

Weaves, webs and flows

Nic Freeman^{*1} and Jan Swart^{†2}

¹School of Mathematics and Statistics, University of Sheffield

²The Czech Academy of Sciences, Institute of Information Theory and Automation.

February 6, 2023

Abstract

We introduce *weaves*, which are random sets of non-crossing càdlàg paths that cover space-time $\overline{\mathbb{R}} \times \overline{\mathbb{R}}$. The Brownian web is one example of a weave, but a key feature of our work is that we do not assume that particle motions have any particular distribution. Rather, we present a general theory of the structure, characterization and weak convergence of weaves.

We show that the space of weaves has a particularly appealing geometry, involving a partition into equivalence classes under which each equivalence class contains a pair of distinguished objects known as a *web* and a *flow*. Webs are natural generalizations of the Brownian web and the flows provide pathwise representations of stochastic flows. Moreover, there is a natural partial order on the space of weaves, characterizing the efficiency with which paths cover space-time, under which webs are precisely minimal weaves and flows are precisely maximal weaves. This structure is key to establishing weak convergence criteria for general weaves, based on weak convergence of finite collections of particle motions.

MSC 2010. Primary: 60D05. Secondary: 60K99.

Acknowledgement. Work sponsored by GAČR grant 22-12790S. This work was initiated during the program Genealogies of Interacting Particle Systems at the IMS, NUS, Singapore.

^{*}n.p.freeman@sheffield.ac.uk

[†]swart@utia.cas.cz

Contents

1	Introduction	1
1.1	Outline of results	2
2	Results	5
2.1	The split real line	5
2.2	The path space	6
2.3	Crossing and ordering	9
2.4	Weaves, webs and flows	10
2.5	Discussion	14
3	Preliminaries	16
3.1	On crossing	16
3.2	On partial orders of sets	21
3.3	On compatibility of order and topology	21
4	Deterministic weaves	25
4.1	Weaves and the non-crossing property	26
4.2	Weaves of bi-infinite paths	30
4.3	Extensions of paths in weaves	33
4.4	The flow map	39
4.5	The web map	42
5	Random weaves	44
5.1	On measurability	44
5.2	On partial ordering of random weaves	46
5.3	Proof of Theorem 2.4.3	50
5.4	Proof of Theorem 2.4.4	50
5.5	Proof of Theorems 2.4.5 and 2.4.6	53
5.6	Proof of Theorem 2.4.8	54
5.7	Proof of Theorem 2.4.9	55
A	Appendices	56
A.1	On the Hausdorff metric	56
A.2	On the M1 topology of Π and $\mathcal{K}(\Pi)$	58
A.3	On measurability in $\mathcal{K}(\Pi)$	61
A.4	Proof of Lemma 2.5.1	64

1 Introduction

In this article we introduce a rich and natural class of objects that generalize the Brownian web. We call these objects *weaves*. Informally, a weave is a random set of non-crossing càdlàg paths, such that each point of space-time is almost surely touched by at least one path. The paths take values in $\overline{\mathbb{R}}$ and each path runs until time $+\infty$, but paths may begin at any point of space-time. It is important to note what is missing: we do *not* require that the paths follow any particular distribution (in the example of the Brownian web, they follow coalescing Brownian motions).

We will establish a framework for weak convergence (i.e. in law) of general weaves, akin to the modern theory of weak convergence for real valued stochastic processes. As the example of the Brownian web shows, individual weaves may display a rich internal geometry. We will see that *space of weaves* also has an interesting structure in its own right. This structure has major implications for the characterization of weaves, thus also for weak convergence.

We use the term *half-infinite* for paths that, after beginning anywhere within space-time, continue until time $+\infty$. If such a path begins at time $-\infty$ then it is said to be *bi-infinite*. Weaves consisting exclusively of bi-infinite paths provide natural pathwise representations of (sufficiently regular) stochastic flows, but can also represent more complicated structures of branching-coalescing paths. We refer to weaves of bi-infinite paths as *flows*, although formally we will first give a different definition and later show equivalence to this.

Stochastic flows have been studied for many decades, as detailed in the book of Kunita (1997). It is remarkable that, despite their long history, stochastic flows have struggled to give rise to a viable theory of their own weak convergence. A key problem is that stochastic flows have traditionally been given a ‘pointwise’ representation, where for each pair of times $-\infty < s < t < \infty$ a random function $X_{s,t} : \mathbb{R} \rightarrow \mathbb{R}$ represents the movement of particles during $[s, t]$. More precisely, $X_{s,t}(x)$ denotes the position at time t of the particle that, at time s , was at location x . The concept of a flow is therefore encapsulated by the consistency condition $X_{s,t} \stackrel{\text{a.s.}}{=} X_{s,u} \circ X_{u,t}$, required at deterministic times. This representation is analogous to the old-fashioned representation of a real valued stochastic process as an infinite family of random variables $(X_t)_{t \geq 0}$, where $X_t \in \mathbb{R}$ denotes the position of the particle at time $t \geq 0$.

The modern perspective is to view a stochastic process as a single random variable, whose value is a random path. Such a representation is known as a ‘pathwise’ representation. Skorohod (1956) introduced a suitable state space \mathcal{D} , whose elements are càdlàg paths, and the resulting theory is detailed within the now ubiquitous texts of Ethier and Kurtz (1986) and Billingsley (1995). From an analytical point of view it is far more convenient to work with convergence of one random càdlàg path, than with convergence of infinitely many \mathbb{R} valued random variables. To abstract this principle a little further, it is better to define a single random variable within a highly structured state space, than to work with infinitely many ‘smaller’ random variables in a more straightforward state space.

The same principle will apply to random sets of càdlàg paths, however such objects have not yet made an analogous transition – with the exception of the Brownian web. The present article seeks to remedy this situation. The Brownian web is a pathwise representation of the stochastic flow of Arratia (1979), in which particles perform independent Brownian motions until they meet, after which particles remain coalesced for all remaining time. Loosely, one such particle begins at each point of space-time.

The modern study of the Brownian web began with Tóth and Werner (1998), who were first to understand its rich internal structure. Based on this work, Fontes et al. (2004) represented the Brownian web as a (single) random variable whose value is a random set of continuous paths, and introduced the term *Brownian web*. In this representation they gave the first conditions for weak convergence to the Brownian web, based on the forwards-in-time motions of finite sets of

particles. A large body of literature has since emerged, leading to the refined criteria available in the survey of Schertzer et al. (2017). Close relatives of the Brownian web have been investigated in similar style and the Brownian web is understood to be the scaling limit of a large and diverse universality class.

The key to this success has been the availability of good criteria for characterization and weak convergence. Such criteria must strike a careful balance: a type of convergence that preserves less information is often easier to prove, and is more often true, but is also less meaningful. One possible approach, used by Berestycki et al. (2015) and Cannizzaro and Hairer (2021) for the case of continuous coalescing paths, is to map sets of paths to other objects (respectively, to sets of ‘tubes’ and real trees) in order to induce a topology that may be used as a basis for weak convergence.

In the present work we handle sets of càdlàg paths directly, in the style that has become popular within the literature of the Brownian web. We give criteria for characterization and convergence of general weaves, with no requirement that the particle motions follow any particular distribution. We must also introduce a suitable state space; a version of Skorohod’s space \mathcal{D} suitable for random sets of càdlàg paths begun at arbitrary points of space-time. The state space constructed by Fontes et al. (2004) is a subset of our own, with matching induced subspace topology.

Let us now briefly comment on the significance of webs. Our exploration of the space of weaves will uncover a natural partition into equivalence classes. Each equivalence class features two distinguished elements, one of which is a flow (as discussed above) and the other of which we will refer to as a *web*. We will see that the property of being a web is equivalent to what remains if one takes the usual definition of the Brownian web and *removes* the requirement that the particle motions have a particular distribution. Webs and flows are in bijective correspondence; moreover they are the extremal points, respectively minima and maxima, within a structure that we will shortly describe.

Within much of the literature on the Brownian web, the proofs rely heavily on the distribution of coalescing Brownian motions. Consequently our own arguments have little in common. Despite this, we remark that what is known about the Brownian web has been invaluable in writing the present article, and the Brownian web is a canonical example of a weave. In fact the majority of our results are new even in the special case of the Brownian web.

1.1 Outline of results

In Section 2 we will introduce our state space and, following that, give rigorous statements of our main results. Setting up the state space requires some significant work, so we will give here a non-rigorous presentation of our main results and the ideas that led to them.

We require that càdlàg paths are allowed to jump *at their initial times*. Naturally, this requires some supporting structure, which we delay for now and appeal instead to the readers intuition. Our concept of a càdlàg path is precisely equivalent to the classical càdlàg path $f : [t, \infty] \rightarrow \overline{\mathbb{R}}$ that is right-continuous with left limits, plus a possible jump at the initial time t . See Figure 1.1.1 for an example showing why this augmentation is necessary.

The theory of weak convergence of real valued stochastic processes is normally presented in Skorohod’s J1 topology. We require the (slightly coarser) Skorohod M1 topology. The reasons for this are rather technical, but roughly speaking our use of non-crossing paths makes it natural to consider jumps as part of the path, rather than as an empty region of space that the path jumps over. The former perspective corresponds to Skorohod’s M1 topology, the latter to J1. An example of càdlàg paths such that $f_n \rightarrow f$ in M1 but not in J1 is $f_n(t) = 0 \vee nt \wedge 1$ with limit $f(t) = \mathbb{1}_{\{t \geq 0\}}$, both defined for all $t \in \overline{\mathbb{R}}$.




Figure 1.1.1: In both figures, time runs upwards and the spatial axis is horizontal. In each figure a weave is depicted via solid lines, with the corresponding flow depicted via including dotted lines.

On the left: A weave \mathcal{A} , featuring a càdlàg path jumping at its initial time. Order the blue circles from bottom to top. Consider the particle motion f_n starting within the n^{th} blue circle, which then follows the red dotted line. The limiting path f is a trajectory that jumps rightwards at its initial time t . We require that weaves are closed sets and we require that weaves have particle motions; consequently we require that the path f exists.

On the right: A warning example related to Theorems 2.4.5 and 2.4.6. A weave \mathcal{A}_ϵ is depicted, along with the corresponding flow \mathcal{F}_ϵ of bi-infinite paths that do not cross \mathcal{A}_ϵ . Space-time points within the horizontal arrows, and forwards in time continuations thereof, are ramified. In the limit as $\epsilon \rightarrow 0$ the red area vanishes; the weaves \mathcal{A}_ϵ converge to a pervasive system of paths that contains crossing (jumping in both directions at t); the sequence of flows $(\mathcal{F}_\epsilon)_{\epsilon>0}$ are not relatively compact (due to paths that jump left-right-left between t and $t + \epsilon$); whilst the m -particle motions $\mathcal{A}_\epsilon|_z = \mathcal{F}_\epsilon|_z$ from finite sets z of non-ramified points converge to those of a weave (which contains only the leftwards jump at t).

A key insight from the Brownian web is that we should consider random *compact* sets of paths; we do so within a suitable version of the M1 topology. We say that a set of càdlàg paths is *pervasive* if each space-time point $z = (x, t)$ is contained within at least one path, including jumps.

A central concept is the *m-particle motion* of a weave. Loosely, if we choose a point $z = (x, t)$ in space-time, we may place a particle at the point within the weave and then watch how it moves, forwards in time. For most deterministic points of space-time (in fact, Lebesgue almost all) this operation is well defined and an almost surely unique forwards in time motion exists. This motion is a single random càdlàg path with initial time t . If we do the same for $m \in \mathbb{N}$ space-time points at once, then we obtain the *m-particle motion* of the weave.

It is clear a priori that a pair of càdlàg paths might cross each other. The meaning is clear for continuous paths and, for now, we appeal to the readers intuition. When we come to define crossing rigorously some clarification will be required, to handle cases where càdlàg paths jump over each other at their initial times. We say that a set of paths is *non-crossing* if none of its elements cross each other.

We are now in a position to describe our main results concerning weaves. Formally, a *weave* is a probability measure on M1-compact sets of half-infinite càdlàg paths, that is almost surely non-crossing and pervasive. We also use the term weave for a random variable with such a law. We adopt the convention of using calligraphic letters, such as \mathcal{A} and \mathcal{B} , for weaves. We also remind the reader that within a general partial order, a typical element might sit below anything from none to infinitely many maxima; similarly for minima.

1. There exists a natural partial order \preceq_d on the space of weaves. Informally, the statement $\mathcal{A} \preceq_d \mathcal{B}$ means: there exists a coupling under which \mathcal{B} covers space-time more efficiently than \mathcal{A} i.e. with fewer and/or longer paths.

By definition, we say that a weave is a *web* if it is minimal (within the space of all weaves) with respect to \preceq_d . We say that a weave is a *flow* if it is maximal.

2. The space of weaves is partitioned into equivalence classes, each of which has a flow as its unique maximal element and a web as its unique minimal element. We write this equivalence relation as $\mathcal{A} \sim \mathcal{B}$. Elements within the same equivalence class need not be \preceq_d -comparable.
3. There exists a pair of deterministic functions $\text{web}(\cdot)$ and $\text{flow}(\cdot)$ with the following properties.
 - (a) A weave \mathcal{A} is a web if and only if $\text{web}(\mathcal{A}) \stackrel{\text{a.s.}}{=} \mathcal{A}$.
 - (b) A weave \mathcal{A} is a flow if and only if $\text{flow}(\mathcal{A}) \stackrel{\text{a.s.}}{=} \mathcal{A}$.

Moreover, a weave is a flow if and only if it comprises exclusively of bi-infinite paths. Therefore, flows are natural pathwise representations of stochastic flows.

The web operation is a slight generalization of the operator $\mathcal{W} \mapsto \overline{\mathcal{W}(D)}$ that is familiar within the standard characterization of the Brownian web. By definition, $\text{flow}(\mathcal{A})$ is the set of bi-infinite càdlàg paths that do not cross \mathcal{A} . The map $\text{flow}(\cdot)$ is continuous, but $\text{web}(\cdot)$ is not.

4. Two weaves \mathcal{A} and \mathcal{B} satisfy $\mathcal{A} \sim \mathcal{B}$ if and only if the m -particle motions of \mathcal{A} and \mathcal{B} have the same distribution.
5. A weak limit of flows is necessarily a flow. Moreover, for flows, weak convergence is equivalent to tightness plus weak convergence of the m -particle motions.

An analogous result holds for general weaves, at the level of equivalence classes. Here we must include the assumption that weak limit points are non-crossing.

6. Each web \mathcal{W} has an associated dual web $\widehat{\mathcal{W}}$, of càdlàg paths running backwards in time, such that \mathcal{W} and $\widehat{\mathcal{W}}$ are almost surely non-crossing.

The triplet $(\mathcal{W}, \widehat{\mathcal{W}}, \mathcal{F})$ may be reconstructed from any single one of \mathcal{W} , $\widehat{\mathcal{W}}$ and \mathcal{F} . If any one of these three consists exclusively of continuous paths, then they all do.

Underpinning all of these results is a delicate operation that takes a half-infinite path within a weave and extends it, backwards in time, into a bi-infinite path, without inducing crossing and preserving càdlàgness. Moreover such extension may be done to all paths within a weave, without breaking the compactness, to obtain its corresponding flow. Note that weaves are by definition closed sets, so this operation does *not* involve taking a limit of suitable paths within the weave. Let us briefly describe what it does involve.

There is a partial order \subseteq on (individual) càdlàg paths, corresponding to the idea that $f \subseteq g$ if and only if the path f may be extended, forwards and/or backwards in time, to give g . Paths within weaves run until time $+\infty$, so for weaves only extension backwards in time is relevant. For a given weave \mathcal{A} , let \mathcal{A}_{\max} denote the set of maximal elements of (\mathcal{A}, \subseteq) .

It turns out that there is a natural bijection between Dedekind cuts of \mathcal{A}_{\max} and bi-infinite paths that do not cross \mathcal{A} . This relationship is reminiscent of Dedekind's famous construction of \mathbb{R} from \mathbb{Q} , but in our case the operation that connects \mathcal{A}_{\max} to $\text{flow}(\mathcal{A})$ is not a topological closure. Loosely, we may take a bi-infinite path h that does not cross \mathcal{A} , and the corresponding Dedekind cut is all paths $f \in \mathcal{A}_{\max}$ that lie strictly to the left of h . The inverse function of this correspondence is more complicated to define and we do not attempt a description at this

point. To extend half-infinite paths backwards in time, we note that a Dedekind cut may be constructed in the same way from any càdlàg path (not necessarily bi-infinite) that does not cross \mathcal{A} , and then use that inverse function to produce a corresponding bi-infinite path.

The proof of this relationship between half-infinite and bi-infinite paths relies on delicate analysis. It requires a formulation of càdlàg paths where potential jumps at the initial time are an integral part of the path, rather than an afterthought to the otherwise classical definition. We introduce such a formulation in Section 2.1, followed by a description of our state space in Section 2.2. We then introduce key notation concerning crossing and ordering of paths in Section 2.3, at which point we are able to give a rigorous presentation of our main results in Section 2.4. We discuss connections to the Brownian web in Section 2.5.1, and connections to related state spaces in Sections 2.5.2.

The proofs appear in Sections 3-5. In Section 3 we set up machinery to work with general càdlàg paths, and with jumps at their initial times. Section 4 treats the internal structure of weaves and the implications thereof for the space of weaves, which is best studied (initially) in a deterministic context. This includes the key result on path extension. In Section 5 we give the proofs of our main results.

2 Results

2.1 The split real line

A function is said to be *càdlàg*, from the French ‘continue à droite, limite à gauche’, if it is right-continuous with left limits. Kolmogorov (1956) observed that a real càdlàg function together with its left-continuous modification can be viewed as a continuous function on a peculiar topological space, introduced by Alexandroff and Urysohn (1929). This will provide an elegant formulation of our results, as well as being a necessary component of more technical proofs. We give here a brief introduction to this space.

Let $\overline{\mathbb{R}} := [-\infty, \infty]$ denote the extended real line. By definition, for any subset $I \subseteq \overline{\mathbb{R}}$, we let

$$I_{\star} = \{(t, \star); t \in I \text{ and } \star \in \{-, +\}\}.$$

We will almost always write $t\star$ in place of the formal notation (t, \star) . We call \mathbb{R}_{\star} the *split real line* and $\overline{\mathbb{R}}_{\star}$ the *extended split real line*. Loosely, to construct \mathbb{R}_{\star} from \mathbb{R} , each $t \in \mathbb{R}$ has been split into two parts, a left part $t-$ and a right part $t+$. We equip $\overline{\mathbb{R}}_{\star}$ with the lexicographic order, from left to right, that is $t_1\star_1 < t_2\star_2$ if and only if either $t_1 < t_2$ or both $t_1 = t_2$ and $\star_1 = -, \star_2 = +$. We use notation for intervals in $\overline{\mathbb{R}}_{\star}$ similar to the usual notation for the extended real line:

$$(t_1\star_1, t_2\star_2) = \{t\star \in \overline{\mathbb{R}}_{\star} : t_1\star_1 < t\star < t_2\star_2\} \quad (2.1)$$

$$[t_1\star_1, t_2\star_2] = \{t\star \in \overline{\mathbb{R}}_{\star} : t_1\star_1 \leq t\star \leq t_2\star_2\}, \quad (2.2)$$

and analogously for half-open intervals such as $(t_1\star_1, t_2\star_2]$ or $[t_1\star_1, t_2\star_2)$. Note that there is some redundancy in this notation since, for example, $(s-, t+) = [s+, t-]$. We say that a set $A \subseteq \mathbb{R}_{\star}$ is *bounded* if $A \subseteq [-T, T]_{\star}$ for some $T < \infty$.

We equip \mathbb{R}_{\star} and $\overline{\mathbb{R}}_{\star}$ with the *order topology*. Recall that, in a totally ordered space $(S, <)$, the order topology is generated by the open intervals $(a, b) = \{x \in S; a < x < b\}$ where $a < b$. The order topology on \mathbb{R} thus coincides with the usual Euclidean topology. The following lemma records all that we need to know about the order topology on \mathbb{R}_{\star} . Parts 1 and 4 appear respectively as Lemma 2.1 and Proposition 2.3 in Freeman and Swart (2023). Parts 2 and 3 are straightforward consequences of part 1.

Lemma 2.1.1 *The following hold.*

1. A sequence $t_n \star_n$ converges to the limit $t+$ (resp. $t-$) if and only if $t_n \rightarrow t$ in \mathbb{R} and $t_n \star_n \geq t+$ (resp. $t_n \star_n \leq t-$) for all n sufficiently large.
2. Intervals of the form $(t_1 \star_1, t_2 \star_2)$ are open and intervals of the form $[t_1 \star_1, t_2 \star_2]$ are closed.
3. The quotient of \mathbb{R}_s by the relation $s \star \sim t \star \Leftrightarrow s = t$ is homeomorphic to \mathbb{R} .
4. The space \mathbb{R}_s is a Hausdorff topological space. It is separable but not metrisable. For $C \subseteq \mathbb{R}_s$, the following three statements are equivalent: (i) C is compact; (ii) C is sequentially compact; (iii) C is closed and bounded.

The topology on $\overline{\mathbb{R}}_s$ permits an elegant description of càdlàg functions on $\overline{\mathbb{R}}$, developed in Section 2.2 of Freeman and Swart (2023). We require this characterization only for the case of real càdlàg functions on closed intervals, as follows. Let $I = [a, b] \subset \overline{\mathbb{R}}$ be a closed interval and let $f : I_s \rightarrow \overline{\mathbb{R}}$ be a function. The following statements are equivalent:

1. f is continuous with respect to the topology on I_s , as a subset of \mathbb{R}_s ;
2. the function $t \mapsto f(t+)$ defined from $[a, b] \mapsto \overline{\mathbb{R}}$ is right continuous with left limits, and $t \mapsto f(t-)$ defined from $(a, b] \rightarrow \mathbb{R}$ is its left continuous modification;
3. the function $t \mapsto f(t-)$ defined from $[a, b] \mapsto \overline{\mathbb{R}}$ is left continuous with right limits, and $t \mapsto f(t+)$ defined from $[a, b) \rightarrow \mathbb{R}$ is its right continuous modification.

Definition 2.1.2 We refer to a function $f : [a, b]_s \rightarrow \overline{\mathbb{R}}$ satisfying (any of) these criteria as a *càdlàg path*.

Definition 2.1.2 is a minor extension of the classical notion of a càdlàg function on $[a, b] \subseteq \overline{\mathbb{R}}$. Specifically, we attach a formal meaning and value to the ‘left limit’ at $a-$, which is absent in the classical definition. It may take any value, which is to say that the value of $f(a-)$ is not restricted by the values of $f(t \star)$ for $t \star > a-$. This introduces the possibility that $f(a-) \neq f(a+)$, corresponding to a jump at the initial time.

Given a càdlàg path f with domain $[a, b]_s$ we write $I(f) = [a, b] \subseteq \overline{\mathbb{R}}$ and $I(f)_s = [a, b]_s \subseteq \overline{\mathbb{R}}_s$. We write $\sigma_f = a$ for the *initial time* and $\tau_f = b$ for the *final time* of f . Similarly, we call $(f(\sigma_f-), \sigma_f)$ and $(f(\tau_f+), \tau_f)$ respectively the *intial* and *final points* of f . We say that f *begins* and its initial point and *ends* at its final point.

We say that f makes a *jump* at $t \in [\sigma_f, \tau_f]$ if $f(t-) \neq f(t+)$. The jump is said to be *to the left* if $f(t+) < f(t-)$ and *to the right* if $f(t+) > f(t-)$. As mentioned, càdlàg paths may jump at their initial and final times or at any time in between. The number of such jumps is at most countable. If $f(t-) = f(t+)$ then we say that f is *continuous* at t , in which case (and only in this case) we write $f(t) = f(t-) = f(t+)$.

2.2 The path space

In this section we introduce the space Π , whose elements are càdlàg paths defined on closed intervals of $\overline{\mathbb{R}}_s$, and the space $\mathcal{K}(\Pi)$, whose elements are compact subsets of Π . We will refer to Freeman and Swart (2023) which considers a more general setup but also acts as a companion paper providing the topological basis for the present article. Let

$$\Pi = \{f : I_s \rightarrow \overline{\mathbb{R}}; f \text{ is a càdlàg path and } I \subseteq \overline{\mathbb{R}} \text{ is a closed interval}\}. \quad (2.3)$$

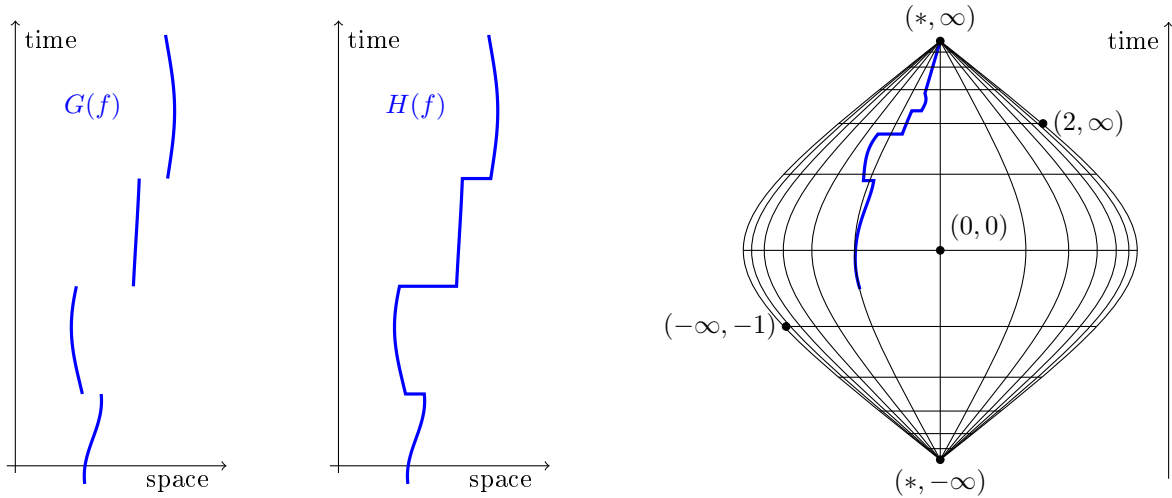


Figure 2.2.1: On the left, the closed graph $G(f)$ and interpolated graph $H(f)$ of a path f . On the right, the compactification \mathbb{R}_c^2 of $\mathbb{R} \times \mathbb{R}$ with the interpolated graph of f , with some points marked for convenience. In Freeman and Swart (2023) \mathbb{R}_c^2 is referred to as a *squeezed space*.

We regard two elements $f, g \in \Pi$ as equivalent if they have the same values outside of times $\pm\infty$. Formally, define the equivalence relation

$$f \stackrel{\Pi}{\sim} g \Leftrightarrow I(f) = I(g) \text{ and } f(t\pm) = g(t\pm) \text{ for all } t \in I(f) \cap \mathbb{R}, \quad (2.4)$$

and work implicitly with the resulting equivalence classes of Π . We abuse notation slightly by continuing to write $f \in \Pi$ for a càdlàg path, but including the notational convention¹ that $f(t\star) = *$ whenever $t = \pm\infty$.

Our main results require Skorohod's M1 topology on Π , which we now introduce. We will introduce the J1 topology at the same time, as it is more widely used and the reader may wish to make a comparison. We define the *closed graph* $G(f)$ and *interpolated graph* $H(f)$ of a càdlàg path $f \in \Pi$ as

$$\begin{aligned} G(f) &= \{(x, t) \in \overline{\mathbb{R}}^2; t \in I(f), x \in \{f(t-), f(t+)\}\}, \\ H(f) &= \{(x, t) \in \overline{\mathbb{R}}^2; t \in I(f), x \in [f(t-), f(t+)\}\}, \end{aligned}$$

where we use the convention $[s, t] := [s \wedge t, s \vee t]$ for $s, t \in \overline{\mathbb{R}}$. See Figure 2.2.1 for a picture displaying the difference between $G(f)$ and $H(f)$: at times $t \in \mathbb{R}$ when the path makes a jump, the line segments between $(f(t-), t)$ and $(f(t+), t)$ are drawn in $H(f)$ but not in $G(f)$.

The reason for (2.4) is that we intend to treat $G(f)$ and $H(f)$ as compact subsets of a suitable space, which in turn will allow us to describe the J1 and M1 topologies. With this in mind, we define

$$\mathbb{R}_c^2 := (\overline{\mathbb{R}} \times \mathbb{R}) \cup \{(*, -\infty), (*, \infty)\}$$

and equip \mathbb{R}_c^2 with a metrisable topology such that the induced subspace topology on $\overline{\mathbb{R}} \times \mathbb{R}$ is the product topology and, as $n \rightarrow \infty$,

$$(x_n, t_n) \rightarrow (*, \pm\infty) \text{ if and only if } t_n \rightarrow \pm\infty.$$

¹To be precise: Freeman and Swart (2023) uses a slightly different compactification procedure to (2.4), which allows the domain of càdlàg paths to be non-interval sets and defines $\overline{\mathbb{R}}_s$ to be a two-point compactification of \mathbb{R}_s . See Section 2.1 within Freeman and Swart (2023) for details. The space Π from the present article is denoted there by Π^1 .

Equation (2.26) of Freeman and Swart (2023) gives an explicit metric $d_{\mathbb{R}_c^2}$ with these properties, under which \mathbb{R}_c^2 is compact. See Figure 2.2.1 for an illustration of $G(f)$, $H(f)$ and the compactification. We endow \mathbb{R}_c^2 with the two-dimensional Lebesgue measure on $\overline{\mathbb{R}} \times \mathbb{R} \subseteq \mathbb{R}_c^2$, placing zero mass at $(*, \pm\infty)$.

Let $f \in \Pi$. There is a natural total order on both $G(f)$ and $H(f)$, which we denote by \sqsubseteq . For $G(f)$ this order is given by $(f(t_1\star_1), t_1) \sqsubseteq (f(t_2\star_2), t_2)$ whenever $t_1\star_1 \leq t_2\star_2$. For $H(f)$ it requires a little more care: we say that $(x_1, t_1) \sqsubseteq (x_2, t_2)$ whenever $t_1 < t_2$, or if $t_1 = t_2$ and x_1 is non-strictly closer to $f(t_1-)$ than x_2 .

Informally, the J1 topology on Π corresponds to convergence of closed graphs, and the M1 topology to convergence of interpolated graphs, with the caveat that (in both cases) the total order \sqsubseteq is preserved by the convergence. Our next step is to formalize this intuition and give a definition of the two topologies.

For a metric space (M, d_M) , let $\mathcal{K}(M)$ denote the space of all nonempty compact subsets of M , equipped with the Hausdorff metric induced by d_M . See Appendix A.1 for a brief introduction to the Hausdorff metric. We define the *second order closed graph* $G^{(2)}(f)$ and the *second order interpolated graph* $H^{(2)}(f)$ of a path $f \in \Pi$ to be

$$\begin{aligned} G^{(2)}(f) &= \{(z_1, z_2); z_i \in G(f) \text{ and } z_1 \sqsubseteq z_2\}, \\ H^{(2)}(f) &= \{(z_1, z_2); z_i \in H(f) \text{ and } z_1 \sqsubseteq z_2\}, \end{aligned}$$

where \sqsubseteq is as defined above. Lemma 3.1 of Freeman and Swart (2023) gives that the sets $G(f)$ and $H(f)$ are compact subsets of \mathbb{R}_c^2 , moreover the sets $G^{(2)}(f)$ and $H^{(2)}(f)$ are compact subsets of $(\mathbb{R}_c^2)^2$. Note that $G^{(2)}(f)$ and $H^{(2)}(f)$ preserve information about the total order \sqsubseteq , whereas $G(f)$ and $H(f)$ do not. It can be seen that $G^{(2)}(f_n) \rightarrow G^{(2)}(f)$ implies $G(f_n) \rightarrow G(f)$, but the converse is not true. Similarly for $H^{(2)}$ and H . If $G(f_n) \rightarrow G(f)$ then $H(f_n) \rightarrow H(f)$, but the converse is not true. Similarly for $G^{(2)}$ and $H^{(2)}$.

Proposition 2.2.1 *In each case, under the metric listed Π is a Polish space.*

1. The J1 topology: $d_{J1}(f, g) = d_{\mathcal{K}((\mathbb{R}_c^2)^2)}(G^{(2)}(f), G^{(2)}(g))$.
2. The M1 topology: $d_{M1}(f, g) = d_{\mathcal{K}((\mathbb{R}_c^2)^2)}(H^{(2)}(f), H^{(2)}(g))$.

We mention also that $d_{J2}(f, g) = d_{\mathcal{K}(\mathbb{R}_c^2)}(G(f), G(g))$ and $d_{M2}(f, g) = d_{\mathcal{K}(\mathbb{R}_c^2)}(H(f), H(g))$ respectively correspond to Skorohod's J2 and M2 topologies. In Freeman and Swart (2023) it is shown that each such metric is equivalent to the corresponding classical Skorohod metric on $\mathcal{D}_{[a,b]} = \{f \in \Pi; I(f) = [a, b], f(a-) = f(a+)\}$, for $a, b \in \mathbb{R}$. See Sections 2.4, 3.2 and 3.4 of that article for details and proofs of these facts.

We now specialize to the case that is relevant to the present article:

from now on the space Π is (implicitly) equipped with the M1 topology.

The same applies to subsets of Π . We write $d_\Pi = d_{M1}$, generating the M1 topology on Π , but we remark that the metric space (Π, d_Π) is not complete. In Appendix A.2 we collate together some results from Freeman and Swart (2023) concerning the M1 topology, including criteria for relative compactness and tightness.

Let $\Pi^\uparrow = \{f \in \Pi; \tau_f = \infty\}$, $\Pi^\downarrow = \{f \in \Pi; \sigma_f = -\infty\}$, and $\Pi^\dagger = \Pi^\uparrow \cap \Pi^\downarrow$ be the subspaces of (respectively) forwards and backwards *half-infinite* paths, and *bi-infinite* paths. Note that Π and Π^\dagger are both symmetric under time reversal. We write $\Pi_c = \{f \in \Pi; f(t-) = f(t+)\}$ for all $t \in I(f)\}$ for the subspace of *continuous* paths. We write $\Pi_c^\dagger := \Pi^\dagger \cap \Pi_c$ and so on. Informally, the induced topology on Π_c may be described as convergence of starting and final times plus locally

uniform convergence of continuous paths. In fact Π_c^\uparrow is the state space introduced by Fontes et al. (2004) for the Brownian web. See Proposition 3.4 of Freeman and Swart (2023) for details.

Our main results will concern systems of half-infinite càdlàg paths and we will tend to state results forwards in time i.e. we will mostly work in Π^\uparrow or its subspace Π^\downarrow . We require Π^\downarrow only for results concerning duality. We write $\mathcal{K}(\Pi)$ for the metric space of compact subsets of Π , where the underlying metric on Π comes from Proposition 2.2.1. It is easily seen from Proposition 2.2.1 that Π^\uparrow , Π^\downarrow and Π^\uparrow are closed subsets of Π , from which it follows that $\mathcal{K}(\Pi^\uparrow)$, $\mathcal{K}(\Pi^\downarrow)$ and $\mathcal{K}(\Pi^\uparrow)$ are closed subsets of $\mathcal{K}(\Pi)$.

2.3 Crossing and ordering

We now introduce notation and terminology associated to Π and Π^\uparrow . We say that a càdlàg path g *extends* a càdlàg path f if $H^{(2)}(f) \subseteq H^{(2)}(g)$. We note that, in all but the trivial case for which $\sigma_f = \tau_f$ or $\sigma_g = \tau_g$ this reduces to the more intuitive condition $H(f) \subseteq H(g)$. The point is that the total ordering of $H(f)$ (see Section 2.2) should coincide with its induced order as a subset of $H(g)$.

We write $f \subseteq g$ to denote that g extends f . It is easily seen that \subseteq is a partial order on Π and we write the corresponding strict order relation as \subset . For $A \subseteq \Pi$ we write

$$A_{\max} = \{g \in A; \text{ for all } f \in A, \text{ if } g \subseteq f \text{ then } g = f\} \quad (2.5)$$

$$A_\uparrow = \{g \in \Pi^\uparrow; g \subseteq f \text{ for some } f \in A\}. \quad (2.6)$$

In words, A_{\max} denotes the set of maximal elements of (A, \subseteq) , whilst A_\uparrow denotes the set of half-infinite paths that may be extended to some $f \in A$.

For sets of paths $A, B \subseteq \Pi^\uparrow$ we define the relation

$$A \preceq B \quad \Leftrightarrow \quad A_\uparrow \cap B \subseteq A \subseteq B_\uparrow, \quad (2.7)$$

which, as a consequence of Lemma 3.2.2, is a partial order on (the set of) subsets of Π^\uparrow . The corresponding strict order relation is written \prec . The intuition behind (2.7) is one of efficient covering of space-time: loosely $A \prec B$ means that B covers more of space-time using longer and/or fewer paths than A .

Note that if A and B are random variables on the same probability space then we can make sense of the event $\{A \preceq B\}$ via (2.7). If A and B are $\mathcal{K}(\Pi^\uparrow)$ valued random variables without an implicit coupling then, with mild abuse of notation, we further extend \preceq by writing $A \preceq_d B$ if and only if there exists a coupling of A and B such that $\mathbb{P}[A \preceq B] = 1$. We write $A \prec_d B$ if $A \preceq_d B$ where A and B do not have the same marginal distributions. In Lemma 5.2.3 we show that \preceq_d is a partial order on the space of (laws of) $\mathcal{K}(\Pi^\uparrow)$ valued random variables.

Definition 2.3.1 We say that paths $f, g \in \Pi$ are *non-crossing* if there exists paths $f', g' \in \Pi^\uparrow$ such that $f \subseteq f'$, $g \subseteq g'$ and $f(t_\star) \leq g(t_\star)$ for all $t_\star \in \overline{\mathbb{R}}_+$.

The precise format of Definition 2.3.1 is motivated by the complication that a pair of càdlàg paths may share the same initial (or final) time and might both jump at this time, with perhaps overlapping jumps. The reader may wish to glance forward at Figure 3.1.1 which depicts some of these complications. We say that a set of paths $A \subseteq \Pi$ is non-crossing if all pairs of elements of A are non-crossing. We say that A and B are non-crossing if $A \cup B$ is non-crossing.

If $f \in \Pi$ and $z \in H(f)$, then we say that f *passes through* the space-time point z . For $f \in \Pi^\uparrow$, if f passes through $z \in \mathbb{R}_c^2$ then we define the *restriction*

$$f|_z \in \Pi^\uparrow \text{ to be the unique } g \subseteq f \text{ such that } (g(\sigma_g-), \sigma_g) = z. \quad (2.8)$$

More generally, we say that f passes through a set $D \subseteq \mathbb{R}_c^2$ if f passes through some point $(x, t) \in D$. For an (unordered) set $D \subseteq \mathbb{R}_c^2$ and $A \subseteq \Pi^\uparrow$ we write

$$A(D) = \{f \in A : f \text{ passes through } D\} \quad (2.9)$$

for the set of paths in \mathcal{A} that pass through D . For convenience, for $z \in \mathbb{R}_c^2$ we write $A(z) = A(\{z\})$. In the same vein we define

$$A|_D = \{f|_z \in \Pi^\uparrow ; f \in \mathcal{A}(z) \text{ and } z \in D\} \quad (2.10)$$

for the set of paths in A that pass through some $z \in B$, with the part prior to z removed. For $z \in \mathbb{R}_c^2$ we write $A|_z = A|_{\{z\}}$.

2.4 Weaves, webs and flows

We are interested in systems of non-crossing paths that touch every point of space-time. More rigorously, we say that $A \subseteq \Pi^\uparrow$ is *pervasive* if $A(z) \neq \emptyset$ for all $z \in \mathbb{R}_c^2$. We remark that, as a consequence of Lemma A.2.2, if $A \in \mathcal{K}(\Pi)$ then it suffices to check that $A(z) \neq \emptyset$ on a dense subset of $z \in \mathbb{R}_c^2$. The key objects studied within the present article are as follows.

Definition 2.4.1 A *weave* is the law of a $\mathcal{K}(\Pi^\uparrow)$ valued random variable that is almost surely pervasive and non-crossing. Let \mathscr{W} denote the set of weaves. A weave that is a minimal element of \mathscr{W} with respect to \preceq_d is known as a *web*. A weave that is a maximal element of \mathscr{W} with respect to \preceq_d is known as a *flow*.

A weave is a probability measure on $\mathcal{K}(\Pi)$, however we mildly abuse terminology in the usual way (c.f. ‘a’ Brownian motion) by saying that a $\mathcal{K}(\Pi^\uparrow)$ valued random variable is a weave if its law satisfies Definition 2.4.1. Similarly for webs and flows. We write

$$\mathscr{W}_{\text{det}} = \{\mathcal{A} \in \mathcal{K}(\Pi^\uparrow) ; \mathcal{A} \text{ is non-crossing and pervasive}\}. \quad (2.11)$$

If \mathcal{A} is a weave then $\mathbb{P}[\mathcal{A} \in \mathscr{W}_{\text{det}}] = 1$. Elements of \mathscr{W}_{det} are said to be *deterministic weaves*. Although \mathscr{W}_{det} is not formally a subset of \mathscr{W} , it may be viewed as such by identifying $\mathcal{A} \in \mathscr{W}_{\text{det}}$ with the probability measure that is a point-mass at \mathcal{A} . The following concept plays a central role for both deterministic and random weaves.

Definition 2.4.2 Let \mathcal{A} be a weave and let $z \in \mathbb{R}_c^2$ be a possibly random point of space-time. We say that z is a *ramification point* of \mathcal{A} if there exists $f, g \in \mathcal{A}(z)$ such that neither $f \subseteq g$ nor $g \subseteq f$. Otherwise, $z \in \mathbb{R}_c^2$ is said to be *non-ramified* in \mathcal{A} .

If z is non-ramified in \mathcal{A} with $f, g \in \mathcal{A}(z)$ then $f \subseteq g$ or $g \subseteq f$. Loosely, ramification points capture where weaves display atypical path behaviour, for example branching or coalescing of paths, or perhaps both. We stress that ‘ z is non-ramified in \mathcal{A} ’ is an event, with some associated probability, and not a deterministic statement. If it is clear from the context which weave is meant then we may simply say that $z \in \mathbb{R}_c^2$ is non-ramified. We say that $D \subseteq \mathbb{R}_c^2$ is non-ramified if all $z \in D$ are non-ramified.

A recurring theme in our results is that behaviour at non-ramified points determines the full behaviour of the weave. This suggests that non-ramified points should be plentiful. In Lemma 5.4.1 we show that for any weave \mathcal{A} the deterministic set

$$\{z \in \mathbb{R}_c^2 ; \mathbb{P}[z \text{ is ramified in } \mathcal{A}] > 0\}$$

has zero Lebesgue measure. Hence if a random point $z \in \mathbb{R}_c^2$ is sampled independently of \mathcal{A} , according to some non-atomic law, then z is almost surely non-ramified in \mathcal{A} .

Let $A \in \mathcal{K}(\Pi)$ and let $D \subseteq \mathbb{R}_c^2$. We define a key pair of deterministic operations as follows:

$$\text{web}_D(A) = \overline{(A|_D)}_{\uparrow} \quad (2.12)$$

$$\text{flow}(A) = \{f \in \Pi^{\uparrow}; f \text{ does not cross } A\}. \quad (2.13)$$

Let us briefly comment on the $\text{web}_D(\cdot)$ operation. The use of $\overline{(\cdot)}$ denotes closure in $\mathcal{K}(\Pi^{\uparrow})$. In Lemma 4.5.3 we will see that, for $\mathcal{A} \in \mathcal{W}_{\text{det}}$, the value of $\text{web}_D(\mathcal{A})$ does not depend upon the choice of dense and non-ramified $D \subseteq \mathbb{R}^2$. Thus (2.12) defines a deterministic function $\text{web}(\cdot)$ with domain \mathcal{W}_{det} , which we write without explicit specification of D .

We are now ready to state our first main result. It shows that extremal points of (\mathcal{W}, \preceq_d) may be characterized as fixed points of the web and flow operators. This leads to a particularly nice description of the structure of \mathcal{W} .

Theorem 2.4.3 *Let \mathcal{A} be a weave.*

1. *The following are equivalent: (a) \mathcal{A} is a web; (b) $\mathcal{A} \stackrel{\text{a.s.}}{=} \text{web}(\mathcal{A})$.*
2. *The following are equivalent: (a) \mathcal{A} is a flow; (b) $\mathcal{A} \stackrel{\text{a.s.}}{=} \text{flow}(\mathcal{A})$; (c) $\mathcal{A} \subseteq \Pi^{\uparrow}$ almost surely.*
3. *Almost surely, $\text{web}(\mathcal{A}) \preceq \mathcal{A} \preceq \text{flow}(\mathcal{A})$.*
4. *There exists a unique (in distribution) web \mathcal{W} and a unique flow \mathcal{F} such that $\mathcal{W} \preceq_d \mathcal{A} \preceq_d \mathcal{F}$, given by $\mathcal{W} \stackrel{d}{=} \text{web}(\mathcal{A})$ and $\mathcal{F} \stackrel{d}{=} \text{flow}(\mathcal{A})$.*

Thus, the space of weaves is partitioned by the equivalence relation

$$\mathcal{A} \sim \mathcal{B} \Leftrightarrow \text{web}(\mathcal{A}) \stackrel{d}{=} \text{web}(\mathcal{B}) \Leftrightarrow \text{flow}(\mathcal{A}) \stackrel{d}{=} \text{flow}(\mathcal{B}), \quad (2.14)$$

under which each equivalence class has a web as its unique minimal element, and a flow as its unique maximal element.

The relation \sim can also be characterized using finite collections of particle motions, for which we now introduce formal notation. Consider a weave \mathcal{A} , a flow \mathcal{F} and a non-ramified point $z \in \mathbb{R}_c^2$. The set $\mathcal{A}|_z$ contains a single path, which begins at z . Similarly, $\mathcal{F}(z)$ contains a single path, which passes through z . This makes it natural to define versions of (2.9) and (2.10) specialized to ordered sets of non-ramified points.

Let $m \in \mathbb{N}$. Given a weave \mathcal{A} and an almost surely non-ramified $\mathbf{z} = (z_i)_{i=1}^m \in (\mathbb{R}_c^2)^m$, we write $\mathcal{A}|_{\mathbf{z}} = (f_1, \dots, f_m)$ where $\{f_i\} \stackrel{\text{a.s.}}{=} \mathcal{A}|_{z_i}$. Similarly, given a flow \mathcal{F} and an almost surely non-ramified $\mathbf{z} = (z_i)_{i=1}^m \in (\mathbb{R}_c^2)^m$, we write $\mathcal{F}(\mathbf{z}) = (f'_1, \dots, f'_m)$ where $f'_i \in \Pi^{\uparrow}$ is the almost surely unique element of $\mathcal{F}(z_i)$. We say that $\mathcal{A}|_{\mathbf{z}}$ is the (*forwards in time*) m -particle motion of \mathcal{A} from \mathbf{z} . They are defined up to almost sure equivalence.

Loosely, the relation \sim also characterizes when two weaves have the same forwards in time m -particle motions, in distribution. With this in mind, we will need to make statements featuring multiple weaves that concern non-ramified points. We adopt the implicit convention that non-ramification is with respect to all weaves featured in the corresponding statement. We are now ready to state our second main result.

Theorem 2.4.4 *Let \mathcal{A}, \mathcal{B} be weaves.*

1. *Suppose that $\mathcal{A} \sim \mathcal{B}$. If $\mathbf{z} \subseteq \mathbb{R}_c^2$ is finite and almost surely non-ramified then $\mathcal{A}|_{\mathbf{z}} \stackrel{d}{=} \mathcal{B}|_{\mathbf{z}}$.*

2. The following are equivalent:

- (a) $\mathcal{A} \sim \mathcal{B}$;
- (b) there exists a coupling of \mathcal{A} and \mathcal{B} such that $\mathcal{A} \cup \mathcal{B}$ is almost surely non-crossing;
- (c) there exists a (deterministic) countable dense $D \subseteq \mathbb{R}^2$, which is almost surely non-ramified, such that $\mathcal{A}|_z \stackrel{d}{=} \mathcal{B}|_z$ for all finite $z \subseteq D$.

Theorems 2.4.3 and 2.4.4 lead towards an appealing limit theory for weaves, which we now develop. We denote weak convergence (i.e. convergence in law) by \xrightarrow{d} . We saw in Section 2.2 that $\mathcal{K}(\Pi^\uparrow)$ is a Polish space, thus weak convergence of $\mathcal{K}(\Pi^\uparrow)$ valued random variables, or equivalently of probability measures on $\mathcal{K}(\Pi)$, is defined in the standard way e.g. as in Section 3.3 of Ethier and Kurtz (1986). This provides a natural sense in which to consider weak convergence of weaves.

Our next theorem shows that weak convergence of flows is equivalent to tightness plus weak convergence of m -particle motions, and explores the same statement in the context of equivalence classes of weaves. Note that convergence of m -particle motions is a statement about weak convergence of $\overline{\mathbb{R}}^m$ valued stochastic processes, which is within the realms of the classical theory in e.g. Ethier and Kurtz (1986).

Theorem 2.4.5 *Let $\mathcal{F}_n, \mathcal{F}$ be flows.*

1. If $\mathcal{F}_n \xrightarrow{d} \mathcal{F}$ then for any $m \in \mathbb{N}$ and non-ramified $z \in (\mathbb{R}_c^2)^m$ we have $\mathcal{F}_n(z) \xrightarrow{d} \mathcal{F}(z)$.
2. Any weak limit point of (\mathcal{F}_n) is a flow. If (\mathcal{F}_n) is tight and for any $m \in \mathbb{N}$ and almost surely non-ramified $z \in (\mathbb{R}^2)^m$ we have $\mathcal{F}_n|_z \xrightarrow{d} \mathcal{F}|_z$, then $\mathcal{F}_n \xrightarrow{d} \mathcal{F}$.
3. Let $\mathcal{A}_n, \mathcal{A}$ be weaves with $\mathcal{A}_n \sim \mathcal{F}_n$ and $\mathcal{A} \sim \mathcal{F}$. If $\mathcal{A}_n \xrightarrow{d} \mathcal{A}$ then $\mathcal{F}_n \xrightarrow{d} \mathcal{F}$. Conversely, if $\mathcal{F}_n \xrightarrow{d} \mathcal{F}$ then any weak limit point \mathcal{B} of (\mathcal{A}_n) is a weave and satisfies $\mathcal{B} \sim \mathcal{F}$.

Lemma 4.4.3, shows that the function $\text{flow}(\cdot)$ is continuous on \mathscr{W}_{det} , which gives the forwards implication of part 3 of Theorem 2.4.5. The function $\text{web}(\cdot)$ is not continuous, as shown by example in Figure 2.5.1, which depicts a sequence of webs (\mathcal{W}_n) converging to a weave \mathcal{A} that is neither a web nor a flow. This suggests that, for purposes of convergence, flows are a more natural representative element of their equivalence class than webs.

Let us now give analogues for general weaves of parts 1 and 2 of Theorem 2.4.5. In part 3 of Theorem 2.4.5 flows provide an overarching structure for weaves in which the non-crossing property is preserved by taking limits of paths. The non-crossing property is preserved when taking limits of bi-infinite paths, but is not necessarily preserved in limits of half-infinite paths. Consequently, if we wish to establish convergence of weaves but also wish to avoid handling their associated flows, then it becomes necessary to check that limit points are non-crossing. See Figure 1.1.1 for a related warning example.

Theorem 2.4.6 *Let $\mathcal{A}_n, \mathcal{A}$ be weaves.*

1. If $\mathcal{A}_n \xrightarrow{d} \mathcal{A}$ then for any $m \in \mathbb{N}$ and non-ramified $z \in (\mathbb{R}_c^2)^m$ we have $\mathcal{A}_n|_z \xrightarrow{d} \mathcal{A}|_z$.
2. If a weak limit point \mathcal{B} of (\mathcal{A}_n) is non-crossing then \mathcal{B} is a weave. If, additionally, for any $m \in \mathbb{N}$ and almost surely non-ramified $z \in (\mathbb{R}^2)^m$ we have $\mathcal{A}_n|_z \xrightarrow{d} \mathcal{A}|_z$, then $\mathcal{A} \sim \mathcal{B}$.

Our next result concerns time-reversed duality, for which we must introduce some more notation. Given $f \in \Pi^\uparrow$, define $f^\circ \in \Pi^\downarrow$ by $f^\circ(t\pm) = -f(-t\mp)$. This operation, which corresponds to a rotation of space by 180 degrees, is applied pointwise to sets of paths as $\mathcal{A}^\circ = \{f^\circ; f \in \mathcal{A}\}$. Clearly $(\mathcal{A}^\circ)^\circ = \mathcal{A}$. Note that \cdot° is an automorphism of Π and Π^\downarrow , and that $(\Pi^\uparrow)^\circ = \Pi^\downarrow$. Proposition A.2.1 implies that $A \subseteq \Pi$ is relatively compact if and only if A° is, and it is trivial to see that the same holds for pervasiveness and the non-crossing property. For $B \subseteq \Pi^\downarrow$ we write

$$B_\downarrow = \{g \in \Pi^\downarrow; g \subseteq f \text{ for some } f \in B\} \quad (2.15)$$

in analogy to (2.6).

A random subset of Π^\downarrow that is compact, pervasive and non-crossing is said to be a *dual weave*. Thus \mathcal{A} is a dual weave if and only if \mathcal{A}° is a weave. We say that $\mathcal{A} \subseteq \Pi^\downarrow$ is a dual web if and only if \mathcal{A}° is a web. Equivalently, we could define a relation on $\mathcal{K}(\Pi^\downarrow)$ akin to (2.7) but with time reversed (i.e. with \cdot_\downarrow in place of \cdot_\uparrow) and then a dual web would be a minimal dual weave with respect to this relation. Note that \mathcal{F} is a flow if and only if \mathcal{F}° is a flow.

Definition 2.4.7 A pair $(\mathcal{W}, \hat{\mathcal{W}})$ is said to be a *double web* if it consists of a web \mathcal{W} and dual web $\hat{\mathcal{W}}$ coupled such that $\mathcal{W} \cup \hat{\mathcal{W}}$ is non-crossing.

Our next result states that each web gives rise to a corresponding double web, in which $\hat{\mathcal{W}}$ essentially contains the extra segments of paths that are required to construct \mathcal{F} directly from \mathcal{W} . There is a subtlety, however: when we come to connect \mathcal{W} and its dual $\hat{\mathcal{W}}$ together to create \mathcal{F} , we must be careful not to introduce crossing. In particular we should be wary of ramification points, at which the multiple in-going and out-going trajectories must be reconnected in such a way that they enter and exit z without crossing each other.

Recall our terminology that a path $f \in \Pi$ begins at the point $(f(\sigma_f-), \sigma_f) \in \mathbb{R}_c^2$ and ends at the point $(f(\tau_f+), \tau_f)$. Given $D \subseteq \mathbb{R}_c^2$ we say that f begins in D if f begins at some point of D , and f ends in D if f ends at some point of D . For $f \in \Pi^\uparrow, g \in \Pi^\downarrow$ are non-crossing, with $\tau_g = \sigma_f$ and $g(\tau_g+) = f(\sigma_f-)$, then we define $h = g \hookrightarrow f \in \Pi^\downarrow$ by

$$h(t\star) = \begin{cases} g(t\star) & \text{for } t\star \leq \sigma_f- \\ f(t\star) & \text{for } t\star \geq \sigma_f+. \end{cases}$$

Thus $g \hookrightarrow f \in \Pi^\downarrow$ is the concatenation of a path $f \in \Pi^\uparrow$ and a path $g \in \Pi^\downarrow$ that (respectively) begin and end at the same point of space-time.

Theorem 2.4.8 *Let \mathcal{A} be a weave and let $\mathcal{W} = \text{web}(\mathcal{A})$, $\mathcal{F} = \text{flow}(\mathcal{A})$. There exists a dual web $\hat{\mathcal{W}}$ on the same probability space such that $(\mathcal{W}, \hat{\mathcal{W}})$ is a double web, and $\hat{\mathcal{W}}$ is unique up to almost sure equivalence. For any $D \subseteq \mathbb{R}^2$ that is dense and almost surely non-ramified,*

$$\hat{\mathcal{W}} = \overline{\{g \in \Pi^\downarrow; g \text{ does not cross } \mathcal{A} \text{ and } g \text{ begins in } D\}_\downarrow} \quad (2.16)$$

$$\mathcal{F} = \overline{\{g \hookrightarrow f \in \Pi^\downarrow; g \in \hat{\mathcal{W}} \text{ ends and } f \in \mathcal{W} \text{ begins at the same point of } D\}}. \quad (2.17)$$

From Theorems 2.4.3 and 2.4.8, if $(\mathcal{W}, \hat{\mathcal{W}}, \mathcal{F})$ is a triplet containing a double web and flow-lines, all coupled to be non-crossing of each other, then given any one element of the triplet we may reconstruct the other two.

Recall that $\Pi_c = \{f \in \Pi; f \text{ is continuous}\}$. Let us end this section by recording that continuity of paths is preserved through all of the various relationships established above. We say that a weave \mathcal{A} is continuous if $\mathcal{A} \subseteq \Pi_c^\uparrow$, and similarly for dual weaves.

Theorem 2.4.9 *Let \mathcal{A} and \mathcal{B} be weaves such that $\mathcal{A} \sim \mathcal{B}$. Then \mathcal{A} is continuous if and only if \mathcal{B} is continuous. If \mathcal{W} is a web then \mathcal{W} is continuous if and only if $\hat{\mathcal{W}}$ is continuous.*




Figure 2.5.1: In all three figures, time runs upwards and the spatial axis is horizontal. In each figure a weave is depicted via solid lines, with the corresponding flow depicted via including dotted lines.

On the left: A weave \mathcal{A} such that $\mathcal{A}(D)$ contains the bi-infinite red path, whereas $\mathcal{A}|_D$ does not. Here $D \subseteq \mathbb{R}^2$ must be dense and non-ramified with respect to the weave \mathcal{A} . In this case we have that $\text{web}(\mathcal{A}) = \overline{(\mathcal{A}|_D)_\uparrow} \prec \mathcal{A}(D)_\uparrow$.

On the center and right: An example related to continuity of the $\text{flow}(\cdot)$ map and lack of continuity of the $\text{web}(\cdot)$ map, as well as to the existence of isolated points within flows and general weaves. In the center, the weave \mathcal{A}_ϵ and corresponding flow $\mathcal{F}_\epsilon = \text{flow}(\mathcal{A}_\epsilon)$ are depicted. The limiting weave $\mathcal{A} = \lim_{\epsilon \rightarrow 0} \mathcal{A}_\epsilon$ is depicted on the right, with $\mathcal{F} = \text{flow}(\mathcal{A})$ again depicted via including dotted lines. The red paths collapse to a single bi-infinite path in the limit. Note that $\mathcal{F}_\epsilon \rightarrow \mathcal{F}$ in accordance with part 3 of Theorem 2.4.5. In this case $\mathcal{W}_\epsilon = \text{web}(\mathcal{A}_\epsilon)$ is equal to $(\mathcal{A}_\epsilon)_\uparrow$. Therefore $\mathcal{W}_\epsilon \rightarrow \mathcal{A}_\uparrow$. On the right, note that $\mathcal{W} = \text{web}(\mathcal{A})$ does not include the bi-infinite red path, and that every point of this path ramified. Consequently $\lim_{\epsilon \rightarrow 0} \mathcal{W}_\epsilon \neq \mathcal{W}$, showing that the map $\text{web}(\cdot)$ is discontinuous at \mathcal{W} . Note that the bi-infinite red path on the right is an isolated point of \mathcal{A} but is not an isolated point of \mathcal{F} .

2.5 Discussion

In this section we first discuss how various objects related to the Brownian web fit into the framework of weaves. We then discuss other topologies that have been introduced for sets of paths and closely related objects. Some open problems are mentioned along the way.

2.5.1 The equivalence class of the Brownian web

We write \mathcal{W}_b for the Brownian web, as defined in (for example) Theorem 2.3 of the survey article of Schertzer et al. (2017). Let us first resolve an apparent conflict in notation. In common with the literature of \mathcal{W}_b , Schertzer et al. (2017) defined $\mathcal{W}_b(D)$ to be the set of paths in \mathcal{W}_b that *begin* at some $z \in D$, where $D \subseteq \mathbb{R}_c^2$. According to (2.9) we reserve $\mathcal{W}_b(D)$ for the set of paths that *pass through* some $z \in D$. This may seem to conflict at first glance, but in fact there is no conflict here, as we now explain.

For the Brownian web, the notation $\mathcal{W}_b(D)$ is widely used when $D \subseteq \mathbb{R}_c^2$ is deterministic and countable, for example in the well known identity $\mathcal{W}_b \stackrel{\text{a.s.}}{=} \overline{\mathcal{W}_b(D)}$. For the Brownian web, almost surely, for each $z \in D$ the set $\mathcal{W}_b(z) = \{f \in \mathcal{W}_b; f \text{ passes through } z\}$ consists of a single path that begins at z . Thus, for the Brownian web, $\mathcal{W}_b(z) \stackrel{\text{a.s.}}{=} \mathcal{W}_b|_z$ and $\mathcal{W}_b(D) \stackrel{\text{a.s.}}{=} \overline{\mathcal{W}_b|_D}$. In general the distinction between $\mathcal{A}(D)$ and $\mathcal{A}|_D$ does matter and (2.12), which features the latter, is required to construct $\text{web}(\mathcal{A})$. See the left part of Figure 2.5.1 for a related example.

Lemma 2.5.1 *It holds that \mathcal{W}_b is a web.*

A short proof of Lemma 2.5.1 is given in Appendix A.4. It rests on combining Theorem 2.4.3 with the key property $\mathcal{W}_b \stackrel{\text{a.s.}}{=} \overline{\mathcal{W}_b(D)}$, from which we may deduce that $\mathcal{W}_b \stackrel{\text{a.s.}}{=} \text{web}(\mathcal{W}_b)$. We refer to the equivalence class of the Brownian web as the class of *Brownian weaves*, which are introduced for the first time in the present article. The *double Brownian web* $(\mathcal{W}_b, \overline{\mathcal{W}_b})$ as

defined in Theorem 2.4 of Schertzer et al. (2017) consists of a pair of coupled random variables, where $(\widehat{\mathcal{W}}_b)^\circ$ and \mathcal{W}_b have the same marginal distributions, such that \mathcal{W}_b and $\widehat{\mathcal{W}}_b$ do not cross. It follows from Theorem 2.4.8 that $(\mathcal{W}_b, \widehat{\mathcal{W}}_b)$ is a double web in our framework.

Let us now comment on the role of special points of the Brownian web versus ramification points of weaves. Within the Brownian web each space-time point $z = (x, t)$ is assigned a ‘type’ denoted $(z_{\text{in}}, z_{\text{out}})$. Here, z_{in} is the number of equivalence classes of incoming paths of $z = (x, t)$ that are distinct under the relation that two paths are equivalent if they are equal on a time interval $[t - \epsilon, t]$ for some $\epsilon > 0$. Similarly for z_{out} , using outgoing paths and $[t, t + \epsilon]$. See Section 2.5 of Schertzer et al. (2017) for further detail. Note that z_{in} and z_{out} are local properties (in space-time) of z , whereas ramification of z is not a local property, because ramification depends on the behaviour of paths within $\mathcal{A}(z)$ for *all* time. The two concepts are related but have different purposes.

Theorem 2.11 of Schertzer et al. (2017) describes the various types of special point within \mathcal{W}_b , plus their associated local geometry and Hausdorff dimension. This provides a highly detailed understanding of the microscopic structure of \mathcal{W}_b . Within the Brownian web points of type $(0, 1)$ have full measure in \mathbb{R}^2 and are non-ramified. Points of all other types are regarded as ‘special’ points of \mathcal{W}_b and are ramified in \mathcal{W}_b . This is something of a coincidence: in general weaves points of type $(0, 1)$ can be ramified (for example, if they are upstream of a branch point) and points of type $(1, 1)$ can be non-ramified (for example, the constant paths in Figure 1.1.1 that do not interact with the jumps). Moreover, for general weaves the most abundant type is not necessarily $(0, 1)$. Within Figures 1.1.1 and 2.5.1 all weaves depicted have points of type $(1, 1)$ with full measure.

Fontes and Newman (2006) explored two examples of Brownian weaves. They considered the *full Brownian web*, which in our terminology is precisely the flow \mathcal{F}_b associated to the Brownian weave, and the *full forwards Brownian web*, which in our terminology is $(\mathcal{F}_b)_\uparrow$, and is an example of a Brownian weave that is neither a flow nor a web. Extension of paths, backwards in time, also features within Fontes and Newman (2006). Their treatment relies fully on the structure of the Brownian web, using the forwards-backwards reflection of Brownian paths established by Soucaliuc et al. (2000). The particular case of our own Theorem 2.4.8 corresponding to $(\mathcal{W}_b, \widehat{\mathcal{W}}_b)$ and \mathcal{F}_b is essentially Proposition 2.5 of Fontes and Newman (2006). They also adapted results of Fontes et al. (2004) and the earlier work of Piterbarg (1998) to give weak convergence criteria for \mathcal{F}_b .

The Brownian net of Sun and Swart (2008) is not a weave, because it contains paths that cross. It is interesting to ask if a generalized form of nets exist, as a family of $\mathcal{K}(\Pi)$ valued random variables corresponding to general weaves, but we do not attempt to answer this question within the present article. It is also interesting to ask if there is a generalization of our results to pervasive systems that permit crossing, such as the α -stable web of Mountford et al. (2019), however at present very few non-trivial examples of such systems are available.

2.5.2 Related topologies

We have already mentioned that the state space $\mathcal{K}(\Pi_c^\uparrow)$ constructed by Fontes et al. (2004), upon which most of the recent work on the Brownian web is based, is a topological subspace of our own state space $\mathcal{K}(\Pi^\uparrow)$. In this section we discuss some other recent works concerning topologies induced upon sets of paths.

Berestycki et al. (2015) mapped sets of paths to sets of ‘tubes’. Loosely, a tube is a subset of space-time that possesses a bottom face, sides and a top face. The so-called *tube topology* is then induced based on which tubes are traversed (i.e. from bottom to top, as time passes, whilst remaining within the sides) by the paths. It is restricted to sets of continuous coalescing paths,

but permits paths to cross.

The state space defined by Berestycki et al. (2015) is compact, which has substantial technical advantages (in particular, tightness becomes automatic) but this comes at the cost of some loss of detail: the tube topology is coarser than that of $\mathcal{K}(\Pi_c^\uparrow)$ and the map from sets of paths to sets of tubes is not injective. In fact, the tube topology regards the sets A and A_\uparrow as identical, where $f \in \Pi^\uparrow$, and similarly for all A' such that $A_\uparrow \preceq A' \preceq A$, meaning that much of the structure displayed in Theorem 2.4.3 is lost. However, the weaker representation makes characterization and convergence easier.

Another piece of detail that is kept visible in $\mathcal{K}(\Pi^\uparrow)$, but is dropped by the tube topology, is behaviour near the start times of paths. This can have implications for universality. For example, Berestycki et al. (2015) showed that systems of coalescing random walkers with heavy tailed jumps (linearly interpolated) will converge in law to the Brownian web, under the tube topology; Newman et al. (2005) had previously shown that such convergence failed within $\mathcal{K}(\Pi_c^\uparrow)$ because jumps attempt to form in the limit at the initial times of some paths. We conjecture that such systems *will converge but not to the Brownian web* if considered as elements of $\mathcal{K}(\Pi^\uparrow)$. Loosely, we expect that the limit of such systems will be a Brownian web that is suitably augmented with jumps at initial times of paths.

Aside from tubes, another possibility is to view the Brownian web as a real tree, where the natural root is a point at time $+\infty$ at which all paths coalesce. In this representation, loosely, each space-time point on a path within the Brownian web becomes a point within the corresponding real tree, and a metric is induced that captures both the natural tree structure and the distances travelled, forwards in time, along individual paths within the Brownian web until coalescence points. Of course, not all sets of paths are suited to such a representation; systems that contain branching are not.

Cannizzaro and Hairer (2021) identify a subset of $\mathcal{K}(\Pi_c^\uparrow)$ that can be naturally represented as real trees. A similar theme underlies the framework of *marked metric measure spaces* introduced by Depperschmidt et al. (2011). Like the tube topology, the setup of Cannizzaro and Hairer (2021) does not distinguish between A_\uparrow and A , in this case by associating real trees with sets of paths of the form A_\uparrow (i.e. decreasing sets under \subseteq). They construct a Polish topology that is shown to be finer than that of $\mathcal{K}(\Pi)$, in particular it enforces that coalescence times of paths are preserved when taking limits.

3 Preliminaries

We now begin the proofs. In this section we develop some underlying concepts that we require for our theory of weaves, related to the structure of $\mathcal{K}(\Pi^\uparrow)$. In Section 3.1 we relate our notion of crossing to a relation \triangleleft that describes when one path lies to the left of another. In Section 3.2 we show that the relation \preceq is a partial order. Finally, in Section 3.3 we examine the interaction between order relations and topology, including to what extent the relations \subseteq , \triangleleft and \preceq are preserved by taking limits. These results are technical in nature. Readers wishing to gloss over technical issues may prefer to note the key results and definitions (at minimum, Definition 3.1.1 and Lemma 3.1.6), then proceed to Section 4.

3.1 On crossing

In this section we study the interaction between crossing and the idea of one path staying to the left (or right) of another. The results in this section have straightforward extensions to Π but for brevity we will state results covering Π^\uparrow . Consequently we must handle jumps at the initial

time σ_f of $f \in \Pi^\uparrow$ but our compactification of space-time, in Figure 2.2.1, means that jumps do not occur at time $\tau_f = +\infty$.

Definition 3.1.1 For $f, g \in \Pi^\uparrow$, we write $f \triangleleft g$ if there exist $f', g' \in \Pi^\uparrow$ with $f \subseteq f'$ and $g \subseteq g'$ such that $f'(t_\star) \leq g'(t_\star)$ for all $t_\star \in \mathbb{R}_s$. We write $f \not\triangleleft g$ if this property fails to hold.

The statement $f \triangleleft g$ should be interpreted as ‘ f lies to the left of g ’. There is some subtlety involved here. If $f \triangleleft g$ then $f(t_\star) \leq g(t_\star)$ for all $t_\star \geq \sigma+$, but no such guarantee exists concerning $g(\sigma-)$ and $f(\sigma-)$. An example of $f \triangleleft g$ with $g(\sigma-) < f(\sigma-)$ appears as (i) in Figure 3.1.1.

The relation \triangleleft does not define a partial order on Π^\uparrow . Antisymmetry fails because if $g \subseteq f$ with $g \neq f$ then we have both $g \triangleleft f$ and $f \triangleleft g$. Transitivity fails too, for example if $f(t_\star) = -\frac{1}{2} + \mathbb{1}\{t_\star = \sigma_f-\}$, $g(t_\star) = 0$, $h(t_\star) = \frac{1}{2} - \mathbb{1}\{t_\star = \sigma_f-\}$, with $\sigma_f = \sigma_h = 0$ and $\sigma_g = 1$, then $f \triangleleft g$ and $g \triangleleft h$ but $f \not\triangleleft h$ (we leave it as an exercise for the reader to construct an example where transitivity fails and $\{f, g, h\}$ is also non-crossing!). For this reason we will not use the symbols \triangleright and $\not\triangleright$ in this article: it would be possible to define them as analogous concepts to \triangleleft and $\not\triangleleft$ with the roles of right and left swapped, but the notation would be unintuitive since $\not\triangleleft$ and \triangleright would not be equivalent. In Lemma 3.1.6 we will establish a more restricted setting in which \triangleleft is better behaved.

Definition 3.1.1 is intuitive and interacts well with Definition 2.3.1, but it is helpful to have a more explicit characterization of when one path lies to the left of another. We define subsets $L(f)$ and $R(f)$ of $\overline{\mathbb{R}} \times \mathbb{R}_s$, and $L_{t_\star}(f)$ and $R_{t_\star}(g)$ of $\overline{\mathbb{R}}$ by:

$$\begin{aligned} L_{t_\star}(f) &= \begin{cases} \emptyset & \text{if } t_\star < \sigma_f-, \text{ or if } t_\star = \sigma_f - \text{ and } f(\sigma_f-) \leq f(\sigma_f+), \\ [-\infty, f(t_\star)) & \text{if } t_\star \geq \sigma_f+, \text{ or if } t_\star = \sigma_f - \text{ and } f(\sigma_f+) < f(\sigma_f-), \end{cases} \\ R_{t_\star}(f) &= \begin{cases} \emptyset & \text{if } t_\star < \sigma_f-, \text{ or if } t_\star = \sigma_f - \text{ and } f(\sigma_f+) \leq f(\sigma_f-), \\ (f(t_\star), \infty] & \text{if } t_\star \geq \sigma_f+, \text{ or if } t_\star = \sigma_f - \text{ and } f(\sigma_f-) < f(\sigma_f+), \end{cases} \\ L(f) &= \bigcup_{t_\star \in \mathbb{R}_s} L_{t_\star}(f) \times \{t_\star\}, & R(f) &= \bigcup_{t_\star \in \mathbb{R}_s} R_{t_\star}(f) \times \{t_\star\}. \end{aligned} \quad (3.1)$$

Note that $L(f)$ and $R(f)$ are subsets of $\overline{\mathbb{R}} \times \mathbb{R}_s$. The significance of $L(f)$ is that if $g(t_\star) \in L_{t_\star}(f)$ then g must stay ‘to the left’ of f in order to avoid crossing it. Similarly if $g(t_\star) \in R_{t_\star}(g)$, to the right. We will formalize this intuition in Lemma 3.1.4. See Figure 3.1.1 for a picture.

For $t \geq \sigma_f+$, the sets $L_{t_\star}(f)$ and $R_{t_\star}(f)$ are, respectively, the set of points strictly to the left and right of $f(t_\star)$. However, for $t = \sigma_f$ we must take into account the presence and direction of a jump at time σ_f . Note that t_\star with $t = \pm\infty$ are excluded from $L(f)$ and $R(f)$, in accordance with the compactification of space-time in Figure 2.2.1.

Lemma 3.1.2 *Let $f, g \in \Pi^\uparrow$. Then $f \triangleleft g$ if and only if $L(f) \cap R(g) = \emptyset$.*

PROOF: We prove the forwards and backwards implications in turn, beginning with the former. Suppose that $f \triangleleft g$. Then there exists $f', g' \in \Pi^\uparrow$ with $f \subseteq f'$ and $g \subseteq g'$ such that $f'(t_\star) \leq g'(t_\star)$ for all $t_\star \in \mathbb{R}_s$. It follows from (3.1) that $L(f') \cap R(g') = \emptyset$, which since $L(f) \subseteq L(f')$ and $R(g) \subseteq R(g')$ implies that $L(f) \cap R(g) = \emptyset$.

For the reverse implication, let $f, g \in \Pi^\uparrow$ with $L(f) \cap R(g) = \emptyset$. Let $\sigma = \sigma_f \vee \sigma_g$. It follows immediately from (3.1) that $f(t_\star) \leq g(t_\star)$ for all $t_\star \geq \sigma+$. We will construct explicit $f', g' \in \Pi^\uparrow$ such that $f \subseteq f'$, $g \subseteq g'$ and $f'(t_\star) \leq g'(t_\star)$ for all $t_\star \in \mathbb{R}_s$. Note that we have nothing to prove if $\sigma_f = \sigma_g = -\infty$, and that we do not need to define f' or g' at t_\star for $t = \pm\infty$. We consider three cases, at least one of which must occur.




Figure 3.1.1: *On the left:* A schematic depiction of the subsets $L(f)$ and $R(f)$ of $\bar{\mathbb{R}} \times \mathbb{R}_s$. Time is running upwards, with values taken by $f(t^\star)$ shown as a thick black line. To help visualize, when f jumps at time t we depict time as split into t^- and t^+ via thin horizontal lines. The path f makes a jump at its starting time σ_f , at the very bottom of the figure. Note that $L_{\sigma_f^-}(f) = \emptyset$ and $R_{\sigma_f^-}(f) = (f(\sigma_f^-), \infty]$ because the jump at σ_f is rightwards.

On the right: Four examples of sets $\{f, g\}$ containing two paths. In each example, the initial point of g lies to the left of the initial point of f . The paths in (i) satisfy $f \triangleleft g$ and do not cross, but in examples (ii)–(iv) we have that f crosses g from right to left.

1. Consider if $L_{\sigma^-}(f) = \emptyset$. By (3.1) we have $\sigma = \sigma_f \geq \sigma_g$ and $f(\sigma^-) \leq f(\sigma^+)$. In this case set

$$f'(t^\star) = \begin{cases} f(t^\star) & \text{if } t^\star \geq \sigma_f^+ \\ -\infty & \text{if } t^\star \leq \sigma_f^- \end{cases} \quad g'(t^\star) = \begin{cases} g(t^\star) & \text{if } t^\star \geq \sigma_g^+ \\ g(\sigma_{g^-}) & \text{if } t^\star \leq \sigma_g^- \end{cases}.$$

2. Consider if $R_{\sigma^-}(g) = \emptyset$. By (3.1) we have $\sigma = \sigma_g \geq \sigma_f$ and $g(\sigma^+) \leq g(\sigma^-)$. In this case set

$$f'(t^\star) = \begin{cases} f(t^\star) & \text{if } t^\star \geq \sigma_f^+ \\ f(\sigma_{f^-}) & \text{if } t^\star \leq \sigma_f^- \end{cases} \quad g'(t^\star) = \begin{cases} g(t^\star) & \text{if } t^\star \geq \sigma_g^+ \\ \infty & \text{if } t^\star \leq \sigma_g^- \end{cases}.$$

3. Consider if $L_{\sigma^-}(f) = [-\infty, f(\sigma^-)]$ and $R_{\sigma^-}(g) = (g(\sigma^-), \infty]$. Using that $L(f) \cap R(g) = \emptyset$ we have $f(\sigma^-) < g(\sigma^-)$. If $\sigma = \sigma_f \geq \sigma_g$ then we set

$$f'(t^\star) = \begin{cases} f(t^\star) & \text{if } t^\star \geq \sigma_f^+ \\ f(\sigma_{f^-}) + g(t^\star) - g(\sigma_{f^-}) & \text{if } t^\star \in [\sigma_{g^+}, \sigma_{f^-}] \\ f(\sigma_{f^-}) + g(\sigma_{g^-}) - g(\sigma_{f^-}) & \text{if } t^\star \leq \sigma_{g^-} \end{cases} \quad g'(t^\star) = \begin{cases} g(t^\star) & \text{if } t^\star \geq \sigma_{g^+} \\ g(\sigma_{g^-}) & \text{if } t^\star \leq \sigma_{g^-} \end{cases}.$$

Note that f' copies the increments of g' during $t^\star \in [\sigma_{g^+}, \sigma_{f^-}]$, backwards in time starting from the condition $f'(\sigma_f) = f(\sigma_{f^-}) < g(\sigma_{f^-}) = g'(\sigma_f)$, and then backwards from time σ_{g^-} both paths remain constant. This ensures that $f'(t^\star) \leq g'(t^\star)$ for all t^\star . If $\sigma = \sigma_g \geq \sigma_f$ then we may employ a similar strategy, where g copies the increments of f during $[\sigma_{f^+}, \sigma_{g^-}]$, backwards in time starting from $g(\sigma_{g^-})$.

In all cases it is clear that $f \subseteq f'$, $g \subseteq g'$ and $f'(t^\star) \leq g'(t^\star)$ for all $t^\star \in \mathbb{R}_s$. ■

Lemma 3.1.3 *Let $f, g \in \Pi^\uparrow$. The following statements are equivalent: (i) f and g are non-crossing; (ii) $L(f) \cap R(g) = \emptyset$ or $R(f) \cap L(g) = \emptyset$; (iii) $f \triangleleft g$ or $g \triangleleft f$.*

PROOF: Equivalence of (ii) and (iii) follows immediately from Definition 2.3.1 and Lemma 3.1.2. It follows trivially from Definitions 2.3.1 and 3.1.1 that (iii) implies (i). Let us now show

that (i) implies (ii). If f and g are non-crossing then we have $f', g' \subseteq \Pi^\dagger$ such that $f \subseteq f'$, $g \subseteq g'$ and $f'(t_\star) \leq g'(t_\star)$ for all $t_\star \in \mathbb{R}_s$ or $g'(t_\star) \leq f'(t_\star)$ for all $t_\star \in \mathbb{R}_s$. In the former case by (3.1) we have $L(f') \cap R(g') = \emptyset$, which implies $L(f) \cap R(g) = \emptyset$. In the latter case by (3.1) we have $L(g') \cap R(f') = \emptyset$, which implies $L(g) \cap R(f) = \emptyset$. Thus we have (ii). \blacksquare

By Lemma 3.1.3, if f and g cross, then there must exist $t_{1\star_1} \neq t_{2\star_2}$ such that $R_{t_{1\star_1}}(\pi) \cap L_{t_{1\star_1}}(\pi') \neq \emptyset$ and $L_{t_{2\star_2}}(\pi) \cap R_{t_{2\star_2}}(\pi') \neq \emptyset$. If this happens for $t_{1\star_1} < t_{2\star_2}$ (resp. $t_{1\star_1} > t_{2\star_2}$), then we say that f crosses g from left to right (resp. from right to left). Of course, it can happen that f crosses g from left to right and also from right to left. See Figure 3.1.1 for a picture.

Lemma 3.1.4 *Suppose that $f, g \in \Pi^\dagger$ are non-crossing and let $I_s = I(f)_s \cap I(g)_\pm$. If there exists $t_\star \in I_s$ such that $(f(t_\star), t_\star) \in L(g)$ then $f \triangleleft g$. If there exists $t_\star \in I_s$ such that $(f(t_\star), t_\star) \in R(g)$ then $g \triangleleft f$. Moreover, in either case $f \neq g$.*

PROOF: We will establish the first claim first: let f, g be as given and suppose $(f(t_\star), t_\star) \in L(g)$. Hence $(f(t_\star), t_\star) \in L_{t_\star}(g)$ so $L_{t_\star}(g) = [-\infty, g(t_\star))$ and $f(t_\star) < g(t_\star)$. Let $\sigma = \sigma_f \vee \sigma_g$. Consider if $s_\bullet \geq \sigma+$ and $f(s_\bullet) < g(s_\bullet)$ for some $s_\bullet \geq \sigma-$. Then, from Definition 3.1.1, we have $L_{s_\bullet}(g) = [-\infty, g(s_\bullet))$ and $R_{s_\bullet}(f) = (f(s_\bullet), \infty]$, which implies $L(g) \cap R(f) \neq \emptyset$. Lemma 3.1.3 thus implies $L(f) \cap R(g) = \emptyset$, from which Lemma 3.1.2 gives $f \triangleleft g$. In particular, if $t_\star \geq \sigma_f+$ then $f \triangleleft g$.

It remains only to consider the case of $t_\star = \sigma_f-$ and, from what we have shown in the paragraph above, in this case we may assume without loss of generality that $f(s_\bullet) \leq g(s_\bullet)$ for all $s_\bullet \geq \sigma+$. We have $f(\sigma_f-) < g(\sigma_f-)$ and $L_{\sigma_f-}(g) = [-\infty, g(t_\star))$. If $f(\sigma_f-) < f(\sigma_f+)$ then $R_{\sigma_f-}(f) = [-\infty, f(\sigma_f-))$ and hence $L(g) \cap R(f) \neq \emptyset$, so here also Lemmas 3.1.3 and 3.1.2 imply $f \triangleleft g$. Otherwise, $f(\sigma_f+) \leq f(\sigma_f-)$, in which case it is immediate from Definition 3.1.1 that $f \triangleleft g$.

The second claim, regarding the case $(f(t_\star), t_\star) \in R(g)$, follows by symmetry (consider space $\overline{\mathbb{R}}$ reflected about the origin). Lastly, the fact that $f \neq g$ follows from noting that points of the form $(g(t_\star), t_\star) \in \mathbb{R}_c^2$ are not elements of $L(g)$ or $R(g)$. \blacksquare

Lemma 3.1.5 *The following hold:*

1. *Suppose $f, g \in \Pi^\dagger$ with $f \triangleleft g$ and write $\sigma = \sigma_f \vee \sigma_g$. Precisely one of following occurs:*

- (i) $f \subseteq g$ or $g \subseteq f$;
- (ii) $f(t_\star) < g(t_\star)$ for some $t_\star \geq \sigma+$;
- (iii) $f(t_\star) = g(t_\star)$ for all $t_\star \geq \sigma+$, $f(\sigma-) < f(\sigma+) = g(\sigma+) < g(\sigma-)$;
- (iv) $f(t_\star) = g(t_\star)$ for all $t_\star \geq \sigma+$, $f(\sigma-) < g(\sigma-) \leq f(\sigma+) = g(\sigma+)$ and $\sigma_g < \sigma_f = \sigma$;
- (v) $f(t_\star) = g(t_\star)$ for all $t_\star \geq \sigma+$, $f(\sigma+) = g(\sigma+) \leq f(\sigma-) < g(\sigma-)$ and $\sigma_f < \sigma_g = \sigma$.

In cases (ii)-(v) we have $g \not\triangleleft f$.

2. *Let $f, g \in \Pi^\dagger$. Then $f \triangleleft g$ and $g \triangleleft f$ if and only if $f \subseteq g$ or $g \subseteq f$.*

PROOF: Let us begin with the first statement. It is clear that all five cases are distinct. Suppose neither of (ii), (iii) (iv) and (v) occurs, and we will seek to prove that (i) holds. Since (ii) fails we have $g(t_\star) \leq f(t_\star)$ for all $t_\star \geq \sigma+$. Since $f \triangleleft g$ we also have $f(t_\star) \leq g(t_\star)$ for all such t_\star , hence in fact we have equality for $t_\star \geq \sigma+$, in particular at $t_\star = \sigma+$. Since (iii) fails, $f(t_-)$ and $g(t_-)$ lie (non-strictly) on the same side of $f(t_+) = g(t_+)$.

Consider first when they both lie to the left, that is $f(t_-) \vee g(t_-) \leq f(t_+) = g(t_+)$. We divide into three cases based upon whether $\sigma_f = \sigma_g$, $\sigma_f < \sigma_g$ or $\sigma_g < \sigma_f$.

- If $\sigma_f = \sigma_g$ then $\sigma_f = \sigma_g = \sigma$, in which case (i) occurs.
- If $\sigma_f < \sigma_g$ then $\sigma_g = \sigma$, so we have $L_{\sigma-}(f) = [-\infty, f(\sigma-))$ and $R_{\sigma-}(g) = (g(\sigma-), \infty]$. By Lemma 3.1.2 we have $L_{\sigma-}(f) \cap R_{\sigma-}(g) = \emptyset$ so $f(\sigma-) \leq g(\sigma-)$. Hence $g \subseteq f$.
- If $\sigma_g < \sigma_f$ then $\sigma_f = \sigma$ and as (iv) does not occur we must have $g(\sigma-) \leq f(\sigma-)$, which means $f \subseteq g$.

In all three cases we have that (i) occurs. It remains to consider when both $f(\sigma-)$ and $g(\sigma-)$ lie to the right of $f(\sigma+) = g(\sigma+)$. A symmetric argument, in which (v) takes the place of (vi), shows that (i) also occurs. Note that in case (ii) we have $L_{t\star}(f) \cap R_{t\star}(g) \neq \emptyset$. In cases (iii), (vi) and (v) we have $L_{\sigma-}(g) \cap R_{\sigma-}(f) \neq \emptyset$. Thus, in all of cases (ii)-(v) Lemma 3.1.2 gives $g \not\subseteq f$. For the second claim of the present lemma, the reverse implication is trivial from Definition 3.1.1 and the forwards implication follows from part 1 of the present lemma. \blacksquare

Part 1 of Lemma 3.1.5 makes explicit the difficulties inherent to paths that may jump at their initial times. We will often use arguments that give some special attention to the initial times of paths. The sets $L(f), R(g)$ and the relation \triangleleft allow us to do so without having to work through cases (i)-(v) in turn.

Lemma 3.1.6 *Let $A \subseteq \Pi^\uparrow$ be non-crossing. Then $(A_{\max}, \triangleleft)$ is a totally ordered space.*

PROOF: Recall that A_{\max} was defined in (2.5). By Lemma 3.1.3 all pairs $f, g \in A$ satisfy $f \triangleleft g$ or $g \triangleleft f$, which holds in particular for A_{\max} . It is clear from Definition 3.1.1 that $f \triangleleft f$ for all $f \in \Pi^\uparrow$, thus also for all $f \in A_{\max}$. If $f, g \in A_{\max}$ satisfy $f \triangleleft g$ and $g \triangleleft f$, then Lemma 3.1.5 tells us that $f \subseteq g$ or $g \subseteq f$. By maximality, this implies $f = g$. We have now shown that \triangleleft is reflexive and antisymmetric, and that all pairs of elements are comparable. It remains to show that \triangleleft is transitive.

Let $f, g, h \in A_{\max}$ with $f \triangleleft g$ and $g \triangleleft h$. If $f = g$ or $g = h$ then it is trivial that $f \triangleleft h$. If $f = h$ then we have $f \triangleleft g$ and $g \triangleleft h = f$, so Lemma 3.1.5 implies that $f \subseteq g$ or $g \subseteq f$, which by maximality implies $f = g$, hence also $f \triangleleft g$. Thus we may assume without loss of generality that f, g, h are distinct elements of A_{\max} . By Lemma 3.1.2 we have $L(f) \cap R(g) = \emptyset$ and $L(g) \cap R(h) = \emptyset$, and we must show that $L(f) \cap R(h)$ is also empty.

We will argue by contradiction. Suppose that $(x, t\star) \in L(f) \cap R(h)$, which implies that $L_{t\star}(f) = [-\infty, f(t\star))$, $R_{t\star}(h) = (h(t\star), \infty]$ and $h(t\star) < x < f(t\star)$. Consider first if $t\star \geq \sigma_g-$. If $g(t\star) \leq x$ then $g(t\star) \in L_{t\star}(f)$, Lemma 3.1.4 gives that $g \triangleleft f$, in which case part 2 of Lemma 3.1.5 and maximality gives that $g = f$, which is a contradiction to our assumptions. Similarly, if $g(t\star) \geq x$ then $g(t\star) \in R_{t\star}(h)$, Lemma 3.1.4 gives that $g \triangleleft h$, from which part 2 of Lemma 3.1.5 and maximality give $g = h$, which is again a contradiction.

It remains to consider when $t < \sigma_g$. In this case $\sigma_f \vee \sigma_h < \sigma_g$. Definition 3.1.1 thus implies that for all $s\star \geq \sigma_g+$ we have $f(s\star) \leq g(s\star) \leq h(s\star)$. Applying part 1 of Lemma 3.1.5 to $f \triangleleft g$ and $g \triangleleft h$, and noting that these are distinct maximal paths (so case (i) of that lemma may not occur), we obtain that for some $s, s' \geq \sigma_g$ it holds that $f(s-) < g(s-) < h(s-)$ and $g(s'-) < h(s'-)$. If $s = s' = \sigma_g$ then $f(s-) < g(s-) < h(s-)$. If $s > \sigma_g$ then $f(s-) < g(s-) \leq h(s-)$, similarly if $s' > \sigma_g$ then $f(s'-) \leq g(s'-) < h(s'-)$. In all three cases we have $u \geq \sigma$ such that $f(u-) < h(u-)$. We have $\sigma_f \vee \sigma_h < \sigma_g \leq s$, so $L_{u-}(h) = [-\infty, h(u-))$ and $R_{u-}(f) = (f(u-), \infty]$, meaning that $L(h) \cap R(f) \neq \emptyset$. However now both $L(f) \cap R(h)$ and $L(h) \cap R(f)$ are non-empty, which by Lemma 3.1.3 implies that f and h cross, which is contradiction. \blacksquare

Lemma 3.1.7 *Let A be non-crossing and let B be non-crossing, both subsets of Π^\uparrow . Suppose $A \preceq B$. Then $A \cup B$ is non-crossing.*

PROOF: Let $f \in A$, $g \in B$ and assume that $A \preceq B$. Since $A \subseteq B_{\uparrow}$ there exists $h \in B_{\uparrow}$ such that $f \subseteq h$. We have that $g, h \in B$ and B is non-crossing, so $\{g, h\}$ is non-crossing, thus also $\{f, h\}$ is non-crossing. The result follows. \blacksquare

3.2 On partial orders of sets

For the duration of Section 3.2 let (E, \leq) denote a partially ordered set. We now recall some standard notation associated to partial orders. For $A \subseteq E$, we write $A_{\leq} := \{e \in E : e \leq e' \text{ for some } e' \in A\}$ for the *downset* of A . Note that $(A_{\leq})_{\leq} = A_{\leq}$. The *upset* A_{\geq} of A is defined in the same way as the downset A_{\leq} , but for the reversed order. A *maximal element* of a subset $A \subseteq E$ is an element $e \in A$ such that there exists no $e' \in A$ with $e < e'$. We write $A_{\max} = \{e \in A : e \text{ is a maximal element of } A\}$. A *minimal element* is defined in the same way but for the reversed order. As usual, we write $e < e'$ if $e \leq e'$ and $e \neq e'$.

Remark 3.2.1 For $A \in \mathcal{K}(\Pi)$, we have specified in (2.5) that A_{\max} refers to the maximal elements of (A, \subseteq) . In (2.6) we defined the set A_{\uparrow} to be the downset of $A \subseteq \Pi^{\uparrow}$ in $(\Pi^{\uparrow}, \subseteq)$. In Remark 3.3.4 we will note that $A \subseteq (A_{\max})_{\downarrow}$ when A is compact i.e. each path in A extends to at least one element of A_{\max} . Note that $B_{\downarrow} \subseteq \Pi^{\downarrow}$ from (2.15) is the downset of B in $(\Pi^{\downarrow}, \subseteq)$.

The following lemma puts the relation \preceq introduced in (2.7) into a wider framework. It is a natural concept that might have been studied elsewhere but have not been able to locate a reference. With slight abuse of notation we will briefly use the notation \preceq in the more abstract setting. We noted in Section 2.3 that \preceq on Π is related to how efficiently paths cover space. In the abstract setting there is a somewhat clearer interpretation. Specifically, under \leq both of the following two operations will make a set $A \subseteq E$ strictly increase: inserting a new element to A that (once inserted) is a \leq -maximal element; removing an existing element that (prior to removal) is not a \leq -maximal element.

Lemma 3.2.2 *Let (E, \leq) be a partially ordered set. For $A, B \subseteq E$, write $A \preceq B$ to mean that $A_{\leq} \cap B \subseteq A \subseteq B_{\leq}$. Then \preceq is a partial order on the set of all subsets of E .*

PROOF: Clearly $A \preceq A$. Also, $A \preceq B \preceq A$ implies $A \subseteq B_{\leq} \cap A \subseteq B$ and by a symmetric argument also $B \subseteq A$, so to complete the proof we must show that the relation \preceq is transitive. The relations $A \preceq B \preceq C$ say that

$$(i) A_{\leq} \cap B \subseteq A, \quad (ii) A \subseteq B_{\leq}, \quad (iii) B_{\leq} \cap C \subseteq B, \quad (iv) B \subseteq C_{\leq}.$$

When we apply one of these facts we will indicate which with a superscript above the corresponding \subseteq . This implies (v) $A_{\leq} \stackrel{(ii)}{\subseteq} B_{\leq}$, (vi) $B_{\leq} \stackrel{(iv)}{\subseteq} C_{\leq}$, and (vii) $B_{\leq} \cap C \stackrel{(iii)}{\subseteq} B_{\leq} \cap B$, from which we get

$$A_{\leq} \cap C \stackrel{(v),(vii)}{\subseteq} A_{\leq} \cap B \stackrel{(i)}{\subseteq} A \quad \text{and} \quad A \stackrel{(ii)}{\subseteq} B_{\leq} \stackrel{(vi)}{\subseteq} C_{\leq},$$

proving that $A \preceq C$. \blacksquare

A set A is said to be *decreasing* if $A = A_{\leq}$ and *increasing* if $A = A_{\geq}$. If A and B are decreasing sets, then $A \preceq B$ if and only if $A \subseteq B$. The proof is trivial and is left to the reader.

3.3 On compatibility of order and topology

We now turn our attention to the interaction between orders and limits. We have now introduced several partial orders related to càdlàg paths: the ‘path extension’ order \subseteq on Π , the ‘coverage

efficiency' order \preceq on $\mathcal{K}(\Pi^\uparrow)$, and the 'leftwards of' relation \triangleleft that was shown in Lemma 3.1.6 to be a total order on A_{\max} , when A is closed. In general \triangleleft is not even a partial order. With these situations in mind we make the following general definition.

Definition 3.3.1 Let (E, \mathcal{T}) be a topological space. We say that a binary relation \leq on E is *compatible* if $\{(e, f) \in E^2 : e \leq f\}$ is a closed subset of E^2 (in the product topology). With mild abuse of notation we also apply this terminology to metric spaces (E, d_E) where d_E is a metric generating the topology on E .

The point is that compatibility implies that if $e_n \rightarrow e$ and $f_n \rightarrow f$ with $e_n \leq f_n$, all elements of E , then we may conclude that $e \leq f$. We wish to study this notion in the three situations listed above, but as per our comments above we do not always wish to assume (in particular, for \triangleleft) that (E, \leq) is partially ordered. Our first result in this direction is negative:

Remark 3.3.2 It is straightforward to see that $(\mathcal{K}(\Pi^\uparrow), d_{\mathcal{K}(\Pi)}, \preceq)$ is not compatible. For example, let $f(t\pm) = 0$ with $\sigma_f = 0$ and let $g_n(t\pm) = 1/n$ with $\sigma_{g_n} = 1$. Then $\{f\} \prec \{f, g_n\}$ and $\{f, g\} \prec \{f\}$. Examples also exist where $A_n \prec B_n$ but the limits of (A_n) and (B_n) are incomparable; we leave this as an exercise for the reader. From a purely abstract point of view, the non-trivial interaction of \preceq with taking limits in $\mathcal{K}(\Pi^\uparrow)$ is the main source of interesting structure within the space of weaves.

Lemma 3.3.3 *The partial order \subseteq is compatible with (Π, d_Π) .*

PROOF: Proposition 2.2.1 gives that convergence in d_Π is equivalent to convergence under the metric $(f, g) \mapsto d_{\mathcal{K}((\mathbb{R}^2)^2)}(H^{(2)}(f), H^{(2)}(g))$. By definition (see Section 2.3) $f \subseteq g$ means that $H^{(2)}(f) \subseteq H^{(2)}(g)$. With these facts in mind the stated result follows from Lemma A.1.4, which asserts that the partial order of set inclusion is compatible with the Hausdorff metric. ■

Remark 3.3.4 Using Lemma 3.3.3 it is straightforward to check that if $A \in \mathcal{K}(\Pi)$ then for all $f \in A$ there exists $g \in A_{\max}$ such that $f \subseteq g$. We will use this fact repeatedly, without referring back to this remark, from now on.

We now consider the relation \triangleleft , which will require rather more work. The following lemma shows that any lack of compatibility between \triangleleft and (Π^\uparrow, d_Π) must involve a jump at the initial time of a limiting path. The reader should bear in mind examples like $g_n(t\star) = \mathbb{1}\{t\star \geq \frac{1}{n}+\}$, $g(t\star) = \mathbb{1}\{t\star \geq 0+\}$, $f(t) = \mathbb{1}\{t\star = 0-\}$, where $\sigma_{g_n} = \frac{1}{n}$ and $\sigma_f = \sigma_g = 0$. Note that $f \triangleleft g_n$ and $g_n \rightarrow g$, but f and g cross by jumping over each other in opposite directions at time 0. This example shows that \triangleleft is not compatible with (Π^\uparrow, d_Π) .

Lemma 3.3.5 *Let $f, g, f_n, g_n \in \Pi^\uparrow$ with $f_n \rightarrow f$ and $g_n \rightarrow g$. Write $\sigma_n = \sigma_{f_n} \vee \sigma_{g_n}$ and $\sigma = \sigma_f \vee \sigma_g$. Suppose that $L_{s\bullet}(f_n) \cap R_{s\bullet}(g_n) = \emptyset$ for all $n \in \mathbb{N}$ and $s\bullet \geq \sigma_n+$. Then also $L_{s\bullet}(f) \cap R_{s\bullet}(g) = \emptyset$ for all $s\bullet \geq \sigma+$.*

PROOF: We will argue by contradiction. Suppose that $L_{s\bullet}(f_n) \cap R_{s\bullet}(g_n) = \emptyset$ for all $n \in \mathbb{N}$ and $s\bullet \geq \sigma_n+$, and that $L_{t\star}(f) \cap R_{t\star}(g) \neq \emptyset$ where $t\star \geq \sigma+$. By the càdlàg property of f and g , without loss of generality we may take $t > \sigma$ and assume that both f and g are continuous at t . Let $\epsilon \in (0, t - \sigma)$ and $N \in \mathbb{N}$ be large enough that $|(\sigma_{f_n} \vee \sigma_{g_n}) - \sigma| \leq \epsilon/2$ for all $n \geq N$. Our assumption that $L_{s\bullet}(f_n) \cap R_{s\bullet}(g_n) = \emptyset$ implies that $f_n(s\bullet) \leq g_n(s\bullet)$ for all $s\bullet \geq \sigma_n+$, which in particular for $N \geq n$ includes all $s\bullet \geq (t - \epsilon/2)+$.

Let $t_n \star_n \rightarrow t \star$. It follows from Lemma 2.1.1 that $t_n \star_n \geq (\sigma \vee \sigma_n) +$ for all sufficiently large n , so let us pass to a subsequence and assume that this holds for all $n \in \mathbb{N}$. By Lemma A.2.3 and continuity of f, g at t we have $f_n(t_n \star_n) \rightarrow f(t \star)$ and $g_n(t_n \star_n) \rightarrow g(t \star)$. As $g(t \star) < f(t \star)$ we obtain that there exists $\delta > 0$ such that $g_n(t_n \star_n) \leq f_n(t_n \star_n) - \delta$. For sufficiently large n we have $|t_n - t| < \epsilon/2$, which implies $t_n \star_n \geq (t - \epsilon/2) +$. This contradicts the result of the previous paragraph. \blacksquare

Lemma 3.3.6 *The relation \triangleleft is compatible with (Π^\dagger, d_Π) . Moreover: let $f_n, g_n, f, g \in \Pi^\dagger$ with $f_n \rightarrow f$ and $g_n \rightarrow g$. If $\{f_n, g_n\}$ is non-crossing for each $n \in \mathbb{N}$, then $\{f, g\}$ is non-crossing.*

PROOF: The first claim follows from Lemmas 3.3.5 and 3.1.2, noting that for $f, g \in \Pi^\dagger$ the relation $f \triangleleft g$ is equivalence to requiring that $L_{s+}(f) \cap R_{s+}(g) = \emptyset$ for all $s \in \mathbb{R}$. For the second claim, by part 2 of Lemma 3.1.5 for each $n \in \mathbb{N}$ we have that $f_n \triangleleft g_n$ or $f_n \triangleleft g_n$. At least one of these two possibilities must hold for infinitely many n . From what we have already proved, it follows that $f \triangleleft g$ or $g \triangleleft f$, from which Lemma 3.1.3 gives that f and g do not cross. \blacksquare

Lemma 3.3.7 *The relation \triangleleft is compatible with (Π_c^\dagger, d_Π) . Moreover: let $f_n, g_n, f, g \in \Pi_c^\dagger$ with $f_n \rightarrow f$ and $g_n \rightarrow g$. If $\{f_n, g_n\}$ is non-crossing for each $n \in \mathbb{N}$, then $\{f, g\}$ is non-crossing.*

PROOF: The first claim follows from Lemmas 3.3.5 and 3.1.2, noting all $f \in \Pi_c^\dagger$ satisfy $f(\sigma_f -) = f(\sigma_f +)$ so that $L_{\sigma_f -}(f) = R_{\sigma_f -}(f) = \emptyset$. The proof of the second claim is essentially the same as that of Lemma 3.3.6. \blacksquare

The remainder of this section concerns conditions under which \triangleleft is preserved in limits $f_n \triangleleft g_n$ with $f_n \rightarrow f$ and $g_n \rightarrow g$ for half-infinite càdlàg paths. We will see, in Lemma 3.3.10 that the key (extra) condition is that $\{f, g\}$ must be non-crossing. From Lemma 3.3.5 if any crossing is to occur in such a limit, it must take place at time $\sigma = \sigma_f \vee \sigma_g$. It is helpful to introduce another relation, which quantifies the ‘amount that f and g cross by at σ ’, whilst f and g are otherwise non-crossing.

Definition 3.3.8 Let $\epsilon > 0$. Let $f, g \in \Pi^\dagger$ and write $\sigma = \sigma_f \vee \sigma_g$. We write $f \blacktriangleleft_\epsilon g$ if $\sigma \in \mathbb{R}$ and $L_{s\bullet}(f) \cap R_{s\bullet}(g) = \emptyset$ for all $s\bullet \geq \sigma +$, as well as $g(\sigma -) + \epsilon \leq f(\sigma -)$ and $f(\sigma +) + \epsilon \leq g(\sigma +)$, with $g(\sigma -) + \epsilon \leq g(\sigma +)$ and $f(\sigma +) + \epsilon \leq f(\sigma -)$.

Lemma 3.3.9 *Let $f, g \in \Pi^\dagger$ and let $\sigma = \sigma_f \vee \sigma_g$.*

1. *If $f \blacktriangleleft_\epsilon g$ for some $\epsilon > 0$ then f and g cross.*
2. *Suppose that $f_n, g_n \in \Pi^\dagger$ with $f_n \rightarrow f$ and $g_n \rightarrow g$. If $f_n \triangleleft g_n$ for all n and $f \not\triangleleft g$ then there exists $\epsilon > 0$ such that $f \blacktriangleleft_\epsilon g$.*

PROOF: We prove the two claims in turn. For the first, suppose that $f \blacktriangleleft_\epsilon g$. Then $g(\sigma -) < g(\sigma +)$, $f(\sigma +) < f(\sigma -)$ and $g(\sigma -) < f(\sigma -)$. Hence $L_{\sigma -}(f) \cap R_{\sigma -}(g) \neq \emptyset$. Since $f(\sigma +) < g(\sigma +)$ we have $L_{\sigma +}(g) \cap R_{\sigma +}(f) \neq \emptyset$. By Lemma 3.1.3, f and g cross.

For the second claim, let $f_n \rightarrow f$, $g_n \rightarrow g$ with $f_n \triangleleft g_n$ and $f \not\triangleleft g$. From Lemma 3.3.5 we have $L_{t\star}(f) \cap R_{t\star}(g) = \emptyset$ and $f(t\star) \leq g(t\star)$ for all $t\star \geq \sigma +$. Since $f \not\triangleleft g$, by Lemma 3.1.2 we must have $L(f) \cap R(g) \neq \emptyset$, which implies that $L_{\sigma -}(f) \cap R_{\sigma -}(g) \neq \emptyset$. Therefore $L_{\sigma -}(f) = [-\infty, f(\sigma -))$, $R_{\sigma -}(g) = (g(\sigma -), \infty]$ with $g(\sigma -) < f(\sigma -)$.

If $g(\sigma +) \leq g(\sigma -)$ then $f(\sigma +) \leq g(\sigma +) \leq g(\sigma -) < f(\sigma -)$ which implies $f \triangleleft g$, so this may not occur. Similarly, if $f(\sigma -) \leq f(\sigma +)$ then $g(\sigma -) < f(\sigma -) \leq f(\sigma +) \leq g(\sigma +)$, which implies $f \triangleleft g$, so this may not occur either. Thus $g(\sigma -) < g(\sigma +)$ and $f(\sigma +) < f(\sigma -)$.

It remains only to show that $f(\sigma+) < g(\sigma+)$. Recall that we have $f(\sigma+) \leq g(\sigma+)$, so we need only eliminate the case $f(\sigma+) = g(\sigma+)$. We will argue by contradiction. We thus assume $g(\sigma-) < g(\sigma+) = f(\sigma+) < f(\sigma-)$. Our strategy is to show that compactness must fail for (f_n) or (g_n) , because in avoiding crossing each other the paths f_n and g_n must become too erratic in a short time interval near σ .

Let

$$\kappa = \min\{g(\sigma+) - g(\sigma-), f(\sigma-) - f(\sigma+)\} \quad (3.2)$$

By right continuity of f and g at $\sigma+$, there exists $\delta > 0$ such that

$$|f(t\star) - f(\sigma+)| \vee |g(t\star) - g(\sigma+)| \leq \kappa/4 \quad \text{for all } t\star \in [\sigma+, (\sigma + \delta)+]. \quad (3.3)$$

By Lemma A.2.4 there exists $s_n \bullet_n$ and $t_n \star_n$ such that $s_n \rightarrow \sigma$, $t_n \rightarrow \sigma$ and $f_n(s_n \bullet_n) \rightarrow f(\sigma-)$ and $g_n(t_n \star_n) \rightarrow g(\sigma-)$. Without loss of generality (or consider the following argument with the roles of f_n and g_n swapped) we may assume that $s_n \bullet_n \leq t_n \star_n$ for infinite many $n \in \mathbb{N}$, and let us pass to a subsequence upon which $s_n \bullet_n \leq t_n \star_n$ holds for all n . In particular, $\sigma_n - \leq s_n \bullet_n$. Without loss of generality we may pass to a further subsequence and assume that $s_n \bullet_n \leq t_n \star_n \leq (\sigma + \delta/3)+$ for all $n \in \mathbb{N}$. Let $u = \sigma + \delta$. By Lemma A.2.4 there exists $u_n \diamond_n$ such that $u_n \rightarrow u$ and $f_n(u_n \diamond_n) \rightarrow f(u+)$. Without loss of generality we may pass to a further subsequence and assume that $u_n \geq \sigma + 2\delta/3$, which implies that

$$s_n \bullet_n \leq t_n \star_n < u_n \diamond_n. \quad (3.4)$$

Again, without loss of generality we may pass to a further subsequence and assume that

$$\begin{aligned} |f_n(s_n \bullet_n) - f(\sigma-)| &\leq \kappa/4 \\ |g_n(t_n \star_n) - g(\sigma-)| &\leq \kappa/4 \\ |f_n(u_n \diamond_n) - f(u+)| &\leq \kappa/4 \end{aligned}$$

for all n . It follows from the above equations, (3.2) and (3.3) that

$$f_n(s_n \bullet_n) \geq f(\sigma+) + 3\kappa/4 \quad (3.5)$$

$$g_n(t_n \star_n) \leq f(\sigma+) - 3\kappa/4 \quad (3.6)$$

$$|f_n(u_n \diamond_n) - f(\sigma+)| \leq \kappa/2 \quad (3.7)$$

for all n .

We must now briefly divide into two cases. If $g_n(t_n-) < g_n(t_n+)$ then $L_{t_n-}(g_n) = [-\infty, g_n(t_n-))$ and by Lemma 3.1.4 we have $f_n(t_n-) \leq g_n(t_n-) < g_n(t_n+)$. Alternatively, if $g_n(t_n+) \leq g_n(t_n-)$ then we have $f_n(t_n+) \leq g_n(t_n+) < g_n(t_n-)$. In either case we have $\circ \in \{-, +\}$ such that $f_n(t_n \circ_n) \leq g_n(t_n-) \wedge g_n(t_n+) \leq g_n(t_n \star_n)$. From (3.6) we thus obtain

$$f_n(t_n \circ_n) \leq f(\sigma+) - 3\kappa/4. \quad (3.8)$$

We must again briefly divide into two cases. If $\star_n = -$ then (3.4) gives $s_n < t_n$, so trivially $s_n \bullet_n \leq t_n \circ_n \leq u_n \diamond_n$. Alternatively, if $\star_n = +$ then we have $f_n(t_n \star_n) \leq g_n(t_n \star_n)$, which from (3.5) and (3.6) we have that $s_n \bullet_n \neq t_n \star_n$. From (3.4) we thus have $s_n \bullet_n \leq t_n -$, so in this case too we obtain that

$$s_n \bullet_n \leq t_n \circ_n < u_n \diamond_n. \quad (3.9)$$

From Proposition A.2.1 and (3.5), (3.7), (3.8), (3.9) we obtain that the sequence (f_n) is not relatively compact. This is a contradiction, and completes the proof. \blacksquare

Lemma 3.3.10 *If $A \subseteq \Pi^\uparrow$ is non-crossing and closed then the relation \triangleleft is compatible with (A, d_Π) . Moreover: suppose that $f_n \rightarrow f$ and $g_n \rightarrow g$, where $f, g, f_n, g_n \in \Pi^\uparrow$ and $f_n \triangleleft g_n$ for all n . If f and g do not cross each other then $f \triangleleft g$.*

PROOF: Note that the second claim is a stronger statement than the first. The second claim permits paths within the sequence (f_n) to cross each other, and paths within the sequence (g_n) to cross each other, requiring only that $f_n \triangleleft g_n$ for each $n \in \mathbb{N}$. We will prove the second claim. Let $f_n \rightarrow f$ and $g_n \rightarrow g$, where $f, g, f_n, g_n \in \Pi^\uparrow$ and $f_n \triangleleft g_n$ for all n . By (both parts of) Lemma 3.3.9, if $f \not\triangleleft g$ then f and g cross. The result follows. \blacksquare

We commented below Definition 3.1.1 that the relation \triangleleft is (in general) not a partial order. In Lemma 3.1.6 we showed that \triangleleft is a total order on A_{\max} , provided $A \subseteq \Pi^\uparrow$ is non-crossing. However, the set A_{\max} is typically not a closed subset of Π^\uparrow , even if A is a closed subset of Π^\uparrow . For example, consider when A contains the paths $f(t) = k$ for $k \in [0, 1]$ and $t \geq \sigma_f = 0$, plus the single path $g(t) = 0$ for $t \geq \sigma_g = -1$. For this reason Lemmas 3.1.6 and 3.3.10 are both important to us, but we must take care when using them together. The following technical lemma will be used in Section 5.1 to help prove that \mathscr{W}_{\det} is measurable.

Lemma 3.3.11 *Let $\epsilon > 0$. Suppose that $f_n \rightarrow f$ and $g_n \rightarrow g$, where $f, g, f_n, g_n \in \Pi^\uparrow$, with $\sigma_f \vee \sigma_g \in \mathbb{R}$. If $f_n \triangleleft_\epsilon g_n$ for all n then $f \triangleleft_\epsilon g$.*

PROOF: We remark that the condition $\sigma_f \vee \sigma_g \in \mathbb{R}$ is necessary, in fact this is all that prevents \triangleleft_ϵ from being compatible with (Π^\uparrow, d_Π) . Let $f, g, f_n, g_n \in \Pi^\uparrow$ be as given and write $\sigma = \sigma_f \vee \sigma_g$. Suppose that $f_n \triangleleft_\epsilon g_n$ for all n . Lemma 3.3.5 gives that $L_{s\bullet}(f) \cap R_{s\bullet}(g) = \emptyset$ for all $s\bullet \geq \sigma +$. Noting that $\sigma_n \rightarrow s$, for all sufficiently large n we have $\sigma_n \in \mathbb{R}$. Let us pass to a subsequence and assume that $\sigma_n \in \mathbb{R}$ for all n .

We have $g_n(\sigma_n -) + \epsilon \leq g_n(\sigma_n +)$. It follows from Lemma A.2.2 that $g(\sigma -) + \epsilon \leq g(\sigma +)$. Similarly, it follows from $f_n(\sigma_n +) + \epsilon \leq f_n(\sigma_n -)$ that $f(\sigma +) + \epsilon \leq f(\sigma -)$. From $g_n(\sigma_n -) + \epsilon \leq f_n(\sigma_n -)$ we obtain

$$\liminf_{n \rightarrow \infty} g_n(\sigma_n -) + \epsilon \leq \limsup_{n \rightarrow \infty} f_n(\sigma_n -). \quad (3.10)$$

By Lemma A.2.2 we have that $\liminf_{n \rightarrow \infty} g_n(\sigma_n -)$ lies between $g(\sigma -)$ and $g(\sigma +)$. We have already shown that $g(\sigma -) < g(\sigma +)$, so in fact $g(\sigma -) \leq \liminf_{n \rightarrow \infty} g_n(\sigma_n -)$. Similarly, Lemma A.2.2 gives that $\limsup_{n \rightarrow \infty} f_n(\sigma_n -)$ lies between $f(\sigma -)$ and $f(\sigma +)$. We have already shown that $f(\sigma +) < f(\sigma -)$, so in fact $\limsup_{n \rightarrow \infty} f_n(\sigma_n -) \leq f(\sigma -)$. From these facts and (3.10) we obtain $g(\sigma -) + \epsilon \leq f(\sigma -)$.

By compactness and Lemma A.2.3 the sequence $(f_n(\sigma_n +))$ has a limit point a between $f(\sigma -)$ and $f(\sigma +)$. We have shown that f jumps leftwards at σ , thus $f(\sigma +) \leq a$. Similarly, $(g_n(\sigma_n +))$ has a limit point b between $g(\sigma -)$ and $g(\sigma +)$. We have shown that g jumps rightwards at σ , thus $b \leq g(\sigma +)$. Using that $f_n(\sigma_n) + \epsilon \leq g_n(\sigma_n +)$ we obtain that $a + \epsilon \leq b$, hence $f(\sigma +) + \epsilon \leq g(\sigma +)$. This completes the proof. \blacksquare

4 Deterministic weaves

Recall the space \mathscr{W}_{\det} introduced in (2.11). An element of \mathscr{W}_{\det} is known as a *deterministic weave* and is, by definition, a deterministic element of $\mathcal{K}(\Pi^\uparrow)$ that is pervasive and non-crossing. We study deterministic weaves in this section, although some results will involve probability within their proofs. The results in this section will feed into the proofs of our main results, in Section 5. Our long term strategy is to establish what can be said in general about the

internal structure of deterministic weaves, to translate this information into statements about the geometric structure of \mathcal{W}_{det} , and finally lift such results into \mathcal{W} .

Definition 2.4.1 defines webs and flows as, respectively, minimal and maximal elements of the space of weaves \mathcal{W} under \preceq_{d} . Recall that elements of \mathcal{W} are formally probability measures on $\mathcal{K}(\Pi)$. We identify \mathcal{W}_{det} with the subset of \mathcal{W} consisting of point-mass measures. It is not immediately clear what Definition 2.4.1 means for deterministic weaves: extremal points of $(\mathcal{W}_{\text{det}}, \preceq)$ are not a priori extremal points of $(\mathcal{W}, \preceq_{\text{d}})$, nor vice versa. We will resolve these difficulties in Lemma 5.2.4, which shows that a random weave \mathcal{W} is a web (resp. flow) if and only if $\mathbb{P}[\mathcal{W} \text{ is almost surely minimal (resp. maximal) in } (\mathcal{W}_{\text{det}}, \preceq)] = 1$. The proof of Lemma 5.2.4 will rely on key results established in Section 4. Therefore, until we have proved Lemma 5.2.4 we will avoid calling any deterministic or random elements of $\mathcal{K}(\Pi)$ a ‘web’ or ‘flow’.

However, we will use the deterministic maps $A \mapsto \text{web}_D(A)$ and $A \mapsto \text{flow}(A)$ defined in (2.12) and (2.13) from this point on. The meaning of these maps on \mathcal{W}_{det} is clear, where in the former case $D \subseteq \mathbb{R}^2$ is also taken to be deterministic.

4.1 Weaves and the non-crossing property

Weaves provide a structure inside of which càdlàg paths behave rather better than within arbitrary subsets of $\mathcal{K}(\Pi)$. We remark that if $f \in \Pi^\uparrow$ does not cross a weave \mathcal{A} then it is trivial to see that $\mathcal{A} \cup \{f\}$ is also a weave. Combining these two facts, to some extent weaves are able to control the behaviour of paths that do not cross them. We begin to explore this idea within the present section. We start with a key technical lemma that uses all of the defining properties of weaves: compactness, pervasiveness and the non-crossing property. It captures what happens when we approximate the middle of a jump with paths beginning earlier in time.

Lemma 4.1.1 *Let \mathcal{A} be a deterministic weave and let $f \in \Pi^\uparrow$ be a path that does not cross \mathcal{A} .*

1. *Suppose $f(t-) < f(t+)$. Then there exists $h \in \mathcal{A}$ such that $h(t-) \leq f(t-)$ and $f \triangleleft h$.*
2. *Suppose $f(t+) < f(t-)$. Then there exists $h \in \mathcal{A}$ such that $f(t-) \leq h(t-)$ and $h \triangleleft f$.*

PROOF: The second statement follows from the first by considering space reflected about the origin (and is written out in full for clarity) so we will prove only the first statement. Suppose $f(t-) < f(t+)$. We will now argue that it suffices to prove that

$$\text{for any } x \in (f(t-), f(t+)) \text{ there exists } h \in \mathcal{A} \text{ such that } h(t-) \leq f(t-) \text{ and } f \triangleleft h. \quad (4.1)$$

With (4.1) in hand, let us write $h^{(x)}$ for the path h generated from x , and note that compactness of \mathcal{A} implies the existence of a subsequential limit $h^{(x)} \rightarrow h'$ as $x \searrow f(t-)$. As $f \triangleleft h^{(x)}$ we have $h^{(x)}(t-) \leq f(t-) < f(t+) \leq h^{(x)}(t+)$, so Lemma A.2.2 ensures that $h'(t-) \leq f(t-)$. Lemma 3.3.10 ensures that $f \triangleleft h'$. Thus h' has the desired properties.

It remains to establish (4.1). Let $x' \in \mathbb{R}$ be such that $f(t-) < x' < f(t+)$, and suppose that f does not cross \mathcal{A} . Let $(a_n) \subseteq (0, \infty)$ be such that $a_n \rightarrow 0$. Let $z_n = (x, t - a_n)$ and by pervasiveness let $h_n \in \mathcal{A}(z_n)$. By compactness of \mathcal{A} , pass to a subsequence and assume without loss of generality that $h_n \rightarrow h \in \mathcal{A}$. By Lemma A.2.3 we have

$$h(t-) \wedge h(t+) \leq x' \leq h(t-) \vee h(t+). \quad (4.2)$$

We have that f and $h_n \in \mathcal{A}$ do not cross. By Lemma 3.1.3 this means that $f \triangleleft h_n$ or $h_n \triangleleft f$. We now consider two cases.

Consider first if $f \triangleleft h_n$ for infinitely many $n \in \mathbb{N}$. Then Lemma 3.3.10 implies that $f \triangleleft h$. In this case $f(t+) \leq h(t+)$, so $x < h(t+)$ which by (4.2) implies $h(t-) \leq x'$. Hence $h(t-) \leq y$, and we have established all the required properties of h .

If the above case does not occur then there exists $N \in \mathbb{N}$ such that $h_n \triangleleft f$ for all $n \geq N$. Without loss of generality we may pass to a subsequence and assume $h_n \triangleleft f$ for all $n \in \mathbb{N}$. As $\sigma_{h_m} < t$ and $L_{t-}(f) = [-\infty, f(t-))$, Lemma 3.1.4 implies that $h_n(t-) \leq f(t-)$. We have also that $h_n((t - a_n)-) \vee h_n((t - a_n)+) \geq x$. Clearly $(t - a_n)\pm < t-$ for all n . Hence by Lemma A.2.3 we have $h(t-) \geq x > f(t-) \geq h(t+)$, which as $f(t-) < f(t+)$ means that f and h cross (by Definition 2.3.1). This is a contradiction, so in fact this case does not occur. This completes the proof. \blacksquare

Lemma 4.1.2 *Let \mathcal{A} be a deterministic weave. Let $f, h \in \Pi^\uparrow$ be such that $\mathcal{A} \cup \{f\}$ is non-crossing, and $\mathcal{A} \cup \{h\}$ is non-crossing. Let $\sigma = \sigma_f \vee \sigma_h$. Suppose that $f(t_\star) < h(t_\star)$ for some $t_\star \geq \sigma+$. Then for all $\epsilon > 0$ there exists $g \in \mathcal{A}$ with $\sigma_g \leq t + \epsilon$ such that $f \triangleleft g$ and $g \triangleleft h$.*

Further, we may choose g such that $g \not\subseteq f$, $f \not\subseteq g$, $g \not\subseteq h$ and $h \not\subseteq g$.

PROOF: By the càdlàg property of f and h for any $\epsilon > 0$ there exists t' such that $|t - t'| < \epsilon/2$ and $f(t'+) < h(t'+)$. Hence we may choose $\epsilon > 0$ chosen sufficiently small that $f(s\pm) < h(s\pm)$ for all $s \in [t', t' + \epsilon/2)$, with $|t - t'| < \epsilon/2$ and $t' > \sigma$. By the càdlàg property of f and h , choose $s \in (t', t' + \epsilon)$ such that both f and h are continuous at s . Note that $s \leq t + \epsilon$. Recall that when f is continuous at s we write $f(s) = f(s\pm)$. Let $x \in (f(s), h(s))$ and by pervasiveness of \mathcal{A} let $g \in \mathcal{A}(x, s)$. Hence $g(s-) \wedge g(s+) \leq x \leq g(s-) \vee g(s+)$. By continuity of f at s we have $L_{s-}(f) = L_{s+}(f) = [-\infty, f(s))$. Note that $f(s) < h(s)$ implies that at least one of $g(s-)$ and $g(s+)$ is strictly greater than $f(t_\star)$. By Lemma 3.1.4 we thus have $f \triangleleft g$, and as $\sigma_f < s$ this means $f \not\subseteq g$ and $g \not\subseteq f$. A symmetrical argument (reflect space about the origin) shows that $g \triangleleft h$, with $g \not\subseteq h$ and $h \not\subseteq g$. \blacksquare

Lemma 4.1.3 *Let $\mathcal{A} \subseteq \Pi^\uparrow$ be a deterministic weave. If $B, C \subseteq \Pi^\uparrow$ are such that $\mathcal{A} \cup B$ is non-crossing, and $\mathcal{A} \cup C$ is non-crossing, then $\mathcal{A} \cup B \cup C$ is non-crossing.*

PROOF: Note that our conditions imply that B is non-crossing, and C is non-crossing. It suffices to prove the case where $B = \{f\}$ and $C = \{h\}$ are singletons, from which the general case follows immediately. To this end, suppose that $f, g \in \Pi^\uparrow$ are such that $\mathcal{A} \cup \{f\}$ is non-crossing and $\mathcal{A} \cup \{h\}$ is non-crossing.

We will argue by contradiction. Suppose that $\mathcal{A} \cup \{f\} \cup \{h\}$ contains a pair of paths that cross. From our assumptions, the only possibility is that f and h cross. By Lemma 3.1.3 f and h cross if and only if $L(g) \cap R(f) \neq \emptyset$ and $L(f) \cap R(h) \neq \emptyset$. Hence there exists $t_\star, s_\bullet \in \mathbb{R}_s$ such that $t_\star < s_\bullet$, with

$$\begin{aligned} L_{t_\star}(h) &= [-\infty, h(t_\star)), & R_{t_\star}(f) &= (f(t_\star), \infty], \\ L_{s_\bullet}(f) &= [-\infty, f(s_\bullet)), & R_{s_\bullet}(h) &= (h(s_\bullet), \infty], \end{aligned}$$

$f(t_\star) < h(t_\star)$ and $h(s_\bullet) < f(s_\bullet)$. Let $\sigma = \sigma_f \vee \sigma_h$.

Consider, first, if $t_\star \geq \sigma+$. In this case, Lemma 4.1.2 implies that there exists $g \in \mathcal{A}$ such that $f \triangleleft g$ and $g \triangleleft h$, with $\sigma_g < s_\bullet$. As $s_\bullet \geq \sigma+$ this means $f(s_\bullet) \leq g(s_\bullet) \leq h(s_\bullet)$, which is a contradiction to $h(s_\bullet) < f(s_\bullet)$.

It remains to consider the case $t_\star = \sigma-$. In this case $\sigma_f = \sigma$ or $\sigma_h = \sigma$. Without loss of generality (or consider space reflected about the origin) let us assume that $\sigma_f = \sigma$. As $R_{\sigma-} = (f(\sigma-), \infty]$ and $\sigma_f = \sigma$ it follows from (3.1) that $f(\sigma-) < f(\sigma+)$. Lemma 4.1.1 implies

the existence of $g \in \mathcal{A}$ such that $f \triangleleft g$ and $g(\sigma-) < f(\sigma+)$. Hence $g(\sigma-) < h(\sigma-)$. If $\sigma_h < \sigma$ then it follows immediately by Lemma 3.1.4 that $g \triangleleft h$. Alternatively, if $\sigma_h = s$ then $L_{t\star}(h) = [-\infty, h(\sigma-))$ so in this case too we have $g \triangleleft h$. We now have $f \triangleleft g$ and $g \triangleleft h$. As $s\bullet \geq \sigma+$ this means $f(s\bullet) \leq g(s\bullet) \leq h(s\bullet)$, which is a contradiction to $h(s\bullet) < f(s\bullet)$. ■

Lemma 4.1.4 *Let \mathcal{A}, \mathcal{B} be deterministic weaves and suppose that $\mathcal{A} \cup \mathcal{B}$ is non-crossing. Let $C \subseteq \Pi^\dagger$ be non-crossing. Then $\mathcal{A} \cup C$ is non-crossing if and only if $\mathcal{B} \cup C$ is non-crossing.*

PROOF: It suffices to consider the case $C = \{f\}$ where $f \in \Pi^\dagger$, from which the general case follows immediately. Assume that $f \in \Pi^\dagger$ does not cross \mathcal{A} . Let $b \in \mathcal{B}$. By assumption b does not cross \mathcal{A} . By Lemma 4.1.3, we have that $\mathcal{A} \cup \{b\} \cup \{f\}$ is non-crossing, so in particular f and b do not cross each other. Since $b \in \mathcal{B}$ was arbitrary, f does not cross \mathcal{B} . ■

In the proof of Lemma 4.1.3 we saw that Lemma 4.1.2 was a natural counterpart to Lemma 4.1.1. The underlying principle is as follows. If two paths f and g are such that $f \not\subseteq g$ and $g \not\subseteq f$ then either: $f(t+) \neq g(t+)$ for some t , or at least one of f and g has made a jump at its initial time, in a direction away from the other. Lemma 4.1.2 applied to the former case, Lemma 4.1.1 to the latter, resulting in a path h that lay between f and g . Variations upon this theme will feature in the proof of several future results, including the next lemma.

The following lemma is stated for paths $f, h \in \Pi^\dagger$ that do not cross a weave \mathcal{A} , but at this stage it is perhaps best understood by considering the special case $f, h \in \mathcal{A}_{\max}$. Note that for maximal paths, the condition $f \not\subseteq h$ and $h \not\subseteq f$ is simply the requirement that $f \neq h$. In this case Lemma 4.1.5 provides a key piece of information about the geometric structure of $(\mathcal{A}_{\max}, \triangleleft)$, namely that any two distinct points, within the total order, will always have another point strictly in between them. This lemma will be a key tool in Section 4.3.

Lemma 4.1.5 *Let \mathcal{A} be a deterministic weave. Suppose $f, h \in \Pi^\dagger$ do not cross \mathcal{A} , with $f \triangleleft h$, $f \not\subseteq h$ and $h \not\subseteq f$. Then there exists $g \in \mathcal{A}_{\max}$ such that $f \triangleleft g$ and $g \triangleleft h$, and $f \not\subseteq g$, $g \not\subseteq f$, $g \not\subseteq h$, $h \not\subseteq g$.*

PROOF: Let f, h be as given in the lemma and set $\sigma = \sigma_f \vee \sigma_h$. By part 2 of Lemma 3.1.5 our conditions on f and h imply that $h \not\triangleleft f$. From Lemma 3.1.2 we thus have $L(h) \cap R(f) \neq \emptyset$. In particular there exists $t\star \geq \sigma-$ such that $L_{t\star}(h) = [-\infty, h(t\star))$ and $R_{t\star}(f) = (f(t\star), \infty]$ with $f(t\star) < h(t\star)$.

Consider first if $t\star \geq \sigma+$. Then Lemma 4.1.2 implies the existence of $g \in \mathcal{A}$ with $f \triangleleft g$ and $g \triangleleft f$, also $f \not\subseteq g$, $g \not\subseteq f$, $g \not\subseteq h$ and $h \not\subseteq g$. Without loss of generality we may take $g \in \mathcal{A}_{\max}$, which completes the proof in this case.

It remains to consider the case $t\star = \sigma-$. In this case we have $\sigma = \sigma_f$ or $\sigma = \sigma_h$. Without loss of generality let us assume that $\sigma = \sigma_f$ (or consider space reflected about the origin). As $R_{\sigma-}(f)$ is non-empty this implies that $f(t-) < f(t+)$.

Remark 4.1.6 Let us briefly comment on the strategy for the remainder of the proof. Although f jumps at σ , Lemma 4.1.1 is not suitable for use here because (if used to construct g) it allows the possibility that $g \subseteq f$. Instead, we require a more sophisticated version of the approximation scheme used in the proof of Lemma 4.1.1, but the path we are looking for here is *not* the limiting path; rather it is some path that occurs sufficiently close to the limit. We require a path with several different properties. To find it, we will repeatedly show that one such desired property can fail only for finitely many n , then (without loss of generality) pass to a subsequence on which the property holds for all n .

As $t\star = \sigma-$ we have $f(\sigma-) < h(\sigma-)$, so

$$f(\sigma-) < f(\sigma+) \wedge h(\sigma-). \quad (4.3)$$

Let $\epsilon > 0$ be such that $f(\sigma-) + 2\epsilon \leq f(\sigma+) \wedge h(\sigma-)$ and let $(a_n) \subseteq (0, \infty)$ be such that $a_n \rightarrow 0$. Let $(y, s_n) = (f(\sigma-) + \epsilon, \sigma - a_n)$ and note that $(y, s_n) \rightarrow (y, \sigma)$ where $y = f(\sigma-) + \epsilon$. By pervasiveness of \mathcal{A} let $g_n \in \mathcal{A}((y, s_n))$ so that

$$g_n(s_n-) \wedge g_n(s_n+) \leq y \leq g_n(s_n-) \vee g_n(s_n+). \quad (4.4)$$

Without loss of generality we may take $g_n \in \mathcal{A}_{\max}$. By compactness of \mathcal{A} we may pass to a subsequence and assume that $g_n \rightarrow g \in \mathcal{A}$.

Consider if $f \subseteq g_n$ for infinitely many n . For such n , noting from (4.3) that f jumps rightwards at σ , we have $g_n(\sigma-) \leq f(\sigma-)$. From (4.4) we have $g_n(s_n-) \vee g_n(s_n+) \geq y = f(\sigma-) + \epsilon$, and $s_n \pm < \sigma-$ with $s_n \rightarrow \sigma$, so Lemma A.2.3 gives that g jumps leftwards at σ , from right of $f(\sigma-) + \epsilon$ to left of $f(\sigma-)$. This would make f and $g \in \mathcal{A}$ cross, which is a contradiction. Hence in fact $f \not\subseteq g_n$ for at most finitely many n , so we may pass to a subsequence and assume $f \not\subseteq g_n$ for all n . If $g_n \subseteq f$ then we would have $\sigma_{g_n} \geq \sigma_f$, which is not the case because $\sigma_{g_n} \leq s_n < \sigma_f$. Hence we also have $g_n \not\subseteq f$ for all n .

Consider if $g_n \triangleleft f$ for infinitely many $n \in \mathbb{N}$. For such n , noting that we have $f \not\subseteq g_n$ and $g_n \not\subseteq f$, by part 2 of Lemma 3.1.5 we have $f \not\triangleleft g_n$. Hence, for such n , using that $R_{\sigma-}(f) = (f(\sigma-), \infty]$, by Lemma 3.1.4 we must have $g_n(\sigma-) \leq f(\sigma-)$. From (4.4) we have $y \leq g_n(s_n-) \vee g_n(s_n+)$, where $s_n < \sigma$ with $s_n \rightarrow \sigma$. By Lemma A.2.3, taking a limit along a subsequence of such n would result in $g(\sigma-) \geq y > f(\sigma-) \geq g(\sigma+)$, in which case f and g cross (by jumping over each other in opposite directions at time σ). This may not occur. Hence in fact $g_n \triangleleft f$ for at most finitely many n . By Lemma 3.1.3 for all n we have $f \triangleleft g_n$ or $g_n \triangleleft f$. We may thus pass to a subsequence and assume that $f \triangleleft g_n$ for all n .

We now have $f \triangleleft g_n$, $f \not\subseteq g_n$ and $g_n \not\subseteq f$. We will move on to establishing properties of g_n with h . Here we divide into two cases, based upon whether $h(\sigma-) \leq h(\sigma+)$ or $h(\sigma+) < h(\sigma-)$.

- Firstly, consider if $h(\sigma+) < h(\sigma-)$.

Consider if $h(\sigma-) \leq g_n(\sigma-)$ for infinitely many n . From (4.4) we have $g_n(s_n-) \wedge g_n(s_n+) \leq y$ and $s_n \pm < \sigma-$ with $s_n \rightarrow \sigma$, so by Lemma A.2.3 we obtain that g jumps rightwards at σ , from left of y to right of $h(\sigma-)$. This means that $g \in \mathcal{A}$ crosses h , which is a contradiction. Therefore we may pass to a subsequence and assume that $g_n(\sigma-) < h(\sigma-)$ for all n .

We have $\sigma_g \leq s_n < \sigma$ and $\sigma_h \leq \sigma$. If $h \subseteq g_n$ then, noting that h jumps leftwards at σ , we would have $h_n(\sigma-) \leq g_n(\sigma-)$, which is a contradiction. Hence $h \not\subseteq g_n$. Similarly, if $g_n \subseteq h$ then $\sigma_h \leq \sigma_g < \sigma$, so we would have $g_n(\sigma-) = h_n(\sigma-)$, which is a contradiction, so $g_n \not\subseteq h$.

Consider if $h \triangleleft g_n$. We have already seen that $h \not\subseteq g_n$ and $g_n \not\subseteq h$, so by part 2 of Lemma 3.1.5 we have $g_n \not\triangleleft h$. We have $L_{\sigma-}(h) = [-\infty, h(\sigma-))$ so Lemma 3.1.4 gives that $h(\sigma-) \leq g_n(\sigma-)$, which again may not occur. Hence in fact $g_n \triangleleft h$.

- Secondly, consider if $h(\sigma-) \leq h(\sigma+)$.

Our assumption $f \triangleleft h$ implies that $f(\sigma\bullet) \leq h(\sigma\bullet)$ for all $\sigma\bullet \geq \sigma+$, so from what we have already proved (in the case $t\star \geq \sigma+$) we may assume without loss of generality that $f(\sigma\bullet) = h(\sigma\bullet)$ for all $\sigma\bullet \geq \sigma+$. Using that $h(\sigma-) \leq h(\sigma+)$ we must therefore have $\sigma_h < \sigma$, as otherwise by (4.3) we would have $h \subseteq f$.

By left continuity of h at $\sigma-$ there exists some $\delta > 0$ such that $h(s\bullet) \geq h(\sigma-) - \epsilon/2$ for all $s\bullet \in [\sigma-, (\sigma - \delta)+]$. From (4.4) for all sufficiently large n we have $s_n \in (\sigma - \delta, s)$ and

$$g_n(s_n-) \wedge g_n(s_n+) \leq y < h(\sigma-) - \epsilon/2 \leq h(s_n\pm). \quad (4.5)$$

It follows immediately that, for such n , $g_n \not\leq h$ and $h \not\leq g_n$. For sufficiently large n we also have $s_n > \sigma_h$, in which case (4.5) gives $g_n(s_n-) \in L_{s_n-}(h)$ or $g_n(s_n+) \in L_{s_n+}(h)$. Hence Lemma 3.1.4 gives $g_n \triangleleft h$. We may thus pass to a subsequence and assume that $g_n \not\leq h$, $h \not\leq g_n$ and $g_n \triangleleft h$ for all n .

In both cases we have now shown, for all g_n within the subsequence that we have passed into, that $f \triangleleft g_n$, $g_n \triangleleft h$, and that none of f, g_n, h are \subseteq -comparable with each other. Therefore, any such g_n has the required properties and the proof of the present lemma is complete. \blacksquare

4.2 Weaves of bi-infinite paths

In this section we establish some geometric properties of deterministic weaves that comprise entirely of bi-infinite paths. One key result is that if $\mathcal{A} \subseteq \Pi^\uparrow$ is a deterministic weave then set of ramification points of \mathcal{A} has zero (two dimensional) Lebesgue measure. This result will be later extended to all deterministic weaves, in Lemma 4.4.6. Whenever we refer to Lebesgue measure in this section, we mean two dimensional Lebesgue measure on \mathbb{R}_c^2 .

Lemma 4.2.1 *For each $f \in \Pi$, the set $H(f)$ has zero Lebesgue measure*

PROOF: This lemma is almost self-evident but in view of the example of Jordan curves with positive Lebesgue measure we will give a short proof. Since càdlàg functions have only countably many discontinuities, the result holds for the part of $H(f)$ corresponding to jumps of f . The remaining part of $H(f)$ can be shown to have zero Lebesgue measure via Fubini's theorem. We leave the details to the reader. \blacksquare

Lemma 4.2.2 *Let $\mathcal{A} \subseteq \Pi^\uparrow$ be a deterministic weave and $h \in \mathcal{A}$.*

1. *If $h(t\star) > -\infty$ for some $t\star \in \mathbb{R}_s$ then there exist a strictly monotone sequence (f_n) with $f_n \triangleleft f_{n+1}$ such that $f_n \rightarrow h$.*
2. *If $h(t\star) < \infty$ for some $t\star \in \mathbb{R}_s$ then there exist a strictly monotone sequence (g_n) with $g_{n+1} \triangleleft g_n$ such that $g_n \rightarrow h$.*

PROOF: We will show only the existence and properties of (f_n) . The corresponding statements for (g_n) follow by symmetry. By Lemma 3.1.6 $(\mathcal{A}, \triangleleft)$ is totally ordered. Let $h \in \mathcal{A}$ and set $L = \{f' \in \mathcal{A}; f' \triangleleft h, f' \neq h\}$. As $h(\star) > -\infty$, by the càdlàg property of $h \in \Pi^\uparrow$ there exists some $s \in \mathbb{R}$ such that h is continuous at s and $h(s) > -\infty$. Taking $g \in \mathcal{A}((h(s)/2, s))$ gives $g \in L$, so L is non-empty. By compactness of \mathcal{A} the set \bar{L} is compact, which by Lemma 3.3.10 implies that (\bar{L}, \triangleleft) contains a unique maximal element f . By Lemma 3.3.10 we have $f \triangleleft h$.

Suppose, in preparation for an argument by contradiction, that $f \neq h$. Since f, h are both bi-infinite there exists $t \in \mathbb{R}$ such that $f(t+) < h(t+)$. By Lemma 4.1.2 and using that $\mathcal{A} \subseteq \Pi^\uparrow$, there exists $g \in \mathcal{A}$ such that $f \neq g$ and $g \neq h$. This is a contradiction to maximality of f in \bar{L} .

We thus have $f = h$, which by definition of L implies that there exists $(f_n) \subseteq L$ such that $f_n \rightarrow h$, with $f_n \neq h$ for all n . Without loss of generality we may choose a strictly monotone subsequence, which completes the proof. \blacksquare

Lemma 4.2.2 fails for general deterministic weaves, which may contain paths that are isolated points (from the left, right or both). For example see the weave \mathcal{A} on the right hand side of Figure 2.5.1. We will shortly show, as a consequence of Lemma 4.2.4, that if a deterministic weave consists entirely of bi-infinite paths then it does not contain any isolated points.

Lemma 4.2.3 *Let $\mathcal{A} \subseteq \Pi^\uparrow$ be a deterministic weave. The order topology induced on \mathcal{A} by the total order \triangleleft coincides with its topology as a subspace of Π .*

PROOF: By Lemma 3.1.6, $(\mathcal{A}, \triangleleft)$ is totally ordered. Recall that the order topology on \mathcal{A} is generated by the open rays

$$\begin{aligned} R_f &= \{g \in \mathcal{A}; g \triangleleft f \text{ and } f \neq g\}, \\ R'_f &= \{g \in \mathcal{A}; f \triangleleft g \text{ and } f \neq g\} \end{aligned}$$

where $f \in \mathcal{A}$. We will show that R_f is open in the M1 topology on \mathcal{A} . The same result follows for R'_f by a symmetrical argument. Note that if f is the bi-infinite path with constant value at $-\infty$ then $R_f = \emptyset$ and $R'_f = \mathcal{A} \setminus \{f\}$, which are automatically open. Similar considerations apply to if f is the bi-infinite path with constant value ∞ . We may therefore restrict to $f \in \mathcal{A}$ such that $f(t_\star) > -\infty$ for some $t_\star \in \mathbb{R}_g$.

By Lemma 3.1.6 \mathcal{A} is totally ordered, from which it follows that $\mathcal{A} \setminus R_g = \{f \in \mathcal{A}; g \triangleleft f\}$. By Lemma 3.3.10 this is a closed subset of \mathcal{A} in the M1 topology, thus R_g is an open subset of \mathcal{A} in the M1 topology. It follows that any subset of \mathcal{A} that is open in the order topology is also open in the M1 topology, and it remains to prove the converse.

It suffices to show that if $B \subseteq \mathcal{A}$ is closed in the M1 topology, then it is also closed in the order topology. Let us write $\xrightarrow{M1}$ for convergence in the M1 topology and $\xrightarrow{\triangleleft}$ for convergence in the order topology. Let $B \subseteq \mathcal{A}$ be closed in the M1 topology i.e. if $f_n \in B$ and $f_n \xrightarrow{M1} f \in \mathcal{A}$ then $f \in B$. Suppose that $h_n \in B$ and $h_n \xrightarrow{\triangleleft} h \in \mathcal{A}$. By Lemma 4.2.2 there exists $f_n, g_n \in \mathcal{A}$ such that $f_n \triangleleft f_{n+1}$, $g_{n+1} \triangleleft g_n$ for all n , and $f_n \xrightarrow{M1} h$, $g_n \xrightarrow{M1} h$ as $n \rightarrow \infty$. For each $m \in \mathbb{N}$ the set

$$O_m = \{h' \in \mathcal{A}; f_m \triangleleft h', h' \triangleleft g_m, f_m \neq h', h' \neq g_m\}$$

is an open interval in the order topology, hence there exists $N_m \in \mathbb{N}$ such that for all $n \geq N_m$ we have $h_n \in O_{N_m}$. Without loss of generality we may assume $N_m \rightarrow \infty$ as $m \rightarrow \infty$. We thus have that for all $n \geq N_m$

$$f_m \triangleleft h_n \quad \text{and} \quad h_n \triangleleft g_m. \tag{4.6}$$

Let h' be any limit point of (h_n) in the M1 topology, thus $h' \in \Pi^\uparrow$. Letting $m \rightarrow \infty$ in (4.6), by Lemma 3.3.10 we obtain $h \triangleleft h'$ and $h' \triangleleft h$. Since both are bi-infinite, we have $h = h'$. It follows that $h_n \xrightarrow{M1} h$, which (since B is closed in the M1 topology) shows that $h \in B$, and thus completes the proof. \blacksquare

Lemma 4.2.4 *Let $\mathcal{A} \subseteq \Pi^\uparrow$ be a deterministic weave. There exists an order preserving homeomorphism ϕ between the totally ordered spaces $(\mathcal{A}, \triangleleft)$ and $([0, 1], \leq)$.*

PROOF: Note that the result of Lemma 4.2.3 is implicit in the statement of the present lemma. Throughout the proof we will use the result of Lemma 3.1.6, that $(\mathcal{A}, \triangleleft)$ is totally ordered. For $f \in \Pi^\uparrow$ we define

$$\begin{aligned} H^-(f) &= \{(x, t); t \in x < f(t-) \wedge f(t+)\}, \\ H^+(f) &= \{(x, t); t \in x > f(t-) \vee f(t+)\}. \end{aligned}$$

Note that $\overline{\mathbb{R}} \times \overline{\mathbb{R}} = H^-(f) \cup H(f) \cup H^+(f)$ and that this union is disjoint. Recall that $d_{\mathbb{R}_c^2}$ is a metric that generates the topology on \mathbb{R}_c^2 and recall that for $A \subset \mathbb{R}_c^2$ the open ϵ -expansion of A is given by $A^{(\epsilon)} = \{z \in \mathbb{R}_c^2 : \text{dist}_{\mathbb{R}_c^2}(z, A) < \epsilon\}$, where $\text{dist}(z, A) = \inf_{a \in A} d_{\mathbb{R}_c^2}(z, a)$. Let μ be a measure on \mathbb{R}_c^2 that is absolutely continuous with respect to Lebesgue measure, with full support. Let $\phi(f) = \mu(H^-(f))$.

It is immediate that ϕ is non-decreasing and that $\phi(f_{-\infty}) = 0$, $\phi(f_{\infty}) = 1$ where $f_{\pm\infty}$ are the constant paths at $\pm\infty$ (it is trivial to check that $f_{\pm\infty} \in \mathcal{A}$). We next show that ϕ is continuous. Assume that $f_n, f \in \mathcal{A}$ and that $f_n \rightarrow f$. From our remarks above Proposition 2.2.1 we have $H(f_n) \rightarrow H(f)$ in the Hausdorff metric induced by d . Hence, for each $\epsilon > 0$, for sufficiently large n we have both $H^-(f_n) \subseteq H^-(f)^{(\epsilon)}$ and $H^+(f_n) \subseteq H^+(f)^{(\epsilon)}$. Thus

$$\limsup_{n \rightarrow \infty} \mu(H^-(f_n)) \leq \mu(H^-(f)^{(\epsilon)}) \quad \text{and} \quad \limsup_{n \rightarrow \infty} \mu(H^+(f_n)) \leq \mu(H^+(f)^{(\epsilon)}).$$

Letting $\epsilon \downarrow 0$ and using the fact that $\phi_\mu(\pi) = \mu(H^-(\pi)) = 1 - \mu(H^+(\pi))$ by Lemma 4.2.1, we see that $\phi_\mu(\pi_n) \rightarrow \phi_\mu(\pi)$ as $n \rightarrow \infty$.

Our next goal is to show that ϕ_μ is a bijection. From what we have already proved, ϕ is surjective and non-decreasing, so it suffices to prove that if $f \triangleleft g$ with $f \neq g$ then $\phi(f) < \phi(g)$. If $f \triangleleft g$ are not equal then, since both are bi-infinite, there exists $t \in \mathbb{R}$ such that $f(t+) < g(t+)$, from which it follows by right continuity that the set $O = \{(x, s); f(t+) < x < g(t+)\}$ has non-empty interior, and thus positive Lebesgue measure. Since $O \subseteq H^-(g) \setminus H^-(f)$ and μ is absolutely continuous with Lebesgue measure, we have $\phi(f) < \phi(g)$, as required.

We have now shown that ϕ is a continuous bijection from the compact space \mathcal{A} to (the Hausdorff topological space) $[0, 1]$, which implies that ϕ is a homeomorphism. The fact that ϕ is non-decreasing and bijective implies that ϕ is order preserving. \blacksquare

Lemma 4.2.5 *The function $(\mathcal{A}, z) \rightarrow \mathbb{1}\{z \text{ is ramified in } \mathcal{A}\}$ is measurable from $\mathcal{K}(\Pi) \times \mathbb{R}_c^2 \rightarrow \{0, 1\}$. Moreover, for any deterministic weave $\mathcal{A} \subseteq \Pi^\uparrow$ the set of ramification points of \mathcal{A} has zero Lebesgue measure.*

PROOF: Let $\mathcal{A} = \{B \in \mathcal{K}(\Pi); \exists f \in \mathcal{A} \text{ with } B \subseteq \{f\}_\uparrow\}$ and let $\text{ram}(\mathcal{A}) \subseteq \mathbb{R}_c^2$ denote the set of ramification points of \mathcal{A} . Note that $z \in \mathbb{R}_c^2$ is non-ramified if and only if $\mathcal{A}(z) \in \mathcal{A}$. It is straightforward to check that $\mathcal{A} \subseteq \mathcal{K}(\Pi)$ is closed, as a consequence of Lemma A.2.2. From Lemma A.3.1 the map $(\mathcal{A}, z) \mapsto \mathcal{A}(z)$ from $\mathcal{K}(\Pi) \times \mathbb{R}_c^2 \rightarrow \mathcal{K}(\Pi)$ is measurable. We have $\mathbb{1}\{z \in \text{ram}(\mathcal{A})\} = \mathbb{1}\{\mathcal{A}(z) \in \mathcal{A}\}$, hence that the map $(\mathcal{A}, z) \mapsto \mathbb{1}\{z \in \text{ram}(\mathcal{A})\}$ is measurable. It follows immediately that $\text{ram}(\mathcal{A})$ is a measurable subset of \mathbb{R}_c^2 , for any $\mathcal{A} \in \mathcal{W}_{\text{det}}$.

It remains to show that the measure of $\text{ram}(\mathcal{A})$ is zero. Take $\phi : \mathcal{A} \rightarrow [0, 1]$ as in the statement of Lemma 4.2.4, and let $z \in \overline{\mathbb{R}} \times \overline{\mathbb{R}}_5$. By Lemma A.2.2 the set $\mathcal{A}(z)$ is closed. By definition of \triangleleft (and the fact that $\mathcal{A} \subseteq \Pi^\uparrow$) we have that $\mathcal{A}(z)$ is an interval of the totally ordered space $(\mathcal{A}, \triangleleft)$. Thus $\mathcal{A}(z)$ is a closed interval of $(\mathcal{A}, \triangleleft)$ and $\phi(\mathcal{A}(z))$ is a closed interval of $[0, 1]$.

Let U be uniformly distributed on $[0, 1]$. Note that $z \in \mathbb{R}_c^2$ is ramified if and only if the closed interval $\mathcal{A}(z)$ is more than just a single point, which occurs if and only if $\mathbb{P}[U \in \phi(\mathcal{A}(z))] > 0$. Let μ be a measure on \mathbb{R}_c^2 that is absolutely continuous with respect to Lebesgue measure. Let Z be a random variable with law μ , independently of U . By Lemma A.3.1 $\mathcal{A}(Z)$ is a $\mathcal{K}(\Pi)$ valued random variable. Then $\mathbb{P}[U \in \phi(\mathcal{A}(Z))] = \mathbb{P}[\phi^{-1}(U) \in \mathcal{A}(Z)]$ is the probability that the random path $\phi^{-1}(U)$ passes through the random point Z . Recalling that μ is absolutely continuous with respect to Lebesgue measure, by Lemma 4.2.1 this probability is zero. Thus Z is almost surely not ramified, which implies that the set of ramification points has zero Lebesgue measure. \blacksquare

Lemma 4.2.6 *Let $\mathcal{A} \subseteq \Pi^\uparrow$ be a deterministic weave. If $h \in \Pi^\uparrow$ does not cross \mathcal{A} then $h \in \mathcal{A}$.*

PROOF: We will argue by contradiction. Suppose that $h \in \Pi^\dagger$ does not cross \mathcal{A} and that $h \notin \mathcal{A}$. Since \mathcal{A} is closed, this means that h is an isolated point of $\mathcal{A} \cup \{h\}$. Since h does not cross \mathcal{A} it is straightforward to check that $\mathcal{A} \cup \{h\}$ is a deterministic weave. Thus, from Lemma 4.2.4 we have that $\mathcal{A} \cup \{h\}$ does not contain any isolated points. This is a contradiction, which completes the proof. ■

Remark 4.2.7 By Lemma 4.2.6 a deterministic weave $\mathcal{A} \subseteq \Pi^\dagger$ always contains a constant path at spatial location $-\infty$, similarly at $+\infty$. These are the minimal and maximal elements of $(\mathcal{A}, \triangleleft)$.

Lemma 4.2.8 Let $\mathcal{A} \subseteq \Pi^\dagger$ be a deterministic weave. If $D \subseteq \mathbb{R}^2$ is dense then $\mathcal{A} = \overline{\mathcal{A}(D)}$.

PROOF: It is immediate that $\overline{\mathcal{A}(D)} \subseteq \mathcal{A}$. Thus $\overline{\mathcal{A}(D)}$ is compact. Since D is dense, it is easily seen from Lemma A.2.2 that $\overline{\mathcal{A}(D)}$ is pervasive, and since $\overline{\mathcal{A}(D)} \subseteq \mathcal{A}$ it is also non-crossing. Hence $\overline{\mathcal{A}(D)}$ is a deterministic weave. If $h \in \mathcal{A}$ then since $\overline{\mathcal{A}(D)} \subseteq \mathcal{A}$ we have that h does not cross $\overline{\mathcal{A}(D)}$. Hence by Lemma 4.2.6 we have $h \in \overline{\mathcal{A}(D)}$. Thus $\mathcal{A} = \overline{\mathcal{A}(D)}$. ■

4.3 Extensions of paths in weaves

Our results on weaves rely fundamentally on the fact that, within a weave, half-infinite paths may be extended into bi-infinite paths, without inducing crossing. We also require that such extensions preserve compactness; this point will be addressed later on in Lemma 4.4.2. As one might expect, bi-infinite extensions of paths are closely related to maximal paths under the partial order \subseteq of path extension. From Lemma 3.1.6, if \mathcal{A} is a weave then $(\mathcal{A}_{\max}, \triangleleft)$ is totally ordered. It is tempting to hope that bi-infinite extensions of paths could be constructed via taking a suitable limit of paths in \mathcal{A} , but in general this is not possible because deterministic weaves are closed sets. A more delicate operation is required.

Definition 4.3.1 Let \mathcal{A} be a deterministic weave. We say that a subset $X \subseteq \mathcal{A}_{\max}$ is a *Dedekind cut* of \mathcal{A}_{\max} if (i) whenever $f, g \in \mathcal{A}_{\max}$ with $f \triangleleft g$ and $g \in X$ we have $f \in X$ and (ii) X has no maximal element.

Let us outline the main results within this section. Recall that $\text{flow}(\mathcal{A})$ denotes the set of bi-infinite càdlàg paths that do not cross \mathcal{A} . Dedekind cuts are best known as part of Dedekind's construction of \mathbb{R} from \mathbb{Q} . A related situation presents itself here, in which \mathcal{A}_{\max} plays the role of \mathbb{Q} and $\text{flow}(\mathcal{A})$ plays the role of \mathbb{R} . Specifically: in Lemma 4.3.2 we show that if $h \in \Pi^\dagger$ does not cross \mathcal{A} then the paths of \mathcal{A}_{\max} that lie strictly to the left of h are a Dedekind cut of \mathcal{A}_{\max} . In Theorem 4.3.9, which will be proved across several lemmas, we show that Dedekind cuts of \mathcal{A}_{\max} are in bijective correspondence with bi-infinite paths that do not cross \mathcal{A} . Thus, each $f \in \Pi^\dagger$ gives rise to a Dedekind cut, which in turn gives rise to a bi-infinite path h , extending f without crossing \mathcal{A} .

Lemma 4.3.2 Let \mathcal{A} be a deterministic weave and let $h \in \Pi^\dagger$ be a path that does not cross \mathcal{A} . Then

$$X_h = \{f \in \mathcal{A}_{\max}; f \triangleleft h, f \not\subseteq h, h \not\subseteq f\} \quad (4.7)$$

is a Dedekind cut of \mathcal{A}_{\max} . If $h, h' \in \Pi^\dagger$ do not cross \mathcal{A} , and are such that $h \not\subseteq h'$ and $h' \not\subseteq h$, then $X_h \neq X_{h'}$.

PROOF: Let us first check that X_h is a Dedekind cut of \mathcal{A}_{\max} . We must check that X_h satisfies conditions (i) and (ii) of Definition 4.3.1. For (i), suppose that $f \in X$ and $g \in \mathcal{A}_{\max}$ with $g \triangleleft f$. We have $f \triangleleft h$. Lemma 3.1.6 (which includes that \triangleleft is transitive) gives that $g \triangleleft h$, as required. For (ii), let $f \in X_h$. Lemma 4.1.5 implies that there exists $g \in \mathcal{A}_{\max}$, such that $f \not\triangleleft g$, $g \not\triangleleft h$, $g \not\triangleleft h, h \not\triangleleft g$, with $f \triangleleft g$ and $g \triangleleft h$. As $f, g, h \in \mathcal{A}_{\max}$ this implies that f, g, h are all distinct. Thus f is not a maximal element of X , as required.

It remains to check that $X_h \neq X_{h'}$ whenever $h, h' \in \Pi^\uparrow$ do not cross \mathcal{A} and satisfy $h \not\triangleleft h'$ and $h' \not\triangleleft h$. Noting that both h and h' do not cross \mathcal{A} , by Lemma 4.1.3 h and h' also do not cross each other. Lemma 3.1.3 gives that $h \triangleleft h'$ or $h' \triangleleft h$. Without loss of generality suppose that $h \triangleleft h'$. Then by Lemma 4.1.5 there exists $g \in \mathcal{A}_{\max}$ such that $g \not\triangleleft h$, $h \not\triangleleft g$, $g \not\triangleleft h'$, $h' \not\triangleleft g$ with $h \triangleleft g$ and $g \triangleleft h'$. It follows that $g \in X_h$ and $g \notin X_{h'}$, as required. \blacksquare

We now consider the inverse map to (4.7), in that we seek to reconstruct a bi-infinite path from its corresponding Dedekind cut. This part is rather technical and will involve the topology introduced in Section 2.1 on $\overline{\mathbb{R}}_5$. which we invite the reader to recall at this point. In particular recall from Lemma 2.1.1 that $t_n \star_n \rightarrow t+$ if and if and only if $t_n \rightarrow t$ and $t_n \star_n \geq t+$ for all sufficiently large n ; similarly $t_n \star_n \rightarrow t-$ if and if and only if $t_n \rightarrow t$ and $t_n \star_n \leq t-$ for all sufficiently large n .

If \mathcal{A} is a weave and X is a Dedekind cut of \mathcal{A}_{\max} then we set

$$\mathcal{L}(X) = \bigcup_{f \in X} L(f) \quad (4.8)$$

where $L(f)$ is defined in (3.1). Note that $\mathcal{L}(X)$ is a subset of $\overline{\mathbb{R}} \times \mathbb{R}_5$, which we equip with the product topology. Recall also that $L_{t\star}(f) = \{x \in \overline{\mathbb{R}}; (x, t\star) \in \mathcal{L}(X)\}$ which, according to (3.1) is either empty or equal to $[-\infty, f(t\star))$. Roughly, our strategy is to show that the right-hand boundary of $\mathcal{L}(X)$ is the graph (in space-time $\overline{\mathbb{R}} \times \mathbb{R}_5$) of a càdlàg path. With this in mind, given a Dedekind cut X of \mathcal{A}_{\max} let $\mathcal{P}_X : \mathbb{R}_5 \rightarrow \overline{\mathbb{R}}$ be given by

$$\mathcal{P}_X(t\star) = \sup \left\{ x \in \overline{\mathbb{R}}; (x, t\star) \in \overline{\mathcal{L}(X)} \right\}. \quad (4.9)$$

Taking the closure of $\mathcal{L}(X)$ in (4.9) is crucial, because being a càdlàg path corresponds to being a continuous function on \mathbb{R}_5 , and being a continuous function corresponds to having a closed graph.

Lemma 4.3.3 *Let \mathcal{A} be a deterministic weave and let X be a Dedekind cut of \mathcal{A}_{\max} . If $(y, t\star) \in \overline{\mathcal{L}(X)}$ and $x \leq y$ then $(x, t\star) \in \overline{\mathcal{L}(X)}$.*

PROOF: Let $(y, t\star) \in \overline{\mathcal{L}(X)}$ and $x < y$. By (4.8) there exists $f_n \in X$ and $(y_n, t_n \star_n) \in \overline{\mathbb{R}} \times \mathbb{R}_5$ such that $(y_n, t_n \star_n) \rightarrow y$ and $y_n \in L_{t_n \star_n}(f_n)$. Thus $x \leq y_n$ for all sufficiently large $n \in \mathbb{N}$. For such n we have $x_n \in L_{t_n \star_n}(f_n)$, which implies $(x_n, t_n \star_n) \in \mathcal{L}(X)$, hence $(x, t\star) \in \overline{\mathcal{L}(X)}$. \blacksquare

Lemma 4.3.4 *Let \mathcal{A} be a deterministic weave and let X be a Dedekind cut of \mathcal{A}_{\max} . The following hold:*

1. *Let $f \in X$ and $g \in \mathcal{A}_{\max} \setminus X$. Then $f \triangleleft g$.*
2. *Suppose that $(x, t\star) \in (\overline{\mathbb{R}} \times \mathbb{R}_5) \setminus \overline{\mathcal{L}(X)}$. Then for all $\epsilon > 0$ there exists $t' \in \mathbb{R}$ and $g \in \mathcal{A}_{\max} \setminus X$ such that $|t - t'| \leq \epsilon$ and $g \in \mathcal{A}((x, t'))$.*

PROOF: For the first part, take $f \in X$ and $g \in \mathcal{A}_{\max} \setminus X$. Note that $f, g \in \mathcal{A}$ so f and g may not cross, which by Lemma 3.1.3 implies that $f \triangleleft g$ or $g \triangleleft f$. If $f \subseteq g$ or $g \subseteq f$ then by maximality $f = g$, in which case $f \triangleleft g$. Alternatively, if both $f \not\subseteq g$ and $g \not\subseteq f$ then $g \triangleleft f$ would imply $g \in X$, because X is a Dedekind cut; so we must have $f \triangleleft g$. Thus, in all cases we have $f \triangleleft g$.

Let us now consider the second claim. First consider the case $t\star = t+$. Let $\epsilon > 0$. As $(x, t+) \notin \overline{\mathcal{L}(X)}$ there exists $\epsilon_0 > 0$ such that

$$([x - \epsilon_0, x + \epsilon_0] \times [t+, (t + \epsilon_0)+]) \cap \overline{\mathcal{L}(X)} = \emptyset. \quad (4.10)$$

Without loss of generality, assume $\epsilon \in (0, \epsilon_0)$. By pervasiveness of \mathcal{A} there exists $g \in \mathcal{A}((x, t'))$ where $t' = t + \epsilon/2$. Without loss of generality we may take $g \in \mathcal{A}_{\max}$. It is clear that $|t - t'| \leq \epsilon$.

Consider if $g \in X$. Note that we have $g(t'+) \vee g(t'-) \geq x$. If $g(t'+) \geq x$ then $[-\infty, x] \subseteq L_g(t'+)$ which would imply $(x, t'+) \in \overline{L(g)} \subseteq \overline{\mathcal{L}(X)}$, contradicting (4.10), so this may not happen. The remaining case is that $g(t'+) < x \leq g(t'-)$, in which case $[-\infty, x] \subseteq L_{t'-}(g)$, implying that $(x, t') \in \overline{L(g)} \subseteq \overline{\mathcal{L}(X)}$, contradicting (4.10), so this may not happen either. We conclude that $g \notin X$. This completes the proof of the case $t\star = t+$. The case $t = t-$ is similar, using in place of (4.10) that for some $\epsilon_0 > 0$ we have $[x - \epsilon_0, x + \epsilon_0] \times [(t - \epsilon_0)-, t-] \cap \overline{\mathcal{L}(X)} = \emptyset$. ■

Lemma 4.3.5 *Let \mathcal{A} be a deterministic weave and let X be a Dedekind cut of \mathcal{A}_{\max} . Then \mathcal{P}_X is a bi-infinite càdlàg path.*

PROOF: Let us write $h = \mathcal{P}_X$ for the duration of this proof. We must show that h is a continuous map from \mathbb{R}_s to $\overline{\mathbb{R}}$. By the closed graph theorem, the function $h : \mathbb{R}_s \rightarrow \overline{\mathbb{R}}$ is continuous if and only if its graph $\mathcal{H} = \{(h(t\star), t\star); t\star \in \mathbb{R}_s\}$ is a closed subset of $\overline{\mathbb{R}} \times \mathbb{R}_s$. Let $t_n\star_n \rightarrow t\star$ in \mathbb{R}_s . By compactness of $\overline{\mathbb{R}}$ the sequence $(h(t_n\star_n), t_n\star_n)$ is relatively compact. Let $(x, t\star)$ be a limit point of this sequence, and (with slight abuse of notation) let us pass to a subsequence such that $h(t_n\star_n) \rightarrow x$. To establish the present lemma we must show that $x = h(t\star)$.

By (4.9), for each $n \in \mathbb{N}$ there exists a sequence $(x_{n,m}, t_{n,m}\star_{n,m})_{m \in \mathbb{N}} \subseteq \mathcal{L}(X)$ such that $(x_{n,m}, t_{n,m}\star_{n,m}) \rightarrow (h(t_n\star_n), t_n\star_n)$ as $m \rightarrow \infty$. By a diagonal argument there exists a strictly increasing function $m : \mathbb{N} \rightarrow \mathbb{N}$ such that $(x_{n,m(n)}, t_{n,m(n)}\star_{n,m(n)}) \rightarrow (x, t\star)$. Hence $(x, t\star) \in \overline{\mathcal{L}(X)}$, which implies that $x \leq h(t\star)$. If $h(t\star) = -\infty$ then we now have $x = h(t\star)$, so in what follows we may assume that $-\infty < h(t\star)$.

We will now argue by contradiction: suppose that $x < h(t\star)$. Let $\epsilon > 0$ be such that $x + 3\epsilon \leq h(t\star)$, and note that $x + \epsilon < h(t\star) - \epsilon$. We consider the cases $t\star = t+$ and $t\star = t-$ in turn.

Suppose, first, that $t\star = t+$. Let us briefly outline the strategy. We will construct a sequence of $f_n \in X$ that come close to the space-time point $(h(t+), t+)$, and a sequence of $g_n \in \mathcal{A}_{\max} \setminus X$ that come close to $(h(t_n\star_n), t_n\star_n) \approx (x, t+)$. Note that $x < h(t+)$, whilst Lemma 4.3.4 gives $f_j \triangleleft g_k$ for all $j, k \in \mathbb{N}$. This combination causes (f_n) and (g_n) to become tangled up in each other, so much so that their limit points $f, g \in \mathcal{A}$ will cross, by jumping over each other in opposite directions at time t , resulting in a contradiction. We now proceed with the proof.

By Lemma 2.1.1, the fact that $t_n\star_n \rightarrow t+$ implies that for sufficiently large n we must have $t_n\star_n \geq t+$. As $h(t_n\star_n) \rightarrow x < h(t+)$, in fact for sufficiently large n we have $t_n\star_n > t+$, and also $h(t_n\star_n) \leq x + \epsilon$. Without loss of generality we pass to a subsequence and assume that both these properties hold for all n .

By (4.9) there exists $(y_n, s_n \bullet_n)_{n \in \mathbb{N}} \subseteq \mathcal{L}(X)$ such that $(y_n, s_n \bullet_n) \rightarrow (h(t+), t+)$. The fact that $s_n \bullet_n \rightarrow t+$ implies that for sufficiently large n we must have $s_n \bullet_n \geq t+$, and for sufficiently large n we also have $y_n \geq h(t\star) - \epsilon$. so without loss of generality we pass to a subsequence and assume that both these properties hold for all n .

Consider if

$$h(t+) = \sup\{x; (x, t+) \in \mathcal{L}(X)\}. \quad (4.11)$$

In this case there exists $f \in X$ such that $L_{t+}(f) = [-\infty, f(t+))$ and $0 < h(t+) - f(t+) \leq \epsilon$, so $f(t+) \geq x + 2\epsilon$. Right continuity (i.e. forwards in time) of f thus implies $h(t_n \star_n) \geq x + \epsilon$ for all sufficiently large n , which contradicts the fact that $h(t_n \star_n) \rightarrow x$. So this case may not occur.

Therefore, $h(t+) \neq \sup\{x; (x, t+) \in \mathcal{L}(X)\}$, which implies that for sufficiently large n we have $s_n \bullet_n \neq t+$ (because $(y_n, s_n \bullet_n) \rightarrow (h(t+), t+)$ and $(y_n, s_n \bullet_n) \in \mathcal{L}(X)$). We have already seen that $s_n \bullet_n \geq t+$, so without loss of generality we pass to a subsequence and assume that $s_n \bullet_n > t+$ for all n .

As $(y_n, s_n \bullet_n) \in \mathcal{L}(X)$ there exists $f_n \in X$ such that $(y_n, s_n \bullet_n) \in L(f_n)$. Hence $f_n(s_n \bullet_n) \geq y_n \geq h(t\star) - \epsilon$. By compactness of \mathcal{A} , without loss of generality we pass to a subsequence and assume that $f_n \rightarrow f \in \mathcal{A}$.

Let $x_n = h(t_n \star_n) + 2^{-n}$. As $h(t_n \star_n) \rightarrow x$, without loss of generality we pass to a subsequence and assume that $x_n \leq x + \epsilon$ for all n . Lemma 4.3.3 gives that $(x_n, t_n \star_n) \notin \overline{\mathcal{L}(X)}$. Thus, by the second part of Lemma 4.3.4, for each $n \in \mathbb{N}$ there exists $t'_n \in \mathbb{R}_\mathfrak{s}$ and $g_n \in \mathcal{A}_{\max} \setminus X$ such that $|t_n - t'_n| \leq 2^{-n}$, $t'_n \star_n > t+$ and $g_n \in \mathcal{A}((x_n, t'_n))$. By compactness of \mathcal{A} , without loss of generality we pass to a subsequence and assume that $g_n \rightarrow g \in \mathcal{A}$. By the first part of Lemma 4.3.4 we have $f_i \triangleleft g_j$ for all $i, j \in \mathbb{N}$.

We have that $t'_n \star_n$ and $s_n \bullet_n$ are both strictly greater than $t+$, and both tend to $t+$ as $n \rightarrow \infty$. Consequently, passing to further subsequences, there exists a strictly increasing function $n \mapsto n'$ such that

$$s_{n+1} \bullet_{n+1} < t'_{n'} \star_{n'} < s_n \bullet_n \quad (4.12)$$

for all n .

We now examine the sequence (g_n) as $n \rightarrow \infty$. We will show that

$$g_{n'}(t'_{n'}-) \vee g_{n'}(t'_{n'}+) \leq x_{n'} \leq x + \epsilon, \quad (4.13)$$

$$g_{n'}(s_n \bullet_n) \geq y_n \geq h(t\star) - \epsilon. \quad (4.14)$$

Equation (4.13) follows because $g_{n'} \in \mathcal{A}((x_{n'}, t'_{n'}))$. To see equation (4.14): we have that $f_n \triangleleft g_{n'}$ and that $[-\infty, y_n] \subseteq L_{s_n \bullet_n}(f_n)$. Lemma 3.1.2 implies $L(f_n) \cap R(g_{n'}) = \emptyset$, and $\sigma_{g_{n'}} < s_n \bullet_n$ so we must have $g_{n'}(s_n \bullet_n) \geq y_n$.

From Lemma A.2.2, combined with (4.12), (4.13) and (4.14) we obtain that the limit g makes a rightwards jump at time t , from below $x + \epsilon$ at time $t-$ to above $h(t\star) - \epsilon$ at time $t+$.

We now turn our attention to (f_n) , in similar style. Here, we show that

$$f_{n+1}(s_{n+1} \bullet_{n+1}) \geq y_{n+1} \geq h(t\star) - \epsilon \quad (4.15)$$

$$f_{n+1}(t'_{n'}-) \wedge f_{n+1}(t'_{n'}+) \leq x_n \leq x + \epsilon \quad (4.16)$$

Equation (4.15) follows from the fact that $(y_n, s_n \bullet_n) \in L(f_n)$. To see equation (4.16): we have that $f_{n+1} \triangleleft g_{n'}$. If $g_{n'}(t'_{n'}+) \leq g_{n'}(t'_{n'}-)$ then we have $f_{n+1}(t'_{n'}+) \leq g_{n'}(t'_{n'}+) \leq x_n$. Alternatively, if $g_{n'}(t'_{n'}-) < g_{n'}(t'_{n'}+)$ then we have $[-\infty, x_n] \subseteq R'_{t'_{n'}-}(g_{n'})$, and Lemma 3.1.2 gives $L(f_{n+1}) \cap R(g_{n'}) = \emptyset$, which implies $f_{n+1}(t'_{n'}-) \leq x_n$. In both cases we have (4.16).

From Lemma A.2.2 combined with (4.12), (4.15) and (4.16) we obtain that the limit f makes a leftwards jump at time t , from above $h(t\star) - \epsilon$ at time $t-$ to below $x + \epsilon$ at time $t+$. Thus

f and g cross (by jumping in opposite directions over each other at time t). As both $f, g \in \mathcal{A}$, this is a contradiction. This completes the proof of the case $t_\star = t+$.

It remains to consider the case $t_\star = t-$. The argument is essentially the same, except that Lemma 2.1.1 requires that we now approach $t-$ from the left (i.e. from backwards in time) rather than $t+$ from the right. In outline: construct a sequence of $f_n \in X$ that come close to the space-time point $(h(t-), t-)$, and a sequence of $g_n \in \mathcal{A}_{\max}$ that come close to $(h(t_n \star_n), t_n \star_n) \approx (x, t-)$. Note that $x < h(t-)$, whilst Lemma 4.3.4 gives $f_j \triangleleft g_k$ for all $j, k \in \mathbb{N}$. This combination causes (f_n) and (g_n) to become entangled with each other, so that once again their limit points $f, g \in \mathcal{A}$ will cross – resulting in a contradiction.

There is one point at which a difference worthy of comment emerges. This concerns (4.11). If $h(t-) = \sup\{x; (x, t-) \in \mathcal{L}(X)\}$ then there exists $f \in X$ such that $L_{t-}(f) = [-\infty, f(t-))$ and $0 < h(t-) - f(t-) \leq \epsilon$.

- If $\sigma_f < t$ then a similar argument to that in the same paragraph as (4.11) applies, using left continuity of f instead of right continuity; this reaches a contradiction.
- If $\sigma_f = t$ then we require a new step within the argument, one that features only here because of the ‘extra’ behaviour of $L_{t_\star}(f)$ when $t_\star = \sigma_f-$, see (3.1). In particular, for $t = \sigma_f$ the fact that $L_{t-}(f) = [-\infty, f(t-))$ implies that $f(t+) < f(t-)$. We then proceed as before to construct $g \in \mathcal{A}$ such that $g(t-) \leq x + \epsilon$ and $g(t+) \geq h(t-) - \epsilon$. Thus f and g cross, reaching a contradiction.

If $h(t-) \neq \{x; (x, t-) \in \mathcal{L}(X)\}$ then we can (and moreover can only) approximate $(h(t-), t-) \in \mathcal{L}(X)$ using space-time points in $\mathcal{L}(X)$ with times strictly less than $t-$. In this case we may proceed as before. This completes the proof. \blacksquare

Lemma 4.3.6 *Let \mathcal{A} be a deterministic weave and let X be a Dedekind cut of \mathcal{A}_{\max} . Then \mathcal{P}_X does not cross \mathcal{A} . Moreover, if $f \in X$ then $f \triangleleft \mathcal{P}_X$, and if $g \in \mathcal{A}_{\max} \setminus X$ then $\mathcal{P}_X \triangleleft g$.*

PROOF: Let us write $h = \mathcal{P}_X$ for the duration of this proof. From Lemma 4.3.5 we have $h \in \Pi^\uparrow$. Note that h crosses \mathcal{A} if and only if h crosses \mathcal{A}_{\max} . We will show, in turn, that (a) $f \in X \Rightarrow f \triangleleft h$ and (b) $g \in \mathcal{A}_{\max} \setminus X \Rightarrow h \triangleleft g$. With this in hand it follows from Lemma 3.1.3 that h does not cross \mathcal{A}_{\max} , thus h does not cross \mathcal{A} .

We begin with (a). Let $f \in X$ and $t_\star \geq \sigma_f-$. If $L_{t_\star}(f) = [-\infty, f(t_\star))$ then $(f(t_\star), t_\star) \in \mathcal{L}(X)$ and hence $f(t_\star) \leq h(t_\star)$, which implies $L_{t_\star}(f) \cap R_{t_\star}(h) = \emptyset$. Alternatively, if $L_{t_\star}(f) = \emptyset$ then it is immediate that $L_{t_\star}(f) \cap R_{t_\star}(h) = \emptyset$. Hence $L(f) \cap R(g) = \emptyset$, which by Lemma 3.1.4 implies that $f \triangleleft h$.

We now move on to (b). Let $g \in \mathcal{A}_{\max} \setminus X$. We will argue by contradiction. Suppose that $h \not\triangleleft g$. Then by Lemma 3.1.2 we have $L(h) \cap R(g) \neq \emptyset$. In particular, for some $t_\star \geq \sigma_g-$ we have $L_{t_\star}(h) = [-\infty, h(t_\star))$ and $R_{t_\star}(g) = (g(t_\star), \infty]$ with $g(t_\star) < h(t_\star)$.

Consider first if $t_\star \geq \sigma_g+$. Then, using the càdlàg property of g and h there exists $\epsilon > 0$ an interval $[a+, b-]$ with $a < b$ such that

$$g(s_\bullet) + \epsilon \leq h(s_\bullet) \tag{4.17}$$

for all $s_\bullet \in [a-, b+]$ (to see this: if $\star = +$ take $a = t$, if $\star = -$ take $b = t$). Let us briefly note our strategy here: we will use (4.17) to show that g lies to the left of some path in X . Fix some $s_\bullet \in \mathbb{R}_\epsilon$ with $a+ < s_\bullet < b-$. By (4.9) there exists $(y_n, s_{n \bullet_n}) \in \mathcal{L}(X)$ such that $(y_n, s_{n \bullet_n}) \rightarrow (h(s_\bullet), s_\bullet)$ as $n \rightarrow \infty$. For all sufficiently large $n \in \mathbb{N}$ we have

$$h(s_\bullet) - \epsilon/2 \leq y_n \quad \text{and} \quad a- \leq s_{n \bullet_n} \leq b+ . \tag{4.18}$$

By Lemma 4.3.5 we have $h(s_n \bullet_n) \rightarrow h(s \bullet)$, so for all sufficiently large n we also have

$$|h(s \bullet) - h(s_n \bullet_n)| \leq \epsilon/2. \quad (4.19)$$

Fix $n \in \mathbb{N}$ large enough that (4.18) and (4.19) both hold. By (4.8) there exists $f_n \in X$ such that $(y_n, s_n \bullet_n) \in L_{s_n \bullet_n}(f)$, which means that $y_n < f(s_n \bullet_n)$. Combining this inequality with (4.17), (4.18) and (4.19) we obtain that

$$g(s_n \bullet_n) \leq h(s_n \bullet_n) - \epsilon \leq h(s \bullet) - \epsilon/2 \leq y_n < f_n(s_n \bullet_n).$$

Hence $g(s_n \bullet_n) \in L_{s_n \bullet_n}(f)$, which by Lemma 3.1.4 means that $g \triangleleft f_n$. As X is a Dedekind cut and $f_n \in X$, we thus have $g \in X$, which is a contradiction.

It remains to consider the case $t \star = \sigma_g -$. In this case by (3.1) we have $R_{\sigma-}(g) = (g(\sigma-), \infty]$ and $g(\sigma-) < g(\sigma+)$. We have also that $g(\sigma-) < h(\sigma-)$. By (4.9) there exists $(y_n, s_n \bullet_n) \in \mathcal{L}(X)$ such that $(y_n, s_n \bullet_n) \rightarrow (h(\sigma-), \sigma-)$, and $f_n \in X$ such that $y_n \in L_{s_n \bullet_n}(f) = [-\infty, f_n(s_n \bullet_n))$. By compactness of \mathcal{A} we may pass to a subsequence and assume that $f_n \rightarrow f \in \mathcal{A}$. Without loss of generality we may assume that $s_n \bullet_n \leq \sigma-$, so $\sigma_{f_n} \leq \sigma_g$.

Suppose that $f_n(\sigma-) \geq g(\sigma-)$: then $f_n(\sigma-) \in R_{\sigma-}(g)$ which, by Lemma 3.1.4 would give $g \triangleleft f_n$, and as X is a Dedekind cut this would give $g \in X$, which is false. Hence in fact $f_n(\sigma-) < g(\sigma-)$. We now have $s_n \bullet_n \leq \sigma-$ with $\sigma_{f_n} \rightarrow \sigma-$, along with $f_n(s_n \bullet_n) \geq y_n$ and $f_n(\sigma-) \leq g(\sigma-)$. By Lemma A.2.3, this implies that f jumps leftwards at σ , from right of $h(\sigma-)$ to left of $g(\sigma-)$. This implies that g and f cross, by jumping over each other in opposite directions as σ , which is a contradiction as both $f, g \in \mathcal{A}$. This completes the proof. \blacksquare

Lemma 4.3.7 *Let \mathcal{A} be a deterministic weave and suppose that $h \in \Pi^\dagger$ does not cross \mathcal{A} . Define $X = X_h$ according to (4.7). Then $\mathcal{P}_X = h$.*

PROOF: Let us write $h' = \mathcal{P}_X$ for the duration of this proof. We must show that $h = h'$. We will argue by contradiction. Noting that h is bi-infinite, if $h \neq h'$ then there exist $t+ \in \mathbb{R}_s$ such that $h(t+) \neq h'(t+)$. Without loss of generality (or consider space reflected about the origin) we may assume that $h(t+) < h'(t+)$. Thus $h \triangleleft h'$. By Lemma 4.1.2, again using that h and h' are bi-infinite, there exists $g \in \mathcal{A}$ such that $h \triangleleft g$ and $g \triangleleft h'$ with $g \not\subseteq h$ and $g \not\subseteq h'$. Without loss of generality we may take $g \in \mathcal{A}_{\max}$. Hence either $g \in X$ or $g \in \mathcal{A}_{\max} \setminus X$. We consider these two cases separately.

If $g \in X$ then, recalling that $X = \{f \in \mathcal{A}_{\max}; f \triangleleft h \text{ and } f \not\subseteq h\}$, we have $g \triangleleft h$. From Lemma 3.1.3, noting that h is bi-infinite, we thus obtain $g \subseteq h$, which is a contradiction. If $g \in \mathcal{A}_{\max} \setminus X$ then by Lemma 4.3.6 we have $h' \triangleleft g$. From Lemma 3.1.3, noting that h' is bi-infinite, we thus obtain $g \subseteq h'$, which is a contradiction. Having reached a contradiction in both cases, we conclude that in fact $h = h'$. \blacksquare

Lemma 4.3.8 *Let \mathcal{A} be a deterministic weave and let X be a Dedekind cut of \mathcal{A}_{\max} . There exists $h \in \Pi^\dagger$ such that $X = X_h$, where X_h is given by (4.7).*

PROOF: Let us write $h = \mathcal{P}_X$ for the duration of this proof. By Lemma 4.3.5 we have $h \in \Pi^\dagger$. Define X_h as in (4.7), that is $X_h = \{f \in \mathcal{A}_{\max}; f \leq h \text{ and } f \not\subseteq h\}$. Note that since $h \in \Pi^\dagger$ we may discard the condition $h \not\subseteq f$, because if $h \subseteq f$ then $f = h$ and thus also $f \subseteq h$. We must show that $X = X_h$.

Let $f \in X$. Lemma 4.3.6 gives that $f \triangleleft h$. Moreover Lemma 4.3.6 gives that $g \triangleleft h$ for all $g \in X$. Thus, if $f \subseteq h$ then we would also have $g \triangleleft f$ for all $g \in X$, which would make f a

maximal element of X ; this a contradiction as X is a Dedekind cut. Hence $f \not\subseteq h$. We thus have $X \subseteq X_h$. We now move on to the reverse inclusion.

Let $f \in X_h$, so we have $f \in \mathcal{A}_{\max}$, $f \triangleleft h$ and $f \not\subseteq h$. If $f \notin X$ then Lemma 4.3.6 gives that $h \triangleleft f$, from which part 2 of Lemma 3.1.5 implies that $f \subseteq h$ or $h \subseteq f$; this is a contradiction. Hence $X_h \subseteq X$, so $X = X_h$ as required. \blacksquare

Theorem 4.3.9 *Let \mathcal{A} be a deterministic weave. Let \mathcal{X} denote the set of Dedekind cuts of \mathcal{A}_{\max} .*

1. *The map $h \mapsto X_h$ given by (4.7) is a bijection between $\text{flow}(\mathcal{A})$ and \mathcal{X} . The inverse map $X \mapsto \mathcal{P}_X$ is given by (4.9).*
2. *For any $f \in \Pi^\uparrow$ that does not cross \mathcal{A} , there exists $h \in \text{flow}(\mathcal{A})$ such that $f \subseteq h$.*

PROOF: Recall from (2.13) that by definition $\text{flow}(\mathcal{A}) = \{h \in \Pi^\uparrow; h \text{ does not cross } \mathcal{A}\}$. Lemma 4.3.2 gives that the range of the map $h \mapsto X_h$ (with domain $\text{flow}(\mathcal{A})$) is within \mathcal{X} and that this map is injective. Lemma 4.3.8 gives that $h \mapsto X_h$ has range \mathcal{X} . Lemmas 4.3.5 and 4.3.6 ensure that the range of the map $X \mapsto \mathcal{P}_X$ (with domain \mathcal{X}) is within $\text{flow}(\mathcal{A})$, so Lemma 4.3.7 gives that $h \mapsto X_h$ and $X \mapsto \mathcal{P}_X$ are inverses of each other, between $\text{flow}(\mathcal{A})$ and \mathcal{X} . This establishes the first claim of the present theorem.

To see the second claim, let $f \in \Pi^\uparrow$ and suppose that f does not cross \mathcal{A} . By Lemma 4.3.2 we have $X_f \in \mathcal{X}$, from which part 1 of the present theorem gives that $h = \mathcal{P}_{X_f} \in \text{flow}(\mathcal{A})$ does not cross \mathcal{A} . It remains to show that $f \subseteq h$. We will argue by contradiction.

Suppose that $f \not\subseteq h$, which as $h \in \Pi^\uparrow$ implies that $h \not\subseteq f$. We have that $\mathcal{A} \cup \{f\}$ is non-crossing and that $\mathcal{A} \cup \{h\}$ is non-crossing. It follows by Lemma 4.1.3 that $\mathcal{A} \cup \{f, h\}$ is non-crossing, so in particular f does not cross h . By Lemma 3.1.3 we have $f \triangleleft h$ or $h \triangleleft f$. We treat these two cases in turn.

Consider, first, if $f \triangleleft h$. Then Lemma 4.1.5 gives $g \in \mathcal{A}_{\max}$ such that $f \triangleleft g$, $g \triangleleft h$, with f, g, h all incomparable under \subseteq . Hence $g \in \mathcal{A}_{\max} \setminus X_f$, which by Lemma 4.3.6 gives $h \triangleleft g$. By Lemma 3.1.6 we thus have $g \subseteq h$ or $h \subseteq g$, which is a contradiction.

The argument when $h \triangleleft f$ is similar. Now Lemma 4.1.5 gives $g \in \mathcal{A}_{\max}$ such that $h \triangleleft g$, $g \triangleleft f$ with f, g, h all incomparable under \subseteq . Hence $g \in X_f$, which by Lemma 4.3.6 gives $g \triangleleft h$, and Lemma 3.1.6 again arrives at a contradiction. This completes the proof. \blacksquare

4.4 The flow map

In this section we establish further properties of the flow operation defined in (2.13), including Lemma 4.4.3 which shows that $\mathcal{A} \mapsto \text{flow}(\mathcal{A})$ is continuous on \mathcal{W}_{det} . Another key result, contained within Lemma 4.4.1, is that the path extension of Theorem 4.3.9 preserves compactness. Here, for the first time, we see interaction between all the main concepts of the present article: relative compactness in $\mathcal{K}(\Pi^\uparrow)$, path extension, and the pervasiveness and non-crossing properties of weaves.

Lemma 4.4.1 *Let $\mathcal{A} \subseteq \mathcal{K}(\Pi^\uparrow)$ be a relatively compact subset of $\mathcal{K}(\Pi^\uparrow)$, where each $\mathcal{A} \in \mathcal{A}$ is a deterministic weave. Suppose that any limit point of \mathcal{A} is non-crossing. Then $\{\text{flow}(\mathcal{A}); \mathcal{A} \in \mathcal{A}\}$ is a relatively compact subset of $\mathcal{K}(\Pi^\uparrow)$.*

PROOF: As Π^\uparrow is a closed subset of Π^\uparrow , also $\mathcal{K}(\Pi^\uparrow)$ is a closed subset of $\mathcal{K}(\Pi^\uparrow)$. It therefore suffices to consider relative compactness in $\mathcal{K}(\Pi^\uparrow)$. By Lemma A.1.1 the set $\{\text{flow}(\mathcal{A}); \mathcal{A} \in \mathcal{A}\}$

is a relatively compact subset of $\mathcal{K}(\Pi^\uparrow)$ if and only if $\mathcal{F} = \bigcup_{\mathcal{A} \in \mathcal{A}} \text{flow}(\mathcal{A})$ is a relatively compact subset of Π^\uparrow . We will argue by contradiction.

Suppose that \mathcal{F} is not relatively compact. Then, by Proposition A.2.1 there exists $T, \kappa > 0$ and sequences $(\delta_n) \subseteq (0, 1)$, $(f_n) \subseteq \mathcal{F}$ such that $\delta_n \searrow 0$ and $w_{T, \delta_n}(f_n) \geq \kappa$. For convenience, recall from (A.2) that

$$w_{T, \delta}(f) = \sup \left\{ d_{\mathbb{R}}(f(t_2 \star_2), [f(t_1 \star_1), f(t_3 \star_3)]) ; t_1 \star_1, t_2 \star_2, t_3 \star_3 \in I_s(f), \right. \\ \left. -T < t_1 < t_2 < t_3 < T, t_3 - t_1 < \delta \right\}.$$

So, we have $t_i^n \star_i^n \in I_s(f_n)$ such that $-T < t_1^n < t_2^n < t_3^n < T$ and $t_3^n - t_1^n < \delta$, with

$$d_{\mathbb{R}}(f_n(t_2^n \star_2^n), [f_n(t_1^n \star_1^n), f_n(t_3^n \star_3^n)]) \geq \kappa. \quad (4.20)$$

By the càdlàg property of the bi-infinite path f_n , we may assume without loss of generality (reducing $\kappa > 0$ if necessary) that f_n is continuous at $t_i^n \star_i^n$ for $i = 1, 2, 3$. Using that \mathbb{R}_c^2 is compact, without loss of generality we may pass to a subsequence and assume additionally that $f_n(t_i^n) \rightarrow y_i \in \mathbb{R}$ in \mathbb{R}_c^2 . Using that $[-T, T]$ is compact, we may pass to a further subsequence and assume additionally that $t_i^n \rightarrow t_i$, with $t_i \in [-T, T]$. Since $0 \leq t_3^n - t_1^n < \delta_n$ in fact $t_1 = t_2 = t_3 = t$. To summarise, we thus have

$$t_1^n < t_2^n < t_3^n \text{ for all } n \text{ and } t_i^n \rightarrow t \in [-T, T] \text{ as } n \rightarrow \infty. \quad (4.21)$$

Without loss of generality (or consider the same setup with space reflected about the origin) we pass to a further subsequence and assume additionally that $f_n(t_1^n) \leq f_n(t_2^n)$, which by (4.20) implies that

$$f_n(t_i^n) + \kappa \leq f_n(t_2^n) \quad \text{for } i = 1, 3. \quad (4.22)$$

Note that $f_n \in \mathcal{A}$, for some $\mathcal{A} \in \mathcal{A}$, and let us write \mathcal{A}_n for such \mathcal{A} . By pervasiveness of \mathcal{A}_n there exist

$$g_n \in \mathcal{A}_n \left((t_1^n, f_n(t_1^n) + \frac{\kappa}{3}) \right) \quad \text{and} \quad h_n \in \mathcal{A}_n \left((t_2^n, f_n(t_2^n) - \frac{\kappa}{3}) \right). \quad (4.23)$$

Since \mathcal{A}_n is a weave, g_n and h_n do not cross each other. By Lemma A.1.1, relative compactness of \mathcal{A} gives that $\mathcal{B} = \bigcup_{\mathcal{A} \in \mathcal{A}} \mathcal{A}$ is a relatively compact subset of Π^\uparrow . Hence we may pass to a further subsequence and assume additionally that $g_n \rightarrow g \in \overline{\mathcal{B}}$ and $h_n \rightarrow h \in \overline{\mathcal{B}}$. The sequence (\mathcal{A}_n) is a subset of \mathcal{A} and therefore is relatively compact, so we may pass to a subsequence and assume that $\mathcal{A}_n \rightarrow \mathcal{A}$. Hence $g, h \in \mathcal{A}$. The set \mathcal{A} is a limit point of \mathcal{A} , therefore (as a hypothesis of the present lemma) \mathcal{A} is non-crossing. Hence g and h may not cross.

Let us briefly comment on the strategy for the remaining part of the proof: we will establish a contradiction through showing that g and h cross at time t , at which time they will jump past each other in opposite directions. By Lemma 3.1.3 we have that g_n, h_n are comparable under \triangleleft to f_n . Since f_n is continuous at t_i^n , by (4.23) and Lemma 3.1.2 we have $f_n \triangleleft g_n$ and $h_n \triangleleft f_n$. Hence $f_n(t_2^n) \leq g_n(t_2^n \pm)$ and $h_n(t_3^n \pm) \leq f_n(t_3^n)$. By definition of g_n, h_n we have $g_n(t_1^n -) \wedge g_n(t_1^n +) \leq f_n(t_1^n) + \frac{\kappa}{3}$ and $f_n(t_2^n) - \frac{\kappa}{3} \leq h_n(t_2^n -) \vee h_n(t_2^n +)$. Combining these facts with (4.22) we thus have $t_i^n \bullet_i^n$ such that

$$g_n(t_1^n \bullet_1^n) + \frac{2\kappa}{3} \leq f_n(t_1^n) + \kappa \leq f_n(t_2^n) \leq g_n(t_2^n \bullet_2^n) \\ h_n(t_3^n \bullet_3^n) + \kappa \leq f_n(t_3^n) + \kappa \leq f_n(t_2^n) \leq h_n(t_2^n \bullet_2^n) + \frac{\kappa}{3}.$$

Hence,

$$g_n(t_1^n \bullet_1^n) + \frac{\kappa}{6} \leq f_n(t_2^n) - \frac{\kappa}{2} \leq g_n(t_2^n \bullet_2^n) - \frac{\kappa}{2}$$

$$h_n(t_3^{\bullet 3}) + \frac{\kappa}{2} \leq f_n(t_2^n) - \frac{\kappa}{2} \leq h_n(t_2^{\bullet 2}) - \frac{\kappa}{6}.$$

Using that $[-T-, T+] \subseteq \mathbb{R}_5$ is compact (which follows from Lemma 2.1.1) we may pass to a subsequence and assume that for $i = 1, 2, 3$ the sequence $(t_i^{\bullet i})$ converges as $n \rightarrow \infty$. We will now send $n \rightarrow \infty$. Recalling that $f_n(t_2^n) \rightarrow y_2$, we thus obtain from Lemma A.2.2 and (4.21) that

$$\begin{aligned} g(t-) &< y_2 - \frac{\kappa}{2} < g(t+) \\ h(t+) &< y_2 - \frac{\kappa}{2} < h(t-), \end{aligned}$$

which implies that g and h cross. This is a contradiction since $g, h \in \mathcal{A}$. \blacksquare

The M1 topology and particular form of $w_{T,\delta}$ is crucial to the above argument. We do not know of an analogous result for the J1 topology.

Lemma 4.4.2 *Let \mathcal{A} be a deterministic weave. Then $\text{flow}(\mathcal{A})$ is a deterministic weave and $\text{flow}(\mathcal{A}) \subseteq \Pi^\uparrow$.*

PROOF: Note that equation (2.13) gives that $\text{flow}(\mathcal{A}) \subseteq \Pi^\uparrow$. Let us first establish that $\text{flow}(\mathcal{A})$ is compact. By Lemma 4.4.1 (taking $\mathcal{A} = \{\mathcal{A}\}$) we have that $\text{flow}(\mathcal{A})$ is relatively compact. We will show that $\text{flow}(\mathcal{A})$ is closed (and thus compact). Suppose that $g_n \rightarrow g$ where $g_n \in \text{flow}(\mathcal{A})$. For any $f \in \mathcal{A}$ there exists $f' \in \text{flow}(\mathcal{A})$ such that $f \subseteq f'$. Thus f' and g_n do not cross, for all n . Lemma 3.3.6 gives that f' and g do not cross, which implies that f and g do not cross. Thus f does not cross \mathcal{A} , which implies that $f \in \text{flow}(\mathcal{A})$.

To check that $\text{flow}(\mathcal{A})$ is a weave, it remains to show that $\text{flow}(\mathcal{A})$ is pervasive and non-crossing. If $g, h \in \text{flow}(\mathcal{A})$ then by (2.13) we have that $\mathcal{A} \cup \{g\}$ is non-crossing and $\mathcal{A} \cup \{h\}$ is non-crossing. By Lemma 4.1.3 we thus have that $\{g, h\}$ is non-crossing, so in fact $\text{flow}(\mathcal{A})$ is non-crossing. If $z \in \mathbb{R}_c^2$ then pervasiveness of \mathcal{A} implies that there exists some $f \in \mathcal{A}$ such that $f \in H(z)$. Theorem 4.3.9 implies that there exists $f' \in \text{flow}(\mathcal{A})$ with $f \subseteq f'$, hence $z \in H(f')$. Thus $\text{flow}(\mathcal{A})$ is pervasive. \blacksquare

Lemma 4.4.3 *Let $\mathcal{A}_n, \mathcal{A}$ be deterministic weaves with $\mathcal{A}_n \rightarrow \mathcal{A}$. Then $\text{flow}(\mathcal{A}_n) \rightarrow \text{flow}(\mathcal{A})$.*

PROOF: The set $\mathcal{A} = \{\mathcal{A}_n; n \in \mathbb{N}\}$ is a relatively compact subset of $\mathcal{K}(\Pi^\uparrow)$. The only limit point of \mathcal{A} is \mathcal{A} , which is non-crossing. Thus by Lemma 4.4.1 the set $\{\text{flow}(\mathcal{A}_n); n \in \mathbb{N}\}$ is a relatively compact subset of $\mathcal{K}(\Pi^\uparrow)$. Suppose that we have $\mathcal{A}_n \rightarrow \mathcal{F}$ along some subsequence and pass to this subsequence, with mild abuse of notation. To prove the present lemma we must show that $\mathcal{F} = \text{flow}(\mathcal{A})$.

By Lemma 4.4.3 we have that $\text{flow}(\mathcal{A})$ is a deterministic weave. Let us show that \mathcal{F} is also a deterministic weave, with $\mathcal{F} \subseteq \Pi^\uparrow$. From the previous paragraph we have $\mathcal{F} \in \mathcal{K}(\Pi^\uparrow)$. As $\mathcal{A}_n \rightarrow \mathcal{F}$ and each \mathcal{A}_n is a subset of the closed set Π^\uparrow , we have $\mathcal{F} \subseteq \Pi^\uparrow$. If $f, g \in \mathcal{F}$ then we have $f_n, g_n \in \text{flow}(\mathcal{A}_n)$ such that $f_n \rightarrow f$ and $g_n \rightarrow g$. We have that f_n and g_n do not cross, so by Lemma 3.3.6 f and g do not cross. Thus \mathcal{F} is non-crossing. For all $z \in \mathbb{R}_c^2$ there exists $f_n \in \mathcal{A}_n$ such that $z \in H(f_n)$. By Lemma A.1.1, relative compactness of (\mathcal{A}_n) implies relative compactness of (f_n) , thus we may pass to a subsequence and assume $f_n \rightarrow f \in \mathcal{F}$. Lemma A.2.2 gives that $z \in H(f)$. Thus \mathcal{F} is pervasive. We have now shown that $\mathcal{F} \subseteq \Pi^\uparrow$ is a deterministic weave.

Our next goal is to show that \mathcal{F} and $\text{flow}(\mathcal{A})$ are non-crossing. Let $f \in \mathcal{F}$ and $g \in \mathcal{A}$. Then there exists $f_n \in \text{flow}(\mathcal{A}_n)$ and $g_n \in \mathcal{A}_n$ such that $f_n \rightarrow f$ and $g_n \rightarrow g$. By Theorem 4.3.9, there exists $g'_n \in \text{flow}(\mathcal{A}_n)$ such that $g_n \subseteq g'_n$. We have $\text{flow}(\mathcal{A}_n) \rightarrow \mathcal{F}$, so by Lemma A.1.1 the set

$\mathcal{F} \cup (\bigcup_{n \in \mathbb{N}} \text{flow}(\mathcal{A}_n))$ is compact, which implies that $\{g'_n; n \in \mathbb{N}\}$ is relatively compact. Hence there exists $g' \in \mathcal{F}$ such that $g'_n \rightarrow g'$. By Lemma A.2.2 we have $g \subseteq g'$. Both f_n and g'_n are elements of \mathcal{A}_n , hence they do not cross each other. By Lemma 3.3.6, f and g' do not cross each other. As $g \subseteq g'$ this means that f and g do not cross.

We now have that \mathcal{F} and \mathcal{A} do not cross, and that both are deterministic weaves consisting entirely of bi-infinite paths. Lemma 4.2.6 gives that $\mathcal{F} = \mathcal{A}$. \blacksquare

Lemma 4.4.4 *Let \mathcal{A}, \mathcal{B} be deterministic weaves and assume that $\mathcal{A} \cup \mathcal{B}$ is non-crossing. Then $\mathcal{A} \cup \mathcal{B}$ is a deterministic weave and $\text{flow}(\mathcal{A}) = \text{flow}(\mathcal{B})$.*

PROOF: It is trivial to check that $\mathcal{A} \cup \mathcal{B}$ is a deterministic weave. Lemma 4.4.2 gives that $\text{flow}(\mathcal{A})$ and $\text{flow}(\mathcal{B})$ are deterministic weaves, composed entirely of bi-infinite paths. Lemma 4.1.4 gives that a path $f \in \Pi^\uparrow$ crosses \mathcal{A} if and only if it crosses \mathcal{B} , so by Lemma 4.2.6 we have $\text{flow}(\mathcal{A}) = \text{flow}(\mathcal{B})$. \blacksquare

Lemma 4.4.5 *Let \mathcal{A} be a deterministic weave. Then $\text{flow}(\mathcal{A})$ is a maximal element of $(\mathcal{W}_{\text{det}}, \preceq)$ and $\mathcal{A} \preceq \text{flow}(\mathcal{A})$.*

PROOF: The reader may wish to check (2.7) for the definition of the partial order \preceq . Note that Lemma 4.4.2 gives that $\text{flow}(\mathcal{A}) \in \mathcal{W}_{\text{det}}$. Let us first show that $\mathcal{A} \preceq \text{flow}(\mathcal{A})$. Theorem 4.3.9 gives that $\mathcal{A} \subseteq (\text{flow}(\mathcal{A}))_\uparrow$. Now consider $f \in \mathcal{A}_\uparrow \cap \text{flow}(\mathcal{A})$. For such f we have $f \in \Pi^\uparrow$, which implies $f \in \mathcal{A}$. Thus $\mathcal{A} \preceq \text{flow}(\mathcal{A})$.

It remains to show that $\text{flow}(\mathcal{A})$ is a maximal element of \mathcal{W}_{det} . Lemma 4.4.2 gives that $\text{flow}(\mathcal{A}) \in \mathcal{W}_{\text{det}}$. Suppose that $\mathcal{B} \in \mathcal{W}_{\text{det}}$, comparable under \preceq to $\text{flow}(\mathcal{A})$. We must show that $\mathcal{B} \preceq \text{flow}(\mathcal{A})$. From Lemma 3.1.7 we have that $\text{flow}(\mathcal{A}) \cup \mathcal{B}$ is non-crossing. From what we have already proved we have that $\mathcal{A} \preceq \text{flow}(\mathcal{A})$, so using Lemma 3.1.7 again gives that $\mathcal{A} \cup \text{flow}(\mathcal{A})$ is non-crossing. Thus by Lemma 4.1.3 we have that $\mathcal{A} \cup \mathcal{B}$ is non-crossing. Lemma 4.4.4 gives that $\text{flow}(\mathcal{A}) = \text{flow}(\mathcal{B})$. From what we have already proved we now have that $\mathcal{B} \preceq \text{flow}(\mathcal{B}) = \text{flow}(\mathcal{A})$, as required. This completes the proof. \blacksquare

Lemma 4.4.6 *Let \mathcal{A} be a deterministic weave. The set of ramification points of \mathcal{A} has Lebesgue measure zero.*

PROOF: By Lemma 4.4.2 we have that $\text{flow}(\mathcal{A}) = \{f \in \Pi^\uparrow; f \text{ does not cross } \mathcal{A}\}$ is a weave. Lemma 4.2.5 gives that the set of ramification points of $\text{flow}(\mathcal{A})$ has measure zero. By Theorem 4.3.9 any ramification point of \mathcal{A} is also a ramification point of \mathcal{F} . The result follows. \blacksquare

In view of Lemma 4.4.6, for any deterministic weave \mathcal{A} there exists a dense countable subset $D \subseteq \mathbb{R}_c^2$ that is non-ramified. In Lemma 5.4.1 we will address, for random waves, how such a subset may be chosen to be a random variable.

4.5 The web map

In this section we study properties of the web operation defined in (2.12). Some results in this section are analogues of properties that were proven (for the flow operation) in Section 4.4. We also address the dependence, or rather the lack thereof, of (2.12) on the set D . The web operation involves both the ‘downset’ operation $A_\uparrow = \{f \in \Pi^\uparrow; f \subseteq g \text{ for some } g \in A\}$ and taking closure in Π^\uparrow , so we begin with the interaction between these two operations.

Lemma 4.5.1 *If $A \subseteq \Pi^\uparrow$ is relatively compact then A_\uparrow is relatively compact and $(\overline{A})_\uparrow = \overline{(A_\uparrow)}$.*

PROOF: The first claim follows immediately from Proposition A.2.1. It remains to establish that $(\overline{A})_\uparrow = \overline{(A)_\uparrow}$. To this end, suppose that $f \in (\overline{A})_\uparrow$. Then there exists $(g_n) \subseteq A$ such that $g_n \rightarrow g \in \Pi^\uparrow$ and $f \subseteq g$. Let $z \in \mathbb{R}_c^2$ denote the initial point of f . By Lemma A.2.2 $z \in \overline{\cup_n H(g_n)}$. Hence, we may pass to a subsequence of the (g_n) and choose $z_n \in H(g_n)$ such that $z_n \rightarrow z$. It follows from Lemma A.2.2 that $g_n|_{z_n} \rightarrow g|_z = f$, so $f \in (\overline{A})_\uparrow$. Thus $(\overline{A})_\uparrow \subseteq \overline{(A)_\uparrow}$.

In preparation for proving the reverse inclusion, let us first show that if $B \subseteq \Pi$ is compact then B_\uparrow is closed. Take such an B , and let $(f_n) \subseteq \mathcal{B}_\uparrow$ with $f_n \rightarrow f$. We have $(g_n) \subseteq B$ such that $f_n \subseteq g_n$. By compactness, and passing to a subsequence of (g_n) , we have that $g_n \rightarrow g \in B$. By Lemma 3.3.3 we have $f \subseteq g$, thus $f \in B_\uparrow$, which establishes that B_\uparrow is closed.

We now show the reverse inclusion. Suppose that $f \in \overline{(A)_\uparrow}$. Then there exists $(g_n) \subseteq A$ and $f_n \subseteq g_n$ such that $f_n \rightarrow f$. In particular, $g_n \in \overline{A}$ which implies $f_n \in (\overline{A})_\uparrow$. From the previous paragraph we have that $(\overline{A})_\uparrow$ is closed, so $f \in (\overline{A})_\uparrow$. Thus $(\overline{A})_\uparrow \supseteq \overline{(A)_\uparrow}$, as required. \blacksquare

Lemma 4.5.2 *Suppose that \mathcal{A} is a deterministic weave. Then \mathcal{A}_\uparrow is a deterministic weave and $\mathcal{A}_\uparrow \preceq \mathcal{A}$.*

PROOF: Relative compactness of \mathcal{A}_\uparrow is given by Lemma 4.5.1. Since \mathcal{A} is a weave, \mathcal{A} is closed, so from Lemma 4.5.1 we have that \mathcal{A}_\uparrow is closed, which thus implies compactness. The non-crossing property of \mathcal{A}_\uparrow is inherited from \mathcal{A} by Definition 2.3.1. Using that $\mathcal{A}_\uparrow = (\mathcal{A}_\uparrow)_\uparrow$ it is straightforward to check that (2.7) gives $\mathcal{A}_\uparrow \preceq \mathcal{A}$. \blacksquare

Lemma 4.5.3 *Let \mathcal{A} be a deterministic weave and D, D' be dense non-ramified subsets of \mathbb{R}^2 . Define $\text{web}_D(\mathcal{A})$ as in (2.12). Then $\text{web}_D(\mathcal{A}) = \text{web}_{D'}(\mathcal{A})$.*

PROOF: Let \mathcal{A} be a weave and let D, D' be dense non-ramified subsets of \mathbb{R}^2 . We note that it suffices to prove that

$$(\mathcal{A}|_D)_\uparrow \subseteq \overline{(\mathcal{A}|_{D'})_\uparrow}. \quad (4.24)$$

With (4.24) in hand, by symmetry we also have that $(\mathcal{A}|_{D'})_\uparrow \subseteq \overline{(\mathcal{A}|_D)_\uparrow}$. Taking closures shows that $\text{web}_D(\mathcal{A}) = \text{web}_{D'}(\mathcal{A})$.

Let f be an element of the left hand side of (4.24). Then there exists $z \in D$ and $g \in \mathcal{A}(z)$ with $f \subseteq g|_z$. Using that D' is dense, take $z_n \in D'$ such that $z_n \rightarrow z$, and using that \mathcal{A} is pervasive take $g'_n \in \mathcal{A}(z_n)$. By compactness of \mathcal{A} we may pass to a subsequence and assume that $g'_n \rightarrow g \in \mathcal{A}(z)$, which by Lemma A.2.2 implies that $g'_n|_{z_n} \rightarrow g|_z$.

Since $g, g' \in \mathcal{A}(z)$ and z is non-ramified, we have that $g|_z = g'|_z$. Thus $f \subseteq g'|_z$. Let w denote the initial point of f . Using that $g'_n|_{z_n} \rightarrow g'|_z$, by Lemma A.1.1 we thus have

$$w \in H(g'|_z) \subseteq \overline{\bigcup_{n=1}^{\infty} H(g'_n|_{z_n})}.$$

If $w \in H(g'_n|_{z_n})$ for some $n \in \mathbb{N}$ then we have $f \subseteq g'_n|_{z_n} \in \mathcal{A}|_{D'}$, and we are done. Otherwise, there exists a subsequence of n such that $w_n \in H(g'_n|_{z_n})$ and $w_n \rightarrow w$, and without loss of generality we pass to this subsequence. We thus have $g'_n|_{w_n} \subseteq g'_n|_{z_n}$, so $g'_n|_{w_n} \in (\mathcal{A}_{D'})_\uparrow$. Noting that $w_n \rightarrow w$ and $g'_n \rightarrow g$, it follows from Lemma A.2.2 that $g'_n|_{w_n} \rightarrow f$. Hence $f \in \overline{(\mathcal{A}_{D'})_\uparrow}$. This establishes (4.24) and completes the proof. \blacksquare

Remark 4.5.4 From the point onwards we will often invoke Lemma 4.5.3 implicitly, through writing $\text{web}_D(\mathcal{A}) = \text{web}(\mathcal{A})$. Lemmas 4.4.6 and Lemma 4.5.3 combine to show that $\text{web} : \mathcal{W}_{\text{det}} \rightarrow \mathcal{W}_{\text{det}}$ is a deterministic function.

Lemma 4.5.5 *Let \mathcal{A}, \mathcal{B} be deterministic weaves and assume that $\mathcal{A} \cup \mathcal{B}$ is non-crossing. Then $\text{web}(\mathcal{A}) = \text{web}(\mathcal{B})$.*

PROOF: Since $\mathcal{A} \cup \mathcal{B}$ is non-crossing, it is trivial to check that $\mathcal{A} \cup \mathcal{B}$ is a weave. By Lemma 4.4.6 there exists $D \subseteq \mathbb{R}^2$ that is dense and non-ramified with respect to $\mathcal{A} \cup \mathcal{B}$. Note that this implies D is also non-ramified with respect to both \mathcal{A} and \mathcal{B} . Fix some $z \in D$ and consider $f \in \mathcal{A}(z)$ and $g \in \mathcal{B}(z)$. Note that $(\mathcal{A} \cup \mathcal{B})(z) = \mathcal{A}(z) \cup \mathcal{B}(z)$. We have $f, g \in (\mathcal{A} \cup \mathcal{B})(z)$ and z is non-ramified, so we have $f|_z = g|_z$. Since z, f, g were arbitrary, by (2.12) this implies $\text{web}_D(\mathcal{A}) = \text{web}_D(\mathcal{B})$. ■

Lemma 4.5.6 *Let \mathcal{A} be a deterministic weave. Then $\text{web}(\mathcal{A})$ is a deterministic weave.*

PROOF: Fix a dense non-ramified $D \subseteq \mathbb{R}^2$. By Lemma 4.5.3 it suffices to check that $\mathcal{W} = \text{web}_D(\mathcal{A})$ is a weave. Noting that $(\mathcal{A}|_D)_\uparrow \subseteq \mathcal{A}_\uparrow$, compactness of \mathcal{A} and Lemma 4.5.1 implies compactness of \mathcal{W} . Similarly, \mathcal{W} inherits the non-crossing property from \mathcal{A} by Lemma 4.5.2. Lastly, for any $z \in \mathbb{R}^2$ there exists $(z_n) \subseteq D$ such that $z_n \rightarrow z$. Take $f_n \in \mathcal{A}(z)$ and note $g_n = f_n|_{z_n} \in \mathcal{W}$. By compactness g_n has a sub-sequential limit point $g \in \mathcal{W}$, and by Lemma A.2.2 we have $g \in \mathcal{W}(z)$. Thus \mathcal{W} is pervasive, so we have that \mathcal{W} is a weave. ■

Lemma 4.5.7 *Let \mathcal{A} be a deterministic weave. Then $\text{web}(\mathcal{A})$ is a minimal element of $(\mathscr{W}_{\text{det}}, \preceq)$ and $\text{web}(\mathcal{A}) \preceq \mathcal{A}$.*

PROOF: Let $D \subseteq \mathbb{R}^2$ be non-ramified with respect to \mathcal{A} and let us write $\mathcal{W} = \text{web}_D(\mathcal{A})$. Note that Lemma 4.5.6 gives that $\mathcal{W} \in \mathscr{W}_{\text{det}}$. Let us first show that $\mathcal{W} \preceq \mathcal{A}$. It is immediate from (2.12) that $\mathcal{W} \subseteq \overline{(\mathcal{A}_\uparrow)}$, so by Lemma 4.5.1 and compactness of \mathcal{A} we have $\mathcal{W} \subseteq \mathcal{A}_\uparrow$. Proposition A.2.1 implies that $\mathcal{A}|_D$ is relatively compact, thus from Lemma 4.5.1 and (2.12) we have $\mathcal{W}_\uparrow = \overline{((\mathcal{A}|_D)_\uparrow)_\uparrow} = \overline{(\mathcal{A}|_D)_\uparrow} = \mathcal{W}$, so trivially $\mathcal{W}_\uparrow \cap \mathcal{A} \subseteq \mathcal{W}$. According to (2.7) we now have $\mathcal{W} \preceq \mathcal{A}$.

Suppose that \mathcal{B} is a deterministic weave, comparable to $\text{web}(\mathcal{A})$. We must show that $\text{web}(\mathcal{A}) \preceq \mathcal{B}$. The argument is analogous to that of Lemma 4.4.5. From Lemma 3.1.7 we have that $\text{web}(\mathcal{A}) \cup \mathcal{B}$ is non-crossing. From what we have already proved we have that $\text{web}(\mathcal{A}) \preceq \mathcal{A}$ so using Lemma 3.1.7 again gives that $\mathcal{A} \cup \text{web}(\mathcal{A})$ is non-crossing. Thus by Lemma 4.1.3 we have that $\mathcal{A} \cup \mathcal{B}$ is non-crossing. Lemma 4.5.5 now gives that $\text{web}(\mathcal{A}) = \text{web}(\mathcal{B})$. From what we have already proved we have that $\text{web}(\mathcal{A}) = \text{web}(\mathcal{B}) \preceq \mathcal{B}$, as required. ■

5 Random weaves

We now turn our attention to random weaves. We will give the proof of our main results (stated in Section 2.4) in Sections 5.3–5.7. We require some technical matters to be dealt with first, largely concerning measurability, before we are in a position to rigorously work with random weaves. The proofs in Sections 5.1 and 5.2 are not necessary for the reader wishing to understand the proofs of our main results in later sections.

5.1 On measurability

To make sense of the statements of our main results in Section 2.4 we require that several objects are measurable. For example, if \mathcal{A} and \mathcal{B} are random weaves the we need $\{\mathcal{A} \preceq \mathcal{B}\}$ to be an event and we need $\text{web}(\mathcal{A})$ and $\text{flow}(\mathcal{A})$ to be random variables. In Appendix A.3 we show

that several basic maps associated to the space $\mathcal{K}(\Pi^\uparrow)$ are measurable, whereas here we consider measurability with a focus specific to weaves. We have also seen in Lemma 4.4.3 that the map $\mathcal{A} \mapsto \text{flow}(\mathcal{A})$ is a continuous map from \mathscr{W}_{det} to itself. The map $\text{web}(\cdot)$ is not continuous on \mathscr{W}_{det} , as shown by example in Figure 2.5.1. We defined web and flow as maps with domain \mathscr{W}_{det} but we check measurability in terms of the topology on $\mathcal{K}(\Pi^\uparrow)$, so the following lemma is a technical necessity.

Lemma 5.1.1 *It holds that \mathscr{W}_{det} is a measurable subset of $\mathcal{K}(\Pi^\uparrow)$.*

PROOF: We have $\mathscr{W}_{\text{det}} = \{A \in \mathcal{K}(\Pi^\uparrow); A \text{ is pervasive}\} \cap \{A \in \mathcal{K}(\Pi^\uparrow); A \text{ is non-crossing}\}$. Using Lemma A.2.2 it is straightforward to check that $\{A \in \mathcal{K}(\Pi^\uparrow); A \text{ is pervasive}\}$ is closed. It remains to show that $\{A \in \mathcal{K}(\Pi^\uparrow); A \text{ is non-crossing}\}$ is measurable.

Consider if $A_n, A \in \mathcal{K}(\Pi)$ are such that $A_n \rightarrow A$, and A_n is non-crossing for each n . Suppose that A fails to be non-crossing, in particular suppose that $f, g \in A$ cross each other. We have $f_n, g_n \in \mathcal{A}_n$ such that $f_n \rightarrow f$ and $g_n \rightarrow g$. By Lemma 3.1.3 we have $f_n \triangleleft g_n$ or $g_n \triangleleft f_n$, at least one of which must hold for infinitely many n . It follows by part 2 of Lemma 3.3.9 that A contains a pair of paths f', g' such that $f' \triangleleft_\epsilon g'$, for some $\epsilon > 0$. Writing

$$\begin{aligned} N &= \{A \in \mathcal{K}(\Pi^\uparrow); A \text{ is non-crossing}\}, \\ M_\epsilon &= \{A \in \mathcal{K}(\Pi^\uparrow); \text{there exists } f, g \in A \text{ with } f \triangleleft_{1/n} g\} \\ L &= \{A \in \mathcal{K}(\Pi^\uparrow); \text{there exists } f \in A \text{ with } \sigma_f = -\infty \text{ or } \sigma_f = +\infty\}. \end{aligned}$$

we thus obtain $\overline{N} = N \cup (\overline{N} \cap \bigcup_{n \in \mathbb{N}} M_{1/n})$. By part 1 of Lemma 3.3.9 we have $N \cap M_\epsilon = \emptyset$, hence $N = \overline{N} \setminus (\bigcup_{n \in \mathbb{N}} M_{1/n})$. Lemma 3.3.11 implies that $M_{1/n} \cup L$ is closed, for each $n \in \mathbb{N}$, so noting that L is also closed we obtain that $M_{1/n} = (M_{1/n} \cup L) \setminus L$ is measurable. Hence N is measurable. This completes the proof. \blacksquare

Lemma 5.1.2 *The map $\mathcal{A} \mapsto \text{web}(\mathcal{A})$ is measurable from \mathscr{W}_{det} to itself.*

PROOF: From Lemma 5.1.1 we have that \mathscr{W}_{det} is measurable. Recall that in Remark 4.5.4 we noted that for $\mathcal{A} \in \mathscr{W}_{\text{det}}$ the value of $\text{web}_D(\mathcal{A})$ does not depend upon D , provided that $D \subseteq \mathbb{R}_c^2$ is dense and non-ramified. We thus write $\text{web}(\mathcal{A}) = \text{web}_D(\mathcal{A})$.

Let μ be a measure on \mathbb{R}_c^2 with full support and no atoms. Let $(z_i)_{i=1}^\infty$ be a sequence of independent random variables with distribution μ , on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The following argument is somewhat unusual, so let us give an outline. We will show that the map from $\mathcal{A} \in \mathscr{W}_{\text{det}}$ to the law of $\text{web}_D(\mathcal{A})$ is a measurable map, and that this law is precisely the probability measure with a point-mass at $\text{web}(\mathcal{A})$. The stated result then follows, using that the map from point-mass measures, to their associated points, is a measurable map.

From Lemma A.3.2 the map $(A, \omega) \mapsto A|_{z_i(\omega)}$ is measurable from $\mathcal{K}(\Pi) \times \Omega \rightarrow \mathcal{K}(\Pi)$. It follows from Lemma A.3.3 that

$$(A, \omega) \mapsto M_n(A, \omega) = \left(\bigcup_{i=1}^n A|_{z_i(\omega)} \right)_\uparrow \quad (5.1)$$

is measurable, for each $n \in \mathbb{N}$, as a function $M_n : \mathscr{W}_{\text{det}} \times \Omega \rightarrow \mathscr{W}_{\text{det}}$.

For all $\mathcal{A} \in \mathscr{W}_{\text{det}}$ and $\omega \in \Omega$ we have $M_n(\mathcal{A}, \omega) \subseteq M_{n+1}(\mathcal{A}, \omega) \subseteq \mathcal{A}_\uparrow$. Lemma 4.5.1 gives that \mathcal{A}_\uparrow is a compact subset of Π , which implies that the sequence $(M_n(\mathcal{A}, \omega))_{n=1}^\infty$ is a relatively compact subset of $\mathcal{K}(\Pi)$. Let \mathcal{B} be a limit point in $\mathcal{K}(\Pi^\uparrow)$ of $(M_n(\mathcal{A}, \omega))$ as $n \rightarrow \infty$.

We aim to show that $\mathbb{P}[B = \text{web}(\mathcal{A})] = 1$. For each $\mathcal{A} \in \mathcal{W}_{\text{det}}$ we have that

$$\mathbb{P}[(z_i) \text{ is dense in } \mathbb{R}^2 \text{ and non-ramified in } \mathcal{A}] = 1. \quad (5.2)$$

We write $D = (z_i(\omega))_{i \in \mathbb{N}}$. Let us condition on the event in (5.2) occurring. Then, if $g \in B$ we have $z_{i_n} \subseteq \mathbb{R}^2$, $f_n \in \mathcal{A}(z_{i_n})$ and $g_n \in \Pi^\uparrow$ such that $g_n \subseteq f_n|_{z_{i_n}}$ and $g_n \rightarrow g$. It follows immediately that $g \in \text{web}_D(\mathcal{A})$, thus $B \subseteq \text{web}_D(\mathcal{A})$. Similarly, if $g \in \text{web}_D(\mathcal{A})$ then there exists $z_i \in \mathbb{R}^2$, $f_i \in \mathcal{A}(z_i)$ and $g_i \in \Pi^\uparrow$ such that $g_i \subseteq f_i$ and $g_i \rightarrow g$. Hence for all i there exists $n \in \mathbb{N}$ such that $g_i \in M_n(\mathcal{A}, \omega)$. Thus $g \in B$ and $B = \text{web}_D(\mathcal{A})$, which by Lemma 4.5.3 is equal to $\text{web}(\mathcal{A})$. We thus have $\mathbb{P}[B = \text{web}(\mathcal{A})] = 1$.

Let $\mathcal{P}(\mathcal{W}_{\text{det}})$ denote the space of probability measures on \mathcal{W}_{det} and let $\mathcal{P}_0(\mathcal{W}_{\text{det}})$ denote the closed subspace of point-mass probability measures. Let $\delta_A \in \mathcal{P}_0(\mathcal{W}_{\text{det}})$ be the probability measure that is a point-mass on $A \in \mathcal{W}_{\text{det}}$. Let $\mathcal{L}_n^A \in \mathcal{P}(\mathcal{W}_{\text{det}})$ denote the law of the random variable $\omega \mapsto M_n(\mathcal{A}, \omega)$. From what we have proved, it follows that \mathcal{L}_n^A converges weakly to $\delta_{\text{web}(\mathcal{A})}$ the probability measure on $\mathcal{K}(\Pi^\uparrow)$ that is a point-mass on $\text{web}(\mathcal{A})$.

For measurable $S \subseteq \mathcal{W}_{\text{det}}$ we have

$$\mathcal{L}_n^A(S) = \mathbb{P}[M_n(\mathcal{A}, \cdot) \in S] = \int_{\omega \in \Omega} \mathbb{1}_S(M_n(\mathcal{A}, \omega)) d\mathbb{P}(\omega),$$

from which it follows that $\mathcal{A} \mapsto \mathcal{L}_n^A$ is a measurable function from \mathcal{W}_{det} to $\mathcal{P}_0(\mathcal{W}_{\text{det}})$. Hence $\mathcal{A} \mapsto \delta_{\text{web}(\mathcal{A})}$ is also measurable. It is easily seen that the map $m : \mathcal{P}_0(\mathcal{W}_{\text{det}}) \rightarrow \mathcal{W}_{\text{det}}$ given by $\delta_A \mapsto A$ is continuous, and thus measurable. Compositions of measurable functions are measurable, hence the map $A \mapsto m(\delta_{\text{web}(\mathcal{A})}) = \text{web}(\mathcal{A})$ is measurable. \blacksquare

Lemma 5.1.3 *The set $\{(A, B) \in \mathcal{K}(\Pi)^2; A \preceq B\}$ is a measurable subset of $\mathcal{K}(\Pi)^2$.*

PROOF: Recall the definition of \preceq on $\mathcal{K}(\Pi)$ from (2.7). It suffices to show that

$$\begin{aligned} C_1 &= \{(A, B) \in \mathcal{K}(\Pi)^2; A \subseteq B_\uparrow\} \\ C_2 &= \{(A, B) \in \mathcal{K}(\Pi)^2; B \cap A_\uparrow \subseteq A\} \end{aligned}$$

are both measurable subsets of $\mathcal{K}(\Pi)^2$, where $\mathcal{K}(\Pi)^2$ is equipped with the product topology and corresponding Borel σ -field. We note that A_\uparrow is closed whenever $A \in \mathcal{K}(\Pi)$ is closed. Using Lemma A.2.2 it is straightforward to check that the set $C_1 = \{(A, B) \in \mathcal{K}(\Pi)^2; \forall f \in A \exists g \in B \text{ such that } f \subseteq g\}$ is closed. We now move on to C_2 . To this end note that $B \cap A_\uparrow \not\subseteq A$ if and only if there exists $f \in A_\uparrow$ and $g \in B$ such that $g \subseteq f$ and $f \not\subseteq A$. The condition $f \not\subseteq A$ is equivalent to $d_{\mathcal{K}(\Pi)}(\{f\}, A) > 0$ which, as A is closed, is in turn equivalent to $d_{\mathcal{K}(\Pi)}(\{f\}, A) \geq \epsilon$ for some $\epsilon > 0$. We thus have that $C_2 = \bigcup_{n \in \mathbb{N}} S_{1/n}$ where

$$S_\epsilon = \{(A, B) \in \mathcal{K}(\Pi)^2; \exists f \in A_\uparrow, g \in B \text{ such that } g \subseteq f \text{ and } d_{\mathcal{K}(\Pi)}(\{f\}, A) \geq \epsilon\}.$$

Similar to above, using Lemma A.2.2 it is straightforward to check that S_ϵ is closed, for any $\epsilon > 0$. Thus C_2 is measurable. \blacksquare

5.2 On partial ordering of random weaves

In this section we show that \preceq_d is a partial order on (the laws of) random weaves, as defined shortly below (2.7). More precisely, recall that $\mathcal{P}(M)$ denotes the space of probability measures on a metric space M . We have shown in Lemma 3.2.2 that \preceq given by (2.7) defines a partial

order on $\mathcal{K}(\Pi^\uparrow)$. In Section 2.3 we defined an extension of \preceq to $\mathcal{P}(\mathcal{K}(\Pi^\uparrow))$, namely if \mathcal{A} and \mathcal{B} are $\mathcal{K}(\Pi^\uparrow)$ valued random variables then we write $\mathcal{A} \preceq_d \mathcal{B}$ if there exists a coupling of \mathcal{A} and \mathcal{B} such that $\mathbb{P}[\mathcal{A} \preceq \mathcal{B}] = 1$. We aim to show that that \preceq_d is a partial order on $\mathcal{P}(\mathcal{K}(\Pi^\uparrow))$.

If \preceq was compatible with $(\mathcal{K}(\Pi), d_\Pi)$, in the sense of Definition 3.3.1, then we could use a classical result e.g. Theorem 2.4 in Liggett (1985) to obtain the extension to $\mathcal{P}(\mathcal{K}(\Pi))$. However, as we saw in Remark 3.3.2 compatibility fails in this situation. Instead we require an original argument that uses compactness and the precise form of (2.7). We first give a preliminary lemma.

Lemma 5.2.1 *Suppose that $\mathcal{D}, \mathcal{D}'$ are $\mathcal{K}(\Pi^\uparrow)$ valued random variables, with the same marginal distributions, coupled such that $\mathbb{P}[\mathcal{D} \preceq \mathcal{D}'] = 1$. Then $\mathbb{P}[\mathcal{D} = \mathcal{D}'] = 1$.*

PROOF: By Proposition 2.2.1 the metric space Π^\uparrow is separable, which implies that its topology has a countable base: there exists a family $(U_i)_{i \in \mathbb{N}}$ of non-empty open subsets of Π^\uparrow such that any open subset $O \subseteq \Pi$ can be written as $O = \cup_{i \in I} U_i$ for some $I \subseteq \mathbb{N}$.

Assume the conditions of the lemma on $\mathcal{D}, \mathcal{D}'$. The proof comes in two parts, corresponding respectively to the inequalities $\mathbb{P}[\mathcal{D}' \setminus \mathcal{D} \neq \emptyset] > 0$ and $\mathbb{P}[\mathcal{D} \setminus \mathcal{D}' \neq \emptyset] > 0$, each of which will be shown to be impossible through an argument by contradiction.

Part 1. Suppose $\mathbb{P}[\mathcal{D}' \setminus \mathcal{D} \neq \emptyset] > 0$. The reader may wish to glance at Remark 5.2.2, immediately below the present proof, for a toy example to illustrate our strategy here. For $A \subseteq \Pi$, we write

$$A^\circ = \{b \in \Pi^\uparrow; \text{there exists } a \in A \text{ such that } a \subseteq b\}. \quad (5.3)$$

On the event that $\{\mathcal{D}' \setminus \mathcal{D}\}$, let $d' \in \mathcal{D}' \setminus \mathcal{D}$ and let $\mathcal{B}_\epsilon(d')$ be the open ball in Π^\uparrow of radius ϵ about d' . We will now show that, almost surely, $\mathcal{B}_\epsilon(d')^\circ \cap \mathcal{D}$ is empty, for sufficiently small $\epsilon > 0$. Suppose that $\mathcal{B}_\epsilon(d')^\circ \cap \mathcal{D} \neq \emptyset$ for all $\epsilon > 0$. Then, taking $\epsilon = 1/n$, we have sequences $f_n \in \mathcal{B}_{1/n}(d')$ and $g_n \in \mathcal{B}_{1/n}(d')^\circ \cap \mathcal{D}$, with $f_n \subseteq g_n$. By compactness of \mathcal{D} we may pass to subsequence and assume convergence $f_n \rightarrow d' \in \mathcal{D}'$ and $g_n \rightarrow d \in \mathcal{D}$. By Lemma A.2.2 we then have $d' \subseteq d$, so $d' \in \mathcal{D}_\uparrow$. We have $\mathbb{P}[\mathcal{D} \preceq \mathcal{D}'] = 1$ upon which event $\mathcal{D}_\uparrow \cap \mathcal{D}' \subseteq \mathcal{D}$, so $d' \in \mathcal{D}$ which is a contradiction. Thus, almost surely, for some (random) $\epsilon > 0$, we have $\mathcal{B}_\epsilon(d')^\circ \cap \mathcal{D} = \emptyset$. Clearly also $d' \in \mathcal{B}_\epsilon(d')$.

Let

$$O = \begin{cases} \mathcal{B}_\epsilon(d') & \text{on the event that } \mathcal{D}' \setminus \mathcal{D} \neq \emptyset \\ \emptyset & \text{otherwise.} \end{cases}$$

From the previous paragraph have that $O^\circ \cap \mathcal{D} = \emptyset$ and with positive probability $d' \in O$. Since O is open, almost surely we may write $O = \cup_{i \in I} U_i$ for some random $I \subseteq \mathbb{N}$. The set I is non-empty with positive probability, hence there is some deterministic $i \in I$ such that with positive probability $d' \in U_i \subseteq O$. Let us write $U = U_i$ for such an i .

On the event that $d' \in U \subseteq O$ we have that $U^\circ \subseteq O^\circ$, which implies that $U^\circ \cap \mathcal{D} = \emptyset$ (because $O^\circ \cap \mathcal{D} = \emptyset$) and $U^\circ \cap \mathcal{D}' \neq \emptyset$ (because it contains d'). Hence the event $\{U^\circ \cap \mathcal{D}' \neq \emptyset \text{ and } U^\circ \cap \mathcal{D} = \emptyset\}$ has positive probability. We thus have

$$\begin{aligned} 0 < \mathbb{P}[U^\circ \cap \mathcal{D}' \neq \emptyset \text{ and } U^\circ \cap \mathcal{D} = \emptyset] &= \mathbb{P}[U^\circ \cap \mathcal{D}' \neq \emptyset] - \mathbb{P}[U^\circ \cap \mathcal{D}' \neq \emptyset \text{ and } U^\circ \cap \mathcal{D} \neq \emptyset] \\ &= \mathbb{P}[U^\circ \cap \mathcal{D} \neq \emptyset] - \mathbb{P}[U^\circ \cap \mathcal{D}' \neq \emptyset \text{ and } U^\circ \cap \mathcal{D} \neq \emptyset] \\ &= \mathbb{P}[U^\circ \cap \mathcal{D} \neq \emptyset \text{ and } U^\circ \cap \mathcal{D}' = \emptyset]. \end{aligned} \quad (5.4)$$

The second line of (5.4) follows because \mathcal{D} and \mathcal{D}' have the same marginal distribution, and the other steps are elementary.

Consider when the event $\{U^\circ \cap \mathcal{D} \neq \emptyset \text{ and } U^\circ \cap \mathcal{D}' = \emptyset\}$ occurs, which by (5.4) has positive probability. Then we have $h \in U^\circ \cap \mathcal{D}$, but $\mathbb{P}[\mathcal{D} \preceq \mathcal{D}'] = 1$ upon which event we have $\mathcal{D} \subseteq \mathcal{D}'_\uparrow$, hence there exists $h' \in \mathcal{D}'$ such that $h \subseteq h'$. By (5.3) we have that $h' \in U^\circ$, which is a contradiction to $U^\circ \cap \mathcal{D}' = \emptyset$. Hence in fact $\mathbb{P}[\mathcal{D}' \setminus \mathcal{D} \neq \emptyset] = 0$, as required.

Part 2. Suppose $\mathbb{P}[\mathcal{D} \setminus \mathcal{D}' \neq \emptyset] > 0$. The argument is similar to Case 1 but somewhat simpler, and we will make use of Case 1 within it. Note that we should expect an asymmetric argument due to the parity inherent in $\mathbb{P}[\mathcal{D} \preceq \mathcal{D}'] = 1$. To make the comparison clear we will recycle much of our notation.

On the event $\mathcal{D} \setminus \mathcal{D}' \neq \emptyset$, take $d \in \mathcal{D} \setminus \mathcal{D}'$. Suppose that $B_\epsilon(d) \cap \mathcal{D}' \neq \emptyset$ for all $\epsilon > 0$. Taking $\epsilon = 1/n$, there exists $f_n \in \mathcal{D}' \cap B_{1/n}(d)$. By compactness we may pass to a subsequence and assume convergence $f_n \rightarrow f \in \mathcal{D}'$. This implies $f = d$, which is a contradiction since $f \in \mathcal{D}'$. Hence there exists a random $\epsilon > 0$ such that $B_\epsilon(d) \cap \mathcal{D}' = \emptyset$.

Let O be equal to $B_\epsilon(d)$ on the event $\{\mathcal{D} \setminus \mathcal{D}' \neq \emptyset\}$ and $O = \emptyset$ otherwise. Thus $O \cap \mathcal{D}' = \emptyset$ and with positive probability $d \in O$. By the same argument as in Part 1, there exists deterministic $i \in \mathbb{N}$ such that with positive probability $d \in U_i \subseteq O$. Thus, setting $U = U_i$, we have that with positive probability $d \in U \cap (\mathcal{D} \setminus \mathcal{D}')$. We have $\mathbb{P}[\mathcal{D} \preceq \mathcal{D}'] = 1$ upon which event $\mathcal{D} \subseteq \mathcal{D}'_\uparrow$. Thus, when $d \in U \cap (\mathcal{D} \setminus \mathcal{D}')$ there exists $d' \in \mathcal{D}'$ such that $d \subseteq d'$, implying that both $U^\circ \cap (\mathcal{D} \setminus \mathcal{D}') \neq \emptyset$ and $U^\circ \cap \mathcal{D}' \neq \emptyset$. From Part 1 we have that almost surely $\mathcal{D}' \subseteq \mathcal{D}$. Hence with positive probability we have both $U^\circ \cap (\mathcal{D} \setminus \mathcal{D}') \neq \emptyset$ and $U^\circ \cap (\mathcal{D}' \cap \mathcal{D}) \neq \emptyset$. Thus, noting that $\mathcal{D} = (\mathcal{D} \cap \mathcal{D}') \cup (\mathcal{D} \setminus \mathcal{D}')$,

$$\begin{aligned} 0 < \mathbb{P}[U^\circ \cap (\mathcal{D} \cap \mathcal{D}') \neq \emptyset] &< \mathbb{P}[U^\circ \cap \mathcal{D} \neq \emptyset] \\ &= \mathbb{P}[U^\circ \cap \mathcal{D}' \neq \emptyset]. \end{aligned}$$

Here, the second line follows because \mathcal{D} and \mathcal{D}' have identical distribution and U is deterministic. It follows that $\mathbb{P}[U^\circ \cap (\mathcal{D}' \setminus \mathcal{D}) \neq \emptyset] > 0$, but from Part 1 we know that $\mathbb{P}[\mathcal{D}' \setminus \mathcal{D} = \emptyset] = 1$, so we have reached a contradiction. This completes the proof. \blacksquare

Remark 5.2.2 The proof of Lemma 5.2.1 is technical but it has a simple idea at its heart. Consider a toy example: take two uniform random variables X, X' on $S = \{1, 2, 3, 4, 5, 6\}$ and suppose that X and X' are coupled in a way that satisfies $\mathbb{P}[X \leq X'] = 1$. We aim to show that $\mathbb{P}[X = X'] = 1$. For $k \in S$,

$$\begin{aligned} \mathbb{P}[X = k, X' \neq k] &= \mathbb{P}[X = k] - \mathbb{P}[X = k, X' = k] \\ &= \mathbb{P}[X' = k] - \mathbb{P}[X' = k, X = k] \\ &= \mathbb{P}[X' = k, X \neq k]. \end{aligned} \tag{5.5}$$

Note the similarity of (5.5) to (5.4). Taking $k = 1$ and using that $\mathbb{P}[X \leq X'] = 1$ we obtain $\mathbb{P}[X = 1, X' > 1, X \leq X'] = \mathbb{P}[X' = 1, X > 1, X \leq X']$ which becomes $\mathbb{P}[X = 1, X' > 1] = 0$, thus $\mathbb{P}[X = 1] = \mathbb{P}[X = 1, X' = 1]$. This is clearly a step in the right direction and is in similar style to the more complex reasoning involving \preceq below (5.4). The finiteness of S is also helpful here whereas in Lemma 5.2.3 we must rely on second countability of Π^\uparrow . We leave it for the reader to complete this toy example and deduce that $\mathbb{P}[X = X'] = 1$.

Lemma 5.2.3 *The relation \preceq_d is a partial order on the space of $\mathcal{K}(\Pi^\uparrow)$ valued random variables.*

PROOF: We will check that \preceq_d on $\mathcal{P}(\mathcal{K}(\Pi^\uparrow))$ is reflexive, transitive and antisymmetric, in turn. Lemma 3.2.2 has already shown that these properties hold in the deterministic case i.e. \preceq_d is a partial order on \mathcal{W}_{det} .

By reflexivity of \preceq on \mathcal{W}_{det} we have that $\mathbb{P}[\mathcal{A} \preceq \mathcal{A}] = 1$ for any $\mathcal{K}(\Pi^\uparrow)$ valued random variable \mathcal{A} , so \preceq_d is reflexive. For transitivity, let us assume that $\mathcal{A}, \mathcal{B}, \mathcal{C}$ are $\mathcal{K}(\Pi^\uparrow)$ valued

random variables, and that we have couplings $(\mathcal{A}, \mathcal{B})$ and $(\mathcal{B}, \mathcal{C})$ such that $\mathbb{P}[\mathcal{A} \preceq \mathcal{B}] = 1$ and (on a possibly different probability space) $\mathbb{P}[\mathcal{B} \preceq \mathcal{C}] = 1$. It follows that there exists a joint coupling $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ on which $\mathbb{P}[\mathcal{A} \preceq \mathcal{B} \text{ and } \mathcal{B} \preceq \mathcal{C}] = 1$. On the event $\{\mathcal{A} \preceq \mathcal{B} \text{ and } \mathcal{B} \preceq \mathcal{C}\}$ transitivity of \preceq on \mathcal{W}_{det} implies that $\mathcal{A} \preceq \mathcal{C}$, so we obtain $\mathbb{P}[\mathcal{A} \preceq \mathcal{C}] = 1$, as required.

It remains to show antisymmetry. Suppose that $\mathcal{A}, \mathcal{A}', \mathcal{B}, \mathcal{B}'$ are $\mathcal{K}(\Pi^\uparrow)$ valued random variables such that $\mathbb{P}[\mathcal{A} \preceq \mathcal{B}] = 1$ and $\mathbb{P}[\mathcal{B}' \preceq \mathcal{A}'] = 1$, where \mathcal{A} and \mathcal{A}' have the same marginal distribution and \mathcal{B} and \mathcal{B}' have the same marginal distribution. We must show that there exists a coupling under which $\mathbb{P}[\mathcal{A} = \mathcal{B}] = 1$. Since \mathcal{A} and \mathcal{A}' have the same marginal distribution, it follows that there exists a coupling $(\mathcal{A}, \mathcal{A}', \mathcal{B}, \mathcal{B}')$ such that $\mathbb{P}[\mathcal{A} = \mathcal{A}', \mathcal{A} \preceq \mathcal{B}, \mathcal{B}' \preceq \mathcal{A}'] = 1$. By transitivity of \preceq on \mathcal{W}_{det} this means that $\mathbb{P}[\mathcal{B}' \preceq \mathcal{B}] = 1$. By Lemma 5.2.1 we have $\mathbb{P}[\mathcal{B} = \mathcal{B}'] = 1$ which means $\mathbb{P}[\mathcal{A} \preceq \mathcal{B}, \mathcal{B} \preceq \mathcal{A}] = 1$ and by antisymmetry of \preceq on \mathcal{W}_{det} we obtain that $\mathbb{P}[\mathcal{A} = \mathcal{B}] = 1$, as required. \blacksquare

Lemma 5.2.4 *Let \mathcal{A} be a weave. Then \mathcal{A} is a web if and only if $\mathbb{P}[\mathcal{A} \text{ is a minimal element of } \mathcal{W}_{\text{det}}] = 1$. Similarly, \mathcal{A} is a flow if and only if $\mathbb{P}[\mathcal{A} \text{ is a maximal element of } \mathcal{W}_{\text{det}}] = 1$.*

PROOF: Let us first give the argument for webs. Let \mathcal{A} be a weave. We must show that \mathcal{A} is a web if and only if $\mathbb{P}[\mathcal{A} \text{ is a minimal element of } \mathcal{W}_{\text{det}}] = 1$. It is trivial to see that almost sure pervasiveness and the non-crossing property pass from either side of the ‘if and only if’ statement to the other side, so it remains only to handle minimality. To be explicit we must show that for a random weave \mathcal{A} the following statements are equivalent:

1. If \mathcal{B} is a weave and there exists a coupling between \mathcal{A} and \mathcal{B} such that $\mathbb{P}[\mathcal{A} \preceq \mathcal{B}] = 1$ or $\mathbb{P}[\mathcal{B} \preceq \mathcal{A}] = 1$, then $\mathbb{P}[\mathcal{A} \preceq \mathcal{B}] = 1$.
2. $\mathbb{P}[\mathcal{A} \text{ is a minimal element of } \mathcal{W}_{\text{det}}] = 1$.

By the definition of \preceq_d from below (2.7), the first statement is precisely the claim that the law of \mathcal{A} is minimal in $\mathcal{P}(\mathcal{K}(\Pi^\uparrow))$.

Let us first show that (2) implies (1). If \mathcal{A} is almost surely minimal in \mathcal{W}_{det} then for any coupling of \mathcal{A} to another weave \mathcal{B} , $\{\mathcal{A} \text{ and } \mathcal{B} \text{ are comparable}\} \subseteq \{\mathcal{A} \preceq \mathcal{B}\}$, where both the left and right hand side are events. Thus (1) holds.

Conversely, let us assume (1). Let

$$\mathcal{A}' = \begin{cases} \mathcal{A} & \text{on the event } \{\text{web}(\mathcal{A}) \prec \mathcal{A}\}, \\ \text{web}(\mathcal{A}) & \text{on the event } \{\mathcal{A} \preceq \text{web}(\mathcal{A})\}. \end{cases}$$

Lemmas 5.1.2 and 5.2.3 combine to show that \mathcal{A}' is a random variable. By Lemma 4.5.6 we have $\mathbb{P}[\text{web}(\mathcal{A}) \preceq \mathcal{A}'] = 1$, which by (1) implies that $\mathbb{P}[\mathcal{A}' \preceq \text{web}(\mathcal{A})] = 1$, so in fact $\mathbb{P}[\mathcal{A} \preceq \text{web}(\mathcal{A})] = 1$. By Lemma 4.5.6 we thus have $\mathbb{P}[\mathcal{A} = \text{web}(\mathcal{A})] = 1$, from which Lemma 4.5.7 gives (2).

In the case of flows we may use a similar argument, reversing the direction of the sign of \preceq . Lemma 4.5.6 is replaced by Lemma 4.4.2, and Lemma 4.5.7 is replaced by Lemma 4.4.5. We leave the details to the reader. \blacksquare

Remark 5.2.5 Suppose that \mathcal{A} is a weave. Our results in Sections 5.1 and 5.2 justify that $\text{web}(\mathcal{A})$ and $\text{flow}(\mathcal{A})$ are random variables, and that if \mathcal{B} some other random weave, coupled to \mathcal{A} then $\{\mathcal{A} \preceq \mathcal{B}\}$ is an event. Moreover, \preceq_d defines a partial order on the laws of random weaves, via the relationship $\mathcal{A} \preceq_d \mathcal{B}$ if and only if there exists a coupling of \mathcal{A} and \mathcal{B} such that $\mathbb{P}[\mathcal{A} \preceq \mathcal{B}] = 1$. We will use these results freely from now on and will not repeatedly cite them when used within the proofs.

5.3 Proof of Theorem 2.4.3

We are now ready to establish our main results concerning random weaves, which henceforth are simply referred to as weaves. These results were stated in Section 2.4 and the proofs are spread across Sections 5.3–5.7. The statements of Theorems 2.4.3–2.4.6 consist of several (numbered) parts. We will use bold text (see e.g. the next paragraph) to track when each part is addressed. Most of our work in Section 4 leads towards these proofs. We begin with Theorem 2.4.3, of which we prove the four statements of the theorem in turn. Let \mathcal{A} be a weave.

Part 1. Let us first assume (a), that \mathcal{A} is a web. By Lemma 5.2.4 \mathcal{A} is almost surely a minimal element of \mathscr{W}_{det} . By Lemma 4.5.7 we have $\text{web}(\mathcal{A}) \preceq \mathcal{A}$, from which minimality implies that $\mathcal{A} = \text{web}(\mathcal{A})$, which gives (b). Conversely, let us assume (b), that $\mathcal{A} \stackrel{\text{a.s.}}{=} \text{web}(\mathcal{A})$. Lemma 4.5.7 thus gives that \mathcal{A} is almost surely a minimal element of \mathscr{W}_{det} , from which Lemma 5.2.4 gives that \mathcal{A} is a web. Thus $(a) \Leftrightarrow (b)$.

Part 2. Again, let \mathcal{A} be a weave. We will show that $(a) \Rightarrow (b) \Rightarrow (c) \Rightarrow (a)$. Let us first assume (a), that \mathcal{A} is a flow. By Lemma 5.2.4 \mathcal{A} is almost surely a maximal element of \mathscr{W}_{det} . By Lemma 4.4.5 we have $\mathcal{A} \preceq \text{flow}(\mathcal{A})$, from which maximality implies that $\mathcal{A} = \text{flow}(\mathcal{A})$, giving (b). Now let us assume (b), that $\mathcal{A} \stackrel{\text{a.s.}}{=} \text{flow}(\mathcal{A})$. It is immediate from (2.13) that $\text{flow}(\mathcal{A}) \subseteq \Pi^\uparrow$, so we have (c).

Lastly, let us assume (c), that $\mathcal{A} \subseteq \Pi^\uparrow$. Suppose that \mathcal{B} is a weave with a coupling to \mathcal{A} such that $\mathbb{P}[\mathcal{A} \preceq \mathcal{B}] = 1$. To see that \mathcal{A} is a flow we must show that this implies $\mathbb{P}[\mathcal{A} = \mathcal{B}] = 1$. Using that $\mathcal{A} \preceq \mathcal{B}$ almost surely, it follows from (2.7) that almost surely $\mathcal{A}_\uparrow \cap \mathcal{B} \subseteq \mathcal{A} \subseteq \mathcal{B}_\uparrow$. As $\mathcal{A} \in \Pi^\uparrow$ we thus have that almost surely $\mathcal{A} \subseteq \mathcal{B}$. We require the reverse inclusion, so let $f \in \mathcal{B}$. On the almost sure event that $\mathcal{A} \preceq \mathcal{B}$, by Lemma 3.1.7 we have that f does not cross \mathcal{A} , so by Theorem 4.3.9 there exists $f' \in \text{flow}(\mathcal{A})$ such that $f \subseteq f'$. Lemma 4.2.6 gives that $f' \in \mathcal{A}$, which implies that $f \in \mathcal{A}_\uparrow \cap \mathcal{B}$. Thus $\mathcal{A} \stackrel{\text{a.s.}}{=} \mathcal{B}$, as required.

Part 3. Lemmas 4.4.5 and 4.5.7 give that $\mathbb{P}[\text{web}(\mathcal{A}) \preceq \mathcal{A} \preceq \text{flow}(\mathcal{A})] = 1$.

Part 4. The existence claim is established by part 3 of the present proof. It remains to prove the uniqueness claim, which we will give in turn for webs and then flows.

Let $\mathcal{W}, \mathcal{W}'$ be webs and suppose that $\mathcal{W} \preceq_d \mathcal{A}$ and $\mathcal{W}' \preceq_d \mathcal{A}$. Then there exists (pairwise) couplings such that $\mathbb{P}[\mathcal{W} \preceq \mathcal{A}] = 1$ and $\mathbb{P}[\mathcal{W}' \preceq \mathcal{A}] = 1$. We seek to show that $\mathcal{W} \stackrel{d}{=} \mathcal{W}'$. It follows that there exists a three-way coupling of $\mathcal{W}, \mathcal{W}'$ and \mathcal{A} such that $\mathbb{P}[\mathcal{W} \preceq \mathcal{A} \text{ and } \mathcal{W}' \preceq \mathcal{A}] = 1$. By Lemma 3.1.7 we have that $\mathcal{W} \cup \mathcal{A}$ is almost surely non-crossing, and $\mathcal{W}' \cup \mathcal{A}$ is almost surely non-crossing. By Lemma 4.1.3 we have that $\mathcal{W} \cup \mathcal{W}'$ is almost surely non-crossing. By Lemma 4.5.5 we thus have $\text{web}(\mathcal{W}) \stackrel{\text{a.s.}}{=} \text{web}(\mathcal{W}')$. By part 2 of the present proof we thus have $\mathcal{W} \stackrel{\text{a.s.}}{=} \mathcal{W}'$, hence in particular \mathcal{W} and \mathcal{W}' have the same marginal distribution, as required.

It remains to prove a corresponding statement for flows. Let $\mathcal{F}, \mathcal{F}'$ be flows and suppose that $\mathcal{A} \preceq_d \mathcal{F}$ and $\mathcal{A} \preceq_d \mathcal{F}'$. Then, as above, there exists a three-way coupling of $\mathcal{F}, \mathcal{F}'$ and \mathcal{A} such that $\mathbb{P}[\mathcal{A} \preceq \mathcal{F} \text{ and } \mathcal{A} \preceq \mathcal{F}'] = 1$. By the same argument as above, again using Lemmas 3.1.7 and 4.1.3, with Lemma 4.4.4 in place of Lemma 4.5.5, and using part 1 of the present proof in place of part 2, we obtain that $\mathcal{F} \stackrel{\text{a.s.}}{=} \mathcal{F}'$. Hence in particular \mathcal{F} and \mathcal{F}' have the same marginal distribution, as required. This completes the proof.

5.4 Proof of Theorem 2.4.4

We require some preparatory lemmas before giving the proof of Theorem 2.4.4.

The next lemma gives us the ability use a *deterministic* non-ramified dense set of space-time points with (random) weaves. It is a straightforward consequence of Lemma 4.4.6, delayed until now because when we stated Lemma 4.4.6 we were focused on deterministic weaves.

Lemma 5.4.1 *Let \mathcal{A} be a weave. Then the set $\{z \in \mathbb{R}_c^2; \mathbb{P}[z \text{ is ramified in } \mathcal{A}] > 0\}$ has zero Lebesgue measure.*

PROOF: Let $\text{ram}(\mathcal{A})$ denote the set of ramification points of a (deterministic or random) weave \mathcal{A} . We have shown in Lemma 4.4.6 that the map $(\mathcal{A}, z) \mapsto \mathbb{1}\{z \in \text{ram}(\mathcal{A})\}$ is measurable from $\mathcal{W}_{\text{det}} \times \mathbb{R}_c^2 \rightarrow \{0, 1\}$, and that $\text{ram}(\mathcal{A})$ is Lebesgue null for all deterministic weaves. By Fubini's theorem, for any weave \mathcal{A} we have $\int_{\mathbb{R}_c^2} \mathbb{P}[z \in \text{ram}(\mathcal{A})] dz = \mathbb{E}[\int_{\mathbb{R}_c^2} \mathbb{1}\{z \in \text{ram}(\mathcal{A})\} dz] = 0$ and the result follows. \blacksquare

Lemma 5.4.2 *Let $A \subseteq \Pi^\dagger$ be non-crossing and pervasive and let $D \subseteq \mathbb{R}^2$ be dense. Let $f, h \in \mathcal{A}$ with $f \triangleleft h$ and suppose $t\star \in \mathbb{R}_s$ is such that $f(t\star) < h(t\star)$. Then for all $\epsilon > 0$ there exists $(x, s) \in D$ and $g \in A((x, s))$ such that $f, h \notin A((x, s))$, $f \triangleleft g$, $g \triangleleft h$ and $|t - s| < \epsilon$. Moreover, if $\star = -$ then we may take $s < t$ and if $\star = +$ then we may take $s > t$.*

PROOF: By the càdlàg property of f, h and denseness of D , there exists $(x, s) \in D$ such that $|s - t| < \epsilon$, with the desired sign for $t - s$ and with $f(s-) \vee f(s+) < x < h(s-) \wedge h(s+)$. Since \mathcal{A} is pervasive there exists $g \in A((x, s))$. It is immediate that $f, h \notin A((x, s))$, and that $\{f, g, h\}$ is non-crossing. By Lemmas 3.1.3 and 3.1.4 we have $f \triangleleft g$ and $g \triangleleft h$. \blacksquare

Lemma 5.4.2 is a technical lemma used in the proof of our next lemma. Recall that in Lemma 4.4.6 we showed that if $\mathcal{A} \in \mathcal{W}_{\text{det}}$ then $\text{ram}(\mathcal{A})$, the set of ramification points of \mathcal{A} , is a measurable and null subset of \mathbb{R}_c^2 . The following lemma is stated as a result for deterministic weaves, which avoids having to find a suitable state space for random null sets.

Lemma 5.4.3 *Let \mathcal{A} be a deterministic weave. Then \mathcal{A} , $\text{flow}(\mathcal{A})$ and $\text{web}(\mathcal{A})$ all have the same ramification points.*

PROOF: Let us write $\mathcal{W} = \text{web}(\mathcal{A})$ and $\mathcal{F} = \text{flow}(\mathcal{A})$. By Theorem 2.4.3, applied to the weave whose law is a point-mass at \mathcal{A} , we have $\mathcal{W} \preceq \mathcal{A} \preceq \mathcal{F}$. Note that if $\mathcal{B}, \mathcal{B}' \in \mathcal{W}_{\text{det}}$ with $\mathcal{B} \preceq \mathcal{B}'$ then for any $b \in \mathcal{B}$ there exists $b' \in \mathcal{B}'$ such that $b \subseteq b'$. It follows that $\text{ram}(\mathcal{W}) \subseteq \text{ram}(\mathcal{A}) \subseteq \text{ram}(\mathcal{F})$. With this in hand it remains only to show that for any deterministic weave \mathcal{A} we have

$$\text{ram}(\mathcal{F}) \subseteq \text{ram}(\mathcal{W}). \quad (5.6)$$

To this end, let $D \subseteq \mathbb{R}_c^2$ be dense and non-ramified. Suppose that $(x, t) \in \mathbb{R}_c^2$ is ramified in \mathcal{F} . We thus have bi-infinite $f, g \in \mathcal{F}(z)$ that are not comparable under \preceq . It follows that there exists $s\star \in \mathbb{R}_s$ such that $f(s\star) \neq g(s\star)$, and without loss of generality we may assume $f(s\star) < g(s\star)$.

- Consider first if $s\star \geq t+$. Take a sequence $z_n = (x_n, t_n) \in D$ such that $z_n \rightarrow z$ and $x_n \leq f(t_n-) \wedge f(t_n+)$, along with a sequence $w_n = (y_n, s_n) \in D$ such that $w_n \rightarrow z$ and $y_n \geq g(t_n-) \vee g(t_n+)$. It is straightforward to check that such sequences exist. Take $f_n \in \mathcal{F}(z_n)$ and $g_n \in \mathcal{F}(w_n)$. By compactness, passing to a subsequence, we may assume that $f_n \rightarrow f'$ and $g_n \rightarrow g'$ where $f', g' \in \mathcal{F}$. By Lemma A.2.2 we have $f', g' \in \mathcal{F}(z)$. Recall from Lemma 3.1.6 that $(\mathcal{F}, \triangleleft)$ is totally ordered. By Lemma 4.2.4 we have $f' \triangleleft f \triangleleft g \triangleleft g'$, which implies that $f'(s\star) < g'(s\star)$. By Lemma A.2.2 we have $f_n|_{z_n} \rightarrow f'|_z$ and $g_n|_{w_n} \rightarrow g'|_z$. Note that $f'|_z, g'|_z \in \mathcal{W}$ by (2.12), and both pass through z . Since $f'|_z(s\star) = f'(s\star) < g'(s\star) = g'|_z(s\star)$ they cannot be comparable under \preceq . Hence, in this case, we have that $z \in \text{ram}(\mathcal{W})$.
- Next, consider if $s\star \leq t-$. We may assume that $f(u\bullet) = g(u\bullet)$ for all $u\bullet \geq t+$ (or else, the case above applies). By Lemma 5.4.2 there exists $z \in D$ and $h \in \mathcal{A}(z)$ such that

$f, g \notin \mathcal{F}(z)$, $f \triangleleft h \triangleleft g$ and $\sigma_z < t$. We have $f(u\bullet) = g(u\bullet)$ for all $u\bullet \geq t+$, and $f \triangleleft h \triangleleft g$, which means $f(u\bullet) = h(u\bullet) = g(u\bullet)$ for all such $u\bullet$. Since $f(t-) \leq h(t-) \leq g(t-)$ we have $h \in \mathcal{A}(z)$. The properties of h given in Lemma 5.4.2 guarantee that h is not equal to f or g , so there exists some $v_1\bullet_1, v_2\bullet_2 \leq t-$ such that $f(v_1\bullet_1) < h(v_1\bullet_1)$ and $h(v_2\bullet_2) < g(v_2\bullet_2)$.

We apply Lemma 5.4.2 twice more, to (f, h) at $v_1\bullet_1$ and to (h, g) at $v_2\bullet_2$. We thus obtain (respectively) for $i = 1, 2$, $z_i \in D$ and $h_i \in \mathcal{A}(z_i)$ such that $f \triangleleft h_1 \triangleleft h \triangleleft h_2 \triangleleft g$, with $f, h \notin \mathcal{F}(z_1)$, $h, g \notin \mathcal{F}(z_2)$ and $\sigma_{z_1}, \sigma_{z_2} < t$. The same argument as above shows that $h_1, h_2 \in \mathcal{F}(z)$. It is clear that $h_1|_{z_1}$ and $h_2|_{z_1}$ are both elements of \mathcal{W} and both pass through z . To complete the proof, we will show that they are not comparable under \subseteq .

Suppose that $h_1|_{z_1} \subseteq h_2|_{z_2}$. Then $\sigma_{z_2} \leq \sigma_{z_1}$, and the fact that $h_1 \triangleleft h \triangleleft h_2$ implies that $h \in \mathcal{A}(z_1)$, which is a contradiction. Similarly we cannot have $h_2|_{z_2} \subseteq h_1|_{z_1}$, so in this case we also have $z \in \text{ram}(\mathcal{W})$.

This completes the proof of Lemma 5.4.3. ■

We are now ready to give the proof of Theorem 2.4.4. We prove the two parts of the theorem in turn, with part 2 first. Let \mathcal{A} and \mathcal{B} be weaves.

Part 1. Suppose that $\mathcal{A} \sim \mathcal{B}$. Theorem 2.4.3 gives that $\text{flow}(\mathcal{A}) \stackrel{d}{=} \text{flow}(\mathcal{B})$, which implies that there exists a coupling of \mathcal{A} and \mathcal{B} such that $\mathbb{P}[\text{flow}(\mathcal{A}) = \text{flow}(\mathcal{B})] = 1$. Let us write $\mathcal{F} = \text{flow}(\mathcal{A}) \stackrel{\text{a.s.}}{=} \text{flow}(\mathcal{B})$. Suppose that $z \subseteq \mathbb{R}_c^2$ is finite and almost surely non-ramified in both \mathcal{A} and \mathcal{B} . Lemma 5.4.3 gives that, almost surely, z is non-ramified in \mathcal{F} .

We seek to show that $\mathcal{A}|_z \stackrel{\text{a.s.}}{=} \mathcal{F}|_z$. Write $z = (z_1, \dots, z_m)$ and fix $i \leq m$. Let $f = \mathcal{A}|_{z_i}$. By Theorem 4.3.9 there exists $f' \in \mathcal{F}$ with $f \subseteq f'$. Therefore $f'|_{z_i} = f$. As z is almost surely non-ramified in \mathcal{F} , and $\mathcal{F} \subseteq \Pi^\uparrow$, in fact $\mathcal{F}(z_i) \stackrel{\text{a.s.}}{=} \{f'\}$, which implies that almost surely $f'|_{z_i} = \mathcal{F}|_{z_i} = f$. We thus have $\mathcal{A}|_{z_i} \stackrel{\text{a.s.}}{=} \mathcal{F}|_{z_i}$.

A symmetric argument shows that $\mathcal{B}|_{z_i} \stackrel{\text{a.s.}}{=} \mathcal{F}|_{z_i}$, so in fact $\mathcal{A}|_{z_i} \stackrel{\text{a.s.}}{=} \mathcal{B}|_{z_i}$. In particular $\mathcal{A}|_{z_i}$ and $\mathcal{B}|_{z_i}$ have the same marginal distribution, so $\mathcal{A}|_{z_i} \stackrel{d}{=} \mathcal{B}|_{z_i}$, as required.

Part 2. Let us first show that (a) and (b) are equivalent. Assume (a), that $\mathcal{A} \sim \mathcal{B}$. Theorem 2.4.3 gives that $\text{flow}(\mathcal{A}) \stackrel{d}{=} \text{flow}(\mathcal{B})$, which implies that there exists a coupling of \mathcal{A} and \mathcal{B} such that $\mathbb{P}[\text{flow}(\mathcal{A}) = \text{flow}(\mathcal{B})] = 1$. Let us write $\mathcal{F} = \text{flow}(\mathcal{A}) \stackrel{\text{a.s.}}{=} \text{flow}(\mathcal{B})$. From (2.13) we thus have $\mathbb{P}[\mathcal{A} \cup \mathcal{F} \text{ is non-crossing and } \mathcal{B} \cup \mathcal{F} \text{ is non-crossing}] = 1$. Lemma 4.1.3 gives that $\mathbb{P}[\mathcal{A} \cup \mathcal{B} \text{ is non-crossing}] = 1$, obtaining (b).

Conversely suppose (b), that \mathcal{A}, \mathcal{B} are coupled weaves such that $\mathbb{P}[\mathcal{A} \cup \mathcal{B} \text{ is non-crossing}] = 1$. By Lemma 4.4.4 we have $\mathbb{P}[\text{flow}(\mathcal{A}) = \text{flow}(\mathcal{B})] = 1$. In particular $\text{flow}(\mathcal{A})$ and $\text{flow}(\mathcal{B})$ have the same marginal distribution, so $\mathcal{A} \sim \mathcal{B}$.

We will next show that (a) and (c) are equivalent. Assume (a), that $\mathcal{A} \sim \mathcal{B}$. Lemma 5.4.1 implies the existence of a deterministic dense countable $D \subseteq \mathbb{R}^2$ such that D is almost surely non-ramified in both \mathcal{A} and \mathcal{B} . If $z \subseteq D$ is finite then it follows by part 1 of the present proof that $\mathcal{A}|_z$ and $\mathcal{B}|_z$ have the same marginal distribution, which establishes (c).

Conversely suppose (c). Note that the (deterministic) operation $\mathcal{A}' \mapsto \text{web}_D(\mathcal{A}')$ depends only on $\mathcal{A}'|_D$. Enumerate $D = (z_i)_{i \in \mathbb{N}}$ and write $z_m = (z_1, \dots, z_m)$. From (c) we have that $\mathcal{A}|_{z_m}$ and $\mathcal{B}|_{z_m}$ have the same marginal distribution, so by countability of D in fact $\mathcal{A}|_D$ and $\mathcal{B}|_D$ have the same marginal distribution. It follows that $\text{web}_D(\mathcal{A})$ and $\text{web}_D(\mathcal{B})$ have the same marginal distribution. By Lemma 4.5.3 we thus have that $\text{web}(\mathcal{A}) \stackrel{d}{=} \text{web}(\mathcal{B})$, so $\mathcal{A} \sim \mathcal{B}$, which establishes (a). This completes the proof.

5.5 Proof of Theorems 2.4.5 and 2.4.6

We give the proof of **Theorem 2.4.6** before that of 2.4.5, because part 2 of Theorem 2.4.5 will be proven as a specialization of part 2 of Theorem 2.4.6. Suppose that $\mathcal{A}_n, \mathcal{A}$ are weaves. We prove parts 1 and 2 of Theorem 2.4.6 in turn.

Part 1. Suppose that $\mathcal{A}_n \rightarrow \mathcal{A}$ and let $z_n, z \in (\mathbb{R}_c^2)^m$ be non-ramified. By Skorohod's Representation Theorem we may (change probability space, preserving the marginal distributions of each \mathcal{A}_n and \mathcal{A}) and assume that $\mathcal{A}_n \xrightarrow{\text{a.s.}} \mathcal{A}$. Let us write $z_n = (z_{n,1}, \dots, z_{n,m})$ and $z = (z_1, \dots, z_m)$. Due to non-ramification, for each i the sets $\mathcal{A}_n|_{z_{n,i}}$ and $\mathcal{A}|_{z_i}$ almost surely contain a single bi-infinite path, which we write as $f_{n,i}$ and (respectively) f_i . By Lemma A.1.1 the set $\bigcup_{n \in \mathbb{N}} \mathcal{A}_n$ is almost surely relatively compact. Therefore we may pass to a subsequence and assume that $f_{n,i} \xrightarrow{\text{a.s.}} g_i \in \Pi^\uparrow$ as $n \rightarrow \infty$, for all i . Since $\mathcal{A}_n \xrightarrow{\text{a.s.}} \mathcal{A}$ we have $g_i \in \mathcal{A}$. By Lemma A.2.2 we have $g_i \in \mathcal{A}(z_i)$, which implies that $g_i \stackrel{\text{a.s.}}{=} f_i$. Thus $\mathcal{A}_n|_{z_n} \xrightarrow{\text{a.s.}} \mathcal{A}|_z$, on the probability space generated by Skorohod's Representation Theorem, which implies convergence in distribution.

Part 2. Let \mathcal{B} be a weak limit point of (\mathcal{A}_n) , that is $\mathcal{A}_n \xrightarrow{\text{d}} \mathcal{B}$ along a subsequence of n . Let us pass to this subsequence, without loss of generality. Further, suppose that \mathcal{B} is almost surely non-crossing. By Skorohod's Representation Theorem, noting that we are interested to prove distributional properties of \mathcal{B} , without loss of generality we may assume that $\mathcal{A}_n \xrightarrow{\text{a.s.}} \mathcal{B}$. Since \mathcal{B} is assumed to be almost surely non-crossing, to show that \mathcal{B} is a weave we need only show that \mathcal{B} is almost surely pervasive. Let $z \in \mathbb{R}_c$. Almost surely, for all n there exists $f_n \in \mathcal{A}_n(z_n)$. By Lemma A.1.1 the set $\mathcal{B} \cup (\bigcup_{n=1}^\infty \mathcal{A}_n)$ is compact, hence there exists $f \in \Pi^\uparrow$ such that $f_n \rightarrow f$. As $\mathcal{A}_n \xrightarrow{\text{a.s.}} \mathcal{B}$ we have $f \in \mathcal{B}$. By Lemma A.2.2 we have $f \in \mathcal{B}(z)$. Thus \mathcal{B} is pervasive, so \mathcal{B} is a weave.

Suppose additionally that $\mathcal{A}_n|_z \rightarrow \mathcal{A}|_z$ for all almost surely non-ramified $z \in (\mathbb{R}^2)^m$. By Lemma 5.4.1 the set $R = \{z \in \mathbb{R}^2; \mathbb{P}[z \text{ is ramified in } \mathcal{A}, \mathcal{B} \text{ or } \mathcal{A}_n] > 0\}$ is a Lebesgue null subset of \mathbb{R}^2 . For any finite sequence z of points in $D = \mathbb{R}^2 \setminus R$ we have (by assumption) that $\mathcal{A}_n|_z \xrightarrow{\text{d}} \mathcal{A}|_z$. Since R is null, D is dense in \mathbb{R}^2 . From part 1 of the present theorem we have also that $\mathcal{A}_n|_z \xrightarrow{\text{a.s.}} \mathcal{B}|_z$, which implies convergence in distribution. Hence $\mathcal{A}|_z \stackrel{\text{d}}{=} \mathcal{B}|_z$ for all finite $z \subseteq D$. Theorem 2.4.4 now gives that $\mathcal{A} \sim \mathcal{B}$. This completes the proof of Theorem 2.4.6.

We now give the proof of **Theorem 2.4.5**, proving each of the three parts in turn. Suppose that $\mathcal{F}_n, \mathcal{F}$ are flows.

Part 1. Note that if $z \in \mathbb{R}_c$ is non-ramified then $\mathcal{F}(z)$ contains only a single bi-infinite path. With this fact in hand, the argument is essentially the same as that of part 1 of Theorem 2.4.6 (from the start of the present section) and is left to the reader.

Part 2. Suppose that \mathcal{F}' is a weak limit point of (\mathcal{F}_n) . Lemma 3.3.6 gives that \mathcal{F}' is almost surely non-crossing. Hence, by part 2 of Theorem 2.4.6, \mathcal{F}' is a weave. As Π^\uparrow is a closed subset of Π we have $\mathcal{F}' \subseteq \Pi^\uparrow$ almost surely, hence by Theorem 2.4.3 \mathcal{F}' is a flow.

Suppose, additionally, that (\mathcal{F}_n) is tight and $\mathcal{F}_n|_z \xrightarrow{\text{d}} \mathcal{F}|_z$ for all non-ramified $z \in (\mathbb{R}^2)^m$. Part 2 of Theorem 2.4.6 thus gives that $\mathcal{F} \sim \mathcal{F}'$. By Theorem 2.4.3, in particular by the fact that each equivalence class contains a unique maximal element, we have $\mathcal{F} \stackrel{\text{d}}{=} \mathcal{F}'$. We now have that (\mathcal{F}_n) is tight and any weak limit point of (\mathcal{F}_n) is equal in distribution to \mathcal{F} , so we have that $\mathcal{F}_n \rightarrow \mathcal{F}$.

Part 3. Suppose that $\mathcal{F}_n \xrightarrow{\text{d}} \mathcal{F}$ and that $\mathcal{A}_n \sim \mathcal{F}_n$. By Theorem 2.4.4, for each $n \in \mathbb{N}$ there exists a coupling of \mathcal{F}_n to \mathcal{A}_n such that $\mathcal{A}_n \cup \mathcal{F}_n$ is almost surely non-crossing. Hence $\mathcal{A}_n \subseteq (\mathcal{F}_n)_\uparrow$. It follows from Proposition A.2.1 that tightness of (\mathcal{F}_n) implies tightness of (\mathcal{A}_n) .

Let \mathcal{A}' be a weak limit point of (\mathcal{A}_n) . We thus have $(\mathcal{F}_n, \mathcal{A}_n) \xrightarrow{\text{d}} (\mathcal{F}, \mathcal{A}')$. We apply Skoro-

hod's Representation Theorem (to the sequence of pairs $(\mathcal{F}_n, \mathcal{A}_n)_{n \in \mathbb{N}}$) and may therefore assume without loss of generality that $\mathcal{F}_n \xrightarrow{\text{a.s.}} \mathcal{F}$ and $\mathcal{A}_n \xrightarrow{\text{a.s.}} \mathcal{A}'$. Note that this preserves the marginal distributions of $(\mathcal{F}_n, \mathcal{A}_n)$, for each n . We therefore have that $\mathcal{A}_n \subseteq (\mathcal{F}_n)_\uparrow$ almost surely.

If $f \in \mathcal{A}$ then (almost surely) there exists $f_n \in \mathcal{A}_n$ such that $f_n \rightarrow f$. Hence also there exists $g_n \in \mathcal{F}_n$ with $f_n \subseteq g_n$. Lemma A.1.1 gives that $\mathcal{F} \cup (\bigcup_{n \in \mathbb{N}} \mathcal{F}_n)$ is almost surely compact, so we may pass to a subsequence and assume that $g_n \rightarrow g$, where $g \in \mathcal{F}$. Lemma A.2.2 gives that $f \subseteq g$. Thus almost surely $\mathcal{A}' \subseteq \mathcal{F}_\uparrow$, which implies $\mathcal{A}' \cup \mathcal{F}$ is non-crossing. Lemma A.2.2 and almost sure pervasiveness of \mathcal{A}_n imply that \mathcal{A}' is almost surely pervasive. Thus \mathcal{A}' is a weave. By Theorem 2.4.4 we have that $\mathcal{A}' \sim \mathcal{F}$, as required. This completes the proof of Theorem 2.4.5.

5.6 Proof of Theorem 2.4.8

Let \mathcal{A} be a weave, let $\mathcal{W} = \text{web}(\mathcal{A})$ and $\mathcal{F} = \text{flow}(\mathcal{A})$. Let D be a countable, dense and almost surely non-ramified subset of \mathbb{R}_c^2 . By Theorem 2.4.3 the event

$$\{\mathcal{W} = \text{web}_D(\mathcal{A}), \mathcal{W} \preceq \mathcal{A} \preceq \mathcal{F}, \mathcal{A} \cup \mathcal{W} \cup \mathcal{F} \text{ is noncrossing, } D \text{ is non-ramified}\}$$

has probability one. Without loss of generality, for the remainder of Section 5.6 we condition on this event occurring. During the course of this proof we will apply several of our previous results in reverse time, to Π^\downarrow values objects rather than Π^\uparrow valued objects. To assist with this we will use the \cdot° operator, defined above the statement of Theorem 2.4.8, which represents rotation of space-time \mathbb{R}_c^2 about the origin by 180 degrees. In previous sections, for $f \in \Pi^\uparrow$ with $(x, t) \in H(f)$ we have written $f|_{(x,t)}$ for the restriction of f to $[t-, \infty+]$. Here, we need to extend this notation to allow for restriction both forwards and backwards in time. For $f \in \Pi$ we will write $f|_{(x,t)}$ for the restriction of f to $[\sigma_f-, t+]$ and $f|_{(x,t)}$ for the restriction of f to $[t-, \tau_f+]$. For $f \in \Pi^\uparrow$ we will avoid the notation $f|_z$ within this section, writing $f|_z$ instead.

We begin the proof of Theorem 2.4.8 by showing that

$$\{g \in \Pi^\downarrow; g \text{ does not cross } \mathcal{A} \text{ and } g \text{ begins in } D\} = \{f|_z \in \Pi^\downarrow; z \in D \text{ and } f \in \mathcal{F}(z)\} \quad (5.7)$$

To see (5.7), first consider if g does not cross \mathcal{A} and begins in D . By applying Theorem 4.3.9 (in reverse time) we obtain $f \in \Pi^\uparrow$ such that $g \subseteq f$ and f does not cross \mathcal{A} , which implies that $f \in \text{flow}(\mathcal{A}) = \mathcal{F}$. We have $g = f|_z$, so the left hand side of (5.7) is contained within the right hand side. For the reverse inclusion, consider if $z \in D$ and $f \in \mathcal{F}(z)$. Clearly $g = f|_z$ begins at $z \in D$ and does not cross \mathcal{F} , which implies g does not cross \mathcal{A} . We have thus established (5.7).

Let

$$\begin{aligned} \widehat{\mathcal{W}}_D &= \overline{\{g \in \Pi^\downarrow; g \text{ does not cross } \mathcal{A} \text{ and } g \text{ begins in } D\}}_\downarrow \\ \widehat{\mathcal{W}}'_D &= \overline{\{f|_z \in \Pi^\downarrow; z \in D \text{ and } f \in \mathcal{F}(z)\}}_\downarrow \end{aligned} \quad (5.8)$$

Note that (5.8) is (2.16), repeated here for convenience. By Proposition A.2.1, compactness of \mathcal{F} implies relative compactness of (the right hand side of) equation (5.7). With this in hand Lemma 4.5.1, applied in reverse time, gives that $\widehat{\mathcal{W}}_D = \widehat{\mathcal{W}}'_D$. From (2.12) we have that

$$\widehat{\mathcal{W}}'_D = (\text{web}_{D^\circ}(\mathcal{F}^\circ))^\circ \quad (5.9)$$

from which Theorem 2.4.3 gives that $\widehat{\mathcal{W}}'_D$ is both a dual web and (from Remark 4.5.4) does not depend on the choice of dense and almost surely non-ramified subset $D \subseteq \mathbb{R}_c^2$. From this point on let us write $\widehat{\mathcal{W}} = \widehat{\mathcal{W}}_D = \widehat{\mathcal{W}}'_D$. From (5.9) we have that $(\widehat{\mathcal{W}}'_D)^\circ = \text{web}_{D^\circ}(\mathcal{F}^\circ)$, which in words says that $(\widehat{\mathcal{W}}'_D)^\circ$ is the web associated to \mathcal{F}° . Theorem 2.4.3 implies that $(\widehat{\mathcal{W}}'_D)^\circ$ does

not cross \mathcal{F}° , thus $\widehat{\mathcal{W}}'_D$ does not cross \mathcal{F} , which by Lemma 4.1.3 implies it also does not cross \mathcal{A} . Therefore $(\mathcal{W}, \widehat{\mathcal{W}})$ is a double web.

The same argument that led to (5.7), but now used forwards in time, gives that $\{f \in \Pi^\uparrow; f \text{ does not cross } \mathcal{A} \text{ and } f \text{ begins in } D\} = \{f \upharpoonright_{z \in D} \text{ and } f \in \mathcal{F}(z)\}$. It follows from (2.12) and Theorem 2.4.3 that

$$\mathcal{W} = \overline{\{f \in \Pi^\uparrow; f \text{ does not cross } \mathcal{A} \text{ and } f \text{ begins in } D\}}. \quad (5.10)$$

We now turn our attention to (2.17). Let

$$\mathcal{F}' = \overline{\{g \hookrightarrow f \in \Pi^\uparrow; g \in \widehat{\mathcal{W}} \text{ ends and } f \in \mathcal{W} \text{ begins at the same point of } D\}}. \quad (5.11)$$

We must show that $\mathcal{F}' \stackrel{\text{a.s.}}{=} \mathcal{F}$. Let $h \in \mathcal{F}'$. Then there exists $z_n \in D$ such that $f_n \in \mathcal{W}$ begins and $g_n \in \widehat{\mathcal{W}}$ ends at z_n , and $h_n \rightarrow h$ where $h_n = (g_n)_{\hookrightarrow}(f_n)$. Using that z_n is non-ramified, let h'_n be the unique element of $\mathcal{F}(z_n)$. Note that f_n does not cross \mathcal{F} . Hence $((h'_n) \upharpoonright_{z_n})_{\hookrightarrow}(f_n)$ is a path passing through z_n that does not cross \mathcal{F} , which is therefore an element of \mathcal{F} , and therefore equal to h_n . Thus $f_n = (h'_n) \upharpoonright_{z_n}$. By a symmetrical argument, $g_n = (h'_n) \downharpoonright_{z_n}$ which implies that $h'_n = (g_n)_{\hookrightarrow}(f_n) = h_n$. Hence $h_n \in \mathcal{F}$, which implies that $h \in \mathcal{F}$.

To see the reverse inclusion, let $f \in \mathcal{F}$. By Lemma 4.2.8 we have $\mathcal{F} = \overline{\mathcal{F}(D)}$, hence there exists $z_n \in D$ and $h_n \in \mathcal{F}(z_n)$ such that $h_n \rightarrow f$. Let $f_n = (h_n) \upharpoonright_{z_n}$ and $g_n = (h_n) \downharpoonright_{z_n}$. We have that h_n does not cross \mathcal{F} , hence f_n and g_n do not cross \mathcal{A} . From (5.8) and (5.10) we have that $f_n \in \mathcal{W}$ and $g_n \in \widehat{\mathcal{W}}$. Hence $h \in \mathcal{F}'$. We thus have $\mathcal{F} = \mathcal{F}'$ which establishes (2.17).

It remains only to show the uniqueness claim. Let $\widehat{\mathcal{U}}$ be a dual web and suppose that $(\mathcal{W}, \widehat{\mathcal{U}})$ is a double web. Then $\widehat{\mathcal{U}}^\circ$ is a web that almost surely does not cross \mathcal{F}° , and the same is true of $\widehat{\mathcal{W}}^\circ$. It is straightforward to check that \mathcal{F}° is a Π^\uparrow valued random variable that inherits closedness, pervasiveness and the non-crossing property from \mathcal{F} . Proposition A.2.1 implies that relative compactness is also inherited through \cdot° , so \mathcal{F}° is a weave. Lemma 4.1.4 now implies that $\widehat{\mathcal{U}}^\circ \cup \widehat{\mathcal{W}}^\circ$ is almost surely non-crossing, from which Lemma 4.5.5 implies that $\text{web}(\widehat{\mathcal{U}}^\circ) \stackrel{\text{a.s.}}{=} \text{web}(\widehat{\mathcal{W}}^\circ)$. By Theorem 2.4.3 we thus have $\widehat{\mathcal{U}}^\circ \stackrel{\text{a.s.}}{=} \widehat{\mathcal{W}}^\circ$, which implies $\widehat{\mathcal{U}} \stackrel{\text{a.s.}}{=} \widehat{\mathcal{W}}$ as required.

5.7 Proof of Theorem 2.4.9

Recall that $\Pi_c^\uparrow, \Pi_c^\downarrow$ and Π_c^\updownarrow respectively denote the sets of continuous forwards half-infinite, backwards half-infinite, and bi-infinite càdlàg paths. The proof of Theorem 2.4.9 is based on the following lemma.

Lemma 5.7.1 *The following hold.*

1. Let $\mathcal{A} \subseteq \Pi_c^\uparrow$ be a deterministic weave. If $h \in \Pi_c^\updownarrow$ does not cross \mathcal{A} then $h \in \Pi_c^\updownarrow$.
2. Let $\mathcal{F} \subseteq \Pi_c^\updownarrow$ be a deterministic flow. If $f \in \Pi_c^\uparrow$ does not cross \mathcal{F} then $f \in \Pi_c^\uparrow$.

PROOF: We will prove each claim in turn, starting with the first. We argue by contradiction. Let $\mathcal{A} \subseteq \Pi_c$. Suppose that that $h \in \Pi_c^\updownarrow$ does not cross \mathcal{A} and that h is discontinuous at $t \in \mathbb{R}$. Without loss of generality (or consider space reflected about the origin) we may assume that $h(t-) < h(t+)$. By Lemma 4.1.1 there exists $f \in \mathcal{A}$ such that $f(t-) \leq h(t-)$ and $h \triangleleft f$. Hence $h(t+) \leq f(t+)$. We thus have $f(t-) \leq h(t-) < h(t+) \leq f(t+)$, which means that $f \in \mathcal{A}$ is discontinuous at t . This is a contradiction, which completes the proof.

It remains to establish the second claim. If $f \in \Pi^\uparrow$ does not cross \mathcal{F} then by Theorem 4.3.9 there exists $f' \in \Pi^\uparrow$ such that $f \subseteq f'$ and f' does not cross \mathcal{F} . From what we have already proved we have $f' \in \mathcal{F}$, thus $f' \in \Pi_c$ which implies $f \in \Pi_c$. \blacksquare

We now give the proof of Theorem 2.4.9. Note that it suffices to prove the results for deterministic weaves. We will prove the two claims in turn, starting with the first. Let \mathcal{A}, \mathcal{B} be deterministic weaves such that $\mathcal{A} \sim \mathcal{B}$. We seek to show that if $\mathcal{A} \subseteq \Pi_c^\uparrow$ then $\mathcal{B} \subseteq \Pi_c^\uparrow$. It then follows by symmetry that $\mathcal{A} \subseteq \Pi_c^\uparrow$ if and only if $\mathcal{B} \subseteq \Pi_c^\uparrow$.

By Theorem 2.4.4, without loss of generality we may assume that \mathcal{A} and \mathcal{B} are coupled such that $\mathcal{A} \cup \mathcal{B}$ is almost surely non-crossing. On that event, by Lemma 4.1.4, a path $f \in \Pi$ crosses \mathcal{A} if and only if it crosses \mathcal{B} . From (2.13) we thus obtain $\text{flow}(\mathcal{A}) \stackrel{\text{a.s.}}{=} \text{flow}(\mathcal{B})$, which we henceforth refer to as \mathcal{F} . By part 1 of Lemma 5.7.1, as $\mathcal{A} \subseteq \Pi_c^\uparrow$ we have also that $\mathcal{F} \subseteq \Pi_c^\uparrow$. Since $\mathcal{B} \preceq \mathcal{F}$ we have $\mathcal{B} \subseteq \mathcal{F}_\uparrow$, thus $\mathcal{B} \subseteq \Pi_c$. The same applies by symmetry with the roles of \mathcal{A} and \mathcal{B} swapped. This proves the first claim of Theorem 2.4.9.

It remains to prove the second claim. Let \mathcal{A} be a deterministic weave. Let \mathcal{W}, \mathcal{F} denote the corresponding web and flow. It remains to show that $\hat{\mathcal{W}} \subseteq \Pi_c^\uparrow$ if and only if $\mathcal{W} \subseteq \Pi_c^\downarrow$, where $\hat{\mathcal{W}}$ is given by (2.16) in Theorem 2.4.8. By symmetry (or consider reversing the direction of time) it suffices to prove that if $\mathcal{W} \subseteq \Pi_c^\downarrow$ then $\hat{\mathcal{W}} \subseteq \Pi_c^\uparrow$. To this end, suppose that $\mathcal{W} \subseteq \Pi_c$.

From what we have already proved, $\mathcal{F} \subseteq \Pi_c$. From (2.16) we have

$$\hat{\mathcal{W}} = \overline{\{g \in \Pi^\downarrow; g \text{ does not cross } \mathcal{F} \text{ and } (g(\tau_g), \tau_g) \in D\}}_\downarrow. \quad (5.12)$$

Let $\hat{f} \in \hat{\mathcal{W}}$. Then there exists $\hat{f}_n \in \Pi^\downarrow$ such that \hat{f}_n does not cross \mathcal{F} and $\hat{f}_n \rightarrow \hat{f}$. By Theorem 4.3.9 (applied in reverse time) there exists $h_n \in \mathcal{F}$ such that $\hat{f}_n \subseteq h_n$. By compactness of \mathcal{F} we may pass to a convergent subsequence $h_n \rightarrow h \in \mathcal{F}$. By Lemma A.2.2 we have $\hat{f} \subseteq h$. Since h does not cross \mathcal{F} , also \hat{f} does not cross \mathcal{F} . By part 2 of Lemma 5.7.1 (applied in reverse time) we thus have $\hat{f} \in \Pi_c^\uparrow$. Thus $\hat{\mathcal{W}} \subseteq \Pi_c^\uparrow$.

A Appendices

A.1 On the Hausdorff metric

Let (M, d_M) be a metric space and let $\mathcal{K}(M)$ denote the set of compact subsets of M , including the empty set. We write $\text{dist}_M(x, A) = \inf_{a \in A} d_M(x, a)$ for the infimum distance from the point $x \in M$ to $A \subseteq M$. We now state some well known facts relating to $\mathcal{K}(M)$. The function

$$d_{\mathcal{K}(M)}(A_1, A_2) := \sup_{x_1 \in A_1} \text{dist}_M(x_1, A_2) \vee \sup_{x_2 \in A_2} \text{dist}_M(x_2, A_1), \quad (\text{A.1})$$

defines a metric on $\mathcal{K}(M)$ known as the Hausdorff metric, or more precisely the Hausdorff metric with respect to d_M . If d and d' are two metrics generating the same topology on M , then their corresponding Hausdorff metrics generate the same topology on M . This topology is known as the Hausdorff topology. Note the subtle difference between ‘the Hausdorff topology’, which refers to this particular topology, and ‘a Hausdorff topology’, which refers to any topology having the Hausdorff property i.e. that distinct points possess disjoint neighbourhoods. (All topological spaces mentioned within the present article have the Hausdorff property.)

Completeness of M implies completeness of $\mathcal{K}(M)$. The same extension from M to $\mathcal{K}(M)$ also holds for separability, and for compactness. We now establish some more detailed connections.

Lemma A.1.1 *Let (M, d_M) be a complete metric space.*

1. Let $\mathcal{X} \subseteq \mathcal{K}(M)$. Then \mathcal{X} is relatively compact if and only if $\bigcup_{X \in \mathcal{X}} X$ is a relatively compact subset of M .
2. Let $X \subseteq M$. Then X is relatively compact if and only if $\{A \in \mathcal{K}(M); A \subseteq X\}$ is a relatively compact subset of $\mathcal{K}(M)$.

PROOF: The two claims are readily seen to be equivalent: take $\mathcal{X} = \{A \in \mathcal{K}(M); A \subseteq X\}$ to see that the (1) \Rightarrow (2) and take $X = \bigcup_{A \in \mathcal{X}} A$ to see that (2) \Rightarrow (1). We will give proof of (1). As completeness of M implies completeness of $\mathcal{K}(M)$, in both M and $\mathcal{K}(M)$ we have that relative compactness is equivalent to total boundedness.

Suppose that $\mathcal{X} \subseteq \mathcal{K}(M)$ is totally bounded. Then, for each $\epsilon > 0$ there is a finite set X_1, \dots, X_n of elements of $\mathcal{K}(M)$ such that, for any $X \in \mathcal{X}$ there is some X_i such that $d_{\mathcal{K}(M)}(X, X_i) < \epsilon$. Let $Y = \bigcup_{i=1}^n X_i$ and note that $\bigcup_{X \in \mathcal{X}} X \subseteq Y^{(\epsilon)}$. Since each X_i is compact in M , Y is also compact in M , and in particular Y is totally bounded. Hence also $\bigcup_{X \in \mathcal{X}} X$ is totally bounded.

Conversely, suppose that $\mathcal{X} \subseteq \mathcal{K}(M)$ is such that $\bigcup_{X \in \mathcal{X}} X \subseteq M$ is totally bounded. Let $\epsilon > 0$. There exists a finite set $\{x_1, \dots, x_n\} \subseteq M$ and a map $f : \bigcup_{X \in \mathcal{X}} X \rightarrow \{x_1, \dots, x_n\}$, such that for any $x \in \bigcup_{X \in \mathcal{X}} X$ we have $d_M(x, f(x)) < \epsilon$. For any $X \in \mathcal{X}$, the set $f(X) = \{f(x); x \in X\}$ is finite and therefore compact, meaning that $f(X) \in \mathcal{K}(M)$. By construction we have $d_{\mathcal{K}(M)}(X, f(X)) < \epsilon$. Let \mathcal{X}' be the set of subsets of $\{x_1, \dots, x_n\}$, hence \mathcal{X}' is a finite subset of $\mathcal{K}(M)$. We have shown that for any $X \in \mathcal{X}$ there is some $X' \in \mathcal{X}'$ such that $d_{\mathcal{K}(M)}(X, X') < \epsilon$. Thus \mathcal{X} is totally bounded. \blacksquare

The next Lemma is a key ingredient of the proof of Lemma A.3.1. It provides a supply of closed (and consequently measurable) subsets of $\mathcal{K}(M)$. We define

$$\begin{aligned} A^{(\epsilon)} &= \{x \in M; \text{dist}_M(x, A) < \epsilon\} && \text{for } \epsilon > 0, \\ A^{[\epsilon]} &= \{x \in M; \text{dist}_M(x, A) \leq \epsilon\} && \text{for } \epsilon \geq 0, \end{aligned}$$

as the (respectively) open and closed ϵ -expansions of $A \subseteq M$. Note that $A^{[0]} = A$.

Lemma A.1.2 *Let $X, Y \subseteq M$, $\epsilon \geq 0$ and $\mathcal{C} \subseteq \mathcal{K}(M)$. Then:*

1. the set $\{A \in \mathcal{K}(M); X \subseteq A^{[\epsilon]}\}$ is closed;
2. the set $\{A \in \mathcal{K}(M); A \cap (Y \setminus X) = \emptyset\}$ is closed if X is closed and Y is open;
3. the set $\{A \in \mathcal{K}(M); A \cap (X \setminus Y) \neq \emptyset\}$ is closed if X is closed and Y is open;
4. the set $\{A \in \mathcal{K}(M); \exists C \in \mathcal{C}, C \subseteq A\}$ is closed if \mathcal{C} is closed.

PROOF: We prove the claims independently. For the first claim, assume that $A_n \rightarrow A$ in $\mathcal{K}(M)$ and $X \subseteq A_n^{[\epsilon]}$ for all n , where $\epsilon \geq 0$. Let $f \in X$, so $d_{\mathcal{K}(M)}(\{f\}, A_n) \leq \epsilon$. Letting $n \rightarrow \infty$ we obtain $d_{\mathcal{K}(M)}(\{f\}, A) \leq \epsilon$. Since $f \in X$ was arbitrary $X \subseteq A^{[\epsilon]}$, as required.

For the second claim, assume that $A_n \rightarrow A$ in $\mathcal{K}(M)$ and $A_n \cap (Y \setminus X) = \emptyset$ for all n , where Y is open and X is closed. We must show that $A \cap (Y \setminus X)$ is empty. We will argue by contradiction. Suppose there exists $f \in A \cap (Y \setminus X)$, which means that $f \in A \cap Y$ and $f \notin X$. Since $A_n \rightarrow A$ there exists $f_n \in A_n$ such that $f_n \rightarrow f$. For each n , noting that $A_n \cap (Y \setminus X)$ is empty we have (a) $f_n \notin Y$ or (b) $f_n \in X$; thus one of these two alternatives must hold for an infinite subsequence of $n \in \mathbb{N}$. If (a) holds for infinitely many n then along that subsequence we have $f_n \rightarrow f$ with $f_n \notin Y$ and $f \in Y$, which contradicts the fact that Y is open. If (b) holds for

infinitely many n then along that subsequence we have $f_n \rightarrow f$ with $f_n \in X$ and $f \notin X$, which contradicts the fact that X is closed. We thus reach the desired contradiction.

For the third claim, suppose that $A_n \rightarrow A$ in $\mathcal{K}(M)$ and $A_n \cap (X \setminus Y) \neq \emptyset$ for all n , where $X \subseteq M$ is closed and $Y \subseteq M$ is open. Take $f_n \in A_n \cap (X \setminus Y)$, so $f_n \in A_n \cap X$ and $f_n \in M \setminus Y$. By Lemma A.1.1 the set $\cup_n A_n$ is a relatively compact M and contains the sequence (f_n) , so we may pass to a convergence subsequence $f_n \rightarrow f$. Since $A_n \rightarrow A$ we have $f \in A$. We have $f_n \in X$ and $f_n \in M \setminus Y$, and since both X and $M \setminus Y$ are closed we thus have $f \in X$ and $f \in M \setminus Y$. Thus $f \in A \cap (X \setminus Y)$, which completes the proof.

For the final claim, suppose that $A_n \rightarrow A$ in $\mathcal{K}(M)$ and that $C_n \in \mathcal{C}$ with $C_n \subseteq A_n$. Using Lemma A.1.1 the set $A \cup (\cup_n A_n)$ is compact, which implies that $\cup_n C_n$ is relatively compact, so we may pass to a subsequence and assume $C_n \rightarrow C \in \mathcal{K}(M)$. As $C_n \subseteq A_n$ we have $C \subseteq A$ and since \mathcal{C} is closed we have $C \in \mathcal{C}$. ■

Lemma A.1.3 *If $F : M \rightarrow M$ is continuous then the map from $\mathcal{K}(M)$ to itself given by $A \mapsto \{F(f); f \in A\}$ is continuous.*

PROOF: Note that if F is uniformly continuous then the conclusion is clear from the definition of the Hausdorff metric, see (A.1). For general continuous F , note by Lemma A.1.1 that if $A_n \rightarrow A$ in $\mathcal{K}(M)$ then the set $(\cup_n A_n) \cup A$ is compact and thus the restriction of F to this set is uniformly continuous. The result follows. ■

Lemma A.1.4 *Let $A_n, B_n, A, B \in \mathcal{K}(M)$ with $A_n \subseteq B_n$ for all n . If $A_n \rightarrow A$ and $B_n \rightarrow B$ then $A \subseteq B$.*

PROOF: Let $a \in A$. Since $A_n \rightarrow A$ there exists $a_n \in A_n$ such that $a_n \rightarrow a$. Thus $a_n \in B_n$. Since $B_n \rightarrow B$ we thus have $a \in B$. ■

A.2 On the M1 topology of Π and $\mathcal{K}(\Pi)$

The book of Whitt (2002) details relative compactness and weak convergence for real valued stochastic processes (i.e. single càdlàg paths) in all four Skorohod topologies. In Freeman and Swart (2023) we introduce a unified framework for these four topologies, suitable for random sets of càdlàg paths. We recall some properties of the M1 version of this framework here.

Compact subsets of Π play a key role in our main results and, for this reason, we require corresponding criteria for relative compactness (of subsets of Π) and tightness (relating to convergence in law of $\mathcal{K}(\Pi)$ valued random variables). Fix a metric $d_{\overline{\mathbb{R}}}$ generating the topology on $\overline{\mathbb{R}}$ and for $A \subseteq \overline{\mathbb{R}}$ let us write $\text{dist}_{\overline{\mathbb{R}}}(x, A) = \inf\{d_{\overline{\mathbb{R}}}(x, y); y \in A\}$. Relative compactness for sets of continuous paths is often characterised using the modulus of continuity, see for example Theorem 7.2 of Billingsley (1995). The analogous object for the M1 topology is

$$w_{T,\delta}(f) = \sup \left\{ \text{dist}_{\overline{\mathbb{R}}}(f(t_2 \star_2), [f(t_1 \star_1), f(t_3 \star_3)]) ; t_1 \star_1, t_2 \star_2, t_3 \star_3 \in I(f)_s, \right. \\ \left. -T \leq t_1 < t_2 < t_3 \leq T, t_3 - t_1 < \delta \right\}, \quad (\text{A.2})$$

with the conventions that the supremum over the empty set is zero, and $[a, b] = [a \wedge b, a \vee b]$.

Proposition A.2.1 *The following hold:*

1. A subset $\mathcal{A} \subseteq \Pi$ is relatively compact if and only if for all $0 < T < \infty$

$$\limsup_{\delta \rightarrow 0} \sup_{f \in \mathcal{A}} w_{T,\delta}(f) = 0.$$

2. A subset $\mathcal{A} \subseteq \mathcal{K}(\Pi)$ is relatively compact if and only if for all $0 < T < \infty$

$$\lim_{\delta \rightarrow 0} \sup_{\mathcal{A} \in \mathcal{A}} \sup_{f \in \mathcal{A}} w_{T,\delta}(f) = 0.$$

3. A sequence of Π valued random variables (f_n) is tight, in the sense that their laws comprise a relatively compact sequence of probability measures on Π , if and only if for all $0 < T < \infty$ and $\epsilon > 0$ we have

$$\lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} \mathbb{P}[w_{T,\delta}(f_n) \geq \epsilon] = 0.$$

4. A sequence of $\mathcal{K}(\Pi)$ valued random variables (\mathcal{A}_n) is tight, in the sense that their laws comprise a relatively compact sequence of probability measures on $\mathcal{K}(\Pi)$, if and only if for all $0 < T < \infty$ and $\epsilon > 0$ we have

$$\lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} \mathbb{P} \left[\sup_{f \in \mathcal{A}_n} w_{T,\delta}(f) \geq \epsilon \right] = 0. \quad (\text{A.3})$$

PROOF: Part 1 follows from the relative compactness criteria given in Theorem 3.7 of Freeman and Swart (2023). In the language of that theorem, compact containment is automatic as $\overline{\mathbb{R}}$ is compact, and the Skorohod-equicontinuity requirement thus becomes part 1 above. Note that, by Lemma A.1.1, parts 1 and 2 of Proposition A.2.1 are in fact equivalent to each other. Note also that part 3 follows from part 4 by taking $\mathcal{A}_n = \{f_n\}$. Therefore, to complete the proof of Proposition A.2.1 it suffices to deduce part 4 as a consequence of part 2. (In fact, a similar argument deduces part 3 from part 1.)

Suppose first that (\mathcal{A}_n) is tight. That is, for each $\kappa > 0$ there exists a compact set $\mathcal{B} \subseteq \mathcal{K}(\Pi)$ such that $\liminf_{n \rightarrow \infty} \mathbb{P}[\mathcal{A}_n \in \mathcal{B}] \geq 1 - \kappa$. By part 2, for any $\epsilon > 0$ and $T \in (0, \infty)$ there exists $\delta_0 > 0$ such that for all $\delta \in (0, \delta_0)$ we have $\sup_{B \in \mathcal{B}} \sup_{f \in B} w_{T,\delta}(f) \leq \epsilon$. Thus $\liminf_{n \rightarrow \infty} \mathbb{P}[\sup_{f \in \mathcal{A}_n} w_{T,\delta}(f) \leq \epsilon] \geq 1 - \kappa$.

It remains to show the reverse implication. Let (\mathcal{A}_n) satisfy (A.3) and let $\kappa > 0$. Note that $w_{T,\delta}(f)$ is an increasing function of both δ and T . By (A.3), for each $k \in \mathbb{N}$ set $\epsilon = 1/k$ and choose $\delta_k > 0$ and $T_k \in (0, \infty)$ such that

$$\limsup_{n \rightarrow \infty} \mathbb{P}[\mathcal{A}_n \setminus B_k \neq \emptyset] \leq \epsilon 2^{-k} \quad (\text{A.4})$$

where $B_k = \{f \in \Pi; w_{T_k,\delta_k}(f) \leq k^{-1}\}$. Note that $B_{k+1} \subseteq B_k$ and let $B = \bigcap_k B_k$. Then $\lim_{\delta \downarrow 0} \sup_{f \in B} w_{T,\delta}(f) = 0$, so part 2 gives that $\mathcal{B} = \{X \in \mathcal{K}(\Pi); X \subseteq B\}$ is a compact subset of $\mathcal{K}(\Pi)$. We have that

$$\begin{aligned} \limsup_n \mathbb{P}[\mathcal{A}_n \setminus B \neq \emptyset] &= \limsup_n \mathbb{P}[\bigcup_k (\mathcal{A}_n \setminus B_k) \neq \emptyset] \\ &\leq \limsup_n \sum_k \mathbb{P}[\mathcal{A}_n \setminus B_k \neq \emptyset] \\ &\leq \sum_k \limsup_n \mathbb{P}[\mathcal{A}_n \setminus B_k \neq \emptyset] \\ &\leq \epsilon. \end{aligned}$$

In the above, the third line uses the reverse Fatou lemma and the final line uses (A.4). Noting that $\{\mathcal{A}_n \setminus B = \emptyset\} = \{\mathcal{A}_n \subseteq B\} = \{\mathcal{A}_n \in \mathcal{B}\}$ we thus obtain $\limsup_n \mathbb{P}[\mathcal{A}_n \in \mathcal{B}] \geq 1 - \epsilon$. Thus (\mathcal{A}_n) is tight, which completes the proof. \blacksquare

The following three lemmas are consequences of Propositions 2.2.1 and (part 1 of) A.2.1. They cover most of our interaction with the M1 topology within the present article. Recall from Section 2.2 that for $f \in \Pi$ there is a natural total order \sqsubseteq on the interpolated graph $H(f)$. We write the associated strict order relation as \sqsubset . We slightly extend the terminology introduced in (2.8): if $w, z \in H(f)$ with $w \sqsubseteq z$ in the total order on $H(f)$, we write $f|_{[w,z]}$ for the unique $g \sqsubseteq f$ such that $g \in \Pi$ begins at w and ends at z .

Lemma A.2.2 *Let $f_n, f \in \Pi$. Suppose that $f_n \rightarrow f$, and that $w_n, z_n \in H(f_n)$ with $z_n \rightarrow z$ and $w_n \rightarrow w$, and suppose that $w_n \sqsubseteq z_n$ in the induced order from $H(f_n)$. Then $w, z \in H(f)$, with $w \sqsubseteq z$ in the induced order from $H(f)$, and $f_n|_{[w_n, z_n]} \rightarrow f|_{[w, z]}$.*

In particular, if $z_n \in H(f_n)$ with $f_n \rightarrow f$ and $z_n \rightarrow z$ then $z \in H(f)$.

PROOF: The second claim follows immediately from the first, so we will prove the first. Let $g_n = f_n|_{[w_n, z_n]}$ and $g = f|_{[w, z]}$. From the hypothesis of the lemma (f_n) is a relatively compact sequence in $\mathcal{K}(\Pi)$. Since $g_n \sqsubseteq f_n$ it follows from part 1 of Proposition A.2.1 that (g_n) is also relatively compact. To establish the present lemma it therefore suffices to show that any limit point of (g_n) is equal to g .

Let g' be a limit point of (g_n) and, with slight abuse of notation, let us pass to a subsequence and assume that $g_n \rightarrow g'$, in the M1 topology. From Proposition 2.2.1 we thus have $d_{\mathcal{K}((\mathbb{R}^2)^2)}(H^{(2)}(g_n), H^{(2)}(g')) \rightarrow 0$. As we noted in comments above Proposition 2.2.1, it is trivial to see that this implies $d_{\mathcal{K}(\mathbb{R}^2)}(H(g_n), H(g')) \rightarrow 0$. Moreover $(h, h') \mapsto d_{\mathcal{K}(\mathbb{R}^2)}(H(h), H(h'))$ is the Hausdorff metric on Π , generating a coarser topology than the M1 topology; as discussed in Section 3.4 of Freeman and Swart (2023) the topology generated is Skorohod's M2 topology. It follows immediately by uniqueness of limits that if $d_{\mathcal{K}(\mathbb{R}^2)}(H(g_n), H(g)) \rightarrow 0$ then $d_{\mathcal{K}((\mathbb{R}^2)^2)}(H^{(2)}(g_n), H^{(2)}(g)) \rightarrow 0$. Therefore, to establish the present lemma we need only show that $H(g') = H(g)$.

Consider $v \in H(g')$. As $H(g_n) \rightarrow H(g')$ there exists $v_n \in H(g_n)$ such that $v_n \rightarrow v$. Since $g_n \sqsubseteq f_n$ we thus have $v_n \in H(f_n)$ with $w_n \sqsubseteq v_n \sqsubseteq z_n$. Thus $(w_n, v_n), (v_n, z_n) \in H^{(2)}(f_n)$. Since $H^{(2)}(f_n) \rightarrow H^{(2)}(f)$ we thus obtain $(w, v), (v, z) \in H^{(2)}(f)$, which implies that $v \in H(f)$ with $w \sqsubseteq v \sqsubseteq z$. Hence $v \in H(g)$. We thus obtain $H(g') \subseteq H(g)$.

It remains to show the reverse inclusion. Consider $v \in H(g)$. Thus $v \in H(f)$ with $w \sqsubseteq v \sqsubseteq z$, which means that $(w, v), (v, z) \in H^{(2)}(f)$. Hence there exists $(w'_n, v_n), (v_n, z'_n) \in H^{(2)}(f_n)$ such that $(w'_n, v_n) \rightarrow (w, v)$ and $(v_n, z'_n) \rightarrow (v, z)$. If $v = w$ then $w_n \rightarrow v$ and as $w_n \in H(g_n)$ we obtain $v \in H(g')$. Similarly, if $v = z$ then $z \in H'(g)$. Without loss of generality we may therefore assume that $w \sqsubset v \sqsubset z$.

Note that $w_n \vee w'_n$ and $z_n \wedge z'_n$, where \wedge and \vee respectively denote min and max under \sqsubseteq , are elements of $H(f_n)$. Moreover $w_n \vee w'_n \rightarrow w$ and $z_n \wedge z'_n \rightarrow z$. If $v_n \sqsubseteq w_n \vee w'_n$ for infinitely many $n \in \mathbb{N}$ then $(v_n, w_n \vee w'_n) \in H^{(2)}(f_n)$ for such n and we would have $(v, w) \in H^{(2)}(f)$, meaning that $v \sqsubseteq w$, which is a contradiction. Hence $w_n \vee w'_n \sqsubseteq v_n$ for all but finitely many n . Similarly $v_n \sqsubseteq z_n \wedge z'_n$. Thus $w_n \sqsubseteq v_n \sqsubseteq z_n$ for all but finitely many n . For such n we have $v_n \in H(g_n)$, which implies that $v \in H(g)$. We therefore obtain $H(g') \subseteq H(g)$, so in fact $H(g') = H(g)$, which completes the proof. \blacksquare

Lemma A.2.3 *Let $f_n, f \in \Pi$. Suppose that $f_n \rightarrow f$ and $t_n \in I(f_n)$ with $t_n \rightarrow t \in \mathbb{R}$. For each n let $\star_n \in \{-, +\}$. Then $f \in I(f)$ and $f(t-) \wedge f(t+) \leq \liminf_n f_n(t_n \star_n) \leq \limsup_n f_n(t_n \star_n) \leq f(t-) \vee f(t+)$.*

PROOF: Let x be any subsequential limit point of $(f_n(t_n \star_n))_{n \in \mathbb{N}}$, which means that (x, t) is a subsequential limit point of $(f_n(t_n \star_n), t_n)$. Note that $(f_n(t_n \star_n), t_n) \in H(f_n)$. By Lemma A.2.2

we thus have $(x, t) \in H(f)$, which implies that $x \in [f(t-) \wedge f(t+), f(t-) \vee f(t+)]$. The result follows. \blacksquare

Lemma A.2.4 *Let $f_n, f \in \Pi$, with $f_n \rightarrow f$ and $t_\star \in I(f)_s$. There exists $t_n \in \mathbb{R}$ such that $t_n \rightarrow t$ and $f_n(t_n\star) \rightarrow f(t_\star)$.*

PROOF: Noting that $H(f_n) \rightarrow H(f)$, take $(x_n, t_n) \in H(f_n)$ such that $(x_n, t_n) \rightarrow (f(t_\star), t)$. Then $(f_n(t_n-), t_n) \sqsubseteq (x_n, t_n) \sqsubseteq (f_n(t_n+), t_n)$. By compactness of $\overline{\mathbb{R}}$ we may, without loss of generality, pass to a subsequence along which both $f_n(t_n-)$ and $f_n(t_n+)$ converge, say $f_n(t_n-) \rightarrow a$ and $f_n(t_n+) \rightarrow b$. Hence, by Lemma A.2.2 we have $(a, t) \in H(f)$ and $(b, t) \in H(f)$ with $(f(t-), t) \sqsubseteq (a, t) \sqsubseteq (f(t_\star), t) \sqsubseteq (b, t) \sqsubseteq (f(t+), t)$. If $\star = -$ then $(a, t) = (f(t_\star), t)$ and $f_n(t_n-) \rightarrow f(t-)$. If $\star = +$ then $(b, t) = (f(t_\star), t)$ and $f_n(t_n+) \rightarrow f(t+)$. \blacksquare

A.3 On measurability in $\mathcal{K}(\Pi)$

In this section we establish that various basic maps involving $\mathcal{K}(\Pi)$ are measurable. Such things are required to work with $\mathcal{K}(\Pi)$ valued random variables in Section 5. The vast majority of the work involved in this section is Lemma A.3.1, which is also used in the proof of Lemma 4.2.5.

Recall that we use the Borel σ -fields on \mathbb{R}_c^2 , Π and $\mathcal{K}(\Pi)$. These σ -fields are generated in each case by the closed (or equivalently, open) subsets. Due to our focus on compactness it is helpful to work with closed sets whenever possible. Recall that d_Π , generating the M1 topology on Π , is defined via Proposition 2.2.1 and the corresponding Hausdorff metric on $\mathcal{K}(\Pi)$ is defined via (A.1).

Lemma A.3.1 *The map $(A, z) \mapsto A(z)$ from $\mathcal{K}(\Pi) \times \mathbb{R}_c^2 \rightarrow \mathcal{K}(\Pi)$ is measurable.*

PROOF: The proof is rather technical. Recall that $A(z) = \{f \in A; z \in H(f)\}$. Note that Proposition 2.2.1 gives that $A(z)$ is a closed subset of Π , which implies compactness since $A(z)$ is a subset of the compact set A . As both $\mathcal{K}(\Pi)$ and \mathbb{R}_c^2 are separable, in order to establish measurability of $(A, z) \mapsto A(z)$ it suffices to show that the marginal maps $A \mapsto A(z)$ and $z \mapsto A(z)$ are both measurable. We split the proof into these two parts, which will be proved independently. In each part we will show that the pre-image of a closed subset of $\mathcal{K}(\Pi)$ is measurable; we will represent this pre-image explicitly using countably many set operations on measurable subsets. Lemma A.1.2 provides a supply of measurable (in fact, closed) subsets of $\mathcal{K}(\Pi)$. The fact that a closed graph $H(f)$ is measurable (in fact, compact) provides a supply of measurable subsets of \mathbb{R}_c^2 .

Measurability of $A \mapsto A(z)$: Fix $z \in \mathbb{R}_c^2$ and denote this map by $M_z : \mathcal{K}(\Pi) \rightarrow \mathcal{K}(\Pi)$, so $M_z(A) = A(z)$. For $z \in \mathbb{R}_c^2$ let us write $\Pi_z = \{f \in \Pi; z \in H(f)\}$ and $\mathcal{K}(\Pi_z)$ for the set of compact subsets of Π_z . Proposition 2.2.1 implies that Π_z is a closed subset of Π , from which part 1 of Lemma A.1.2 gives that $\mathcal{K}(\Pi_z)$ a closed subset of $\mathcal{K}(\Pi)$. Note that M_z maps into $\mathcal{K}(\Pi_z)$. It therefore suffices to show that the pre-image of a closed subset of $\mathcal{K}(\Pi_z)$ is measurable. To this end, let $\mathcal{C} \subseteq \mathcal{K}(\Pi_z)$ be closed and let \mathcal{C}' be a dense countable subset of \mathcal{C} . We will show that the following are equivalent:

1. $A(z) \in \mathcal{C}$;
2. for all $\epsilon > 0$ there exists $\delta > 0$ and $C \in \mathcal{C}'$ such that $(A \cap \Pi_z^{(\delta)}) \setminus C^{[\epsilon]} = \emptyset$ and $C \subseteq A^{[\epsilon]}$.

We first give the forwards implication (1) \Rightarrow (2). Let us assume $A(z) \in \mathcal{C}$ and suppose that (2) fails, in preparation for an argument by contradiction. Then there exists $\epsilon > 0$ such that for all

$C \in \mathcal{C}'$ and all $\delta > 0$ it holds that $(A \cap \Pi_z^{(\delta)}) \setminus C^{[\epsilon]} \neq \emptyset$ or $C \not\subseteq A^{[\epsilon]}$. We have $A(z) \in \mathcal{C}$ so we may choose $C \in \mathcal{C}'$ such that

$$d_{\mathcal{K}(\Pi)}(C, A(z)) \leq \frac{\epsilon}{2}. \quad (\text{A.5})$$

Hence $C \subseteq A(z)^{[\epsilon]}$, which implies $C \subseteq A^{[\epsilon]}$. Taking $\delta = \frac{1}{n}$ we thus have that there exists an infinite subsequence of $n \in \mathbb{N}$ for which $(A \cap \Pi_z^{(1/n)}) \setminus C^{[\epsilon]} \neq \emptyset$. For such n , take $f_n \in (A \cap \Pi_z^{(1/n)}) \setminus C^{[\epsilon]}$. We thus have $f_n \notin C^{[\epsilon]}$, which by (A.5) implies that $f_n \notin A(z)^{[\epsilon/2]}$. We also have that $f_n \in A \cap \Pi_z^{(1/n)}$ which, by compactness of A implies that (f_n) has a subsequential limit $f \in A(z)$. However, this contradicts the conclusion reached in the previous sentence. Thus (1) \Rightarrow (2).

Let us now establish the reverse implication (2) \Rightarrow (1). Take $\epsilon = \frac{1}{n}$ and take $\delta = \delta_n > 0$ and $C = C_n \in \mathcal{C}'$ as given from (2). That is, we have

$$C_n \subseteq A^{[1/n]} \quad \text{and} \quad A \cap \Pi_z^{(\delta_n)} \subseteq C_n^{[1/n]}. \quad (\text{A.6})$$

From the first statement in (A.6), since $C_n \in \mathcal{K}(\Pi_z)$ we have $C_n \subseteq A^{[1/n]} \cap \Pi_z \subseteq A(z)^{[1/n]}$. It is automatic that $A(z) \subseteq A \cap \Pi_z^{(\delta_n)}$ so from the second statement in (A.6) we obtain $A(z) \subseteq C_n^{[1/n]}$. Putting these together, we obtain $d_{\mathcal{K}(\Pi)}(A(z), C_n) \leq \frac{1}{n}$. Thus $C_n \rightarrow A(z)$ as $n \rightarrow \infty$. Since \mathcal{C} is closed we thus obtain $A(z) \in \mathcal{C}$. Thus (2) \Rightarrow (1).

We now have that (1) \Leftrightarrow (2). It follows that

$$\begin{aligned} M_z^{-1}(\mathcal{C}) &= \bigcap_{\epsilon > 0} \bigcup_{\delta > 0} \bigcup_{C \in \mathcal{C}'} \left\{ B \in \mathcal{K}(\Pi); (B \cap \Pi_z^{(\delta)}) \setminus C^{[\epsilon]} = \emptyset \right\} \cap \left\{ B \in \mathcal{K}(\Pi); C \subseteq B^{[\epsilon]} \right\} \\ &= \bigcap_{n \in \mathbb{N}} \bigcup_{m \in \mathbb{N}} \bigcup_{C \in \mathcal{C}'} \left\{ B \in \mathcal{K}(\Pi); B \cap (\Pi_z^{(1/m)} \setminus C^{[1/n]}) = \emptyset \right\} \cap \left\{ B \in \mathcal{K}(\Pi); C \subseteq B^{[1/n]} \right\} \end{aligned}$$

Note that we have used the set algebraic identity $X \cap (Y \setminus Z) = (X \cap Y) \setminus Z$. By parts 1 and 2 of Lemma A.1.2 the last line of the above consists of countable unions and intersections of closed subsets of $\mathcal{K}(\Pi)$. Thus M_z is measurable.

Measurability of $z \mapsto A(z)$: We will recycle some parts of our notation, to preserve the symmetry between this argument the above. Fix $A \in \mathcal{K}(\Pi)$ and let us denote the map in question by $M_A : \mathbb{R}_c^2 \rightarrow \mathcal{K}(\Pi)$, so $M_A(z) = A(z)$. For any $z \in \mathbb{R}_c^2$ we have that $M_A(z) \in \mathcal{K}(A)$, where $\mathcal{K}(A)$ denotes the set of all compact subsets of A . Part 1 of Lemma A.1.2 gives that $\mathcal{K}(A)$ is a closed subset of $\mathcal{K}(\Pi)$. It is straightforward to check that if $B \in \mathcal{K}(A)$ then

$$M_A^{-1}(B) = \left(\bigcap_{f \in B} H(f) \right) \setminus \left(\bigcup_{f \in A \setminus B} H(f) \right). \quad (\text{A.7})$$

In words, equation (A.7) says that $M_A(z) = B$ if and only if B is precisely the set of $f \in A$ such that $z \in H(f)$. In order to establish measurability of M_A it suffices to fix a closed subset \mathcal{C} of $\mathcal{K}(A)$ and show that $M_A^{-1}(\mathcal{C})$ is a measurable subset of \mathbb{R}_c^2 . Let \mathcal{C}' be a dense countable subset of \mathcal{C} . We will show that the following are equivalent:

3. $A(z) \in \mathcal{C}$;
4. for all $\epsilon > 0$ there exists $C \in \mathcal{C}'$ such that $z \in \left(\bigcap_{f \in C} H(f)^{[\epsilon]} \right) \setminus \left(\bigcup_{f \in A \setminus C^{[\epsilon]}} H(f) \right)$.

We first give the forwards implication (3) \Rightarrow (4). Suppose that $A(z) \in \mathcal{C}$, that is $M_A(z) \in \mathcal{C}$. Then there exists $B \in \mathcal{C}$ such that $z \in M_A^{-1}(B)$. Choose $C \in \mathcal{C}'$ such that $d_{\mathcal{K}(\Pi)}(B, C) \leq \epsilon$.

From (A.7) we have that $z \in \bigcap_{f \in B} H(f)$. It is straightforward to check that $\bigcap_{f \in B} H(f) \subseteq \bigcap_{f \in B^{[\epsilon]}} H(f)^{[\epsilon]}$ and as $C \subseteq B^{[\epsilon]}$ we obtain that $z \in \bigcap_{f \in C} H(f)^{[\epsilon]}$. Similarly, we have $B \subseteq C^{[\epsilon]}$, so $A \setminus C^{[\epsilon]} \subseteq A \setminus B$ and from (A.7) we have $z \notin \bigcup_{f \in A \setminus B} H(f) \supseteq \bigcup_{f \in A \setminus C^{[\epsilon]}} H(f)$. We have thus obtained that (3) \Rightarrow (4).

Let us now establish the reverse implication (4) \Rightarrow (3). Take $\epsilon = \frac{1}{n}$ and let $C = C_n$ be as given from (4). Noting that $C_n \subseteq A$ for all $n \in \mathbb{N}$, Proposition A.2.1 gives that the sequence $(C_n)_{n \in \mathbb{N}}$ is relatively compact (as sequence of elements of $\mathcal{K}(\Pi)$), and thus has a convergence subsequence. With slight abuse of notation let us pass to this convergent subsequence and set $B = \lim_n C_n$. It follows immediately that $B \in \mathcal{C}$.

From (4) we have $z \in \bigcap_{f \in C_n} H(f)^{[\epsilon_n]}$. For each $g \in B$ there exists $g_n \in C_n$ such that $g_n \rightarrow g$. As $g_n \in C_n$ we have $z \in H(g_n)^{[\epsilon_n]}$. Proposition 2.2.1 gives that $H(g_n) \rightarrow H(g)$ in $\mathcal{K}(\mathbb{R}_c^2)$, which implies that $H(g_n)^{[\epsilon_n]} \rightarrow H(g)$. It follows immediately that $z \in H(g)$, and as $g \in B$ was arbitrary we have $z \in \bigcap_{g \in B} H(g)$.

Similarly, from (4) we have $z \notin \bigcup_{f \in A \setminus C_n^{[\epsilon_n]}} H(f)$ for all n . Take $g \in A \setminus B$ and note that because $B \subseteq \Pi$ is closed we have $d_{\mathcal{K}(\Pi)}(\{g\}, B) > 0$. As $C_n \rightarrow B$ we have also that $C_n^{[\epsilon_n]} \rightarrow B$, so for sufficiently large n we have $g \notin C_n^{[\epsilon_n]}$. Hence $z \notin H(g)$. As $g \in A \setminus B$ was arbitrary we thus have $z \notin \bigcup_{g \in A \setminus B} H(g)$. Putting this together with the conclusion of the previous paragraph and (A.7) we obtain that $z \in M_A^{-1}(B)$. Thus (4) \Rightarrow (3).

We now have that (3) \Leftrightarrow (4). It follows that

$$\begin{aligned} M_A^{-1}(\mathcal{C}) &= \bigcap_{\epsilon > 0} \bigcup_{C \in \mathcal{C}'} \left(\left(\bigcap_{f \in C} H(f)^{[\epsilon]} \right) \setminus \left(\bigcup_{f \in A \setminus C^{[\epsilon]}} H(f) \right) \right) \\ &= \bigcap_{n \in \mathbb{N}} \bigcup_{C \in \mathcal{C}'} \left(\left(\bigcap_{f \in C} H(f)^{[1/n]} \right) \setminus \left(\bigcup_{f \in A \setminus C^{[1/n]}} H(f) \right) \right) \end{aligned} \quad (\text{A.8})$$

We now examine the two terms in brackets on the right hand side of (A.8). The set $\bigcap_{f \in C} H(f)^{[1/n]}$ is an intersection of closed sets, and is therefore closed. Note also that

$$\bigcup_{f \in A \setminus C^{[1/n]}} H(f) = \bigcup_{m \in \mathbb{N}} \left(\bigcup_{f \in A \setminus C^{(1/n+1/m)}} H(f) \right). \quad (\text{A.9})$$

We claim that $\bigcup_{f \in E} H(f)$ is closed whenever $E \subseteq \Pi$ is compact; to see this take $x_n \in H(f_n)$ where $f_n \in E$ and $x_n \rightarrow x \in \mathbb{R}_c^2$, pass to a convergence subsequence $f_n \rightarrow f \in E$ and then by Proposition 2.2.1 we have $H(f_n) \rightarrow H(f)$ so $x \in H(f)$. In particular, as $C^{(\delta)}$ is open for all $\delta > 0$ we have that $A \setminus C^{(\delta)}$ is compact, so the right hand side of (A.9) is a countable union of closed sets. We have now shown that (A.8) represents $M_A^{-1}(\mathcal{C})$ using countably many set operations of measurable subsets of \mathbb{R}_c^2 . Thus M_A is measurable. \blacksquare

Lemma A.3.2 *Let $z \in \mathbb{R}_c^2$. The map $(A, z) \mapsto A|_z$ is a measurable map from $\mathcal{K}(\Pi) \times \mathbb{R}_c^2 \rightarrow \mathcal{K}(\Pi)$.*

PROOF: Recall that $A|_z = \{f|_z; f \in A(z)\}$. As both $\mathcal{K}(\Pi)$ and \mathbb{R}_c^2 are separable, in order to establish measurability of $(A, z) \mapsto A|_z$ it suffices to show that the marginal maps $A \mapsto A|_z$ and $z \mapsto A|_z$ are both measurable. Recall $\Pi_z = \{f \in \Pi; z \in H(f)\}$ from the proof of Lemma A.3.1, in which we showed that $\Pi_z \subseteq \Pi$ and $\mathcal{K}(\Pi_z) \subseteq \mathcal{K}(\Pi)$ were both closed. By Lemma A.3.1 the

map $(A, z) \mapsto A(z)$ is measurable, and it clear that it maps into $\mathcal{K}(\Pi_z)$. Recall that for $f \in \Pi_z$, $f|_z$ is the unique $g \subseteq f$ such that $(g(\sigma_g^-), \sigma_g) = z$, and $A|_z = \{f|_z; f \in A \cap \Pi_z\}$. By Lemma A.2.2 the map $f \mapsto f|_z$ defined from Π_z to itself is continuous. It follows from Lemma A.1.3 that $A \mapsto A|_z$ is continuous on $\mathcal{K}(\Pi_z)$. Since $A|_z = (A(z))|_z$, this completes the proof. \blacksquare

Lemma A.3.3 *The map $A \mapsto A_\uparrow$ is a continuous map from $\mathcal{K}(\Pi)$ to itself.*

PROOF: Let $A_n, A \in \mathcal{K}(\Pi)$ with $A_n \rightarrow A$. Then (A_n) is relatively compact and, noting that $g \subseteq f \Rightarrow w_{T,\delta}(f) \leq w_{T,\delta}(g)$, part 2 of Proposition A.2.1 gives that $((A_n)_\uparrow)$ is relatively compact. Let $B \in \mathcal{K}(\Pi)$ be a limit point of $((A_n)_\uparrow)$. We must show that $B = A_\uparrow$. Without loss of generality let us pass to a subsequence and assume that $(A_n)_\uparrow \rightarrow B$.

Let $g \in B$. Then there exists $f_n \in A_n$ with $g_n \subseteq f_n$ and $g_n \rightarrow g$. By Lemma A.1.1 the set $\cup_n A_n$ is relatively compact, so we may pass to a further subsequence and assume that $f_n \rightarrow f$. Thus $f \in A$. Lemma A.2.2 gives that $g \subseteq f$, so $g \in A_\uparrow$. Hence $B \subseteq A_\uparrow$.

It remains to show the reverse inclusion. Let $g \in A_\uparrow$. Then there exists $f \in A$ with $g \subseteq f$. As $A_n \rightarrow A$ there exists $f_n \in A_n$ such that $f_n \rightarrow f$. From Proposition 2.2.1 (and the remarks just above it) we have $H(f_n) \rightarrow H(f)$. In particular, there exists $z_n \in H(f_n)$ such that $z_n \rightarrow (g(\sigma_g^-), \sigma_g) \in H(f)$. Let $g_n = f_n|_{z_n}$. Lemma A.2.2 gives that $g_n \rightarrow g$, so $g \in B$. Hence $A_\uparrow \subseteq B$. Thus $A_\uparrow = B$ and the proof is complete. \blacksquare

A.4 Proof of Lemma 2.5.1

We must show that the Brownian web \mathcal{W}_b satisfies our definition of a web, in Definition 2.4.1. The argument rests on well known properties of the Brownian web. We noted in Section 2.2 that (Π_c^\uparrow, d_Π) is the state space that is in common usage for the Brownian web. As \mathcal{W}_b is a $\mathcal{K}(\Pi_c^\uparrow)$ valued random variable, it is also a $\mathcal{K}(\Pi^\uparrow)$ valued random variable. We will now refer to points (a)-(c) of Theorem 2.3 of Schertzer et al. (2017), which defines \mathcal{W}_b .

We first show that \mathcal{W}_b is a weave. Let $D \subseteq \mathbb{R}^2$ be dense and countable. Point (a) of this definition gives that almost surely, $\mathcal{W}_b(z)$ is non-empty at all $z \in D$, which by Lemma A.2.2 implies that \mathcal{W}_b is almost surely pervasive. Moreover since $\mathcal{W}_b(z)$ is almost surely a singleton, D is almost surely non-ramified. It is well known in the literature that \mathcal{W}_b is almost surely non-crossing; strictly this follows because point (b) of the definition gives that $\mathcal{W}_b(D)$ is non-crossing, point (c) gives $\mathcal{W}_b \stackrel{\text{a.s.}}{=} \overline{\mathcal{W}_b(D)}$, and the non-crossing property is preserved by taking limits of continuous paths (this last implication uses Lemma 3.3.7). We have now shown that \mathcal{W}_b is a weave.

Point (c) of the definition gives that $\mathcal{W}_b \stackrel{\text{a.s.}}{=} \overline{\mathcal{W}_b(D)}$, for any deterministic dense countable $D \subseteq \mathbb{R}^2$. Noting our remarks immediately above the present lemma regarding $\mathcal{W}_b(D)$ and $\mathcal{W}_b|_D$, we have $\mathcal{W}_b \stackrel{\text{a.s.}}{=} \overline{(\mathcal{W}_b|_D)}$. Lemma A.2.2 implies that if $A_n \rightarrow A$ in $\mathcal{K}(\Pi)$ then $(A_n)_\uparrow \rightarrow A_\uparrow$. It is straightforward to combine this fact with the usual system of random walk approximations to the Brownian web (e.g. Figure 9 of Schertzer et al. (2017)) to show that $\mathcal{W}_b \stackrel{\text{a.s.}}{=} (\mathcal{W}_b)_\uparrow$. Hence in particular $\mathcal{W}_b|_D \stackrel{\text{a.s.}}{=} (\mathcal{W}_b|_D)_\uparrow$, from which (2.12) gives $\mathcal{W}_b \stackrel{\text{a.s.}}{=} \text{web}(\mathcal{W}_b)$. Theorem 2.4.3 thus gives that \mathcal{W}_b is a web.

References

- P. Alexandroff and P. Urysohn. Mémoire sur les espaces topologiques compacts, dédié à Monsieur D. Egoroff. *Verhandelingen Amsterdam*, 14(1), 1929.

- R. A. Arratia. *Coalescing Brownian motions on the line*. University of Wisconsin–Madison, 1979.
- N. Berestycki, C. Garban, and A. Sen. Coalescing Brownian flows: A new approach. *The Annals of Probability*, 43(6), 2015.
- P. Billingsley. *Probability and Measure*. Wiley, 1995.
- G. Cannizzaro and M. Hairer. The brownian web as a random \setminus -tree. 2021.
- A. Depperschmidt, A. Greven, and P. Pfaffelhuber. Marked metric measure spaces. *Electronic Communications in Probability*, 16(none), 2011.
- S. N. Ethier and T. G. Kurtz. *Markov processes: characterization and convergence*. Wiley, 1986.
- L. R. Fontes and C. M. Newman. The full Brownian web as scaling limit of stochastic flows. *Stochastics and Dynamics*, 6(02):213–228, 2006.
- L. R. G. Fontes, M. Isopi, C. M. Newman, and K. Ravishankar. The Brownian web: characterization and convergence. *The Annals of Probability*, 32(4):2857–2883, 2004.
- N. Freeman and J. Swart. Skorohod’s topologies on path space. *arXiv:2301.05637*, pages 1–47, 2023.
- A. N. Kolmogorov. On Skorohod convergence. *Theory of Probability and its Applications*, 1: 213–222, 1956.
- H. Kunita. *Stochastic flows and stochastic differential equations*, volume 24. Cambridge university press, 1997.
- T. M. Liggett. *Interacting particle systems*, volume 276. Springer, Berlin, 1985.
- T. Mountford, K. Ravishankar, and G. Valle. A construction of the stable web. *Latin American Journal of Probability and Mathematical Statistics*, 16(1):787, 2019.
- C. Newman, K. Ravishankar, and R. Sun. Convergence of coalescing nonsimple random walks to the Brownian web. *Electronic Journal of Probability*, 10, 2005.
- V. V. Piterbarg. Expansions and contractions of isotropic stochastic flows of homeomorphisms. *The Annals of Probability*, 26(2), 1998.
- E. Schertzer, R. Sun, and J. Swart. The Brownian web, the Brownian net, and their universality. *Advances in disordered systems, random processes and some applications*, pages 270–368, 2017.
- A. Skorohod. Limit theorems for stochastic processes. *Theory of Probability and its Applications*, 1:261–290, 1956.
- F. Soucaliuc, B. Tóth, and W. Werner. Reflection and coalescence between independent one-dimensional brownian paths. In *Annales de l’Institut Henri Poincaré (B) Probability and Statistics*, volume 36, pages 509–545, 2000.
- R. Sun and J. M. Swart. The Brownian net. *The Annals of Probability*, pages 1153–1208, 2008.
- B. Tóth and W. Werner. The true self-repelling motion. *Probability Theory and Related Fields*, 111(3):375–452, 1998.
- W. Whitt. *Stochastic-Process Limits*. Springer New York, 2002.