}_t)$ is almost-surely right-continuous, which implies that (3.6) holds for all $s > 0$.

We prove (3.7) for every non-trivial, possible forest $F \in \mathbb{F}(\Pi)$. Fix one, if it exists, abbreviate $\beta = \beta_\Pi$, $\widetilde{N} = \widetilde{N}_\Pi$, and recall the definition of $G(\boldsymbol{x}, F, \mathrm{d}\tau, \mathrm{d}\xi)$ from (3.13). Then (3.7) is equivalent to $\mathrm{e}^{\beta s} \widetilde{N}(\boldsymbol{x}) G(\boldsymbol{x}, F, \mathrm{d}(\tau + s), \mathrm{d}\xi) = \widetilde{N}(\boldsymbol{y}) G(\boldsymbol{y}, F, \mathrm{d}\tau, \mathrm{d}\xi)$. Plugging (3.6) into (3.11) and rearranging shows that for every $\boldsymbol{x} \in E_\circ^\Pi$ and $s > 0$,

$$\mathrm{e}^{\beta s} \widetilde{N}(\boldsymbol{x}) G(\boldsymbol{x}, F, \mathrm{d}(\tau + s), \mathrm{d}\xi) \, \mathrm{d}\boldsymbol{y} = \widetilde{N}(\boldsymbol{y}) G(\boldsymbol{y}, F, \mathrm{d}\tau, \mathrm{d}\xi) \, \mathrm{d}\boldsymbol{y} \tag{3.17}$$

as measures on $\mathrm{dec}(F) \times E_\circ^\Pi$. Since $G(\boldsymbol{x}, F, \mathrm{d}\tau, \mathrm{d}\xi)$ is a locally finite measure on $\mathrm{dec}(F)$, the LHS in (3.17) is continuous in $s \in [0, \infty)$ w.r.t. the topology of vague convergence on the space of locally finite measures on $\mathrm{dec}(F) \times E_\circ^\Pi$. Thus we can let $s \to 0$ in (3.17) and obtain that it also holds for $s = 0$. Since the σ-algebra on $\mathrm{dec}(F)$ is countably generated,

[5] The sceptical reader may apply Lemma II.F.2 to $\log g(s, \boldsymbol{x}, \boldsymbol{y})$, which is well-defined because $g(s, \boldsymbol{x}, \boldsymbol{y}) > 0$. The assumptions of Lemma II.F.2 are satisfied with $b(\boldsymbol{x}) = \log G_F(\boldsymbol{x}, 0)$, and $A_{\boldsymbol{x}} \coloneqq E_\circ^\Pi \setminus \bigcup_{s,t \in \mathbb{Q} \cap [0,\infty)} N_{\boldsymbol{x},s,t}$.

we can further find for every $\boldsymbol{x} \in E_\circ^\Pi$ and $s \geq 0$ a single null set $N_{\boldsymbol{x},s} \subset E_\circ^\Pi$ such that for all $\boldsymbol{y} \in E_\circ^\Pi \setminus N_{\boldsymbol{x},s}$,

$$\mathrm{e}^{\beta s}\widetilde{N}(\boldsymbol{x})G(\boldsymbol{x}, F, \mathrm{d}(\tau+s), \mathrm{d}\xi) = \widetilde{N}(\boldsymbol{y})G(\boldsymbol{y}, F, \mathrm{d}\tau, \mathrm{d}\xi)$$

as measures on $\mathrm{dec}(F)$, which is (3.7). If $s = 0$ then this equality is symmetric in $\boldsymbol{x}$ and $\boldsymbol{y}$ and must thus already hold for *all* pairs $\boldsymbol{x}, \boldsymbol{y} \in E_\circ^\Pi$. Then for $s > 0$, $\boldsymbol{x}, \boldsymbol{y} \in E_\circ^\Pi$, and $\boldsymbol{y}' \notin N_{\boldsymbol{x},s}$,

$$\mathrm{e}^{\beta s}\widetilde{N}(\boldsymbol{x})G(\boldsymbol{x}, F, \mathrm{d}(\tau+s), \mathrm{d}\xi) = \widetilde{N}(\boldsymbol{y}')G(\boldsymbol{y}', F, \mathrm{d}\tau, \mathrm{d}\xi) = \widetilde{N}(\boldsymbol{y})G(\boldsymbol{y}, F, \mathrm{d}\tau, \mathrm{d}\xi).$$

□

For the rest of this section, we fix an arbitrary choice for $\widetilde{N}_\Pi$, $\Pi \in \mathcal{P}$.

Lemma 3.7. *$\widetilde{N}_\Pi$ is continuous for every $\Pi \in \mathcal{P}$.*

Proof. Fix $\Pi \in \mathcal{P}$. First note that (3.6) implies integrability of $\widetilde{N}_\Pi$ over $E_\circ^\Pi$, because the LHS is integrable, and $\prod_{u\in\Pi} p_s(\boldsymbol{x}_u - \boldsymbol{y}_u)$ can be uniformly bounded below by a positive constant for fixed $s > 0$. Now suppose that $\boldsymbol{x}_n \to \boldsymbol{x}$ in $E_\circ^\Pi$, and define $f \in C_b(\mathcal{X})$ by $f(\Pi', \boldsymbol{x}) \coloneqq \mathbb{1}_{\{\Pi'=\Pi\}}$. Then $S_1 f \in C_b(\mathcal{X})$, and

$$(S_1 f)(\boldsymbol{x}_n) = \int_{E_\circ^\Pi} S_t((\Pi, \boldsymbol{x}_n), (\Pi, \mathrm{d}\boldsymbol{y})) = \mathrm{e}^{-\beta_\Pi}\frac{1}{\widetilde{N}_\Pi(\boldsymbol{x}_n)}\int \widetilde{N}_\Pi(\boldsymbol{y})\prod_{u\in\Pi} p_1(\boldsymbol{y}_u - \boldsymbol{x}_n(u))\,\mathrm{d}\boldsymbol{y}.$$

By dominated convergence and integrability of $\widetilde{N}_\Pi$, we can pull the limit $n \to \infty$ into the integral on the RHS, the LHS converges to $(S_1 f)(\boldsymbol{x})$ because $S_1 f$ is continuous, and all involved quantities are positive, so we must have $\widetilde{N}_\Pi(\boldsymbol{x}_n) \to \widetilde{N}_\Pi(\boldsymbol{x})$. □

Lemma 3.8. *Let $\Pi \in \mathcal{P}$. If $\beta_\Pi > 0$ then the trivial forest is impossible, and if $\beta_\Pi = 0$ then the trivial forest is the only possible forest and $\widetilde{N}_\Pi$ is constant.*

Proof. Fix $\Pi \in \mathcal{P}$ and $\boldsymbol{x} \in E_\circ^\Pi$, and let $F_0 \in \mathbb{F}(\Pi)$ be the trivial forest. If we integrate (3.6) over $\boldsymbol{y} \in E_\circ^\Pi$ and let $s \to \infty$, we obtain

$$\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0) = \lim_{s\to\infty}\frac{\mathrm{e}^{-\beta_\Pi s}}{\widetilde{N}_\Pi(\boldsymbol{x})}\int \widetilde{N}_\Pi(\boldsymbol{y})\prod_{u\in\Pi} p_s(\boldsymbol{x}_u - \boldsymbol{y}_u)\,\mathrm{d}\boldsymbol{y} = \frac{\mathbb{1}_{\{\beta_\Pi=0\}}}{\widetilde{N}_\Pi(\boldsymbol{x})}\int \widetilde{N}_\Pi(\boldsymbol{y})\,\mathrm{d}\boldsymbol{y}, \quad (3.18)$$

where we used that $p_s \to 1$ uniformly as $s \to \infty$. If $\beta_\Pi > 0$ then (3.18) implies $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F_0) = 0$ for all $\boldsymbol{x} \in E_\circ^\Pi$, so F_0 is impossible. Otherwise, $\boldsymbol{z} \mapsto \widetilde{N}_\Pi(\boldsymbol{z})\mathbb{P}^{\boldsymbol{z}}(\mathrm{Fr} = F_0)$ is constant and positive by Lemma 3.6, so (3.18) turns into

$$1 = \int \mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F_0)^{-1}\,\mathrm{d}\boldsymbol{y},$$

which implies that $\mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F_0) = 1$ for almost-all $\boldsymbol{y} \in E_\circ^\Pi$, in particular at least one $\boldsymbol{y}_0 \in E_\circ^\Pi$. Then $\mathbb{P}^{\boldsymbol{y}_0}(\mathrm{Fr} = F) = 0$ for every $F \in \mathbb{F}(\Pi) \setminus \{F_0\}$, so all non-trivial forests are impossible. Thus $\mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F_0) = 1$ for all $\boldsymbol{y} \in E_\circ^\Pi$, so $\widetilde{N}_\Pi(\boldsymbol{y}) = \widetilde{N}_\Pi(\boldsymbol{y})\mathbb{P}^{\boldsymbol{y}}(\mathrm{Fr} = F_0)$ is constant in $\boldsymbol{y} \in E_\circ^\Pi$. □

Lemmas 3.6 to 3.8 together imply Proposition 3.4.

Theorem 3.9. *There exists a unique family $\boldsymbol{\nu} \in R$ such that* (2.12) *holds for every $F \in \mathbb{F}$ and $\boldsymbol{x} \in E_\circ^{\mathrm{lf}(F)}$, that is*

$$P^{\boldsymbol{x}}(F, \mathrm{d}\tau, \mathrm{d}\xi) = \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f_{\boldsymbol{\nu}}(\tau, \xi \,|\, \boldsymbol{x}) \boldsymbol{\nu}_F(\mathrm{d}\xi)\, \mathrm{d}\tau$$

as measures on $\mathrm{dec}(F)$. *For all $\Pi \in \mathcal{P}$, $\beta_\Pi = |\nu_\Pi|$, and $\widetilde{N}_\Pi / N^{\boldsymbol{\nu}}$ is constant on $E_\circ^\Pi$.*

That is, there are choices for $\widetilde{N}_\Pi$, $\Pi \in \mathcal{P}$ such that $\widetilde{N}_\Pi = N^{\boldsymbol{\nu}}|_{E_\circ^\Pi}$. We first prove a helpful consequence of Proposition 3.4. If $F = \{\Pi_0^F, \Pi_1^F, \dots, \Pi_m^F\} \in \mathbb{F}$ with $m > 1$, $G \coloneqq \{\Pi_1^F, \dots, \Pi_m^F\}$, and $(\tau, \xi) \in \mathrm{dec}(G)$, $s > 0$, $\boldsymbol{z} \in E_\circ^{\Pi_1 \setminus \Pi}$, then we define $s\tau \in \mathrm{tm}(F)$ by $\Pi_0 \mapsto s$ and $\Pi_i \mapsto s + \tau(\Pi_i)$ for $i \in [m]$, and $\boldsymbol{z}\xi \in \mathrm{sp}(F)$ in the sense of (2.4). This defines a bijection

$$\begin{aligned} (0, \infty) \times E_\circ^{\Pi_1 \setminus \Pi} \times \mathrm{dec}(G) \quad &\to \quad \mathrm{dec}(F), \\ (s, \boldsymbol{z}, \tau, \xi) \quad &\mapsto \quad (s\tau, \boldsymbol{z}\xi). \end{aligned}$$

Lemma 3.10. *Let $F = \{\Pi, \Pi_1^F, \dots, \Pi_m^F\} \in \mathbb{F}$ be non-trivial, then*

$$H_F(\mathrm{d}\tau, \mathrm{d}\xi) \coloneqq \mathrm{e}^{\beta_\Pi \min \tau} \frac{\widetilde{N}_\Pi(\boldsymbol{x}) P^{\boldsymbol{x}}(F, \mathrm{d}\tau, \mathrm{d}\xi)}{\prod\limits_{\substack{u \in \Pi \\ \mathrm{pr}_F(u) \neq \emptyset}} p(\tau_{\mathrm{pr}_F(u)}, \xi_{\mathrm{pr}_F(u)} - \boldsymbol{x}_u)},$$

defines a finite measure on $\mathrm{dec}(F)$ *whose definition does not depend on $\boldsymbol{x} \in E_\circ^\Pi$. If $m = 1$, then there is a finite measure H°_{Π, Π_1^F} on $E_\circ^{\Pi_1^F \setminus \Pi}$ such that*

$$H_F(\mathrm{d}\tau, \mathrm{d}\xi) = H^\circ_{\Pi, \Pi_1^F}(\mathrm{d}\xi)\, \mathrm{d}\tau.$$

If $m > 1$, so $G \coloneqq \{\Pi_1^F, \dots, \Pi_m^F\}$ is non-trivial, then there exists a finite measure H_F° on $E_\circ^{\Pi_1^F \setminus \Pi} \times \mathrm{dec}(G)$ such that

$$H_F(\mathrm{d}(s\tau), \mathrm{d}(\boldsymbol{z}\xi)) = H_F^\circ(\mathrm{d}\boldsymbol{z}, \mathrm{d}\tau, \mathrm{d}\xi)\, \mathrm{d}s.$$

Proof. Abbreviate $\mathrm{pr} = \mathrm{pr}_F$ and $\mathrm{ch} = \mathrm{ch}_F$. If we multiply (3.7) with $\mathrm{e}^{\beta_\Pi \min \tau}$, we get

$$\mathrm{e}^{\beta_\Pi \min(\tau + s)} \frac{\widetilde{N}_\Pi(\boldsymbol{x}) P^{\boldsymbol{x}}(F, \mathrm{d}(\tau + s), \mathrm{d}\xi)}{\prod\limits_{\substack{u \in \Pi \\ \mathrm{pr}(u) \neq \emptyset}} p(\tau_{\mathrm{pr}(u)} + s, \xi_{\mathrm{pr}(u)} - \boldsymbol{x}_u)} = \mathrm{e}^{\beta_\Pi \min \tau} \frac{\widetilde{N}_\Pi(\boldsymbol{y}) P^{\boldsymbol{y}}(F, \mathrm{d}\tau, \mathrm{d}\xi)}{\prod\limits_{\substack{u \in \Pi \\ \mathrm{pr}(u) \neq \emptyset}} p(\tau_{\mathrm{pr}(u)}, \xi_{\mathrm{pr}(u)} - \boldsymbol{y}_u)} \tag{3.19}$$

for every $\boldsymbol{x}, \boldsymbol{y} \in E_\circ^\Pi$ and $s \geq 0$ as measures on $\mathrm{dec}(F)$. Putting $s = 0$ shows that the definition of H_F does not depend on $\boldsymbol{x}$, and in general we can rewrite (3.19) as

$$H_F(\mathrm{d}(\tau + s), \mathrm{d}\xi) = H_F(\mathrm{d}\tau, \mathrm{d}\xi), \qquad s \geq 0, \tag{3.20}$$

as measures on $\mathrm{dec}(F)$.

Now assume that $m = 1$, then $F = \{\Pi, \Pi_1^F\}$, and H_F is a finite measure on $(0,\infty) \times E_\circ^{\Pi_1^F \setminus \Pi}$ which is translation-invariant in its first component. For $A \subset E_\circ^{\Pi_1^F \setminus \Pi}$ measurable, $H_F(\cdot, A)$ is a translation-invariant measure on $(0,\infty)$, thus a multiple $H^\circ_{\Pi,\Pi_1^F}(A) \in [0,\infty]$ of Lebesgue measure. Then

$$H^\circ_{\Pi,\Pi_1^F}(A) = H_F((1,2) \times A), \qquad A \subset E^{\Pi_1^F \setminus \Pi} \text{ measurable},$$

which implies that H°_{Π,Π_1^F} is a finite measure. Then $H_F(\mathrm{d}s, \mathrm{d}\xi) = H^\circ_{\Pi,\Pi_1^F}(\mathrm{d}\xi)\,\mathrm{d}s$.

Now suppose that $m > 1$, then (3.20) implies that, for $A \subset E_\circ^{\Pi_1^F \setminus \Pi} \times \mathrm{dec}(G)$ measurable,

$$H_F\left(\{(s\tau, \boldsymbol{z}\xi) \in \mathrm{dec}(F) \colon s \in \,\cdot\,, (\boldsymbol{z}, (\tau, \xi)) \in A\}\right)$$

is a translation-invariant measure on $(0,\infty)$ and thus a multiple of Lebesgue measure, which implies with a similar argument that

$$H_F(\mathrm{d}(s\tau), \mathrm{d}(\boldsymbol{z}\xi)) = H_F^\circ(\mathrm{d}\boldsymbol{z}, \mathrm{d}\tau, \mathrm{d}\xi)\,\mathrm{d}s$$

for a finite measure H_F° on $E_\circ^{\Pi_1^F \setminus \Pi} \times \mathrm{dec}(G)$. □

Proof of Theorem 3.9. We use an induction over $|\operatorname{lf}(F)|$ to prove existence of $\boldsymbol{\nu}$, and comment on uniqueness at the end of the proof. If $\Pi \in \mathcal{P}$ with $|\Pi| = 1$ then the only $F \in \mathbb{F}(\Pi)$ is trivial and (2.12) holds automatically. Furthermore $\beta_\Pi = |\nu_\Pi| = 0$ by definition, and both $N^{\boldsymbol{\nu}}|_{E_\circ^\Pi}$ (for any $\boldsymbol{\nu} \in R$ by definition) and $\widetilde{N}_\Pi$ are constant (Lemma 3.8). Let $\Pi \in \mathcal{P}$ with $|\Pi| \geq 2$ and suppose that $\nu_{\Pi',\Pi''}$ for $\Pi'' < \Pi'$ with $|\Pi'| < |\Pi|$ have already been found such that (2.12) holds for all $F \in \mathbb{F}$ with $|\operatorname{lf}(F)| < |\Pi|$, and such that for all $\Pi' \in \mathcal{P}$ with $|\Pi'| < |\Pi|$, $\beta_{\Pi'} = |\nu_{\Pi'}|$, and $\widetilde{N}_{\Pi'}/N^{\boldsymbol{\nu}}$ is constant on $E_\circ^{\Pi'}$ (this is a well-defined statement because $N^{\boldsymbol{\nu}}$ on $E_\circ^{\Pi'}$ only depends on transition measures ν_{Π_1,Π_2} with $\Pi_1 \leq \Pi'$).

Note that it is now enough to find finite measures $\nu_{\Pi,\Pi'}$ for $\Pi' \in \mathcal{P}$ with $\Pi' < \Pi$ such that $|\nu_\Pi| = \beta_\Pi$ and, for some $c_\Pi > 0$,

$$P^{\boldsymbol{x}}(F, \mathrm{d}\tau, \mathrm{d}\xi) = \frac{c_\Pi}{\widetilde{N}_\Pi(\boldsymbol{x})} f_{\boldsymbol{\nu}}(\tau, \xi \,|\, \boldsymbol{x}) \boldsymbol{\nu}_F(\mathrm{d}\xi)\,\mathrm{d}\tau \tag{3.21}$$

as measures on $\mathrm{dec}(F)$ for every $F \in \mathbb{F}(\Pi)$ and $\boldsymbol{x} \in E_\circ^\Pi$. Indeed, the left-hand side defines a probability measure on $\mathrm{dec}(\mathbb{F}(\Pi))$, and $N^{\boldsymbol{\nu}}$ is, by definition, the correct normalisation, so $c_\Pi/\widetilde{N}_\Pi = 1/N^{\boldsymbol{\nu}}$ on $E_\circ^\Pi$.

If $\beta_\Pi = 0$, then all $F \in \mathbb{F}(\Pi)$ except for the trivial forest are impossible, and (2.12) holds for all $F \in \mathbb{F}(\Pi)$ by putting $\nu_{\Pi,\Pi'} = 0$ for all $\Pi' < \Pi$. Then $|\nu_\Pi| = 0 = \beta_\Pi$, and $\widetilde{N}_\Pi$ is constant (Lemma 3.8). Assume for the remainder of the inductive step that $\beta_\Pi > 0$, and let $F = \{\Pi, \Pi_1^F, \dots, \Pi_m^F\} \in \mathbb{F}$ with $m \geq 1$, for if F is trivial then both sides of (2.12) vanish automatically (given the success of the remaining inductive step, so that $|\nu_\Pi| = \beta_\Pi > 0$). If $m = 1$ and $\beta_{\Pi_1^F} > 0$, then $F = \{\Pi, \Pi_1^F\}$ is impossible, and both left- and right-hand side of (2.12) are zero (the latter because $|\nu_{\Pi_1^F}| = \beta_{\Pi_1^F} > 0$ by induction

hypothesis). If $\beta_{\Pi_1^F} = 0$, we let $\nu_{\Pi,\Pi_1^F} = c_\Pi^{-1} H^\circ_{\Pi,\Pi_1^F}$ for a $c_\Pi > 0$ that we specify later. Then rearranging the definition of H_F in Lemma 3.10 implies

$$P^{\boldsymbol{y}}(F, \mathrm{d}\tau, \mathrm{d}\xi) = \mathrm{e}^{-\beta_\Pi \tau} \frac{c_\Pi}{\widetilde{N}_\Pi(\boldsymbol{y})} \prod_{\substack{u \in \Pi \\ \mathrm{pr}(u) \neq \emptyset}} p(\tau_{\mathrm{pr}(u)}, \xi_{\mathrm{pr}(u)} - \boldsymbol{y}_u) \nu_{\Pi,\Pi_1^F}(\mathrm{d}\xi)\, \mathrm{d}\tau,$$

where we abbreviated $\mathrm{pr} = \mathrm{pr}_F$ and $\mathrm{ch} = \mathrm{ch}_F$. Assuming that $|\nu_\Pi| = \beta_\Pi$ after choosing the remaining transition measures and c_Π, this is equivalent to (3.21) (recall the relevant definitions from (2.9) to (2.11)).

Now suppose $m > 1$. The claim has already been proved for $G = \{\Pi_1^F, \dots, \Pi_m^F\}$ by induction hypothesis because $|\Pi_1^F| < |\Pi|$. If G is impossible then by (2.12) for G we must have $\beta_{\mathrm{rt}(G)} = \beta_{\mathrm{rt}(F)} > 0$ or $\nu_{\Pi_i^F,\Pi_{i+1}^F} = 0$ for some $i \in [m-1]$, and since F is also impossible, both left- and right-hand sides of (2.12) for F are zero. Suppose thus that G is possible, so $\beta_{\mathrm{rt}(F)} = 0$ and $\nu_{\Pi_i^F,\Pi_{i+1}^F} \neq 0$ for all $i \in [m-1]$. Let $\Pi_{\mathrm{m}} = \Pi \setminus \Pi_1^F$ and $\Pi_{\mathrm{n}} = \Pi_1^F \cap \Pi$, the leaf nodes that did and did not merge in the first transition, respectively. Rearranging the definition of H_F in Lemma 3.10 gives

$$\begin{aligned}
&P^{\boldsymbol{x}}(F, \mathrm{d}(s\tau), \mathrm{d}(\boldsymbol{z}\xi)) \\
&\quad = \frac{\mathrm{e}^{-\beta_\Pi s}}{\widetilde{N}_\Pi(\boldsymbol{x})} H_F(\mathrm{d}(s\tau), \mathrm{d}(\boldsymbol{z}\xi)) \Bigg[\prod_{\substack{u \in \Pi \\ \mathrm{pr}(u) \neq \emptyset}} p((s\tau)_{\mathrm{pr}(u)}, (\boldsymbol{z}\xi)_{\mathrm{pr}(u)} - \boldsymbol{x}_u) \Bigg] \\
&\quad = \frac{\mathrm{e}^{-\beta_\Pi s}}{\widetilde{N}_\Pi(\boldsymbol{x})} H^\circ_F(\mathrm{d}\boldsymbol{z}, \mathrm{d}\tau, \mathrm{d}\xi)\, \mathrm{d}s \Bigg[\prod_{\substack{u \in \Pi_{\mathrm{n}} \\ \mathrm{pr}(u) \neq \emptyset}} p(\tau_{\mathrm{pr}(u)} + s, \xi_{\mathrm{pr}(u)} - \boldsymbol{x}_u) \prod_{u \in \Pi_{\mathrm{m}}} p_s(\boldsymbol{x}_u - \boldsymbol{z}_{\mathrm{pr}(u)}) \Bigg].
\end{aligned} \tag{3.22}$$

We now evaluate

$$\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F, \mathrm{Tm} \in \mathrm{d}(s\tau), \mathrm{Sp} \in \mathrm{d}(\boldsymbol{z}\xi), \boldsymbol{X}_s|_{\Pi_{\mathrm{n}}} \in \mathrm{d}\boldsymbol{y}), \tag{3.23}$$

a sub-probability measure on $\mathrm{dec}(F) \times E_\circ^{\Pi_{\mathrm{n}}}$, in two ways. On the one hand, by (2.12) it equals

$$\begin{aligned}
&P^{\boldsymbol{x}}(F, \mathrm{d}(s\tau), \mathrm{d}(\boldsymbol{z}\xi)) K_{\boldsymbol{x}}((F, s\tau, z\xi), \{\boldsymbol{X}_s|_{\Pi_{\mathrm{n}}} \in \mathrm{d}\boldsymbol{y}\}) \\
&\quad = P^{\boldsymbol{x}}(F, \mathrm{d}(s\tau), \mathrm{d}(\boldsymbol{z}\xi)) \prod_{u \in \Pi_{\mathrm{n}}} p_s(\boldsymbol{x}_u - \boldsymbol{y}_u) \left[\frac{p(\tau_{\mathrm{pr}(u)}, \boldsymbol{y}_u - \xi_{\mathrm{pr}(u)})}{p(\tau_{\mathrm{pr}(u)} + s, \boldsymbol{x}_u - \xi_{\mathrm{pr}(u)})} \right]_{\mathrm{pr}(u) \neq \emptyset} \mathrm{d}\boldsymbol{y}.
\end{aligned} \tag{3.24}$$

Let $\tau_0 \coloneqq \inf\{t > 0 \colon \Pi_t \neq \Pi_0\}$, and $\tau_{\Pi_1^F} \coloneqq \inf\{t > 0 \colon \Pi_t = \Pi_1^F\}$, which are both $(\mathcal{F}_t^{\boldsymbol{X}})$-stopping times, and write

$$P^{\boldsymbol{x}}_{\Pi,\Pi_1^F}(\mathrm{d}s, \mathrm{d}z, \mathrm{d}\boldsymbol{y}) \coloneqq \mathbb{P}^{\boldsymbol{x}}(\tau_0 = \tau_{\Pi_1^F} \in \mathrm{d}s, \boldsymbol{X}_s \in \mathrm{d}(\boldsymbol{y}\boldsymbol{z})),$$

a sub-probability measure on $(0, \infty) \times E_\circ^\Pi$ whose total mass is the probability that the first jump is $\Pi \to \Pi_1^F$ when starting at $\boldsymbol{x}$. For the second evaluation, by the strong Markov

property at time $\tau_{\Pi_1^F}$, (3.23) equals $P^{\boldsymbol{x}}_{\Pi,\Pi_1^F}(\mathrm{d}s,\mathrm{d}\boldsymbol{z},\mathrm{d}\boldsymbol{y})P^{\boldsymbol{yz}}(G,\mathrm{d}\tau,\mathrm{d}\xi)$, and equating with (3.24) gives

$$\begin{aligned}P^{\boldsymbol{x}}(F,\mathrm{d}(s\tau),\mathrm{d}(\boldsymbol{z}\xi))\,\mathrm{d}\boldsymbol{y}\\ = \frac{P^{\boldsymbol{x}}_{\Pi,\Pi_1^F}(\mathrm{d}s,\mathrm{d}\boldsymbol{z},\mathrm{d}\boldsymbol{y})}{\prod_{u\in\Pi_\mathrm{n}} p_s(\boldsymbol{x}_u-\boldsymbol{y}_u)}P^{\boldsymbol{yz}}(G,\mathrm{d}\tau,\mathrm{d}\xi)\prod_{\substack{u\in\Pi_\mathrm{n}\\ \mathrm{pr}(u)\neq\emptyset}}\frac{p(\tau_{\mathrm{pr}(u)}+s,\boldsymbol{x}_u-\xi_{\mathrm{pr}(u)})}{p(\tau_{\mathrm{pr}(u)},\boldsymbol{y}_u-\xi_{\mathrm{pr}(u)})}.\end{aligned}$$

Combining with (3.22) gives, after some cancellation,

$$\begin{aligned}P^{\boldsymbol{x}}_{\Pi,\Pi_1^F}(\mathrm{d}s,\mathrm{d}\boldsymbol{z},\mathrm{d}\boldsymbol{y})P^{\boldsymbol{yz}}(G,\mathrm{d}\tau,\mathrm{d}\xi) = \frac{\mathrm{e}^{-\beta_\Pi s}}{\widetilde{N}_\Pi(\boldsymbol{x})}H^\circ_F(\mathrm{d}\boldsymbol{z},\mathrm{d}\tau,\mathrm{d}\xi)\,\mathrm{d}s\\ \times\prod_{u\in\Pi_\mathrm{n}}p_s(\boldsymbol{x}_u-\boldsymbol{y}_u)\prod_{u\in\Pi_\mathrm{m}}p_s(\boldsymbol{x}_u-\boldsymbol{z}_{\mathrm{pr}(u)})\prod_{\substack{u\in\Pi_\mathrm{n}\\ \mathrm{pr}(u)\neq\emptyset}}p(\tau_{\mathrm{pr}(u)},\boldsymbol{y}_u-\xi_{\mathrm{pr}(u)})\,\mathrm{d}\boldsymbol{y}.\end{aligned}$$

for fixed $\boldsymbol{x}\in E^\Pi_\circ$. By (2.12) which already holds for G, $P^{\boldsymbol{yz}}(G,\mathrm{d}\tau,\mathrm{d}\xi)$ has a strictly positive density $\frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{yz})}f_{\boldsymbol{\nu}}(\cdot\,|\,\boldsymbol{yz})|_{\mathrm{dec}(G)}$ w.r.t. the finite measure $\boldsymbol{\nu}_G(\mathrm{d}\xi)\,\mathrm{d}\tau$, so we can apply Lemma 3.11 below with $x=(\boldsymbol{z},\boldsymbol{y})$ and $\zeta=(\tau,\xi)$, to obtain existence of a measure $K(\mathrm{d}\boldsymbol{y},\mathrm{d}\boldsymbol{z})$ on $E^{\Pi_1^F}_\circ$ (which may depend on $\boldsymbol{x}$ and F) with

$$K(\mathrm{d}\boldsymbol{y},\mathrm{d}\boldsymbol{z})\,\mathrm{d}s = \mathrm{e}^{\beta_\Pi s}\widetilde{N}_\Pi(\boldsymbol{x})\frac{P^{\boldsymbol{x}}_{\Pi,\Pi_1^F}(\mathrm{d}s,\mathrm{d}\boldsymbol{y},\mathrm{d}\boldsymbol{z})}{\prod_{u\in\Pi_\mathrm{n}}p_s(\boldsymbol{x}_u-\boldsymbol{y}_u)\prod_{u\in\Pi_\mathrm{m}}p_s(\boldsymbol{x}_u-\boldsymbol{z}_{\mathrm{pr}(u)})} \tag{3.25}$$

as measures on $E^{\Pi_1^F}_\circ\times(0,\infty)$, and

$$K(\mathrm{d}\boldsymbol{y},\mathrm{d}\boldsymbol{z})\boldsymbol{\nu}_G(\mathrm{d}\xi)\,\mathrm{d}\tau = H^\circ_F(\mathrm{d}\boldsymbol{z},\mathrm{d}\tau,\mathrm{d}\xi)\frac{N^{\boldsymbol{\nu}}(\boldsymbol{yz})}{f_{\boldsymbol{\nu}}(\tau,\xi\,|\,\boldsymbol{yz})}\prod_{\substack{u\in\Pi_\mathrm{n}\\ \mathrm{pr}(u)\neq\emptyset}}p(\tau_{\mathrm{pr}(u)},\boldsymbol{y}_u-\xi_{\mathrm{pr}(u)})\,\mathrm{d}\boldsymbol{y}, \tag{3.26}$$

as measures on $E^{\Pi_1^F}_\circ\times\mathrm{dec}(G)$. From (3.25) it follows that K depends on F only through (Π,Π_1^F), and from (3.26) it follows that K does not depend on $\boldsymbol{x}$. In other words, had we chosen another $\boldsymbol{x}\in E^\Pi_\circ$ and another possible forest $F\in\mathbb{F}(\Pi)$ whose first transition is $\Pi\to\Pi_1^F$, we would have obtained the same measure $K=K_{\Pi,\Pi_1^F}$. It follows from (3.26) that K is absolutely continuous in its first component, that is there exist measures $K(\boldsymbol{y},\mathrm{d}\boldsymbol{z})$ on $E^{\Pi_1^F\setminus\Pi}_{\boldsymbol{y}}\coloneqq\{\boldsymbol{z}\in E^{\Pi_1^F\setminus\Pi}_\circ\colon\boldsymbol{yz}\in E^{\Pi_1^F}_\circ\}$ such that $\boldsymbol{y}\mapsto K(\boldsymbol{y},A)$ is measurable for measurable A, and such that

$$K(\boldsymbol{y},\mathrm{d}\boldsymbol{z})\boldsymbol{\nu}_G(\mathrm{d}\xi)\,\mathrm{d}\tau = H^\circ_F(\mathrm{d}\boldsymbol{z},\mathrm{d}\tau,\mathrm{d}\xi)\frac{N^{\boldsymbol{\nu}}(\boldsymbol{yz})}{f_{\boldsymbol{\nu}}(\tau,\xi\,|\,\boldsymbol{yz})}\prod_{\substack{u\in\Pi_\mathrm{n}\\ \mathrm{pr}(u)\neq\emptyset}}p(\tau_{\mathrm{pr}(u)},\boldsymbol{y}_u-\xi_{\mathrm{pr}(u)}) \tag{3.27}$$

as measures on $E^{\Pi_1^F\setminus\Pi}_{\boldsymbol{y}}\times\mathrm{dec}(G)$ for Lebesgue-almost all $\boldsymbol{y}\in E^{\Pi_\mathrm{n}}_\circ$. (The fact that we can choose a single null-set outside of which equality holds as measures is because the

σ-algebra of $E_{\boldsymbol{y}}^{\Pi_1^F \setminus \Pi} \times \mathrm{dec}(G)$ is countably generated.) If we rearrange (3.27) for H_F° and plug it into (3.22), we obtain

$$\begin{aligned} &P^{\boldsymbol{x}}(F, \mathrm{d}(s\tau), \mathrm{d}(\boldsymbol{z}\xi)) \\ &\quad = \frac{\mathrm{e}^{-\beta_\Pi s}}{\widetilde{N}_\Pi(\boldsymbol{x})} \Bigg[\prod_{\substack{u \in \Pi_\mathrm{n} \\ \mathrm{pr}(u) \neq \emptyset}} \frac{p(\tau_{\mathrm{pr}(u)}, \xi_{\mathrm{pr}(u)} - \boldsymbol{x}_u)}{p(\tau_{\mathrm{pr}(u)}, \xi_{\mathrm{pr}(u)} - \boldsymbol{y}_u)} \prod_{u \in \Pi_\mathrm{m}} p_s(\boldsymbol{x}_u - \boldsymbol{z}_{\mathrm{pr}(u)}) \Bigg] f_{\boldsymbol{\nu}}(\tau, \xi \,|\, \boldsymbol{y}\boldsymbol{z}) \\ &\qquad\qquad \times \frac{K(\boldsymbol{y}, \mathrm{d}\boldsymbol{z})}{N^{\boldsymbol{\nu}}(\boldsymbol{y}\boldsymbol{z})} \boldsymbol{\nu}_G(\mathrm{d}\xi)\, \mathrm{d}\tau\, \mathrm{d}s \end{aligned}$$

for Lebesgue-almost all $\boldsymbol{y}$. By the definition of $f_{\boldsymbol{\nu}}$ (2.10), this is equal to

$$P^{\boldsymbol{x}}(F, \mathrm{d}(s\tau), \mathrm{d}(\boldsymbol{z}\xi)) = \frac{1}{\widetilde{N}_\Pi(\boldsymbol{x})} f_{\boldsymbol{\nu}}(s\tau, \boldsymbol{z}\xi \,|\, \boldsymbol{x}) \frac{K(\boldsymbol{y}, \mathrm{d}\boldsymbol{z})}{N^{\boldsymbol{\nu}}(\boldsymbol{y}\boldsymbol{z})} \boldsymbol{\nu}_G(\mathrm{d}\xi)\, \mathrm{d}(s\tau). \tag{3.28}$$

Fix an arbitrary $\boldsymbol{y}_0 \in E_\circ^{\Pi_\mathrm{n}}$ for which this holds and put $\nu_{\Pi,\Pi_1^F}(\mathrm{d}\boldsymbol{z}) \coloneqq c_\Pi^{-1} \frac{K(\boldsymbol{y}_0, \mathrm{d}\boldsymbol{z})}{N^{\boldsymbol{\nu}}(\boldsymbol{y}_0 \boldsymbol{z})}$ for the same $c_\Pi > 0$ we introduced before. If we divide by $N^{\boldsymbol{\nu}}(\boldsymbol{y}_0 \boldsymbol{z})$ in (3.27) and integrate over $\xi \in \mathrm{sp}(G)$ and $\tau \in \{1 \leq \tau \leq 2\} \subset \mathrm{tm}(G)$ say it follows that ν_{Π,Π_1^F} is a finite measure. Now (3.28) evaluated for $\boldsymbol{y} = \boldsymbol{y}_0$ becomes

$$P^{\boldsymbol{x}}(F, \mathrm{d}(s\tau), \mathrm{d}(\boldsymbol{z}\xi)) = \frac{c_\Pi}{\widetilde{N}_\Pi(\boldsymbol{x})} f_{\boldsymbol{\nu}}(s\tau, \boldsymbol{z}\xi \,|\, \boldsymbol{x}) \boldsymbol{\nu}_F(\mathrm{d}(\boldsymbol{z}\xi))\, \mathrm{d}(s\tau),$$

which implies (3.21).

We have now defined a finite measure $\nu_{\Pi,\Pi'}$ for every $\Pi' \in \mathcal{P}$ with $\Pi' < \Pi$: if $\beta_{\Pi'} = 0$, then $\nu_{\Pi,\Pi'}$ is defined in the "$m = 1$ case". Otherwise, there exists a possible, non-trivial forest $G \in \mathbb{F}(\Pi')$ so that $\nu_{\Pi,\Pi'}$ is defined in the "$m > 1$ case", in which we explained why the definition of $\nu_{\Pi,\Pi'}$ in that case does not depend on the choice of $G \in \mathbb{F}(\Pi')$.

It remains to show that we can choose $c_\Pi > 0$ such that $|\nu_\Pi| = \beta_\Pi$. Since $\beta_\Pi > 0$ there must be at least one non-trivial, possible forest $F \in \mathbb{F}(\Pi)$, in which case (3.28) necessitates that the associated ν_{Π,Π_1^F} is non-zero, so $|\nu_\Pi| > 0$. Thus there exists a unique choice for $c_\Pi > 0$ (which by definition is a linear factor in every non-zero $\nu_{\Pi,\Pi'}$) for which $|\nu_\Pi| = \beta_\Pi$.

Finally, the construction of $\boldsymbol{\nu}$ satisfying (2.12) in this proof was unique. The sceptical reader can prove uniqueness of the transition measures $\boldsymbol{\nu}$ associated with a Brownian spatial coalescent directly with an induction of a similar structure to this proof. □

Lemma 3.11. *Suppose $\Omega_1, \Omega_2, \Omega_3$ are measurable spaces, and*

$$G(\mathrm{d}s, \mathrm{d}x) f(x, \zeta) \nu(\mathrm{d}\zeta) = g(s, x) F(\mathrm{d}x, \mathrm{d}\zeta) \mu(\mathrm{d}s),$$

where G and F are measures on $\Omega_1 \times \Omega_2$ and $\Omega_2 \times \Omega_3$ respectively, $f \colon \Omega_2 \times \Omega_3 \to \mathbb{R}$ and $g \colon \Omega_1 \times \Omega_2 \to \mathbb{R}$ are strictly positive measurable functions, and ν and μ are measures on Ω_3

and Ω_1 respectively, such that there exist $A \subset \Omega_3$ and $B \subset \Omega_1$ with $\mu(A), \nu(B) \in (0, \infty)$. Then there is a measure H on Ω_2 such that

$$\frac{G(\mathrm{d}s, \mathrm{d}x)}{g(s,x)} = H(\mathrm{d}x)\mu(\mathrm{d}s), \qquad \frac{F(\mathrm{d}\zeta, \mathrm{d}x)}{f(\zeta, x)}\nu(\mathrm{d}\zeta) = H(\mathrm{d}x)\nu(\mathrm{d}\zeta),$$

respectively as measures on $\Omega_1 \times \Omega_2$ and $\Omega_2 \times \Omega_3$.

Proof. Rewrite the equality as an equality of measures

$$\frac{G(\mathrm{d}s, \mathrm{d}x)}{g(s,x)}\nu(\mathrm{d}\zeta) = \frac{F(\mathrm{d}x, \mathrm{d}\zeta)}{f(x,\zeta)}\mu(\mathrm{d}s). \tag{3.29}$$

Let A, B be sets of positive and finite measure w.r.t. μ and ν respectively (by assumption there exists at least one each), and define $H(\mathrm{d}x)$ by integrating (3.29) over $s \in A, \zeta \in B$ and multiplying with $\frac{1}{\mu(A)\nu(B)}$, that is

$$H(\mathrm{d}x) \coloneqq \frac{1}{\mu(A)} \int_{s\in A} \frac{G(\mathrm{d}s, \mathrm{d}x)}{g(s,x)} = \frac{1}{\nu(B)} \int_{\zeta \in B} \frac{F(\mathrm{d}x, \mathrm{d}\zeta)}{f(x,\zeta)}.$$

Note H is independent of B by the first and A by the second equality. To prove the first equality of measures in the claim, take a set A with $\mu(A) \in (0, \infty)$ and $C \subset \Omega_2$. Then,

$$\int_{x\in C}\int_{s\in A} H(\mathrm{d}x)\mu(\mathrm{d}s) = H(C)\mu(A) = \mu(A)\int_C H(\mathrm{d}x) = \int_{x\in C}\int_{s\in A} \frac{G(\mathrm{d}s, \mathrm{d}x)}{g(s,x)}.$$

If A is a μ-null set, then by integrating (3.29) over it and some set of positive ν-measure we find that $\int_A \frac{G(\mathrm{d}s,\mathrm{d}x)}{g(s,x)} = 0$, so equality holds in this case as well. Since μ is σ-finite, we find equality for all measurable sets A. The second equality follows similarly. □

Given a Brownian spatial coalescent, say with transition measures $\boldsymbol{\nu} \in R$, recall that a forest $F \in \mathbb{F}$ is either possible, that is $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F) > 0$ for all $\boldsymbol{x} \in E_\circ^{\mathrm{lf}(F)}$, or impossible, that is $\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F) = 0$ for all such $\boldsymbol{x}$. We call a merge event (Π, Π') *possible* if there exists a possible forest $F = \{\Pi, \Pi', \ldots\} \in \mathbb{F}$, that is if for every $\boldsymbol{x} \in E_\circ^\Pi$ there is a positive probability that the first merge event is (Π, Π'). We collect the following obvious but useful statement.

Lemma 3.12. *A forest $F \in \mathbb{F}$ is possible if and only if $\nu_{\Pi,\Pi'} \neq 0$ for all $(\Pi, \Pi') \in F$ and $|\nu_{\mathrm{rt}(F)}| = 0$. A (Π, Π')-merger is possible if and only if $\nu_{\Pi,\Pi'} \neq 0$.*

In particular, if $|\nu_\Pi| > 0$ for every $\Pi \in \mathcal{P}$ with $|\Pi| \geq 2$, then the only possible forests are trees. This will be the case for sampling consistent Brownian spatial coalescents (except the trivial one), as we will see in Section 3.3.

3.2 Characterisation of Label Invariance

In preparation for the proof of Theorem 1.12, this section is devoted to a characterisation of label invariance for Brownian spatial coalescents.

For $(n, \vec{k}) \in M$ and $\Pi, \Pi' \in \mathcal{P}$ with $n = |\Pi|$ and $\Pi' < \Pi$, we call (Π, Π') an $(n, \vec{k})$-*merger* if Π' can be obtained from Π by coalescing m disjoint sets $I_1, \dots, I_m \subset \Pi$ of sizes k_1 to k_m into one each. A non-spatial coalescent with transition rates $\boldsymbol{\lambda} \in R^\circ$ is label invariant if and only if there are numbers $\lambda_{n,\vec{k}} \geq 0$ for $(n, \vec{k}) \in M$ such that $\lambda_{\Pi,\Pi'} = \lambda_{n,\vec{k}}$ whenever (Π, Π') is an $(n, \vec{k})$ merger (recall Definition 1.1 and the following paragraph). We show that a similar characterisation holds for Brownian spatial coalescents. For $m \in \mathbb{N}$ define

$$E_\circ^m = \left\{ \boldsymbol{x} \in E^m \colon\ [\{i\} \mapsto \boldsymbol{x}_i] \in E_\circ^{\{\{1\},\dots,\{m\}\}} \right\}.$$

In particular, $E_\circ^1 = E$. If $(n, k_1, \dots, k_m) \in M$, choose a bijection $\ell \colon [m] \to \Pi' \setminus \Pi$ such that $\ell(i)$ for $i \in [m]$ is the union of k_i distinct elements of Π, which is unique up to a permutation of pairs (i, j) with $k_i = k_j$. Define $\kappa_{\Pi,\Pi'} \colon E_\circ^{\Pi' \setminus \Pi} \to E_\circ^m$ by $\boldsymbol{x} \mapsto \boldsymbol{x} \circ \ell$ (here and in the following, we occasionally identify a vector $\boldsymbol{x} \in E^m$ with the map $[i \mapsto \boldsymbol{x}_i] \in E^{[m]}$). Recall the definition of $\mathcal{I}(\Pi_0, \Pi_1)$ from the paragraph preceding Definition 2.4. We further extend a map $\iota \in \mathcal{I}(\Pi_0, \Pi_1)$ to

$$\mathbb{F}(\Pi_0) \to \mathbb{F}(\Pi_1); \quad F = \{\Pi_0^F, \dots, \Pi_m^F\} \mapsto \{\iota(\Pi_0^F), \dots, \iota(\Pi_m^F)\},$$

and, for every $F \in \mathbb{F}(\Pi_0)$, to

$$\begin{aligned} \mathrm{sp}(F) &\to \mathrm{sp}(\iota(F)); \quad \xi \mapsto \xi \circ \iota^{-1}, \\ \mathrm{tm}(F) &\to \mathrm{tm}(\iota(F)); \quad \tau \mapsto \tau \circ \iota^{-1}. \end{aligned}$$

Finally, this lets us define ι on $\mathrm{dec}(\mathbb{F}(\Pi_0))$ by $(F, \tau, \xi) \mapsto (\iota(F), \iota(\tau), \iota(\xi))$.

Lemma 3.13. *The Brownian spatial coalescent with transition measures $\boldsymbol{\nu} \in R^\circ$ is label invariant if and only if for all $\Pi_0, \Pi_1 \in \mathcal{P}$ of equal size and $\iota \in \mathcal{I}(\Pi_0, \Pi_1)$,*

$$\forall \Pi' < \Pi_0 \colon \iota \# \nu_{\Pi_0,\Pi'} = \nu_{\iota(\Pi_0),\iota(\Pi')}.$$

In that case, there exists for every $(n, \vec{k}) = (n, k_1, \dots, k_m) \in M$ a finite measure $\nu_{n,\vec{k}}$ on $E_\circ^m$ such that for every $(n, \vec{k})$-merger (Π, Π'),

$$\kappa_{\Pi,\Pi'} \# \nu_{\Pi,\Pi'} = \nu_{n,\vec{k}}.$$

If $k_i = k_j$, then $\nu_{n,\vec{k}}$ is symmetric in the ith and jth coordinate.

The proof is straightforward and in Appendix II.C. Lemma 3.13 justifies that, for a label invariant Brownian spatial coalescent, we identify $\nu_{\Pi,\Pi'}$ for a $(n, \vec{k})$-merger (Π, Π') with $\nu_{n,\vec{k}}$, and by the symmetry of $\nu_{n,\vec{k}}$ the choice of the underlying map $\kappa_{\Pi,\Pi'}$ is irrelevant.

3.3 Characterisation of Sampling Consistency

In this section we prove Theorem 1.12. Throughout this section, fix a label invariant Brownian spatial coalescent, say with transition measures $\boldsymbol{\nu} \in R$. Recall in this context the measures $\nu_{n,\vec{k}}$, and write $|\nu_n| \coloneqq |\nu_\Pi|$ for any $\Pi \in \mathcal{P}$ with $|\Pi| = n$, and $|\nu_1| \coloneqq 0$. The following lemma is a core part of the proof. For two finite measures write $m_1 \sim m_2$ if $m_1 = cm_2$ for some $c > 0$.

Lemma 3.14. *Let $\boldsymbol{\nu}(\mathrm{d}\xi) = \boldsymbol{\lambda}\,\mathrm{d}\xi$ for some $\boldsymbol{\lambda} \in R^\circ$. Then the following are equivalent*

(i) The non-spatial coalescent with transition rates $\boldsymbol{\lambda}$ is sampling consistent.

(ii) $\int_{E_{\boldsymbol{x}}} N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y < \infty$ for every $\boldsymbol{x} \in \mathcal{X}$, and the Brownian spatial coalescent with transition measures $\boldsymbol{\nu}$ is sampling consistent w.r.t. the probability measures $(\mu_{\boldsymbol{x}})$ satisfying $\mu_{\boldsymbol{x}}(\mathrm{d}y) \sim N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y$.

In that case, $\int N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y = N^{\boldsymbol{\nu}}(\boldsymbol{x})$.

Lemma 3.14 is proved in Section 3.3.1, and implies the "if" direction of Theorem 1.12. For the "only if" direction it remains to show that sampling consistency of a Brownian spatial coalescent w.r.t. some family $(\mu_{\boldsymbol{x}})$ of probability measures implies $\boldsymbol{\nu}(\mathrm{d}\xi) \sim \mathrm{d}\xi$, and that $\mu_{\boldsymbol{x}}(\mathrm{d}y) \sim N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y$. That is the content of Section 3.3.2. In the remainder of this section, we set up the notation used in the proofs and make some technical observations related to sampling consistency.

Excluding from now the trivial case where $|\nu_n| = 0$ for all $n \in \mathbb{N}$ (so $\nu_{n,\vec{k}} = 0$ for all $(n, \vec{k})$), sampling consistency necessitates $|\nu_n| > 0$ for all $n \geq 2$, in particular that only trees are possible. This is part of Lemma 3.20 below. We now present a characterisation of all trees G from which a fixed tree F can be "subsampled" by removing one leaf. This characterisation will be at the core of our main arguments.

For every $\Pi_0 \in \mathcal{P}$ we fix some $v_\oplus \subset \mathbb{N}$ disjoint from all members of Π_0 and put $\Pi_0^\oplus \coloneqq \Pi_0 \cup \{v_\oplus\} \in \mathcal{P}$. For $F \in \mathbb{T}(\Pi_0)$ and $G \in \mathbb{T}(\Pi_0^\oplus)$ we say that G *extends* F if $F = G{\downarrow_{\Pi_0}} \coloneqq \{\Pi \setminus\setminus v_\oplus \colon \Pi \in G\}$, and write $\mathbb{T}^\oplus(F)$ for the set of such G. In that case, if $(\tau, \xi) \in \mathrm{dec}(G)$ then $\tau{\downarrow_F} \in \mathrm{tm}(F)$ and $\xi{\downarrow_F} \in \mathrm{sp}(F)$ denote the decorations induced on F, respectively defined by $\Pi \setminus\setminus v_\oplus \mapsto \min\{\tau(\Pi') \colon \Pi' \setminus\setminus v_\oplus = \Pi \setminus\setminus v_\oplus\}$ and $u \setminus v_\oplus \mapsto \xi_u$; write $\tau_\oplus \coloneqq \tau_{\mathrm{pr}_G(v_\oplus)}$ and $\xi_\oplus \coloneqq \xi_{\mathrm{pr}_G(v_\oplus)}$ for time and location of the first merge of $v_\oplus$, see Fig. 3.1.

If $G^\star = (G, \tau, \xi) \in \mathrm{dec}(\mathbb{T})$ and $G{\downarrow_{\Pi_0}} = F$ then we denote $G^\star{\downarrow_{\Pi_0}} = (G{\downarrow_{\Pi_0}}, \tau{\downarrow_F}, \xi{\downarrow_F})$. For $F^\star \in \mathrm{dec}(\mathbb{T})$ write $\mathrm{dec}^\oplus(F^\star) = \{G^\star \colon G^\star{\downarrow_{\Pi_0}} = F^\star\}$, and for fixed G extending F, $\mathrm{dec}^\oplus(F^\star \,|\, G) = \{(G, \tau, \xi) \in \mathrm{dec}^\oplus(F^\star)\}$. Write $\Pi_G^\oplus = \max\{\Pi \in G \colon v_\oplus \notin \Pi\}$ for the state of the tree immediately after the first merge involving $v_\oplus$, and $\Pi_F^\oplus \coloneqq \Pi_G^\oplus \setminus\setminus v_\oplus$ for the same state projected on Π_0. If G extends F, it has to fall within exactly one of the following three classes (see also Fig. 3.1).

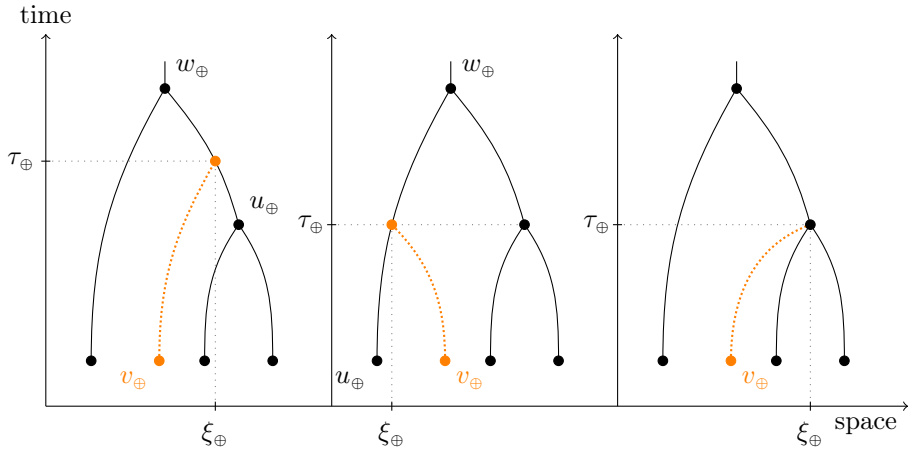

Figure 3.1: There are three ways to extend a fixed tree to an additional leaf. Either the leaf $v_\oplus$ merges binary with another node $u_\oplus$, at a time at which no merge happens in the original tree (left), or simultaneously with an existing merge event (middle); or $v_\oplus$ joins an existing merge event (right). The additional degrees of freedom in the tree decoration given that of the underlying tree are, from right to left, none, the location $\xi_\oplus$, and the time $\tau_\oplus$ and location $\xi_\oplus$ of the additional merge event.

(i) ***Multiple merge***: $v_\oplus$ is part of a multiple merger (Figure 3.1 right). In this case we can identify $\mathrm{dec}(G)$ and $\mathrm{dec}(F)$, so $\mathrm{dec}^\oplus(F^\star \,|\, G)$ is a singleton.

(ii) ***Binary merge***: $v_\oplus$ merges binary with a node $u_\oplus \in \mathrm{nd}(G)$, and no other merge takes place simultaneously (Figure 3.1 left). In this case, for $F^\star = (F, \tau_0, \xi_0) \in \mathrm{dec}(\mathbb{F})$ there is a bijection

$$\begin{array}{rcl} \mathrm{dec}^\oplus(F^\star \,|\, G) & \longleftrightarrow & (\tau_0(\Pi_F^\oplus), \tau_0(\Pi_F^\oplus)^+) \times E, \\ (\tau, \xi) & \mapsto & (\tau_\oplus, \xi_\oplus), \end{array} \tag{3.30}$$

whose inverse we denote $(s, z) \mapsto (s^\oplus \tau_0, z^\oplus \xi_0)$. In this case $w_\oplus := \mathrm{pr}_G(u_\oplus \cup v_\oplus) \neq \emptyset$ except if the merge event involving $v_\oplus$ is the final one, that is $|\Pi_G^\oplus| = 1$.

(iii) ***Simultaneous binary merge***: $v_\oplus$ merges binary with a node $u_\oplus \in \mathrm{nd}(G)$ as part of a simultaneous merge event (Figure 3.1 middle). Then $w_\oplus := \mathrm{pr}_G(u_\oplus \cup v_\oplus) \neq \emptyset$, and we can identify $\mathrm{tm}(G)$ and $\mathrm{tm}(F)$, and for $F^\star = (F, \tau_0, \xi_0) \in \mathrm{dec}(\mathbb{F})$ there is a bijection

$$\mathrm{dec}^\oplus(F^\star \,|\, G) \longleftrightarrow E_{\xi_0|_{\Pi_F^\oplus}}; \quad (\tau, \xi) \mapsto \xi_\oplus, \tag{3.31}$$

(recall (2.5),) the inverse of which we denote $z \mapsto (\tau_0, z^\oplus \xi_0)$.

Denote the sets of such trees by $\mathbb{T}_{\mathrm{m}}^\oplus(F)$, $\mathbb{T}_{\mathrm{b}}^\oplus(F)$, and $\mathbb{T}_{\mathrm{sb}}^\oplus(F)$, respectively. We do not explicitly denote dependence of $u_\oplus$ and $w_\oplus$ etc. on G, but it will always be clear from context.

This allows us to formulate sampling consistency on a more precise technical level. If all transition measures are absolutely continuous (which we will prove is a consequence

of sampling consistency in Lemma 3.24), then with a slight abuse of notation, we write $\nu_{n,\vec{k}}(\mathrm{d}\boldsymbol{z}) = \nu_{n,\vec{k}}(\boldsymbol{z})\,\mathrm{d}\boldsymbol{z}$ and $\boldsymbol{\nu}_F(\mathrm{d}\xi) = \boldsymbol{\nu}_F(\xi)\,\mathrm{d}\xi$ etc., as well as $P^{\boldsymbol{x}}(F^\star) \coloneqq P^{\boldsymbol{x}}(F,\tau,\xi) \coloneqq \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f_{\boldsymbol{\nu}}(\tau,\xi \,|\, \boldsymbol{x})\boldsymbol{\nu}_F(\xi)$, so that $P^{\boldsymbol{x}}(F,\mathrm{d}\tau,\mathrm{d}\xi) = P^{\boldsymbol{x}}(F,\tau,\xi)\,\mathrm{d}\tau\,\mathrm{d}\xi$.

Lemma 3.15. *Let $\boldsymbol{\nu} \in R$ be a non-trivial family of transition measures that are absolutely continuous w.r.t. Lebesgue measure. Then the associated coalescent process is sampling consistent w.r.t. a family of probability measures $(\mu_{\boldsymbol{x}})_{\boldsymbol{x}\in\mathcal{X}}$ if and only if $|\nu_n| > 0$ for all $n \geq 2$, and for every $\Pi_0 \in \mathcal{P}$, $F \in \mathbb{T}(\Pi_0)$, $\boldsymbol{x} \in E_\circ^{\Pi_0}$, and Lebesgue-a.e. $F^\star = (F,\tau,\xi) \in \mathrm{dec}(F)$,*

$$\begin{aligned} K_{\boldsymbol{x}}(F^\star,\cdot)P^{\boldsymbol{x}}(F,\tau,\xi) = K_{\boldsymbol{x}}(F^\star,\cdot)&\sum_{G\in\mathbb{T}_m^\oplus(F)} \int P^{\boldsymbol{x}\boldsymbol{y}}(G,\tau,\xi)\mu_{\boldsymbol{x}}(\mathrm{d}y) \\ &+ \sum_{G\in\mathbb{T}_b^\oplus(F)} \iiint K_{\boldsymbol{x}\boldsymbol{y}}((G,s^\oplus\tau,z^\oplus\xi),\{\mathring{\boldsymbol{X}}\in\cdot\})P^{\boldsymbol{x}\boldsymbol{y}}(G,s^\oplus\tau,z^\oplus\xi)\,\mathrm{d}s\,\mathrm{d}z\mu_{\boldsymbol{x}}(\mathrm{d}y) \\ &+ \sum_{G\in\mathbb{T}_{sb}^\oplus(F)} \iint K_{\boldsymbol{x}\boldsymbol{y}}((G,\tau,z^\oplus\xi),\{\mathring{\boldsymbol{X}}\in\cdot\})P^{\boldsymbol{x}\boldsymbol{y}}(G,\tau,z^\oplus\xi)\,\mathrm{d}z\mu_{\boldsymbol{x}}(\mathrm{d}y), \end{aligned} \tag{3.32}$$

where $(\mathring{\boldsymbol{X}}_t) \coloneqq (\boldsymbol{X}_t \setminus\!\setminus v_\oplus)$, and domains of integration for s and z are as indicated by the bijections (3.30) *and* (3.31).

Proof. The fact that sampling consistency implies $|\nu_n| > 0$ for all $n \geq 2$ is in Lemma 3.20 below. Eq. (3.32) is a direct reformulation of (2.6) using (2.8), the characterisation of $\mathbb{T}^\oplus(F)$ for $F \in \mathbb{T}$, and in order to get an equation pointwise for a.e. $F^\star$, the fact that $K_{\boldsymbol{x}}(F^\star,\{\mathrm{Dec}\in A\}\cap(\cdot)) = \mathbb{1}_{\{F^\star\in A\}}K_{\boldsymbol{x}}(F^\star,\cdot)$ for all $F^\star \in \mathrm{dec}(\mathbb{T})$ and measurable $A \subset \mathrm{dec}(F)$. That it suffices to consider trees is because $|\nu_n| > 0$ for $n \geq 2$ makes all forests that are not trees impossible. □

Before moving on to the proof of Lemma 3.14, we give an explicit description of the law of a non-spatial coalescent when regarded as a random, time decorated forest, which will make it easier to connect with the laws $(P^{\boldsymbol{x}})$ of a Brownian spatial coalescent. It follows from a simple argument about competing exponential clocks.

Lemma 3.16. *Let $\boldsymbol{\lambda} \in R^\circ$ be the transition rates of a non-spatial coalescent $(\mathbb{P}^\Pi)_{\Pi\in\mathcal{P}}$, and denote the total jump rate while at Π by $\lambda_\Pi \geq 0$. We define $f_{\mathrm{tm}}^{\boldsymbol{\lambda}}\colon \mathrm{tm}(\mathbb{F}) \to (0,\infty)$ by $f_{\mathrm{tm}}^{\boldsymbol{\lambda}}\big|_{\mathrm{tm}(F)} \equiv 0$ if $\lambda_{\mathrm{rt}(F)} > 0$, otherwise*

$$f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(F,\tau) \coloneqq \prod_{(\Pi,\Pi')\in F} \lambda_{\Pi,\Pi'}\mathrm{e}^{-\lambda_\Pi(\tau_{\Pi'}-\tau_\Pi)}, \qquad \tau \in \mathrm{tm}(F),$$

which is 1 if F is trivial. Then for $\Pi \in \mathcal{P}$ and $F \in \mathbb{F}(\Pi)$,

$$(\mathrm{Tm}\,\#\mathbb{P}^\Pi)(F,\mathrm{d}\tau) = f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(F,\tau)\,\mathrm{d}\tau.$$

3.3.1 Proof of Lemma 3.14

In this section we fix label invariant rates $\boldsymbol{\lambda} \in R^\circ$ and define $\boldsymbol{\nu} \in R$ by $\nu_{n,\vec{k}}(\mathrm{d}\boldsymbol{z}) \coloneqq \lambda_{n,\vec{k}}\,\mathrm{d}\boldsymbol{z}$. Then the laws $(P^{\boldsymbol{x}})$ that characterise the Brownian spatial coalescent process with transition measures $\boldsymbol{\nu}$ (recall (2.12)) take the simpler form

$$P^{\boldsymbol{x}}(F,\tau,\xi) = \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(F,\tau) f^F_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x}), \tag{3.33}$$

where in this chapter we will explicitly denote dependence of $f^{\boldsymbol{\lambda}}_{\mathrm{tm}}$ and f_{sp} on F. We further remark that a characterisation similar to Lemma 3.15 holds for non-spatial coalescent processes, which are sampling consistent with transition rates $\boldsymbol{\lambda}$ if and only if $\lambda_n > 0$ for all $n \geq 2$ and

$$f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(F,\tau) = \sum_{G \in \mathbb{T}^{\oplus}_{\mathrm{b}}(F)} \int f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(G, s^{\oplus}\tau)\,\mathrm{d}s + \sum_{G \in \mathbb{T}^{\oplus}_{\mathrm{sb}}(F)} f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(G,\tau) + \sum_{G \in \mathbb{T}^{\oplus}_{\mathrm{m}}(F)} f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(G,\tau) \tag{3.34}$$

for every $\Pi_0 \in \mathcal{P}$, $F \in \mathbb{T}(\Pi_0)$, and Lebesgue-a.e. $\tau \in \mathrm{tm}(F)$, where the RHS is simply a density of $\mathbb{P}^{\Pi^{\oplus}_0}(\{(G,\tau')\colon G{\downarrow_{\Pi_0}} = F, \tau'{\downarrow_F} \in \cdot\})$ on $\mathrm{tm}(F)$.

Lemma 3.17. *If the non-spatial coalescent with transition rates $\boldsymbol{\lambda}$ is sampling consistent, then $\int N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y = N^{\boldsymbol{\nu}}(\boldsymbol{x})$ for all $\boldsymbol{x} \in \mathcal{X}$, and the Brownian spatial coalescent with transition measures $\boldsymbol{\nu}$ is sampling consistent w.r.t. the probability measures defined by $\mu_{\boldsymbol{x}} = \frac{N^{\boldsymbol{\nu}}(\boldsymbol{x}y)}{N^{\boldsymbol{\nu}}(\boldsymbol{x})}\,\mathrm{d}y$.*

Proof. In that case $|\nu_n| = \lambda_n > 0$ for $n \geq 2$, so the assumptions of Lemma 3.15 are satisfied and we have to confirm (3.32). Note that passing to the total mass in (3.32) and then integrating over $(\tau,\xi) \in \mathrm{dec}(F)$ gives 1 on both sides (it equates to taking the total mass in (2.6)). In particular, if we prove (3.32) for any family of measures, in this case $\mu_{\boldsymbol{x}}(\mathrm{d}y) = \frac{N^{\boldsymbol{\nu}}(\boldsymbol{x}y)}{N^{\boldsymbol{\nu}}(\boldsymbol{x})}\,\mathrm{d}y$, then they must already be probability measures (in particular finite), that is $\int N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y = N^{\boldsymbol{\nu}}(\boldsymbol{x})$. Let $\Pi_0 \in \mathcal{P}$, $F^\star = (F,\tau,\xi) \in \mathrm{dec}(\mathbb{T}(\Pi_0))$ and $\boldsymbol{x} \in E^{\Pi}_{\circ}$. The laws on both sides of (3.32) are determined by their pushforward under the collection of maps $(\mathrm{Path}^{F^\star}_u)_{u \in \mathrm{nd}(F)}$, which are an independent family with respect to both. Hence it suffices to show that (3.32) holds when applying a pushforward with $\mathrm{Path}_u \coloneqq \mathrm{Path}^{F^\star}_u$ for a fixed but arbitrary $u \in \mathrm{nd}(F)$. Assume u is not the root, so $w \coloneqq \mathrm{pr}_F(u) \neq \emptyset$, otherwise the proof is similar. Then $\mathrm{Path}_u \# K_{\boldsymbol{x}}(F^\star, \cdot) = B^{(\tau_u,(\boldsymbol{x}\xi)_u)\to(\tau_w,\xi_w)}$. If $G^\star = (G, s^{\oplus}\tau, z^{\oplus}\xi) \in \mathbb{T}^{\oplus}_{\mathrm{b}}(F)$ with $u_{\oplus} = u$ then

$$\mathrm{Path}_u \# K_{\boldsymbol{x}y}(G^\star, \{\mathring{\boldsymbol{X}} \in \cdot\}) = B^{(\tau_u,(\boldsymbol{x}\xi)_u)\to(s,z)\to(\tau_w,\xi_w)},$$

and

$$\begin{aligned}
&\iiint B^{(\tau_u,(\boldsymbol{x}\xi)_u)\to(s,z)\to(\tau_w,\xi_w)} P^{\boldsymbol{x}y}(G, s^{\oplus}\tau, z^{\oplus}\xi)\,\mathrm{d}s\,\mathrm{d}z\mu_{\boldsymbol{x}}(\mathrm{d}y) \\
&= \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})}\int f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(G, s^{\oplus}\tau)\left(\iint B^{(\tau_u,(\boldsymbol{x}\xi)_u)\to(s,z)\to(\tau_w,\xi_w)}(\cdot) f^{G}_{\mathrm{sp}}(z^{\oplus}\xi \mid s^{\oplus}\tau, \boldsymbol{x}y)\,\mathrm{d}z\,\mathrm{d}y\right)\mathrm{d}s \qquad (3.35)\\
&= \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})}\left(\int f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(G, s^{\oplus}\tau)\,\mathrm{d}s\right) B^{(\tau_u,(\boldsymbol{x}\xi)_u)\to(\tau_w,\xi_w)}(\cdot) f^{F}_{\mathrm{sp}}(\xi \mid \tau, \boldsymbol{x}),
\end{aligned}$$

where the second equality followed directly from Lemma 3.22 because

$$f^{G}_{\mathrm{sp}}(z^{\oplus}\xi \mid s^{\oplus}\tau, \boldsymbol{x}y) = f^{F}_{\mathrm{sp}}(\xi \mid \tau, \boldsymbol{x})\frac{p_s(y-z)p_{s-\tau_u}(z-(\boldsymbol{x}\xi)_u)p_{\tau_w-s}(\xi_w-z)}{p_{\tau_w-\tau_u}(\xi_w-(\boldsymbol{x}\xi)_u)}. \qquad (3.36)$$

If $u_{\oplus} \neq u$ then the pushforward of $K_{\boldsymbol{x}y}(G^{\star}, \{\mathring{\boldsymbol{X}} \in \cdot\})$ under Path_u is the same as that of $K_{\boldsymbol{x}}(F^{\star}, \cdot)$, and then

$$\iiint P^{\boldsymbol{x}y}(G, s^{\oplus}\tau, z^{\oplus}\xi)\,\mathrm{d}s\,\mathrm{d}z\mu_{\boldsymbol{x}}(\mathrm{d}y) = \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f^{F}_{\mathrm{sp}}(\xi \mid \tau, \boldsymbol{x})\int f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(G, s^{\oplus}\tau)\,\mathrm{d}s$$

follows from (3.35). Similar calculations can be made for $G \in \mathbb{T}^{\oplus}_{\mathrm{sb}}(F)$, and if $G \in \mathbb{T}^{\oplus}_{\mathrm{m}}(F)$ then $f^{G}_{\mathrm{sp}}(\xi \mid \tau, \boldsymbol{x}y) = f^{F}_{\mathrm{sp}}(\xi \mid \tau, \boldsymbol{x})p_{\tau_{\oplus}}(\xi_{\oplus} - y)$ so

$$\int P^{\boldsymbol{x}y}(G, \tau, \xi)\mu_{\boldsymbol{x}}(\mathrm{d}y) = \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(G, \tau) f^{F}_{\mathrm{sp}}(\xi \mid \tau, \boldsymbol{x}).$$

We can thus cancel $\frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f^{F}_{\mathrm{sp}}(\xi \mid \tau, \boldsymbol{x}) B^{(\tau_u,(\boldsymbol{x}\xi)_u\to(\tau_w,\xi_w)}$ to conclude that the pushforward of (3.32) under any $u \in \mathrm{nd}(F)$ is equivalent exactly to (3.34) and thus implied by sampling consistency of the non-spatial coalescent with transition rates $\boldsymbol{\lambda}$. □

Lemma 3.18. *If $\int N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y < \infty$ for all $\boldsymbol{x} \in \mathcal{X}$, and the Brownian spatial coalescent with transition measures $\boldsymbol{\nu}(\mathrm{d}\xi) = \boldsymbol{\lambda}\,\mathrm{d}\xi$ is sampling consistent w.r.t. the probability measures $(\mu_{\boldsymbol{x}})$ defined by $\mu_{\boldsymbol{x}}(\mathrm{d}y) \sim N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y$, then the non-spatial coalescent with transition rates $\boldsymbol{\lambda}$ is sampling consistent and $\int N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y = N^{\boldsymbol{\nu}}(\boldsymbol{x})$.*

Proof. The assumptions of Lemma 3.15 are satisfied, so (3.32) holds with

$$\mu_{\boldsymbol{x}}(\mathrm{d}y) = \frac{N^{\boldsymbol{\nu}}(\boldsymbol{x}y)}{\int N^{\boldsymbol{\nu}}(\boldsymbol{x}y')\,\mathrm{d}y'}\,\mathrm{d}y, \quad \boldsymbol{x} \in \mathcal{X}.$$

By the same calculations made in the proof of Lemma 3.17, this implies

$$\begin{aligned}
&\frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(F, \tau) \\
&= \frac{1}{\int N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y}\left(\sum_{G\in\mathbb{T}^{\oplus}_{\mathrm{b}}(F)}\int f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(G, s^{\oplus}\tau)\,\mathrm{d}s + \sum_{G\in\mathbb{T}^{\oplus}_{\mathrm{sb}}(F)} f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(G, \tau) + \sum_{G\in\mathbb{T}^{\oplus}_{\mathrm{m}}(F)} f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(G, \tau)\right) \qquad (3.37)
\end{aligned}$$

for every $\Pi_0 \in \mathcal{P}$, $F \in \mathbb{T}(\Pi_0)$, and $\tau \in \mathrm{tm}(F)$. Integrating out $\tau \in \mathrm{tm}(F)$ gives $N^{\boldsymbol{\nu}}(\boldsymbol{x}) = \int N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y$, and then (3.37) turns into (3.34) which implies sampling consistency of the non-spatial coalescent with transition rates $\boldsymbol{\lambda}$. □

Lemmas 3.17 and 3.18 together imply Lemma 3.14.

3.3.2 Consequences of Sampling Consistency

In this section we fix some label invariant Brownian spatial coalescent, say with transition measures $\boldsymbol{\nu} \in R$, and assume it is sampling consistent w.r.t. some family of probability measures $(\mu_{\boldsymbol{x}})_{\boldsymbol{x}\in\mathcal{X}}$. The goal is to prove $\boldsymbol{\nu}(\mathrm{d}\xi) \sim \mathrm{d}\xi$, and that $\mu_{\boldsymbol{x}}(\mathrm{d}y) \sim N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y$ for every $\boldsymbol{x} \in \mathcal{X}$.

We write $(n, k_1, \ldots, k_m) \le (n', k'_1, \ldots, k'_{m'})$ if there is a way to sample n particles out of the n' particles in an $(n', \vec{k}')$-merger restricted to which we observe an $(n, \vec{k})$-merger. That is, if $n \le n'$ and $m \le m'$ and there exists an assignment of n particles to buckets of sizes $k'_0, k'_1, \ldots, k_{m'}$ (where $k'_0 = n' - \sum k'_i$) such that the numbers of particles in buckets 1 to m' have sizes k_1 to k_m (ignoring empty buckets, and without maintaining order).

Lemma 3.19. *(i) If $(n, \vec{k}) \le (n', \vec{k}')$ and $(n', \vec{k}')$ is possible, then so is $(n, \vec{k})$.*

(ii) If $(n, \vec{k})$ is possible and $n' > n$, then there exists $\vec{k}'$ such that $(n', \vec{k}') \ge (n, \vec{k})$ and $(n', \vec{k}')$ is possible.

Proof. Suppose $(n, \vec{k}) \le (n', \vec{k}')$ and the latter is possible. Let $\Pi, \Pi' \in \mathcal{P}$ with $|\Pi| = n$, $|\Pi'| = n'$ and $\Pi \subset \Pi'$. There exists a possible forest $F' \in \mathbb{F}(\Pi')$ that starts with an $(n', \vec{k}')$-merger such that the forest $F \coloneqq F'{\downarrow_\Pi}$ induced on Π starts with an $(n, \vec{k})$-merger. It suffices to show that F is possible. By sampling consistency, for some fixed $\boldsymbol{x} \in E_\circ^\Pi$ there is a probability measure μ on $E_\circ^{\Pi'\setminus\Pi}$ such that

$$\mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F) = \int \mathbb{P}^{\boldsymbol{xz}}(\mathrm{Fr}\,|_\Pi = F)\mu(\mathrm{d}\boldsymbol{z}) \ge \int \mathbb{P}^{\boldsymbol{xz}}(\mathrm{Fr} = F')\mu(\mathrm{d}\boldsymbol{z}) > 0,$$

so F is possible. Now suppose $(n, \vec{k})$ is possible and $n' > n$. Let $\Pi, \Pi' \in \mathcal{P}$ with $|\Pi| = n$, $|\Pi'| = n'$ and $\Pi \subset \Pi'$, and let $F \in \mathbb{F}(\Pi)$ such that F is possible and it starts with an $(n, \vec{k})$-merger. Then

$$0 < \mathbb{P}^{\boldsymbol{x}}(\mathrm{Fr} = F) = \int \mathbb{P}^{\boldsymbol{xz}}(\mathrm{Fr}\,|_\Pi = F)\mu(\mathrm{d}\boldsymbol{z}) \le \sum_{\substack{F'\in\mathbb{F}(\Pi')\\ F'|_\Pi = F}} \int \mathbb{P}^{\boldsymbol{xz}}(\mathrm{Fr} = F')\mu(\mathrm{d}\boldsymbol{z}).$$

Thus there exists a possible forest $F' \in \mathbb{F}(\Pi')$ with $F'|_\Pi = F$, whose first merge event satisfies $(n', \vec{k}') \ge (n, \vec{k})$. □

Lemma 3.20. *Exactly one of the following hold.*

(i) All mergers are impossible.

(ii) All of the mergers (n, n) for $n \ge 2$ are possible and all other mergers are impossible.

(iii) For every $n \ge 3$ there exists a possible $(n, \vec{k})$-merger with $\sum_i k_i < n$, and $(2, 2)$ is possible.

In cases (ii) and (iii), $|\nu_n| > 0$ for all $n \geq 2$, in particular the only possible forests are trees.

The extremal cases (i) and (ii) correspond in the non-spatial setting to the Λ-coalescent with Λ the zero-measure, and a Dirac mass at 1, respectively. Lemma 3.20 states that if neither (i) nor (ii) hold, then for every initial condition with at least three particles there is a positive probability that the first merge event leaves at least one particle untouched.

Proof of Lemma 3.20. Suppose that every merger that isn't of the form (n, n) for $n \geq 2$ is impossible. Then either (i) holds, or there exists an $n \geq 2$ such that (n, n) is possible. Then $(k, k) \leq (n, n)$ is possible for all $2 \leq k < n$. If $n' > n$, then $(n', \vec{k}')$ is possible for some $\vec{k}'$, but then it must be of the form (n', n'), so in fact (n, n) is possible for all $n \geq 2$ and all other merge events are impossible, so (ii) holds. In particular $|\nu_n| > 0$ for all $n \geq 2$.

Now assume that there exists $n_0 \geq 3$ such that a merge event $(n_0, \vec{k}_0) \neq (n_0, n_0)$ is possible, and we want to show (iii) holds. Then $(2, 2) \leq (n_0, \vec{k}_0)$ is possible, and for any $n > n_0$, by Lemma 3.19(ii) there exists a possible merge event $(n, \vec{k}) \geq (n_0, \vec{k}_0)$, which cannot be (n, n) because otherwise $(n_0, \vec{k}_0) = (n_0, n_0)$. Thus if $n \geq 3$ arbitrary, we can take $N \coloneqq (3n) \vee n_0$ and a possible merge event $(N, k_1, \ldots, k_m) \neq (N, N)$. If $m = 1$, then $k_1 < N$ so $(n, n-1) \leq (N, \vec{k})$ is possible by Lemma 3.19(i). If $m \geq 2$, then we construct a merger $(n, \vec{k}') \leq (N, \vec{k})$ with $\sum_i k'_i < n$ by assigning one particle to the smallest bucket of size k_m, and spread all other particles over the remaining buckets, which works as long as $n - 1 \leq N - k_1$, but if this wasn't true then $N \geq mk_1 \geq 2(N - n + 1)$ and thus $n \geq N/2$, a contradiction. □

We assume from now that we are in case (iii). In case (i) there is nothing to show, and in case (ii) proofs are easier than and can be directly adapted from those that follow. We will from now explicitly denote dependence of $f_{\boldsymbol{\nu}}$ on F by writing $f^F_{\boldsymbol{\nu}}$. The following lemma serves as a start to an inductive argument.

Lemma 3.21. *There is $\lambda_{2,2} > 0$ such that $\nu_{2,2}(\mathrm{d}z) = \lambda_{2,2}\,\mathrm{d}z$, and for every $x \in E$,*

$$\mu_x(\mathrm{d}y) \sim N^{\boldsymbol{\nu}}(xy)\,\mathrm{d}y.$$

A central ingredient in this and many other proofs in this section is the observation stated in the following lemma. It essentially says that, under some assumptions, if (T, Z) is a $(0, \infty) \times E$ valued random variable such that the law of a Brownian bridge from $(0, 0)$ to (T, Z), followed by a Brownian motion starting at (T, Z), is the same as that of a Brownian motion starting in $(0, 0)$, then $Z \sim \mathcal{N}(0, T)$ conditional on T. It is proved in Appendix II.B.

Lemma 3.22. *Suppose $x_0 \in E$, $s_0 > 0$, $f\colon E \times (s_0, \infty) \to (0, \infty)$ is continuous, and μ is a finite measure on E, and that $\int_{s_0}^{\infty} \int_E f(x,s)\mu(\mathrm{d}x)\,\mathrm{d}s < \infty$. Then*

$$B^{(x_0,s_0)+}(\cdot) \sim \iint B^{(x_0,s_0)\to(x,s)+}(\cdot) f(x,s)\mu(\mathrm{d}x)\,\mathrm{d}s$$

if and only if $f(x,s)\mu(\mathrm{d}x) \sim p_{s-s_0}(x - x_0)\,\mathrm{d}x$ for all $s > s_0$. If further $x_1 \in E$, $s_1 > s_0$, and $f\colon E \times (s_0, s_1) \to (0,\infty)$ is continuous, then

$$B^{(x_0,s_0)\to(x_1,s_1)}(\cdot) \sim \iint B^{(x_0,s_0)\to(x,s)\to(x_1,s_1)}(\cdot) f(x,s)\mu(\mathrm{d}x)\,\mathrm{d}s$$

if and only if $f(x,s)\mu(\mathrm{d}x) \sim p_{s-s_0}(x-x_0)p_{s_1-s}(x_1-x)\,\mathrm{d}x$ for all $s \in (s_0, s_1)$.

Proof of Lemma 3.21. Let $\Pi_0 = \{u, v\} \in \mathcal{P}$ and $x \in E^{\{u\}}$. Abbreviate $\mu(\mathrm{d}y) \coloneqq N^{\boldsymbol{\nu}}(xy)^{-1}\mu_x(\mathrm{d}y)$ and $\nu \coloneqq \nu_{2,2}$, so that $\mu(\mathrm{d}y), \nu(\mathrm{d}y) \sim \mathrm{d}y$ are to be shown. The law of $\boldsymbol{X}_t(u)$ under $\mathbb{P}^x$ is $B^{(0,x)+}$. If $y \in E^{\{v\}}$, then started from xy the only possible tree is $F = \{\Pi_0, \{uv\}\}$, and conditional on $\mathrm{Dec} = (\tau, \xi) \in \mathrm{dec}(F)$, the law of $\mathring{\boldsymbol{X}}_t(u)$ (recall $\mathring{\boldsymbol{X}}$ from Lemma 3.15) under $\mathbb{P}^{xy}$ is $B^{(0,x)\to(\tau,\xi)+}$. Thus by sampling consistency,

$$\begin{aligned} B^{(0,x)+}(\cdot) &= \int_E \int_{\mathrm{dec}(F)} B^{(0,x)\to(\tau,\xi)+}(\cdot) P^{xy}(F, \mathrm{d}\tau, \mathrm{d}\xi)\mu_x(\mathrm{d}y), \\ &= \int_E \int_E \int_0^\infty B^{(0,x)\to(\tau,\xi)+}(\cdot) f^F_{\boldsymbol{\nu}}(\tau, \xi \,|\, xy)\,\mathrm{d}\tau\nu(\mathrm{d}\xi)\mu(\mathrm{d}y), \end{aligned}$$

which by Lemma 3.22 and definition of $f^F_{\boldsymbol{\nu}}(\cdot \,|\, xy)$ implies

$$\left(\int_E p_s(x-z)p_s(y-z)\mu(\mathrm{d}y)\right)\nu(\mathrm{d}z) \overset{!}{\sim} p_s(x-z)\,\mathrm{d}z$$

for Lebesgue-a.a. $s > 0$. That is, for such $s > 0$ there is a constant $c(s) > 0$ such that $\left(\int p_s(y-z)\mu(\mathrm{d}y)\right)\nu(\mathrm{d}z) = c(s)\,\mathrm{d}z$. Then ν must have a positive Lebesgue-density which we can write as $1/g$ for some measurable $g\colon E \to (0,\infty)$, and then $T_s\mu = c(s)g$ Lebesgue-a.e. for Lebesgue-a.a. $s > 0$, where $(T_t)_{t\geq 0}$ denotes the heat semigroup. By Lemma 3.23 below there is a constant $c > 0$ such that $g \equiv c$ a.e. and $\mu(\mathrm{d}z) = c\,\mathrm{d}z$. □

Lemma 3.23. *Suppose that μ is a finite, non-zero measure on E, $g\colon E \to (0,\infty)$ is measurable, $N \subset (0,\infty)$ is a Lebesgue-null set and for $s \in (0,\infty) \setminus N$ there is $c(s) > 0$ such that $T_s\mu = c(s)g$ Lebesgue-a.e. Then there is a constant $c > 0$ such that $g(z) = c$ a.e. and $\mu(\mathrm{d}z) = c\,\mathrm{d}z$.*

Proof. All equalities of functions $E \to \mathbb{R}$ in this proof are meant Lebesgue-a.e. First note that g is bounded because $T_s\mu$ is bounded for any $s > 0$. Let $s_0 \in S$, put $f \coloneqq T_{s_0}\mu = c(s_0)g$, and $\widetilde{c}(s) \coloneqq \frac{c(s+s_0)}{c(s_0)}$ whenever $s \in S \coloneqq (0,\infty) \setminus (N - s_0)$. Then $T_s f = c(s+s_0)g = \widetilde{c}(s)f$ for all $s \in S$, which by the semigroup property implies $\widetilde{c}(t+s) = \widetilde{c}(s)\widetilde{c}(t)$ whenever

$s, t, s+t \in S$. This implies that there exists $\alpha \in \mathbb{R}$ with $\widetilde{c}(s) = \mathrm{e}^{\alpha s}$ for all s in some dense subset A of S. Indeed, the set $N' \coloneqq \bigcup_{r\in\mathbb{Q}_+} \frac{N-s_0}{r}$ is null and its complement $S' \subset S$ has the property that for all $s \in S'$ and $r \in \mathbb{Q}_+$ also $rs \in S'$. Fix some $s_1 \in S'$, then $\widetilde{c}(\frac{p}{q}s_1)^q = \widetilde{c}(ps_1) = \widetilde{c}(s_1)^p$ for all $p, q \in \mathbb{N}$, that is $\widetilde{c}(rs_1) = \widetilde{c}(s_1)^r$ for all $r \in \mathbb{Q}_+$, which implies the claim with $\alpha = \log c(s_1)/s_1$ and $A = s_1\mathbb{Q}_+$. We have proved that $T_s g = \mathrm{e}^{\alpha s} g$ for all $s \in A$. The LHS tends to the constant $c \coloneqq \int g(x)\,\mathrm{d}x$ uniformly on E as $s \to \infty$, which necessitates that $\alpha = 0$ and that $g \equiv c$ a.e. Then $(T_s\mu)(\mathrm{d}z) = c\,\mathrm{d}z$ as measures for all $s \in (0,\infty) \setminus N$, and the LHS tends to μ weakly as $s \to \infty$, so $\mu(\mathrm{d}z) = c\,\mathrm{d}z$. □

Lemma 3.24. *All measures $\nu_{n,\vec{k}}$ are absolutely continuous w.r.t. Lebesgue measure.*

Proof. We prove the claim with an induction over $n \geq 2$, the start of which is due to Lemma 3.21. Let $n \geq 3$, assume that the claim is proved for $(n', \vec{k}')$ for all $n' < n$, and take $(n, k_1, \ldots, k_m)$ possible. Let $\Pi_0 \in \mathcal{P}$ with $|\Pi_0| = n$, and $\Pi_0^\oplus$ etc. as above. If $k_1 \geq 3$ then $(n-1, \vec{k}') \coloneqq (n-1, k_1-1, k_2, \ldots, k_m) \leq (n, \vec{k})$ is possible and there are possible trees $F \in \mathbb{T}(\Pi_0)$, $G \in \mathbb{T}(\Pi^\oplus)$ with $G{\downarrow_{\Pi_0}} = F$ whose first merge events are respectively $(n-1, \vec{k}')$ and $(n, \vec{k})$. If $\xi \in \mathrm{sp}(F)$ write $\xi^{(0)} \coloneqq \xi|_{\Pi_F^1 \setminus \mathrm{lf}(F)}$ for the location(s) of the first merge event, and similarly for $\xi \in \mathrm{sp}(G)$. Then, for fixed $\boldsymbol{x} \in E_\circ^{\Pi_0}$ and $y \in E_{\boldsymbol{x}}$, $P^{\boldsymbol{x}}(F, \xi^{(0)} \in \cdot)$ and $P^{\boldsymbol{x}y}(G, \xi^{(0)} \in \cdot)$ have positive densities w.r.t. $\nu_{n-1,\vec{k}'}$ and $\nu_{n,\vec{k}}$, respectively. Indeed,

$$P^{\boldsymbol{x}}(F, \xi^{(0)} \in \mathrm{d}\boldsymbol{z}) = \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} \Bigg(\int_{\mathrm{tm}(F)} \int_{\mathrm{sp}(F\setminus\Pi_0)} f_{\boldsymbol{\nu}}^F(\tau, \xi'\boldsymbol{z} \,|\, \boldsymbol{x})\,\mathrm{d}\tau\, \boldsymbol{\nu}_{F\setminus\Pi_0}(\mathrm{d}\xi') \Bigg) \nu_{n-1,\vec{k}'}(\mathrm{d}\boldsymbol{z}),$$

and similarly for $P^{\boldsymbol{x}y}(G, \xi^{(0)} \in \mathrm{d}\boldsymbol{z})$. Since $P^{\boldsymbol{x}}(F, \xi^{(0)} \in \cdot) \geq \int P^{\boldsymbol{x}y}(G, \xi^{(0)} \in \cdot)\mu_{\boldsymbol{x}}(\mathrm{d}y)$ by sampling consistency, this implies that $\nu_{n,\vec{k}}$ is absolutely continuous w.r.t. $\nu_{n-1,\vec{k}'}$ and thus has a Lebesgue density. If $k_1 = \ldots = k_m = 2$, then $(n-1, \vec{k}') \coloneqq (n-1, k_1, \ldots, k_{m-1}) \leq (n, \vec{k})$ is possible, and by an analogous argument we obtain that $\nu_{n,\vec{k}}$ is absolutely continuous w.r.t. Lebesgue measure in the first $m-1$ arguments, and thus by symmetry (Lemma 3.13) in all arguments. □

We have now proved the assumptions of Lemma 3.15, of which we use the following corollary.

Lemma 3.25. *Let $\Pi_0 \in \mathcal{P}$, $F \in \mathbb{T}(\Pi_0)$, and $\boldsymbol{x} \in E_\circ^{\Pi_0}$, then*

$$\begin{aligned} K_{\boldsymbol{x}}(F^\star, \cdot) \sim \sum_{G\in\mathbb{T}_b^\oplus(F)} \iiint K_{\boldsymbol{x}y}((G, s^\oplus\tau, z^\oplus\xi), \{\mathring{\boldsymbol{X}} \in \cdot\}) P^{\boldsymbol{x}y}((G, s^\oplus\tau, z^\oplus\xi))\,\mathrm{d}s\,\mathrm{d}z\mu_{\boldsymbol{x}}(\mathrm{d}y) \\ + \sum_{G\in\mathbb{T}_{sb}^\oplus(F)} \iint K_{\boldsymbol{x}y}((G, \tau, z^\oplus\xi), \{\mathring{\boldsymbol{X}} \in \cdot\}) P^{\boldsymbol{x}y}((G, \tau, z^\oplus\xi))\,\mathrm{d}z\mu_{\boldsymbol{x}}(\mathrm{d}y) \quad (3.38) \end{aligned}$$

for Lebesgue-a.e. $F^\star = (F, \tau, \xi) \in \operatorname{dec}(F)$*, where domains of integration for* s *and* z *are as indicated by the bijections* (3.30) *and* (3.31)*. For fixed* $F^\star$*, if* $G \in \mathbb{T}_b^\oplus(F)$ *and* $y \in E_{\boldsymbol{x}}$*,*

$$P^{\boldsymbol{x}y}((G, s^\oplus\tau, z^\oplus\xi)) \sim \frac{\mathrm{e}^{-\beta_G^\oplus s}}{N^{\boldsymbol{\nu}}(\boldsymbol{x}y)} p_s(z-y) p_{s-\tau_{u_\oplus}}(z-(\boldsymbol{x}\xi)_{u_\oplus}) \left[p_{\tau_{w_\oplus}-s}(\xi_{w_\oplus}-z)\right]_{w_\oplus \neq \emptyset} \boldsymbol{\nu}_G(z^\oplus\xi) \tag{3.39}$$

where $\beta_G^\oplus = |\nu_{|\Pi_G^\oplus|}| - |\nu_{|\Pi_G^\oplus|+1}|$*, and if* $G \in \mathbb{T}_{sb}^\oplus(F)$*,*

$$P^{\boldsymbol{x}y}((G, \tau, z^\oplus\xi)) \sim \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x}y)} p_{\tau_\oplus}(z-y) p_{\tau_\oplus-\tau_{u_\oplus}}(z-(\boldsymbol{x}\xi)_{u_\oplus}) p_{\tau_{w_\oplus}-\tau_\oplus}(\xi_{w_\oplus}-z) \boldsymbol{\nu}_G(z^\oplus\xi). \tag{3.40}$$

The constants in both cases only depend on $F^\star$ *and* $\boldsymbol{x}$*.*

Proof. The first claim follows directly from (3.32), and (3.39) and (3.40) follow from the definition of $f_{\boldsymbol{\nu}}$ by discarding terms that don't depend on s, z, y, and z, y, respectively. □

Lemma 3.26. *For every* $\boldsymbol{x} \in \mathcal{X}$*,* $\mu_{\boldsymbol{x}}(\mathrm{d}y) \sim N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y$*.*

Proof. Let $n \geq 3$, $\Pi_0 \in \mathcal{P}$ with $|\Pi_0| = n$, and $\boldsymbol{x} \in E_\circ^{\Pi_0}$. By Lemma 3.20 there exists a possible tree $F = \{\Pi_0, \dots, \Pi_m = \{u_\oplus\}\}$ such that $G \coloneqq \{\Pi_0 \cup \{v_\oplus\}, \dots, \Pi_{m-1} \cup \{v_\oplus\}, \{u_\oplus, v_\oplus\}, \{u_\oplus \cup v_\oplus\}\}$ is possible. Note $G{\downarrow_{\Pi_0}} = F$. In words, in G the first $n-1$ particles merge according to F and then the root of F merges with the nth particle $v_\oplus$. Fix $F^\star = (F, \tau, \xi) \in \operatorname{dec}(F)$ for which (3.38) holds, and denote by $s_0 \coloneqq \tau(u_\oplus)$, $z_0 \coloneqq \xi(u_\oplus)$ time and place of the birth of F's root. Then the law of $\boldsymbol{Y}_\bullet \coloneqq \boldsymbol{X}_{s_0+\bullet}(u_\oplus)$, the motion of F's root, under $K_{\boldsymbol{x}}(F^\star, \cdot)$ is $B^{(s_0,z_0)+}$. For any $G'^\star \in \operatorname{dec}^\oplus(F^\star)$ with $G' \neq G$ the law of $\boldsymbol{Y}$ under $K_{\boldsymbol{x}}(G'^\star, \{\mathring{\boldsymbol{X}} \in \cdot\})$ is the same, because $v_\oplus$ merges into F before s_0. If $G^\star = (G, s^\oplus\tau, z^\oplus\xi) \in \operatorname{dec}^\oplus(F^\star \,|\, G)$, then the law of $\boldsymbol{Y}$ under $K_{\boldsymbol{x}}(G^\star, \{\mathring{\boldsymbol{X}} \in \cdot\})$ is $B^{(s_0,z_0)\to(s,z)+}$, so by Lemma 3.25,

$$\begin{aligned} B^{(s_0,z_0)+}(\cdot) &\sim \int_{E_{\boldsymbol{x}}} \int_E \int_{s_0}^\infty B^{(s_0,z_0)\to(s,z)+}(\cdot) P^{\boldsymbol{x}y}((G, s^\oplus\tau, z^\oplus\xi))\,\mathrm{d}s\,\mathrm{d}z\mu_{\boldsymbol{x}}(\mathrm{d}y) \\ &\sim \iiint B^{(s_0,z_0)\to(s,z)+}(\cdot) p_s(z-y) p_{s-s_0}(z-z_0) \mathrm{e}^{-|\nu_2|s}\,\mathrm{d}s\,\mathrm{d}z\mu(\mathrm{d}y), \end{aligned}$$

where we put $\mu(\mathrm{d}y) \coloneqq \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x}y)}\mu_{\boldsymbol{x}}(\mathrm{d}y)$ and used that $\nu_{2,2}(\mathrm{d}z) \sim \mathrm{d}z$ is already known (that is, $\nu_{2,2}(z)$ is constant in z). By Lemma 3.22, this implies that, for Lebesgue-a.a. $s \in (s_0, \infty)$,

$$\left(\int_{E_{\boldsymbol{x}}} p_s(z-y) p_{s-s_0}(z-z_0)\mu(\mathrm{d}y)\right)\mathrm{d}z \sim p_{s-s_0}(z-z_0)\,\mathrm{d}z,$$

that is $T_s\mu$ is a.e. constant for a.a. $s > 0$, and thus $\mu(\mathrm{d}z) \sim \mathrm{d}z$ by Lemma 3.23. □

Lemma 3.27. *For every* $(n, \vec{k}) = (n, k_1, \dots, k_m) \in M$ *and* $i \in [m]$*,* $\nu_{n,\vec{k}}$ *is constant in the* i*th argument if* $k_i = 2$*. In particular,* $\nu_{n,2}(\mathrm{d}z) = \lambda_{n,2}\,\mathrm{d}z$ *for some* $\lambda_{n,2} \geq 0$ *for each* $n \geq 2$*.*

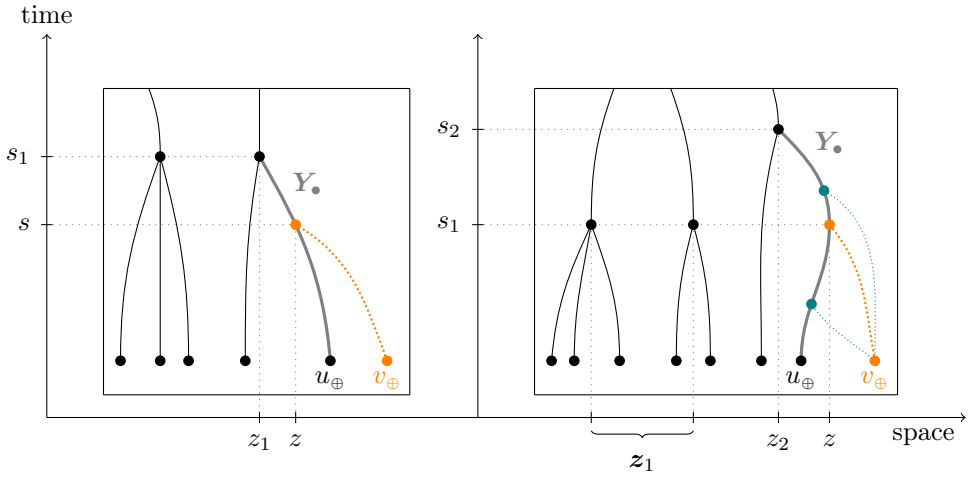

Figure 3.2: Illustrations for the proof of Lemma 3.27.

Proof. We first prove $\nu_{n,2}(\mathrm{d}z) \sim \mathrm{d}z$ for all $n \geq 2$. This is already proved for $n = 2$, so let $n \geq 3$, $(\Pi_0, \boldsymbol{x}) \in \mathcal{X}$ with $|\Pi_0| = n - 1$, and $F \in \mathbb{T}(\Pi_0)$ a possible tree. Let $u_\oplus \in \Pi_0$ be one of the leaves involved in the first merge event, and denote by $G \in \mathbb{T}(\Pi_0^\oplus)$ the tree in which an additional leaf $v_\oplus$ merges with $u_\oplus$ before the first merge of F, so $G{\downarrow_{\Pi_0}} = F$. Fix $F^\star = (F, \tau, \xi) \in \mathrm{dec}(F)$ for which (3.38) holds, and put $s_1 \coloneqq \tau(\mathrm{pr}_F(u_\oplus))$, $z_1 \coloneqq \tau(\mathrm{pr}_F(u_\oplus))$, the time and place of the first merge of $u_\oplus$ in F. See Fig. 3.2 (left). Then the law of $(\boldsymbol{Y}_r \coloneqq \boldsymbol{X}_r(u_\oplus))_{0 \leq r \leq s_1}$ under $K_{\boldsymbol{x}}(F^\star, \cdot)$ is $B^{(0,x_0)\to(s_1,z_1)}$ where $x_0 \coloneqq \boldsymbol{x}(u_\oplus)$. For any $G'^\star \in \mathrm{dec}^\oplus(F^\star)$ with $G' \neq G$ the law of $\boldsymbol{Y}$ under $K_{\boldsymbol{x}}(G'^\star, \{\mathring{\boldsymbol{X}} \in \cdot\})$ is the same, but if $G^\star = (G, s^\oplus\tau, z^\oplus\xi) \in \mathrm{dec}^\oplus(F^\star \,|\, G)$, then the law of $\boldsymbol{Y}$ under $K_{\boldsymbol{x}}(G^\star, \{\mathring{\boldsymbol{X}} \in \cdot\})$ is $B^{(0,x_0)\to(s,z)\to(s_1,z_1)}$. Thus by Lemma 3.25, for some $\beta \in \mathbb{R}$,

$$\begin{aligned} B^{(0,x_0)\to(s_1,z_1)}(\cdot) &\sim \int_{E_{\boldsymbol{x}}} \int_E \int_0^{s_1} B^{(0,x_0)\to(s,z)\to(s_1,z_1)}(\cdot) P^{\boldsymbol{x}y}((G, s^\oplus\tau, z^\oplus\xi))\,\mathrm{d}s\,\mathrm{d}z\mu_{\boldsymbol{x}}(\mathrm{d}y) \\ &\sim \iiint B^{(0,x_0)\to(s,z)\to(s_1,z_1)}(\cdot)\mathrm{e}^{-\beta s} p_s(z-y)p_s(z-x_0) \\ &\qquad\qquad p_{s_1-s}(z_1-z)\nu_{n,2}(z)\,\mathrm{d}s\,\mathrm{d}z\,\mathrm{d}y \\ &= \iint B^{(0,x_0)\to(s,z)\to(s_1,z_1)}(\cdot)\mathrm{e}^{-\beta s} p_s(z-x_0)p_{s_1-s}(z_1-z)\nu_{n,2}(\mathrm{d}z)\,\mathrm{d}s, \end{aligned}$$

which by Lemma 3.22 implies $\nu_{n,2}(\mathrm{d}z) \sim \mathrm{d}z$.

Now let $n \geq 3$ and $(n, k_1, \ldots, k_m)$ possible with $m \geq 2$ and $k_m = 2$. Take $\Pi_0 \in \mathcal{P}$ with $|\Pi_0| = n - 1$, and let $G \in \mathbb{T}(\Pi_0^\oplus)$ be a tree in which the first merge event is $(n, \vec{k})$, and in which the parent of the two leaves $u_\oplus, v_\oplus$ comprising the merge associated to k_m, is involved in the merge event following the first. See Fig. 3.2 (right). Let $F \coloneqq G \setminus\!\setminus v_\oplus \in \mathbb{T}(\Pi_0)$, so $G{\downarrow_{\Pi_0}} = F$. Fix $F^\star = (F, \tau, \xi) \in \mathrm{dec}(F)$ for which (3.38) holds, and denote $s_1 \coloneqq \tau(\Pi_1^F)$, $\boldsymbol{z}_1 \coloneqq \xi|_{\Pi_1^F \setminus \Pi_0}$, $s_2 \coloneqq \tau(\mathrm{pr}_F(u_\oplus))$, $z_2 \coloneqq \xi(\mathrm{pr}_F(u_\oplus))$, then the law of $(\boldsymbol{Y}_r \coloneqq \boldsymbol{X}_r(u_\oplus))_{0 \leq r \leq s_2}$ under $K_{\boldsymbol{x}}(F^\star, \cdot)$ is $B^{(0,x_0)\to(s_2,z_2)}$, where $x_0 = \boldsymbol{x}(u_\oplus)$. If $G'^\star \in \mathrm{dec}^\oplus(F^\star)$ with $G' \neq G$, then the law of $\boldsymbol{Y}$ under $K_{\boldsymbol{x}y}(G'^\star, \{\mathring{\boldsymbol{X}} \in \cdot\})$ is the same,

except in the two cases where G' merges $v_\oplus$ with $u_\oplus$ in $(0, s_1)$ or in (s_1, s_2) (see teal coloured parts of Fig. 3.2). In the former case,

$$\begin{aligned}
&\iiint \boldsymbol{Y} \# K_{\boldsymbol{x}\boldsymbol{y}}(G'^{\star}, \{\mathring{\boldsymbol{X}} \in \cdot\}) P^{\boldsymbol{x}\boldsymbol{y}}((G, s^{\oplus}\tau, z^{\oplus}\xi))\,\mathrm{d}s\,\mathrm{d}z\mu_{\boldsymbol{x}}(\mathrm{d}y) \\
&\quad\sim \iiint B^{(0,x_0)\to(s,z)\to(s_2,z_2)}(\cdot)\mathrm{e}^{-\beta s} p_s(z-y)p_s(z-x_0)p_{s_2-s}(z_2-z)\,\mathrm{d}s\,\mathrm{d}z\,\mathrm{d}y \\
&\quad= \int_0^{s_1} B^{(0,x_0)\to(s,z)\to(s_2,z_2)}(\cdot)\mathrm{e}^{-\beta s} p_s(x_0-z)p_{s_2-s}(z-z_2)\,\mathrm{d}z\,\mathrm{d}s \\
&\quad\sim B^{(0,x_0)\to(s_2,z_2)}(\cdot),
\end{aligned}$$

where we used Lemma 3.22, and similarly if the merge is in (s_1, s_2). Now, if $G^\star = (G, \tau, z^\oplus\xi) \in \mathrm{dec}^\oplus(F^\star \,|\, G)$ (see the orange coloured part of Fig. 3.2, right), then the law of $\boldsymbol{Y}$ under $K_{\boldsymbol{x}\boldsymbol{y}}(G^\star, \{\mathring{\boldsymbol{X}} \in \cdot\})$ is $B^{(0,x_0)\to(s_1,z)\to(s_2,z_2)}$, so by Lemma 3.25,

$$\begin{aligned}
B^{(0,x_0)\to(s_2,z_2)}(\cdot) &\sim \int_{E_{\boldsymbol{x}}}\int_{E_{\boldsymbol{z}_1}} B^{(0,x_0)\to(s_1,z)\to(s_2,z_2)}(\cdot)P^{\boldsymbol{x}\boldsymbol{y}}((G,\tau,z^\oplus\xi))\,\mathrm{d}z\mu_{\boldsymbol{x}}(\mathrm{d}y) \\
&\sim \iint B^{(0,x_0)\to(s_1,z)\to(s_2,z_2)}(\cdot)p_{s_1}(z-x_0)p_{s_1}(z-y) \\
&\qquad\qquad p_{s_2-s_1}(z_2-z)\nu(\boldsymbol{z}_1,z)\,\mathrm{d}z\,\mathrm{d}y \\
&\sim \iint B^{(0,x_0)\to(s_1,z)\to(s_2,z_2)}(\cdot)p_{s_1}(z-x_0)p_{s_2-s_1}(z_2-z)\nu(\boldsymbol{z}_1,z)\,\mathrm{d}z.
\end{aligned}$$

Evaluating both laws at time s_1 gives that $\nu_{n,\vec{k}}(\boldsymbol{z}_1, \cdot)$ is (a.e.) constant. By symmetry of $\nu_{n,\vec{k}}$ it must then be constant in all arguments i with $k_i = 2$. □

This means we can write $\nu_{n,\vec{k}}(\boldsymbol{z}) = \nu_{n,\vec{k}}(z_1, \ldots, z_j)$ if $(n, k_1, \ldots, k_m) \in M$ with $m \geq 2$ and $k_j = \ldots = k_m = 2$. The following lemma finishes this section.

Lemma 3.28. $\boldsymbol{\nu}(\mathrm{d}\xi) \sim \mathrm{d}\xi$.

Proof. We first prove that for every $(n, \vec{k}) \in M$ there exist non-negative constants $(c^{(j)})_{0\leq j\leq m}$ and c^{sb} (whose dependence on $(n, \vec{k})$ we suppress) such that

$$\nu_{n,\vec{k}}(\boldsymbol{z}) = c^{(0)}\nu_{n+1,\vec{k}}(\boldsymbol{z}) + c^{\mathrm{sb}}\nu_{n+1,k_1,\ldots,k_m,2}(\boldsymbol{z}) + \sum_{j=1}^{m} c^{(j)}\nu_{n+1,k_1,\ldots,k_j+1,\ldots,k_m}(\boldsymbol{z}), \tag{3.41}$$

for a.e. $\boldsymbol{z}$, and $c^{(j)} > 0$ for $j \in [m]$. Fix $(n, \vec{k})$ and let F be a possible tree that starts with an $(n, \vec{k})$-merger, and put $F^\circ := F \setminus \mathrm{lf}(F)$. By passing to the total mass in (3.32), for a.e. $\boldsymbol{z} \in E_\circ^{\Pi_1^F \setminus \Pi_0^F}$, $\xi \in \mathrm{sp}(F^\circ)$, and $\tau \in \mathrm{tm}(F)$,

$$\begin{aligned}
P^{\boldsymbol{x}}((F,\tau,\boldsymbol{z}\xi)) = &\sum_{G\in\mathbb{T}_{\mathrm{b}}^{\oplus}(F)} \iiint P^{\boldsymbol{x}\boldsymbol{y}}((G, s^\oplus\tau, z^\oplus\boldsymbol{z}\xi))\,\mathrm{d}s\,\mathrm{d}z\mu_{\boldsymbol{x}}(\mathrm{d}y) \\
&+ \sum_{G\in\mathbb{T}_{\mathrm{sb}}^{\oplus}(F)} \iint P^{\boldsymbol{x}\boldsymbol{y}}((G,\tau,z^\oplus\boldsymbol{z}\xi))\,\mathrm{d}z\mu_{\boldsymbol{x}}(\mathrm{d}y) + \sum_{G\in\mathbb{T}_{\mathrm{m}}^{\oplus}(F)} \int P^{\boldsymbol{x}\boldsymbol{y}}((G,\tau,\boldsymbol{z}\xi))\mu_{\boldsymbol{x}}(\mathrm{d}y).
\end{aligned} \tag{3.42}$$

For the remainder of this proof put $f_F(\tau) \coloneqq \prod_{(\Pi,\Pi')\in F} \mathrm{e}^{-|\nu_\Pi|(\tau_{\Pi'}-\tau_\Pi)}$ and analogously for other trees, so that

$$f^F_{\boldsymbol{\nu}}(\tau, \boldsymbol{z}\xi \,|\, \boldsymbol{x}) = \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f_F(\tau) f^F_{\mathrm{sp}}(\boldsymbol{z}\xi \,|\, \tau, \boldsymbol{x}).$$

Further recall that $P^{\boldsymbol{x}}((F,\tau,\boldsymbol{z}\xi)) = f^F_{\boldsymbol{\nu}}(\tau, \boldsymbol{z}\xi \,|\, \boldsymbol{x})\boldsymbol{\nu}_F(\mathrm{d}(\boldsymbol{z}\xi))$. We will now divide (3.42) by $f^F_{\mathrm{sp}}(\boldsymbol{z}\xi \,|\, \tau, \boldsymbol{x})$ and then integrate it over ξ and τ. On the left we get $\frac{|\boldsymbol{\nu}_{F^\circ}|}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} \left(\int f_F(\tau)\,\mathrm{d}\tau\right) \nu_{n,\vec{k}}(\boldsymbol{z})$. Now let $G \in \mathbb{T}^{\oplus}_{\mathrm{b}}(F)$ such that the binary merge is before the first merge of F, then the corresponding summand in (3.42) is $\lambda_{2,2}/\int N^{\boldsymbol{\nu}}(\boldsymbol{x}y)\,\mathrm{d}y$ times

$$\int \left(\iint f^G_{\mathrm{sp}}(z^{\oplus}\boldsymbol{z}\xi \,|\, s^{\oplus}\tau, \boldsymbol{x}y)\,\mathrm{d}y\,\mathrm{d}z \right) f_G(s^{\oplus}\tau)\nu_{n,\vec{k}}(\boldsymbol{z})\boldsymbol{\nu}_{F^\circ}(\xi)\,\mathrm{d}s$$
$$= f^F_{\mathrm{sp}}(\boldsymbol{z}\xi \,|\, \tau, \boldsymbol{x}) \left(\int f_G(s^{\oplus}\tau)\,\mathrm{d}s \right) \nu_{n,\vec{k}}(\boldsymbol{z})\boldsymbol{\nu}_{F^\circ}(\xi),$$

an equality we've used before, cmp. (3.36). Dividing by $f^F_{\mathrm{sp}}(\boldsymbol{z}\xi \,|\, \tau, \boldsymbol{x})$ and integrating over ξ and τ gives a positive (because F° is possible) multiple of $\nu_{n,\vec{k}}(\boldsymbol{z})$. Similarly, if the merge is after the first of F, we obtain a non-negative (because then G may be impossible) multiple of $\nu_{n+1,\vec{k}}(\boldsymbol{z})$. If $G \in \mathbb{T}^{\oplus}_{\mathrm{sb}}(F)$ then we get non-negative multiples of either $\nu_{n,\vec{k}}(\boldsymbol{z})$ or $\nu_{n,k_1,\dots,k_m,2}(\boldsymbol{z})$, and if $G \in \mathbb{T}^{\oplus}_{\mathrm{m}}(F)$ then we get either a non-negative multiple of $\nu_{n+1,\vec{k}}(\boldsymbol{z})$, or a positive multiple of $\nu_{n+1,k_1\dots k_j+1\dots k_m}(\boldsymbol{z})$ for some $j \in [m]$, and for each such j there is a corresponding $G \in \mathbb{T}^{\oplus}_{\mathrm{m}}(F)$.

This proves the claim surrounding (3.41), from which we conclude with an induction over k in the statement "Every $\nu_{n,\vec{k}}$ is constant in the ith argument if $k_i \leq k$". For $k = 2$ this is Lemma 3.27, now suppose it has been proved for some fixed $k \geq 2$. Take some $(n,\vec{k})$ for which there is $j \in [m]$ with $k_j = k+1$ and $k_i \leq k$ for $i > j$. Put $\vec{k}' \coloneqq (k_1, \dots k_{j-1}, k_j - 1, k_{j+1}, \dots, k_m)$, and invoke (3.41) with $\nu_{n,\vec{k}'}$ on the LHS. Then $\nu_{n,\vec{k}}$ appears as a summand on the RHS (perhaps multiple times), and all other summands and the LHS are, by the induction hypothesis, constant in the jth to mth arguments. Hence also $\nu_{n,\vec{k}}$ must be constant in the jth argument, and by symmetry in all arguments with $k_i = k+1$. □

3.4 Drift Representation

In this section we prove the drift representation Theorem 1.15.

Proof of Theorem 1.15. Denote by $(\Pi_t, \boldsymbol{X}_t)_{t\geq 0}$ the Brownian spatial coalescent with transition rates $\boldsymbol{\nu}(\mathrm{d}\boldsymbol{z}) = \boldsymbol{\lambda}\,\mathrm{d}\boldsymbol{z}$ for some $\boldsymbol{\lambda} \in R^\circ$, started from $\Pi_0 = \{\{1\},\dots,\{n\}\}$ and $\boldsymbol{X}_0 = \boldsymbol{x} \in E^{\Pi_0}_\circ$. Denote $\boldsymbol{Z}_t = (Z^1_t, \dots, Z^n_t) = (\boldsymbol{X}_t(\{1\}), \dots, \boldsymbol{X}_t(\{n\}))$ as long as t is smaller than the random time of the first merge event, that is for $t < \inf\{s > 0 \colon \Pi_s = \Pi_0\}$.

We perform a slightly informal generator calculation, and leave technical details to the interested reader. Recall that the (time-dependent) generator of a Brownian bridge

on the torus E, starting at x_0 at time t_0, going to x_1 at time t_1, is given by

$$A_s f(x) = \frac{1}{2}\Delta f(x) + \nabla_x \log p_{t_1-s}(x_1 - x) \cdot \nabla f(x), \qquad s \in [t_0, t_1).$$

Then the generator of $\boldsymbol{Z}$ is

$$\begin{aligned}
&Af(z_1, \dots, z_n) \\
&\quad = \sum_{i=1}^{n} \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{z})} \int\limits_{\mathrm{dec}(\mathbb{F}(\Pi_0))} \left(\frac{1}{2}\nabla_{z_i}^2 f(\boldsymbol{z}) + \Big[\nabla_{z_i} \log p(\tau_{\mathrm{pr}_F(\{i\})}, \xi_{\mathrm{pr}_F(\{i\})} - z_i)\Big]_{\mathrm{pr}_F(\{i\})\neq\emptyset} \cdot \nabla_{z_i} f(\boldsymbol{z}) \right) \\
&\qquad\qquad \times f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(F, \tau) f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{z}) \,\mathrm{d}\tau \,\mathrm{d}\xi \\
&\quad = \frac{1}{2}\Delta f(\boldsymbol{z}) + \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{z})} \sum_{i=1}^{n} \nabla_{z_i} f(\boldsymbol{z}) \cdot \int\limits_{\mathrm{dec}(\mathbb{F}(\Pi_0))} \Big[\nabla_{z_i} \log p(\tau_{\mathrm{pr}_F(\{i\})}, \xi_{\mathrm{pr}_F(\{i\})} - z_i)\Big]_{\mathrm{pr}_F(\{i\})\neq\emptyset} \\
&\qquad\qquad f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(F, \tau) f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{z}) \,\mathrm{d}\tau \,\mathrm{d}\xi.
\end{aligned}$$

If $(F, \tau, \xi) \in \mathrm{dec}(\mathbb{F}(\Pi_0))$ with $\mathrm{pr}_F(\{i\}) \neq \emptyset$, then $f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{z})$ depends on z_i only through a factor $p(\tau_{\mathrm{pr}_F(\{i\})}, \xi_{\mathrm{pr}_F(\{i\})} - z_i)$, so

$$(\nabla_{z_i} \log p(\tau_{\mathrm{pr}_F(\{i\})}, \xi_{\mathrm{pr}_F(\{i\})} - z_i)) f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{z}) = \nabla_{z_i} f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{z}).$$

If $\mathrm{pr}_F(\{i\}) = \emptyset$, then $f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{z})$ is independent of z_i, so in any case

$$\begin{aligned}
Af(z_1, \dots, z_n) &= \frac{1}{2}\Delta f(\boldsymbol{z}) + \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{z})} \sum_{i=1}^{n} \nabla_{z_i} f(\boldsymbol{z}) \cdot \int\limits_{\mathrm{dec}(\mathbb{F}(\Pi_0))} f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(F, \tau) \nabla_{z_i} f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{z}) \,\mathrm{d}\tau \,\mathrm{d}\xi \\
&= \frac{1}{2}\Delta f(\boldsymbol{z}) + \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{z})} \sum_{i=1}^{n} \nabla_{z_i} f(\boldsymbol{z}) \cdot \nabla_{z_i} N^{\boldsymbol{\nu}}(\boldsymbol{z}) \\
&= \frac{1}{2}\Delta f(\boldsymbol{z}) + \nabla f(\boldsymbol{z}) \cdot \nabla \log N^{\boldsymbol{\nu}}(\boldsymbol{z}).
\end{aligned}$$

To pull the derivative ∇_{z_i} out of the integral in the second step (and thereby showing differentiability of $N^{\boldsymbol{\nu}}$), we require statements analogous to Lemmas II.E.1 and II.E.2 with $f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x})$ replaced by $\nabla_{x_i} f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x})$. This is true using similar methods, with the modification that in $d = 1$, points $\boldsymbol{x}$ which have two identical coordinates have to be excluded (as is already the case in $d \geq 2$); recall also Remark 1.16 (ii) on this issue. □

3.5 Associated Population Models

This section contains proofs of the results presented in Section 1.5 of the introduction.

3.5.1 Ξ-Fleming-Viot Process

Fix a non-zero measure Ξ on $\triangle$, and recall Section 1.5. We begin by constructing a stationary, bi-infinite version of the particle representation $(\boldsymbol{Y}(t))$ of the Ξ-Fleming-Viot process. Let $\mathfrak{N}^{ij}$ for $i, j \in \mathbb{N}$, $i < j$, and $\mathfrak{M}$ be Poisson point processes as described in

Section 1.5. Recall that for $i, j \in \mathbb{N}$ and $i < j$, level j "looks down" to level i at points of $\mathfrak{N}^{ij}$, and at points of $\mathfrak{M}$ if i and j are in the same basket, and that basket contains no levels smaller than i. In particular, the ordered set of times $(\tau_l^j)_{l\in\mathbb{Z}}$ where level j looks down to some level $i < j$ is itself a Poisson point process with finite intensity, so almost-surely countably infinite and without accumulation points. Denote by $k_l^j \in \{1, \dots, j-1\}$ for $l \in \mathbb{Z}$ the level to which particle j looks down at time τ_l^j.

Let $(B^1(t))_{t\in(-\infty,\infty)}$ be a Brownian motion running at stationarity on E, and

$$(B_l^j(t) \colon j \in \mathbb{N}, j \geq 2, l \in \mathbb{Z})_{t\geq 0}$$

be an independent family of Brownian motions on E, independent of B^1 and the Poisson point processes, and all starting at zero. To construct the process $(\boldsymbol{Y}(t))_{t\in(-\infty,\infty)}$, we put $Y_1 = B^1$, and then construct Y_j for $j \geq 2$ inductively. Suppose $(Y_1(t), \dots, Y_{j-1}(t))_{t\in(-\infty,\infty)}$ has already been constructed for some $j \geq 2$. Then for every $l \in \mathbb{Z}$ put

$$Y_j(t) = Y_{k_l^j}(\tau_l^j) + B_l^j(t - \tau_l^j), \qquad t \in [\tau_l^j, \tau_{l+1}^j).$$

In words, whenever j looks down to some level $i < j$ at some time τ, we let Y_j start at the location of Y_i at time τ and evolve as a Brownian motion until the next time it looks down. Since these lookdown times are an infinite, discrete set, this defines $Y_j(t)$ for all $t \in (-\infty, \infty)$. Clearly, for any fixed $t_0 \in \mathbb{R}$, the law of the evolution of the process $(\boldsymbol{Y}(t))_{t\geq t_0}$ is as described in Section 1.5, and since this entire construction is invariant under constant time shifts, the distribution of $\boldsymbol{Y}(t)$ and $\boldsymbol{Y}(s)$ is the same for any fixed $s, t \in (-\infty, \infty)$.

We will now prove Theorem 1.24, which will imply both Proposition 1.20 and Theorem 1.23. The statement is trivial for $n = 1$, so fix $n \geq 2$ for the rest of Section 3.5.1, and put $\Pi_0 = \{\{1\}, \dots, \{n\}\}$. We present a formal construction of the process described in Theorem 1.24. It requires the following measurable maps (their existence follows from a basic measure theoretic argument, see Lemma II.F.1:

(i) a map $h_0 \colon E^n \to E^n$ such that the pushforward of Lebesgue measure $\mathrm{d}\boldsymbol{y}$ on E^n under h_0 is $N^{\Xi}(\boldsymbol{y})\,\mathrm{d}\boldsymbol{y}$.

(ii) for every $m \in \{1, \dots, n-1\}$ a map $h_m \colon E^m \times E^{n-m} \to E^{n-m}$ such that for every $\boldsymbol{x} \in E_\circ^m$, the pushforward of Lebesgue measure $\mathrm{d}\boldsymbol{y}$ on E^{n-m} under $h_m(\boldsymbol{x}, \cdot)$ is $\frac{N^{\Xi}(\boldsymbol{x}\boldsymbol{y})}{N^{\Xi}(\boldsymbol{x})}\,\mathrm{d}\boldsymbol{y}$, and $h_m(\boldsymbol{x}, \cdot) = \mathrm{id}_{E^{n-m}}$ if $\boldsymbol{x} \in E^m \setminus E_\circ^m$.

(iii) for every $F \in \mathbb{T}(\Pi)$ a map $h_F \colon E^n \times \mathrm{dec}(F) \to \mathrm{dec}(F)$ such that for every $\boldsymbol{x} \in E_\circ^n$, the pushforward of $f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(F, \tau)\,\mathrm{d}\tau\,\mathrm{d}\xi$ under $h_F(\boldsymbol{x}, \cdot)$ is $\frac{1}{N_F^{\Xi}(\boldsymbol{x})} f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(F, \tau) f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x})\,\mathrm{d}\tau\,\mathrm{d}\xi$, and $h_F(\boldsymbol{x}, \cdot) = \mathrm{id}_{\mathrm{dec}(F)}$ if $\boldsymbol{x} \in E^n \setminus E_\circ^n$.

The definitions of $h_m(\boldsymbol{x}, \cdot)$ and $h_F(\boldsymbol{x}, \cdot)$ if $\boldsymbol{x} \in E^m \setminus E_\circ^m$ and $\boldsymbol{x} \in E^n \setminus E_\circ^n$, respectively, is irrelevant as long as all maps are measurable.

Now let $(S, \mathcal{A}, \mathbb{P})$ be a probability space that supports the following random variables, all independent.

(i) For every $m \in \{0, \dots, n-1\}$ an i.i.d. sequence $(\omega_m^{(i)})_{i \in \mathbb{N}_0}$ of uniform E^{n-m}-valued random variables.

(ii) For every $F \in \mathbb{T}(\Pi_0)$ an i.i.d. sequence $(\omega_F^{(i)})_{i \in \mathbb{N}_0}$ of random variables in $\mathrm{dec}(F)$ with distribution $f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(F, \tau)\, \mathrm{d}\tau\, \mathrm{d}\xi$.

(iii) For every non-empty $u \subset [n]$ and $i \in \mathbb{N}_0$, a d-dimensional standard Brownian bridge $\omega_B^{(i)} \in C([0,1], E)$ with $\omega_B^{(i)}(0) = \omega_B^{(i)}(1) = 0$.

(iv) An i.i.d. sequence $(\omega_U^{(i)})_{i \in \mathbb{N}_0}$ of uniform $[0,1)$-random variables.

We further fix an ordering $\mathbb{T}(\Pi_0) = \{F^{(1)}, \dots, F^{(|\mathbb{T}(\Pi_0)|)}\}$ of $\mathbb{T}(\Pi_0)$.

Definition 3.29. Let $\boldsymbol{\nu} \in R$. The *n'th level Brownian spatial coalescent with resampling* associated with transition measures $\boldsymbol{\nu}$ is a càglàd (left-continuous with right-limits) E^n-valued process $\boldsymbol{Z}^n(t) = (Z_1^n(t), \dots, Z_n^n(t))$, $t \geq 0$, defined on S as follows. Given a realisation $\omega \in S$:

(i) Put $\zeta^{(0)} \coloneqq h_0(\omega_0^{(0)})$, and $t^{(0)} \coloneqq 0$. (At the end of the construction, $\zeta^{(i)} = \boldsymbol{Z}^n(t^{(i)}+)$.)

(ii) For $i \in \mathbb{N}_0$, assuming $0 = t^{(0)} < \ldots < t^{(i)}$ are defined and $\boldsymbol{Z}^n(t)$ has been constructed for $t \in [0, t^{(i)})$:

(ii.1) If $\zeta^{(i)} \in E^n \setminus E_\circ^n$ put $j = 1$, otherwise let $j \in \{1, \dots, |\mathbb{T}(\Pi_0)|\}$ be the unique number such that

$$\sum_{j'<j} N_{F^{(j')}}^{\Xi}(\zeta^{(i)}) \leq \omega_U^{(i)} N^{\Xi}(\zeta^{(i)}) < \sum_{j' \leq j} N_{F^{(j')}}^{\Xi}(\zeta^{(i)}).$$

Put $F \coloneqq F^{(j)}$.

(ii.2) Let $(\tau, \xi) = h_F(\zeta^{(i)}, \omega_F^{(i)}) \in \mathrm{dec}(F)$, and $t^{(i+1)} \coloneqq \tau(\Pi_1^F)$ the time of the first merge event. For $t \in (t^{(i)}, t^{(i+1)}]$, and $l \in [n]$, put

$$Z_l^n(t) \coloneqq \zeta_l^{(i)} + \frac{t - t^{(i)}}{\tau_{\mathrm{pr}_F(\{l\})}}(\xi_{\mathrm{pr}_F(\{l\})} - \zeta_l^{(i)}) + \sqrt{\tau_{\mathrm{pr}_F(\{l\})}}\,\omega_B^{(i)}\left(\frac{t - t^{(i)}}{\tau_{\mathrm{pr}_F(\{l\})}}\right),$$

which is a Brownian bridge started from $\zeta_l^{(i)}$ at time $t^{(i)}$, going to $\xi_{\mathrm{pr}_F(\{l\})}$ at time $t^{(i)} + \tau_{\mathrm{pr}_F(\{l\})}$, and stopped at time $t^{(i+1)}$.

(ii.3) Let $I = \{\min u \colon u \in \Pi_1^F\}$ be the set of levels that were either not involved in the first merge event, or were minimal among the set of levels with which they merged. Let $J = [n] \setminus I$. For every $l \in I$, put $\zeta_l^{(i+1)} = Z_l^n(t^{(i+1)}-)$, and

$$\zeta_J^{(i)} \coloneqq h_{|J|}(\zeta_I^{(i)}, \omega_{|J|}^{(i+1)}),$$

where $\zeta_J^{(i+1)} = (\zeta_j^{(i+1)} : j \in J)$, and $\zeta_J^{(i+1)}$ similarly.

This defines a measurable map from S into the space of càglàd paths $[0,\infty) \to E^n$.

All definitions dealing with cases where $\boldsymbol{Z}^n(t+) \in E^n \setminus E^n_\circ$ for some $t \geq 0$, which happens with probability zero, are only in place for the map $S \to D([0,\infty), E^n)$ to be well-defined, and do not affect the law of $\boldsymbol{Z}^n$.

Remark 3.30. If the Brownian spatial coalescent with transition measures $\boldsymbol{\nu}$ is sampling consistent (that is, associated to some finite measure Ξ on $\triangle$), then for every $m < n$, the laws of $(Z_1^m, \dots, Z_m^m)$ and $(Z_1^n, \dots, Z_m^n)$ are the same, and we expect that a variant of the Kolmogorov extension theorem can be used to construct an E^∞-valued process $\boldsymbol{Z}(t) = (Z_1(t), Z_2(t), \dots)$ whose first n levels have the same law as $\boldsymbol{Z}^n$ for every $n \in \mathbb{N}$. Making this precise is not necessary to prove our results.

From now, let $\boldsymbol{Z}^n$ be the n'th level Brownian spatial Ξ-coalescent with resampling.

Theorem 3.31. *The laws of $(Z_1^n(t), \dots, Z_n^n(t))_{t\geq 0}$ and $(Y_1(-t), \dots, Y_n(-t))_{t\geq 0}$ are the same. In particular, the stationary distribution of $\boldsymbol{Y}$ is an i.i.d. sample from a random realisation of μ^Ξ.*

Proof. It suffices to show that for every $T > 0$, the laws of $(Z_1^n(t), \dots, Z_n^n(t))_{t\in[0,T]}$, and $(Y_1(T-t), \dots, Y_n(T-t))_{t\in[0,T]}$, which coincides with that of $(Y_1(-t), \dots, Y_n(-t))_{t\in[0,T]}$, are the same. Denote by $\mathbb{Q}^{\boldsymbol{y}}$ for $\boldsymbol{y} \in E^n$ the probability measure on $D([0,T], E^n)$ describing the law of $(Y_1(t), \dots, Y_n(t))_{t\in[0,T]}$ started from (or conditioned on) $(Y_1(0), \dots, Y_n(0)) = \boldsymbol{y}$. Then we show

$$\mathbb{P}((Z_1^n(t), \dots, Z_n^n(t))_{t\in[0,T]} \in \cdot) = \int_{E^n} \mathbb{Q}^{\boldsymbol{y}}((Y_1(T-t), \dots, Y_n(T-t))_{t\in[0,T]} \in \cdot) N^\Xi(\boldsymbol{y})\, \mathrm{d}\boldsymbol{y}. \tag{3.43}$$

In particular, this implies that

$$N^\Xi(\boldsymbol{z})\, \mathrm{d}\boldsymbol{z} = \mathbb{P}((Z_1^n(0), \dots, Z_n^n(0)) \in \mathrm{d}\boldsymbol{z}) = \int_{E^n} \mathbb{Q}^{\boldsymbol{y}}((Y_1(T), \dots, Y_n(T)) \in \mathrm{d}\boldsymbol{z}) N^\Xi(\mathrm{d}\boldsymbol{y})\, \mathrm{d}\boldsymbol{y},$$

that is, $N^\Xi(\boldsymbol{y})\, \mathrm{d}\boldsymbol{y}$ (which equals $\mathbb{E}\left[\mu^\Xi(\mathrm{d}y_1) \dots \mu^\Xi(\mathrm{d}y_n)\right]$, recall (1.7)) is the stationary distribution of $(Y_1(t), \dots, Y_n(t))$. Plugging this back into (3.43) gives

$$\mathbb{P}((Z_1^n(t), \dots, Z_n^n(t))_{t\in[0,T]} \in \cdot) = \mathbb{P}((Y_1(T-t), \dots, Y_n(T-t))_{t\in[0,T]} \in \cdot),$$

which is the claim.

We now prove (3.43). For simplicity we assume that $\xi_j = 0$ for $j \geq 2$ for Ξ-a.e. $\boldsymbol{\xi} \in \triangle$, in which case the evolution of $\boldsymbol{Y}$ can be described in terms of a finite measure Λ on $[0,1]$, see Remark 1.19. The general case requires more notation but no different ideas. Consider the event where $(Y_1(t), \dots, Y_n(t))$, restricted to the first n levels, starts at

$\boldsymbol{y}^0 = (y_1^0, \dots, y_n^0)$, undergoes a branching event at time $s \in (0,T)$ involving $k \geq 2$ levels $J = \{j_1 < \ldots < j_k\}$, where $Y^n(s-) = \boldsymbol{y}^s = (y_1^s, \dots, y_n^s)$, and has no further branching events until time T where it ends in $Y^n(T) = \boldsymbol{y}^T = (y_1^T, \dots, y_n^T)$. The rate at which a branching event evolving exactly indices J happens is $\lambda_{n,k}$, the associated rate of the Λ-coalescent. Indeed, if $k > 2$ then the rate is

$$\int p^k (1-p)^{n-k} \frac{\Lambda(\mathrm{d}p)}{p^2} = \lambda_{n,k},$$

and if $k = 2$ it is $\Lambda(\{0\}) = \int p^{k-2}(1-p)^{n-k}\Lambda(\mathrm{d}p) = \lambda_{n,2}$. Thus the probability (density) of the entire event is

$$\underbrace{N^\Lambda(\boldsymbol{y}^0)\,\mathrm{d}\boldsymbol{y}^0}_{\text{initial sample}} \times \underbrace{\prod_{j=1}^n p_s(y_j^0 - y_j^s)\,\mathrm{d}y_j^s}_{\text{spatial movement in } (0,s)} \times \underbrace{\lambda_{n,k}\mathrm{e}^{-\lambda_n s}}_{\text{branching event}}$$
$$\times \underbrace{\prod_{j\in J} p_{T-s}(y_{j_1}^s - y_j^T)\,\mathrm{d}y_j^T \prod_{j\notin J} p_{T-s}(y_j^s - y_j^T)\,\mathrm{d}y_j^T}_{\text{spatial movement in } (s,T)} \times \underbrace{\mathrm{e}^{-\lambda_n(T-s)}}_{\text{no branching events in } (s,T)}. \quad (3.44)$$

The process $((Z_1^n(T-t), \dots, Z_n^n(T-t)))_{t\in[0,T]}$ lies in the same event if and only if

(i) the initial state is $\boldsymbol{Z}^n(0) = \boldsymbol{y}^T$,

(ii) the first coalescence event of the Brownian spatial Λ-coalescent started from $\boldsymbol{y}^T$ at time 0 is a multiple merger of the lineages with labels J at time $T-s$,

(iii) the resampling at time $T-s$, given the locations $\boldsymbol{y}^{s,1} \coloneqq (y_j^s)_{j\in[n]\setminus J\cup\{j_1\}}$ yields $\boldsymbol{y}^{s,2} \coloneqq (y_j^s)_{j\in J\setminus\{j_1\}}$,

(iv) the Brownian spatial Λ-coalescent started from $\boldsymbol{y}^s$ at time $T-s$ has no coalescence events until time T, where it ends in state $\boldsymbol{y}^0$.

The probability density for this is (recall Lemma 3.1)

$$\underbrace{N^\Lambda(\boldsymbol{y}^T)\,\mathrm{d}\boldsymbol{y}^T}_{\text{(i)}} \times \underbrace{\frac{N^\Lambda(\boldsymbol{y}^{s,1})}{N^\Lambda(\boldsymbol{y}^T)}\lambda_{n,k}\mathrm{e}^{-\lambda_n(T-s)} \prod_{j\in J} p_{T-s}(y_{j_1}^s - y_j^T) \prod_{j\notin J} p_{T-s}(y_j^s - y_j^T)\,\mathrm{d}\boldsymbol{y}^{s,1}}_{\text{(ii)}}$$
$$\times \underbrace{\frac{N^\Lambda(\boldsymbol{y}^s)\,\mathrm{d}\boldsymbol{y}^{s,2}}{N^\Lambda(\boldsymbol{y}^{s,1})}}_{\text{(iii)}} \times \underbrace{\frac{N^\Lambda(\boldsymbol{y}^0)}{N^\Lambda(\boldsymbol{y}^s)}\mathrm{e}^{-\lambda_n s}\prod_{j=1}^n p_s(y_j^0 - y_j^s)\,\mathrm{d}y_j^s}_{\text{(iv)}},$$

which is equal to (3.44) after some cancellations. Conditional on this event, both the realisation of $\boldsymbol{Z}^n$ and $\boldsymbol{Y}$ are obtained by sampling Brownian bridges between $\boldsymbol{y}^0$ and $\boldsymbol{y}^s$, and between $\boldsymbol{y}^s$ and $\boldsymbol{y}^T$. The same argument applies (with no modifications but more notation) to all possible realisations. □

This proves Proposition 1.20 and Theorem 1.24. To deduce Theorem 1.23, we need to show that the coalescent process obtained from $\boldsymbol{Z}^n$ by projecting onto the trajectories of only the initial n particles and their ancestors (equivalently by forgetting about all resampled particles and their ancestors), has the law of a Brownian spatial coalescent. This essentially follows from the definitions, and the strong Markov property of the Brownian spatial coalescent. A formal argument first observes that the statement is true by definition if we stop at the time $t^{(1)}$ of the first coalescence–resampling event, and using an inductive argument we can assume it is true started from the smaller number of particles remaining after the first merge event. Then the strong Markov property of the Brownian spatial coalescent applied at time $t^{(1)}$, and sampling consistency imply that the overall law is the same.

3.5.2 Spatial Cannings Models

Let $N \in \mathbb{N}$, and recall the definitions in Section 1.5.2. The construction of a bi-infinite, stationary version $(\boldsymbol{Y}^N(t), \mathcal{Y}_t^N)_{t\in(-\infty,\infty)}$ is done in exactly the same way as the corresponding construction for the Ξ-Fleming-Viot process, presented at the beginning of Section 3.5.1.

Recall the definition of $(\Pi_n^{N,n}, \boldsymbol{X}_t^{N,n})_{t\geq 0}$ from Section 1.5.2, which for $N \in \mathbb{N}$ and $n \in [N]$ is a random variable with values in Ω.

Lemma 3.32. *For any $n, N \in \mathbb{N}$, $n \leq N$, the distribution of $(\Pi_t^{N,n}, \boldsymbol{X}_t^{N,n})_{t\geq 0}$ is the Brownian spatial coalescent with transition measures*

$$\nu_{n,\vec{k}}(\mathrm{d}\boldsymbol{z}) = T_N p_{n,\vec{k}}^N \,\mathrm{d}\boldsymbol{z}, \quad (n,\vec{k}) \in M,$$

started from $\Pi_0^{N,n} = \{\{1\},\dots,\{n\}\}$ and $(\boldsymbol{X}_0^{N,n}(\{1\}),\dots,\boldsymbol{X}_0^{N,n}(\{n\}))$ a random sample from $N^{\boldsymbol{\nu}}(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}$. Furthermore, the stationary distribution of $(Y_1^N,\dots,Y_n^N)$ is given by $N^{\boldsymbol{\nu}}(\boldsymbol{y})\,\mathrm{d}\boldsymbol{y}$.

Proof. The distribution of $(\Pi_t)_{t\geq 0}$ is a non-spatial coalescent with transition rates

$$\lambda_{n,\vec{k}} = T_N p_{n,\vec{k}}^N, \quad (n,\vec{k}) \in M.$$

This follows by a simple combinatorial argument explained in [33], see equation (3) in that paper. If $(O_1,\dots,O_N) = (1,\dots,1)$ a.s. then $\lambda_{n,\vec{k}} = 0$ for all $(n,\vec{k}) \in M$ and the statement is trivial. Otherwise $\lambda_n \geq \lambda_{n,2} > 0$ for all $n \geq 2$, so that $\lim_{t\to\infty} |\Pi_t| = 1$ almost-surely, and $Z := \mathrm{Dec}((\Pi^{N,n}, \boldsymbol{X}^{N,n}))$ takes values in $\mathrm{dec}(\mathbb{T})$ (recall (2.7)). By definition of the spatial Cannings model, conditional on Z, we can obtain $(\boldsymbol{X}^{N,n})$ by running independent Brownian bridges along the branches of the coalescence tree encoded in Z, and a Brownian motion after the final coalescence event (this is just the time reversal of the first level in the forwards process). Hence it suffices to show that the law of Z, conditional on

$\boldsymbol{X}_0^{N,n} = \boldsymbol{x}$, is equal to $\mathrm{Dec}\,\#\mathbb{P}^{\boldsymbol{x}}$ where $\mathbb{P}^{\boldsymbol{x}}$ is the law of a Brownian spatial coalescent with transition measures $\boldsymbol{\nu}$, started at $\boldsymbol{x}$. We do this by verifying (3.33).

Let $(F, \tau, \xi) \in \mathrm{dec}(\mathbb{T})$ with $\Pi_0^F = \Pi_0 \coloneqq \{\{1\}, \dots, \{n\}\}$, and denote $\pi = (\pi_t)_{t \geq 0} = \mathrm{Tm}^{-1}(F, \tau) \in \Omega_0$ (the unique path π of a non-spatial coalescent satisfying $\mathrm{Tm}(\pi) = (F, \tau)$). Denote by $u_0 = \{1, \dots, n\}$ the root of F (that is, $\mathrm{rt}(F) = \{u_0\}$). Let $\boldsymbol{x} \in E^{\Pi_0}$. Then,

$$\begin{aligned}
&\mathbb{P}(Z \in \mathrm{d}(F, \tau, \xi), \boldsymbol{X}_0^{N,n} \in \mathrm{d}\boldsymbol{x}) \\
&\qquad = \mathbb{P}(\mathrm{Tm}((\Pi_t^{N,n})) \in \mathrm{d}\tau) \mathbb{P}\left(Z \in \mathrm{d}(F, \tau, \xi), \boldsymbol{X}_0^{N,n} \in \mathrm{d}\boldsymbol{x} \,\middle|\, \mathrm{Tm}((\Pi_t^{N,n})) \in \mathrm{d}\tau\right).
\end{aligned}$$

By Lemma 3.16,

$$\mathbb{P}(\mathrm{Tm}((\Pi_t^{N,n})) \in \mathrm{d}\tau) = f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(F, \tau)\, \mathrm{d}\tau,$$

and because the movement of particles forwards in time, conditional on the branching, are independent Brownian motions,

$$\begin{aligned}
&\mathbb{P}(Z \in \mathrm{d}(F, \tau, \xi), \boldsymbol{X}_0^{N,n} \in \mathrm{d}\boldsymbol{x} \mid \mathrm{Tm}((\Pi_t^{N,n})) \in \mathrm{d}\tau) \\
&\qquad = \underbrace{\mathbb{P}(Y_1^N(-\tau_{u_0}) \in \mathrm{d}\xi_{u_0})}_{\substack{\text{location of most recent} \\ \text{common ancestor (MRCA)}}} \times \underbrace{\prod_{u \in \mathrm{nd}(F) \setminus \mathrm{rt}(F)} p(\tau_{\mathrm{pr}_F(u)} - \tau_u, (\boldsymbol{x}\xi)_u - \xi_{\mathrm{pr}_F(u)})\, \mathrm{d}(\boldsymbol{x}\xi)_u}_{\substack{\text{locations of all other particles at branch points} \\ \text{conditional on location of MRCA}}} \\
&\qquad = \mathrm{d}\xi_{u_0} \times f_{\mathrm{sp}}(\xi \mid \tau, \boldsymbol{x}) \left(\prod_{u \in \mathrm{nd}^\circ(F) \setminus \mathrm{rt}(F)} \mathrm{d}\xi_u\right) \mathrm{d}\boldsymbol{x} \\
&\qquad = f_{\mathrm{sp}}(\xi \mid \tau, \boldsymbol{x})\, \mathrm{d}\xi\, \mathrm{d}\boldsymbol{x},
\end{aligned}$$

where we used that the first (in fact, any) level of $\boldsymbol{Y}^N$ at stationarity is uniformly distributed in E. We conclude

$$\mathbb{P}(Z \in \mathrm{d}(F, \tau, \xi), \boldsymbol{X}_0^{N,n} \in \mathrm{d}\boldsymbol{x}) = f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(F, \tau) f_{\mathrm{sp}}(\xi \mid \tau, \boldsymbol{x})\, \mathrm{d}\xi\, \mathrm{d}\tau\, \mathrm{d}\boldsymbol{x}.$$

Integrating out (F, τ, ξ) gives

$$\mathbb{P}(\boldsymbol{X}_0^{N,n} \in \mathrm{d}\boldsymbol{x}) = N^{\boldsymbol{\nu}}(\boldsymbol{x})\, \mathrm{d}\boldsymbol{x}, \tag{3.45}$$

so

$$\mathbb{P}(Z \in \mathrm{d}(F, \tau, \xi) \mid \boldsymbol{X}_0^{N,n} \in \mathrm{d}\boldsymbol{x}) = \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(F, \tau) f_{\mathrm{sp}}(\xi \mid \tau, \boldsymbol{x})\, \mathrm{d}\xi\, \mathrm{d}\tau.$$

Recalling (3.33), this shows that the law of $\mathrm{Dec}((\Pi_t^{N,n}, \boldsymbol{X}_t^{N,n}))$ conditional on $\boldsymbol{X}_0^{N,n} = \boldsymbol{x}$ is equal to $\mathrm{Dec}\,\#\mathbb{P}^{\boldsymbol{x}}$ where $\mathbb{P}^{\boldsymbol{x}}$ is the law of a Brownian spatial coalescent with transition measures $\boldsymbol{\nu}$ started from $\boldsymbol{x}$.

Finally, we already know that $(Y_1^N, \dots, Y_n^N)$ is at stationarity, and $(Y_1^N(0), \dots, Y_n^N(0)) = (\boldsymbol{X}_0^{N,n}(\{1\}), \dots, \boldsymbol{X}_0^{N,n}(\{n\}))$, so the final claim follows from (3.45). □

Lemma 3.33. *Let $n, N \in \mathbb{N}$, $n \leq N$, and denote by $\boldsymbol{Z}^n$ the n'th level Brownian spatial coalescent with resampling associated with transition measures*

$$\nu_{n,\vec{k}}(\mathrm{d}\boldsymbol{z}) = T_N p^N_{n,\vec{k}}\, \mathrm{d}\boldsymbol{z}.$$

Then $(Y_1^N(-t), \ldots, Y_n^N(-t))_{t\geq 0}$ and $\boldsymbol{Z}^n$ have the same law.

Proof. Let $k \in \{2, \ldots, n\}$ and $J \subset \{1, \ldots, n\}$ with $|J| = k$. We want to compute the probability $q_{n,k}$ that at any given reproduction event in the spatial Cannings model, all levels in J look down to level $\min J$, and all levels in $\{1, \ldots, n\} \setminus J$ do not look down. For a given realisation $(O_1, \ldots, O_N) = (o_1, \ldots, o_N)$, and indices $i_J \in [N]$ and $i_j \in [N]$, $j \in [n] \setminus J$, all pairwise distinct, the probability that level j for $j \in [n] \setminus J$ falls into "bucket" i_j, and all levels in J fall into bucket i_J, is

$$\frac{1}{(N)_n} (o_{i_J})_k \prod_{j \notin J} o_j.$$

(Recall that $(N)_n = N(N-1)\ldots(N-n_1).$) There are $(N)_{n'}$ choices for the indices $i_J, (i_j)_{j\notin J}$ where $n' = n - k + 1$, so by exchangeability of $(O_1, \ldots, O_N)$,

$$q_{n,k} = \frac{(N)_{n'}}{(N)_n} \mathbb{E}\left[(O_1)_k O_2 \ldots O_{n'}\right] = p^N_{n,k}.$$

Thus the rate at which such a lookdown event happens when observing $(Y_1^N, \ldots, Y_n^N)$ forwards in time is equal to $T_N p^N_{n,\vec{k}}$. A similar argument holds for general lookdown events. With this observation, and the fact that we know the stationary distribution of $(Y_1^N, \ldots, Y_n^N)$ from Lemma 3.32, the rest of the proof is identical to that of Theorem 3.31. □

The final ingredients in the proof of Theorem 1.26 are the following two lemmas.

Lemma 3.34. *If $\lambda^N_{n,\vec{k}} \to \lambda_{n,\vec{k}}$, then the Brownian spatial coalescent with transition measures $\nu^N_{n,\vec{k}}(\mathrm{d}\boldsymbol{z}) = \lambda^N_{n,\vec{k}}\, \mathrm{d}\boldsymbol{z}$ converges weakly to those with transition measures $\nu_{n,\vec{k}}(\mathrm{d}\boldsymbol{z}) = \lambda_{n,\vec{k}}\, \mathrm{d}\boldsymbol{z}$.*

Proof. For $\boldsymbol{x} \in \mathcal{X}$ and $N \in \mathbb{N}$ put $P^{\boldsymbol{x},N} := \mathrm{Dec}\,\#\mathbb{P}^{\boldsymbol{x},N}$, and $P^{\boldsymbol{x}} := \mathrm{Dec}\,\#\mathbb{P}^{\boldsymbol{x}}$. Suppose we can show that $P^{\boldsymbol{x},N} \implies P^{\boldsymbol{x}}$ for every $\boldsymbol{x} \in \mathcal{X}$. Then, for any $f \in C_b(\Omega)$,

$$\int_\Omega f(\omega)\mathbb{P}^{\boldsymbol{x},N}(\mathrm{d}\omega) = \int_{\mathrm{dec}(\mathbb{F})} K_{\boldsymbol{x}}(F^\star, f) P^{\boldsymbol{x},N}(\mathrm{d}F^\star), \tag{3.46}$$

where $K_{\boldsymbol{x}}(F^\star, f) = \int_\Omega f(\omega) K_{\boldsymbol{x}}(F^\star, \mathrm{d}\omega)$ is continuous in $\boldsymbol{x}$ and $F^\star$ (recall Lemma 2.9), and bounded by $\|f\|_\infty$. Hence the right-hand side of (3.46) converges to the same expression with $P^{\boldsymbol{x},N}$ replaced by $P^{\boldsymbol{x}}$, so to $\int_\Omega f(\omega)\mathbb{P}^{\boldsymbol{x}}(\mathrm{d}\omega)$. This proves $\mathbb{P}^{\boldsymbol{x},N} \implies \mathbb{P}^{\boldsymbol{x}}$ as $N \to \infty$.

Now recall from (3.33) that

$$P^{\boldsymbol{x}}(F, \mathrm{d}\tau, \mathrm{d}\xi) = \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(F, \tau) f^{F}_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x}) \,\mathrm{d}\tau \,\mathrm{d}\xi,$$

and analogously for $P^{\boldsymbol{x},N}$. By Lemma 3.36, $N^{\boldsymbol{\nu}^N}(\boldsymbol{x}) \to N^{\boldsymbol{\nu}}(\boldsymbol{x})$ as $N \to \infty$ for fixed $\boldsymbol{x}$, so

$$\frac{1}{N^{\boldsymbol{\nu}^N}(\boldsymbol{x})} f^{\boldsymbol{\lambda}^N}_{\mathrm{tm}}(F, \tau) f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x}) \to \frac{1}{N^{\boldsymbol{\nu}}(\boldsymbol{x})} f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(F, \tau) f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x}).$$

Hence both $P^{\boldsymbol{x},N}$ and $P^{\boldsymbol{x}}$ have densities w.r.t. the same (uniform) measure, and the densities converge pointwise, so Scheffé's lemma implies $P^{\boldsymbol{x},N} \implies P^{\boldsymbol{x}}$ for every $\boldsymbol{x} \in \mathcal{X}$. □

Lemma 3.35. *If $\lambda^N_{n,\vec{k}} \to \lambda_{n,\vec{k}}$, then the n'th level Brownian spatial coalescent with resampling associated with transition measures $\nu^N_{n,\vec{k}}(\mathrm{d}\boldsymbol{z}) = \lambda^N_{n,\vec{k}} \,\mathrm{d}\boldsymbol{z}$ converges weakly to that associated with transition measures $\nu_{n,\vec{k}}(\mathrm{d}\boldsymbol{z}) = \lambda_{n,\vec{k}} \,\mathrm{d}\boldsymbol{z}$.*

Proof. Exactly by the same argument that proved Lemma 3.34. The resampling causes no issues because $N^{\boldsymbol{\nu}^N} \to N^{\boldsymbol{\nu}}$ pointwise (see Lemma 3.36). □

Proof of Theorem 1.26. The first statement has been proved in Lemma 3.32. Statement (ii) implies (iii), and (iii) and (iv) are equivalent by [33]. Statement (iv) implies (ii) by Lemmas 3.32 and 3.34. The coalescent process $(\Pi^{N,n}, \boldsymbol{X}^{N,n})$ is a continuous image of $(Y^N_1, \ldots, Y^N_n)$, so (i) implies (ii) by the continuous mapping theorem. On the other hand, by Lemma 3.33 the time reversal $\boldsymbol{Z}^{N,n}$ of $(Y^N_1, \ldots, Y^N_n)$ is the Brownian spatial coalescent with resampling associated with the transition measures $\boldsymbol{\nu}^N$ defined in Theorem 1.26. Thus if (iv) holds, then by Lemma 3.35 the time reversal of, and thus $(Y^N_1, \ldots, Y^N_n)$ itself converge weakly. This proves equivalence of (i) to (iv).

Now suppose that (iii) holds. Then the limit of $T_N p^N_{n,\vec{k}}$ must be the rates of a Ξ-coalescent by [33], so by Lemmas 3.32 and 3.34 the limit of $(\Pi^{N,n}, \boldsymbol{X}^{N,n})$ is a Brownian spatial Ξ-coalescent (with initial condition drawn from $N^{\Xi}(\boldsymbol{x}) \,\mathrm{d}\boldsymbol{x}$). Similarly, by Lemmas 3.33 and 3.35 the time reversal of $(Y^N_1, \ldots, Y^N_n)$ converges to the Brownian spatial Ξ-coalescent with resampling, which is itself the time reversal of the first n levels $(Y_1, \ldots, Y_n)$ of the Ξ-Fleming-Viot process by Theorem 3.31, so $(Y^N_1, \ldots, Y^N_n)$ converges weakly to $(Y_1, \ldots, Y_n)$. □

Lemma 3.36. *Suppose that $\boldsymbol{\lambda}^N \to \boldsymbol{\lambda}$ in R°, in the sense that $\lambda^N_{n,\vec{k}} \to \lambda_{n,\vec{k}}$ for every $(n, \vec{k}) \in M$. Define $\boldsymbol{\nu}^N \in R$ by $\nu_{n,\vec{k}}(\mathrm{d}\xi) = \lambda_{n,\vec{k}} \,\mathrm{d}\xi$ for $(n, \vec{k}) \in M$ and $N \in \mathbb{N}$, and $\boldsymbol{\nu}$ similarly. Then,*

$$N^{\boldsymbol{\nu}^N}(\boldsymbol{x}) \to N^{\boldsymbol{\nu}}(\boldsymbol{x})$$

for every $\boldsymbol{x} \in \mathcal{X}$.

Proof. Recall that for $(\Pi, \boldsymbol{x}) \in \mathcal{X}$,

$$N^{\boldsymbol{\nu}}(\boldsymbol{x}) = \sum_{F \in \mathbb{F}(\Pi)} \int_{\mathrm{tm}(F)} \int_{\mathrm{sp}(F)} f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(\tau) f_{\mathrm{sp}}(\xi \mid \tau, \boldsymbol{x}) \, \mathrm{d}\xi \, \mathrm{d}\tau.$$

The map $\boldsymbol{\lambda} \mapsto f^{\boldsymbol{\lambda}}_{\mathrm{tm}}(\tau)$ is continuous for every $\tau \in \mathrm{tm}(F)$, and it is easy to see that there exists $\boldsymbol{\lambda}' \in R^\circ$ and $N_0 \in \mathbb{N}$ such that $\sup_{N \geq N_0} f^{\boldsymbol{\lambda}^N}_{\mathrm{tm}} \leq f^{\boldsymbol{\lambda}'}_{\mathrm{tm}}$. Then the claim follows from the dominated convergence theorem. □

Acknowledgements

I would like to thank my supervisor Alison Etheridge for numerous helpful discussions, valuable feedback, guidance, and encouragement. I would further like to thank Michal Bassan and Félix Foutel-Rodier for many insightful conversations. I further acknowledge support by the EPSRC grant EP/W523781/1.

Bibliography

[1] Nick H Barton, Alison M Etheridge, and Amandine Véber. "A new model for evolution in a spatial continuum". In: *Electronic Journal of Probability* 15 (2010).

[2] Nick H Barton, Alison M Etheridge, and Amandine Véber. "Modelling evolution in a spatial continuum". In: *Journal of Statistical Mechanics: Theory and Experiment* 2013.01 (2013), P01002.

[3] Julien Berestycki, Nathanaël Berestycki, and Jason R Schweinsberg. "The genealogy of branching Brownian motion with absorption". In: *The Annals of Probability* 41.2 (2013), pp. 527–618.

[4] Jean Bertoin and Jean-Françcois Le Gall. "Stochastic flows associated to coalescent processes". In: *Probability Theory and Related Fields* 126.2 (2003), pp. 261–288.

[5] Matthias Birkner, Huili Liu, and Anja Sturm. "Coalescent results for diploid exchangeable population models". In: *Electronic Journal of Probability* 23.none (2018), pp. 1–44. DOI: 10.1214/18-EJP175.

[6] Matthias Birkner et al. "A modified lookdown construction for the Xi-Fleming-Viot process with mutation and populations with recurrent bottlenecks". In: *Alea* 6 (2009), pp. 25–61.

[7] Erwin Bolthausen and Alain-Sol Sznitman. "On Ruelle's probability cascades and an abstract cavity method". In: *Communications in mathematical physics* 197 (1998), pp. 247–276.

[8] Éric Brunet and Bernard Derrida. "Genealogies in simple models of evolution". In: *Journal of Statistical Mechanics: Theory and Experiment* 2013.01 (2013), P01006. DOI: 10.1088/1742-5468/2013/01/P01006.

[9] Vincent Calvez et al. "Dynamics of lineages in adaptation to a gradual environmental change". In: *Annales Henri Lebesgue* 5 (2022), pp. 729–777.

[10] Fernando Cordero et al. "Λ-coalescents arising in a population with dormancy". In: *Electronic Journal of Probability* 27.none (2022), pp. 1–34. DOI: 10.1214/22-EJP739.

[11] Dimitrios Diamantidis et al. "Bursts of coalescence within population pedigrees whenever big families occur". In: *bioRxiv* (2023). DOI: 10.1101/2023.10.17.562743. eprint: https://www.biorxiv.org/content/early/2023/10/21/2023.10.17.562743.1.full.pdf.

[12] Peter Donnelly and Thomas G Kurtz. "A countable representation of the Fleming-Viot measure-valued diffusion". In: *The Annals of Probability* 24.2 (1996), pp. 698–742. DOI: 10.1214/aop/1039639359.

[13] Peter Donnelly and Thomas G Kurtz. "Particle representations for measure-valued population models". In: *The Annals of Probability* 27.1 (1999), pp. 166–205.

[14] Bjarki Eldon and John Wakeley. "Coalescent Processes When the Distribution of Offspring Number Among Individuals Is Highly Skewed". In: *Genetics* 172.4 (2006), pp. 2621–2633. ISSN: 1943-2631. DOI: 10.1534/genetics.105.052175. eprint: https://academic.oup.com/genetics/article-pdf/172/4/2621/42069578/genetics2621.pdf.

[15] Alison M Etheridge. "Almost sure convergence of measure-valued branching processes: a critical exponent". In: *Proceedings of the Royal Society of Edinburgh Section A: Mathematics* 124.4 (1994), pp. 811–823.

[16] Alison M Etheridge. *An introduction to superprocesses*. 20. American Mathematical Soc., 2000.

[17] Alison M Etheridge. "Asymptotic behaviour of measure-valued critical branching processes". In: *Proceedings of the American Mathematical Society* 118.4 (1993), pp. 1251–1261.

[18] Alison M Etheridge. "Drift, draft and structure: some mathematical models of evolution". In: *Banach center publications* 1.80 (2008), pp. 121–144.

[19] Stewart N Ethier and Thomas G Kurtz. "The infinitely-many-neutral-alleles diffusion model". In: *Advances in Applied Probability* 13.3 (1981), pp. 429–452.

[20] Steven N Evans. "Coalescing Markov labelled partitions and a continuous sites genetics model with infinitely many types". In: *Annales de l'Institut Henri Poincare (B) Probability and Statistics*. Vol. 33. 3. Elsevier. 1997, pp. 339–358.

[21] Nic Freeman. "The segregated Λ-coalescent". In: *The Annals of Probability* 43.2 (2015), pp. 435–467.

[22] Andreas Greven et al. "Renormalisation of hierarchically interacting Cannings processes". In: *arXiv preprint* (2012). arXiv:1209.1856.

[23] Kenji Handa. "Quasi-invariance and reversibility in the Fleming-Viot process". In: *Probability theory and related fields* 122.4 (2002), pp. 545–566.

[24] Samuel G. G. Johnston, Andreas Kyprianou, and Tim Rogers. "Multitype Lambda-coalescents". In: *The Annals of Applied Probability* 33.6A (2023), pp. 4210–4237. DOI: 10.1214/22-AAP1891.

[25] Samuel G. G. Johnston and Amaury Lambert. "The coalescent structure of uniform and Poisson samples from multitype branching processes". In: *arXiv preprint* (2019). arXiv:1912.00198.

[26] John FC Kingman. "The coalescent". In: *Stochastic processes and their applications* 13.3 (1982), pp. 235–248.

[27] John FC Kingman, Giorgio Koch, and Fabio L Spizzichino. "Exchangeability and the evolution of large populations". In: *Exchangeability in probability and statistics* 91 (1982), p. 112.

[28] Zenghu Li, Tokuzo Shiga, and Lihua Yao. "A reversibility problem for Fleming-Viot processes". In: *Electronic Communications in Probability* 4 (1999), pp. 65–76.

[29] Vlada Limic and Anja Sturm. "The spatial Λ-coalescent". In: *Electronic Journal of Probability* 11.none (2006), pp. 363–393. DOI: 10.1214/EJP.v11-319.

[30] Martin Möhle. "On multi-type Cannings models and multi-type exchangeable coalescents". In: *arXiv preprint* (2023). arXiv:2304.05809.

[31] Martin Möhle. "Robustness Results for the Coalescent". In: *Journal of Applied Probability* 35.2 (1998), pp. 438–447. ISSN: 00219002.

[32] Martin Möhle. "Weak Convergence to the Coalescent in Neutral Population Models". In: *Journal of Applied Probability* 36.2 (1999), pp. 446–460. ISSN: 00219002.

[33] Martin Möhle and Serik Sagitov. "A Classification of Coalescent Processes for Haploid Exchangeable Population Models". In: *The Annals of Probability* 29.4 (2001), pp. 1547–1562. DOI: 10.1214/aop/1015345761.

[34] Martin Möhle and Serik Sagitov. "Coalescent patterns in diploid exchangeable population models". In: *Journal of mathematical biology* 47.4 (2003), pp. 337–352.

[35] Ludger Overbeck and Michael Röckner. "Geometric aspects of finite and infinite-dimensional Fleming-Viot processes". In: *Random Operators and Stochastic Equations* 5.1 (1997), pp. 35–58. DOI: doi:10.1515/rose.1997.5.1.35.

[36] Jim Pitman. "Coalescents With Multiple Collisions". In: *The Annals of Probability* 27.4 (1999), pp. 1870–1902. DOI: 10.1214/aop/1022874819.

[37] L Chris G Rogers and David Williams. *Diffusions, Markov processes, and martingales*. Vol. 2. Cambridge University Press, 2000.

[38] Serik Sagitov. "Convergence to the coalescent with simultaneous multiple mergers". In: *Journal of Applied Probability* 40.4 (2003), pp. 839–854. DOI: 10.1239/jap/1067436085.

[39] Serik Sagitov. "The General Coalescent with Asynchronous Mergers of Ancestral Lines". In: *Journal of Applied Probability* 36.4 (1999), pp. 1116–1125. ISSN: 00219002.

[40] Ori Sargsyan and John Wakeley. "A coalescent process with simultaneous multiple mergers for approximating the gene genealogies of many marine organisms". In: *Theoretical Population Biology* 74.1 (2008), pp. 104–114. ISSN: 0040-5809. DOI: https://doi.org/10.1016/j.tpb.2008.04.009.

[41] Jason R Schweinsberg. "Coalescent processes obtained from supercritical Galton–Watson processes". In: *Stochastic Processes and their Applications* 106.1 (2003), pp. 107–139. ISSN: 0304-4149. DOI: https://doi.org/10.1016/S0304-4149(03)00028-0.

[42] Jason R Schweinsberg. *Coalescents with simultaneous multiple collisions*. University of California, Berkeley, 2001.

[43] Tokuzo Shiga. "A stochastic equation based on a Poisson system for a class of measure-valued diffusion processes". In: *Journal of Mathematics of Kyoto University* 30.2 (1990), pp. 245–279.

[44] Arno Siri-Jégousse and Alejandro H Wences. "Exchangeable coalescents beyond the Cannings class". In: *arXiv preprint* (2023). arXiv:2212.02154.

[45] Aaron Stephen Andrew Smith. "Probability models of population structure and gene flow". PhD thesis. University of Oxford, 2021.

Appendix

II.A Maps and Spaces

In this section we prove measurability of the maps Fr, Tm, Sp, and Dec. Recall that Ω and Ω_0 are equipped with the product topology and (consequently) the product σ algebra.

Lemma II.A.1. *For $N \in \mathbb{N}$ define $\mathrm{Fr}^{(N)} \colon \Omega_0 \to \mathbb{F}$ by*

$$\mathrm{Fr}^{(N)}(\omega) := \{\omega(t) \colon t \in \mathcal{T}_N\},$$

where $\mathcal{T}_N \subset (0,\infty)$ are chosen such that $\mathcal{T}_{N+1} \supset \mathcal{T}_N$ for all $N \in \mathbb{N}$, and $\mathcal{T}_N$ is an $(1/N)$-cover of $[0,N]$, that is $\min_{t\in\mathcal{T}_N} |t-s| < 1/N$ for all $s \in [0,N]$. Then $\mathrm{Fr}^{(N)}(\omega) \to \mathrm{Fr}(\omega)$ as $N \to \infty$ in $\mathbb{F}$ for all $\omega \in \Omega$. Since $\mathrm{Fr}^{(N)}$ is measurable, this implies that Fr is measurable.

Proof. Let $\omega \in \Omega$ with $\mathrm{Fr}(\omega) = (\Pi_0^F, \dots, \Pi_m^F)$ for some $m \in \mathbb{N}$, and denote the jump times of ω by $0 =: t_0 < t_1 < \ldots < t_m$. Then there exists $N \in \mathbb{N}$ such that $(t_{i-1}, t_i) \cap \mathcal{T}_N \neq \emptyset$ for all $i \in [N]$, and therefore $\mathrm{Fr}^{(N)}(\omega) = \mathrm{Fr}(\omega)$. Since the sequence $(\mathcal{T}_N)$ is nested, $\mathrm{Fr}^{(N')}(\omega) = \mathrm{Fr}(\omega)$ for all $N' \geq N$. □

Lemma II.A.2. *For $F \in \mathbb{F}$, the map*

$$\mathrm{Tm}_F \colon \Omega_0 \supset \{\mathrm{Fr} = F\} \to \mathrm{tm}(F); \quad (\Pi_t) \mapsto \left[\Pi \mapsto \inf\{t > 0 \colon \Pi_t = \Pi\}\right]$$

is measurable.

Proof. Say $F = (\Pi_0^F, \dots, \Pi_m^F)$ with $m \geq 1$. By definition of the topology on $\mathrm{tm}(F)$, measurability of Tm_F is equivalent to measurability of $\mathrm{Tm}_F(\cdot)(\Pi_i^F)$ for all $i \in [m]$. Then for $i \in [m]$ and $t > 0$,

$$\left\{\omega \colon \mathrm{Tm}_F(\omega)(\Pi_i^F) \leq t\right\} = \left\{\omega(t) \leq \Pi_i^F\right\} = \bigcup_{\Pi \leq \Pi_i^F} \{\omega(t) = \Pi\},$$

is measurable as a finite union of measurable sets. □

With this definition, $\mathrm{Tm} \colon \Omega_0 \to \mathrm{tm}(\mathbb{F})$ is given by $\omega \mapsto (\mathrm{Fr}(\omega), \mathrm{Tm}_{\mathrm{Fr}(\omega)}(\omega))$.

Lemma II.A.3. *The map $\mathrm{Tm} \colon \Omega_0 \to \mathrm{tm}(\mathbb{F})$ is bijective and bimeasurable.*

Proof. Bijectivity is obvious. By definition of the topology on $\mathrm{tm}(\mathbb{F})$, measurability of Tm is equivalent to measurability of Tm_F for all $F \in \mathbb{F}$, which was proved above. For measurability of the inverse, recall that the topology on Ω_0 is generated by cylinders, i.e. sets of the form $\{\omega \colon \omega(t) = \Pi\}$ for $t \geq 0$ and $\Pi \in \mathcal{P}$. The preimage of such a set under the inverse of Tm is

$$\bigcup_{\substack{F_1 \in \mathbb{F} \\ \mathrm{rt}(F_1) = \Pi}} \bigcup_{\substack{F_2 \in \mathbb{F} \\ \mathrm{lf}(F_2) = \Pi}} \{(F_1 F_2, (\tau_1/t/\tau_2)) \colon \tau_1 \in \mathrm{tm}(F_1), \tau_2 \in \mathrm{tm}(F_2), \tau_1(\mathrm{rt}(F_1)) \leq t\}$$

(recall the definition of $(\tau_1/t/\tau_2)$ from (3.1) in the main paper), which is a countable union over sets that are measurable by the definition of the topologies on $\mathrm{tm}(\mathbb{F})$ and $\mathrm{tm}(F)$, $F \in \mathbb{F}$. □

Lemma II.A.4. *The map* $\mathrm{Sp} \colon \Omega \to \bigcup_{F \in \mathbb{F}} \mathrm{sp}(F)$ *is measurable.*

Proof. It is enough to show that for any $F \in \mathbb{F}$ and $u \in \mathrm{nd}^\circ(F)$, the map $\{\mathrm{Fr} = F\} \to E$ given by $\omega = (\boldsymbol{x}_t) \mapsto \boldsymbol{x}_{\mathrm{Tm}_F(\omega)_u}(u)$ is measurable. It can be written as the composition of the three maps

$$\begin{array}{ccccccc} \{\mathrm{Fr} = F\} & \to & \{\mathrm{Fr} = F\} \times (0,\infty) & \to & \mathcal{X} & \to & E \\ \hline \omega & \mapsto & (\omega, \mathrm{Tm}_F(\omega)(u)) & & & & \\ & & (\omega, t) & \mapsto & \omega(t) & & \\ & & & & (\Pi, \boldsymbol{x}) & \mapsto & \begin{cases} \boldsymbol{x}(u), & u \in \Pi, \\ x_0, & \text{else}, \end{cases} \end{array}$$

for some fixed, arbitrary $x_0 \in E$. The first map is measurable by Lemma II.A.2, the second by definition of the product topology, and the third because it is continuous. □

Lemma II.A.5. *The map* $\mathrm{Dec} \colon \Omega \to \mathrm{dec}(\mathbb{F})$ *is measurable.*

Proof. Follows from Lemmas II.A.2 and II.A.4. □

II.B Disintegration of Brownian Motion (Lemma 3.22)

A central ingredient in the characterisation of sampling consistency in Section 3 of the main paper was Lemma 3.22. It essentially says that, under some assumptions, if (T, Z) is a $(0,\infty) \times E$ valued random variable such that the law of a Brownian bridge from $(0,0)$ to (T,Z), followed by a Brownian motion starting at (T,Z), is the same as that of a Brownian motion starting in $(0,0)$, then $Z \sim \mathcal{N}(0,T)$ conditional on T (see Fig. II.B.1 for an illustration). Recall that for finite measures m_1, m_2, we write $m_1 \sim m_2$ if $m_1 = c m_2$ for some $c > 0$. To keep the appendix self-contained, we restate the Lemma here.

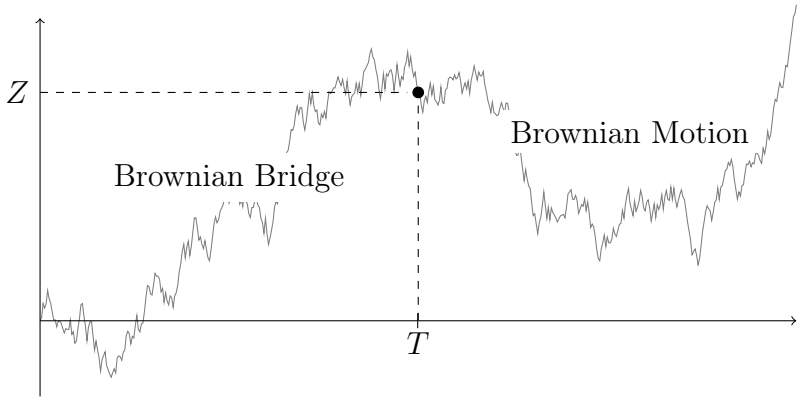

Figure II.B.1: Illustration for Lemma 3.22.

Lemma 3.22. *Suppose $x_0 \in E$, $s_0 > 0$, $f\colon E \times (s_0, \infty) \to (0, \infty)$ is continuous, and μ is a finite measure on E, and that $\int_{s_0}^{\infty} \int_E f(x,s)\mu(\mathrm{d}x)\,\mathrm{d}s < \infty$. Then*

$$B^{(x_0,s_0)+}(\cdot) \sim \iint B^{(x_0,s_0)\to(x,s)+}(\cdot) f(x,s)\mu(\mathrm{d}x)\,\mathrm{d}s$$

if and only if $f(x,s)\mu(\mathrm{d}x) \sim p_{s-s_0}(x - x_0)\,\mathrm{d}x$ for all $s > s_0$. If further $x_1 \in E$, $s_1 > s_0$, and $f\colon E \times (s_0, s_1) \to (0, \infty)$ is continuous, then

$$B^{(x_0,s_0)\to(x_1,s_1)}(\cdot) \sim \iint B^{(x_0,s_0)\to(x,s)\to(x_1,s_1)}(\cdot) f(x,s)\mu(\mathrm{d}x)\,\mathrm{d}s$$

if and only if $f(x,s)\mu(\mathrm{d}x) \sim p_{s-s_0}(x - x_0)p_{s_1-s}(x_1 - x)\,\mathrm{d}x$ for all $s \in (s_0, s_1)$.

Proof. The "if" direction is clear in both cases: for fixed $s > s_0$, if we sample $z \sim \mathcal{N}(0, s)$ and, independent of z, a Brownian bridge from (s_0, x_0) to (s, z) followed by a Brownian motion started from (s, z), then the law is the same as that of a Brownian motion started from (s_0, x_0). Similarly for the second claim. From now assume that $d = 1$, so E is the one-dimensional torus. The proof is analogous in higher dimensions. We start with the first statement. Without loss of generality $x_0 = s_0 = 0$ and $\iint f(x,s)\mu(\mathrm{d}x)\,\mathrm{d}s = 1$. A convenient reformulation is obtained by putting $g(s) \coloneqq \int f(x,s)\mu(\mathrm{d}x)$ and $\nu_s(\mathrm{d}x) \coloneqq f(x,s)\mu(\mathrm{d}x)/g(s)$. Dominated convergence can be applied to show that g is continuous, because f assumes a unique finite maximum on the compact set $E \times [s,t]$ for any $0 < s < t$. By a similar argument, ν_s is continuous in s w.r.t. the topology of weak convergence of probability measures. An argument bijecting continuous paths in E and $\mathbb{R}$ shows that it suffices to prove the analogous statement in $\mathbb{R}$, where the probability measure ν_s on the torus is replaced by the probability measure $\frac{\widetilde{p}_s(x)}{p_s(\pi(x))}\nu_s(\mathrm{d}\pi(x))$ on $\mathbb{R}$, which is continuous in s w.r.t. all Wasserstein distances. Here $\pi\colon \mathbb{R} \to E$ is the natural projection characterised by $\pi(x) - x \in \mathbb{Z}$, and $\widetilde{p}_s$ is the heat kernel on $\mathbb{R}$.

So let $g\colon (0,\infty)\to(0,\infty)$ with $\int g(s)\,\mathrm{d}s = 1$ be continuous, ν_s for $s>0$ be probability measures on $\mathbb{R}$, continuous in s w.r.t. the 1-Wasserstein distance, such that

$$B^{(0,0)+}(\cdot) = \iint B^{(0,0)\to(x,s)+}(\cdot)\nu_s(\mathrm{d}x)g(s)\,\mathrm{d}s.$$

Denote by $\varphi_s(u) \coloneqq \int \mathrm{e}^{\mathrm{i}ux}\nu_s(\mathrm{d}x)$ the characteristic function of ν_s, which is bounded and differentiable in u, and both φ_s and $\partial_u\varphi_s$ are jointly continuous in u and s by assumptions on ν_s. This will justify all interchanges of differentiation and integration that follow. Evaluating the equation at time $t>0$ and taking characteristic functions on both sides gives, after a short calculation,

$$\mathrm{e}^{-tu^2/2} = \int_0^t \varphi_s(u)\mathrm{e}^{-(t-s)u^2/2}g(s)\,\mathrm{d}s + \int_t^\infty \varphi_s(ut/s)\mathrm{e}^{-t(s-t)u^2/(2s)}g(s)\,\mathrm{d}s.$$

Cancelling $\mathrm{e}^{-tu^2/2}$ on both sides and taking a derivative in t gives

$$\begin{aligned} 0 &= \int_t^\infty \partial_t\left(\varphi_s\left(\tfrac{ut}{s}\right)\mathrm{e}^{(ut)^2/(2s)}\right)g(s)\,\mathrm{d}s \\ &= \frac{u}{t}\partial_u\left(\int_t^\infty \varphi_s\left(\tfrac{ut}{s}\right)\mathrm{e}^{(ut)^2/(2s)}g(s)\,\mathrm{d}s\right). \end{aligned} \tag{II.B.1}$$

This implies that $\int_t^\infty \varphi_s(ut/s)\mathrm{e}^{(ut)^2/(2s)}g(s)\,\mathrm{d}s$ is constant in u, so letting $u\to 0$ for fixed $t>0$ implies by dominated convergence that

$$\int_t^\infty g(s)\,\mathrm{d}s = \int_t^\infty \varphi_s(ut/s)\mathrm{e}^{(ut)^2/(2s)}g(s)\,\mathrm{d}s,$$

for every $u,t>0$. Taking a derivative in t for fixed u gives, reusing (II.B.1),

$$-g(t) = -g(t)\varphi_t(u)\mathrm{e}^{u^2t/2} + \underbrace{\int_t^\infty \partial_t\left(\varphi_s(ut/s)\mathrm{e}^{(ut)^2/(2s)}\right)g(s)\,\mathrm{d}s}_{=0},$$

that is $\varphi_t(u) = \mathrm{e}^{-u^2t/2}$ and therefore $\nu_t = \mathcal{N}(0,t)$.

Consider the second statement with $x_0 = x_1 = s_0 = 0$ and $s_1 = 1$, then by similar arguments we need to prove that

$$B^{(0,0)\to(0,1)}(\cdot) = \int_0^1\int_{\mathbb{R}} B^{(0,0)\to(x,s)\to(1,0)}(\cdot)\nu_s(\mathrm{d}x)g(s)\,\mathrm{d}s \tag{II.B.2}$$

implies $\nu_s = \mathcal{N}(0, s(1-s))$, with assumptions on ν_s and g analogous to before. Consider the measurable map $C([0,1],\mathbb{R})\to C([0,\infty),\mathbb{R})$ that maps a path y to $t\mapsto(1+t)y(\frac{t}{1+t})$; it sends $B^{(0,0)\to(0,1)}$ to $B^{(0,0)+}$, and $B^{(0,0)\to(x,s)\to(0,1)}$ to $B^{(0,0)\to(\frac{x}{1-s},\frac{s}{1-s})+}$ for $x\in\mathbb{R}$, $s\in(0,1)$. The easiest way to see this is that the pushforward in both cases is a Gaussian process with the correct mean and covariance. Applying this map to (II.B.2) and substituting $u = \frac{s}{1-s}$ gives

$$B^{(0,0)+}(\cdot) = \int_0^\infty\int_{\mathbb{R}} B^{(0,0)\to((1+u)x,u)+}(\cdot)\nu_{u/(1+u)}(\mathrm{d}x)\frac{g(\frac{u}{1+u})}{(1+u)^2}\,\mathrm{d}u.$$

By what we've already proved this implies $\nu_{u/(1+u)}(\mathrm{d}x) \sim p_u(\frac{x}{1+u})\,\mathrm{d}x$, that is

$$\nu_s(\mathrm{d}x) \sim p_{s/(1-s)}(\frac{x}{1-s})\,\mathrm{d}x \sim p_{s(1-s)}(x)\,\mathrm{d}x.$$

□

II.C Characterisation of Label Invariance (Lemma 3.13)

Recall the following characterisation of label invariance of a Brownian spatial coalescent, stated as Lemma 3.13 in the main paper.

Lemma 3.13. *The Brownian spatial coalescent with transition measures $\boldsymbol{\nu} \in R^\circ$ is label invariant if and only if for all $\Pi_0, \Pi_1 \in \mathcal{P}$ of equal size and $\iota \in \mathcal{I}(\Pi_0, \Pi_1)$,*

$$\forall \Pi' < \Pi_0 \colon \iota \# \nu_{\Pi_0,\Pi'} = \nu_{\iota(\Pi_0),\iota(\Pi')}.$$

In that case, there exists for every $(n, \vec{k}) = (n, k_1, \ldots, k_m) \in M$ a finite measure $\nu_{n,\vec{k}}$ on $E_\circ^m$ such that for every $(n, \vec{k})$-merger (Π, Π'),

$$\kappa_{\Pi,\Pi'} \# \nu_{\Pi,\Pi'} = \nu_{n,\vec{k}}.$$

If $k_i = k_j$, then $\nu_{n,\vec{k}}$ is symmetric in the ith and jth coordinate.

Proof. Let $\boldsymbol{\nu} \in R$ be the transition measures of a Brownian spatial coalescent. We assume that $|\nu_\Pi| > 0$ for all $\Pi \in \mathcal{P}$ with $|\Pi| \geq 2$ to slightly simplify the proof. Let $\Pi_0, \Pi_1 \in \mathcal{P}$ be of equal size. Note that the right-hand condition is equivalent by a simple induction to $\iota \# \boldsymbol{\nu}_F = \boldsymbol{\nu}_{\iota(F)}$ for all $\iota \in \mathcal{I}(\Pi_0, \Pi_1)$ and $F \in \mathbb{F}(\Pi_0)$. Let $\iota \in \mathcal{I}(\Pi_0, \Pi_1)$, then (2.3) from the main paper (the equation characterising label invariance) for $\boldsymbol{x} \in E_\circ^{\Pi_0}$ is equivalent to $P^{\iota(\boldsymbol{x})}(F^\star \in \cdot) = P^{\boldsymbol{x}}(\iota(F^\star) \in \cdot)$, that is

$$f_{\boldsymbol{\nu}}(\tau, \xi \,|\, \iota(\boldsymbol{x}))\boldsymbol{\nu}_F(\mathrm{d}\xi)\,\mathrm{d}\tau = f_{\boldsymbol{\nu}}(\iota^{-1}(\tau), \iota^{-1}(\xi) \,|\, \boldsymbol{x})(\iota \# \boldsymbol{\nu}_{\iota^{-1}(F)})(\mathrm{d}\xi)\,\mathrm{d}\tau$$

as measures on $\mathrm{dec}(F)$ for every $F \in \mathbb{F}(\Pi_1)$, which in turn is equivalent to

$$\prod_{(\Pi,\Pi') \in F} \mathrm{e}^{-|\nu_\Pi|(\tau_{\Pi'} - \tau_\Pi)} \boldsymbol{\nu}_F(\mathrm{d}\xi)\,\mathrm{d}\tau = \prod_{(\Pi,\Pi') \in F} \mathrm{e}^{-|\nu_{\iota^{-1}(\Pi)}|(\tau_{\Pi'} - \tau_\Pi)} (\iota \# \boldsymbol{\nu}_{\iota^{-1}(F)})(\mathrm{d}\xi)\,\mathrm{d}\tau \tag{II.C.1}$$

for all $F \in \mathbb{F}(\Pi_1)$. Clearly $\boldsymbol{\nu}_F = \iota \# \boldsymbol{\nu}_{\iota^{-1}(F)}$ for all $F \in \mathbb{F}(\Pi_1)$ implies (II.C.1). If label invariance holds, then integrating out $\xi \in \mathrm{sp}(F)$ and varying τ in (II.C.1) implies $|\nu_\Pi| = |\nu_{\iota^{-1}(\Pi)}|$ for all $(\Pi, \Pi') \in F$, which back into (II.C.1) implies $\iota \# \boldsymbol{\nu}_{\iota^{-1}(F)} = \boldsymbol{\nu}_F$.

Now suppose label invariance holds, and let $n \geq 2$ and $k_1 \geq \ldots \geq k_m \geq 2$ with $\sum_i k_i \leq n$. For existence of $\nu_{n,\vec{k}}$ it suffices to show that $\kappa_{\Pi,\Pi'} \# \nu_{\Pi,\Pi'}$ is the same for every $(n, \vec{k})$-merger (Π, Π'). Let (Π_0, Π_0') and (Π_1, Π_1') be $(n, \vec{k})$-mergers, and abbreviate

$\kappa_i \coloneqq \kappa_{\Pi_i,\Pi_i'}$ and $\nu_i \coloneqq \nu_{\Pi_i,\Pi_i'}$ for $i = 1, 2$. Fix some $\iota \in \mathcal{I}(\Pi_0, \Pi_1)$ for which $\iota(\Pi_0') = \Pi_1'$, then

$$\kappa_1 \# \nu_1 = \kappa_1 \# (\iota \# \nu_0) = (\kappa_1 \circ \iota) \# \nu_0. \tag{II.C.2}$$

Here $\kappa_1 \circ \iota \colon E_\circ^{\Pi_0' \setminus \Pi_0} \to E_\circ^m$ depends on the choices made for κ_0 and ι, but is always of the same form as κ_0, that is $\boldsymbol{x} \mapsto \boldsymbol{x} \circ \ell$ for some bijection $\ell \colon [m] \to \Pi' \setminus \Pi$ such that $\ell(i)$ is the union of k_i disjoint elements of Π_0. If we show that $(\,\cdot \circ \ell) \# \nu_0$ is the same for each such choice for ℓ, then both equality of (II.C.2) to $\kappa_0 \# \nu_0$, and the claimed symmetry of $\nu_{n,\vec{k}}$ follow. If two such bijections $\ell, \ell' \colon [m] \to \Pi' \setminus \Pi$ are given, then there is a permutation π of $\Pi' \setminus \Pi$ such that $\pi \circ \ell = \ell'$, and $\pi(u)$ for $u \in \Pi' \setminus \Pi$ is the result of merging as many blocks from Π as u.

□

II.D Conditional Kernel of the Brownian Spatial Coalescent (Lemma 2.9)

Recall Lemma 2.9 from the main paper, which states that the kernel $K_{\boldsymbol{x}}(F^\star, \cdot)$, which describes the law of the Brownian spatial coalescent started at $\boldsymbol{x}$ conditioned on $\mathrm{Dec} = F^\star$, is well-defined and well-behaved.

Lemma 2.9. *Given $F^\star = (F, \tau, \xi) \in \mathrm{Dec}(\mathbb{F})$ and $\boldsymbol{x} \in E_\circ^{\mathrm{lf}(F)}$, there is a unique law $K_{\boldsymbol{x}}(F^\star, \cdot) \in \mathcal{M}_1(\Omega)$ under which $\boldsymbol{X}_0 = \boldsymbol{x}$ and $\mathrm{Dec} = F^\star$ a.s., and*

$$\big(\mathrm{Path}_u^{F^\star}(\boldsymbol{X}) \colon u \in \mathrm{nd}(F)\big)$$

is a family of independent random variables such that the law of $\mathrm{Path}_u^{F^\star}(\boldsymbol{X})$ is $B^{(\tau_u, (\boldsymbol{x}\xi)_u)+}$ if $\mathrm{pr}_F(u) = \emptyset$, and $B^{(\tau_u, (\boldsymbol{x}\xi)_u) \to (\tau_{\mathrm{pr}_F(u)}, \xi_{\mathrm{pr}_F(u)})}$ if $\mathrm{pr}_F(u) \neq \emptyset$. There exists an extension of these laws to a family $(K_{\boldsymbol{x}}(F^\star, \cdot) \colon \boldsymbol{x} \in \mathcal{X}, F^\star \in \mathrm{dec}(\mathbb{F}))$ such that

$$\mathcal{X} \times \mathrm{dec}(\mathbb{F}) \to \mathcal{M}_1(\Omega); \quad (\boldsymbol{x}, F^\star) \mapsto K_{\boldsymbol{x}}(F^\star, \cdot)$$

is continuous. In particular, $(\boldsymbol{x}, F^\star) \mapsto K_{\boldsymbol{x}}(F^\star, A)$ is measurable for any measurable $A \subset \Omega$.

Proof of Lemma 2.9. Uniqueness follows because the collection of maps $(\mathrm{Path}_u^{F^\star} \colon u \in \mathrm{nd}(F))$ uniquely determines an element of $\{\mathrm{Dec} = F^\star\}$. For existence and the remaining properties, let

$$\Big((B_t^u)_{t \in [0,1]} \colon u \subset \mathbb{N}\Big), \qquad \Big((W_t^u)_{t \geq 0} \colon u \subset \mathbb{N}\Big),$$

be independent families of, respectively, i.i.d. standard Brownian bridges and i.i.d. standard Brownian motions, defined on some common probability space $(S, \mathcal{A}, \mathbb{P})$. For any

$\boldsymbol{x}$ and $F^\star = (F, \tau, \xi) \in \mathrm{dec}(\mathbb{F})$, we define a random element $(\Pi_t, \boldsymbol{X}_t) \in \{\mathrm{Dec} = F^\star\}$ with $\boldsymbol{X}_0 = \boldsymbol{x}$ in the following way: $(\Pi_t) \coloneqq \mathrm{Tm}^{-1}(F, \tau)$, and if $u \in \mathrm{nd}(F)$ and $t \in [\tau_u, \tau_{\mathrm{pr}_F(u)})$,

$$\boldsymbol{X}_t(u) \coloneqq \begin{cases} (\boldsymbol{x}\xi)_u + (\xi_{\mathrm{pr}_F(u)} - (\boldsymbol{x}\xi)_u) B^u_{(t-\tau_u)/(\tau_{\mathrm{pr}_F(u)} - \tau_u)}, & \mathrm{pr}_F(u) \neq \emptyset, \\ (\boldsymbol{x}\xi)_u + W^u_{t-\tau_u}, & \text{else.} \end{cases}$$

Then $\boldsymbol{X}_t \in E^{\Pi_t}$ for all $t \geq 0$, and by standard properties of Brownian motion, there is a $\mathbb{P}$-null set $N_{\boldsymbol{x}, F^\star} \in \mathcal{A}$ outside of which $\boldsymbol{X}_t \in E_\circ^{\Pi_t}$ for all $t \geq 0$. On $N_{\boldsymbol{x}, F^\star}$, we let $(\Pi_t, \boldsymbol{X}_t)$ be defined by $(\Pi_t) = \mathrm{Tm}^{-1}(F, \tau)$, and $\boldsymbol{X}_t$ linearly interpolates between nodes of the forest (any other particular element of $\{\mathrm{Dec} = F^\star\}$ starting at $\boldsymbol{x}$ would do). Then the law $K_{\boldsymbol{x}}(F^\star, \cdot)$ of $(\Pi_t, \boldsymbol{X}_t)$ has the stated properties.

Let $F = (\Pi_0, \dots, \Pi_m) \in \mathbb{F}$ be fixed, $\tau^n \to \tau$ in $\mathrm{tm}(F)$ (recall from Definition 2.7 that this is in the upper limit topology), $\xi^n \to \xi$ in $\mathrm{sp}(F)$, and $\boldsymbol{x}_n \to \boldsymbol{x}$ in $E_\circ^{\mathrm{lf}(F)}$, $F_n^\star \coloneqq (F, \tau^n, \xi^n)$, $F^\star \coloneqq (F, \tau, \xi)$, and denote the associated paths by $(\Pi_t^n, \boldsymbol{X}_t^n)$, $n \in \mathbb{N}$, and $(\Pi_t, \boldsymbol{X}_t)$. This gives a coupling of the laws $K_{\boldsymbol{x}_n}(F_n^\star, \cdot)$, $n \in \mathbb{N}$, and $K_{\boldsymbol{x}}(F^\star, \cdot)$ on the probability space $(S, \mathcal{A}, \mathbb{P})$, so almost-sure convergence on S implies weak convergence of the laws. Put $\tau_i \coloneqq \tau(\Pi_i)$ for $i \in \{0, \dots, m\}$, and $\tau_{m+1} \coloneqq \infty$, and similarly for τ_i^n, $n \in \mathbb{N}$. We show that $(\Pi^n, \boldsymbol{X}^n) \to (\Pi, \boldsymbol{X})$ in Ω holds outside $N_{\boldsymbol{x}, F^\star} \cup \bigcup_{n \in \mathbb{N}} N_{\boldsymbol{x}_n, F_n^\star}$. Let $t > 0$. By construction, the only discontinuities of $(\Pi, \boldsymbol{X})$ are the jump times of (Π_t). If t is no such jump time, say $\tau_i < t < \tau_{i+1}$ for some $i \in \{0, \dots, m\}$, then for large n also $\tau_i^n < t < \tau_{i+1}^n$, so $\Pi^n(t) = \Pi_i = \Pi(t)$. If t is a jump time, say $t = \tau_i$ for some $i \in [m]$, then (recalling that $\mathrm{tm}(F)$ is equipped with the upper limit topology) $\tau_i^n \leq \tau_i = t < \tau_{i+1}^n$, and hence $\Pi^n(t) = \Pi(t)$, for sufficiently large n. Now say $\Pi \coloneqq \Pi(t)$, and without loss of generality $\Pi^n(t) = \Pi$ for all $n \in \mathbb{N}$. It remains to show that then $\boldsymbol{X}^n(t)_u \to \boldsymbol{X}(t)_u$ for any $u \in \Pi$. If $\mathrm{pr}_F(u) = \emptyset$, then

$$|\boldsymbol{X}^n(t)_u - \boldsymbol{X}(t)_u| \leq |(\boldsymbol{x}_n \xi^n)_u - (\boldsymbol{x}\xi)_u| + \left| W^u_{t-\tau_u^n} - W^u_{t-\tau_u} \right| \overset{n \to \infty}{\longrightarrow} 0,$$

and similarly if $\mathrm{pr}_F(u) \neq \emptyset$.

We have proved that whenever $F_n^\star \to F^\star$ in $\mathrm{dec}(\mathbb{F})$, and $\boldsymbol{x}_n \to \boldsymbol{x}$ in $E_\circ^{\mathrm{lf}(F)}$, then $K_{\boldsymbol{x}_n}(F_n^\star, \cdot) \to K_{\boldsymbol{x}}(F^\star, \cdot)$ weakly. Now fix arbitrary choices $(\boldsymbol{x}_F \in E_\circ^{\mathrm{lf}(F)})_{F \in \mathbb{F}}$, and put for $F \in \mathbb{F}$ and $\boldsymbol{x} \in \mathcal{X}$ with $\boldsymbol{x} \notin E_\circ^{\mathrm{lf}(F)}$,

$$K_{\boldsymbol{x}}(F^\star, \cdot) \coloneqq K_{\boldsymbol{x}_F}(F^\star, \cdot).$$

This defines $K_{\boldsymbol{x}}(F, \cdot)$ for any $\boldsymbol{x} \in \mathcal{X}$ and $F \in \mathbb{F}$. Now let $F_n^\star \to F^\star$ in $\mathrm{dec}(\mathbb{F})$, and $\boldsymbol{x}_n \to \boldsymbol{x}$ in $\mathcal{X}$. If $\boldsymbol{x} \in E_\circ^{\mathrm{lf}(F)}$, then for large n we must have $\mathrm{dom}(\boldsymbol{x}_n) = \mathrm{dom}(\boldsymbol{x}) = \mathrm{lf}(F)$ and $F_n = F$, so also $\boldsymbol{x}_n \in E_\circ^{\mathrm{lf}(F_n)}$, in which case $K_{\boldsymbol{x}_n}(F_n^\star, \cdot) \to K_{\boldsymbol{x}}(F^\star, \cdot)$ is already proved. If $\boldsymbol{x} \notin E_\circ^{\mathrm{lf}(F)}$, then by the same argument also $\boldsymbol{x}_n \notin E_\circ^{\mathrm{lf}(F_n)}$ and $F_n = F$ for large n, so

$$K_{\boldsymbol{x}_n}(F_n^\star, \cdot) = K_{\boldsymbol{x}_{F_n}}(F_n^\star, \cdot) = K_{\boldsymbol{x}_F}(F_n^\star, \cdot) \overset{n \to \infty}{\longrightarrow} K_{\boldsymbol{x}_F}(F^\star, \cdot) = K_{\boldsymbol{x}}(F^\star, \cdot),$$

by what we have already proved, since $\boldsymbol{x}_F \in E_\circ^{\mathrm{lf}(F)} = E_\circ^{\mathrm{lf}(F_n)}$ for large n.

It remains to show that $(\boldsymbol{x}, F^\star) \mapsto K_{\boldsymbol{x}}(F^\star, A)$ is measurable for measurable $A \subset \Omega$. The set of such A is a λ-system, so by the π-λ lemma it suffices to show the claim for A taken from a π-system that generates the σ-algebra on Ω, such as the family of closed sets. By what we have already showed and the Portmanteau theorem, if $A \subset \Omega$ is closed, and $\boldsymbol{x}_n \to \boldsymbol{x}$ and $F_n^\star \to F^\star$,

$$\overline{\lim_{n\to\infty}}\, K_{\boldsymbol{x}_n}(F_n^\star, A) \le K_{\boldsymbol{x}}(F^\star, A),$$

which implies that $(\boldsymbol{x}, F^\star) \mapsto K_{\boldsymbol{x}}(F^\star, A)$ is upper semi-continuous and hence measurable. □

II.E N^ν is finite and continuous (Lemma 2.12)

In this section we prove Lemma 2.12 from the main paper.

Lemma 2.12. *If $\boldsymbol{\nu}(\mathrm{d}\xi) \sim \mathrm{d}\xi$ then $N^{\boldsymbol{\nu}}$ is continuous (in particular finite).*

Throughout the section, fix $\boldsymbol{\nu} \in R$ with $\boldsymbol{\nu}(\mathrm{d}\xi) = \boldsymbol{\lambda}\,\mathrm{d}\xi$ for some $\boldsymbol{\lambda} \in R^\circ$, that is $\nu_{\Pi,\Pi'}(\mathrm{d}\xi) = \lambda_{\Pi,\Pi'}\,\mathrm{d}\xi$ for every $\Pi, \Pi' \in \mathcal{P}$, $\Pi' < \Pi$. Then $N_{\boldsymbol{\nu}}^F$ for $F \in \mathbb{F}$ is a map $E_\circ^{\mathrm{lf}(F)} \to (0,\infty)$ defined by

$$N_{\boldsymbol{\nu}}^F(\boldsymbol{x}) = \int_{\mathrm{tm}(F)} \int_{\mathrm{sp}(F)} f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(\tau) f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x})\,\mathrm{d}\xi\,\mathrm{d}\tau.$$

The goal of this section is to show that $N_{\boldsymbol{\nu}}^F$ is finite and continuous.

Lemma II.E.1. *Let $F \in \mathbb{F}$. For fixed $\tau \in \mathrm{tm}(F)$, the map*

$$E_\circ^{\mathrm{lf}(F)} \to (0,\infty); \quad \boldsymbol{x} \mapsto \int_{\mathrm{sp}(F)} f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x})\,\mathrm{d}\xi$$

is continuous.

Proof. For fixed $\tau \in \mathrm{tm}(F)$, $f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x})$ is continuous in $\boldsymbol{x}$ for every $\xi \in \mathrm{sp}(F)$, and bounded uniformly in $\boldsymbol{x}$ and ξ, so the statement follows from the dominated (or bounded) convergence theorem. □

Let $\Pi \in \mathcal{P}$. Then $E_\circ^\Pi$ is an open subset of E^Π, so the distance between some $\boldsymbol{x} \in E^\Pi$ to $E^\Pi \setminus E_\circ^\Pi$ is well-defined and denoted $\varepsilon_d(\boldsymbol{x})$. More explicitly,

$$\varepsilon_d(\boldsymbol{x}) = \begin{cases} \min_{u\neq v\in\Pi} \rho(\boldsymbol{x}_u, \boldsymbol{x}_v), & d \ge 2, \\ \max_{u_0,v_0\in\Pi} \min_{u\neq v\in\Pi,\{u,v\}\neq\{u_0,v_0\}} \rho(\boldsymbol{x}_u, \boldsymbol{x}_v), & d = 1, \end{cases} \tag{II.E.1}$$

where ρ denotes the Euclidean metric on E, see (2.1). If $d \ge 2$ then $\varepsilon_d(\boldsymbol{x})$ is the smallest, if $d = 1$ it is the second smallest distance among all pairs of two particles.

Lemma II.E.2. *Let $F \in \mathbb{F}$. Then there exists $C = C(F, \boldsymbol{\lambda}, d) > 0$ such that, for any $A \subset E_\circ^{\mathrm{lf}(F)}$,*

$$\int_{\mathrm{tm}(F)} f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(\tau) \sup_{\boldsymbol{x} \in A} \left(\int_{\mathrm{sp}(F)} f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x}) \, \mathrm{d}\xi \right) \mathrm{d}\tau \le C \left[1 + \left(\inf_{\boldsymbol{x} \in A} \varepsilon_d(\boldsymbol{x}) \right)^{-C} \right]. \tag{II.E.2}$$

In particular, $N_{\boldsymbol{\nu}}^F \colon E_\circ^{\mathrm{lf}(F)} \to (0, \infty)$ is continuous and, specialising A in (II.E.2) *to a singleton,*

$$N_{\boldsymbol{\nu}}^F(\boldsymbol{x}) \le C \left(1 + \varepsilon_d(\boldsymbol{x})^{-C} \right), \qquad \boldsymbol{x} \in E_\circ^{\mathrm{lf}(F)}.$$

The continuity of $N_{\boldsymbol{\nu}}^F$ follows from (II.E.2) and Lemma II.E.1 by the dominated convergence theorem. We start by proving a version of (II.E.2) where the torus is replaced by $\mathbb{R}^d$, where a certain heat kernel trick can be used that isn't available on the torus. Then we transfer the result back to the torus. Write $\widetilde{E} = \mathbb{R}^d$, and denote by $\widetilde{p}_t \colon \mathbb{R}^d \to (0, \infty)$ for $t > 0$ the heat kernel, defined by

$$\widetilde{p}_t(x) = (2\pi t)^{-d/2} \exp\left(-\frac{|x|^2}{2t} \right), \qquad x \in \mathbb{R}^d.$$

Define $\widetilde{E}_\circ^\Pi$ for $\Pi \in \mathcal{P}$ analogously to $E_\circ^\Pi$, and for a forest $F \in \mathbb{F}$, denote by $\widetilde{\mathrm{sp}}(F)$ the set of maps $\xi \colon \mathrm{nd}^\circ(F) \to \widetilde{E}$ such that $\xi\big|_\Pi \in \widetilde{E}_\circ^\Pi$ for all $\Pi \in F$. For $\xi \in \widetilde{\mathrm{sp}}(F)$, $\tau \in \mathrm{tm}(F)$ and $\boldsymbol{x} \in \widetilde{E}_\circ^{\mathrm{lf}(F)}$, write

$$\widetilde{f_{\mathrm{sp}}}(\xi \,|\, \tau, \boldsymbol{x}) \coloneqq \prod_{u \in \mathrm{nd}(F) \setminus \mathrm{rt}(F)} \widetilde{p}(\tau_{\mathrm{pr}_F(u)} - \tau_u, (\boldsymbol{x}\xi)_u - \xi_{\mathrm{pr}_F(u)}).$$

For every $x, y \in E$ there exists a $\vec{k}(x, y) \in \{-1, 0\}$ (not necessarily unique) such that $\rho(x, y) = |x - y + \vec{k}(x, y)|$ (on the left is the Euclidean metric on the torus E, on the right is the Euclidean norm on $\mathbb{R}^d$). Here and in the following, if we treat an element $x \in E$ as an element of $\mathbb{R}^d$, then we identify it with its unique representative in $[0, 1)^d$. Then, for some $C = C(d) > 0$,

$$\begin{aligned} p_t(x - y) &\le C(1 + \sqrt{t})^d \widetilde{p}_t(x - y + \vec{k}(x, y)) \\ &\le C(1 + \sqrt{t})^d \sum_{\vec{k} \in \{-1,0,1\}^d} \widetilde{p}_t(x - y + \vec{k}), \end{aligned} \tag{II.E.3}$$

where the first inequality follows e.g. from [1], Theorem 8.

Lemma II.E.3. *For every $F = (\Pi_0^F, \dots, \Pi_m^F) \in \mathbb{F}$ there is $C = C(d, F) > 0$ such that for every $\tau \in \mathrm{tm}(F)$ and $\boldsymbol{x} \in E_\circ^{\mathrm{lf}(F)}$,*

$$f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x}) \le C \left(1 + \sqrt{\max \tau} \right)^C \sum_{\substack{\vec{\boldsymbol{\ell}} \colon \mathrm{nd}^\circ(F) \to \{-m, \dots, m\}^d \\ \vec{\boldsymbol{k}} \colon \mathrm{lf}(F) \to \{-m, \dots, m\}^d}} \widetilde{f_{\mathrm{sp}}}(\xi + \vec{\boldsymbol{\ell}} \,|\, \tau, \boldsymbol{x} + \vec{\boldsymbol{k}}) \tag{II.E.4}$$

for all $\xi \in \mathrm{sp}(F)$, and

$$\int_{\mathrm{sp}(F)} f_{\mathrm{sp}}(\xi \,|\, \tau, \boldsymbol{x}) \, \mathrm{d}\xi \le C(1 + \sqrt{\max \tau})^C \sum_{\vec{\boldsymbol{k}} \colon \mathrm{lf}(F) \to \{-m, \dots, m\}^d} \int_{\widetilde{\mathrm{sp}}(F)} \widetilde{f_{\mathrm{sp}}}(\xi \,|\, \tau, \boldsymbol{x} + \vec{\boldsymbol{k}}) \, \mathrm{d}\xi. \tag{II.E.5}$$

Proof. Denote by $C > 0$ a constant depending only on d and F whose value may increase from line to line. Let $v \in \mathrm{rt}(F)$ with $\mathrm{ch}_F(v) \neq \emptyset$, and suppose for simplicity that $\mathrm{ch}_F(u) \neq \emptyset$ for all $u \in \mathrm{ch}_F(v)$. Then, using (II.E.3),

$$\begin{aligned}\prod_{u\in\mathrm{ch}_F(v)} p_{\tau_v-\tau_u}(\xi_u-\xi_v) &\leq \prod_{u\in\mathrm{ch}_F(v)} \left(C(1+\sqrt{\tau_v-\tau_u})^d \sum_{\vec{k}_u\in\{-1,0,1\}^d} \widetilde{p}_{\tau_v-\tau_u}(\xi_u-\xi_v+\vec{k}_u)\right)\\ &\leq C(1+\sqrt{\max\tau})^C \prod_{u\in\mathrm{ch}_F(v)} \sum_{\vec{k}_u\in\{-1,0,1\}^d} \widetilde{p}_{\tau_v-\tau_u}[(\xi_u+\vec{k}_u)-\xi_v].\end{aligned}$$

We put $\vec{\boldsymbol{\ell}}_u \coloneqq \vec{k}_u$ for $u \in \mathrm{ch}_F(v)$. Now let $u \in \mathrm{ch}_F(v)$ and suppose again for simplicity that $\mathrm{ch}_F(w) \neq \emptyset$ for all $w \in \mathrm{ch}_F(u)$. Then

$$\begin{aligned}&\prod_{w\in\mathrm{ch}_F(u)} p_{\tau_u-\tau_w}(\xi_w-\xi_u)\\ &\quad= \prod_{w\in\mathrm{ch}_F(u)} p_{\tau_u-\tau_w}[\xi_w+\vec{\boldsymbol{\ell}}_u-(\xi_u+\vec{\boldsymbol{\ell}}_u)]\\ &\quad\leq C(1+\sqrt{\max\tau})^C \prod_{w\in\mathrm{ch}_F(u)} \sum_{\vec{k}_w\in\{-1,0,1\}^d} \widetilde{p}_{\tau_u-\tau_w}[(\xi_w+\vec{k}_w+\vec{\boldsymbol{\ell}}_u)-(\xi_u+\vec{\boldsymbol{\ell}}_u)].\end{aligned}$$

We put $\vec{\boldsymbol{\ell}}_w \coloneqq \vec{k}_w + \vec{\boldsymbol{\ell}}_u$ for $w \in \mathrm{ch}_F(u)$. Proceeding inductively in this way down towards the leaves yields (II.E.4), from which the second claim easily follows (note that the integral on the left-hand side is w.r.t. Lebesgue measure on the torus, and on the right-hand side w.r.t. Lebesgue measure on all of $\mathbb{R}^d$).

□

This reduction lets us make use of the following statement. Define $\widetilde{\varepsilon}_d(\boldsymbol{x})$ for $\boldsymbol{x} \in \widetilde{E}^\Pi$ for some $\Pi \in \mathcal{P}$ analogously to ε_d, by replacing $\rho(\cdot,\cdot)$ in (II.E.1) with the Euclidean distance $|\cdot - \cdot|$ on $\mathbb{R}^d$.

Lemma II.E.4. *Let $F \in \mathbb{F}$ and $K > 0$. Then there exists $C = C(F, \boldsymbol{\lambda}, d, K) > 0$ such that, for any $A \subset \widetilde{E}_\circ^{\mathrm{lf}(F)}$,*

$$\int_{\mathrm{tm}(F)} f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(\tau)(1+\sqrt{\max\tau})^K \sup_{\boldsymbol{x}\in A}\left(\int_{\widetilde{\mathrm{sp}}(F)} \widetilde{f_{\mathrm{sp}}}(\xi\,|\,\tau,\boldsymbol{x})\,\mathrm{d}\xi\right)\mathrm{d}\tau \leq C\left[1+\left(\inf_{\boldsymbol{x}\in A}\widetilde{\varepsilon}_d(\boldsymbol{x})\right)^{-C}\right].$$

The proof of Lemma II.E.4 hinges on an explicit representation of $\int \widetilde{f_{\mathrm{sp}}}(\xi\,|\,\tau,\boldsymbol{x})\,\mathrm{d}\xi$ based on the following heat kernel trick.

Lemma II.E.5. *Let $x_1,\dots,x_n \in \mathbb{R}^d$ and $s_1,\dots,s_n > 0$. Then,*

$$\prod_{i=1}^n \widetilde{p}_{s_i}(x_i - z) = \frac{\widetilde{p}_s(z-\overline{x})}{\widetilde{p}_s(0)} \prod_{i=1}^n \widetilde{p}_{s_i}(x_i-\overline{x}),$$

where $\frac{1}{s} = \sum_{i=1}^n \frac{1}{s_i}$ and $\overline{x} = \sum_{i=1}^n \frac{s}{s_i} x_i$.

Proof. By definition of the heat kernel,

$$\prod_{i=1}^{n} \widetilde{p}_{s_i}(x_i - z) = \left(\prod_{i=1}^{n} (2\pi s_i)^{-d/2}\right) \exp\left(-\frac{1}{2}\sum_{i=1}^{n} \frac{(x_i - z)^2}{s_i}\right), \tag{II.E.6}$$

and

$$\begin{aligned}\sum_{i=1}^{n} \frac{(x_i - z)^2}{s_i} = \frac{z^2}{s} - 2\frac{\overline{x}z}{s} + \sum_{i=1}^{n} \frac{x_i^2}{s_i} &= \frac{1}{s}(z - \overline{x})^2 - \frac{1}{s}\overline{x}^2 + \sum_{i=1}^{n} \frac{x_i^2}{s_i} \\ &= \frac{1}{s}(z - \overline{x})^2 + \sum_{i=1}^{n} \frac{(x_i - \overline{x})^2}{s_i},\end{aligned} \tag{II.E.7}$$

where in the last step we used that

$$\sum_{i=1}^{n} \frac{(x_i - \overline{x})^2}{s_i} = \sum_{i=1}^{n} \frac{x_i^2}{s_i} - 2\overline{x}\sum_{i=1}^{n} \frac{x_i}{s_i} + \frac{1}{s}\overline{x}^2 = \sum_{i=1}^{n} \frac{x_i^2}{s_i} - \frac{1}{s}\overline{x}^2.$$

Plugging (II.E.7) back into (II.E.6) finishes the proof. □

Lemma II.E.6. *For a non-trivial forest $F \in \mathbb{F}$ and $\boldsymbol{x} \in \widetilde{E}_\circ^{\mathrm{lf}(F)}$, $\tau \in \mathrm{tm}(F)$,*

$$\int_{\widetilde{\mathrm{sp}}(F)} \widetilde{f_{\mathrm{sp}}}(\xi \,|\, \tau, \boldsymbol{x})\, \mathrm{d}\xi = \prod_{v \in \mathrm{nd}^\circ(F)} \frac{1}{\widetilde{p}_{r_v}(0)} \prod_{u \in \mathrm{ch}_F(v)} \widetilde{p}_{r_u + \tau_v - \tau_u}(\overline{x}_v - \overline{x}_u), \tag{II.E.8}$$

where $\overline{x}_u = \boldsymbol{x}_u$ and $r_u = 0$ for $u \in \mathrm{lf}(F)$ and, inductively for $v \in \mathrm{nd}^\circ(F)$,

$$\frac{1}{r_v} \coloneqq \sum_{u \in \mathrm{ch}_F(v)} \frac{1}{r_u + \tau_v - \tau_u}, \qquad \overline{x}_v \coloneqq \sum_{u \in \mathrm{ch}_F(v)} \frac{r_v}{r_u + \tau_v - \tau_u}\overline{x}_u.$$

The following hold.

(i) If $v \in \mathrm{nd}^\circ(F)$ with $\mathrm{ch}(v) \subset \mathrm{lf}(F)$, then

$$r_v = \frac{\tau_v}{|\,\mathrm{ch}_F(v)|}, \qquad \overline{x}_v = \frac{1}{|\,\mathrm{ch}_F(v)|}\sum_{u \in \mathrm{ch}_F(v)} \boldsymbol{x}_u.$$

(ii) There is $c = c(F) > 0$ such that $c\tau_v \le r_v < \tau_v$ for all $v \in \mathrm{nd}^\circ(F)$.

(iii) If $v \in \mathrm{nd}^\circ(F)$ then

$$\max_{u \in \mathrm{ch}(v)} |\overline{x}_u - \overline{x}_v| \ge \frac{1}{2} \max_{u,u' \in \mathrm{ch}(v)} |\overline{x}_u - \overline{x}_{u'}|.$$

Proof. Eq. (II.E.8) follows by applying Lemma II.E.5 inductively from leaves to root, which is tedious but straightforward.

(i) is obvious from the definitions.

(ii) If v is not a leaf and $r_u \le \tau_u$ holds for all children, then

$$\frac{1}{r_v} \ge \sum_{u\in \mathrm{ch}_F(v)} \frac{1}{\tau_u + \tau_v - \tau_u} = \frac{|\,\mathrm{ch}_F(v)|}{\tau_v} \ge \frac{2}{\tau_v} > \frac{1}{\tau_v}.$$

If v is not a leaf and $r_u \ge c\tau_u$ holds for all children, then

$$\frac{1}{r_v} \le \sum_{u\in \mathrm{ch}_F(v)} \frac{1}{c\tau_u + \tau_v - \tau_u} \le \sum_{u\in \mathrm{ch}_F(v)} \frac{1}{c(\tau_u + \tau_v - \tau_u)} = \frac{|\,\mathrm{ch}_F(v)|}{c}\frac{1}{\tau_v}.$$

(iii) If $v \in \mathrm{nd}^\circ(F)$ and $u_1, u_2 \in \mathrm{ch}(v)$, then

$$|\overline{x}_{u_1} - \overline{x}_{u_2}| \le |\overline{x}_{u_1} - \overline{x}_v| + |\overline{x}_v - \overline{x}_{u_2}| \le 2 \max_{u\in\mathrm{ch}(v)} |\overline{x}_u - \overline{x}_v|.$$

□

Proof of Lemma II.E.4. Fix a non-trivial forest $F = (\Pi_0^F, \dots, \Pi_m^F) \in \mathbb{F}$, and abbreviate $\mathrm{ch} = \mathrm{ch}_F$, $\mathrm{pr} = \mathrm{pr}_F$ and, for $\tau \in \mathrm{tm}(F)$ and $i \in [m]$, $\tau_i = \tau_{\Pi_i^F}$. Throughout the proof, $c, C > 0$ are constants that only depend on F, d, $\boldsymbol{\lambda}$, and K, and whose value may respectively decrease and increase from line to line. First let $d \ge 2$, then for $\boldsymbol{x} \in E_\circ^{\mathrm{lf}(F)}$ and $\tau \in \mathrm{tm}(F)$, using (II.E.8),

$$\int \widetilde{f_{\mathrm{sp}}}(\xi \,|\, \tau, \boldsymbol{x})\,\mathrm{d}\xi \le C \prod_{i=1}^m \underbrace{\prod_{v\in\Pi_i^F\setminus\Pi_{i-1}^F} r_v^{d/2} \prod_{u\in\mathrm{ch}(v)} \widetilde{p}_{r_u+\tau_v-\tau_u}(\overline{x}_v - \overline{x}_u)}_{=:G_i(\tau,\boldsymbol{x})}. \tag{II.E.9}$$

Denote $G^{(i)}(\tau,\boldsymbol{x}) := \prod_{j=i+1}^m G_j(\tau,\boldsymbol{x})$, and $f_{\mathrm{tm}}^{\boldsymbol{\lambda}\,(i)}(\tau) := \prod_{j=i+1}^m \lambda_{\Pi_{j-1}^F,\Pi_j^F} \mathrm{e}^{-\lambda_{\Pi_j^F}(\tau_j-\tau_{j-1})}$ for $i \in \{0,\dots,m\}$, so $f_{\mathrm{tm}}^{\boldsymbol{\lambda}\,(0)} = f_{\mathrm{tm}}^{\boldsymbol{\lambda}}$ and $f_{\mathrm{tm}}^{\boldsymbol{\lambda}\,(m)} \equiv 1$.

We show inductively that for $i \in \{1,\dots,m\}$ there exists $b_i \ge 0$ such that

$$\int f_{\mathrm{tm}}^{\boldsymbol{\lambda}\,(i)}(\tau)(1+\sqrt{\tau_m})^K \sup_{\boldsymbol{x}\in A} G^{(i)}(\tau,\boldsymbol{x})\,\mathrm{d}\tau_{i+1}\dots\mathrm{d}\tau_m \le C\tau_1^{-b_i}(1+\tau_i)^C. \tag{II.E.10}$$

For $i = m$ the LHS is equal to $(1+\sqrt{\tau_m})^K \le 2^K(1+\tau_m)^K$ and we can choose $b_m = 0$. Suppose we have proved the claim for some fixed $i \in \{2,\dots,m\}$, then by induction hypothesis

$$\begin{aligned}\int f_{\mathrm{tm}}^{\boldsymbol{\lambda}\,(i-1)}(\tau)(1+\sqrt{\tau_m})^K \sup_{\boldsymbol{x}\in A} G^{(i-1)}(\tau,\boldsymbol{x})\,\mathrm{d}\tau_i\dots\mathrm{d}\tau_m \\ \le C\tau_1^{-b_i}\int_{\tau_{i-1}}^\infty \mathrm{d}\tau_i\,(1+\tau_i)^C \mathrm{e}^{-\lambda_i(\tau_i-\tau_{i-1})} \sup_{\boldsymbol{x}\in A} \prod_{v\in\Pi_i^F\setminus\Pi_{i-1}^F} r_v^{d/2} \prod_{u\in\mathrm{ch}(v)} \widetilde{p}_{r_u+\tau_v-\tau_u}(\overline{x}_v - \overline{x}_u),\end{aligned}$$

where we abbreviated $\lambda_i := \lambda_{\Pi_i^F}$. We bound $r_v \le \tau_v = \tau_i$, and $\widetilde{p}_{r_u+\tau_v-\tau_u}(\overline{x}_v - \overline{x}_u) \le C(r_u+\tau_v-\tau_u)^{-d/2}$, and by Lemma II.E.6,

$$r_u + \tau_v - \tau_u \ge c\tau_u + \tau_v - \tau_u \ge c\tau_v \ge c\tau_1. \tag{II.E.11}$$

Let l be the total number of children of nodes $v \in \Pi_i^F \setminus \Pi_{i-1}^F$. Then our bound is no larger than $\tau_1^{-b_i}$ times

$$\begin{aligned} C \int_{\tau_{i-1}}^{\infty} \mathrm{d}\tau_i \, (1+\tau_i)^C \mathrm{e}^{-\lambda(\tau_i - \tau_{i-1})} \tau_i^{|\Pi_i^F \setminus \Pi_{i-1}^F| d/2} (c\tau_1)^{-ld/2} \\ \leq C\tau_1^{-ld/2} \int_0^{\infty} \mathrm{e}^{-\lambda_i s} (1 + \tau_{i-1} + s)^C \, \mathrm{d}s \\ \leq C\tau_1^{-ld/2} (1+\tau_{i-1})^C, \end{aligned}$$

which completes the inductive step. Using (II.E.10) with $i = 1$, recalling that $r_u = \tau_u = \tau_1$ for $u \in \mathrm{lf}(F)$, and putting $s \coloneqq \tau_1, b \coloneqq b_1, \lambda \coloneqq \lambda_1$, (II.E.9) becomes

$$\begin{aligned} \int f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(\tau)(1+\sqrt{\tau_m})^K \sup_{\boldsymbol{x} \in A} \int \widetilde{f_{\mathrm{sp}}}(\xi \mid \tau, \boldsymbol{x}) \, \mathrm{d}\xi \, \mathrm{d}\tau \\ \leq C \int_0^{\infty} s^{-b} (1+s)^C \mathrm{e}^{-\lambda s} \sup_{\boldsymbol{x} \in A} \prod_{v \in \Pi_1^F \setminus \Pi_0^F} s^{d/2} \prod_{u \in \mathrm{ch}(v)} \widetilde{p}_{r_u + s_v - s_u}(\overline{x}_v - \overline{x}_u) \, \mathrm{d}s \\ \leq C \int_0^{\infty} s^{-b} (1+s)^C \mathrm{e}^{-\lambda s} \sup_{\boldsymbol{x} \in A} \prod_{v \in \Pi_1^F \setminus \Pi_0^F} \prod_{u \in \mathrm{ch}(v)} \widetilde{p}_s(\overline{x}_v - \overline{x}_u) \, \mathrm{d}s. \end{aligned}$$

Fix any $v_0 \in \Pi_1^F \setminus \Pi_0^F$ and pick, for fixed $\boldsymbol{x} \in A$, $u_0 \in \mathrm{ch}(v)$ for which $|\overline{x}_{v_0} - \boldsymbol{x}_{u_0}|$ is largest, so that $|\boldsymbol{x}_u - \boldsymbol{x}_{u'}| \leq 2|\boldsymbol{x}_{u_0} - \overline{x}_{v_0}|$ for all $u, u' \in \mathrm{ch}(v)$, in particular $2|\boldsymbol{x}_{u_0} - \overline{x}_{v_0}| \geq \inf_{\boldsymbol{x} \in A} \min_{u \neq v} |\boldsymbol{x}_u - \boldsymbol{x}_v| =: D$. Bound $\widetilde{p}_s(\overline{x}_v - \overline{x}_u) \leq Cs^{-d/2}$ if $(u, v) \neq (u_0, v_0)$, so we can further bound by

$$C \int_0^{\infty} s^{-b'} (1+s)^C \mathrm{e}^{-\lambda s} \mathrm{e}^{-cD^2/s} \, \mathrm{d}s \leq C \left(1 + D^{-2b'}\right), \tag{II.E.12}$$

for some $b' > b$, using Lemma II.E.8.

Now let $d = 1$. The approach above fails when the first transition $\Pi_0^F \to \Pi_1^F$ is a binary merger of two identical particles, say u and u' with parent $v = u \cup u'$, in which case $\overline{x}_v = \overline{x}_u = \overline{x}_{u'} = \boldsymbol{x}_u$, so $\widetilde{p}_{\tau_1}(\overline{x}_v - \overline{x}_u) = \widetilde{p}_{\tau_1}(\overline{x}_v - \overline{x}_{u'}) = (2\pi\tau_1)^{-d/2}$, and the final integral (of the form (II.E.12) but without the factor $\mathrm{e}^{-cD^2/s}$) diverges. We split A into a finite number of sets according to which pair of leaves are closest to each other (and potentially identical):

$$A_{uv} \coloneqq \left\{ \boldsymbol{x} \in A \colon |\boldsymbol{x}_u - \boldsymbol{x}_v| = \min_{\substack{u', v' \in \mathrm{lf}(F) \\ \text{distinct}}} |\boldsymbol{x}_{u'} - \boldsymbol{x}_{v'}| \right\},$$

for distinct $u, v \in \mathrm{lf}(F)$. There might be overlap due to ties, but it will always be true that in A_{uv} all particle locations are distinct except possibly $\boldsymbol{x}_u = \boldsymbol{x}_v$, and

$$\forall \boldsymbol{x} \in A_{uv} \colon \widetilde{\varepsilon_1}(\boldsymbol{x}) = \min_{\substack{u' \neq v' \in \mathrm{lf}(F) \\ \{u', v'\} \neq \{u, v\}}} |\boldsymbol{x}_{u'} - \boldsymbol{x}_{v'}|.$$

Since we can bound $\sup_{\boldsymbol{x}\in A} \le \sum_{u\neq v} \sup_{\boldsymbol{x}\in A_{uv}}$, it suffices to restrict our attention to $A_{u_0v_0}$ for fixed but arbitrary, distinct $u_0, v_0 \in \mathrm{lf}(F)$. With a slight abuse of notation we denote $A = A_{u_0v_0}$ from now. If the first merge event is not a pure binary merge of u_0 and v_0, that is $\Pi_1^F \setminus \Pi_0^F \neq \{u_0 \cup v_0\}$, then we can copy the proof of the $d \geq 2$ case, and in the final step we choose $v \in \Pi_1^F \setminus \Pi_0^F$ with $v \neq u_0 \cup v_0$, so there are $u, u' \in \mathrm{ch}(v)$ with $\{u, u'\} \neq \{u_0, v_0\}$ and thus $|\boldsymbol{x}_u - \boldsymbol{x}_{u'}| \geq \widetilde{\varepsilon_1}(\boldsymbol{x})$ for all $\boldsymbol{x} \in A$. Now suppose the first event is a binary merge of u_0 and v_0, that is $\Pi_1^F \setminus \Pi_0^F$ is a singleton set containing their parent $w_0 \coloneqq u_0 \cup v_0$, so

$$G_1(\tau, \boldsymbol{x}) = (\tau_1/2)^{1/2} \widetilde{p}_{\tau_1}(\boldsymbol{x}_{u_0} - \overline{x}_{w_0}) \widetilde{p}_{\tau_1}(\boldsymbol{x}_{v_0} - \overline{x}_{w_0}) \leq C\tau_1^{-1/2} \tag{II.E.13}$$

for all $\boldsymbol{x} \in A$. Then as before, we can prove inductively that for $i \in \{2, \dots, m\}$ there is $b_i \geq 0$ with

$$\int f_{\mathrm{tm}}^{\boldsymbol{\lambda}\ (i)}(\tau)(1 + \sqrt{\tau_m})^K \sup_{\boldsymbol{x}\in A} G^{(i)}(\tau, \boldsymbol{x}) \,\mathrm{d}\tau_{i+1} \dots \mathrm{d}\tau_m \leq C\tau_2^{-b_i}(1 + \tau_i)^C, \tag{II.E.14}$$

where we only need to replace the lower bound $c\tau_1$ by $c\tau_2$ in (II.E.11). Let $s \coloneqq \tau_1$ and $r \coloneqq \tau_2 - \tau_1$, then plugging (II.E.13) and (II.E.14) into (II.E.9) gives

$$\begin{aligned} &\int f_{\mathrm{tm}}^{\boldsymbol{\lambda}}(\tau) \sup_{\boldsymbol{x}\in A} \int \widetilde{f_{\mathrm{sp}}}(\xi \mid \tau, \boldsymbol{x}) \,\mathrm{d}\xi \,\mathrm{d}\tau \\ &\qquad \leq C \int_0^\infty \mathrm{d}s\, \mathrm{e}^{-\lambda_1 s} s^{-1/2} \int_0^\infty \mathrm{d}r\, (s + r)^{-b_2} (1 + s + r)^C \mathrm{e}^{-\lambda_2 r} \sup_{\boldsymbol{x}\in A} G_2(\tau, \boldsymbol{x}). \end{aligned} \tag{II.E.15}$$

If the second merge involves at least two leaves, say $w \in \Pi_2^F \setminus \Pi_1^F$ with $u, v \in \mathrm{ch}(w) \cap \mathrm{lf}(F)$, then necessarily $\{u, v\} \cap \{u_0, v_0\} = \emptyset$, so $\max_{u'\in\mathrm{ch}(w)} |\boldsymbol{x}_{u'} - \overline{x}_w| \geq \frac{1}{2}|\boldsymbol{x}_u - \boldsymbol{x}_v| \geq \widetilde{\varepsilon_1}(\boldsymbol{x})$ for all $\boldsymbol{x} \in A$ and τ, so we can bound

$$\sup_{\boldsymbol{x}\in A} G_2(\tau, \boldsymbol{x}) \leq C(1 + \tau_2)^C \tau_2^{-b_1} \mathrm{e}^{-cD^2/\tau_2}$$

for some $b_1 \geq 0$, where $D = \inf_{\boldsymbol{x}\in A} \widetilde{\varepsilon_1}(\boldsymbol{x})$. With $b \coloneqq b_1 + b_2$ we can further bound (II.E.15) by

$$\begin{aligned} &C \int_0^\infty \mathrm{d}s\, \mathrm{e}^{-\lambda_1 s} s^{-1/2} \int_0^\infty \mathrm{d}r\, \underbrace{(s + r)^{-b} \mathrm{e}^{-cD^2/(s+r)}}_{\leq CD^{-2b}\ \forall s,r>0} (1 + s + r)^C \mathrm{e}^{-\lambda_2 r} \\ &\qquad \leq CD^{-2b} \int_0^\infty \mathrm{d}s\, \mathrm{e}^{-\lambda_1 s} s^{-1/2} \int_0^\infty \mathrm{d}r\, (1 + s + r)^C \mathrm{e}^{-\lambda_2 r} \\ &\qquad \leq CD^{-2b} \int_0^\infty \mathrm{e}^{-\lambda_1 s} s^{-1/2} (1 + s)^C \,\mathrm{d}s \\ &\qquad \leq CD^{-2b}. \end{aligned} \tag{II.E.16}$$

If the second merge involves at most one additional leaf, then it must be a binary merge of w_0 and a leaf $u \notin \{u_0, v_0\}$ with parent $w_1 = u \cup w_0$. Then $|\boldsymbol{x}_u - \overline{x}_{w_1}| \geq \frac{1}{2}|\boldsymbol{x}_u - \overline{x}_{w_0}| \geq$

$c|\boldsymbol{x}_u - \boldsymbol{x}_{u_0}| \geq c\varepsilon_1(\boldsymbol{x})$ for all $\boldsymbol{x} \in A$ and τ, where we used Lemma II.E.9 in the second step. Then we can bound

$$G_2(\tau, \boldsymbol{x}) = \underbrace{r_{w_1}^{1/2}}_{\leq \tau_2^{1/2}} \underbrace{\widetilde{p}_{\tau_2 - \tau_1/2}(\overline{x}_{w_0} - \overline{x}_{w_1})}_{\leq C\tau_2^{-1/2}} \underbrace{\widetilde{p}_{\tau_2}(\boldsymbol{x}_u - \overline{x}_{w_1})}_{\leq C\tau_2^{-1/2}\mathrm{e}^{-cD^2/\tau_2}} \leq C\tau_2^{-1/2}\mathrm{e}^{-cD^2/\tau_2},$$

and finish the proof as in (II.E.16).

□

The following lemma finishes the proof of Lemma II.E.2 together with (II.E.5) and Lemma II.E.4.

Lemma II.E.7. *For every $\Pi \in \mathcal{P}$, $\boldsymbol{x} \in E_\circ^\Pi$, and $\vec{\boldsymbol{k}} \colon \Pi \to \mathbb{Z}^d$,*

$$\widetilde{\varepsilon}_d(\boldsymbol{x} + \vec{\boldsymbol{k}}) \geq \varepsilon_d(\boldsymbol{x}).$$

Proof. If $u, v \in \Pi$, then

$$|\boldsymbol{x}_u + \vec{\boldsymbol{k}}_u - (\boldsymbol{x}_v + \vec{\boldsymbol{k}}_v)| \geq \inf\{|\boldsymbol{x}_u - \boldsymbol{x}_v + \vec{\ell}| \colon \ : \vec{\ell} \in \mathbb{Z}^d\} = \rho(\boldsymbol{x}_u, \boldsymbol{x}_v).$$

With that, the claim follows directly from the definitions of ε_d and $\widetilde{\varepsilon}_d$. □

We finish by proving some technical lemmas used in the proof of Lemma II.E.4.

Lemma II.E.8. *For $a, b, \lambda, y > 0$,*

$$\int_0^\infty \mathrm{e}^{-\lambda s - y^2/s}(1+s)^a s^{-b}\,\mathrm{d}s \leq C\left(1 + y^{-2b}\right),$$

where $C = C(a, b, \lambda) > 0$.

Proof. First note that

$$\sup_{s>0}\left(s^{-b}\mathrm{e}^{-y^2/s}\right) = y^{-2b}\sup_{s>0}\left(s^{-b}\mathrm{e}^{-1/s}\right) = Cy^{-2b},$$

so

$$\int_0^\infty \mathrm{e}^{-\lambda s - y^2/s}(1+s)^a s^{-b}\,\mathrm{d}s \leq 2^a \int_0^1 s^{-b}\mathrm{e}^{-y^2/s}\,\mathrm{d}s + \int_1^\infty \mathrm{e}^{-\lambda s}(1+s)^a\,\mathrm{d}s \leq Cy^{-2b} + C.$$

□

Lemma II.E.9. *If $x_1, \ldots, x_n \in \mathbb{R}^d$ and $|x_1 - x_2| \leq |x_i - x_j|$ for all distinct $i, j \in [n]$, then*

$$|x_i - \frac{x_1 + x_2}{2}| \geq \frac{3}{\sqrt{2}}|x_i - x_1|$$

for all $i \in \{3, \ldots, n\}$.

Proof. Put $a \coloneqq |x_1 - x_2|$ and $\overline{x} \coloneqq (x_1 + x_2)/2$, and let $i \in \{3, \ldots, n\}$. Then $x_i \notin B(x_1, a) \cup B(x_2, a)$, which means that $|x_i - \overline{x}|$ is minimised if $x_i \in B(x_1, a) \cap B(x_2, a)$, in which case $|x_i - x_1| = |x_i - x_2| = |x_1 - x_2| = a$, so by Pythagoras $|x_i - \overline{x}| = \frac{3}{\sqrt{2}}a$. □

II.F Minor Technical Lemmas

This section contains some minor technical lemmas used in the proofs of the main paper.

Lemma II.F.1. *Let A be a measurable space and B a Borel space, m a probability measure on B with no atoms, and $f\colon A\times B\to(0,\infty)$ measurable with $\int_B f(a,b)m(\mathrm{d}b)=1$ for every $a\in A$, so that $m_a(\mathrm{d}b)\coloneqq f(a,b)m(\mathrm{d}b)$ defines a probability measure on B for every $a\in A$. Then there exists a measurable function $h\colon A\times B\to B$ such that*

$$\forall a\in A\colon h(a,\cdot)\#m=m_a.$$

Proof. Let $\varphi\colon B\to[0,1]$ be a Borel isomorphism,

$$f'\colon A\times[0,1]\to(0,\infty);\,(a,x)\mapsto f(a,\varphi^{-1}(x)),$$

and $m'\coloneqq\varphi\#m$. Then $\int_{[0,1]}f'(a,x)m'(\mathrm{d}x)=\int_B f(a,b)m(\mathrm{d}b)=1$ for every $a\in A$, so $m'_a(\mathrm{d}x)\coloneqq f'(a,x)m'(\mathrm{d}x)$ is a probability measure on $[0,1]$, and

$$\int_{[0,1]}g(x)m'_a(\mathrm{d}x)=\int_{[0,1]}g(x)f(a,\varphi^{-1}(x))m'(\mathrm{d}x)=\int_B g(\varphi(b))f(a,b)m(\mathrm{d}x),$$

so $m'_a=\varphi\#m_a$. Now if there exists a measurable $h'\colon A\times[0,1]\to[0,1]$ with $h'(a,\cdot)\#m'=m'_a$ for every $a\in A$, then $h\colon A\times B\to B$ defined by $(a,b)\mapsto\varphi^{-1}(h'(a,\varphi(b)))$ satisfies

$$\begin{aligned}\int_B g(b)(h(a,\cdot)\#m)(\mathrm{d}b)=\int_B g(\varphi^{-1}(h'(a,\varphi(b))))m(\mathrm{d}b)&=\int_{[0,1]}g(\varphi^{-1}(h'(a,x)))m'(\mathrm{d}x)\\&=\int_{[0,1]}g(\varphi^{-1}(x))m'_a(\mathrm{d}x)\\&=\int_B g(b)m_a(\mathrm{d}b),\end{aligned}$$

so $h(a,\cdot)\#m=m_a$. This shows that it suffices to show the statement if $B=[0,1]$.

Define $F\colon A\times[0,1]\to[0,1]$ by $F(a,y)=\int_0^y f(a,x)m(\mathrm{d}x)$, which is measurable by Tonelli's theorem. Then $F^{-1}(a,u)\coloneqq\inf\{y\in[0,1]\colon F(a,y)\geq u\}$ defines a function $A\times[0,1]\to[0,1]$ which is measurable because F is càdlàg. Furthermore, if λ denotes Lebesgue measure on $[0,1]$, then $F^{-1}(a,\cdot)\#\lambda=m_a$. Let $G\colon[0,1]\to[0,1]$ be defined by $G(y)=\int_0^y m(\mathrm{d}x)$, so that $G\#m=\lambda$ because m has no atoms. Then $h(a,x)\coloneqq F^{-1}(a,G(x))$ is a measurable map $A\times[0,1]\to[0,1]$, and

$$h(a,\cdot)\#m=(F^{-1}(a,\cdot)\circ G)\#m=F^{-1}(a,\cdot)\#(G\#m)=F^{-1}(a,\cdot)\#\lambda=m_a.$$

□

Lemma II.F.2. *Suppose $h\colon(\mathbb{Q}\cap[0,\infty))\times A\times A\to\mathbb{R}$ is a function for a set A, and that there is a family of sets $(A_x\subset A)_{x\in A}$ such that $A_x\cap A_{x'}\neq\emptyset$ for all $x,x'\in A$, and for every $s,t,\in\mathbb{Q}\cap[0,\infty)$, $x,y\in A$, $z\in A_x$,*

$$g(t+s,x,y)=g(s,x,z)+g(t,z,y).$$

Suppose further that $g(0,x,y) = b(x) - b(y)$ *for some function* $b\colon A \to \mathbb{R}$. *Then there is* $\beta \in \mathbb{R}$ *such that for all* $t \in \mathbb{Q} \cap [0,\infty)$, $x \in A$ *and* $z \in A_x$,

$$g(t,x,z) = \beta t + b(x) - b(z).$$

Proof. Fix $x \in A$. Then for $t, s \in \mathbb{Q} \cap [0,\infty)$ and $z \in A_x$ we have $g(t+s,x,z) = g(s,x,z) + g(t,z,z)$, and thus

$$g(t,z,z) + g(s,z,z) = \Big[g(t,x,z) - g(0,x,z)\Big] + \Big[g(t+s,x,z) - g(t,x,z)\Big] = g(t+s,z,z),$$

for any $t, s, \in \mathbb{Q} \cap [0,\infty)$ and $z \in A_x$. This implies that there exists $\beta(z) \in \mathbb{R}$ such that $g(t,z,z) = \beta(z)t$ for $t \in \mathbb{Q} \cap [0,\infty)$. Then for all $z \in A_x$ and $t \in \mathbb{Q} \cap [0,\infty)$,

$$g(t,x,z) = g(0,x,z) + g(t,z,z) = b(x) - b(z) + \beta(z)t. \tag{II.F.1}$$

We show that $\beta(z) = \beta(z')$ for all $z, z' \in \bigcup_{x\in A} A_x$. By assumption and (II.F.1), for every $z \in A_x$ we have $g(t,x,x) = g(t,x,z) + g(0,z,x) = \beta(z)t$, in particular $\beta(z) = \beta(z') =: \beta(x)$ for all $z, z' \in A_x$.

Finally, if $x, x' \in A$ then there is $z \in A_x \cap A_{x'}$, so $\beta(x) = \beta(z) = \beta(x')$, so in fact there is one $\beta \in \mathbb{R}$ such that $\beta(z) = \beta$ for all $z \in \bigcup_{x\in A} A_x$. □

Chapter III

Repulsion-Diffusion Equations

1 Introduction

1.1 Motivation—The Pain in the Torus

The phenomenon that has been infamously dubbed the "pain in the torus" by Felsenstein in 1985 [23] can be summarised as follows: *in one and two dimensions, spatial population models with independent critical branching and diffusive spatial motion concentrate in increasingly few well-separated clumps, until they eventually die out.* The latter fact is simply due to the criticality of the branching, and can be overcome by conditioning on survival or by starting with an infinite population. The "clumping" issue is more inherently connected to the interplay of the branching mechanism and the diffusive motion in low dimensions. We look at this problem in more detail at the example of a (critical) *superBrownian motion*, though the analysis, on a high level, follows the same line of reasoning for many related population models. A superBrownian motion (see e.g. [20, 35]) is a measure-valued process on $\mathbb{R}^d$, $d \in \mathbb{N}$, that can be obtained as the scaling limit of an individual based model in which an initial number of particles of order N[1] move according to independent Brownian motions, and branch at rate N into a random number of offspring with mean $1+o(1/N)$ and variance $\gamma/N+o(1/N)$, where $\gamma \geq 0$ is the *branching variance* and N is the scaling parameter which tends to infinity. Formally, a superBrownian motion (SBM) is the solution to the stochastic partial differential equation (SPDE)

$$\mathrm{d}X_t = \frac{1}{2}\Delta X_t \,\mathrm{d}t + \sqrt{\gamma X_t}\,\mathrm{d}\mathcal{W}_t, \tag{1.1}$$

for a space-time white noise $\mathcal{W}$. If $d = 1$, then this SPDE is well-posed, but otherwise X_t is singular w.r.t. Lebesgue measure, and (1.1) is ill-posed and has to be replaced with a martingale problem. Note if $\gamma = 0$ then the SBM is just the solution to the heat equation. Alternatively, we can recover the heat equation by taking a scaling limit of

[1] or, for infinite populations, an initial spatial distribution of particles whose density is locally of order N.

the same population model except that the branching rate and variance of the offspring distribution remain constant as $N \to \infty$.

A heuristic for the clumping problem

We will now heuristically explain why and how the "clumping problem" arises in a SBM started from infinite mass in dimensions $d \leq 2$. In short, for large times t, an increasingly small number of individuals that were alive at time zero will be ancestral to the entire population at time t, and if $d \leq 2$ then the diffusion is not "fast enough" to disperse and spread these large families, and they form well-separated clumps.

Consider a SBM started from Lebesgue measure. We cut $\mathbb{R}^d$ into a grid of unit sized cubes, and regard the initial mass in each of them as one family. Due to the independent branching, we can let each of the families evolve independently from each other, and obtain the process started from Lebesgue measure as their superposition (this is called the branching property, see e.g. [20, p. 2]). Each family is a critical branching process started from finite mass, so the probability that it is still alive at time n is proportional to $1/n$, and, if alive, its expected size is proportional to n [35, Thm. II.1.1]. Hence, in expectation, after n units of time all but every n'th family has gone extinct, and each of the living families consists of order n individuals. Due to their diffusive movement, each family will have spread over an area of radius $\sim \sqrt{n}$, hence the population density of any surviving family is $\sim n^{1-d/2}$. In the critical case $d = 2$, this crude heuristic misses a factor $\log n$ (c.f. (1.2) below), so the density of the surviving families diverges as $n \to \infty$ if and only if $d \leq 2$, in which case they form separated clumps.

This first attempt at a heuristic calculation reveals that the dichotomy between stable long-term dynamics and clumping can really be understood as the dichotomy between unbounded and bounded asymptotic population density of a single surviving family, that is, a finite mass SBM conditioned on survival. It is a classical result due to Evans [22] that the distribution of this process is that of a single "immortal particle" that follows the path of a Brownian motion and throws off mass at a constant rate, which then evolves like an ordinary SBM, independent of the immortal particle. This is not unexpected: In the unconditioned process, at large times the entire population will have descended from increasingly few ancestors that were alive at time zero, until eventually none remain and the process goes extinct; the conditioning imposes that one of those ancestors—the immortal particle—will never perish. If (Z_t) is a Brownian motion that denotes the path of the immortal particle, then the SBM (X_t) conditioned on survival, which we may now as well start from zero, formally satisfies the SPDE

$$\mathrm{d}X_t = \left(\frac{1}{2}\Delta X_t + \gamma\delta_{Z_t}\right)\mathrm{d}t + \sqrt{\gamma X_t}\,\mathrm{d}\mathcal{W}_t, \quad X_0 = 0,$$

where δ_z is the Dirac delta at $z \in \mathbb{R}^d$. Then the mean measure conditional on (Z_t) has a density ρ_t that solves $\partial_t \rho = \frac{1}{2}\Delta\rho + \gamma\delta_{Z_t}$, so $\rho_t = \gamma \int_0^t G_{t-s}(\cdot - Z_s)\,\mathrm{d}s$, where $(G_s \colon \mathbb{R}^d \to \mathbb{R})_{s>0}$ is the heat kernel. If we centre the process around the immortal particle by putting

$$\widetilde{\rho}_t(\cdot) := \rho_t(\cdot + Z_t) = \gamma \int_0^t G_{t-s}(\cdot + Z_t - Z_s)\,\mathrm{d}s,$$

which has expectation $\gamma \int_0^t G_{2s}(\cdot)\,\mathrm{d}s$, then we can find

$$\mathbb{E}\left[X_t(B(Z_t,1))\right] = \int_{B(0,1)} \mathbb{E}\left[\widetilde{\rho}_t(x)\right]\mathrm{d}x = \gamma \int_0^t \int_{B(0,1)} G_{2(t-s)}(x)\,\mathrm{d}x\,\mathrm{d}s \propto \begin{cases} \sqrt{t}, & d = 1, \\ \log t, & d = 2, \\ 1, & d \geq 3. \end{cases} \tag{1.2}$$

Therefore, at least in expectation, mass accumulates in the vicinity of the immortal particle for large times if $d \leq 2$, and remains bounded if $d \geq 3$. This recovers the picture painted in the beginning of the section: In dimensions $d \leq 2$, a SBM started from infinite mass concentrates in increasingly few large clumps (centred around the ancestors of the surviving families), with space in between growing increasingly empty.

Introducing repulsion

As a model for a spatially evolving population (for which $d = 2$ is the most natural dimension), this is very unrealistic, and a better model should reflect stable long-term population dynamics. One of the most obvious reasons this clumping phenomenon does not occur in nature is that real individuals do not behave independently from surrounding individuals, as is assumed in the model underlying SBM. Indeed, a high population density leads to resource scarcity and decreases the average number of offspring, and causes migration away from the overcrowded area. The former effect has already been successfully integrated into the SBM model and been shown to lead to stable long-term dynamics [21]. The latter, however, has not yet been studied in this context. A natural way to implement this is to introduce a pairwise repulsion between individuals: in the individual based prelimiting model of SBM, in addition to the Brownian movement, let a particle at x experience a drift $-\nabla W(x - y)$ if there is another particle at y. Here $W \colon \mathbb{R}^d \to \mathbb{R}$ is the *interaction potential*, which must be decreasing in order for the interaction to be repulsive. The limiting process in this case will be a measure-valued process (X_t) which is formally described by the SPDE

$$\mathrm{d}X_t = \left(\frac{1}{2}\Delta X_t + \nabla \cdot (X_t \nabla W \star X_t)\right)\mathrm{d}t + \sqrt{\gamma X_t}\,\mathrm{d}\mathcal{W}_t,$$

where $\star$ denotes convolution. If we again consider a single surviving family, that is a finite mass superprocess conditioned on survival, then this will be described by an immortal

particle that constantly immigrates mass into the system, and that is repulsed from the mass it throws off. Formally, we arrive at

$$\begin{aligned} \mathrm{d}X_t &= \left(\frac{1}{2}\Delta X_t + \nabla \cdot (X_t \nabla W \star X_t) + \gamma \delta_{Z_t}\right) \mathrm{d}t + \sqrt{\gamma X_t}\, \mathrm{d}\mathcal{W}_t, \\ \mathrm{d}Z_t &= \mathrm{d}B_t - \nabla W \star X_t(Z_t)\, \mathrm{d}t, \end{aligned}$$

for independent space-time white noise $\mathcal{W}$ and Brownian motion B. This is a very complicated process, and a natural first step is to study it without the noise; if the equation were linear, this would be the same as taking expectations conditional on (Z_t). If we also replace the Dirac immigration with a bounded function centred on Z_t—which should not change the behaviour of the system with regards to clumping, but makes the equation more regular—then, conditional on Z_t, we arrive at the partial differential equation

$$\partial_t \rho = \frac{1}{2}\Delta \rho + \nabla \cdot (\rho \nabla W \star \rho) + f_t, \tag{1.3}$$

where $f_t = f(\cdot - Z_t)$ for some $f \geq 0$ with $\int f(x)\, \mathrm{d}x = 1$. Given our motivation, the question we want to answer is under what assumptions on the repulsion does the solution of (1.3) exhibit bounded long-term behaviour in dimensions one and two, and this chapter is devoted to attacking this question in two different ways. We will from now assume that the *immigration* term f_t in (1.3) does not depend on time (i.e. we "freeze" the immortal particle). This simplifies the analysis and, at least in the case where there is no repulsion, is easily checked not to change the clumping-stability dichotomy. That said, about half of our results apply with almost no modification to the case of time-dependent immigration, see Remark 2.6 below.

1.2 Summary of results

As motivated in the previous section, we study the following partial differential equation,

$$\partial_t \rho = \frac{1}{2}\Delta \rho + \nabla \cdot (\rho \nabla W \star \rho) + f, \tag{1.4}$$

on $\mathbb{R}^d$ for any dimension $d \in \mathbb{N}$, started from a non-negative, bounded and integrable initial condition, where

(i) f is the *immigration*, which is bounded, integrable, and non-negative,

(ii) $W \in C^2(\mathbb{R}^d \setminus \{0\})$ is the *interaction potential*, which we assume to be eventually decreasing.

Our first class of results, which form Section 2, gives sufficient conditions on the repulsion under which ρ is guaranteed to remain globally bounded for all times for any bounded

initial condition. More precisely, we find sharp conditions on W under which the following global boundedness property holds:

$$\forall \rho_0 \colon \|\rho_0\|_\infty \leq M \implies \sup_{t\geq 0} \|\rho_t\|_\infty \leq M, \tag{1.5}$$

for an explicit value $M > 0$. Under the same conditions on W, and for the same value M, a form of the maximum principle holds:

$$\forall \rho_0 \colon \|\rho_0\|_\infty \geq M \implies \max_{t\geq 0} \|\rho_t\|_\infty = \|\rho_0\|_\infty . \tag{1.6}$$

Together, (1.5) and (1.6) give sufficient conditions on W under which

$$\sup_{t\geq 0} \|\rho_t\|_\infty < \infty \tag{1.7}$$

for any bounded initial condition. It will turn out that the critical strength of the repulsion for (1.5) and (1.6) to hold is $|\nabla W(x)| \sim |x|^{1-d}$, and W must be singular at the origin.

Note that, even though these results are sharp, they give much stronger assertions than we need; what we really care about is that (1.7) holds for any bounded initial condition, which we would expect to be true under weaker assumptions. Indeed, in $d \geq 3$ it holds without any repulsion at all. Our second class of results, which form Section 3, take a different angle on the problem. In general, we would expect (1.7) to correlate with the existence of (non-negative) stationary states of (1.4). For example, without repulsion, (1.4) reduces to the heat equation with immigration $\partial_t \rho = \frac{1}{2}\Delta\rho + f$, whose solution started from zero is given by

$$\rho_t = \int_0^t G_s \star f \, \mathrm{d}s,$$

where $G_s(x) = (2\pi s)^{-d/2} \mathrm{e}^{-x^2/(2s)}$, $s > 0$, is the heat kernel. Provided that $f \not\equiv 0$:

(i) If $d \leq 2$, then there are no stationary distributions and $\rho_t \uparrow \infty$ locally uniformly.

(ii) If $d \geq 3$, then there is a unique non-negative stationary distribution ρ, and $\rho_t \uparrow \rho$ locally uniformly.

This is proved in Lemma 2.19. Our second class of results gives necessary conditions on the strength of the repulsion for stationary states of (1.4) to exist, and identify their tail behaviour if they do. The precise results are in Section 3 below, particularly Theorems 3.4 to 3.6, and we present here the conjecture to which they give rise. See Fig. 1.1 for an illustration.

Conjecture 1.1. *Suppose that* $|\nabla W(x)| \sim |x|^{-\alpha}$ *for some* $\alpha \geq 0$.

(i) $d = 1$.

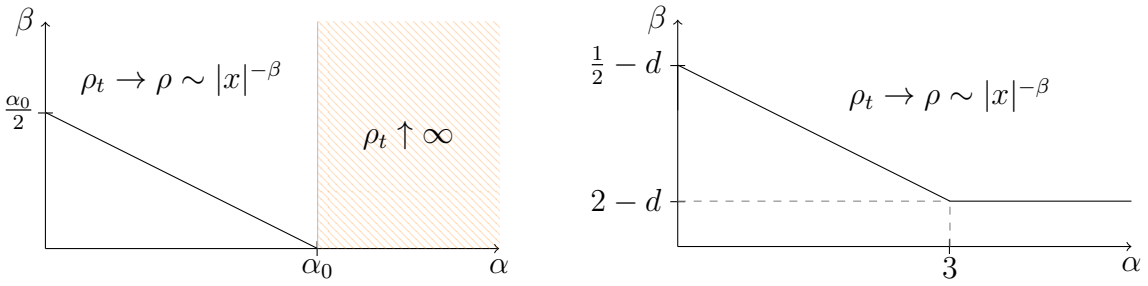

Figure 1.1: Illustration for Conjecture 1.1 for $d = 1, 2$ (left) and $d \geq 3$ (right). Here, $\alpha_0 = 1$ if $d = 1$ and $\alpha_0 = 3$ if $d = 2$.

(i.1) If $\alpha \in [0, 1)$, then ρ_t converges to a stationary distribution with $\rho(x) \sim |x|^{-\frac{1-\alpha}{2}}$.

(i.2) If $\alpha > 1$, then $\rho_t \uparrow \infty$ locally uniformly.

(ii) $d = 2$.

(ii.1) If $\alpha \in [0, 3)$, then ρ_t converges to a stationary distribution with $\rho(x) \sim |x|^{-\frac{3-\alpha}{2}}$.

(ii.2) If $\alpha > 3$, then $\rho_t \uparrow \infty$ locally uniformly.

(iii) If $d \geq 3$, then ρ_t converges to a stationary distribution with

$$\rho(x) \sim \begin{cases} |x|^{\frac{1+\alpha}{2}-d}, & \alpha \in [0, 3), \\ |x|^{2-d}, & \alpha \geq 3. \end{cases}$$

1.3 Related work

SuperBrownian motion with immigration

A lot of work has been done on superBrownian motion with immigration (and no interaction). Amongst the known results for this process are central limit theorems [26] and large deviation principles for large [39, 40] and small times [41]. Further results on this and more general measure-valued diffusions with immigration can be found in [31] and references therein.

Branching Brownian motion with interaction

Some work has also been done on branching Brownian motion (BBM) with interaction. Classical BBM, without interaction, is a system of particles that move as independent Brownian motions, and branch at constant rates into exactly two offspring. The number of particles grows exponentially and there is no chance of extinction, which makes it quite different from the BPS with critical branching considered here. Questions of interest in this setting are often of extremal type, such as the structure of the process close to the furthest particle from the origin [1, 2], or large time limits of the population's

empirical measure if scaled appropriately [18, 19]. Some authors have studied BBM with repulsive or attractive interactions. Engländer considered the case where particles have an Ornstein-Uhlenbeck-type attraction or repulsion (that is, $W(x) = bx^2$ for $b \in \mathbb{R}$) to or from their common centre of mass [18], and from each other [19]. Similar results have been obtained in the context of supercritical superBrownian motion [24]. This particular interaction often allows for explicit calculations because $\nabla W(x) \propto x$ is linear. Note also that in light of our motivation, it is not the most natural choice of repulsion, since its strength grows, rather than decays with distance. In another recent paper [6], authors study a BBM in which they introduce short-range pairwise repulsion through a change of measure that penalises the total time that particles spend within close range of each other. They show that the dominant effect of the penalisation is a drastic reduction in branching rate, and that this model is well-approximated by a simplified model in which only branching events are penalised, and there is no repulsion between individuals once they are born.

Aggregation-diffusion equations

Equation (1.4) is also related to a well-studied class of non-local, nonlinear partial differential equations known as *aggregation-diffusion equations*,

$$\partial_t \rho = \frac{1}{2}\Delta\rho + \nabla \cdot (\rho \nabla W \star \rho), \tag{1.8}$$

which differ from (1.4) in that there is no immigration, and the interaction potential is attractive at long range rather than repulsive. We refer to [10, 27] for reviews of this class of equations. Variants of (1.8) also exist with non-linear diffusion [33, 38]; when the diffusion is linear as above, (1.8) is commonly called a *McKean Vlasov equation*. Aggregation-diffusion equations have attracted significant interest in the literature because they describe the large scale dynamics of a wide variety of interacting particle systems arising in biology, physics, social and other life sciences, which are often driven by long-range attraction and short range repulsion. Examples include chemotaxis, bacteria orientation, or motion of human crowds, see [29, 34, 10]. Significant mathematical interest is further due to the delicate competition between aggregation and diffusion, which leads to a dichotomy between well-posedness and finite time blowup [5, 3, 4, 9, 36, 42, 25, 14, 15]. For sufficiently weak interaction, the diffusion dominates and the solution asymptotically simplifies to the solution of the heat equation [11], while a balance between diffusion and aggregation can lead to the existence of non-trivial steady states [8, 13, 12, 28, 32].

Even though (1.4) looks similar to (1.8), its behaviour is markedly different. Where the competition between aggregation and diffusion decides the behaviour of (1.8), that of (1.4) is decided by the competition between the immigration against the diffusion *and* the repulsion, which both work to spread the immigrated mass.

2 Dynamical Properties

In the following section we will give precise statements of the first class of our results, and Section 2.2 contains the proofs. The appendix III.A contains a brief definition of fractional Sobolev spaces and some their properties that are useful in the proofs.

2.1 Results

We require the following regularity assumptions. By $\mathcal{W}^{\gamma,p}$ we denote the usual Sobolev spaces on $\mathbb{R}^d$, where $p \in [1,\infty]$, and $\gamma \geq 0$ may not be an integer. If $\gamma = 0$ then $\mathcal{W}^{\gamma,p} = L^p$ is the usual L^p space. See Section III.A for a brief and [17, 37, 7] for a comprehensive introduction to fractional Sobolev spaces. We further write $g_+ = g \vee 0$ and $g_- = (-g) \vee 0$ for the positive and negative part of a function $g\colon \mathbb{R}^d \to \mathbb{R}$, respectively.

Assumption (A). (i) *There is $\gamma_f \in (0,1)$ such that $f \in \mathcal{W}^{\gamma_f,1} \cap \mathcal{W}^{\gamma_f,\infty}$.*

(ii) ∇W and ΔW are bounded on $\mathbb{R}^d \setminus B(0,r)$ for all $r > 0$, and locally integrable.

(iii) At least one of $(\Delta W)_-$ and $(\Delta W)_+$ is integrable.

Assumption (A)(i) is required to ensure that the mild solution is also a classical solution. A simple sufficient criterion is that f is compactly supported and Hölder continuous (with any positive exponent). Assumption (A)(iii) ensures that $\int \Delta W := \int_{\mathbb{R}^d} \Delta W(x)\,\mathrm{d}x \in [-\infty,\infty]$ exists. As we will see shortly, the fact that W is eventually decreasing implies $\int \Delta W < \infty$ and hence that in fact $(\Delta W)_+$ must be integrable. Here and later, the integral is in a strict, classical sense; for example the Newtonian potential has $\Delta W_N \equiv 0$ outside the null set $\{0\}$, so $\int \Delta W_N = 0$. The following well-posedness result only requires (A)(i), and (A)(ii) for ∇W, but no assumptions on ΔW.

Theorem 2.1 (Well-posedness)**.** *For every initial condition there exists a unique mild solution to* (1.4) *up to a maximal existence time $T^\star \in (0,\infty]$. It is non-negative and solves* (1.4) *in the classical sense on $(0,T^\star)$.*

In this statement and the remainder of Section 2.1, we only consider initial conditions that are non-negative, bounded, and integrable, without mentioning this explicitly. Precise statements are in Theorems 2.9 and 2.10 below. We note that it is not possible to prove $T^\star = \infty$ under the current assumptions, because the short-range behaviour of W could lead to finite-time blowup.

The main result of this work is the identification of an index $c_W \in \mathbb{R}$ that plays a critical role in the behaviour of solutions to (1.4). We first state our main results in terms of c_W, and discuss its definition and properties afterwards.

Proposition 2.2. *If $c_W > 0$, then $T^\star = \infty$ for every initial condition.*

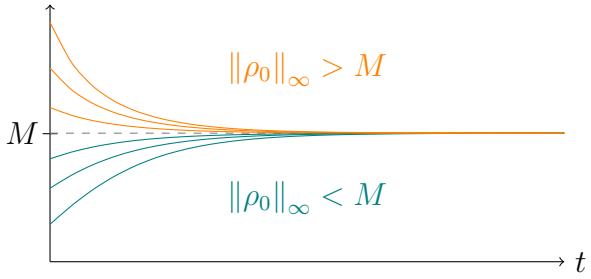

Figure 2.1: Evolution of the upper bound on $\|\rho_t\|_\infty$ in Corollary 2.16 for $c_W > 0$ and different values of $\|\rho_0\|_\infty$.

Theorem 2.3 (Maximum Principle)**.** *If $c_W > 0$ and $M = \sqrt{\|f\|_\infty / c_W}$, then*

$$\|\rho_0\|_\infty \geq M \implies \max_{0 \leq t < T^\star} \|\rho_t\|_\infty = \|\rho_0\|_\infty \tag{2.1}$$

for every initial condition ρ_0. If $M < \sqrt{\|f\|_\infty / c_W}$, or $c_W < 0$ and $M > 0$ arbitrary, then there exists an initial condition ρ_0 for which (2.1) is false.

Theorem 2.4 (Global Boundedness)**.** *If $c_W > 0$ and $M = \sqrt{\|f\|_\infty / c_W}$, then*

$$\|\rho_0\|_\infty \leq M \implies \sup_{0 \leq t < T^\star} \|\rho_t\|_\infty \leq M \tag{2.2}$$

for every initial condition ρ_0. If $M < \sqrt{\|f\|_\infty / c_W}$, or $c_W < 0$ and $M > 0$ arbitrary, then there exists an initial condition ρ_0 such that (2.2) is false.

Both theorems are consequences of the following result. For a function $g \colon [0, T) \to [0, \infty)$ and $t \in [0, T)$, write $\partial_t^+ g(t) \coloneqq \overline{\lim}_{h \downarrow 0} \frac{g(t+h) - g(t)}{h}$.

Theorem 2.5. *For every initial condition and $t \in [0, T^\star)$,*

$$\partial_t^+ \|\rho_t\|_\infty \leq \|f\|_\infty - c_W \|\rho_t\|_\infty^2 . \tag{2.3}$$

For every $c > c_W$, there exists an initial condition such that (2.3) is false at time zero if c_W is replaced by c.

The differential inequality (2.3) can be turned into an explicit upper bound, which implies the positive assertions of Theorems 2.3 and 2.4. See Corollary 2.16 for a precise statement, and Fig. 2.1 for an illustration. We remark that counterexamples for ρ_0 in the negative statements of all three theorems can be chosen to be infinitely differentiable and compactly supported.

Remark 2.6. If the immigration f is time-dependent with sufficient regularity, then existence and uniqueness for solutions to (1.4) still hold, and Theorem 2.5 remains true with (2.3) replaced by

$$\partial_t^+ \|\rho_t\|_\infty \leq \|f_t\|_\infty - c_W \|\rho_t\|_\infty^2 , \quad t \in [0, T^\star).$$

In particular, $T^\star = \infty$ and $\sup_{t \geq 0} \|\rho_t\|_\infty < \infty$ for any initial condition as long as $c_W > 0$ and $\sup_{t \geq 0} \|f_t\|_\infty < \infty$.

We now define the index c_W. Denote by

$$\langle \nabla W \rangle (R) \coloneqq \frac{1}{|\partial B(0,R)|} \int\limits_{\partial B(0,R)} \nabla W \cdot \mathrm{d}\widehat{n} \tag{2.4}$$

the average radial part of ∇W at a distance $R > 0$, where $|\partial B(0,R)| = c_d R^{d-1}$ is the surface area of $B(0,R)$. An important reference will be the *Newtonian potential*,

$$W_N(x) = c_d^{-1} \begin{cases} \frac{1}{d-2}|x|^{2-d}, & d \neq 2, \\ -\log|x|, & d = 2, \end{cases} \tag{2.5}$$

where c_d is the surface area of the unit ball in $\mathbb{R}^d$. The Newtonian potential is the Green's function of the Laplace equation (that is, $\Delta W_N = -\delta_0$ in a distributional sense), and has the physical interpretation of the electrodynamical repulsive potential in $\mathbb{R}^d$.

Lemma 2.7. *The following limits exist,*

$$\eta_W \coloneqq \lim_{r \to 0} \frac{\langle \nabla W \rangle (r)}{\langle \nabla W_N \rangle (r)}, \qquad \alpha_W \coloneqq \lim_{R \to \infty} \frac{\langle \nabla W \rangle (R)}{\langle \nabla W_N \rangle (R)}, \tag{2.6}$$

with $\eta_W \in \mathbb{R}$ and $\alpha_W \in [0, \infty]$, and $\alpha_W = \eta_W - \int(\Delta W)$.

The quantities η_W and α_W compare the strength of the repulsion to that of the Newtonian potential at short and long range, respectively. As we will see in Lemma 2.12 below, η_W also determines the singular behaviour of W at the origin, in that ΔW contains, in a distributional sense, a multiple $-\eta_W$ of the Dirac delta at zero. Note that Lemma 2.7 implies $\int(\Delta W)_+ < \infty$, because $\eta_W \in \mathbb{R}$ and $\alpha_W \geq 0$. Thus,

$$c_W \coloneqq \eta_W - \int (\Delta W)_+ \tag{2.7}$$

is well-defined. If $\alpha_W < \infty$ then $c_W = \alpha_W - \int(\Delta W)_-$ by Lemma 2.7, so in fact

$$c_W = \frac{1}{2}\left(\alpha_W + \eta_W - \int |\Delta W|\right), \quad \text{if } \alpha_W < \infty. \tag{2.8}$$

Our main results apply positively only to potentials with $c_W > 0$. From the different representations of c_W it follows that $\alpha_W > 0$ and $\eta_W > 0$ are both necessary conditions. That is, the critical strength of the repulsion is $|\nabla W(x)| \sim |x|^{1-d}$, and the repulsion potential has to have a (Newtonian) singularity at the origin. The latter is not entirely surprising, since we would not generally expect to be able to find strong bounds on the supremum norm (like (2.1) to (2.3)) of the solution to a PDE with smooth non-local interaction. Before we give examples for W, we note that (2.7) and linearity of η_W in W imply that

$$c_{W+W'} \geq c_W + c_{W'}, \qquad c_{aW} = a c_W, \tag{2.9}$$

for any $a > 0$ and potentials W, W' satisfying Assumption (A). This implies that the class of potentials with $c_W > 0$ is closed under positive linear combinations, as well as small perturbations: If $c_W > 0$ and $\widetilde{W}$ is a perturbative potential, then $c_{W+\varepsilon\widetilde{W}} > 0$ for small $\varepsilon > 0$, in fact as long as $\varepsilon < c_W |c_{\widetilde{W}}|^{-1}$. An example for a perturbation could be a smooth potential that decays faster than Newtonian (so that $\eta_W = \alpha_W = 0$), in which case $c_{\widetilde{W}} = -\frac{1}{2}\int |\Delta \widetilde{W}|$ by (2.8).

Examples for interaction potentials that satisfy Assumption (A) include repulsive power laws

$$W(x) = -P_A(x) \coloneqq -\begin{cases} \frac{|x|^A}{A}, & A \neq 0, \\ \log|x|, & A = 0, \end{cases}$$

with $1 \geq A \geq 2 - d$ (for Theorem 2.1 to hold it is sufficient if $A > 1 - d$), and Morse potentials [16],

$$W(x) = -C_A \mathrm{e}^{-|x|/\ell_A} + C_R \mathrm{e}^{-|x|/\ell_R},$$

with $C_A, C_R, \ell_A, \ell_B > 0$, which are repulsive at long range if $\ell_R < \ell_A$. Of these the most natural example with $c_W > 0$ is the Newtonian potential with $\alpha_W = \eta_W = c_W = 1$, which in this notation is $W_N = -c_d^{-1} P_{2-d}$. Another family of examples is a mixture of repulsive power laws

$$W(x) = W_N(x) - P_A(x), \qquad 1 \geq A > 2 - d,$$

which have $|\nabla W(x)| \sim |x|^{1-d}$ at short and $|\nabla W(x)| \sim |x|^{A-1}$ at long range, and $\alpha_W = \infty$, $\eta_W = c_W = 1$. Morse potentials have $\eta_W = \alpha_W = 0$ and $c_W < 0$, but could be used as a perturbation, see the earlier discussion following (2.9).

2.2 Proofs

We write C for an unimportant positive constant whose value may change from one appearance to the next. For $\gamma \geq 0$,

$$\mathcal{X}^\gamma \coloneqq \mathcal{W}^{\gamma,1} \cap \mathcal{W}^{\gamma,\infty}, \quad \|\!|\!|\cdot|\!|\!|_\gamma \coloneqq \|\cdot\|_{\mathcal{W}^{\gamma,1}} + \|\cdot\|_{\mathcal{W}^{\gamma,\infty}}, \tag{2.10}$$

as well as $\mathcal{X} \coloneqq \mathcal{X}^0 = L^1 \cap L^\infty$ and $\|\!|\!|\cdot|\!|\!| \coloneqq \|\!|\!|\cdot|\!|\!|_0 = \|\cdot\|_{L^1} + \|\cdot\|_\infty$, and $\mathcal{X}_+ \coloneqq \{f \in \mathcal{X} \colon f \geq 0\}$. If $g \colon \mathbb{R}^d \to \mathbb{R}^m$ for some $m \in \mathbb{N}$, then $\|g\|_{L^p} \coloneqq \sum_{i=1}^m \|g_i\|_{L^p}$, similarly for other norms. Recall (2.10), and note that by Lemma III.A.2, $\mathcal{X}^\gamma$ with this norm is a Banach space which embeds continuously into $C^{\lfloor\gamma\rfloor-1,1}(\mathbb{R}^d)$ if $\gamma \geq 1$, and into $C^{\lfloor\gamma\rfloor,\gamma-\lfloor\gamma\rfloor}(\mathbb{R}^d)$ if $\gamma \in (0,\infty) \setminus \mathbb{N}$. Here, $C^{k,\beta}(\mathbb{R}^d)$ for $k \in \mathbb{N}_0$ and $\beta \in (0,1]$ is the space of functions $g \colon \mathbb{R}^d \to \mathbb{R}$ which are k times continuously differentiable and for which

$$\|g\|_{C^{k,\beta}} \coloneqq \sum_{|\alpha|<k} \|\partial^\alpha g\|_{L^\infty} + \sum_{|\alpha|=k} \sup_{x \neq y} \frac{|\partial^\alpha g(x) - \partial^\alpha g(y)|}{|x-y|^\beta} < \infty, \tag{2.11}$$

with the usual notational conventions for multi-indices α.

2.2.1 Existence and regularity of solutions

This section contains proofs of local in time well-posedness and regularity of solutions to (1.4). They only require Assumption (A)(i), and (A)(ii) for ∇W, but no assumptions on ΔW. Note that (A)(i) becomes $f \in \mathcal{X}^{\gamma_f}$ in the notation (2.10). Many of the ideas in this section were inspired by arguments in Section 2 of [11].

We begin by establishing well-posedness of a weak version of (1.4). More precisely, by Duhamel's formula we can formally rewrite (1.4) with initial condition $\rho_0 \in \mathcal{X}$ as an integral equation

$$\rho_t = G_t \star \rho_0 + \int_0^t G_s \star f \, \mathrm{d}s + \int_0^t \nabla G_{t-s} \star (\rho_s \nabla W \star \rho_s) \, \mathrm{d}s, \tag{2.12}$$

where $(G_s \colon \mathbb{R}^d \to (0,\infty))_{s>0}$ denotes the heat kernel, and $F \star H \coloneqq \sum_{i=1}^d F_i \star H_i$ for vector fields F and H.

Lemma 2.8. *If $\gamma \geq 0$ then there exists $C > 0$ such that for any $g \in \mathcal{X}^\gamma$,*

$$\|\nabla W \star g\|_{\mathcal{W}^{\gamma,\infty}} \leq C |\!|\!| g |\!|\!|_\gamma.$$

In particular, $|\!|\!| g \nabla W \star g |\!|\!| \leq C |\!|\!| g |\!|\!|^2$

Proof. Denote $F_1 \coloneqq \nabla W \mathbb{1}_{B(0,1)} \in L^1$ and $F_2 \coloneqq \nabla W \mathbb{1}_{\mathbb{R}^d \setminus B(0,1)} \in L^\infty$. Then, using the fractional version of Young's convolution inequality (see Theorem A.1 in [11]),

$$\begin{aligned}
\|\nabla W \star g\|_{\mathcal{W}^{\gamma,\infty}} &\leq \|F_1 \star g\|_{\mathcal{W}^{\gamma,\infty}} + \|F_2 \star g\|_{\mathcal{W}^{\gamma,\infty}} \\
&\leq C(\|F_1\|_{L^1} \|g\|_{\mathcal{W}^{\gamma,\infty}} + \|F_2\|_{L^\infty} \|g\|_{\mathcal{W}^{\gamma,1}}) \\
&\leq C(\|F_1\|_{L^1} + \|F_2\|_{L^\infty}) |\!|\!| g |\!|\!|_\gamma.
\end{aligned}$$

For the additional claim, if $p \in \{1, \infty\}$, then

$$\|g \nabla W \star g\|_{L^p} \leq \|g\|_{L^p} \|\nabla W \star g\|_{L^\infty} \leq C |\!|\!| g |\!|\!|^2.$$

□

We now prove well-posedness of (2.12) locally in time. Standard facts that we will use repeatedly are that, for $\gamma \geq 0$ and $s > 0$,

$$\|G_s\|_{\mathcal{W}^{\gamma,1}} \leq C s^{-\gamma/2}, \qquad \|\nabla G_s\|_{\mathcal{W}^{\gamma,1}} \leq C s^{-(1+\gamma)/2}. \tag{2.13}$$

(See, for example, [11, p.6].)

Theorem 2.9 (Local in time well-posedness)**.** *Given $\rho_0 \in \mathcal{X}$, there is a $T^\star \in (0,\infty]$ and a $\rho \in C([0,T^\star), \mathcal{X})$ which solves (2.12) started at ρ_0 and such that any solution $\widetilde{\rho} \in C([0,T), \mathcal{X})$ of (2.12) starting at ρ_0 satisfies $T \leq T^\star$ and coincides with ρ on $[0,T)$. Furthermore, if $T^\star < \infty$ then $\|\rho_t\|_\infty \to \infty$ as $t \uparrow T^\star$. If ρ_0, f, and W are radially symmetric, then so is ρ_t for all $t \in [0,T^\star)$.*

Proof. For fixed $T > 0$ define $\mathcal{F}\colon C([0,T],\mathcal{X}) \to C([0,T],\mathcal{X})$ by the right-hand side (RHS) of (2.12), that is,

$$\mathcal{F}[\rho]_t = G_t \star \rho_0 + \int_0^t G_s \star f \,\mathrm{d}s + \int_0^t \nabla G_{t-s} \star (\rho_s \nabla W \star \rho_s)\,\mathrm{d}s, \quad t \in [0,T].$$

We write $|||\rho||| \coloneqq \sup_{0\le s\le T} |||\rho_s|||$ for $\rho \in C([0,T],\mathcal{X})$. For $p \in \{1,\infty\}$, by Young's convolutional inequality and Lemma 2.8,

$$\begin{aligned}
\|\mathcal{F}[\rho]_t\|_{L^p} &\le \|G_t\|_{L^1}\,\|\rho_0\|_{L^p} + \int_0^t \|G_s\|_{L^1}\,\|f\|_{L^p}\,\mathrm{d}s \\
&\qquad + \int_0^t \|\nabla G_{t-s}\|_{L^1}\,\|\rho_s\|_{L^p}\,\|\nabla W \star \rho_s\|_{L^\infty}\,\mathrm{d}s \\
&\le \|\rho_0\|_{L^p} + C\left(t + \int_0^t (t-s)^{-1/2}\,\|\rho_s\|_{L^p}\,|||\rho_s|||\,\mathrm{d}s\right) \\
&\le |||\rho_0||| + C\left(t + \sqrt{t}\sup_{s\le t}|||\rho_s|||^2\right).
\end{aligned}$$

This implies

$$|||\mathcal{F}[\rho]||| \le |||\rho_0||| + C\left(T + \sqrt{T}|||\rho|||^2\right). \tag{2.14}$$

(In particular, $\mathcal{F}[\rho]_t \in \mathcal{X}$ so $\mathcal{F}$ is well-defined.) By a similar argument, if $\rho^1, \rho^2 \in C([0,T],\mathcal{X})$, then

$$\begin{aligned}
&\big\|\mathcal{F}[\rho^1]_t - \mathcal{F}[\rho^2]_t\big\|_{L^p} \\
&\quad\le \int_0^t \|\nabla G_{t-s}\|_{L^1}\,\big\|\rho_s^1\nabla W \star \rho_s^1 - \rho_s^2 \nabla W \star \rho_s^2\big\|_{L^p}\,\mathrm{d}s \\
&\quad= \int_0^t \|\nabla G_{t-s}\|_{L^1}\,\big\|\rho_s^1\nabla W \star (\rho_s^1 - \rho_s^2) + (\rho_s^1 - \rho_s^2)\nabla W \star \rho_s^2\big\|_{L^p}\,\mathrm{d}s \\
&\quad\le C\int_0^t (t-s)^{-1/2}\Big[\,\big\|\rho_s^1\big\|_{L^p}\,|||\rho_s^1 - \rho_s^2||| + \big\|\rho_s^1 - \rho_s^2\big\|_{L^p}\,|||\rho_s^2|||\Big]\,\mathrm{d}s \\
&\quad\le C\sqrt{t}\sup_{s\le t}\left(|||\rho_s^1||| + |||\rho_s^2|||\right)|||\rho_s^1 - \rho_s^2|||,
\end{aligned}$$

so

$$|||\mathcal{F}[\rho^1] - \mathcal{F}[\rho^2]||| \le C\sqrt{T}\left(|||\rho^1||| + |||\rho^2|||\right)|||\rho^1 - \rho^2|||. \tag{2.15}$$

Put $\mathcal{Y} = \{g \in \mathcal{X}\colon |||g||| \le |||\rho_0||| + 1\}$. Then by (2.14) and (2.15), $T > 0$ can be chosen small enough so that $\mathcal{F}$ maps to itself and is a contraction on the space $C([0,T],\mathcal{Y})$. This shows that

$$T^\star \coloneqq \sup\{t \ge 0\colon \text{there is a solution in } C([0,T],\mathcal{X}) \text{ to } (2.12)\} > 0.$$

Suppose now that ρ and $\widetilde{\rho}$ are two different solutions to (2.12) on $[0,t]$ for some $t \in (0,T^\star]$, starting at ρ_0. Let

$$t_0 \coloneqq \sup\{s \in [0,t]\colon \rho(r) = \widetilde{\rho}(r)\ \forall r \le s\} \in [0,t].$$

Suppose that $t_0 < t$, let $t_1 \coloneqq t_0 + \delta \in (t_0, t)$ for some $\delta \in (0, t - t_0)$, and define $\mathcal{F}\colon C([t_0, t_1], \mathcal{X}) \to C([t_0, t_1], \mathcal{X})$ by

$$\mathcal{F}[u]_s = G_{s-t_0} \star \rho(t_0) + \int_{t_0}^{s} G_{s-r} \star f \,\mathrm{d}r + \int_{t_0}^{s} \nabla G_{s-r} \star (u_r \nabla W \star u_r) \,\mathrm{d}r,$$

for $s \in [t_0, t_1]$. Put $K \coloneqq |||\rho||| \vee |||\widetilde{\rho}|||$. Similarly to before we can show that $|||\mathcal{F}[u]_s||| \leq K + C(\delta + \sqrt{\delta}|||u|||^2)$ for some $C > 0$ and all $s \in [t_0, t_1]$, hence for sufficiently small $\delta > 0$, $\mathcal{F}$ maps to itself and, by an argument identical to that leading to (2.15), is a contraction on $C([t_0, t_1], \mathcal{Y})$, where $\mathcal{Y} = \{g \in \mathcal{X}\colon |||g||| \leq K + 1\}$, so it has a unique fixed point. Since the restrictions of both ρ and $\widetilde{\rho}$ to $[t_0, t_1]$ are fixed points of $\mathcal{F}$, we conclude they must coincide on $[t_0, t_1]$, contradicting the definition of t_0.

We have proved that all solutions must coincide at all times where both are defined, in particular there exists a $\rho \in C([0, T^\star), |||\cdot|||)$ which solves (2.12), and such that any solution $\widetilde{\rho} \in C([0, T), |||\cdot|||)$ of (2.12) satisfies $T \leq T^\star$ and coincides with ρ on $[0, T)$.

Now suppose that $T^\star < \infty$, in which case we show that

$$|||\rho(t)||| \longrightarrow \infty, \quad t \uparrow T^\star.$$

This already implies that $\|\rho(t)\|_\infty \to \infty$ as $t \uparrow T^\star$ because $\sup_{t \in [0, T^\star)} \|\rho(t)\|_{L^1} < \infty$ as a consequence of Lemma 2.11 below. Assume for contradiction that there is a $K > 0$ and a sequence $0 < T_n \uparrow T^\star$ such that $|||\rho(T_n)||| \leq K$ for all $n \in \mathbb{N}$. Then we define $\mathcal{F}_n\colon C([T_n, T_n + \delta], \mathcal{X}) \to C([T_n, T_n + \delta], \mathcal{X})$ by

$$\begin{aligned}\mathcal{F}_n[u]_t = G_{t-T_n} \star \rho(T_n) &+ \int_{T_n}^{t} G_{t-s} \star f \,\mathrm{d}s \\ &+ \int_{T_n}^{t} \nabla G_{t-s} \star (u_s \nabla W \star u_s) \,\mathrm{d}s, \qquad T_n \leq t \leq T_n + \delta,\end{aligned}$$

for some $\delta > 0$. Similarly to before we show $|||\mathcal{F}_n[u]||| \leq K + C(\delta + \sqrt{\delta}|||u|||^2)$ for some $C > 0$, and (2.15) with T replaced by δ. Now choose δ small enough that $\mathcal{F}_n$ (maps to itself and) is a contraction on the space $C([T_n, T_n + \delta], \mathcal{Y})$ where $\mathcal{Y} = \{g \in \mathcal{X}\colon g\colon |||g||| \leq K + 1\}$. This choice of δ can be made independent of n, so there exists $n \in \mathbb{N}$ with $T_n + \delta > T^\star$, and we can concatenate $\rho\big|_{[0, T_n]}$ with the fixed point of $\mathcal{F}_n$ to obtain a solution to (2.12) defined on $[0, T_n + \delta] \supsetneq [0, T^\star]$, a contradiction. This also implies that, if $T^\star < \infty$, there cannot be a solution defined on $[0, T^\star]$.

If f, W, and ρ_0 are radially symmetric, then $\mathcal{F}$ preserves radial symmetry, so the fixed point iteration started at the constant in time function $(\rho_t \equiv \rho_0\colon t \in [0, T))$ is radially symmetric at every step and converges to ρ uniformly on $\mathbb{R}^d$, so ρ is radially symmetric. □

For the remainder of this section, we assume some $\rho_0 \in \mathcal{X}_+$ to be given, and denote by $\rho \in C([0, T^\star), \mathcal{X})$ the unique solution to (2.12) started at ρ_0.

Theorem 2.10 (Regularity)**.** *For any $\gamma \in (0, 2+\gamma_f)$,*

$$\rho \in C([0,T^\star), \mathcal{X}_+) \cap C((0,T^\star), \mathcal{X}^\gamma)$$

In particular, ρ has bounded $C^{2,\gamma_f}(\mathbb{R}^d)$-norm on compact subsets of $(0,T^\star)$, and solves (1.4) in the classical sense on $(0,T^\star)$. If $\|\!|\!|\rho_0|\!|\!|_\gamma < \infty$ (in particular if $\rho_0 \equiv 0$), then $\rho \in C([0,T^\star), \mathcal{X}^\gamma)$.

Proof. Fix $T \in (0,T^\star)$. Then $\sup_{0\le s\le T} \|\!|\!|\rho_s|\!|\!| < \infty$, so it suffices to show that for any $\delta > 0$, $\gamma \in [\gamma_f, 2+\gamma_f)$,

$$\sup_{\delta\le s\le T} \|\!|\!|\rho_s|\!|\!|_{\gamma-\gamma_f} < \infty \implies \sup_{2\delta\le s\le T} \|\!|\!|\rho_s|\!|\!|_\gamma < \infty. \tag{2.16}$$

Indeed, by iteration this implies that $\rho \in C((t_0,t_1), \mathcal{X}^\gamma)$ for all $0 < t_0 < t_1 < T^\star$ and hence $\rho \in C((0,T^\star), \mathcal{X}^\gamma)$, for all $\gamma \in [0, 2+\gamma_f)$.

So let $\gamma \in [\gamma_f, 2+\gamma_f)$, $\delta > 0$, and assume that $\sup_{\delta\le s\le T} \|\!|\!|\rho_s|\!|\!|_{\gamma-\gamma_f} < \infty$. Let $p \in \{1,\infty\}$ and $t \in [2\delta, T]$. Then,

$$\begin{aligned}\|\rho_t\|_{\mathcal{W}^{\gamma,p}} \le \|G_t\|_{\mathcal{W}^{\gamma,1}} \|\rho_0\|_{L^p} + \int_0^t \|G_s\|_{\mathcal{W}^{\gamma-\gamma_f,1}} \|f\|_{\mathcal{W}^{\gamma_f,p}} \,\mathrm{d}s \\ + \int_0^t \|\nabla G_{t-s} \star (\rho_s \nabla W \star \rho_s)\|_{\mathcal{W}^{\gamma,p}} \,\mathrm{d}s.\end{aligned} \tag{2.17}$$

By (2.13),

$$\begin{aligned}\|G_t\|_{\mathcal{W}^{\gamma,1}} \|\rho_0\|_{L^p} + \int_0^t \|G_s\|_{\mathcal{W}^{\gamma-\gamma_f,1}} \|f\|_{\mathcal{W}^{\gamma_f,p}} \,\mathrm{d}s \\ \le C\left(t^{-\gamma/2} + \int_0^t s^{-(\gamma-\gamma_f)/2} \,\mathrm{d}s\right) \\ \le C\left(\delta^{-\gamma/2} + \int_0^T s^{-(\gamma-\gamma_f)/2} \,\mathrm{d}s\right),\end{aligned} \tag{2.18}$$

which is a constant independent of t given δ, γ, T, and finite because $\gamma - \gamma_f < 2$. If $\gamma = \gamma_f$, in particular $\gamma < 1$, then by Young's fractional convolution inequality and Lemma 2.8,

$$\begin{aligned}\int_0^t \|\nabla G_{t-s} \star (\rho_s \nabla W \star \rho_s)\|_{\mathcal{W}^{\gamma,p}} \,\mathrm{d}s &\le \int_0^T \|\nabla G_{t-s}\|_{\mathcal{W}^{\gamma,1}} \|\rho_s \nabla W \star \rho_s\|_{L^p} \,\mathrm{d}s \\ &\le C\left(\int_0^T s^{-(1+\gamma)/2} \,\mathrm{d}s\right) \sup_{0\le s\le T} \|\!|\!|\rho_s|\!|\!|^2,\end{aligned}$$

which is a finite constant depending only on T and γ_f. If $\gamma > \gamma_f$, then we fix $0 < \varepsilon < (1-\gamma_f) \wedge (\gamma - \gamma_f)$, and bound, using Young's fractional convolution inequality,

$$\begin{aligned}\int_0^t \|\nabla G_{t-s} \star (\rho_s \nabla W \star \rho_s)\|_{\mathcal{W}^{\gamma,p}} \,\mathrm{d}s \\ \le \int_0^\delta \|\nabla G_{t-s}\|_{\mathcal{W}^{\gamma,1}} \|\rho_s \nabla W \star \rho_s\|_{L^p} \,\mathrm{d}s \\ + \int_\delta^t \|\nabla G_{t-s}\|_{\mathcal{W}^{\gamma_f+\varepsilon,1}} \|\rho_s \nabla W \star \rho_s\|_{\mathcal{W}^{\gamma-\gamma_f-\varepsilon,p}} \,\mathrm{d}s.\end{aligned} \tag{2.19}$$

By Lemmas III.A.1, III.A.3 and 2.8, putting $\beta \coloneqq \gamma - \gamma_f$,

$$\begin{aligned}\|\rho_s \nabla W \star \rho_s\|_{\mathcal{W}^{\beta-\varepsilon,p}} &\le \|\rho_s\|_{\mathcal{W}^{\beta,p}} \|\nabla W \star \rho_s\|_{\mathcal{W}^{\beta,\infty}} \\ &\le C \|\rho_s\|_{\mathcal{W}^{\beta,p}} |||\rho_s|||_\beta \\ &\le C |||\rho_s|||_\beta^2.\end{aligned}$$

Thus,

$$\begin{aligned}&\int_\delta^t \|\nabla G_{t-s}\|_{\mathcal{W}^{\gamma_f+\varepsilon,1}} \|\rho_s \nabla W \star \rho_s\|_{\mathcal{W}^{\gamma-\gamma_f-\varepsilon,p}} \,\mathrm{d}s \\ &\qquad \le C \sup_{\delta \le s < T} |||\rho_s|||_{\gamma-\gamma_f} \int_0^T s^{-(1+\gamma_f+\varepsilon)/2} \,\mathrm{d}s \le C \sup_{\delta \le s \le T} |||\rho_s|||_{\gamma-\gamma_f}. \quad (2.20)\end{aligned}$$

By Lemma 2.8,

$$\int_0^\delta \|\nabla G_{t-s}\|_{\mathcal{W}^{\gamma,1}} \|\rho_s \nabla W \star \rho_s\|_{L^p} \,\mathrm{d}s \le C \sup_{0 \le s \le T} |||\rho_s|||^2 \cdot \int_\delta^T s^{-(1+\gamma)/2} \,\mathrm{d}s < \infty,$$

which, given δ, T, γ, is a constant independent of t. Combining this with (2.17) to (2.20) gives

$$\|\rho_t\|_{\mathcal{W}^{\gamma,p}} \le C \left(1 + \sup_{\delta \le s \le T} |||\rho_s|||_{\gamma-\gamma_f}^2\right) < \infty,$$

for any $t \in [2\delta, T]$ and $p \in \{1, \infty\}$, for $C > 0$ that does not depend on t given δ, T, and γ. We conclude

$$\sup_{2\delta \le s \le T} |||\rho_s|||_\gamma < \infty$$

as claimed. Then recall from Lemma III.A.2 that $\mathcal{X}^{2+\gamma_f} \hookrightarrow C^{2,\gamma_f}(\mathbb{R}^d)$.

If $\gamma \in [0, 2+\gamma_f)$ and $|||\rho_0|||_\gamma < \infty$, we can bound $|||G_t \star \rho_0|||_\gamma \le \|G_t\|_{L^1} |||\rho_0|||_\gamma = |||\rho_0|||_\gamma < \infty$ which is uniform in $t \in [0, T^\star)$, in contrast to the bound $|||G_t \star \rho_0||| \le Ct^{-\gamma/2}$ that we used in (2.17). This lets us prove (2.16) with $\delta = 0$, giving $\rho \in C([0, T^\star), \mathcal{X}^\gamma)$.

We now show that ρ solves (1.4) in the classical sense on $(0, T^\star)$. Clearly $\rho(0, \cdot) = \rho_0$. It suffices to show now that (1.4) is satisfied at a fixed $t \in (0, T^\star)$ and $x \in \mathbb{R}^d$. Because of the instant regularisation of ρ, we can assume without loss of generality that already $\rho_0 \in \mathcal{X}^\gamma$ for all $\gamma \in [0, 2+\gamma_f)$, so that $\sup_{0 \le s \le t} |||\rho_s|||_\gamma < \infty$ by the above, and thus

$$\sup_{0 \le s \le t} |||\rho_s \nabla W \star \rho_s|||_\gamma < \infty, \quad \gamma \in [0, 2+\gamma_f) \quad (2.21)$$

by Lemmas III.A.3 and 2.8. In particular $\rho_s \nabla W \star \rho_s \in C^2(\mathbb{R}^d)$ for all $s \in [0, t]$, so we can rewrite (2.12) as

$$\rho(t) = G_t \star \rho_0 + \int_0^t G_{t-s} \star f \,\mathrm{d}s + \int_0^t G_{t-s} \star \nabla \cdot (\rho_s \nabla W \star \rho_s) \,\mathrm{d}s.$$

Recall that $G\colon [0,t]\times\mathbb{R}^d\to\mathbb{R}$ is smooth with globally bounded derivatives of any order, and $\partial_t G_t = \frac{1}{2}\Delta G_t$. To show that we can pull the time derivative into the integrals, we establish the following bounds.

$$\begin{aligned}\|(\partial_t G_{t-s})\star f\|_\infty = \frac{1}{2}\|(\Delta G_{t-s})\star f\|_\infty &\le C\|G_{t-s}\star f\|_{\mathcal{W}^{2,\infty}}\\ &\le C\|G_{t-s}\|_{\mathcal{W}^{2-\gamma_f,1}}\|f\|_{\mathcal{W}^{\gamma_f,\infty}}\\ &\le C(t-s)^{-(2-\gamma_f)/2}|||f|||_{\gamma_f},\end{aligned}$$

which is integrable over $(0,t)$. Furthermore,

$$\begin{aligned}\|(\partial_t G_{t-s})\star\nabla\cdot(\rho_s\nabla W\star\rho_s)\|_\infty = \frac{1}{2}\|(\Delta G_{t-s})\star\nabla\cdot(\rho_s\nabla W\star\rho_s)\|_\infty &\le C\|G_{t-s}\star\rho_s\nabla W\star\rho_s\|_{\mathcal{W}^{3,\infty}}\\ &\le C\|G_{t-s}\|_{\mathcal{W}^{1,1}}\|\rho_s\nabla W\star\rho_s\|_{\mathcal{W}^{2,\infty}}\\ &\le C(t-s)^{-1/2}\left(\sup_{0\le s\le t}|||\rho_s|||_2^2\right),\end{aligned}$$

where we recalled (2.21) and Lemma 2.8 in the final step, which is also integrable over $(0,t)$. Hence, $\rho(\cdot,x)$ is differentiable at t and

$$\begin{aligned}\partial_t\rho_t &= (\partial_t G_t)\star\rho_0 + f + \int_0^t(\partial_t G_{t-s})\star f\,\mathrm{d}s + \nabla\cdot(\rho_t\nabla W\star\rho_t)\\ &\qquad + \int_0^t(\partial_t G_{t-s})\star\nabla\cdot(\rho_s\nabla W\star\rho_s)\,\mathrm{d}s\\ &= \frac{1}{2}\Delta\rho_t + \nabla\cdot(\rho_t\nabla W\star\rho_t) + f.\end{aligned}$$

Non-negativity will be proved in the following lemma. □

Lemma 2.11. *For any $t\in[0,T^\star)$, ρ_t is non-negative and $\|\rho_t\|_{L^1} = \|\rho_0\|_{L^1} + t\|f\|_{L^1}$.*

Proof. The idea is to show that $\int\rho(t,x)_-\,\mathrm{d}x = 0$ for all $t\ge 0$. For that purpose, let $(j_\varepsilon)_{\varepsilon>0}$ be a family of smooth and convex functions such that $j_\varepsilon(s) = (-s)\vee 0$ on $\mathbb{R}\backslash[-\varepsilon,0]$, and $0\le j_\varepsilon''\le 2/\varepsilon$ in $[-\varepsilon,0]$. Then for any $\varepsilon>0$,

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}\int j_\varepsilon(\rho(t,x))\,\mathrm{d}x &= \int_{\mathbb{R}^d} j_\varepsilon'(\rho(t,x))\left(\frac{1}{2}\Delta\rho_t + \nabla\cdot(\rho_t\nabla W\star\rho_t) + f\right)(x)\,\mathrm{d}x\\ &\le -\int_{\mathbb{R}^d} j_\varepsilon''(\rho(t,x))\nabla\rho_t(x)\cdot\left(\frac{1}{2}\nabla\rho_t + \rho_t\nabla W\star\rho_t\right)(x)\,\mathrm{d}x\\ &\le -\int_{\mathbb{R}^d} J_\varepsilon(\rho(t,x))\nabla\rho_t(x)\cdot(\nabla W\star\rho_t)(x)\,\mathrm{d}x\\ &=: H_\varepsilon(t),\end{aligned}\tag{2.22}$$

where in the second step we used that $j_\varepsilon'\le 0$ and $f\ge 0$, and in the third step we used that $j_\varepsilon''(\rho)\,|\nabla\rho|^2\ge 0$ and put $J_\varepsilon(s) := j_\varepsilon''(s)s$. To understand why this was helpful,

formally replace $j_\varepsilon(s)$ by $j(s) \coloneqq (-s) \vee 0$, so that the left-hand side (LHS) turns into $\frac{\mathrm{d}}{\mathrm{d}t} \int \rho(t,x)_- \,\mathrm{d}x$, and the RHS is zero because $j''(x) = \delta(x)$, so $J(\rho) = \delta(\rho)\rho = 0$.

More precisely, (2.22) implies $\int j_\varepsilon(\rho(t,x))\,\mathrm{d}x \le j_\varepsilon(\rho_0(x)) + \int_0^t H_\varepsilon(s)\,\mathrm{d}s$ for all $\varepsilon > 0$ and $t \in [0, T^\star)$. The LHS converges by dominated convergence to $\int \rho(t,x)_- \,\mathrm{d}x$, and the RHS converges by dominated convergence to zero. Indeed, $J_\varepsilon \to 0$ pointwise on $\mathbb{R} \setminus \{0\}$ as $\varepsilon \to 0$, and $0 \le J_\varepsilon(s) \le 2$ for all $s \in \mathbb{R}$ and $\varepsilon > 0$ by assumption on j_ε. Hence, $\int \rho(t,x)_- \,\mathrm{d}x = 0$ for all $t \in [0, T^\star)$, so by continuity of $\rho(t)$, it is non-negative for all $t \in [0, T^\star)$.

For the second claim,

$$\begin{aligned}\int \rho(t,x)\,\mathrm{d}x &= \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} G_t(x-y)\rho_0(y)\,\mathrm{d}y\,\mathrm{d}x + \int_0^t \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} G_s(x-y) f(y)\,\mathrm{d}y\,\mathrm{d}x\,\mathrm{d}s \\ &\qquad + \int_0^t \int\int (\nabla G_{t-s})(x-y)(\rho_s \nabla W \star \rho_s)(y)\,\mathrm{d}y\,\mathrm{d}x\,\mathrm{d}s \\ &= \int \rho_0(x)\,\mathrm{d}x + t\left(\int_{\mathbb{R}^d} f(x)\,\mathrm{d}x\right),\end{aligned}$$

where we used that $\int G_s = 1$ and $\int \nabla G_s = 0$ for all $s > 0$. □

2.2.2 Maximum principle and global boundedness

The main goal of this section is to prove Theorems 2.3 to 2.5. For clarity of presentation, we will postpone the proofs of some auxiliary lemmas until after we show Theorems 2.3 to 2.5.

The representation of c_W used mainly in the proofs is $c_W = \eta_W - \int (\Delta W)_+$, and a crucial first step is to establish the following meaning of η_W.

Lemma 2.12. *There is $\xi_W \in \mathbb{R}^d$ such that for any $g \in \mathcal{X}^1 \cap C^1(\mathbb{R}^d)$ (in particular any $g \in \mathcal{X}^\gamma$ with $\gamma > 1$),*

$$\nabla \cdot (\nabla W \star g) = (\Delta W) \star g - \eta_W g + \xi_W \cdot \nabla g.$$

If $W(x) = W(-x)$ for all $x \in \mathbb{R}^d \setminus \{0\}$, then $\xi_W = 0$.

That is, ΔW can be viewed in a distributional sense as a sum of the function $x \mapsto \Delta W(x)$ on $\mathbb{R}^d \setminus \{0\}$, and the distribution $-\eta_W \delta_0 - \xi_W \cdot \nabla \delta_0$, where δ_0 denotes the Dirac mass at the origin. The significance of this lies in the following application: If g is a non-negative function with a global maximum at some x_0, then $\nabla g(x_0) = 0$ and $g(x_0) = \|g\|_\infty$, so

$$\begin{aligned}\nabla \cdot (\nabla W \star g)(x_0) &= (\Delta W \star g)(x_0) - \eta_W g(x_0) \\ &\le \|g\|_\infty \int (\Delta W)_+ - \eta_W g(x_0) \\ &= -c_W g(x_0).\end{aligned}$$

This already proves the first assertion of the following lemma. Denote by $C^\infty_{c+}(\mathbb{R}^d)$ the set of infinitely differentiable, compactly supported, non-negative functions on $\mathbb{R}^d$.

Lemma 2.13. *If $g \in \mathcal{X}^1 \cap C^1(\mathbb{R}^d)$ is non-negative and has a global maximum at $x_0 \in \mathbb{R}^d$, then*

$$\nabla \cdot (\nabla W \star g)(x_0) \leq -c_W g(x_0). \tag{2.23}$$

For any $c > c_W$ and $x_0 \in \mathbb{R}^d$, there exists $g \in C^\infty_{c+}(\mathbb{R}^d)$ such that g has a global maximum of any given height at x_0, $\Delta g(x_0) = 0$, and

$$\nabla \cdot (\nabla W \star g)(x_0) > -c g(x_0).$$

Note that if $W = W_N$ is the Newtonian potential, then $\nabla \cdot (\nabla W \star g) = -g$ which trivially implies the assertion of Lemma 2.13. Now if $t > 0$ and ρ_t has a global maximum at some $x_0 \in \mathbb{R}^d$, then $\Delta \rho_t(x_0) \leq 0$, $\nabla \rho_t(x_0) = 0$, and $\rho_t(x_0) = \|\rho_t\|_\infty$, so evaluating the PDE (1.4) at (t, x_0) and using Lemma 2.12 gives

$$\begin{aligned} \partial_t \rho_t(x_0) &\leq \rho_t(x_0) \nabla \cdot (\nabla W \star \rho_t)(x_0) + f(x_0) \\ &\leq \|f\|_\infty - c_W \|\rho_t\|_\infty^2 . \end{aligned} \tag{2.24}$$

This is almost the differential inequality in Theorem 2.5, where the LHS is $\partial_t^+ \|\rho_t\|_\infty$, and the following two lemmas close that gap.

Lemma 2.14. *Suppose $T > 0$ and $g \in C([0,T], \mathcal{X}_+)$ such that $g(t, \cdot)$ is Lipschitz continuous for all $t \in (0,T]$, and g is differentiable in time on $(0,T]$ with $\partial_t g \colon (0,T] \times \mathbb{R}^d \to \mathbb{R}$ jointly continuous. Suppose further that there is a continuous function $C \colon [0,T] \to [0,\infty)$ such that for all $t \in (0,T]$, if g_t has a global maximum at x,*

$$\partial_t g(t,x) \leq C(t).$$

Then for all $t \in [0,T)$,

$$\partial_t^+ \|g_t\|_\infty \leq C(t).$$

Lemma 2.15. *For any $\rho_0 \in \mathcal{X}_+$ and $T \in (0, T^\star)$, $\rho \colon [0,T] \times \mathbb{R}^d \to [0,\infty)$ satisfies the assumptions of Lemma 2.14.*

We can now prove Theorem 2.5.

Proof of Theorem 2.5. The differential inequality (2.3) follows from (2.24) and Lemmas 2.14 and 2.15, so it remains to prove sharpness. Let $c > c_W$ and $h > 0$, put $c' \coloneqq (c + c_W)/2 > c_W$, and choose $x_0 \in \mathbb{R}^d$ with $f(x_0) \geq \|f\|_\infty - (c - c_W)h^2/2$. Then by Lemma 2.13, there is $\rho_0 \in C^\infty_{c+}(\mathbb{R}^d)$ such that ρ_0 has a global maximum of height h at

x_0, $\Delta\rho_0(x_0) = 0$, $\rho_0(x_0) = \|\rho_0\|_\infty = h$, and $\nabla \cdot (\nabla W \star \rho_0)(x_0) > -c'\rho_0(x_0) = -c'h$. Then, since $\nabla\rho_0(x_0) = 0$,

$$\begin{aligned}\partial_t \rho(t, x_0)\Big|_{t=0} &= \rho_0(x_0)\nabla \cdot (\nabla W \star \rho_0)(x_0) + f(x_0) \\ &> -c'h^2 + \|f\|_\infty - (c - c_W)h^2/2 \\ &= \|f\|_\infty - c\,\|\rho_0\|_\infty^2 .\end{aligned}$$

In particular,

$$\begin{aligned}\partial_t^+ \|\rho_t\|_\infty \Big|_{t=0} = \varlimsup_{t\downarrow 0} \frac{\|\rho_t\|_\infty - \|\rho_0\|_\infty}{t} &\geq \varlimsup_{t\downarrow 0} \frac{\rho(t, x_0) - \rho_0(x_0)}{t} \\ &= \partial_t \rho(t, x_0)\Big|_{t=0} \\ &> \|f\|_\infty - c\,\|\rho_0\|_\infty^2 .\end{aligned}$$

□

As we mentioned in Section 2.1, the differential inequality (2.3) can be turned into an explicit upper bound, recall also Fig. 2.1 for an illustration.

Corollary 2.16. *If $c_W > 0$ and $\rho_0 \in \mathcal{X}_+$, then $T^\star = \infty$ and for all $t \geq 0$,*

$$\|\rho_t\|_\infty \leq M \begin{cases} \coth(Mc_W(t + t_0)), & \|\rho_0\|_\infty > M, \\ \tanh(Mc_W(t + t_0)), & \|\rho_0\|_\infty < M, \\ 1, & \|\rho_0\|_\infty = M, \end{cases} \tag{2.25}$$

where $M = \sqrt{\|f\|_\infty / c_W}$, and $t_0 \in \mathbb{R}$ is such that the right-hand side at $t = 0$ is $\|\rho_0\|_\infty$.

Proof. Suppose $c_W > 0$ and let $M \coloneqq \sqrt{\|f\|_\infty / c_W}$. Then the unique solution to the ODE

$$\begin{cases} G'(t) = \|f\|_\infty - c_W G(t)^2, & t \geq 0, \\ G(0) = \|\rho_0\|_\infty , \end{cases}$$

is given by

$$G(t) = M \begin{cases} \coth\left(c_W M(t + t_0)\right), & G(0) > M, \\ \tanh\left(c_W M(t + t_0)\right), & G(0) < M, \\ 1, & G(0) = M, \end{cases}$$

where $t_0 \in \mathbb{R}$ is such that $G(0) = \|\rho_0\|_\infty$. Then (2.3) implies $\partial_t^+ \|\rho_t\|_\infty \leq G'(t)$ for all $t \in [0, T^\star)$ and thus $\|\rho_t\|_\infty \leq G(t)$ for all $t \in [0, T^\star)$. In particular, if $T^\star$ were finite then $\|\rho_t\|_\infty \leq G(t) \to G(T^\star) < \infty$ as $t \uparrow T^\star$, contradicting Theorem 2.9. □

This also proves Proposition 2.2. The positive assertions of Theorems 2.3 and 2.4 are now straightforward consequences of Corollary 2.16, and the sharpness follows from Theorem 2.5.

Proof of Theorems 2.3 and 2.4. Suppose that $c_W > 0$, put $M \coloneqq \sqrt{\|f\|_\infty / c_W}$, and let $\rho_0 \in \mathcal{X}_+$. If $\|\rho_0\|_\infty = M$, then (2.25) implies $\|\rho_t\|_\infty \leq M = \|\rho_0\|_\infty$ for all $t \in [0, T^\star)$. If $\|\rho_0\|_\infty < M$, then (2.25) implies

$$\|\rho_t\|_\infty \leq M \tanh(c_W M(t + t_0)) < M,$$

for all $t \in [0, T^\star)$. If $\|\rho_0\|_\infty > M$, then

$$\|\rho_t\|_\infty \leq M \coth(c_W(t + t_0)) =: G(t),$$

for all $t \in [0, T^\star)$. Since $G(0) > M$, we must have $t_0 > 0$. Thus, G is strictly decreasing on $[0, \infty)$, so

$$\|\rho_t\|_\infty \leq G(t) < G(0) = \|\rho_0\|_\infty, \quad t \in (0, T^\star).$$

For the sharpness statements, suppose first that $c_W > 0$ and $M < \sqrt{\|f\|_\infty / c_W}$. Then it suffices to show that there exists $\rho_0 \in \mathcal{X}_+$ with $\|\rho_0\|_\infty = M$ and $\partial_t^+ \|\rho_t\|_\infty \big|_{t=0}$ > 0. For that purpose note that $\|f\|_\infty - c_W M^2 > 0$, so there is $c > c_W$ such that still $\|f\|_\infty - cM^2 > 0$. Then there exists by Theorem 2.5 a $\rho_0 \in C_{c+}^\infty(\mathbb{R}^d)$ with $\|\rho_0\|_\infty = M$ and

$$\partial_t^+ \|\rho_t\|_\infty \Big|_{t=0} > \|f\|_\infty - c \|\rho_0\|_\infty^2 = \|f\|_\infty - cM^2 > 0.$$

Now suppose that $c_W < 0$ and $M > 0$. Then again it suffices to show that there exists $\rho_0 \in \mathcal{X}_+$ with $\|\rho_0\|_\infty = M$ and $\partial_t^+ \|\rho_t\|_\infty > 0$. Let $c_W < c < 0$, then by Theorem 2.5 there exists $\rho_0 \in C_{c+}^\infty(\mathbb{R}^d)$ with $\|\rho_0\|_\infty = M$ and

$$\partial_t^+ \|\rho_t\|_\infty \Big|_{t=0} > \|f\|_\infty - c \|\rho_0\|_\infty^2 \geq \|f\|_\infty \geq 0.$$

□

In the following subsections, we prove Lemmas 2.12 to 2.15, as well as Lemma 2.7 (stated in Section 2.1) which establishes existence of the limits defining η_W and α_W. Finally, we prove the claim made in the introduction regarding the asymptotics of (1.4) with $W \equiv 0$.

Proofs of Lemmas 2.7 and 2.12

Using the same argument as in the proof of Lemma 2.8, we can show that Assumption (A)(ii) implies

$$\|\Delta W \star g\|_{\mathcal{W}^{\gamma,\infty}} \leq C \|\!|g|\!\|_\gamma \tag{2.26}$$

for any $\gamma \geq 0$ and $g \in \mathcal{X}^\gamma$ (where $C > 0$ depends on γ but not g). We further need the following technical lemma.

Lemma 2.17. *There is $C > 0$ such that for any $g \in \mathcal{W}^{1,\infty}$ with $g(0) = 0$ and any $\varepsilon > 0$, there exists $\widetilde{g} \in \mathcal{W}^{1,\infty}$ with $\widetilde{g} = g$ on $B(0, \varepsilon)$, $\widetilde{g} \equiv 0$ on $\mathbb{R}^d \setminus B(0, 2\varepsilon)$, and $\|\widetilde{g}\|_{\mathcal{W}^{1,\infty}} \leq C \|g\|_{\mathcal{W}^{1,\infty}(B(0,\varepsilon))}$.*

Proof. Define $\widetilde{g}$ in radial coordinates by

$$\widetilde{g}(r,\varphi) \coloneqq \begin{cases} g(r,\varphi), & r < \varepsilon, \\ g(2\varepsilon - r, \varphi), & \varepsilon \le r < 2\varepsilon, \\ 0, & 2\varepsilon \le r, \end{cases}$$

where φ stands collectively for all $d-1$ angular variables. Then $\widetilde{g}$ is continuous, $\|\widetilde{g}\|_{L^\infty} \le \|g\|_{L^\infty(B(0,\varepsilon))}$ and $\|\widetilde{g}\|_{\mathrm{Lip}} \le \|g\|_{\mathrm{Lip}(B(0,\varepsilon))}$, where

$$\|h\|_{\mathrm{Lip}(\Omega)} \coloneqq \sup_{\substack{x,y\in\Omega \\ x\neq y}} \frac{|h(x)-h(y)|}{|x-y|}$$

for $h\colon \mathbb{R}^d \to \mathbb{R}$ and $\Omega \subset \mathbb{R}^d$. Then recall the well-known result that $\mathcal{W}^{1,\infty} = C^{0,1}(\mathbb{R}^d)$, the space of bounded Lipschitz functions. □

The following lemma is identical to Lemma 2.12 except that η_W is replaced by a possibly different value η'_W. We then show together with Lemma 2.7 (which proves existence of α_W and η_W) that $\eta_W = \eta'_W$.

Lemma 2.18. *There exist $\eta'_W \in \mathbb{R}$ and $\xi_W \in \mathbb{R}^d$ such that for any $g \in \mathcal{X}^1 \cap C^1(\mathbb{R}^d)$ (in particular any $g \in \mathcal{X}^\gamma$ with $\gamma > 1$),*

$$\nabla \cdot (\nabla W \star g) = (\Delta W) \star g - \eta'_W g + \xi_W \cdot \nabla g.$$

If $W(x) = W(-x)$ for all $x \in \mathbb{R}^d \setminus \{0\}$, then $\xi_W = 0$.

Proof. We prove the claim at a fixed but arbitrary $x_0 \in \mathbb{R}^d$, and assume without loss of generality that $x_0 = 0$. Recall that $\nabla W \star g \in \mathcal{W}^{1,\infty}$ for any $g \in \mathcal{X}^1$ by Lemma 2.8, so it is bounded and globally Lipschitz continuous. This remains true if we replace g by $g\varphi_\varepsilon$ for a smooth radially symmetric bump function φ_ε with $\mathbb{1}_{B(0,\varepsilon)} \le \varphi_\varepsilon \le \mathbb{1}_{B(0,2\varepsilon)}$. Put $\psi_\varepsilon \coloneqq 1 - \varphi_\varepsilon$. Then,

$$\nabla \cdot (\nabla W \star g) = \nabla \cdot (\nabla W \star g\psi_\varepsilon) + \nabla \cdot (\nabla W \star g\varphi_\varepsilon).$$

Now,

$$\nabla \cdot (\nabla W \star g\psi_\varepsilon)(0) = \nabla \cdot \left(\int \nabla W(\cdot - y) g(y) \psi_\varepsilon(y) \,\mathrm{d}y \right)(0) = (\Delta W \star g\psi_\varepsilon)(0),$$

because ΔW is bounded away from zero. More precisely, as long as $x \in B(0,\varepsilon/2)$, say, then $|x-y| \ge |y| - |x| \ge \varepsilon/2$ for any $y \in \mathbb{R}^d$ with $\psi_\varepsilon(y) > 0$, and $\sup_{B(0,\varepsilon/2)^c} \Delta W < \infty$ by Assumption (A)(ii). By (2.26), $\Delta W \star g$ is well-defined, hence by dominated convergence, $(\Delta W \star g\psi_\varepsilon)(0) \to (\Delta W \star g)(0)$ as $\varepsilon \to 0$. Thus,

$$\mathcal{F}[g] \coloneqq \lim_{\varepsilon\to 0} \nabla \cdot (\nabla W \star g\varphi_\varepsilon)(0) = \nabla \cdot (\nabla W \star g)(0) - (\Delta W \star g)(0)$$

exists and is finite for any $g \in \mathcal{X}^1$. Inspecting the RHS above we see that

$$\mathcal{F} \text{ is linear and } |\mathcal{F}[g]| \le C|||g|||_1.$$

Indeed, $|\mathcal{F}[g]| \le \|\nabla \cdot (\nabla W \star g)\|_{L^\infty} + \|\Delta W \star g\|_{L^\infty}$, and

$$\|\nabla \cdot (\nabla W \star g)\|_{L^\infty} \le \|\nabla W \star g\|_{\mathcal{W}^{1,\infty}} \le C|||g|||_1$$

by Lemma 2.8, and $\|\Delta W \star g\|_{L^\infty} \le C|||g|||_1$ by (2.26). Furthermore, from the definition it is immediate that $\mathcal{F}[g]$ depends on g only locally around 0, in the sense that $\mathcal{F}[g_1] = \mathcal{F}[g_2]$ if $g_1, g_2 \in \mathcal{X}^1$ coincide in a neighbourhood of 0.

Now suppose that $g \in \mathcal{X}^1$ with $g(0) = 0$. Let $\varepsilon \in (0,1)$, then by Lemma 2.17 there is $\widetilde{g} \in \mathcal{X}^1$ such that $\widetilde{g} = g$ on $B(0,\varepsilon)$, $\widetilde{g} \equiv 0$ on $\mathbb{R}^d \setminus B(0,2\varepsilon)$ and $|||\widetilde{g}|||_1 \le C\,\|\widetilde{g}\|_{\mathcal{W}^{1,\infty}} \le C\,\|g\|_{\mathcal{W}^{1,\infty}(B(0,\varepsilon))}$ with constants independent of g and ε (the first inequality holds because $\widetilde{\rho}$ is supported in $B(0,\varepsilon) \subset B(0,1)$). This implies that $|\mathcal{F}[g]| = |\mathcal{F}[\widetilde{g}]| \le C|||\widetilde{g}|||_1 \le C\,\|g\|_{\mathcal{W}^{1,\infty}(B(0,\varepsilon))}$. Letting $\varepsilon \to 0$ implies

$$\forall g \in \mathcal{X}^1, g(0) = 0\colon\ |\mathcal{F}[g]| \le C \varliminf_{\varepsilon \to 0} \|g\|_{\mathcal{W}^{1,\infty}(B(0,\varepsilon))}\,.$$

In particular, if $g \in \mathcal{X}^1 \cap C^1(\mathbb{R}^d)$ and $g(0) = \nabla g(0) = 0$ then $\mathcal{F}[g] = 0$. By linearity, there must be $\eta'_W \in \mathbb{R}$ and $\xi_W \in \mathbb{R}^d$ with

$$\mathcal{F}[g] = -\eta'_W g(0) + \xi_W \cdot \nabla g(0)$$

for all $g \in \mathcal{X}^1 \cap C^1(\mathbb{R}^d)$. Now suppose that $W(x) = W(-x)$ for all $x \in \mathbb{R}^d$. To show that $\xi_W = 0$ it suffices to prove that whenever $g \in \mathcal{X}^1 \cap C^1(\mathbb{R}^d)$ with $g(0) = 0$, then $\mathcal{F}[g] = 0$. Since we already know that $\mathcal{F}[g]$ does not change when we replace g by a function in $\mathcal{X}^1 \cap C^1(\mathbb{R}^d)$ that has the same value and gradient at 0, we may assume that $g(x) = \beta \cdot x$ for all $x \in B(0,1)$, where $\beta = \nabla g(0)$. Then,

$$\begin{aligned}\nabla \cdot (\nabla W \star g\varphi_\varepsilon)(0) &= \beta \cdot (\nabla W \star \varphi_\varepsilon)(0) + (\nabla W \star g \nabla \varphi_\varepsilon)(0) \\ &= \beta \cdot \int \nabla W(x)\varphi_\varepsilon(x)\,\mathrm{d}x + \beta \cdot \int x(\nabla W(x) \cdot \nabla \varphi_\varepsilon(x))\,\mathrm{d}x = 0,\end{aligned}$$

because both integrands are antisymmetric, hence $\mathcal{F}[g] = 0$. □

We now prove Lemma 2.7, which establishes existence of η_W and α_W, and show that $\eta_W = \eta'_W$. Note that $\langle \nabla W_N \rangle(r) = -c_d^{-1} r^{1-d}$, so equivalently to (2.6) we could write

$$\eta_W = -c_d \lim_{r \to 0} r^{d-1} \langle \nabla W \rangle(r), \qquad \alpha_W = -c_d \lim_{R \to \infty} R^{d-1} \langle \nabla W \rangle(R).$$

Proof of Lemma 2.7. Let $(g_\varepsilon)_{\varepsilon>0}$ be a smooth approximation to unity with g_ε supported in $B(0,\varepsilon)$ for all $\varepsilon > 0$. Then by Lemma 2.18 and (2.4), and recalling that $c_d > 0$ denotes the surface area of the unit ball in $\mathbb{R}^d$,

$$\begin{aligned}\int\limits_{\partial B(0,R)} (\nabla W \star g_\varepsilon) \cdot \mathrm{d}\widehat{n} &= \int_{B(0,R)} \nabla \cdot (\nabla W \star g_\varepsilon)(x)\,\mathrm{d}x \\ &= -\eta'_W \int_{B(0,R)} g_\varepsilon(x)\,\mathrm{d}x + \xi_W \cdot \int_{B(0,R)} \nabla g_\varepsilon(x)\,\mathrm{d}x \\ &\qquad + \int_{B(0,R)} (\Delta W \star g_\varepsilon)(x)\,\mathrm{d}x \\ &= -\eta'_W + \int_{B(0,R)} (\Delta W \star g_\varepsilon)(x)\,\mathrm{d}x,\end{aligned} \tag{2.27}$$

where we used that $\int_{B(0,R)} \nabla g_\varepsilon(x)\,\mathrm{d}x = 0$ by integration by parts. We want to take $\varepsilon \to 0$ for fixed $R > 0$. On the LHS we obtain

$$\int\limits_{\partial B(0,R)} (\nabla W \star g) \cdot \mathrm{d}\widehat{n} \longrightarrow \int\limits_{\partial B(0,R)} \nabla W \cdot \mathrm{d}\widehat{n} = c_d R^{d-1} \langle \nabla W \rangle (R),$$

by dominated convergence, which is applicable because ∇W is bounded away from the origin by Assumption (A)(ii). For the RHS, we cannot directly apply dominated or monotone convergence without further assumptions on ΔW. Instead, we write $\Delta W \star g_\varepsilon = (\Delta W)_+ \star g_\varepsilon - (\Delta W)_- \star g_\varepsilon$, and

$$\begin{aligned}\int_{B(0,R)} [(\Delta W)_+ \star g_\varepsilon](x)\,\mathrm{d}x &= \int_{B(0,R)} \int_{B(0,\varepsilon)} (\Delta W)_+(x-y) g_\varepsilon(y)\,\mathrm{d}y\,\mathrm{d}x \\ &= \int_{B(0,\varepsilon)} g_\varepsilon(y) \left(\int_{B(y,R)} (\Delta W)_+(x)\,\mathrm{d}x \right) \mathrm{d}y.\end{aligned} \tag{2.28}$$

Now, for any $y \in B(0,\varepsilon)$, if say $\varepsilon < R/2$,

$$\begin{aligned}\left| \int_{B(0,R)} \Delta W(x)_+\,\mathrm{d}x - \int_{B(y,R)} \Delta W(x)_+\,\mathrm{d}x \right| &\le \int\limits_{B(0,R)\Delta B(y,R)} |\Delta W(x)|\,\mathrm{d}x \\ &\le C\varepsilon R^{d-1} \sup_{|z|\ge R/2} |\Delta W(z)|,\end{aligned}$$

where $A \Delta B = (A \setminus B) \cup (B \setminus A)$ for $A, B \subset \mathbb{R}^d$. This goes to zero as $\varepsilon \to 0$, so continuing in (2.28),

$$\begin{aligned}&\int_{B(0,\varepsilon)} g_\varepsilon(y) \left(\int_{B(y,R)} (\Delta W)_+(x)\,\mathrm{d}x \right) \mathrm{d}y \\ &\qquad = \int_{\mathbb{R}^d} g_\varepsilon(y) \left(\int_{B(0,R)} (\Delta W)_+(x)\,\mathrm{d}x + O(\varepsilon) \right) \mathrm{d}y \\ &\qquad \longrightarrow \int_{B(0,R)} (\Delta W)_+(x)\,\mathrm{d}x,\end{aligned}$$

as $\varepsilon \to 0$. We can proceed similarly with $(\Delta W)_-$ and combine the results to obtain

$$\int_{B(0,R)} (\Delta W \star g_\varepsilon)(x)\,\mathrm{d}x \longrightarrow \int_{B(0,R)} \Delta W(x)\,\mathrm{d}x$$

as $\varepsilon \to 0$. We have now shown that letting $\varepsilon \to 0$ in (2.27) yields

$$-\eta'_W + \int_{B(0,R)} \Delta W(x)\,\mathrm{d}x = c_d R^{d-1} \langle \nabla W \rangle (R) = -\frac{\langle \nabla W \rangle (R)}{\langle \nabla W_N \rangle (R)},$$

where we recalled $\langle \nabla W_N \rangle (R) = -c_d^{-1} R^{1-d}$. If we let $R \to 0$, then the LHS tends to $-\eta'_W$ because ΔW is locally integrable. Thus the limit on the RHS exists, so η_W in (2.6) is well-defined and equals η'_W. If we let $R \to \infty$, then the LHS tends to $-\eta_W + \int \Delta W$ by Assumption (A)(iii), which implies that α_W in (2.6) is well-defined and that $\alpha_W = \eta_W - \int(\Delta W)$ holds. □

Proof of Lemma 2.13

We already proved (2.23) just before stating Lemma 2.13, so it remains to prove sharpness.

Proof of Lemma 2.13, Sharpness. Let $c > c_W$, $x_0 \in \mathbb{R}^d$ and $M > 0$. We will construct a function $g \in C^\infty_{c+}(\mathbb{R}^d)$ such that g has a global maximum of height M at x_0 and $\nabla \cdot (\nabla W \star g)(x_0) > -cg(x_0)$. Because this inequality is linear in g, we may assume $M = 1$. Put $U \coloneqq \big\{x \in \mathbb{R}^d \setminus \{0\} : \Delta W(x) > 0\big\}$, which is an open, possibly empty subset of $\mathbb{R}^d \setminus \{0\}$. For $\varepsilon > 0$ let

$$U_\varepsilon \coloneqq \{x \in U \colon d(x, U^c) > \varepsilon\},$$

which is also open and $\bigcup_{\varepsilon>0} U_\varepsilon = U$ (note that we shrink, not grow U by ε). Put $h_\varepsilon \coloneqq \varphi_{\varepsilon/2} \star \mathbb{1}_{U_\varepsilon \cap B(0,\varepsilon^{-1})}$ where $\varphi \in C^\infty_{c+}(B(0,1))$ with $\int \varphi = 1$ and $\varphi_\varepsilon(x) \coloneqq \varepsilon^{-d}\varphi(\varepsilon^{-1}x)$. Then $h_\varepsilon \in C^\infty_{c+}(\mathbb{R}^d)$ and

$$\mathbb{1}_{U_{2\varepsilon} \cap B(0,(2\varepsilon)^{-1})} \le h_\varepsilon \le \mathbb{1}_{U_{\varepsilon/2}}. \tag{2.29}$$

In particular $0 \le h_\varepsilon \uparrow \mathbb{1}_U$ as $\varepsilon \to 0$, and, since h_ε is supported on U for all $\varepsilon > 0$, and by monotone convergence,

$$\int_{\mathbb{R}^d} \Delta W(x) h_\varepsilon(x)\,\mathrm{d}x = \int_{\mathbb{R}^d} (\Delta W)_+(x) h_\varepsilon(x)\,\mathrm{d}x \uparrow \int (\Delta W)_+, \quad \varepsilon \to 0. \tag{2.30}$$

Furthermore, by (2.29), $h_\varepsilon \equiv 0$ on $B(0, \varepsilon/2)$, in particular $h_\varepsilon(0) = \Delta h_\varepsilon(0) = 0$.

Now let $f_\varepsilon \in C^\infty_{c+}(B(x_0, \varepsilon/2))$ with $f_\varepsilon(x_0) = 1$, $\Delta f_\varepsilon(x_0) = 0$, $0 \le f_\varepsilon \le 1$. In particular $f_\varepsilon \to 0$ a.e. as $\varepsilon \to 0$, and by dominated convergence (recall that ΔW is locally integrable by Assumption (A)(ii)),

$$(\Delta W \star f_\varepsilon)(x) \to 0, \quad \varepsilon \to 0, \tag{2.31}$$

for any $x \in \mathbb{R}^d$. Now put

$$g_\varepsilon \coloneqq h_\varepsilon(x_0 - \cdot) + f_\varepsilon \in C^\infty_{c+}(\mathbb{R}^d), \qquad \varepsilon > 0.$$

Recall that $h_\varepsilon(x_0 - \cdot)$ is supported in $B(x_0, \varepsilon/2)^c$ and f_ε is supported in $B(x_0, \varepsilon/2)$, and both are upper bounded by 1, so $0 \le g_\varepsilon \le 1$. Furthermore,

$$g_\varepsilon(x_0) = h_\varepsilon(0) + f_\varepsilon(x_0) = 1, \qquad \Delta g_\varepsilon(x_0) = \Delta h_\varepsilon(0) + \Delta f_\varepsilon(x_0) = 0,$$

so g_ε attains a global maximum at x_0, and by (2.30) and (2.31),

$$(\Delta W \star g_\varepsilon)(x_0) = (\Delta W \star h_\varepsilon)(0) + (\Delta W \star f_\varepsilon)(x_0) \to \int (\Delta W)_+, \quad \varepsilon \to 0.$$

Now choose $\varepsilon > 0$ so small that $(\Delta W \star g_\varepsilon)(x_0) > \int (\Delta W)_+ - (c - c_W)$, then

$$\begin{aligned} \nabla \cdot (\nabla W \star g_\varepsilon) &= (\Delta W \star g_\varepsilon)(x_0) - \eta_W g_\varepsilon(x_0) \\ &= (\Delta W \star g_\varepsilon)(x_0) - \int (\Delta W)_+ - c_W \\ &> -c \\ &= -c g_\varepsilon(x_0), \end{aligned}$$

where we used $\eta_W = c_W + \int (\Delta W)_+$, see (2.7). □

Proof of Lemmas 2.14 and 2.15

Proof of Lemma 2.14. It suffices to show that for any fixed $t_0 \in [0, T)$, and every $\varepsilon, \varepsilon' > 0$,

$$\|g_t\|_\infty \le \|g_{t_0}\|_\infty + \int_{t_0}^t \Big(C(s) + \varepsilon\Big)\, \mathrm{d}s + \varepsilon', \quad t \in [t_0, T]. \tag{2.32}$$

Indeed, taking $\varepsilon' \downarrow 0$, this implies that (2.32) holds for all $t \in [t_0, T]$, $\varepsilon > 0$, and $\varepsilon' = 0$. Thus, for every $\varepsilon > 0$,

$$\partial_t^+ \|g_t\|_\infty \Big|_{t=t_0} = \varlimsup_{t \downarrow t_0} \frac{\|g_t\|_\infty - \|g_{t_0}\|_\infty}{t - t_0} \le \varlimsup_{t \downarrow t_0} \frac{1}{t - t_0} \int_{t_0}^t \Big(C(s) + \varepsilon\Big)\, \mathrm{d}s = C(t_0) + \varepsilon,$$

which implies the claim at t_0 by taking $\varepsilon \downarrow 0$.

Now fix $t_0 \in [0, T)$, $\varepsilon, \varepsilon' > 0$, and we show (2.32). We may assume that $t_0 = 0$. We want to show that

$$F(t) \coloneqq \|g_t\|_\infty - \|g_0\|_\infty - \int_0^t \Big(C(s) - \varepsilon\Big)\, \mathrm{d}s - \varepsilon' \le 0 \tag{2.33}$$

for all $t \in [0, T]$. Define

$$t_1 \coloneqq \sup\{t \in [0, T] \colon F(t) \le 0\}.$$

F is continuous because $g \in C([0, T], L^\infty)$, and $F(0) = -\varepsilon' < 0$, so $t_1 > 0$. Now if $F(t) \le 0$ does not hold for all $t \in [0, T]$, then $t_1 < T$, so there would be $T > t_n \downarrow t_1 > 0$

such that $F(t) \le 0$ for all $0 \le t \le t_1$ and $F(t_n) > 0$ for all $n \in \mathbb{N}$. Since g_{t_n} is Lipschitz continuous and integrable it must vanish at infinity, so it attains its maximum at some $x_n \in \mathbb{R}^d$ and we have

$$F(t_n) = g(t_n, x_n) - \|g_0\|_\infty - \int_0^{t_n} \Big(C(s) + \varepsilon\Big)\,\mathrm{d}s - \varepsilon' > 0, \quad n \in \mathbb{N}. \tag{2.34}$$

In particular, $g(t_n, x_n) \ge \varepsilon'$ for all $n \in \mathbb{N}$, so $g(t_1, x_n) \ge \varepsilon'/2$ for all $n \ge n_0$ for some $n_0 \in \mathbb{N}$, because $g \in C([0,T], L^\infty)$. Then by Lipschitz continuity of $g(t_1, \cdot)$ in space, $\inf_{n\in\mathbb{N}} \int_{B(x_n,1)} g(t_1, y)\,\mathrm{d}y > 0$, which would contradict integrability of $g(t_1, \cdot)$ if (x_n) were unbounded. Hence (x_n) must be bounded, without loss of generality already convergent to some $x_1 \in \mathbb{R}^d$. Then, using (2.33) and (2.34),

$$F(t_1) = \|g_{t_1}\|_\infty - \|g_0\|_\infty - \int_0^{t_1} \Big(C(s) + \varepsilon\Big)\,\mathrm{d}s - \varepsilon',$$

and

$$F(t_1) = \lim_{n\to\infty} F(t_n) = g(t_1, x_1) - \|g_0\|_\infty - \int_0^{t_1} \Big(C(s) + \varepsilon\Big)\,\mathrm{d}s - \varepsilon',$$

so $\|g_{t_1}\|_\infty = g(t_1, x_1)$. Thus by assumption $\partial_t g(t_1, x_1) \le C(t_1)$, so

$$\forall x \in \mathbb{R}^d\colon g(t_1, x) - \|g_0\|_\infty - \int_0^{t_1} \Big(C(s) + \varepsilon\Big)\,\mathrm{d}s - \varepsilon' \le F(t_1) = 0,$$

$$\partial_t \Big[g(\cdot, x_1) - \|g_0\|_\infty - \int_0^t \Big(C(s) + \varepsilon\Big)\,\mathrm{d}s - \varepsilon'\Big]_{t=t_1} \le C(t_1) - \Big(C(t_1) + \varepsilon\Big) = -\varepsilon < 0.$$

But this implies, since $\partial_t g$ is jointly continuous, that the second inequality holds in $[t_1, t_1 + \delta) \times B(x_1, \delta)$ for some $\delta > 0$, so in fact $g(t,x) < \|g_0\|_\infty + \int_0^t \Big(C(s) + \varepsilon\Big)\,\mathrm{d}s + \varepsilon'$ for all $(t,x) \in (t_1, t_1 + \delta) \times B(x_1, \delta)$, which contradicts the fact (2.34) that $g(t_n, x_n) > \|g_0\|_\infty + \int_0^t \Big(C(s) + \varepsilon\Big)\,\mathrm{d}s + \varepsilon'$ for all $n \in \mathbb{N}$. □

Proof of Lemma 2.15. Let $\rho_0 \in \mathcal{X}_+$, $T \in (0, T^\star)$, and fix $\gamma \in (2, 2 + \gamma_f)$. Then by Theorem 2.10, $\rho \in C([0,T], \mathcal{X}_+) \cap C((0,T], \mathcal{X}^\gamma)$, in particular ρ_t is Lipschitz continuous for $t \in (0,T]$. It remains to prove that $\partial_t \rho$ is jointly continuous on $[\delta, T] \times \mathbb{R}^d$ for any $\delta > 0$. By the instant regularisation proved in Theorem 2.10, $\rho_\delta \in \mathcal{X}^\gamma$, so we may as well assume that $\rho_0 \in \mathcal{X}^\gamma$ and $\delta = 0$. Then we know that ρ solves (1.4) on $[0,1]$, so

$$\partial_t \rho_t = \Delta \rho_t + \nabla \cdot (\rho_t (\nabla W \star \rho_t)) + f. \tag{2.35}$$

Using $\rho \in C([0,1], \mathcal{X}^\gamma)$ and Lemmas 2.8 and 2.12 and (2.26), it is straightforward to confirm through direct calculation that the RHS in (2.35) and therefore $\partial_t \rho \in C([0,1], \mathcal{W}^{\gamma-2,\infty})$, where $\mathcal{W}^{\gamma-2,\infty} = C^{0,\gamma-2}(\mathbb{R}^d)$ (see (III.A.1) in Section III.A) with $\gamma - 2 \in (0,1)$. In particular, $\partial_t \rho$ is continuous in time, and $(\gamma-2)$-Hölder continuous in space uniformly on $[0,T]$. This implies joint continuity □

Asymptotics of solution in absence of repulsion

We close by presenting a proof of a simple claim made in the introduction regarding the asymptotics of (1.4) with $W \equiv 0$.

Lemma 2.19. *Let $f \in \mathcal{X}_+$ not a.e. zero, and put $\rho_t \coloneqq \int_0^t G_s \star f \,\mathrm{d}s$ for $t \geq 0$.*

(i) If $d \leq 2$, then $\rho_t \uparrow \infty$ locally uniformly.

(ii) If $d \geq 3$, then $\rho_t \uparrow \rho$ where $\rho \in L^\infty$.

Proof. (i) We may assume $\int_{B(0,1)} f(x)\,\mathrm{d}x \geq 1$. Then, for any $x \in \mathbb{R}^d$,

$$\begin{aligned}\rho_t(x) = \int_0^t G_s \star f(x)\,\mathrm{d}s &\geq \int_0^t \int_{B(0,1)} f(y) G_s(x-y)\,\mathrm{d}y\,\mathrm{d}s \\ &\geq \int_0^t (2\pi s)^{-d/2} \mathrm{e}^{-(|x|+1)^2/(2s)}\,\mathrm{d}s \\ &\geq \mathrm{e}^{-(|x|+1)^2/2} \int_1^t (2\pi s)^{-d/2}\,\mathrm{d}s,\end{aligned}$$

which if $d \leq 2$ goes to ∞ locally uniformly in x as $t \to \infty$.

(ii) We have $\rho_t \uparrow \rho \coloneqq \int_0^\infty G_s \star f \,\mathrm{d}s = G \star f$ as $t \to \infty$, where

$$G(x) \coloneqq \int_0^\infty G_s(x)\,\mathrm{d}s = c(d)\,|x|^{2-d}, \quad x \in \mathbb{R}^d,$$

for some $c(d) > 0$ is the well-known Green's function of the Laplace equation in $d \geq 3$. Then G is integrable at the origin and bounded away from the origin, and f is both bounded and integrable, so $G \star f$ is bounded. □

3 Stationary Distributions

In this section we are concerned with results regarding the (non-)existence and asymptotic behaviour of stationary distributions of (1.4). The following section contains results, and Section 3.2 contains the proofs.

3.1 Results

We will make the following assumptions.

Assumption (B). *(i) f is radially symmetric.*

(ii) There exists $\alpha > 0$ such that $W'(r) \sim -r^{-\alpha}$ and $W''(r) \sim r^{-1-\alpha}$.

Here we slightly abuse notation by writing $W(x) = W(|x|)$ and $\widehat{x}\cdot\nabla W(x) = W'(|x|) = (\partial_r W)(|x|)$ for the radial part and radial derivative of W, respectively, where $\widehat{x} = x/|x|$ for $x \neq 0$. We do so similarly for other radially symmetric functions.

We proceed by making precise what we mean by a stationary distribution associated with (1.4). Suppose there is a non-negative, radially symmetric solution (ρ_t) to (1.4) started from zero – we prove later in Section 2.2.1 that there is a unique solution to (1.4) which is non-negative and radially symmetric for all times, given mild conditions on W and f – and that ρ_t converges as $t \to \infty$ to a bounded limiting shape ρ. Then ρ will be non-negative and radially symmetric, and for every $t \geq 0$ and $x_0 \in \mathbb{R}^d$, by radial symmetry of W,

$$\nabla W \star \rho_t(x_0) = \int \nabla W(x)\rho_t(x_0 - x)\,\mathrm{d}x = \frac{1}{2}\int \nabla W(x)\Big(\rho_t(x_0 - x) - \rho_t(x_0 + x)\Big)\,\mathrm{d}x. \tag{3.1}$$

This quantity may converge thanks to cancellation even if $\nabla W \star \rho$ does not exist as a Lebesgue integral. In that case, we would expect to be able to retrieve the limit of (3.1) as $t \to \infty$ from ρ as $\lim_{N\to\infty}\int_{B(0,N)} \nabla W(x)\rho(x_0 - x)$, which motivates the following definition.

Definition 3.1. For $\rho \in L^\infty(\mathbb{R}^d)$, we denote by $\nabla W \,\widetilde{\star}\, \rho$ the locally uniform limit of $\int_{B(0,N)} \nabla W(x)\rho(\,\cdot\, - x)\,\mathrm{d}x$ as $N \to \infty$, if it exists.

We remark that if $\nabla W \star \rho$ does exist as a Lebesgue integral, then $\nabla W \,\widetilde{\star}\, \rho$ exists and is equal to $\nabla W \star \rho$, see Lemma 3.12.

Now we would expect the limit ρ of ρ_t as $t \to \infty$ to be a stationary distribution in the sense of the following definition.

Definition 3.2. We say that a non-negative, bounded, radially symmetric function $\rho \in C^2(\mathbb{R}^d)$ is a *stationary distribution associated with* (1.4) if $\nabla W \,\widetilde{\star}\, \rho$ exists and is in $C^1(\mathbb{R}^d)$, and the following equation is satisfied in the classical sense:

$$0 = \frac{1}{2}\Delta\rho + \nabla\cdot(\rho\nabla W \,\widetilde{\star}\, \rho) + f. \tag{3.2}$$

For some of our results we will need to assume that a stationary solution ρ satisfies the following.

Assumption (C). *A stationary distribution ρ satisfies this assumption if there is $\beta > 0$ such that $\rho(r) \sim r^{-\beta}$.*

Remark 3.3. Many of the arguments would simplify greatly if we would make the stronger assumption that $\rho'(r) \sim -r^{-1-\beta}$, but we do not in general expect ρ_t and therefore ρ to be radially decreasing: If say the repulsion ∇W is continuous at the origin, then $\nabla W(0) = \lim_{|x|\to\infty} \nabla W(x) = 0$, so there is positive distance $d_0 \in (0, \infty)$ at which the repulsion is maximal, and we believe this can cause non-monotonic fluctuations on the length scale of d_0.

Theorem 3.4 ($d = 1$). *(i) If $\alpha > 2$, then there is no stationary solution.*

(ii) If $\alpha \in [1, 2]$, then there is no stationary solution that satisfies (C).

(iii) If $\alpha \in (0, 1)$ and ρ is a stationary solution that satisfies (C), then $\beta = \frac{1-\alpha}{2}$.

Theorem 3.5 ($d = 2$). *(i) If $\alpha > 3$, then there is no stationary solution.*

(ii) If $\alpha = 3$, then there is no stationary solution that satisfies (C).

(iii) If $\alpha \in (0, 3)$, and ρ is a stationary solution that satisfies (C), then $\beta = \frac{3-\alpha}{2}$.

Theorem 3.6 ($d \geq 3$). *If ρ is a bounded non-negative stationary solution that satisfies (C), then*

$$\beta = \begin{cases} d - 2, & \alpha \geq 3, \\ d - \frac{1+\alpha}{2}, & \alpha < 3. \end{cases}$$

Remark 3.7. (i) Theorem 3.4(i) and Theorem 3.5(i) also hold if we replace Assumption (B)(ii) by integrability of W.

(ii) In the boundary cases $\alpha = 1$ in one, and $\alpha = 3$ in two dimensions, we suspect there might be stationary distributions with $\rho(r) \sim 1$, that is $\lim_{r\to\infty} \rho(r) > 0$, which don't satisfy Assumption (C) because we require $\beta > 0$. This is why we do not make conjectures in those cases.

3.2 Proofs

Recall that in this section we work with Assumption (B). Further recall that for a radially symmetric function $g\colon \mathbb{R}^d \to \mathbb{R}$, we slightly abuse notation by writing $g(x) = g(|x|)$ and $\widehat{x} \cdot \nabla g(x) = g'(|x|) = (\partial_r g)(|x|)$ for its radial part and radial derivative, respectively. We write c for and unimportant positive constant whose value may change from line to line. If it benefits clarity, we will say explicitly what c depends on. We write $g_1(x) \lesssim g_2(x)$ if $g_1(x) \leq c g_2(x)$ for all $|x| \geq r_0$ for constants $c, r_0 > 0$. We only use this notation if the dependence of the constants r_0 and c is clear from context.

We begin by proving Theorem 3.4(i) and Theorem 3.5(i), where we claim non-existence of bounded stationary solutions without requiring Assumption (C). The idea is that $\alpha > d + 1$ implies integrability of W and ∇W, and hence that $(\nabla W) \mathbin{\widetilde{\star}} \rho = (\nabla W) \star \rho = \nabla(W \star \rho)$. Then we use radial symmetry of ρ, f, W to rewrite (3.2) as an ODE for the radial part of ρ, which can be rewritten into a fixed point equation (3.6) for ρ in terms of $\int_0^r (W \star \rho)'(s)\,\mathrm{d}s = (W \star \rho)(r) - (W \star \rho)(0)$, which leads to a contradiction of non-negativity of ρ if $d \leq 2$.

Proof of Theorem 3.4(i) and Theorem 3.5(i). It suffices to show that the assumptions $\alpha > d + 1$, $d \leq 2$, and the existence of a bounded non-negative solution ρ to (3.2) lead to a contradiction.

Firstly, $\alpha > d+1 > d$ implies integrability of ∇W and hence that $\nabla W \star \rho$ exists and is equal to $\nabla W \widetilde{\star} \rho$, see Lemma 3.12. Moreover, $\alpha > d+1$ implies that $\lim_{|x|\to\infty} W(x) =: c$ exists, and that $\widetilde{W} := W - c$ is integrable. Replacing W by $\widetilde{W}$ has no effect on (3.2), so we may assume that $W = \widetilde{W}$ is integrable. Then $W \star \rho$ exists and, recalling Lemma 3.12,

$$\nabla W \widetilde{\star} \rho = \nabla W \star \rho = \nabla (W \star \rho). \tag{3.3}$$

Then, rearranging (3.2) gives

$$-f = \nabla \cdot \left(\frac{1}{2} \nabla \rho + \rho \nabla W \star \rho \right).$$

Integrate this over $B(0,r)$ for $r > 0$ to obtain, using Gauss' integral theorem,

$$\begin{aligned} -\int_{B(0,r)} f(x)\,\mathrm{d}x &= \int_{B(0,r)} \nabla \cdot \left(\frac{1}{2} \nabla \rho + \rho \nabla W \star \rho \right) \mathrm{d}x \\ &= \int_{\partial B(0,r)} \left(\frac{1}{2} \nabla \rho + \rho \nabla W \star \rho \right) \cdot \mathrm{d}\widehat{n} \\ &= c_d r^{d-1} \left(\frac{1}{2} \rho'(r) + \rho(r)(W \star \rho)'(r) \right), \end{aligned}$$

where in the last step we recalled radial symmetry of ρ and W, and (3.3). Putting

$$g(r) := \frac{2}{c_d} \left(\int_{B(0,r)} f(x)\,\mathrm{d}x \right) r^{1-d}, \quad r > 0, \tag{3.4}$$

we obtain

$$0 = \frac{1}{2} \rho' + \rho (W \star \rho)' + g. \tag{3.5}$$

Solving this equation, we see that there exists $K > 0$ such that

$$\rho(r) = \mathrm{e}^{-2\int_0^r (W\star\rho)'(s)\,\mathrm{d}s} \left(K - \int_0^r g(s) \mathrm{e}^{2\int_0^s (W\star\rho)'(u)\,\mathrm{d}u}\,\mathrm{d}s \right). \tag{3.6}$$

Then, $\int_0^r (W \star \rho)'(s)\,\mathrm{d}s = W \star \rho(r) - W \star \rho(0)$, so

$$\rho(r) = \mathrm{e}^{-2W\star\rho(r)} \left(K - \int_0^r g(s) \mathrm{e}^{2W\star\rho(s)}\,\mathrm{d}s \right).$$

Now $\|\rho \star W\|_\infty \le \|\rho\|_\infty \|W\|_{L^1} < \infty$, and $\int_0^\infty g(s)\,\mathrm{d}s = \infty$ because $g(s) \sim s^{1-d}$ by (3.3) and integrability of f, so

$$\int_0^r g(s) \mathrm{e}^{2W\star\rho(s)}\,\mathrm{d}s \to \infty, \quad r \to \infty,$$

which contradicts non-negativity of ρ. □

The basic idea behind the proof of the remaining statements in Theorems 3.4 to 3.6 is to compare the asymptotics of the different terms in (3.2) under the assumption that $\rho(r) \sim r^{-\beta}$. For illustrative purposes assume for the moment also that $\rho'(r) \sim -r^{-1-\beta}$. Then integrating (3.2) over $x \in B(0,r)$ will yield with Gauss' theorem that, in radial coordinates,

$$1 \asymp -r^{d-1}\left(\rho'(r) + \rho(r)(\nabla W \mathbin{\widetilde{\star}} \rho)(r)\right) \implies r^{1-d} \asymp r^{-1-\beta} - r^{-\beta}(\nabla W \mathbin{\widetilde{\star}} \rho)(r). \tag{3.7}$$

See (3.11) in the proof below for a derivation of the first equation. Here $\asymp$ denotes asymptotic equality, a precise definition is below. Then if we know the asymptotics of $\nabla W \mathbin{\widetilde{\star}} \rho$ in terms of α and β, we can write consistency equations for α, β, d by requiring that the largest positive and negative terms on the right hand side of (3.7) be of the same order of magnitude as $r \to \infty$. This lets us express β in terms of d and α, and arrive at contradictions if this would necessitate $\beta \leq 0$.

Since we do not actually make assumptions on ρ' but only require that $\rho(r) \sim r^{-\beta}$ for some $\beta > 0$, we will integrate the left hand equation in (3.7) again over $r \in (0,R)$ to replace all appearances of ρ' by ρ, and it will suffice to find the asymptotics of

$$F(R) \coloneqq -\int_{B(0,R)} \widehat{x} \cdot \nabla W \mathbin{\widetilde{\star}} \rho(x)\,\mathrm{d}x, \quad R > 0. \tag{3.8}$$

We write $f(R) \asymp g(R)$ for non-negative f, g to mean that there are constants $C, R_0 > 0$ (which may depend on α, β, d) such that $\frac{1}{C} f(R) \leq g(R) \leq C f(R)$ for $R \geq R_0$. If we write say $f(R) \asymp h(R) \pm g(R)$ we mean that $\frac{1}{C}(h(R) - g(R)) \leq f(R) \leq C(h(R) + g(R))$. For a statement S and $r > 0$, we write

$$(\log r)_S \coloneqq \begin{cases} \log r, & S \text{ is true,} \\ 1, & \text{else.} \end{cases}$$

Lemma 3.8. *Suppose $\rho \in L^1 \cap L^\infty$ is non-negative, radially symmetric, and there is $\beta > 0$ such that $\rho(r) \sim r^{-\beta}$. If $\alpha + \beta \leq d - 1$, then $F(R) = \infty$ for large R, otherwise*

$$\begin{aligned} F(R) \asymp \pm\big(1 &+ R^{d-1-\beta}(\log R)_{\beta=d-1}(\log R)_{\alpha=d+1}) \\ &+ R^{d-\alpha}(\log R)_{\alpha=d} \\ &+ R^{2d-\alpha-\beta}(\log R)_{\alpha+\beta=2d}(\log R)_{\beta=d}(\log R)_{\alpha=d+1}. \end{aligned} \tag{3.9}$$

Remark 3.9. Under the stronger assumption that $\rho'(r) \sim -r^{-1-\beta}$, one can much more easily handle cancellations of the form $\rho(x + x_0) - \rho(x - x_0)$ as they appear in (3.1), and show that

$$-\widehat{x} \cdot \nabla W \mathbin{\widetilde{\star}} \rho(x) \asymp \pm R^{-1-\beta} + R^{-\alpha} + R^{d-\alpha-\beta}(\log R)_{\beta=d}(\log R)_{\alpha=d+1},$$

which immediately implies

$$\begin{aligned} F(R) \asymp \pm(1 &+ R^{d-1-\beta}(\log R)_{\beta=d-1}) \\ &+ R^{d-\alpha}(\log R)_{\alpha=d} \\ &+ R^{2d-\alpha-\beta}(\log R)_{\alpha+\beta=2d}(\log R)_{\beta=d}(\log R)_{\alpha=d+1}. \end{aligned} \tag{3.10}$$

This is the same as (3.9) with the exception of the $(\log R)_{\alpha=d+1}$ term at $R^{d-1-\beta}$. We believe that a refinement of the following proof yields (3.10) under the assumptions of Lemma 3.8.

Before proving Lemma 3.8, we show how it implies the remaining statements in Theorems 3.4 to 3.6 using the idea explained above. We write $o_x(1)$ to mean a function $g = g(x)$ that tends to zero as $x \to \infty$, and may change from line to line.

Proof of Theorem 3.4(ii) and (iii), Theorem 3.5 (ii) and (iii), Theorem 3.6. Suppose that ρ is a bounded non-negative stationary solution, see Definition 3.2, that satisfies Assumption (C). Then, for $r > 0$,

$$\begin{aligned} 1 \asymp \int_{B(0,r)} f(x)\,\mathrm{d}x &= -\int_{B(0,r)} \nabla \cdot \Big(\nabla\rho + \rho\nabla W \mathbin{\widetilde{\star}} \rho\Big)\,\mathrm{d}x \\ &= -\int_{\partial B(0,r)} \Big(\nabla\rho + \rho\nabla W \mathbin{\widetilde{\star}} \rho\Big) \cdot \mathrm{d}\widehat{n} \\ &\asymp -r^{d-1}\Big(\rho'(r) + \rho(r)\widehat{r} \cdot \nabla W \mathbin{\widetilde{\star}} \rho(r)\Big). \end{aligned} \tag{3.11}$$

Thus, if $\rho(r) \neq 0$,

$$\frac{1}{\rho(r)} \asymp -r^{d-1}(\log\rho)' - r^{d-1}\widehat{r} \cdot \nabla W \mathbin{\widetilde{\star}} \rho(r). \tag{3.12}$$

Recall $\rho(r) \sim r^{-\beta}$, say $\rho(r) = (1 + o_r(1))ar^{-\beta}$. Then for $r_0 > 0$, if $d \neq 1$,

$$\begin{aligned} \int_{r_0}^{R} r^{d-1}(\log\rho)'\,\mathrm{d}r &= r^{d-1}\log\rho(r)\Big|_{r_0}^{R} - \int_{r_0}^{R}(d-1)r^{d-2}\log\rho(r)\,\mathrm{d}r \\ &= -\beta\left[r^{d-1}\log(cr)\Big|_{r_0}^{R} - \int_{r_0}^{R}(d-1)r^{d-2}\log(cr)\,\mathrm{d}r\right] + o_{r_0}(1)R^{d-1} \\ &= -\frac{\beta}{d-1}R^{d-1} + o_{r_0}(1)R^{d-1}. \end{aligned}$$

If $d = 1$, then $\int_{r_0}^{R}(\log\rho)'\,\mathrm{d}r = \log\rho(R) - \log\rho(r_0) \asymp -\log R$, hence we can choose $r_0 > 0$ so large that $\int_{r_0}^{R} r^{d-1}(\log\rho)'\,\mathrm{d}r \asymp -R^{d-1}(\log R)_{d=1}$. Furthermore, $\int_{r_0}^{R}\frac{1}{\rho(r)}\,\mathrm{d}r \asymp R^{1+\beta}$, so by integrating (3.12) from r_0 to R,

$$R^{1+\beta} \asymp R^{d-1}(\log R)_{d=1} + F(R)$$

Then Lemma 3.8 necessitates $\alpha + \beta > d - 1$, and we claim it further implies that

$$\beta = \begin{cases} d - \frac{1+\alpha}{2}, & \alpha < 3 \wedge (d+1), \\ d-2, & \text{else.} \end{cases} \tag{3.13}$$

Indeed, if $\alpha \geq d+1$, then $|F(R)| \lesssim 1 + R^{d-1-\beta}(\log R)^3$, so $\beta = d-2$ (and we have a contradiction in $d=1$). If $\alpha < d+1$ and $\beta > d$, then $F(R) \asymp \pm 1 + R^{d-\alpha}(\log R)_{\alpha=d}$ so $R^{1+\beta} \asymp R^{d-1}(\log R)_{d=1} + R^{d-\alpha}(\log R)_{\alpha=d}$, so $\beta = d-2$ or $\beta = d-1-\alpha$, both contradict $\beta > d$. A similar argument excludes the possibility that $\beta = d$, so $\beta < d$ and $F(R) \asymp \pm 1 + R^{2d-\alpha-\beta}(\log R)_{\alpha+\beta=2d}$, so

$$R^{1+\beta} \asymp R^{d-1}(\log R)_{d=1} + R^{2d-\alpha-\beta}(\log R)_{\alpha+\beta=2d}.$$

Then if $2d-\alpha-\beta > d-1$, i.e. $\alpha+\beta < d+1$ (in particular $\alpha+\beta \neq 2d$), we get $1+\beta = 2d-\alpha-\beta$, i.e. $\beta = d - \frac{1+\alpha}{2}$, which satisfies $\alpha+\beta < d+1$ iff $\alpha < 3$. Otherwise, we get $\beta = d-2$, which satisfies $\alpha+\beta \geq d+1$ iff $\alpha \geq 3$. This proves (3.13).

(i) If $d=1$, then $\beta = d-2 = -1 < 0$ cannot be true, so $\alpha < 3 \wedge (d+1) = 2$ and $\beta = 1 - \frac{1+\alpha}{2} = \frac{1-\alpha}{2}$. This contradicts $\beta > 0$ if $\alpha \geq 1$.

(ii) If $d=2$, then $\beta = d-2 = 0$ cannot be true, so $\alpha < 3 \wedge (d+1) = 3$ and $\beta = 3 - \frac{1+\alpha}{2} = \frac{3-\alpha}{2}$.

(iii) If $d \geq 3$, then $3 \wedge (d+1) = 3$ so

$$\beta = \begin{cases} d - \frac{1+\alpha}{2}, & \alpha < 3, \\ d-2, & \alpha \geq 3. \end{cases}$$

□

3.2.1 Proof of Lemma 3.8

This section is dedicated to the (long and technical) proof of Lemma 3.8.

Proof of Lemma 3.8. We assume that $d \geq 2$; the arguments simplify considerably if $d=1$. Recall that for a radially symmetric function such as W we write $W'(|x|) = W'(x) = \widehat{x} \cdot \nabla W(x)$ for $x \in \mathbb{R}^d$. By replacing W with $-W$, we can drop the minus sign in (3.8) and assume $W'(r) \sim +r^{-\alpha}$. Then,

$$\begin{aligned}
F(R) &= \int_{B(0,R)} \widehat{x}_0 \cdot \nabla W \mathbin{\widetilde{\star}} \rho(x_0)\,\mathrm{d}x_0 \\
&= \int_{B(0,R)} \mathrm{d}x_0 \lim_{N\to\infty} \int_{B(0,N)} \mathrm{d}x\, \mathbb{1}_{\{x\cdot x_0 > 0\}}\, \widehat{x}_0 \cdot \nabla W(x) \Big(\rho(x_0 - x) - \rho(x_0 + x) \Big) \\
&= \lim_{N\to\infty} \int_{B(0,N)} \mathrm{d}x\, W'(x) \int_{B(0,R)\cap\{x\cdot x_0 > 0\}} \cos(x, x_0) \Big(\rho(x_0 - x) - \rho(x_0 + x) \Big)\,\mathrm{d}x_0 \\
&= c_d \lim_{N\to\infty} \int_0^N \mathrm{d}r\, r^{d-1} W'(r) \int_{B(0,R)\cap\{e_1\cdot x_0 > 0\}} \cos(e_1, x_0) \Big(\rho(x_0 - re_1) - \rho(x_0 + re_1) \Big)\,\mathrm{d}x_0 \\
&= c_d \lim_{N\to\infty} \int_0^N \mathrm{d}r\, r^{d-1} W'(r) \left(S^-(r,R) - S^+(r,R) \right), \qquad (3.14)
\end{aligned}$$

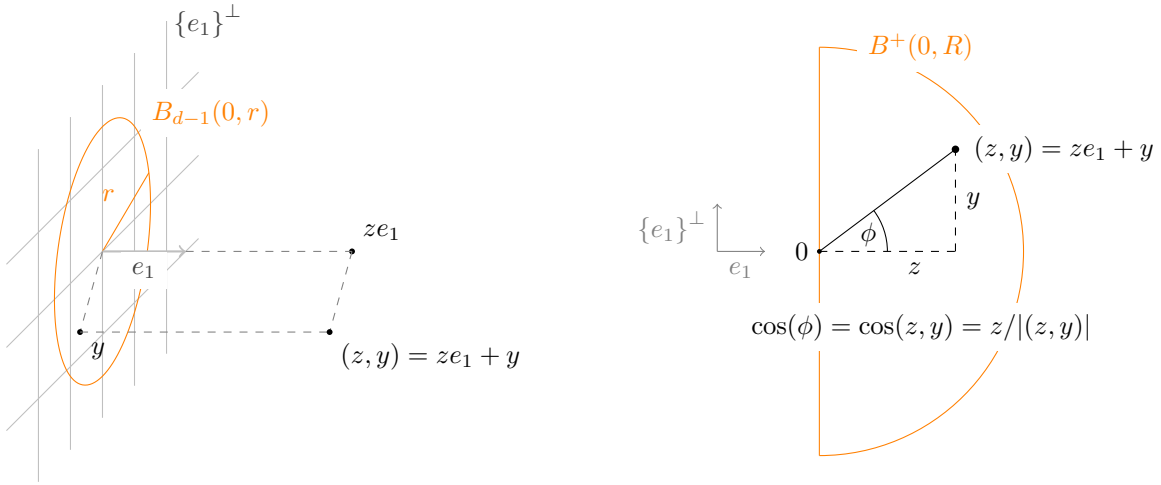

Figure 3.1: Illustration for the notations $B_{d-1}(0,R)$ for $R > 0$, as well as $(z,y) = ze_1 + y$ and $\cos(z,y)$ for $z > 0$ and $y \in \{e_1\}^\perp$.

where we could apply the dominated convergence theorem in the third line because the convergence in Definition 3.1 is uniform on $B(0,R)$, and used Fubini's theorem in the same step. Here $\cos(x,y) = \widehat{x} \cdot \widehat{y}$ for $x, y \in \mathbb{R}^d$, and $e_1 \in \mathbb{R}^d$ denotes the first unit vector, and

$$\begin{aligned} S^+(r,R) &= \int_{B(0,R)\cap\{e_1\cdot x_0>0\}} \cos(e_1,x_0)\rho(re_1+x_0)\,\mathrm{d}x_0 \\ &= \int_0^R \mathrm{d}z \int_{B_{d-1}(0,\sqrt{R^2-z^2})} \mathrm{d}y\, \cos(e_1,ze_1+y)\rho((r+z)e_1+y) \\ S^-(r,R) &= \int_{B(0,R)\cap\{e_1\cdot x_0>0\}} \cos(e_1,x_0)\rho(re_1-x_0)\,\mathrm{d}x_0 \\ &= \int_0^R \mathrm{d}z \int_{B_{d-1}(0,\sqrt{R^2-z^2})} \mathrm{d}y\, \cos(e_1,ze_1+y)\rho((r-z)e_1+y). \end{aligned} \tag{3.15}$$

Here $B_{d-1}(0,r)$ denotes the ball in $\mathbb{R}^{d-1}$ centred at the origin with radius r, which we identify with $\{x \in \mathbb{R}^d \colon |x| < r,\, x \cdot e_1 = 0\} = B(0,R) \cap \{e_1\}^\perp \subset \mathbb{R}^d$. If $z > 0$ and $y \in \{e_1\}^\perp$, we write $(z,y) \coloneqq ze_1 + y \in \mathbb{R}^d$, and abbreviate $B^+(0,R) \coloneqq \{x \in B(0,R) \colon x \cdot e_1 > 0\}$, and $\int_{B^+(0,R)} \mathrm{d}(z,y) \coloneqq \int_0^R \mathrm{d}z \int_{B_{d-1}(0,\sqrt{R^2-z^2})} \mathrm{d}y$, and $\cos(z,y) \coloneqq \cos(e_1, ze_1+y) = z/|(z,y)|$. This notation is illustrated in Fig. 3.1, and an illustration for the representation of $F(R)$ derived in (3.14) is in Fig. 3.2. Then,

$$F(R) = c_d \lim_{N\to\infty} \int_{B^+(0,R)} \mathrm{d}(z,y)\cos(z,y) \int_0^N \mathrm{d}r\, r^{d-1} W'(r) \Big(\rho((r-z)e_1+y) - \rho((r+z)e_1+y)\Big). \tag{3.16}$$

By multiplying W and ρ by constants, we may assume that $r^\alpha W'(r) \to 1$, $r^{1+\alpha}W''(r) \to -\alpha$, and $r^\beta \rho(r) \to 1$ as $r \to \infty$. We introduce the following constants.

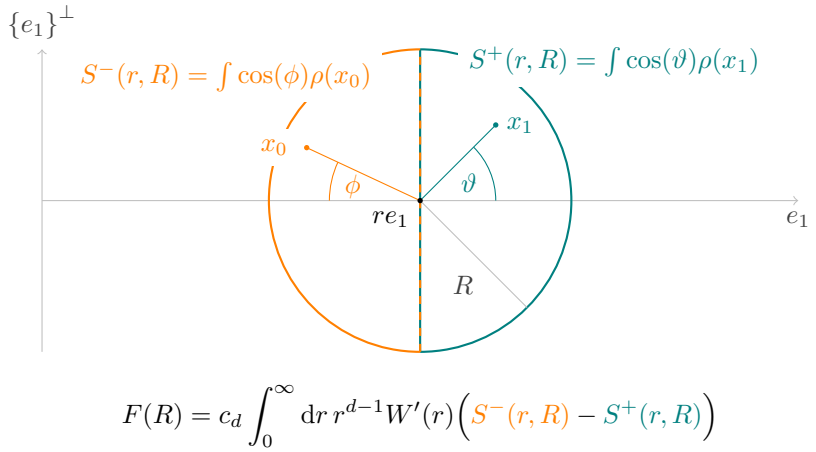

Figure 3.2: Illustration for (3.14).

(1) Choose $r_0 > 0$ such that for all $x \in \mathbb{R}^d$ with $|x| \geq r_0/2$,

$$\begin{aligned} \frac{1}{2}|x|^{-\beta} &\leq \quad \rho(x) \quad &&\leq 2|x|^{-\beta}, \\ \frac{1}{2}|x|^{-\alpha} &\leq \quad W'(x) \quad &&\leq 2|x|^{-\alpha}, \\ \frac{1}{2}|x|^{-1-\alpha} &\leq -\frac{1}{\alpha} W''(x) &&\leq 2|x|^{-1-\alpha}. \end{aligned}$$

(2) Choose $K > 1$ large enough that

$$a(K) := \frac{\alpha + \beta + 1 - d}{\alpha + 1 - d} \left(1 + \frac{2}{K}\right)^{d-1-\alpha} \left(\frac{K^2}{(K+1)^2 + 1}\right)^{\beta/2} > 1.$$

(3) Choose $\varepsilon = \varepsilon(K) \in (0,1)$ as explained in (ii.3) later in the proof.

(4) Choose $\eta = \eta(\varepsilon, K) \in (0,1)$ as explained in (iv) later in the proof.

We now explain the high level idea of the proof. The challenge is that one has to be very precise when weighing the positive and negative contributions to $F(R)$ in (3.14) (see also Fig. 3.2), because they will have the same order of magnitude, and the fact that ρ may not be decreasing makes direct comparisons difficult. We will split the integral over $r \in (0, \infty)$ in (3.16) into four regions.

(i) If $K > 0$ is sufficiently large, then for $r \in (KR, \infty)$, naive lower and upper bounds respectively on the positive and negative contributions will suffice to show that the positive contribution dominates.

(ii) In the region $r \in (r_0, \varepsilon R)$, we can use radial symmetry of ρ to rewrite $S^-(r,R)$ in a way that the domain of integration overlaps with that of $S^+(r,R)$ (see Fig. 3.3), and such that in the overlapping region, the positive contribution dominates (see Fig. 3.4). Then the only remaining negative contribution can be shown to be of the

same order in R as (i), and proportional to ε, so choosing $\varepsilon = \varepsilon(K)$ small enough lets us absorb it in (i).

(iii) The region $r \in (0, r_0)$ can be treated similarly to (ii), but because we make no assumptions on ∇W close to the origin, the sign of these contributions cannot be controlled (hence the contributions to (3.9) with indeterminate sign).

(iv) As mentioned previously, ρ may not be decreasing, but $\rho(x) \sim |x|^{-\beta}$ implies that ρ is decreasing w.r.t multiplicative increments, in the sense that $\rho(ax) \leq (1 + o_{|x|}(1))a^{-\beta}\rho(x)$ for $a > 1$. Then we can find for any $\eta > 0$ a number $\gamma(\varepsilon, K, \eta) > 1$ such that $\rho(x) \geq \gamma\rho(y)$ for all $|x|, |y| \in (\varepsilon R, KR)$ with say $|x - y| \geq \eta R$. Hence when we exclude in the integrals defining $S^{\pm}(r, R)$ points that are within distance ηR from each other, then the positive contribution is at least γ times the negative (see Fig. 3.5 for a visualisation). The contributions we excluded can be absorbed by (i) if we choose $\eta = \eta(\varepsilon, K) > 0$ small enough.

We recall and introduce some notation for the proof. We write $o_x(1)$ to mean a function $g = g(x)$ that tends to zero as $x \to \infty$ that may depend on $\alpha, \beta, d, r_0, K, \varepsilon, \eta$, but nothing else, and may change from line to line. We use c for positive constants that can change from line to line, and we denote dependencies on K, ε, η, but not α, β, d, r_0. All inequalities we write will assume that R is sufficiently large, for example we will assume $\varepsilon R > r_0$ without comment; if it matters, we explicitly denote dependence of R_0 on any variables other than α, β, d, r_0.

As explained, we proceed by splitting the integral over $r \in (0, \infty)$ in (3.16). Denote for $J \subset (0, \infty)$,

$$F_J(R) \coloneqq c_d \lim_{N\to\infty} \int_{B^+(0,R)} \mathrm{d}(z, y) \cos(z, y) \int_{J\cap(0,N)} r^{d-1} W'(r) \Big(\rho((r-z)e_1+y) + \rho((r+z)e_1+y) \Big). \tag{3.17}$$

(i) **The contribution** (KR, ∞)

For fixed $(z, y) \in B^+(0, R)$ and $N > 2KR$ say,

$$\begin{aligned}
&\int_{KR}^{N} r^{d-1} W'(r) \Big(\rho((r - z)e_1 + y) - \rho((r + z)e_1 + y) \Big) \, \mathrm{d}r \\
&\qquad = \int_{KR-z}^{KR+z} (r + z)^{d-1} W'(r + z) \rho(re_1 + y) \, \mathrm{d}r && \text{(i.1)} \\
&\qquad\quad + \int_{KR+z}^{N-z} \rho(re_1 + y) \Big[(r + z)^{d-1} W'(r + z) - (r - z)^{d-1} W'(r - z) \Big] \, \mathrm{d}r && \text{(i.2)} \\
&\qquad\quad - \int_{N-z}^{N+z} \rho(re_1 + y)(r - z)^{d-1} W'(r - z) \, \mathrm{d}r. && \text{(i.3)}
\end{aligned}$$

Note that (i.1) is always positive. In fact, if $(z,y) \in B^+(0,R)$, $r \in (KR-z, KR+z)$, then $r+z \in (KR,(K+2)R)$ and $|re_1+y| \asymp R$, so

$$\text{(i.1)} \asymp R^{d-1-\alpha-\beta} \int_{KR-z}^{KR+z} \mathrm{d}r = 2zR^{d-1-\alpha-\beta}, \tag{3.18}$$

thus

$$\int_{B^+(0,R)} \cos(z,y)\text{(i.1)}\,\mathrm{d}(z,y) \asymp R^{d-1-\alpha-\beta} \int_{B^+(0,R)} z\cos(z,y)\,\mathrm{d}(z,y) \asymp R^{2d-\alpha-\beta}. \tag{3.19}$$

In contrast, the sign of (i.2) depends on whether $w(r) \coloneqq r^{d-1}W'(r)$, $r > 0$, is increasing or decreasing. This depends on α:

$$\begin{aligned} w'(r) = (d-1)r^{d-2}W'(r) + r^{d-1}W''(r) &= (1+o_r(1))(d-1)r^{d-2-\alpha} - (1+o_r(1))\alpha r^{d-2-\alpha} \\ &= \begin{cases} (1+o_r(1))(d-1-\alpha)r^{d-2-\alpha}, & \alpha \neq d-1, \\ o_r(1)r^{-1}, & \alpha = d-1. \end{cases} \end{aligned} \tag{3.20}$$

Finally, if $r \in (N-z, N+z) \subset (N-R, N+R)$, then $\rho(re_1+y) = (1+o_{N/R}(1))N^{-\beta}$, and $W'(r-z) = (1+o_{N/R}(1))N^{-\alpha}$, so

$$|\text{(i.3)}| = (1+o_{N/R}(1))(2z)N^{d-1-\alpha-\beta}. \tag{3.21}$$

- **Case $\alpha > d-1$:**

  In this case, $\lim_{N\to\infty}\text{(i.3)} = 0$ by (3.21), and w is eventually decreasing by (3.20), so (i.2) is negative. Then for $r > r_0$, $z > 0$,

  $$0 \le w(r) - w(r+z) = -\int_r^{r+z} w'(s)\,\mathrm{d}s \le z|w'(r)| = (1+o_r(1))z(\alpha-d+1)r^{d-2-\alpha}.$$

  Hence,

  $$\begin{aligned} |\text{(i.2)}| &\le \int_{KR+z}^{\infty} \rho(re_1+y)|w(r+z)-w(r-z)|\,\mathrm{d}r \\ &\le (1+o_R(1))(2z)(\alpha-d+1)\int_{KR+z}^{\infty} \rho(re_1+y)(r-z)^{d-2-\alpha}\,\mathrm{d}r. \end{aligned}$$

  Now, for $(z,y) \in B^+(0,R)$ and $r \in (KR+z,\infty)$ we have $|re_1+y| = \sqrt{r^2+|y|^2} \ge r \ge r-z$, so

  $$\begin{aligned} |\text{(i.2)}| &\le (1+o_R(1))(2z)(\alpha-d+1)\int_{KR}^{\infty} r^{d-2-\alpha-\beta}\,\mathrm{d}r \\ &= (1+o_R(1))(2z)\frac{\alpha+1-d}{\alpha+\beta+1-d}(KR)^{d-1-\alpha-\beta}. \end{aligned}$$

To compare with (i.1), we produce a more precise lower bound corresponding to the asymptotics derived in (3.18). If $(z,y) \in B^+(0,R)$, and $r \in (KR,\infty)$, then $r^2+|y|^2 \le ((K+1)^2+1)R^2$, and $r+z \le (K+2)R$, so (and because $s \mapsto s^{d-1-\alpha}$ is decreasing),

$$\text{(i.1)} \ge (1-o_R(1))(2z)R^{d-1-\alpha-\beta}(K+2)^{d-1-\alpha}((K+1)^2+1)^{-\beta/2}.$$

By our choice of K, this implies that $|\text{(i.2)}| \le (1+o_R(1))a(K)\text{(i.1)}$. Hence, and by (3.19),

$$F_{(KR,\infty)}(R) = c_d \int_{B^+(0,R)} \cos(z,y)\Big[\text{(i.1)}+\text{(i.2)}\Big]\,\mathrm{d}(z,y)$$
$$\asymp \int_{B^+(0,R)} \cos(z,y)\text{(i.1)}\,\mathrm{d}(z,y) \asymp R^{2d-\alpha-\beta}. \quad (3.22)$$

- **Case $\alpha < d-1$:**

  In this case, w is eventually increasing by (3.20), so (i.2) is positive. Then for $z>0$ and $r>r_0$,

  $$w(r+z)-w(r) = \int_r^{r+z} w'(s)\,\mathrm{d}s = (1+o_r(1))z(d-1-\alpha)r^{d-2-\alpha}.$$

  For fixed $(z,y) \in B^+(0,R)$,

  $$\text{(i.2)} = \int_{KR+z}^{N} \rho(re_1+y)(f(r+z)-f(r-z))\,\mathrm{d}r$$
  $$= (1+o_R(1))(2z)(d-1-\alpha)\int_{KR}^{N} r^{d-2-\alpha-\beta}\,\mathrm{d}r$$

  If $\alpha+\beta > d-1$, then $\text{(i.2)} \asymp zR^{d-1-\alpha-\beta}$, and $\lim_{N\to\infty}\text{(i.3)} = 0$ by (3.21), so $F_{(KR,\infty)}(R) \asymp R^{2d-\alpha-\beta}$. If $\alpha+\beta < d-1$, then

  $$\text{(i.2)} = (1+o_R(1)+o_{N/R}(1))(2z)\underbrace{\frac{d-1-\alpha}{d-1-\alpha-\beta}}_{>1} N^{d-1-\alpha-\beta}.$$

  Comparing with (3.21), we see that for sufficiently large R we have

  $$\text{(i.2)}+\text{(i.3)} \ge (2z)(1+o_R(1)+o_{N/R}(1))\frac{\beta}{d-1-\alpha-\beta}N^{d-1-\alpha-\beta} \to \infty, \quad N\to\infty.$$

  Hence in this case $F_{(KR,\infty)}(R) = \infty$ for large R. If $\alpha+\beta = d-1$, then $|\text{(i.3)}|$ is bounded as $N\to\infty$ for R fixed, and $\text{(i.2)} \asymp \log N \to \infty$ as $N\to\infty$, so $F_{(KR,\infty)}(R) = \infty$ for large R in this case as well.

- **Case $\alpha = d-1$:**

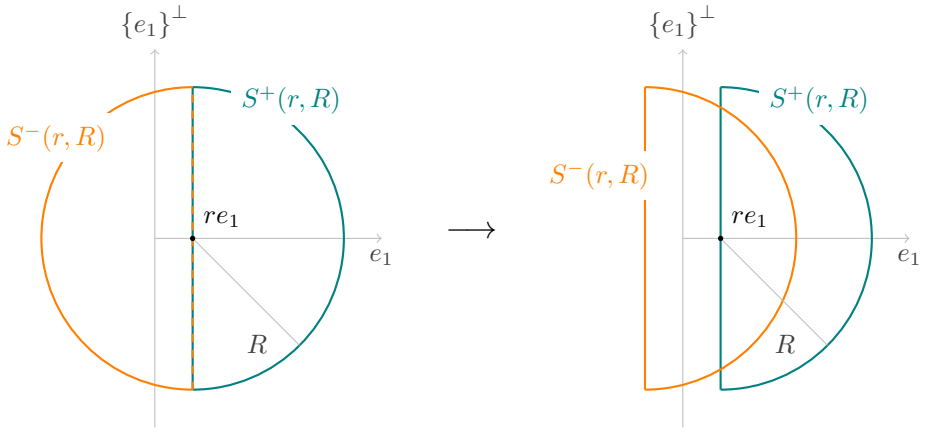

Figure 3.3: Illustration for the first step in (3.24) using radial symmetry of ρ.

In this case, $\lim_{N\to\infty}$(i.3) $= 0$ by (3.21) and because $\beta > 0$. Furthermore $w(r) \to 1$ as $r \to \infty$, and for $r > r_0, z > 0$, by (3.20),

$$|w(r+z) - w(r)| = \left| \int_r^{r+z} w'(s)\,\mathrm{d}s \right| \le o_r(1)\frac{z}{r}.$$

Hence, for fixed $(z, y) \in B^+(0, R)$,

$$|\text{(i.2)}| \le o_R(1) \int_{KR+z}^{\infty} r^{-\beta}\frac{z}{r}\,\mathrm{d}r = o_R(1) z R^{-\beta} = o_R(1) z R^{d-1-\alpha-\beta}.$$

Hence $|$(i.2)$| \le o_R(1)$(i.1) by (3.18), so (3.22) holds in this case as well.

Note that $\alpha < d-1$ is necessary for $\alpha + \beta \le d-1$ to hold, so we can summarise

$$F_{(KR,\infty)}(R) \begin{cases} = \infty, & \alpha+\beta \le d-1, \\ \asymp R^{2d-\alpha-\beta}, & \text{else.} \end{cases} \tag{3.23}$$

(ii) **The contribution** $(r_0, \varepsilon R)$

Note that for fixed $r \in (r_0, \varepsilon R)$, and every $z > 0$, $y \in B_{d-1}(0, R)$, we have $\rho((r - z)e_1 + y)$ by radial symmetry of ρ, and thus

$$\begin{aligned} &S^-(r,R) - S^+(r,R) && (3.24)\\ &\quad= \int_{B_{d-1}(0,R)} \mathrm{d}y \int_0^{R'(y)} \mathrm{d}z\, \cos(z,y)\Big(\rho((-r+z)e_1+y) - \rho((r+z)e_1+y)\Big) \\ &\quad= \int_{B_{d-1}(0,R)} \mathrm{d}y \int_0^{2r\wedge R'(y)} \mathrm{d}z\, \cos(z,y)\rho((-r+z)e_1+y) && \text{(ii.1)}\\ &\qquad+ \int_{B_{d-1}(0,R)} \mathrm{d}y \int_0^{(R'(y)-2r)_+} \mathrm{d}z\, \rho((r+z)e_1+y)(\cos(z+2r,y) - \cos(z,y)) \\ &\qquad && \text{(ii.2)}\\ &\qquad- \int_{B_{d-1}(0,R)} \mathrm{d}y \int_{(R'(y)-2r)_+}^{R'(y)} \mathrm{d}z\, \rho((r+z)e_1+y)\cos(z,y), && \text{(ii.3)} \end{aligned}$$

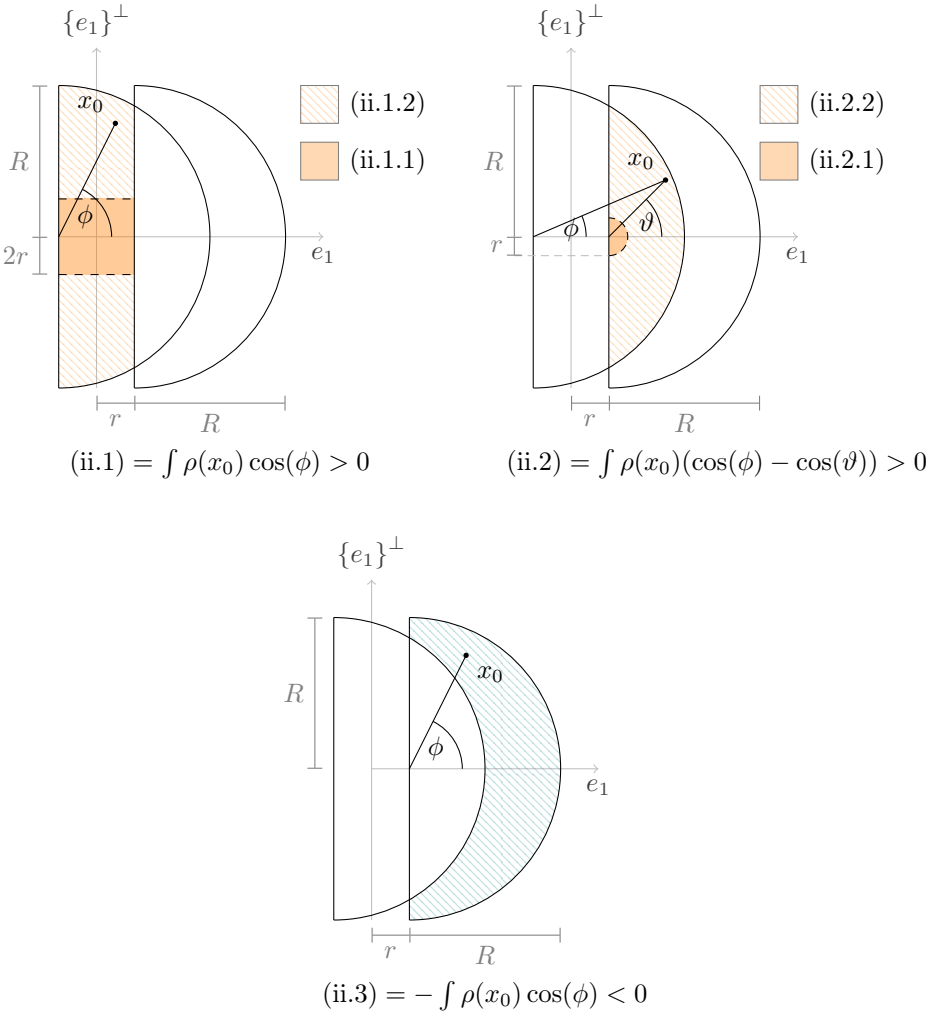

(ii.1) $= \int \rho(x_0)\cos(\phi) > 0$

(ii.2) $= \int \rho(x_0)(\cos(\phi) - \cos(\vartheta)) > 0$

(ii.3) $= -\int \rho(x_0)\cos(\phi) < 0$

Figure 3.4: An illustration of how the integral (ii) in (3.24) is being split into (ii.1), (ii.2), and (ii.3).

where we put $R'(y) \coloneqq \sqrt{R^2 - |y|^2}$ for $y \in B_{d-1}(0, R)$. An illustration of the first step using the radial symmetry of ρ is in Fig. 3.3. We will address each of the three integrals (ii.1)-(ii.3) in turn. Supporting illustrations are in Fig. 3.4.

(ii.1) We split the integral over y into $B_{d-1}(0, 2r)$, denoted by (ii.1.1), and $B_{d-1}(0, R) \cap \{|y| > 2r\}$, denoted (ii.1.2), see also Fig. 3.4. For the former, note that if $|y| < 2r$ then $R'(y) \geq 2r$ as long as $\varepsilon \leq 1/\sqrt{8}$, so

$$\text{(ii.1.1)} = \int_{B_{d-1}(0,2r)} \mathrm{d}y \int_0^{2r} \mathrm{d}z \cos(z, y)\rho((-r + z)e_1 + y).$$

The integrand is non-negative, so we obtain a lower bound by restricting the integration to $(z, y) = ze_1 + y \in re_1 + B(0, r/2)$. In that case,

$$\cos(z, y) = \frac{z}{\sqrt{z^2 + |y|^2}} \geq \frac{r/2}{\sqrt{(r/2)^2 + (r/2)^2}} = \frac{1}{\sqrt{2}},$$

so

$$\begin{aligned}(\text{ii.1.1}) \geq \frac{1}{\sqrt{2}} \int_{B(0,r/2)} \rho(x)\,\mathrm{d}x &\asymp \int_{B(0,r_0/2)} \rho(x)\,\mathrm{d}x + \int_{r_0/2}^{r/2} s^{d-1}\rho(s)\,\mathrm{d}s \\ &\asymp 1 + \int_{r_0/2}^{r/2} s^{d-1-\beta}\,\mathrm{d}s \\ &\asymp 1 + r^{d-\beta}(\log r)_{\beta=d}.\end{aligned}$$

A matching upper bound can be obtained using $\cos(e_1, ze_1 + y) \leq 1$. Then, if $\beta \neq d$,

$$\begin{aligned}\int_{r_0}^{\varepsilon R} r^{d-1}W'(r)(1 + r^{d-\beta})\,\mathrm{d}r &\asymp \int_{r_0}^{\varepsilon R} r^{d-1-\alpha}(1 + r^{d-\beta})\,\mathrm{d}r \\ &\asymp 1 + R^{d-\alpha}(\log R)_{\alpha=d} + R^{2d-\alpha-\beta}(\log R)_{\alpha+\beta=2d}.\end{aligned}$$

If $\beta = d$, then by Lemma 3.11

$$\int_{r_0}^{\varepsilon R} r^{d-1}W'(r)\log r\,\mathrm{d}r \asymp \int_{r_0}^{\varepsilon R} r^{d-1-\alpha}\log r\,\mathrm{d}r \asymp 1 + R^{d-\alpha}\log R(\log R)_{\alpha=d}.$$

We summarise

$$\int_{r_0}^{\varepsilon R} r^{d-1}W'(r)(\text{ii.1.1})\,\mathrm{d}r \asymp 1 + R^{d-\alpha}(\log R)_{\beta=d}(\log R)_{\alpha=d} + R^{2d-\alpha-\beta}(\log R)_{\alpha+\beta=2d}. \tag{3.25}$$

We turn to (ii.1.2). If $y \in B_{d-1}(0,R)$ with $|y| > 2r$ and $z \in (0, 2r)$, then $z < 2r < |y|$ so $\frac{z}{\sqrt{2}|y|} \leq \frac{z}{\sqrt{z^2+y^2}} \leq \frac{z}{|y|}$, and

$$\rho((-r+z)e_1 + y) \geq \frac{1}{2}((-r+z)^2 + y^2)^{-\beta/2} \geq 2^{-1-\beta/2}\,|y|^{-\beta}\,, \tag{3.26}$$

and similarly $\rho((-r+z)e_1 + y) \leq 2|y|^{-\beta}$. Further note that $R'(y) \geq r$ for all $|y| \leq R/2$ given that $\varepsilon \leq 1/\sqrt{2}$, so

$$\int\limits_{A_{d-1}(r,R/2)} \mathrm{d}y \int_0^r \mathrm{d}z \leq \int\limits_{A_{d-1}(r,R)} \mathrm{d}y \int_0^{2r\wedge R'(y)} \mathrm{d}z \leq \int\limits_{A_{d-1}(r,R)} \mathrm{d}y \int_0^{2r} \mathrm{d}z. \tag{3.27}$$

Here $A_{d-1}(a,b) = \{y \in B_{d-1}(0,b)\colon |y| > a\} = \big\{y \in \mathbb{R}^d\colon a < |y| < b,\, y \cdot e_1 = 0\big\}$. Hence,

$$\begin{aligned}(\text{ii.1.2}) &= \int\limits_{A_{d-1}(r,R)} \mathrm{d}y \int_0^{2r\wedge R'(y)} \mathrm{d}z\,\frac{z}{\sqrt{z^2+y^2}}\rho((-r+z)e_1 + y) \\ &\leq 2\int_{A_{d-1}(r,R)} \mathrm{d}y \int_0^{2r} \mathrm{d}z \frac{z}{|y|}|y|^{-\beta} \\ &\asymp \int_r^R \mathrm{d}s\, s^{d-2} \int_0^{2r} \mathrm{d}z\, z s^{-1-\beta} \\ &\asymp r^2(r^{d-2-\beta} + R^{d-2-\beta})(\log\frac{R}{r})_{\beta=d-2},\end{aligned} \tag{3.28}$$

and the same asymptotic lower bound can be obtained using (3.26) and (3.27). Then, if $\beta > d-2$, $\int_{r_0}^{\varepsilon R} r^{d-1} W'(r)(\text{ii.1.2})\,\mathrm{d}r \asymp 1 + R^{2d-\alpha-\beta}(\log R)_{\alpha+\beta=2d}$. If $\beta = d-2$, then

$$\int_{r_0}^{\varepsilon R} r^{d-1} W'(r)(\text{ii.1.2})\,\mathrm{d}r \asymp \int_{r_0}^{\varepsilon R} r^{d+1-\alpha} \log \frac{R}{r}\,\mathrm{d}r \asymp \log R + R^{d+2-\alpha}(\log R)^2_{\alpha=d+2}.$$

If $\beta < d-2$, then

$$\int_{r_0}^{\varepsilon R} r^{d-1} W'(r)(\text{ii.1.2})\,\mathrm{d}r \asymp R^{d-2-\beta} \int_{r_0}^{\varepsilon R} r^{d+1-\alpha}\,\mathrm{d}r \asymp R^{d-2-\beta} + R^{2d-\alpha-\beta}(\log R)_{\alpha=d+2}.$$

We summarise

$$\int_{r_0}^{\varepsilon R} r^{d-1} W'(r)(\text{ii.1.2})\,\mathrm{d}r \asymp 1 + R^{d-2-\beta}(\log R)_{\beta=d-2}(\log R)_{\alpha=d+2} + R^{2d-\alpha-\beta}(\log R)_{\alpha+\beta=2d,\beta>d-2}, \quad (3.29)$$

so with (3.25)

$$\int_{r_0}^{\varepsilon R} r^{d-1} W'(r)(\text{ii.1})\,\mathrm{d}r \asymp 1 + R^{d-2-\beta}(\log R)_{\beta=d-2}(\log R)_{\alpha=d+2} + R^{d-\alpha}(\log R)_{\beta=d}(\log R)_{\alpha=d} + R^{2d-\alpha-\beta}(\log R)_{\alpha+\beta=2d}. \quad (3.30)$$

(ii.2) Fix $r \in (r_0, \varepsilon R)$. If $y \in B_{d-1}(0, \sqrt{R^2-4r^2})$, so that $R'(y) - 2r > 0$, and $z \in (0, R'(y)-2r)$, then

$$\begin{aligned} 0 \le \Delta_{\cos}(z,y) \coloneqq \cos(z+2r,y) - \cos(z,y) &= \frac{z+2r}{\sqrt{(z+2r)^2+y^2}} - \frac{z}{\sqrt{z^2+y^2}} \\ &\le \frac{2r}{\sqrt{(2r)^2+z^2+y^2}} \\ &\le \frac{2r}{|(z,y)|} \wedge 1. \end{aligned} \quad (3.31)$$

We split the integral (ii.2) into $|(z,y)| < r$ and $|(z,y)| > r$, respectively denoted (ii.2.1) and (ii.2.2), see also Fig. 3.4. If $|(z,y)| < r$, then $R'(y) \ge \sqrt{R^2-r^2} \ge R\sqrt{1-\varepsilon^2} \ge 2\varepsilon R \ge 2r$, since $\varepsilon \le 1/\sqrt{5}$, so

$$\begin{aligned} (\text{ii.2.1}) &= \int_{B_{d-1}(0,R)} \mathrm{d}y \int_0^{(R'(y)-2r)_+} \mathrm{d}z\, \mathbb{1}_{\{|(z,y)|<r\}} \rho((r+z)e_1+y)\Delta_{\cos}(z,y) \\ &= \int_{B^+(0,r)} \mathrm{d}(z,y)\, \rho((r+z)e_1+y)\Delta_{\cos}(z,y). \end{aligned} \quad (3.32)$$

Using $\Delta_{\cos}(z,y) \le 1$, we immediately obtain an upper bound

$$(\text{ii.2.1}) \le \int_{B^+(0,r)} \rho((r+z)e_1+y)\,\mathrm{d}(z,y) = \int_{re_1+B^+(0,r)} \rho(x)\,\mathrm{d}x \asymp r^{d-\beta}.$$

For a matching lower bound, note that if $|(z,y)| < r$ then

$$\cos(z+2r,y) = \frac{z+2r}{\sqrt{(z+2r)^2+y^2}} \geq \frac{z+2r}{\sqrt{r^2+4zr+(2r)^2}} \geq \frac{2r}{\sqrt{5r^2}} = \frac{2}{\sqrt{5}},$$

and, if additionally $\theta \coloneqq \measuredangle(e_1, ze_1+y) \geq \pi/4$, say, that is $\cos(z,y) = \cos(\theta) \leq \frac{1}{\sqrt{2}}$, then $1 \geq \Delta_{\cos}(z,y) \geq \frac{2}{\sqrt{5}} - \frac{1}{\sqrt{2}} > 0$. Hence, again using (3.32),

$$\begin{aligned}
\text{(ii.2.1)} &\geq c \int_{B^+(0,r)} \mathbb{1}_{\{\cos(e_1, ze_1+y) \leq 1/\sqrt{2}\}} \rho((r+z)e_1+y)\,\mathrm{d}(z,y) \\
&\asymp r^{-\beta} \operatorname{vol}\left(\left\{x \in B^+(0,r)\colon \cos(e_1,x) \leq \frac{1}{\sqrt{2}}\right\}\right) \\
&\asymp r^{d-\beta}.
\end{aligned}$$

We conclude

$$\int_{r_0}^{\varepsilon R} r^{d-1} W'(r) \text{(ii.2.1)}\,\mathrm{d}r \asymp \int_{r_0}^{\varepsilon R} r^{2d-1-\alpha-\beta}\,\mathrm{d}r \asymp 1 + R^{2d-\alpha-\beta}(\log R)_{\alpha+\beta=2d}.$$

We turn to (ii.2.2). First note that $R - 2r \geq (1-2\varepsilon)R \geq R/2$ since $\varepsilon < 1/4$, so

$$\int_{A^+(r,R/2)} \mathrm{d}(y,z) \leq \int\limits_{B_{d-1}(0,R)} \mathrm{d}y \int\limits_0^{(R'(y)-2r)_+} \mathbb{1}_{\{|(z,y)|>r\}}\,\mathrm{d}z \leq \int_{A^+(r,R)} \mathrm{d}(y,z). \tag{3.33}$$

Here $A^+(a,b) = \big\{x \in \mathbb{R}^d\colon a < |x| < b,\, x \cdot e_1 > 0\big\}$. Furthermore, if $(z,y) \in B^+(0,R)$ with $|(z,y)| \geq r$, then

$$\begin{aligned}
&\rho((r+z)e_1+y) \geq \frac{1}{2}((r+z)^2+|y|^2)^{-\beta/2} \geq 2^{-1-\beta}|(z,y)|^{-\beta} \\
&\rho((r+z)e_1+y) \leq 2((r+z)^2+|y|^2)^{-\beta/2} \leq 2|(z,y)|^{-\beta}.
\end{aligned} \tag{3.34}$$

Furthermore $\Delta_{\cos}(z,y) \leq \frac{2r}{|(z,y)|}$ by (3.31), so with (3.33) and (3.34) we obtain an upper bound

$$\begin{aligned}
\text{(ii.2.2)} &= \int\limits_{B_{d-1}(0,R)} \mathrm{d}y \int\limits_0^{(R'(y)-2r)_+} \mathrm{d}z\, \mathbb{1}_{\{|(z,y)|>r\}} \rho((r+z)e_1+y)\Delta_{\cos}(z,y) \\
&\leq \int_{A^+(r,R)} 2|(z,y)|^{-\beta} \frac{2r}{|(y,z)|}\,\mathrm{d}(y,z) \\
&\asymp r \int_r^R s^{d-2-\beta}\,\mathrm{d}s \\
&\asymp r(r^{d-1-\beta} + R^{d-1-\beta})(\log \frac{R}{r})_{\beta=d-1}.
\end{aligned} \tag{3.35}$$

For a matching lower bound we first prove a lower bound on $\Delta_{\cos}(z,y)$ if $z \leq |y|$ (that is, if $\cos(z,y) \leq 1/\sqrt{2}$). Let for fixed $y > 0$, $h(x) \coloneqq x/\sqrt{x^2+y^2}$ for

$x > 0$. Then $h'(x) = y^2/(\sqrt{x^2+y^2})^3$ is non-negative and decreasing in x. Now if $|(z,y)| > r$ and $z \le |y|$ then $|y| \ge |(z,y)|/\sqrt{2} > r/\sqrt{2}$, so by Taylor's theorem

$$\Delta_{\cos}(z,y) = h(z+2r) - h(z) \ge 2rh'(z+2r) \ge 2rh'(|y| + \sqrt{2}|y|) = c\frac{r}{|y|} \ge c\frac{r}{|(z,y)|}.$$

Combining this with the lower bounds in (3.33) and (3.34) we obtain as in (3.35),

$$\text{(ii.2.2)} \ge cr \int_{A^+(r,R)} \mathbb{1}_{\{\cos(x,e_1) \le 1/\sqrt{2}\}} |x|^{-1-\beta} \,\mathrm{d}x \asymp r(r^{d-1-\beta} + R^{d-1-\beta})(\log \frac{R}{r})_{\beta=d-1}.$$

If $\beta > d-1$ then $\int_{r_0}^{\varepsilon R} r^{d-1} W'(r) \text{(ii.2.2)} \,\mathrm{d}r \asymp 1 + R^{2d-\alpha-\beta}(\log R)_{\alpha+\beta=2d}$. If $\beta < d-1$, then

$$\int_{r_0}^{\varepsilon R} r^{d-1} W'(r) \text{(ii.2.2)} \,\mathrm{d}r \asymp R^{d-1-\beta} \int_{r_0}^{\varepsilon R} r^{d-\alpha} \,\mathrm{d}r \asymp R^{d-1-\beta}(1 + R^{d+1-\alpha}(\log R)_{\alpha=d+1}).$$

Finally if $\beta = d-1$ then, by Lemma 3.11,

$$\int_{r_0}^{\varepsilon R} r^{d-1} W'(r) \text{(ii.2.2)} \,\mathrm{d}r \asymp \int_{r_0}^{\varepsilon R} r^{d-\alpha} \log \frac{R}{r} \,\mathrm{d}r = \log R + R^{d-\alpha+1}(\log R)^2_{\alpha=d+1}.$$

Combining this with (3.30) we obtain

$$\begin{aligned}\int_{r_0}^{\varepsilon R} r^{d-1} W'(r) \Big[\text{(ii.1)}+\text{(ii.2)}\Big] \,\mathrm{d}r \asymp 1 &+ R^{d-\alpha}(\log R)_{\beta=d}(\log R)_{\alpha=d} + R^{d-1-\beta}(\log R)_{\beta=d-1} \\ &+ R^{2d-\alpha-\beta}\Big[(\log R)_{\alpha=d+1}(\log R)_{\beta=d-1,\alpha=d+1} + (\log R)_{\alpha+\beta=2d}\Big]. \quad (3.36)\end{aligned}$$

(ii.3) First note that

$$\begin{aligned}|\text{(ii.3)}| &= \int_{B^+(0,R)\setminus(B^+(-2re_1,R)} \mathrm{d}(y,z)\, \rho((r+z)e_1 + y) \cos(z,y) \\ &= \int_0^R \mathrm{d}z \int_{A_{d-1}((\sqrt{R^2-(z+2r)^2})_+, \sqrt{R^2-z^2})} \mathrm{d}y\, \rho((r+z)e_1 + y) \cos(z,y). \end{aligned} \quad (3.37)$$

If $(z,y) \in B^+(0,R) \setminus B^+(-2re_1, R)$, then $x' := (r+z)e_1 + y \in B^+(re_1, R) \setminus B^+(-re_1, R)$ and $R \ge |x'| \ge R - r \ge R/2$ (since $\varepsilon \le 1/2$), so $\rho(x') \asymp R^{-\beta}$. Furthermore,

$$\mathrm{vol}(B^+(0,R) \setminus B^+(-2re_1, R)) \le \mathrm{vol}(A^+(R-2r, R)) \asymp \int_{R-2r}^R s^{d-1} \,\mathrm{d}s \asymp rR^{d-1},$$

and thus

$$|\text{(ii.3)}| \le crR^{d-1-\beta}. \quad (3.38)$$

Then if $\alpha \geq d+1$,

$$\int_{r_0}^{\varepsilon R} r^{d-1} W'(r) |(\text{ii.3})| \, \mathrm{d}r \leq c R^{d-1-\beta} (\log R)_{\alpha=d+1}.$$

If $\alpha < d+1$, then

$$\int_{r_0}^{\varepsilon R} r^{d-1} W'(r) |(\text{ii.3})| \, \mathrm{d}r \leq c \varepsilon^{d+1-\alpha} R^{2d-\alpha-\beta}.$$

By (3.23), if $\alpha \geq d+1$ and $\varepsilon = \varepsilon(K) > 0$ is small enough then $\int_{r_0}^{\varepsilon R} r^{d-1} W'(r) |(\text{ii.3})| \, \mathrm{d}r \leq \frac{1}{3} F_{(KR,\infty)}$ for large R.

Hence, and by (3.23) and (3.36), if $\alpha + \beta > d-1$ (otherwise $F_{(KR,\infty)}(R) = \infty$),

$$\begin{aligned} F_{(r_0,\varepsilon R)}(R) + F_{(KR,\infty)}(R) \asymp 1 \pm R^{d-1-\beta} (\log R)_{\alpha=d+1} (\log R)_{\beta=d-1} + R^{d-\alpha} (\log R)_{\beta=d} (\log R)_{\alpha=d} \\ + R^{2d-\alpha-\beta} \Big[(\log R)_{\alpha=d+1} (\log R)_{\beta=d-1, \alpha=d+1} + (\log R)_{\alpha+\beta=2d} \Big]. \quad (3.39) \end{aligned}$$

(iii) **The contribution** $(0, r_0)$

We decompose the same way as we did in $r_0 \to \varepsilon R$. That is, for fixed $r \in (0, r_0)$,

$$\begin{aligned} &\int_{B_{d-1}(0,R)} \mathrm{d}y \int_0^{R'(y)} \mathrm{d}z \, \cos(e_1, z e_1 + y) \Big(\rho((-r+z)e_1 + y) - \rho((r+z)e_1 + y) \Big) \\ &= \int_{B_{d-1}(0,R)} \mathrm{d}y \int_0^{2r \wedge R'(y)} \mathrm{d}z \cos(e_1, z e_1 + y) \rho((-r+z)e_1 + y) \qquad (\text{iii.1}) \\ &\quad + \int_{B_{d-1}(0,R)} \mathrm{d}y \int_0^{(R'(y)-2r)_+} \mathrm{d}z \, \rho((r+z)e_1 + y) (\cos(e_1, (z+2r)e_1 + y) - \cos(e_1, z e_1 + y)) \\ &\qquad\qquad (\text{iii.2}) \\ &\quad - \int_{B_{d-1}(0,R)} \mathrm{d}y \int_{(R'(y)-2r)_+}^{R'(y)} \mathrm{d}z \, \rho((r+z)e_1 + y) \cos(e_1, z e_1 + y). \qquad (\text{iii.3}) \end{aligned}$$

Recall for this purpose also Figs. 3.3 and 3.4. For (iii.1), we split into $|y| < 2r_0$ and $|y| > 2r_0$, respectively denoted (iii.1.1) and (iii.1.2). The former is non-negative and bounded as $R \to \infty$. For the latter, by arguments identical to those used for (ii.1.2), see e.g. (3.26) to (3.28), we find

$$(\text{iii.1.2}) \asymp \int_{r_0}^{R} \mathrm{d}s \, s^{d-2} \int_0^{2r} \mathrm{d}z \, z s^{-1-\beta} \asymp r^2 \int_{r_0}^{R} s^{d-3-\beta} \, \mathrm{d}s \asymp r^2 (1 + R^{d-2-\beta}) (\log R)_{\beta=d-2}.$$

Then,

$$\begin{aligned} \left| \int_0^{r_0} r^{d-1} W'(r) (\text{iii.1.2}) \, \mathrm{d}r \right| &\leq c (1 + R^{d-2-\beta}) (\log R)_{\beta=d-2} \int_{B(0,r_0)} |x|^2 |\nabla W(x)| \, \mathrm{d}x \\ &\leq c (1 + R^{d-2-\beta}) (\log R)_{\beta=d-2}, \qquad (3.40) \end{aligned}$$

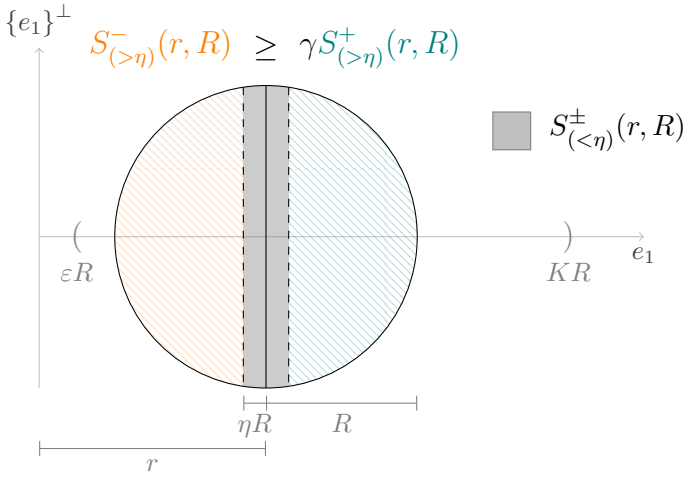

Figure 3.5: Illustration for proof idea for $F_{(\varepsilon R, KR)}(R)$.

because ∇W is locally integrable by assumption.

We split (iii.2) into $|(z,y)| < r$ and $|(z,y)| > r$, respectively denoted (iii.2.1) and (iii.2.2). The former is non-negative and bounded as $R \to \infty$. For the latter, by arguments identical to those used for (ii.2.2), see e.g. (3.33) to (3.36)

$$\text{(iii.2.2)} \asymp r \int_{r_0}^{R} s^{d-2-\beta} \,\mathrm{d}s \asymp r(1 + R^{d-1-\beta})(\log R)_{\beta=d-1}.$$

Then similar to (3.40) we conclude by local integrability of ∇W that

$$\left| \int_0^{r_0} r^{d-1} W'(r) \text{(iii.2.2)} \,\mathrm{d}r \right| \leq c(1 + R^{d-1-\beta})(\log R)_{\beta=d-1}. \tag{3.41}$$

For (iii.3) we obtain identically to how we derived (3.38) that

$$|\text{(iii.3)}| \leq crR^{d-1-\beta}.$$

Again local integrability of ∇W then implies $\int_0^{r_0} r^{d-1} W'(r) \text{(iii.3)} \,\mathrm{d}r \leq cR^{d-1-\beta}$. We summarise

$$\left| F_{(0,r_0)}(R) \right| \leq c(1 + R^{d-1-\beta}(\log R)_{\beta=d-1}).$$

(iv) **The contribution** $(\varepsilon R, KR)$

Recall from (3.14) and (3.17) that

$$F_{(\varepsilon R, KR)}(R) = \int_{\varepsilon R}^{KR} r^{d-1} W'(r) \Big(S^-(r,R) - S^+(r,R) \Big) \,\mathrm{d}r.$$

For a definition and visualisation of $S^\pm(r,R)$ recall (3.15) and Fig. 3.2. For $r \in (\varepsilon R, KR)$, denote

$$B^+_{(>\eta)}(0,R) \coloneqq \left\{ x \in B^+(0,R) \colon x \cdot e_1 > \eta R \right\} = \left\{ ze_1 + y \colon (z,y) \in B^+(0,R), z > \eta R \right\},$$

$$S^\pm_{(>\eta)}(r,R) \coloneqq \int_{B^+_{(>\eta)}(0,R)} \cos(e_1, ze_1 + y) \rho((r \pm z)e_1 + y) \,\mathrm{d}(y,z),$$

and $B^+_{(<\eta)}$, $S^\pm_{(<\eta)}(r,R)$ similarly. See also Fig. 3.5. Then, $S^+_{(<\eta)}(r,R) \le c(\varepsilon)\eta R^{d-\beta}$, so

$$\int_{\varepsilon R}^{KR} r^{d-1} W'(r) S^+_{(<\eta)} \,\mathrm{d}r \le c(\varepsilon, K)\eta R^{2d-\alpha-\beta}, \tag{3.42}$$

which by (3.23) is smaller than $F_{(KR,\infty)}(R)/3$ for large R if $\eta = \eta(\varepsilon, K) > 0$ is small enough. Now we show that there is $\gamma = \gamma(\eta, \varepsilon, K) > 1$ such that for large R and $r \in (\varepsilon R, KR)$, $S^-_{(>\eta)}(r,R) \ge \gamma S^+_{(>\eta)}(r,R)$. If $r \in (\varepsilon R, KR)$ and $x = (r+z)e_1 + y \in S^+_{(>\eta)}(r,R)$, $x' := (r-z)e_1 + y \in S^-_{(>\eta)}(r,R)$, then by Lemma 3.10,

$$\frac{|x|^2}{|x'|^2} = \frac{(r+z)^2 + |y|^2}{(r-z)^2 + |y|^2} \ge \frac{1 + (z/R)^2}{1 + (1 - \frac{r}{z} \wedge \frac{z}{r})^2 (z/R)^2} \ge \frac{1+\eta^2}{1 + (1 - \varepsilon \wedge \frac{\eta}{K})^2 \eta^2} =: a_1(\varepsilon, \eta, K)^2 > 1.$$

Hence $|x| \ge (\varepsilon R) \vee (a_1 |x'|)$, so if $a_0 \in (0,1)$ and $\gamma > 1$ with $a_0 \gamma^{-1} \in (a_1^{-\beta}, 1)$, then for $R \ge R_0(\varepsilon, a_0, \gamma)$,

$$\rho(x) \le (\gamma^{-1} a_0 |x'|^{-\beta}) \wedge (c(\varepsilon) R^{-\beta}).$$

Now let $r_1 = r_1(a_0) > 0$ be large enough that $\rho(\xi) \ge a_0 |\xi|^{-\beta}$ for all $|\xi| \ge r_1$. Then

$$\rho(x') \ge \Big(\inf_{|\xi| \le r_1} \rho(\xi) \Big) \vee \big(a_0 |x'|^{-\beta} \big) \ge \gamma \rho(x),$$

where we were justified to assume that $\rho > 0$ in $B(0, r_1)$ (and thus everywhere); indeed, $F(R)$ is linear in ρ, and if ρ is compactly supported than it is easy to check that $F(R) \asymp 1 + R^{d-\alpha} (\log R)_{\alpha=d}$, which is absorbed by $F_{(r_0, \varepsilon R)}(R)$, see (3.39). Thus, for $R \ge R_0(\varepsilon, a_0, \gamma) = R_0(\varepsilon, \eta, K)$, and $r \in (\varepsilon R, KR)$,

$$\begin{aligned} S^-_{(>\eta)}(r,R) &= \int_{B^+_{(>\eta)}(0,R)} \cos(e_1, ze_1 + y) \rho((r-z)e_1 + y) \,\mathrm{d}(y,z) \\ &\ge \gamma \int_{B^+_{(>\eta)}(0,R)} \cos(e_1, \mathrm{e}_1 + y) \rho((r+z)e_1 + y) \,\mathrm{d}(y,z) \\ &= \gamma S^+_{(>\eta)}(r,R). \end{aligned}$$

We conclude

$$\begin{aligned} \int_{\varepsilon R}^{KR} r^{d-1} W'(r) \Big(S^-(r,R) - S^+_{(>\eta)}(r,R) \Big) \,\mathrm{d}r &\asymp \int_{\varepsilon R}^{KR} r^{d-1} W'(r) S^-(r,R) \,\mathrm{d}r \\ &\asymp R^{d-\alpha} + R^{2d-\alpha-\beta} (\log R)_{\beta=d}, \end{aligned}$$

where we used that, if say $r < R/2$ then $S^-(r,R) \asymp 1 + R^{d-\beta} (\log R)_{\beta=d}$, if $r > 2R$ then $S^-(r,R) \asymp R^{d-\beta}$, and certainly $S^-(r,R) \le c(\varepsilon)(1 + R^{d-\beta}(\log R)_{\beta=d})$ for any $r > \varepsilon R$. Recalling (3.42) and the subsequent comment, we summarise

$$F_{(\varepsilon R, \infty)}(R) \asymp R^{d-\alpha} + R^{2d-\alpha-\beta} (\log R)_{\beta=d},$$

as long as $\alpha + \beta > d - 1$ (otherwise $F_{(KR,\infty)}(R) = \infty$).

Combining (i) to (iv), we conclude that $F(R) = \infty$ for large R if $\alpha+\beta \le d-1$, otherwise

$$\begin{aligned} F(R) &\asymp \pm\,(1 + R^{d-1-\beta}(\log R)_{\beta=d-1}(\log R)_{\alpha=d+1}) + R^{d-\alpha}(\log R)_{\beta=d}(\log R)_{\alpha=d} \\ &\quad + R^{2d-\alpha-\beta}\Big[(\log R)_{\alpha=d+1}(\log R)_{\beta=d-1,\alpha=d+1} + (\log R)_{\alpha+\beta=2d} + (\log R)_{\beta=d}\Big] \\ &\asymp \pm(1 + R^{d-1-\beta}(\log R)_{\beta=d-1}(\log R)_{\alpha=d+1} + R^{d-\alpha}(\log R)_{\alpha=d}) \\ &\quad + R^{2d-\alpha-\beta}(\log R)_{\alpha+\beta=2d}(\log R)_{\beta=d}(\log R)_{\alpha=d+1}, \end{aligned}$$

where the second step above can be confirmed with a tedious but straightforward case distinction. □

Lemma 3.10. *If $R > 0$, and $z, r > 0$, $y \in B_{d-1}(0, R)$, then*

$$\frac{(r+z)^2 + y^2}{(r-z)^2 + y^2} \ge \frac{1 + (z/R)^2}{1 + (1 - \frac{r}{z} \wedge \frac{z}{r})^2 (z/R)^2}.$$

Proof. Without loss of generality we assume $R = 1$. Then, since $(r-z)^2 \le (r+z)^2$,

$$D := \frac{(r+z)^2 + y^2}{(r-z)^2 + y^2} \ge \frac{(r+z)^2 + 1}{(r-z)^2 + 1}.$$

If $z = r$ the statement is trivial. If $z > r$ then $(r+z)/(z-r) \ge z/(z-r) = 1/(1-\frac{r}{z})$, so

$$D \ge \frac{1 + (r+z)^2}{1 + (1-\frac{r}{z})^2 (r+z)^2} \ge \frac{1+z^2}{1 + (1-\frac{r}{z})^2 z^2}.$$

If $z < r$ then $(r+z)/(r-z) \ge r/(r-z) = 1/(1-\frac{z}{r})$, and the same argument yields the desired statement. □

Lemma 3.11. *For fixed $a \in \mathbb{R}$, $r_0 > 0$, and $c \in (0,1)$,*

$$\int_{r_0}^{cR} r^a \log \frac{R}{r}\,\mathrm{d}r \asymp \begin{cases} \log R + R^{a+1}, & a \ne -1, \\ (\log R)^2, & a = -1, \end{cases} \qquad R \to \infty,$$

and

$$\int_{r_0}^{cR} r^a \log r\,\mathrm{d}r \asymp \begin{cases} 1 + R^{a+1}\log R, & a \ne -1, \\ (\log R)^2, & a = -1. \end{cases}$$

Proof.

$$\begin{aligned} \int_{r_0}^{cR} r^a \log\left(\frac{R}{r}\right)\mathrm{d}r &= \frac{r^{a+1}}{(a+1)^2}\left((a+1)\log\frac{R}{r} + 1\right)\Big|_{r=r_0}^{cR} \\ &= \frac{(cR)^{a+1}}{(a+1)^2}\left((a+1)\log\frac{1}{c} + 1\right) - \frac{r_0^{a+1}}{(a+1)^2}\left((a+1)\log\frac{R}{r_0} + 1\right) \\ &\asymp R^{a+1}. \end{aligned}$$

The second statement is similarly straightforward. □

Finally, we prove a small lemma referred to in Section 3.1

Lemma 3.12. *If $\rho \in L^\infty$ is non-negative and $\nabla W \star \rho$ exists as a Lebesgue integral, then $\nabla W \mathbin{\widetilde{\star}} \rho$ exists in the sense of Definition 3.1, and $\nabla W \star \rho = \nabla W \mathbin{\widetilde{\star}} \rho$.*

Proof. Because ρ is bounded and ∇W locally integrable, existence of $\nabla W \star \rho$ as a Lebesgue integral is equivalent to

$$\int_{\mathbb{R}^d \setminus B(0,1)} |x|^{-\alpha} \rho(x) \, \mathrm{d}x < \infty. \tag{3.43}$$

It now suffices to show that

$$\delta_N(x_0) \coloneqq \left| \int_{B(0,N)^c} \nabla W(x) \rho(x_0 - x) \, \mathrm{d}x \right|$$

converges to zero locally uniformly as $N \to \infty$. Indeed,

$$\begin{aligned} \delta_N(x_0) = \left| \int_{B(x_0,N)^c} \rho(x) \nabla W(x_0 - x) \, \mathrm{d}x \right| &\leq c \int_{B(x_0,N)^c} \rho(x) |x_0 - x|^{-\alpha} \, \mathrm{d}x \\ &\leq c \int_{B(x_0,N)^c} \rho(x) (|x|^{-\alpha} + |x_0||x|^{-\alpha-1}) \, \mathrm{d}x. \end{aligned} \tag{3.44}$$

Then by (3.43) and (3.44), for any $R > 0$,

$$\sup_{|x_0| \leq R} \delta_N(x_0) \leq \left(1 + \frac{R}{N-R}\right) \int_{B(0,N-R)^c} \rho(x) |x|^{-\alpha} \, \mathrm{d}x \to 0, \quad N \to \infty.$$

□

Acknowledgments

I would like to thank my D.Phil. advisor Alison Etheridge for many lively discussions, guidance, encouragement, and helpful feedback. I would like to thank José Carrillo for many insightful conversations, valuable feedback, for pointing out useful literature, and for advice regarding publication of the project. I thank the referees of [30]—the paper corresponding to Section 2—particularly for helping to improve the presentation, and for pointing out important missing references. I further acknowledge support by the EPSRC grant EP/W523781/1.

Bibliography

[1] Elie Aidékon et al. "Branching Brownian motion seen from its tip". In: *Probab. Theory and Related Fields* 157 (2013), pp. 405–451.

[2] Louis-Pierre Arguin, Anton Bovier, and Nicola Kistler. "The extremal process of branching Brownian motion". In: *Probab. Theory Related Fields* 157.3-4 (2013), pp. 535–574.

[3] Jacob Bedrossian and Nancy Rodriguez. "Inhomogeneous Patlak–Keller–Segel models and aggregation equations with nonlinear diffusion in Rd". In: *Discrete Contin. Dyn. Syst., Ser. B* 19.5 (2014), p. 1279.

[4] Jacob Bedrossian, Nancy Rodriguez, and Andrea L. Bertozzi. "Local and global well-posedness for aggregation equations and Patlak–Keller–Segel models with degenerate diffusion". In: *Nonlinearity* 24.6 (2011), p. 1683.

[5] Adrien Blanchet, Jean Dolbeault, and Benoit Perthame. "Two-dimensional Keller-Segel model: Optimal critical mass and qualitative properties of the solutions." In: *Electron. J. Differential Equations* 2006 (2006).

[6] Anton Bovier and Lisa Hartung. "Branching Brownian Motion with Self-Repulsion". In: *Ann. Henri Poincaré*. Springer. 2022, pp. 1–26.

[7] Lorenzo Brasco, David Gómez-Castro, and Juan L. Vázquez. "Characterisation of homogeneous fractional Sobolev spaces". In: *Calc. Var. Partial Differ. Equ.* 60.2 (2021), pp. 1–40.

[8] Martin Burger, Marco Di Francesco, and Marzena Franek. "Stationary states of quadratic diffusion equations with long-range attraction". In: *Commun. Math. Sci.* 11.3 (2013), pp. 709–738.

[9] Vincent Calvez and José A. Carrillo. "Volume effects in the Keller–Segel model: energy estimates preventing blow-up". In: *J. Math. Pures Appl. (9)* 86.2 (2006), pp. 155–175.

[10] José A Carrillo, Katy Craig, and Yao Yao. "Aggregation-Diffusion Equations: Dynamics, Asymptotics, and Singular Limits". In: *Active Particles, Volume 2* (2019), p. 65.

[11] José A Carrillo et al. "Asymptotic simplification of Aggregation-Diffusion equations towards the heat kernel". In: *Arch. Ration. Mech. Anal.* 247.1 (2023), p. 11.

[12] José A. Carrillo et al. "Ground states in the diffusion-dominated regime". In: *Calc. Var. Partial Differ. Equ.* 57.5 (2018), pp. 1–28.

[13] José A. Carrillo et al. "Nonlinear aggregation-diffusion equations: radial symmetry and long time asymptotics". In: *Invent. Math.* 218.3 (2019), pp. 889–977.

[14] Li Chen, Jian-Guo Liu, and Jinhuan Wang. "Multidimensional Degenerate Keller–Segel System with Critical Diffusion Exponent 2n/(n+2)". In: *SIAM J. Math. Anal.* 44.2 (2012), pp. 1077–1102.

[15] Li Chen and Jinhuan Wang. "Exact Criterion for Global Existence and Blow Up to a Degenerate Keller–Segel System". In: *Doc. Math.* 19 (2014), pp. 103–120.

[16] Maria R D'Orsogna et al. "Self-propelled particles with soft-core interactions: patterns, stability, and collapse". In: *Phys. Rev. Lett.* 96.10 (2006), p. 104302.

[17] Eleonora Di Nezza, Giampiero Palatucci, and Enrico Valdinoci. "Hitchhiker's guide to the fractional Sobolev spaces". In: *Bull. Sci. Math.* 136.5 (2012), pp. 521–573. ISSN: 0007-4497. DOI: https://doi.org/10.1016/j.bulsci.2011.12.004.

[18] János Engländer. "The Center of Mass for Spatial Branching Processes and an Application for Self-Interaction". In: *Electronic J. Probab.* 15 (2010), pp. 1938–1970. DOI: 10.1214/EJP.v15-822.

[19] János Engländer and Liang Zhang. "Branching diffusion with particle interactions". In: *Electron. J. Probab.* 21 (2016), pp. 1–25. DOI: 10.1214/16-EJP4782.

[20] Alison M. Etheridge. *An introduction to superprocesses*. 20. American Mathematical Soc., 2000.

[21] Alison M. Etheridge. "Survival and extinction in a locally regulated population". In: *Ann. Appl. Probab.* 14.1 (2004), pp. 188–214. DOI: 10.1214/aoap/1075828051.

[22] Steven N. Evans. "Two representations of a conditioned superprocess". In: *Proc. R. Soc. Edinb., Sect. A, Math.* 123.5 (1993), pp. 959–971. DOI: 10.1017/S0308210500029619.

[23] Joseph Felsenstein. "A Pain in the Torus: Some Difficulties with Models of Isolation by Distance". In: *The American Naturalist* 109.967 (1975), pp. 359–368. ISSN: 00030147, 15375323.

[24] Hardeep Gill. "A Super Onstein-Uhlenbeck process interacting with its center of mass". In: *Ann. Probab.* 41.2 (2013), pp. 989–1029.

[25] Miguel A. Herrero and Juan J.L. Velázquez. "Singularity patterns in a chemotaxis model". In: *Math. Ann.* 306 (1996), pp. 583–623. DOI: 10.1007/BF01445268.

[26] Wen-Ming Hong and Zeng-Hu Li. "A central limit theorem for super-Brownian motion with super-Brownian immigration". In: *J. Appl. Probab.* 36.4 (1999), pp. 1218–1224.

[27] Dirk Horstmann. "From 1970 until present: the Keller–Segel model in chemotaxis and its consequences". In: (2003).

[28] Gunnar Kaib. "Stationary states of an aggregation equation with degenerate diffusion and bounded attractive potential". In: *SIAM J. Math. Anal.* 49.1 (2017), pp. 272–296.

[29] Evelyn F. Keller and Lee A. Segel. "Model for chemotaxis". In: *J. Theor. Biol.* 30.2 (1971), pp. 225–234.

[30] Peter Koepernik. "On a repulsion-diffusion equation with immigration". In: *Discrete and Continuous Dynamical Systems* 44.4 (2024), pp. 1106–1133. DOI: 10.3934/dcds.2023140.

[31] Zeng-Hu Li. "Immigration structures associated with Dawson-Watanabe superprocesses". In: *Stochastic Process. Appl.* 62.1 (1996), pp. 73–86.

[32] Elliott H. Lieb and Horng-Tzer Yau. "The Chandrasekhar theory of stellar collapse as the limit of quantum mechanics". In: *Commun. Math. Phys.* 112.1 (1987), pp. 147–174.

[33] Karl Oelschläger. "Large systems of interacting particles and the porous medium equation". In: *J. Differ. Equations* 88.2 (1990), pp. 294–346. ISSN: 0022-0396. DOI: https://doi.org/10.1016/0022-0396(90)90101-T.

[34] Clifford S. Patlak. "Random walk with persistence and external bias". In: *Bull. Math. Biol.* 15 (1953), pp. 311–338.

[35] Edwin A. Perkins. "Dawson-Watanabe superprocesses and measure-valued diffusions". In: *Lectures on probability theory and statistics (Saint-Flour, 1999)*. Vol. 1781. Lecture Notes in Math. Springer, Berlin, 2002, pp. 125–324.

[36] Yoshie Sugiyama. "Time global existence and asymptotic behavior of solutions to degenerate quasi-linear parabolic systems of chemotaxis". In: *Differ. Integral Equ.* 20.2 (2007), pp. 133–180.

[37] Félix del Teso, David Gómez-Castro, and Juan L. Vázquez. "Estimates on translations and Taylor expansions in fractional Sobolev spaces". In: *Nonlinear Anal.* 200 (2020), p. 111995.

[38] Chad M. Topaz, Andrea L. Bertozzi, and Mark A. Lewis. "A nonlocal continuum model for biological aggregation". In: *Bull. Math. Biol.* 68.7 (2006), pp. 1601–1623.

[39] Hong Wenming. "Limiting behavior of the super-Brownian motion with super-Brownian immigration". In: *CR Math. Rep. Acad. Sci. Canada* 25.1 (2003), pp. 1–6.

[40] Mei Zhang. "Large deviations for super-Brownian motion with immigration". In: *J. Appl. Probab.* 41.1 (2004), pp. 187–201.

[41] Mei Zhang. "Some scaled limit theorems for an immigration super-Brownian motion". In: *Science in China Series A: Mathematics* 51.2 (2008), pp. 203–214.

[42] Yuming Paul Zhang. "On a class of diffusion-aggregation equations". In: *Discrete Contin. Dyn. Syst.* 40.2 (2020), p. 907.

Appendix

III.A Fractional Sobolev spaces

We give a minimal definition of fractional Sobolev spaces, and refer the reader to [3, 4, 1] for detailed introductions. For $p \in [1, \infty)$, $s \in (0, 1)$, and $u \colon \mathbb{R}^d \to \mathbb{R}$ measurable let

$$[u]_{\mathcal{W}^{s,p}} \coloneqq \left(s(1-s) \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \frac{|u(x) - u(y)|^p}{|x-y|^{d+sp}} \, \mathrm{d}x \, \mathrm{d}y \right)^{1/p},$$

and

$$[u]_{\mathcal{W}^{s,\infty}} = \sup_{x \neq y} \frac{|u(x) - u(y)|}{|x-y|^s}. \tag{III.A.1}$$

For $k \in \mathbb{N}$, $s \in (k, k+1)$, and $p \in [1, \infty]$, let

$$[u]_{\mathcal{W}^{s,p}} \coloneqq \sum_{|\alpha| = k} [\partial^\alpha u]_{\mathcal{W}^{s-k,p}},$$

with the usual notation for multi-indices α. Then

$$\|u\|_{\mathcal{W}^{s,p}} = \left(\|u\|_{\mathcal{W}^{\lfloor s \rfloor, p}}^p + [u]_{\mathcal{W}^{s,p}}^p \right)^{1/p}$$

for $p \in [1, \infty)$, and

$$\|u\|_{\mathcal{W}^{s,\infty}} = \|u\|_{\mathcal{W}^{\lfloor s \rfloor, \infty}} + [u]_{\mathcal{W}^{s,\infty}}, \tag{III.A.2}$$

define the fractional Sobolev norms.

For two normed spaces write $A \hookrightarrow B$ if $A \subset B$ with continuous inclusion.

Lemma III.A.1. *If $0 \leq \gamma \leq \beta$ and $p \in [1, \infty]$, then $\mathcal{W}^{\beta,p} \hookrightarrow \mathcal{W}^{\gamma,p}$.*

Proof. Assume $0 < \gamma < \beta$, otherwise there is nothing to show. Then the claim follows because $\mathcal{W}^{\gamma,p}$ can be written as interpolation space between L^p and $\mathcal{W}^{\beta,p}$, so $\|\cdot\|_{\mathcal{W}^{\gamma,p}} \leq C(\|\cdot\|_{L^p} + \|\cdot\|_{\mathcal{W}^{\beta,p}}) \leq C \|\cdot\|_{\mathcal{W}^{\beta,p}}$. See [2, Appendix A] for details on interpolation spaces in the context of Sobolev norms. □

Lemma III.A.2. (i) *If $\gamma > 0$, $\gamma \notin \mathbb{N}$, then $\mathcal{X}^\gamma \hookrightarrow \mathcal{W}^{\gamma,\infty} = C^{\lfloor \gamma \rfloor, \gamma - \lfloor \gamma \rfloor}(\mathbb{R}^d)$.*

(ii) If $\gamma \geq 1$, then $\mathcal{X}^\gamma \hookrightarrow \mathcal{W}^{\gamma,\infty} \hookrightarrow C^{\lfloor \gamma \rfloor - 1, 1}(\mathbb{R}^d)$.

Proof. First note that $\mathcal{X}^\gamma \hookrightarrow \mathcal{W}^{\gamma,\infty}$ by definition of $\mathcal{X}^\gamma$, see (2.10).

(i) For any $k \in \mathbb{N}_0$ and $\gamma \in (k, k+1)$, if $f\colon \mathbb{R}^d \to \mathbb{R}$ is k-times differentiable then $\|f\|_{\mathcal{W}^{\gamma,\infty}} = \|f\|_{C^{k,\gamma-k}(\mathbb{R}^d)}$ by definition, see (III.A.1), (III.A.2) and (2.11). Hence we only have to prove that any $f \in \mathcal{W}^{\gamma,\infty}$ for $\gamma > 0$ is $\lfloor\gamma\rfloor$-times differentiable.

There is nothing to prove for $\gamma \in (0,1)$, so suppose the claim is true for $\gamma \in (0,k)\setminus\mathbb{N}$ for some $k \in \mathbb{N}_0$ and let $f \in \mathcal{W}^{\gamma,\infty}$ for some $\gamma \in (k, k+1)$. Put $\beta \coloneqq \gamma - \lfloor\gamma\rfloor$. Then f is continuous because $\mathcal{W}^{\gamma,\infty} \subset \mathcal{W}^{1,\infty} = C^{0,1}(\mathbb{R}^d)$ (the last equality is a well-known theorem). Furthermore $f \in \mathcal{W}^{\gamma,\infty} \subset \mathcal{W}^{1+\beta,\infty}$, so f has partial weak derivatives of first order that are β-Hölder continuous, so they are in fact proper derivatives [2]. By the induction hypothesis, all first partial derivatives are themselves $(\lfloor\gamma\rfloor - 1)$ times differentiable.

(ii) Since $\mathcal{W}^{\gamma,\infty} \hookrightarrow \mathcal{W}^{\lfloor\gamma\rfloor,\infty}$, it suffices to show that $\mathcal{W}^{k,\infty} \hookrightarrow C^{k-1,1}(\mathbb{R}^d)$ for all $k \in \mathbb{N}$. This is a well-known theorem for $k = 1$. Suppose now it is proved for some $k \in \mathbb{N}$, and let $f \in \mathcal{W}^{k+1,\infty}$. Then by (i), f has proper first partial derivatives, and they are in $\mathcal{W}^{k,\infty} \hookrightarrow C^{k-1,1}(\mathbb{R}^d)$, so in fact f is k-times differentiable and, by the induction hypothesis,

$$\begin{aligned}\|f\|_{\mathcal{W}^{k+1,\infty}} &= \|f\|_\infty + \sum_{i=1}^d \|\partial_i f\|_{\mathcal{W}^{k,\infty}} \\ &\le \|f\|_\infty + C\sum_{i=1}^d \|\partial_i f\|_{C^{k-1,1}} \\ &\le C\,\|f\|_{C^{k,1}}\,.\end{aligned}$$

□

Lemma III.A.3. *Let $0 \le \gamma < \beta$. Then there is $C > 0$ such that the following hold.*

(i) If $f, g \in \mathcal{W}^{\gamma,\infty}$, then $fg \in \mathcal{W}^{\gamma,\infty}$ and $\|fg\|_{\mathcal{W}^{\gamma,\infty}} \le C\,\|f\|_{\mathcal{W}^{\gamma,\infty}}\,\|g\|_{\mathcal{W}^{\gamma,\infty}}$,

(ii) If $f \in \mathcal{W}^{\gamma,1}, g \in \mathcal{W}^{\beta,\infty}$, then $fg \in \mathcal{W}^{\gamma,1}$ and $\|fg\|_{\mathcal{W}^{\gamma,1}} \le C\,\|f\|_{\mathcal{W}^{\gamma,1}}\,\|g\|_{\mathcal{W}^{\beta,\infty}}$.

Proof. If $\gamma \in \mathbb{N}_0$ and $p \in \{1,\infty\}$, then for any multi-index $|\alpha| \le \gamma$,

$$\begin{aligned}\|\partial^\alpha(fg)\|_{L^p} &= \left\|\sum_{\alpha_1+\alpha_2=\alpha} (\partial^{\alpha_1} f)(\partial^{\alpha_2} g)\right\|_{L^p} \le \sum_{\alpha_1+\alpha_2=\alpha} \|\partial^{\alpha_1} f\|_{L^p}\,\|\partial^{\alpha_2} g\|_{L^\infty} \\ &\le C\,\|f\|_{\mathcal{W}^{\gamma,p}}\,\|g\|_{\mathcal{W}^{\gamma,\infty}}\,.\end{aligned}$$

[2] Let $d = 1$ and φ_ε an approximation to unity in $\mathbb{R}$, $x > 0$, and put $\psi_\varepsilon \coloneqq \varphi_\varepsilon \star \mathbb{1}_{[0,x]}$, then

$$\int_0^x f'(y)\,\mathrm{d}y \leftarrow \int_{\mathbb{R}} \psi_\varepsilon(y) f'(y)\,\mathrm{d}y = -\int_{\mathbb{R}} \psi_\varepsilon'(y) f(y)\,\mathrm{d}y = f(x) - f(0),$$

so $f(x) = f(0) + \int_0^x f'(y)\,\mathrm{d}y$. Analogously for $x < 0$. Since f' is continuous, f is in fact differentiable in the classical sense with derivative f'. This also works in $d \ge 2$.

and both claims follow. If $\gamma \notin \mathbb{N}_0$, we need to show the claim with the LHS replaced by $[fg]_{\mathcal{W}^{\gamma,p}}$. By a similar application of the product rule it suffices to consider $\gamma \in (0,1)$. If $p = \infty$, then

$$\begin{aligned}
[fg]_{\mathcal{W}^{\gamma,\infty}} &= \sup_{x \neq y} \frac{|f(x)g(x) - f(y)g(y)|}{|x-y|^\gamma} \\
&\leq \|f\|_{L^\infty} [g]_{\mathcal{W}^{\gamma,\infty}} + \|g\|_{L^\infty} [f]_{\mathcal{W}^{\gamma,\infty}} \\
&\leq 2 \|f\|_{\mathcal{W}^{\gamma,\infty}} \|g\|_{\mathcal{W}^{\gamma,\infty}} .
\end{aligned}$$

If $p = 1$, then

$$\begin{aligned}
[fg]_{\mathcal{W}^{\gamma,1}} &= \gamma(1-\gamma) \int\int \frac{|f(x)g(x) - f(y)g(y)|}{|x-y|^{d+\gamma}} \,\mathrm{d}x \,\mathrm{d}y \\
&\leq C \|g\|_{L^\infty} [f]_{\mathcal{W}^{\gamma,1}} + \int\int |f(y)| \frac{|g(x) - g(y)|}{|x-y|^{d+\gamma}} \,\mathrm{d}x \,\mathrm{d}y.
\end{aligned}$$

We split the integral according to $|x-y| \leq 1$ or > 1. The former contribution can be bounded by

$$[g]_{\mathcal{W}^{\beta,\infty}} \iint\limits_{|x-y|\leq 1} \frac{|f(y)|}{|x-y|^{d+\gamma-\beta}} \,\mathrm{d}x \,\mathrm{d}y \leq C [g]_{\mathcal{W}^{\gamma,\infty}} \|f\|_{L^1} ,$$

where we used that $\int_{B(0,1)} |x|^{-d+\beta-\gamma} \,\mathrm{d}x \leq C \int_0^1 r^{-1+\beta-\gamma} \,\mathrm{d}r < \infty$. The contribution with $|x-y| > 1$ can be bounded by

$$2 \|g\|_{L^\infty} \iint_{|x-y|>1} \frac{|f(y)|}{|x-y|^{d+\gamma}} \,\mathrm{d}x \,\mathrm{d}y \leq C \|g\|_{L^\infty} \int |f(y)| \,\mathrm{d}y = \|g\|_{L^\infty} \|f\|_{L^1} ,$$

where we used that $\int_{\mathbb{R}^d \setminus B(0,1)} |x|^{-\gamma-d} \,\mathrm{d}x \leq C \int_1^\infty r^{-1-\gamma} \,\mathrm{d}r < \infty$. □