

**Minerva Lectures at Columbia University,
February – March 2019:
Universality in models of local random growth
via Gibbs resamplings**

Alan Hammond

A. HAMMOND, DEPARTMENTS OF MATHEMATICS AND STATISTICS, U.C. BERKELEY, 899 EVANS HALL, BERKELEY, CA, 94720-3840, U.S.A.

E-mail address: alanmh@berkeley.edu

ABSTRACT. An important technique for understanding a random system is to find a higher dimensional random system that enjoys an attractive and tractable structure and that has the system of interest as a marginal; and to analyse the new structure to make inferences about the original system. For example, the Airy_2 process is an important and natural random process, mapping the real line to itself, since it offers, rigorously in certain examples and putatively in very many more, a scaled description at advanced time of a random interface whose growth is stimulated by local randomness and which is subject to restoring forces such as surface tension. The Airy_2 process may be embedded in a canonical way as the uppermost curve in a richer random object, the Airy line ensemble - an ordered system of random continuous curves. This richer object has an attractive probabilistic property not apparent in the Airy_2 process itself - it is, with suitable boundary conditions, an infinite system of mutually avoiding Brownian motions; and, as such, it enjoys a natural resampling probability called the Brownian Gibbs property. The Brownian Gibbs property of the Airy line ensemble is a key probabilistic technique by which aspects of the concerned Kardar-Parisi-Zhang universality class of random growth models may be investigated. This short series of lectures will explain how, harnessed with limited but essential inputs of integrable origin, the property has been exploited in the recent work [Ham17] to make very strong inferences regarding the locally Brownian nature of the Airy_2 process; about the scaled coalescence behaviour of geodesics in last passage percolation growth models; and about the structure of the scaled interface when these models are initiated from very general initial conditions.

Chapter 1. Lecture One: Exponents, curvature and Gaussianity via resampling	3
Chapter 2. Lecture Two: the Gibbs property of Brownian LPP, and the main theorems	15
2.1. Brownian last passage percolation [LPP]	15
2.2. Mutually avoiding staircase collections – and ensembles of energy profiles	16
2.3. The Brownian Gibbs property	19
2.4. Stage One: Brownian Gibbs basics	26
Chapter 3. Lecture Three: Local fluctuation results via the Wiener candidate	30
3.1. Reprising this for real – with a regular ensemble	31
Chapter 4. Lecture Four: proving Brownian bridge regularity via the jump ensemble	38
4.1. More promising than the Wiener candidate – the jump ensemble	41
4.2. Proving Brownian bridge regularity via the realized promise of the jump curve	42
4.3. Sketch of proof of Proposition 4.4	46
Bibliography	51
Contents	

CHAPTER 1

Lecture One: Exponents, curvature and Gaussianity via resampling

Important stochastic processes – such as Brownian motion or the Airy_2 process – describe the scaled, universal, behaviour of physical systems.

Sometimes, these processes can profitably be analysed by embedding them in higher dimensional random systems in such a way that the bigger system has an attractive probabilistic property which may act as a tool, being brought to bear in the rigorous analysis of the original object of interest. This probabilistic property may be a *Gibbs resampling* rule – a rule which permits a bounded part of the larger system to be resampled, given the form of the remaining part, according to an explicit conditional distribution. The latter distribution may permit analysis to be made when this given part of the larger system is resampled, with insight so gathered about the big system than being transmitted to the original process which is embedded within it.

For example, the Airy_2 process \mathcal{A} is a random continuous function $\mathcal{A} : \mathbb{R} \rightarrow \mathbb{R}$ which may be embedded as the lowest indexed, and uppermost, curve $\mathcal{A}(1, \cdot)$ in the *Airy line ensemble* [AiryLE], which is a random \mathbb{N} -indexed collection $\mathcal{A} : \mathbb{N} \times \mathbb{R} \rightarrow \mathbb{R}$ of continuous curves. With suitable interpretation that copes with boundary conditions for an infinite system of curves, the AiryLE is a system of Brownian motions conditioned on mutual avoidance – see Figure 1.1. This has the implication that, if the top few curves are resampled on a given compact real interval, then their conditional distribution is given by a finite system of Brownian bridges, attached suitably to the endpoints generated by the curves' removal, conditioned on mutual avoidance and on avoidance of the lower boundary curve. This *Brownian Gibbs* property will be precisely specified later, and will form a central element in this exposition.

For now, we give two simple examples of random systems which demonstrate four characteristics of the Kardar-Parisi-Zhang [KPZ] universality class, namely:

- an exponent of two-thirds for the spatial scale;
- an exponent of one-third for the scale of height;
- and interfaces that are *locally Brownian ...*
- as well as *globally parabolic*.

The second of the two examples also illustrates how resampling a random system can be a useful tool for understanding it.

Example 1: a baby version of the Airy_2 process. Let $B : \mathbb{R} \rightarrow \mathbb{R}$, $B(0) = 0$, be standard two-sided Brownian motion. Let $t > 0$ be large. Set $X : \mathbb{R} \rightarrow \mathbb{R}$, $X(t) = B(x) - x^2/t$.

Further consider $\mathcal{C} = \mathcal{C}_X : \mathbb{R} \rightarrow \mathbb{R}$ to be the least concave majorant of $X : \mathbb{R} \rightarrow \mathbb{R}$.

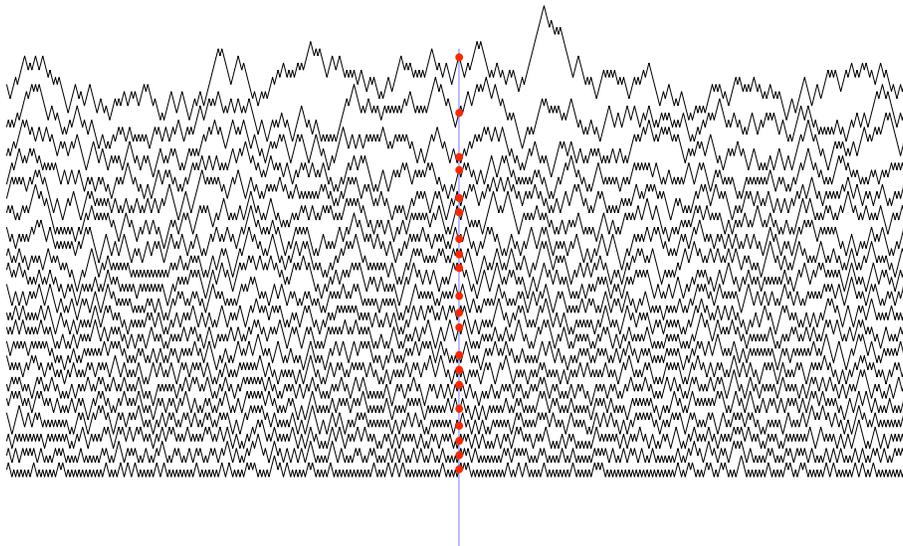


FIGURE 1.1. A discrete variant of the Airy line ensemble whose domain is cyclic. When the mesh that specifies the depicted walkers tends to zero and the Airy line ensemble is obtained, the indicated point process will converge towards the process of random statistics of the largest eigenvalues of a large matrix drawn from the Gaussian unitary ensemble. Simulation by Judit Zádor.

Naturally, the process X is locally Brownian but globally parabolic. In fact, it also evinces the characteristic KPZ exponents of $1/3$ and $2/3$. We may identify these exponents by studying how closely X follows the global contour offered by \mathcal{C}_X . The process \mathcal{C}_X has a graph, a closed subset of \mathbb{R}^2 , which is comprised of the union of a countable collection of closed planar line segments. The *facet through zero* may be defined to be the planar segment among these that passes through the vertical axis through the origin. The *facet length* is the random length of the interval occupied by x -coordinates of points in the facet through zero.

The facet length's scaling in t as $t \nearrow \infty$ is a measure of persistence of randomness in the horizontal scale. To find a counterpart in the vertical direction, we may consider the *inward deviation* or *local roughness* of the random interface X at the origin. Indeed, we may define the local roughness (at zero) to be the distance between $X(0)$ and $\mathcal{C}_X(0)$ – so that local roughness measures the fluctuation of the interface away from its convex hull at a typical location. Figure 1.2 illustrates this.

The claims that validate the exponent values for this model are that

in mean value: facet length scales as $t^{2/3}$ as $t \nearrow \infty$

and

in mean value: local roughness scales as $t^{1/3}$ as $t \nearrow \infty$.

To gesture towards the proofs of these claims, the facet length ℓ may be identified as the horizontal scale

above which *parabolic curvature* is dominant

and

below which *Brownian fluctuation* dominates.

The scale of ℓ may be identified by equating the two effects:

Brownian fluctuation \approx parabolic curvature;

or

$$\ell^{1/2} \approx \frac{\ell^2}{T}.$$

Which leads to $\ell \approx T^{2/3}$.

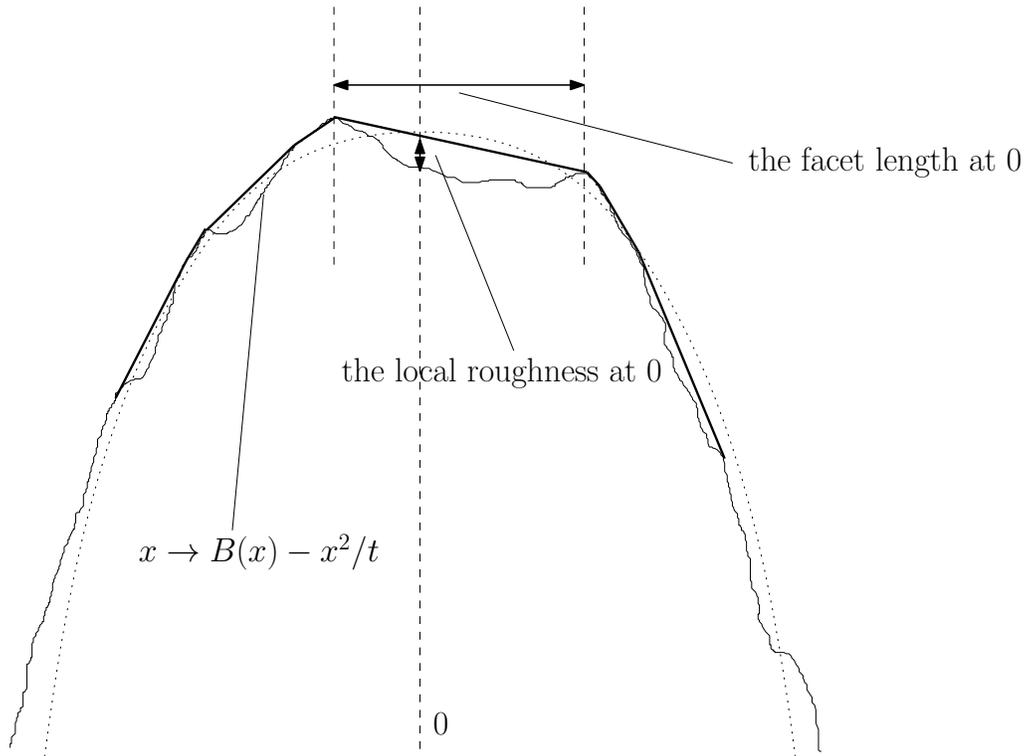


FIGURE 1.2. A toy model for the Airy_2 process: Brownian motion with a parabola subtracted whose coefficient t^{-1} is small and positive.

Regarding local roughness, the inward deviation made by X as it describes the journey from one extreme point of the graph of its convex hull to the next may be expected, in view of Gaussian fluctuation, to have scale given by the square-root of the horizontal extent of that journey. Which is to say: the local roughness will behave in the form $(t^{2/3})^{1/2} = t^{1/3}$.

Example 2: oriented random walk constrained by area trap.

Consider a state space Λ that consists of oriented (meaning downright) paths in the first quadrant of \mathbb{Z}^2 that begin on the y -axis and that end on the x -axis.

Let $\gamma \in \Lambda$. Then the length $L(\gamma)$ of γ is defined to be the number of steps taken by γ , which is also the sum of the vertical coordinate at which γ begins and the horizontal coordinate at which it ends.

The area $A(\gamma)$ is the area of the finite region enclosed by the graph of γ and the two coordinate axes – see Figure 1.3.

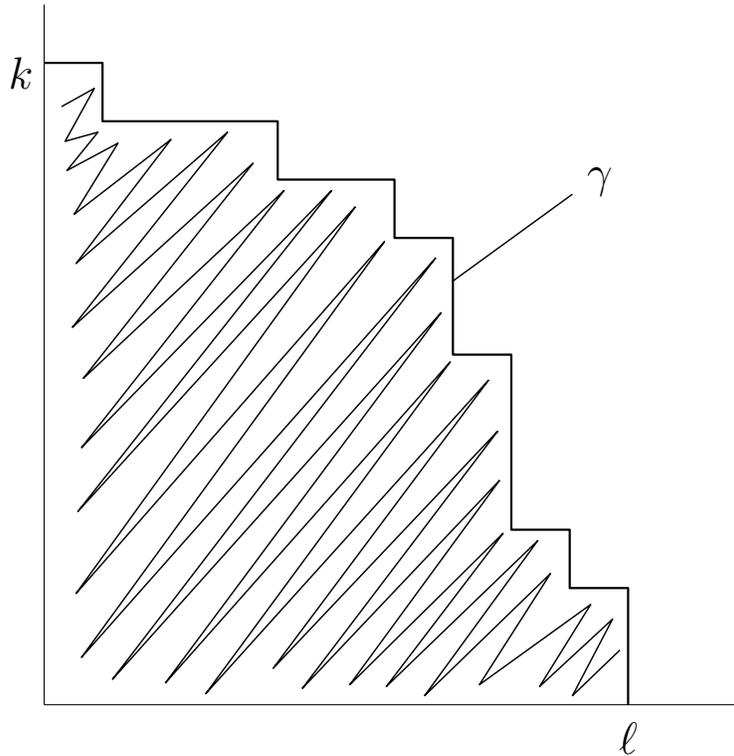


FIGURE 1.3. An element $\gamma \in \Lambda$ with $L(\gamma) = k + l$ and whose area $A(\gamma)$ is the area of the hatched region.

If we set

$$\Lambda_n = \{\gamma \in \Lambda : L(\gamma) = n\},$$

then $|\Lambda_n| = 2^n$. This is because the number of length n downright paths that begin at the origin is equal to 2^n ; and each such path may be translated directly upwards to a unique location at which it forms an element in Λ_n .

Now let $\lambda \in (0, 1/2)$. We define a probability measure \mathbb{P}_λ on Λ by insisting that

$$\mathbb{P}_\lambda(\gamma) = Z^{-1} \lambda^{L(\gamma)},$$

where $Z \in (0, \infty)$ is a normalization that ensures that \mathbb{P}_λ has unit mass.

Since $\lambda < 1/2$, Z is finite.

What is the distribution of length of the path obtained by sampling the law \mathbb{P}_λ ? It verifies

$$\mathbb{P}_\lambda(\Gamma \text{ has length } n) = Z^{-1}(2\lambda)^n,$$

where we write Γ for the sampled path. This equality is a condition of exponential tail decay because $2\lambda < 1$.

The law \mathbb{P}_λ may be viewed as a background or reference measure on which we now impose a constraint that involves area capture.

To do so, let $N \in \mathbb{N}$. The area-constrained random model that we study in this example is

$$\mathbb{P}_\lambda^{N^2} := \mathbb{P}_\lambda(\cdot \mid A(\Gamma) \geq N^2).$$

See Figure 1.4.

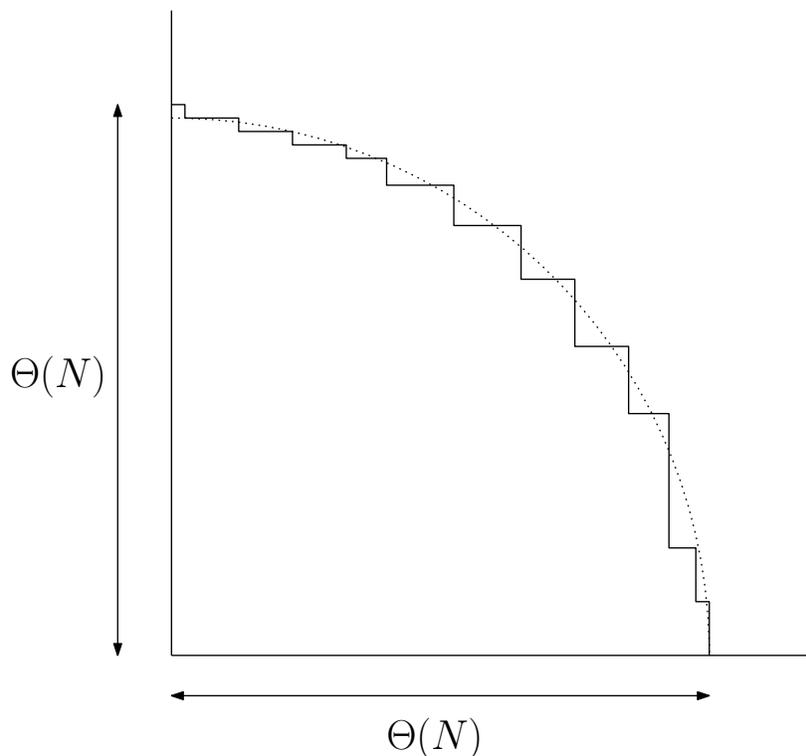


FIGURE 1.4. A caricature of a typical sample of $\mathbb{P}_\lambda^{N^2}$.

Here are two natural questions, which are not quite at the heart of our brief study of this model, but rather will aid the study: Under $\mathbb{P}_\lambda^{N^2}$, what is the typical value of

- the length of the sampled path Γ ;
- and the excess area of Γ , namely $\text{ExcessArea}(\Gamma) := A(\Gamma) - N^2$?

The answers are:

- that $L(\Gamma)$ typically has order N ;

- and that $\text{ExcessArea}(\Gamma)$ typically also has order N ; or more precisely,

$$\mathbb{P}_\lambda^{N^2}(\text{ExcessArea}(\Gamma) \geq tN) \leq C \exp\{-ct\}, \quad (1)$$

where C and c are two positive constants, and $t > 0$ is supposed to satisfy $t \geq C \log N$.

Why are these the correct answers?

Regarding the first, note that the event $A(\Gamma) \geq N^2$ on which we condition in specifying the law $\mathbb{P}_\lambda^{N^2}$ is typically satisfied by randomly chosen elements in Γ_n , provided that the index n is a large constant multiple of N . Since the random length of a sample of \mathbb{P}_λ has an exponential tail, we see that the answer to first query is in its scale the correct one.

To see that $\text{ExcessArea}(\Gamma) \asymp N$, note that, given a typical sample of $\mathbb{P}_\lambda^{N^2}$, we may perform an operation in which the sampled path Γ is shifted one unit to the right, with a horizontal edge added to abut its starting point, in order that a new element of the state space Γ be obtained. A gain in area of order N typically results from this manoeuvre – see Figure 1.5(left). But the new path and the old one have the same probability up to a constant factor to be sampled under \mathbb{P}_λ , because only one edge was added in the formation of the new path. In this way, the excess area is seen to have positive probability under the conditioned law $\mathbb{P}_\lambda^{N^2}$ to be of order at most N . It is the opposing bound, and its more precise form displayed above, which will interest us more. But this form may be demonstrated, roughly speaking, by a reversal of the argument just given – Figure 1.5(right) illustrates. If the excess of area is at least Nt , we may sample a path that realizes this excess, shift it to the left by an order of t units, and cut off the part that was pushed into the second quadrant; and we will be left with a new path, which continues to realize the basic requirement that the captured area be at least N^2 – but which does so with an order of t fewer edges than in the originally sampled path. The new path is thus preferred to the old in the underlying law \mathbb{P}_λ by a factor which grows exponentially in t . From this, then, we see that the part of the state space in which the excess area of the sampled path is at least tN must be exponentially small in t under the law $\mathbb{P}_\lambda^{N^2}$ – which is the stated estimate in the second answer.

These answers will be valuable as we derive the next result – which asserts that the powers of two-thirds and one-third are present in the constrained area-trap model in a manner similar to in our first example. We define \mathcal{C}_Γ to be the least concave majorant of the graph of Γ . See Figure 1.6 for the specifications of facet length and local roughness.

THEOREM 1.1. *Under $\mathbb{P}_\lambda^{N^2}$,*

the mean facet length scales as $N^{2/3}$

and

the mean local roughness scales as $N^{1/3}$.

We will prove this result – or explain the substantive elements of the proof – using a tool of *resampling*. We will find a random surgical procedure which takes as input a sample Γ of $\mathbb{P}_\lambda^{N^2}$, modifies a certain part of the sampled path, and returns as output another element of the state space Λ . The returned element Γ^{re} (the resampled path) will verify $A(\Gamma^{\text{re}}) \geq N^2$. Moreover, and crucially, the random surgery will be such that, if the input is chosen to have law $\mathbb{P}_\lambda^{N^2}$, then so will the output – so that the procedure may indeed be viewed as a resampling of this law, with the input path being distributed according to the law and the output path also doing so.

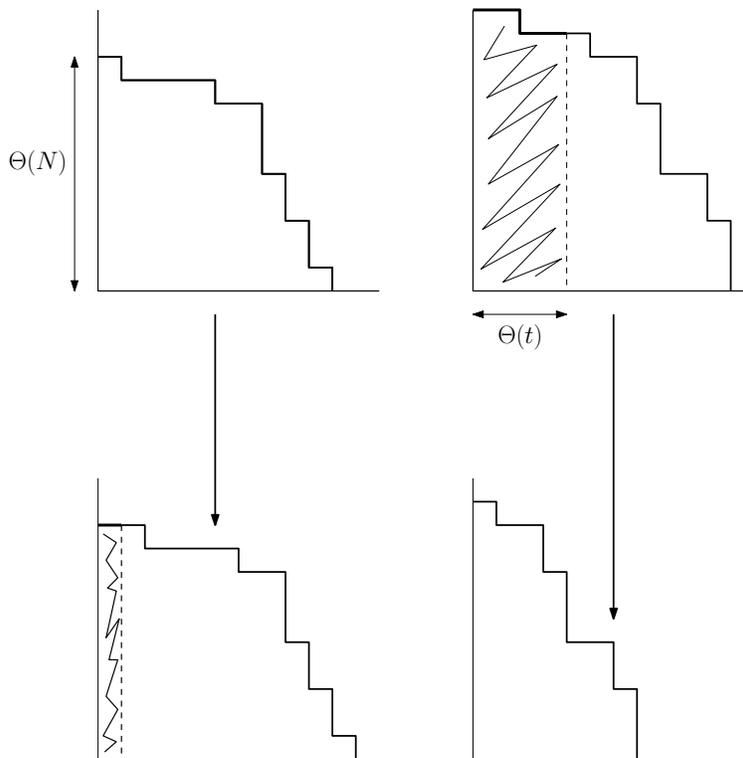


FIGURE 1.5. *Left:* An area gain of order N is attained by a unit rightward displacement, with the outcome as likely as the input up to a unit-order factor. *Right:* An area loss of order Nt alongside a length saving of order t .

To prepare for the construction of the random surgery, we first note a basic resampling property enjoyed by the underlying path measure \mathbb{P}_λ .

Take a sample Γ of \mathbb{P}_λ . Choose two points a and b on Γ . These points are chosen to be at certain given deterministic distances along Γ from its starting location on the y -axis – but in fact they may also be chosen according to a procedure that scans for the more counterclockwise of the points in a sweep centred at the origin that proceeds clockwise from vertical; and for the other point, by such a sweep proceeding counterclockwise from horizontal. The reason for this restriction on the selection of the pair of points is that it ensures that the subpath of Γ between them is not investigated during the selection, so that its randomness remains undisturbed.

The points a and b thus selected, the path Γ may be viewed as a concatenation of three pieces: a first piece, abutting the y -axis; a second, from a to b ; and a third, leading from b to the x -axis.

Remove Γ 's middle section, between a and b . Then ask:

what is the conditional distribution of the removed piece, given the form of what remains?

The answer is simple enough: this piece-to-be-added has the distribution of a uniformly chosen downright path from a to b .

We may pose a similar question about resampling a given section of Γ under $\mathbb{P}_\lambda^{N^2}$.

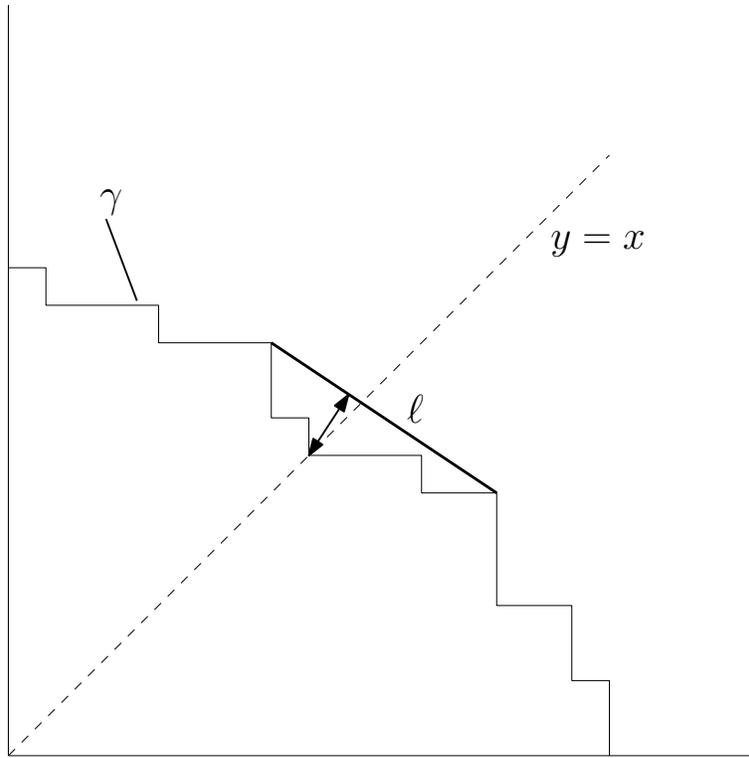


FIGURE 1.6. *Left:* The facet length (of the facet through the diagonal) of the down-right path γ is the length of the planar line segment ℓ . The local roughness (of the vertex in γ lying on the diagonal) is the length of the arrow.

First, select two points a and b on Γ , with b more clockwise than a – subject to the same constraints on the search as before. Again remove Γ between the two selected points. And ask the question posed above.

The answer is that the piece-to-be-added has the uniform law on downright paths from a to b , *subject to the condition that the resulting overall path traps an area that is at least N^2* – see Figure 1.7. This area constraint may be expressed in terms of the added path from a to b as the condition that the area trapped between the proposed path and the lower and left sides of the rectangle with northwest corner at a and with southeast corner at b is at least a *certain level* – where this level is such that an added path that traps exactly that much area within the rectangle forms part of an overall path that traps an area of n^2 on its lower-side in the first quadrant.

We now present a resampling argument to establish that mean facet length under $\mathbb{P}_\lambda^{N^2}$ scales at most as $N^{2/3+o(1)}$. We will argue that

$$\mathbb{P}_\lambda^{N^2} \left(\text{the facet length of } \Gamma \text{ is at least } N^{2/3+\varepsilon} \right) \rightarrow 0 \quad (2)$$

as $N \rightarrow \infty$ for any given $\varepsilon > 0$. The excess area bound (1) will play a central role in deriving this result.

The symbol \mathbb{P} will denote the randomness of the procedure that will enable this derivation. The procedure begins simply by sampling the law $\mathbb{P}_\lambda^{N^2}$. Call the sample Γ^{in} . This is our *input* configuration; we will randomly modify it to form the output.

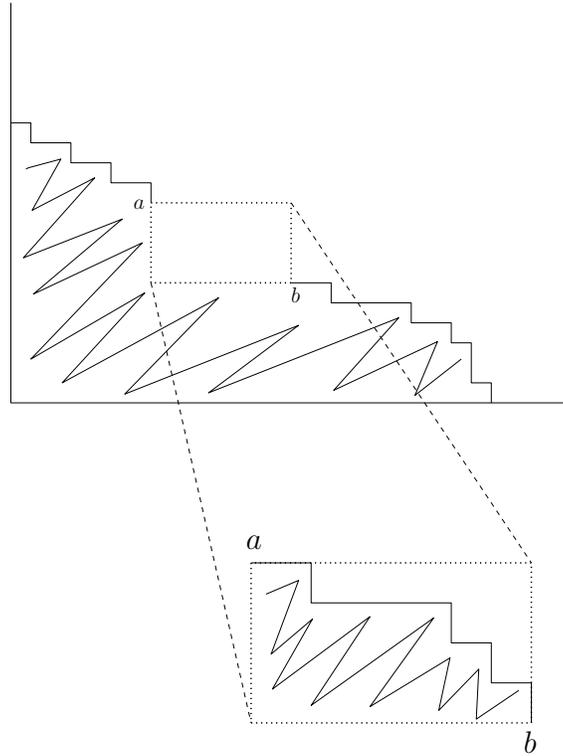


FIGURE 1.7. A realization Γ of $\mathbb{P}_\lambda^{N^2}$ is resampled between its visits to vertices a and b . The sum of the areas of the hatched regions – the lower one indicating the resampled curve – must be least N^2 in order that the random operation maintains the law $\mathbb{P}_\lambda^{N^2}$.

Define the event

$$\text{BigFacet} = \left\{ \text{the facet length of } \Gamma^{\text{in}} \text{ is at least } N^{2/3+\varepsilon} \right\}.$$

We pick vertices a and b along Γ , making a choice uniformly among those in which b is clockwise of a . The part of Γ between a and b is removed, and replaced by a new downright path $\Gamma_{a \rightarrow b}^{\text{resample}}$ connecting a to b , in such a way that the resulting overall path – running along Γ from the start of this path to a ; then along $\Gamma_{a \rightarrow b}^{\text{resample}}$; then from b along Γ to this path's end – which we will call Γ^{out} , retains the law of $\mathbb{P}_\lambda^{N^2}$. Recall that the rule for sampling the random downright path is the uniform choice with an input-determined area constraint discussed before this derivation began.

Recalling that the facet of a downright path is the concave boundary facet that cuts through the diagonal $y = x$, we set

$$\text{GoodHit} = \left\{ \text{the endpoints of the facet of } \Gamma^{\text{in}} \text{ are } a \text{ and } b \right\};$$

then $\mathbb{P}(\text{GoodHit}) = O(1)N^{-2}$, since the mean of $L(\Gamma^{\text{in}})$ has order N .

Our aim is to show that, when the input verifies **BigFacet**, then experiencing the modest good fortune of the occurrence of **GoodHit** and a little further serendipity in the action of the resampling, we obtain an output that has an excess in area that is atypical according to (1) – an atypicality which can only mean that **BigFacet** is a rare event under the input configuration law $\mathbb{P}_\lambda^{N^2}$.

The remaining element which needs to be made precise is to describe which further serendipity it is exactly on the part of the random resampling which will force this structure in the output. The event will be called **MoreOutThanIn**. To specify it, note that $\Gamma_{a \rightarrow b}^{\text{resample}}$ partitions the rectangle Rect whose northwest corner is a and whose southeast corner is b into two regions, which may naturally be called the lower-left region and the upper-right region. The event **MoreOutThanIn** occurs when the area of the lower-left region is at least the sum of one-half the area of Rect and the quantity $\|a - b\|^{3/2}$.

We may find a small constant $c > 0$ such that, for all $N \in \mathbb{N}$, the bound $\mathbb{P}(\text{MoreOutThanIn}) \geq c$ holds. Indeed, we could take $c = 1/2$ were we merely to demand that the lower-left region trap one-half of Rect 's area; and, since the resampling of Γ in this rectangle experiences a Gaussian fluctuation, of order given by the square root of the distance between a and b , we see that the further area gain of $\|a - b\|^{3/2}$ is achieved with constant probability.

Our central claim about the random resampling asserts that

$$\text{BigFacet} \cap \text{GoodHit} \cap \text{MoreOutThanIn} \subseteq \left\{ A(\Gamma^{\text{out}}) \geq N^2 + cN^{1+3\epsilon/2} \right\} \quad (3)$$

and is illustrated by Figure 1.8.

To verify the claim, consider the path Γ' which is formed by the replacement of the section of Γ about a and b by the planar line segment that connects these two endpoints. On the event **GoodHit**, we have that $A(\Gamma^{\text{in}}) \leq A(\Gamma')$ – indeed, in this circumstance, a and b are endpoints of a facet of the input path, so the formation of Γ' involves the substitution of the subpath of the input path in its journey between these consecutive extreme points with the affine route directly between the endpoints of the removed journey. Furthermore, the event **BigFacet** \cap **GoodHit** \cap **MoreOutThanIn** entails that $A(\Gamma^{\text{out}}) \geq A(\Gamma') + cN^{1+3\epsilon/2}$. Indeed, $A(\Gamma^{\text{out}}) - A(\Gamma') \geq c\|a - b\|^{3/2} \geq cN^{1+3\epsilon/2}$, where it is the occurrence of **MoreOutThanIn** that ensures the first inequality and that of **BigFacet** \cap **GoodHit** which ensures the second. Thus do we obtain (3).

Since $\mathbb{P}(\text{BigFacet})$ is equal to

$$\mathbb{P}_\lambda^{N^2} \left(\text{the facet length of } \Gamma \text{ is at least } N^{2/3+\epsilon} \right),$$

while $\mathbb{P}(\text{GoodHit} | \text{BigFacet}) = O(1)N^{-2}$ and

$$\mathbb{P} \left(\text{MoreOutThanIn} \mid \text{BigFacet} \cap \text{GoodHit} \right) \geq c,$$

we see that

$$\begin{aligned} & \mathbb{P}_\lambda^{N^2} \left(\text{the facet length of } \Gamma \text{ is at least } N^{2/3+\epsilon} \right) \times O(1)N^{-2} \times c \\ & \leq \mathbb{P} \left(\text{BigFacet} \cap \text{GoodHit} \cap \text{MoreOutThanIn} \right) \\ & \leq \mathbb{P} \left(A(\Gamma^{\text{out}}) \geq N^2 + cN^{1+3\epsilon/2} \right) \end{aligned}$$

where the second inequality makes use of the claim (3). Since Γ^{out} under \mathbb{P} is $\mathbb{P}_\lambda^{N^2}$ -distributed, we see that the term in the last displayed line equals

$$\mathbb{P}_\lambda^{N^2} \left(\text{ExcessArea}(\Gamma) \geq N \cdot cN^{3\epsilon/2} \right)$$

which by (1) is at most $C \exp \{ -c^2 N^{3\epsilon/2} \}$. Thus, we learn that

$$\mathbb{P}_\lambda^{N^2} \left(\text{the facet length of } \Gamma \text{ is at least } N^{2/3+\epsilon} \right) \leq O(1)N^2 c^{-1} C \exp \{ -c^2 N^{3\epsilon/2} \},$$

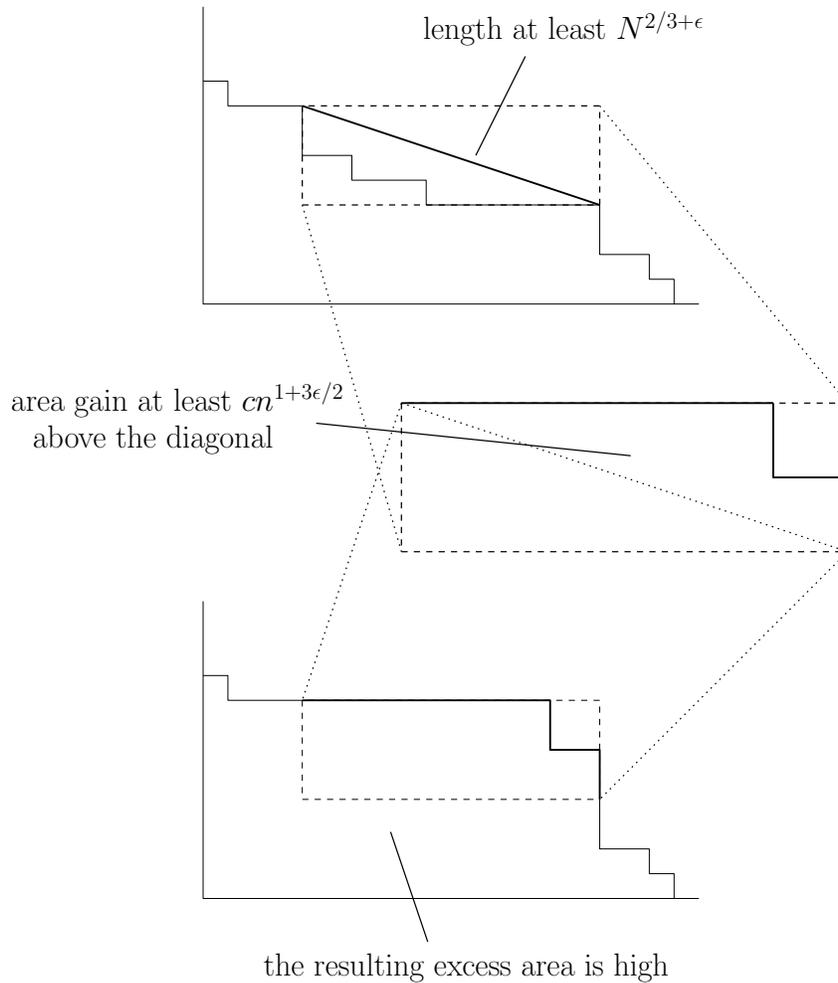


FIGURE 1.8. When the original downright path verifies $\text{BigFacet} \cap \text{GoodHit}$ (depicted at the top), and the resampling depicted in the middle verifies MoreInThanOut , the area gain relative to the diagonal in the depicted rectangle is secured for the resampled downright path at the bottom – so that this path has an area excess of at least $cN^{1+3\epsilon/2}$.

so that (2) is indeed obtained.

Our principal result about the area-trap path model, Theorem 1.1, makes two assertions. First, mean facet length is said to have a typical order of $N^{2/3}$; and second, mean local roughness is said to have order $N^{1/3}$. We have explained the reason for the upper bound on mean facet length. We will not discuss much the accompanying lower bound, nor the bounds needed on local roughness – beyond saying that all these bounds arise in essence from the Gaussian fluctuation that is inherent in the area-constrained path, a path which is largely free from cares dictated by area on spatial scales below $N^{2/3}$.

We have seen how resampling the area trap path model has betrayed its characteristic KPZ features – the two exponents; its global curvature; and its local Gaussianity. The model has thus largely served the purpose of exposition that we have intended for it. It is perhaps worth mentioning,

however, that the basic resampling technique of proof leads to stronger conclusions. We may define $\text{MFL}(\Gamma)$ to be the *maximum facet length* of a downright path $\Gamma \in \Lambda$ – the maximum length of any of the planar line segments that comprise the graph of the concave majorant \mathcal{C}_Γ . Further we may denote by MLR the *maximum local roughness* of Γ – this being the maximum over vertices in Γ of the distance from the vertex to the graph of \mathcal{C}_Γ .

The resampling technique may be pursued to obtain the next result, in which the poly-logarithmic corrections to the one-third and two-third powers are exhibited for these maximum statistics.

THEOREM 1.2. *There exist positive constants C and c such that*

$$\mathbb{P}_\lambda^{N^2} \left(c \leq \frac{\text{MLR}(\Gamma)}{n^{1/3}(\log n)^{2/3}} \leq C \right) \rightarrow 1$$

and

$$\mathbb{P}_\lambda^{N^2} \left(c \leq \frac{\text{MFL}(\Gamma)}{n^{2/3}(\log n)^{1/3}} \leq C \right) \rightarrow 1$$

as $N \rightarrow \infty$.

The area trap path model is in a ‘baby’ KPZ universality class – it evinces exponents and qualitative features of the richer universality, but presumably without its rich distributional scaled structure including the GUE Tracy-Widom distribution. The model is a useful testing ground for some aspects of KPZ universality – the polylogarithmic corrections in Theorem 1.2 are also found in KPZ, in problems concerning the maximal scaled energy, and the maximum fluctuation, of short geodesics.

Lecture Two: the Gibbs property of Brownian LPP, and the main theorems

Enough of baby KPZ – we begin the lecture by formulating the model in the KPZ universality class which will become our object of study.

2.1. Brownian last passage percolation [LPP]

2.1.1. The model’s definition. This model was introduced by [GW91] and further studied in [OY02]; we will call it Brownian LPP. On a probability space carrying a law labelled \mathbb{P} , we let $B : \mathbb{Z} \times \mathbb{R} \rightarrow \mathbb{R}$ denote an ensemble of independent two-sided standard Brownian motions $B(k, \cdot) : \mathbb{R} \rightarrow \mathbb{R}$, $k \in \mathbb{Z}$.

Let $i, j \in \mathbb{Z}$ with $i \leq j$. We denote the integer interval $\{i, \dots, j\}$ by $\llbracket i, j \rrbracket$. Further let $x, y \in \mathbb{R}$ with $x \leq y$. With these parameters given, we consider the collection of non-decreasing lists $\{z_k : k \in \llbracket i + 1, j \rrbracket\}$ of values $z_k \in [x, y]$. With the convention that $z_i = x$ and $z_{j+1} = y$, we associate an energy to any such list, namely $\sum_{k=i}^j (B(k, z_{k+1}) - B(k, z_k))$. We may then define the maximum energy, $M_{(x,i) \rightarrow (y,j)}^1$, to be the supremum of the energies of all such lists.

We use a simpler notation in a special case – that of narrow wedge initial data. For $x > 0$, we set $M_n^1(x) = M_{(0,1) \rightarrow (x,n)}^1$, so that $M_n^1 : [0, \infty) \rightarrow \mathbb{R}$ is a random energy profile of geodesics emanating from the origin to a variable location. (In fact, these geodesics begin at $(0, 1)$ rather than at $(0, 0)$ – a minor discrepancy that ensures that n horizontal coordinates are traversed by such paths.)

2.1.2. A geometric view: staircases. In order to make a study of those lists that attain the maximum energy, we begin by noting that the lists are in bijection with certain subsets of $[x, y] \times [i, j] \subset \mathbb{R}^2$ that we call *staircases*. Staircases offer a geometric perspective on Brownian LPP and perhaps help in visualizing the problems in question.

The staircase associated to the non-decreasing list $\{z_k : k \in \llbracket i + 1, j \rrbracket\}$ is specified as the union of certain horizontal planar line segments, and certain vertical ones. The horizontal segments take the form $[z_k, z_{k+1}] \times \{k\}$ for $k \in \llbracket i, j \rrbracket$. Here, the convention that $z_i = x$ and $z_{j+1} = y$ is again adopted. The right and left endpoints of each consecutive pair of horizontal segments are interpolated by a vertical planar line segment of unit length. It is this collection of vertical line segments that form the vertical segments of the staircase.

The resulting staircase may be depicted as the range of an alternately rightward and upward moving path from starting point (x, i) to ending point (y, j) . The set of staircases with these starting and ending points will be denoted by $SC_{(x,i) \rightarrow (y,j)}$. Such staircases are in bijection with the collection of non-decreasing lists considered above. Thus, any staircase $\phi \in SC_{(x,i) \rightarrow (y,j)}$ is assigned an energy $E(\phi) = \sum_{k=i}^j (B(k, z_{k+1}) - B(k, z_k))$ via the associated z -list.

2.2. Mutually avoiding staircase collections – and ensembles of energy profiles

Let $n \in \mathbb{N}$ and $\ell \in \llbracket 1, n \rrbracket$.

For $x > 0$, let $SC_{(0,1) \rightarrow (x,n)}^\ell$ denote the collection of ℓ -tuples $\bar{\phi} = (\phi_1, \dots, \phi_\ell)$, where

- $\phi_j \in SC_{(0,j) \rightarrow (x,n-\ell+j)}$ for $j \in \llbracket 1, \ell \rrbracket$;
- and the union of the horizontal planar segments of the ϕ_j are pairwise disjoint.

We define the energy $E(\bar{\phi})$ of such a collection $\bar{\phi}$ by setting

$$E(\bar{\phi}) = \sum_{j=1}^{\ell} E(\phi_j),$$

where note that the summand has already been defined.

We now define the maximum ℓ -tuple energy

$$M_n^\ell(x) = \sup \left\{ E(\bar{\phi}) : \bar{\phi} \in SC_{(0,1) \rightarrow (x,n)}^\ell \right\}.$$

DEFINITION 2.1. The n -indexed Brownian LPP line ensemble $L_n : \llbracket 1, n \rrbracket \times [0, \infty) \rightarrow \mathbb{R}$ is defined by insisting that, for $\ell \in \llbracket 1, n \rrbracket$,

$$M_n^\ell(x) = \sum_{i=1}^{\ell} L_n(i, x).$$

In this way, $L_n(i, x)$ is the energy gain that arises when the number of disjoint staircases from $(0, 1)$ to (x, n) is increased from $i - 1$ to i . It is thus not hard to see that $L_n(i, x) < L_n(i - 1, x)$ whenever $i \geq 2$ and $x > 0$.

DEFINITION 2.2. Let $n \in \mathbb{N}$. The Dyson Brownian motion line ensemble $\text{DysonBM}_n : \llbracket 1, n \rrbracket \times [0, \infty) \rightarrow \mathbb{R}$ is, formally, a system of n standard Brownian motions on $[0, \infty)$ conditioned never to intersect (and recorded in decreasing order).

The conditioning in this definition is singular and something needs to be done to make sense of it. The theory of the Doob- h transform could be invoked. Or we could take the weak limit as $\varepsilon \searrow 0$ and $K \nearrow \infty$ of a system of Brownian motions each begun at distinct locations of absolute value at least ε and conditioned on absence of intersection during $[0, K]$.

The fundamental relationship between Brownian LPP and mutually avoiding Brownian motions is next.

THEOREM 2.3 (O’Connell-Yor 2002). *Let $n \in \mathbb{N}$. The Brownian LPP and Dyson Brownian motion line ensembles, L_n and DysonBM , both map $\llbracket 1, n \rrbracket \times [0, \infty)$ to \mathbb{R} . They are equal in law.*

Note that $L_n(1, n)$ is the geodesic energy for the journey $(0, 1) \rightarrow (x, n)$. How does this random variable behave for high n ? Since $L_n(1, n)$ is equal in law to $\text{DysonBM}_n(1, n)$, a connection of Dyson Brownian motion to random matrix theory offers a means of furnishing an answer.

THEOREM 2.4. *Let $t > 0$. The point process indicated in the Figure 2.1, namely $\{\text{DysonBM}_n(i, t) : i \in \llbracket 1, n \rrbracket\}$, is equal in law to the set of eigenvalues $\{\lambda_n(i, t) : i \in \llbracket 1, n \rrbracket\}$ of a matrix sampled from the Gaussian unitary ensemble with entry variance t .*

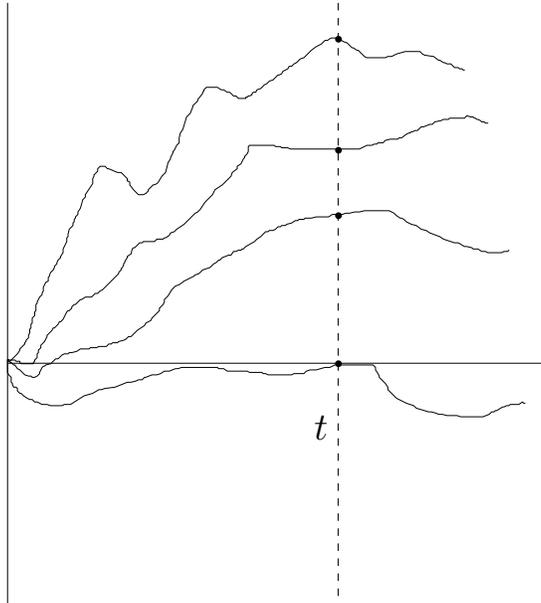


FIGURE 2.1. The dots indicate the point process of values of Dyson Brownian motion at time t .

Using this connection, and tail bounds on the top GUE eigenvalue due to Aubrun and Ledoux, it is understood that $\text{DysonBM}_n(1, n)$ behaves to first order as $2n$; has standard deviation $n^{1/3}$; and upper bounds on the scaled tail of the form $C \exp \{ -ct^{3/2} \}$.

That is, setting

$$\mathcal{L}_n^{\text{sc}}(1, 0) = 2^{-1/2} n^{-1/3} (L_n(1, n) - 2n),$$

we have:

PROPOSITION 2.5. *There exist constants $C, c > 0$ such that for every $n \in \mathbb{N}$ it is the case that*

(1) for $s \in [0, 2^{1/2} n^{1/3}]$,

$$\mathbb{P}(\mathcal{L}_n^{\text{sc}}(1, 0) \leq -s) \leq C \exp \{ -cs^{3/2} \},$$

(2) and, for $s \geq 0$,

$$\mathbb{P}(\mathcal{L}_n^{\text{sc}}(1, 0) \geq s) \leq C \exp \{ -cs^{3/2} \}.$$

Proof. The first statement is due to Ledoux and the second to Aubrun in view of the connections between GUE and the Dyson Brownian motion marginal, and between Dyson Brownian motion and Brownian LPP. \square

As the notation $\mathcal{L}_n^{\text{sc}}(1, \cdot)$ suggests, we want to specify a stochastic process $\mathcal{L}_n^{\text{sc}}(1, \cdot)$ that describes the scaled energy of a geodesic from the origin to a further endpoint that is in a variable location.

In view of the consideration described in Figure 2.2, we define the *scaled* Brownian last passage percolation line ensemble

$$\mathcal{L}_n^{\text{sc}} : \llbracket 1, n \rrbracket \times \left[-\frac{1}{2} n^{1/3}, \infty \right) \rightarrow \mathbb{R}$$

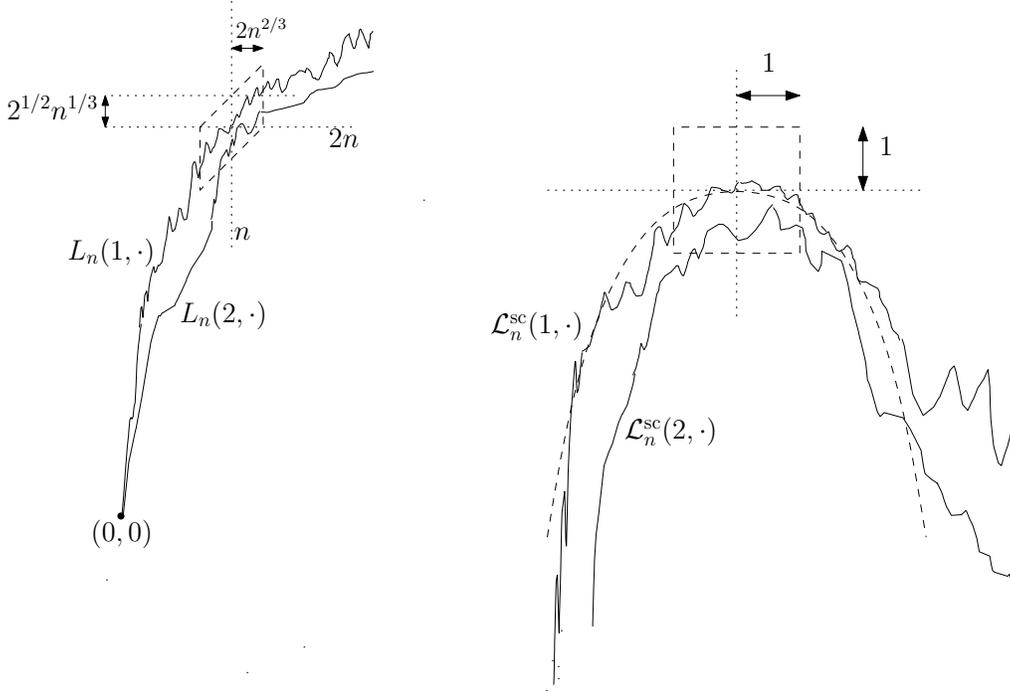


FIGURE 2.2. The parallelogram on the left is mapped to the square on the right by an affine change of coordinates. The size of the parallelogram – its height of order $n^{1/3}$ and its width of order $n^{2/3}$ – are, in view of Proposition 2.5 and Brownian scaling, suitable for the specification of the scaled ensemble as the image curve collection on the right.

by setting, for $(i, x) \in \llbracket 1, n \rrbracket \times [-\frac{1}{2}n^{1/3}, \infty)$,

$$\mathcal{L}_n^{\text{sc}}(i, x) = 2^{-1/2}n^{-1/3} \left(L_n(i, n + 2n^{2/3}x) - 2n - 2n^{2/3}x \right). \quad (4)$$

What does the scaled ensemble describe? A geodesic that runs from $(0, 1)$ to (n, n) has energy of the form $2n + \Theta(1)n^{1/3}$, where the random quantity $\Theta(1)$ is this geodesic's *scaled energy* or *weight*. This weight is $\mathcal{L}_n^{\text{sc}}(1, 0)$. If instead the geodesic runs from $(0, 1)$ to $(n + 2n^{2/3}, n)$, then the associated weight is $\mathcal{L}_n^{\text{sc}}(1, 1)$. In the case of a triple of disjoint staircases, the weight is $\sum_{i=1}^3 \mathcal{L}_n^{\text{sc}}(i, 0)$.

Indeed, the affine scaling map $R_n : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by

$$R_n(v_1, v_2) = \left(2^{-1}n^{-2/3}(v_1 - v_2), v_2/n \right).$$

(and depicted in Figure 2.3) sends staircases to n -zigzags and geodesics to *polymers* in such a way that

$$\rho_{n;(0,0)}^{(x,1)} := \text{the image under } R_n \text{ of the geodesic from } (0, 1) \text{ to } R_n^{-1}(x, 1),$$

which is the polymer whose journey is between the scaled coordinate locations $(0, 0)$ and $(x, 1)$, may be viewed as having weight $\mathcal{L}_n^{\text{sc}}(1, x)$; while the multi-polymer seen in the middle sketch of the upcoming Figure 2.6 – the scaled counterpart to the disjoint staircase triple mentioned a moment ago – has weight $\sum_{i=1}^3 \mathcal{L}_n^{\text{sc}}(i, 0)$.

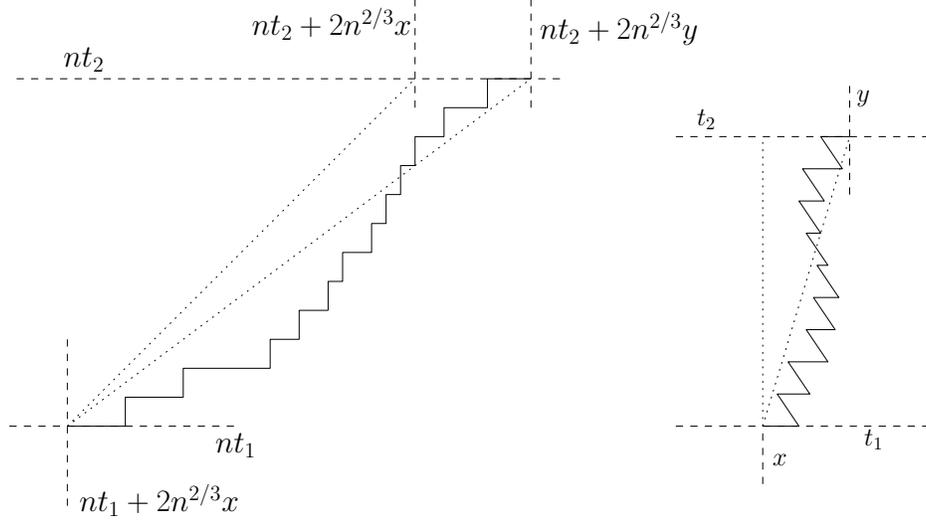


FIGURE 2.3. Let $n \in \mathbb{N}$, $t_1 < t_2$ be reals and also let $x, y \in \mathbb{R}$. The endpoints of the geodesic in the left sketch are such that, when the scaling map R_n is applied to produce the right sketch, the result is an n -polymer from (x, t_1) to (y, t_2) . We work with the special case $t_1 = 0$ and $t_2 = 1$ – but scaling considerations dictate the behaviour of the general case from the special one.

2.3. The Brownian Gibbs property

This property is a crucial probabilistic tool for analysing mutually avoiding systems of Brownian motions, such as $L_n : \llbracket 1, n \rrbracket \times [0, \infty) \rightarrow \mathbb{R}$ or indeed $\mathcal{L}_n^{\text{sc}} : \llbracket 1, n \rrbracket \times [0, \infty) \rightarrow \mathbb{R}$.

First a little notation. Let $k \in \mathbb{N}$. Write \bar{x} for a vector in \mathbb{R}^k . Such a vector $\bar{x} = (x_1, x_2, \dots, x_k)$ is a k -decreasing list if $x_i > x_{i+1}$ for $1 \leq i \leq k-1$. Write $\mathbb{R}_{>}^k \subseteq \mathbb{R}^k$ for the set of k -decreasing lists.

Recall Brownian bridge $B : [a, b] \rightarrow \mathbb{R}$, $B(a) = x \in \mathbb{R}$, $B(b) = y \in \mathbb{R}$ – this is Brownian motion $W : [a, b] \rightarrow \mathbb{R}$, $W(a) = x$, conditioned on $W(b) = y$.

DEFINITION 2.6. Let $k \in \mathbb{N}$, $a, b \in \mathbb{R}$ with $a < b$, and $\bar{x}, \bar{y} \in \mathbb{R}_{>}^k$. Write $\mathcal{B}_{k; \bar{x}, \bar{y}}^{[a, b]}$ for the law of the ensemble $B : \llbracket 1, k \rrbracket \times [a, b] \rightarrow \mathbb{R}$ whose constituent curves $B(i, \cdot) : [a, b] \rightarrow \mathbb{R}$, $i \in \llbracket 1, k \rrbracket$, are independent Brownian bridges that satisfy $B(i, a) = x_i$ and $B(i, b) = y_i$.

Let $f : [a, b] \rightarrow \mathbb{R} \cup \{-\infty\}$ be a measurable function such that $x_k > f(a)$ and $y_k > f(b)$. Define the non-touching event on an interval $A \subset [a, b]$ with lower boundary data f by

$$\text{NoTouch}_f^A = \left\{ \text{for all } x \in A, B(i, x) > B(j, x) \text{ whenever } 1 \leq i < j \leq k, \text{ and } B(k, x) > f(x) \right\}.$$

We omit the subscript f in the case that it equals $-\infty$ throughout $[a, b]$ (and thus plays no role). We omit the superscript A in the case that $A = [a, b]$. With this convention, the event NoTouch always imposes *internal* curve avoidance, but only imposes *external* avoidance of the lower boundary condition when this is indicated in the subscript.

The conditional measure $\mathcal{B}_{k; \bar{x}, \bar{y}}^{[a, b]}(\cdot | \text{NoTouch}_f)$ is the *mutually avoiding Brownian bridge ensemble on the interval $[a, b]$ with entrance data \bar{x} , exit data \bar{y} and lower boundary condition f* .

We will sometimes refer to the *acceptance probability*, which is defined to be $\mathcal{B}_{k; \bar{x}, \bar{y}}^{[a, b]}(\text{NoTouch}_f)$.

The law $\mathcal{B}_{k;\bar{x},\bar{y}}^{[a,b]}(\cdot \mid \text{NoTouch})$ is a prototypical example of a line ensemble that enjoys the Brownian Gibbs property.

DEFINITION 2.7. Let $\mathcal{L} : \llbracket 1, n \rrbracket \times [a, b] \rightarrow \mathbb{R}$ be an ensemble. It satisfies the Brownian Gibbs property if, whenever $k \in \llbracket 1, n \rrbracket$ and $[u, v] \subseteq [a, b]$, the conditional distribution of \mathcal{L} restricted to $\llbracket 1, k \rrbracket \times [u, v]$ given the data \mathcal{L} on the remainder – on $\llbracket 1, k \rrbracket \times ([a, b] \setminus [u, v]) \cup \llbracket k+1, n \rrbracket \times [a, b]$ – equals $\mathcal{B}_{k;\bar{x},\bar{y}}^{[u,v]}(\cdot \mid \text{NoTouch}_f)$, where $\bar{x} = (\mathcal{L}(1, u), \dots, \mathcal{L}(k, u))$, $\bar{y} = (\mathcal{L}(1, v), \dots, \mathcal{L}(k, v))$ and $f = \mathcal{L}(k+1, \cdot)|_{\cdot \in [u, v]}$. Here, we take $f = -\infty$ if $k = n$.

See Figure 2.4.

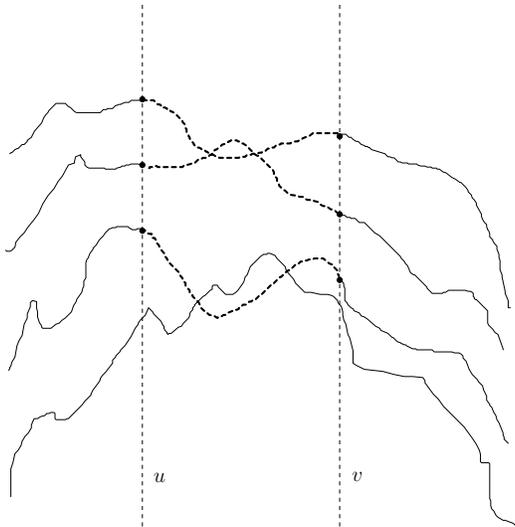


FIGURE 2.4. In a rejection sampling view of the Brownian Gibbs resampling of the top three curves of an ensemble on the interval $[u, v]$, independent Brownian bridges on this interval with the respective bead endpoints are consecutively sampled until the outcome verifies the concerned avoidance constraints. The three bold dotted curves in the depicted attempt fail due to both internal and external violations of avoidance.

Back to the *scaled* Brownian last passage percolation line ensemble $\mathcal{L}_n^{\text{sc}} : \llbracket 1, n \rrbracket \times [-\frac{1}{2}n^{1/3}, \infty) \rightarrow \mathbb{R}$. Recall Figure 2.2, which offered a schematic depiction of the highest two curves in the unscaled and scaled ensembles for a high value of $n \in \mathbb{N}$. The dashed parallelogram on the left transforms into the dashed square on the right under the affine change of coordinates in (4) by which $\mathcal{L}_n^{\text{sc}}$ is formed from L_n . The right sketch depicts when n is large the highest curves in $\mathcal{L}_n^{\text{sc}}$. These surge upwards until, far to the left of the origin, – at scale $-n^{-1/9}$ – they join a bounded channel about the parabola $-2^{-1/2}x^2$, which they then typically inhabit until far beyond the origin on the right, – at scale $n^{1/9}$ – when the parabola drops away beneath them. We may expect then that the scaled ensemble $\mathcal{L}_n^{\text{sc}}$:

- satisfies the Brownian Gibbs property;
- satisfies one-point tails on the top curve inherited from Aubrun and Ledoux;
- and in the large is parabolic.

The next definition captures all this.

By a *Brownian Gibbs ensemble*, we mean a line ensemble that satisfies the Brownian Gibbs property.

DEFINITION 2.8. Consider a Brownian Gibbs ensemble that has the form

$$\mathcal{L} : \llbracket 1, n \rrbracket \times [-z_n, \infty) \rightarrow \mathbb{R},$$

and which is defined on a probability space under the law \mathbb{P} . The number n of ensemble curves and the absolute value z_n of the finite endpoint may take any values in \mathbb{N} and $[0, \infty)$. (In fact, we may also take $z_n = \infty$, except that we would then take the domain of definition of \mathcal{L}_n to be $\llbracket 1, n \rrbracket \times \mathbb{R}$.)

For two constants $C \geq c > 0$, the ensemble \mathcal{L} is called (c, C) -regular – or, in practice, simply regular – if:

(1) **Endpoint escape.** $z_n \geq cn^{1/3}$.

(2) **One-point lower tail.** If $z \geq -z_n$ satisfies $|z| \leq cn^{1/9}$, then

$$\mathbb{P}\left(\mathcal{L}(1, z) + 2^{-1/2}z^2 \leq -s\right) \leq C \exp\{-cs^{3/2}\}$$

for all $s \in [1, n^{1/3}]$.

(3) **One-point upper tail.** If $z \geq -z_n$ satisfies $|z| \leq cn^{1/9}$, then

$$\mathbb{P}\left(\mathcal{L}(1, z) + 2^{-1/2}z^2 \geq s\right) \leq C \exp\{-cs^{3/2}\}$$

for all $s \in [1, \infty)$.

PROPOSITION 2.9. *There exist choices of the positive constants c and C such that each of the scaled Brownian LPP line ensembles $\mathcal{L}_n^{\text{sc}} : \llbracket 1, n \rrbracket \times [-\frac{1}{2}n^{1/3}, \infty) \rightarrow \mathbb{R}$, $n \in \mathbb{N}$, is (c, C) -regular.*

Proof. Details are omitted, but in essence this is the Brownian Gibbs property, which is inherited from DysonBM_n via O’Connell-Yor, along with Aubrun’s and Ledoux’s tail bounds on the top eigenvalue in GUE_n . \square

Now $\mathcal{L}_n^{\text{sc}}$ enjoys integrable features – for example, there are determinantal expressions for probabilities such as that of the event depicted in Figure 2.5, in which each of the very short vertical intervals is visited by one of the curves in the ensemble. Prähofer and Spohn [PS02] introduced in 2002 what they called the multi-line Airy process – in essence, the high- n limit of the ensemble $\mathcal{L}_n^{\text{sc}}$ when the limit is taken in the sense of finite dimensional distributions.

The first application of the Brownian Gibbs property was in joint work with Ivan Corwin:

THEOREM 2.10 ([CH14]). *We may couple the ensembles $\mathcal{L}_n^{\text{sc}}$, $n \in \mathbb{N}$, on a single probability space so that the restrictions $\mathcal{L}_n^{\text{sc}} : \llbracket 1, k \rrbracket \times I \rightarrow \mathbb{R}$ (where here $k \in \mathbb{N}$ is any positive integer, and $I \subset \mathbb{R}$ any compact interval) converge uniformly to the restriction to this domain of a non-intersecting ensemble $\mathcal{L} : \mathbb{N} \times \mathbb{R} \rightarrow \mathbb{R}$ (with $\mathbb{N} = \{1, 2, \dots\}$) that satisfies the Brownian Gibbs property.*

Indeed, setting $\mathcal{A}(i, x) = 2^{1/2}\mathcal{L}(i, x) + x^2$, $\mathcal{A} : \mathbb{N} \times \mathbb{R} \rightarrow \mathbb{R}$ is an ensemble that is stationary in x – this is the Airy line ensemble.

We will not explain the proof, but the basic idea is to use finite dimensional convergence, with the Brownian Gibbs property furnishing the necessary regularity. We will explain now at least a little of the technical aspect of this convergence. Consider a Brownian Gibbs resampling of $\mathcal{L}_n^{\text{sc}}$, i.e., a random map sending the law $\mathcal{L}_n^{\text{sc}}$ to itself in which say the top two curves of this ensemble are

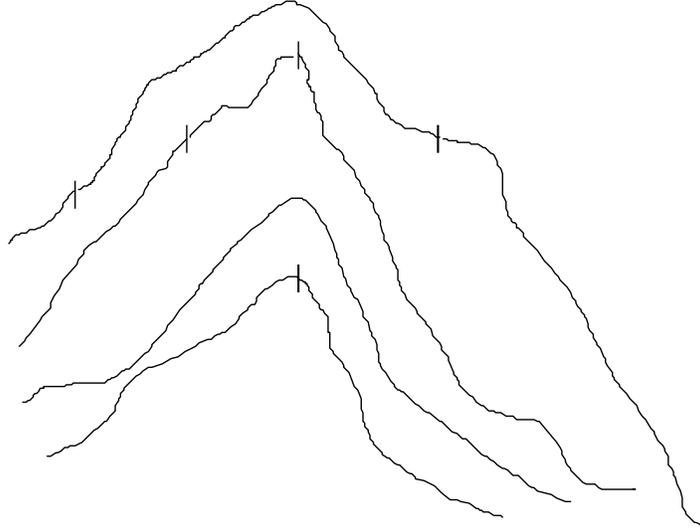


FIGURE 2.5. Imagine that each of the depicted vertically aligned intervals has an infinitesimal length dx . The depicted curves of $\mathcal{L}_n^{\text{sc}}$ collectively visit all these intervals. The probability of this happening has the form $h(dx)^K$, where K denotes the number of the intervals. The factor h depends on the intervals' locations and has a determinantal expression.

resampled on the interval $[-1, 1]$. The conditional distribution in the update step is $\mathcal{B}(\cdot | \text{NoTouch}_f)$, where $\mathcal{B} = \mathcal{B}_{2; x_1, x_2, y_1, y_2}$ and $f = \mathcal{L}_n^{\text{sc}}(3, \cdot)$, $x_1 = \mathcal{L}_n^{\text{sc}}(1, -1)$, $x_2 = \mathcal{L}_n^{\text{sc}}(2, -1)$, $y_1 = \mathcal{L}_n^{\text{sc}}(1, 1)$ and $y_2 = \mathcal{L}_n^{\text{sc}}(2, 1)$.

Recall the acceptance probability $\mathcal{A}_n = \mathcal{B}(\text{NoTouch}_f)$. The law \mathcal{B} is unadulterated Brownian randomness – with square-root regularity for curve fluctuation, enough for tightness given finite dimensional distributional control. In the resample that maps $\mathcal{L}_n^{\text{sc}}$ to itself, it is the possible rarity of NoTouch_f – the possibly low value of the acceptance probability – which may disrupt the presence of this unadulterated Brownian randomness in the image copy of $\mathcal{L}_n^{\text{sc}}$. So the key technical proposition in [CH14] states:

PROPOSITION 2.11. *For all $\varepsilon > 0$, there exists $\delta > 0$ such that, for $n \in \mathbb{N}$,*

$$\mathbb{P}(\mathcal{A}_n \geq \delta) \geq 1 - \varepsilon.$$

We will not return to establish this, but we will extensively discuss ideas that can rather quickly give a proof of this proposition.

A principal consequence of this main theorem of [CH14] is:

THEOREM 2.12. *Let $\mathcal{L} : \mathbb{N} \times \mathbb{R} \rightarrow \mathbb{R}$ denote the weak limit of $\mathcal{L}_n^{\text{sc}}$ as $n \rightarrow \infty$. Let $k \in \mathbb{N}$ and let $I = [x, y] \subset \mathbb{R}$ be a compact interval. Then the random function $I \rightarrow \mathbb{R} : z \rightarrow \mathcal{L}(k, z) - \mathcal{L}(k, x)$ is absolutely continuous with respect to Brownian motion.*

We will consistently reserve the symbol \mathcal{L} to denote this ensemble $\mathcal{L} : \mathbb{N} \times \mathbb{R} \rightarrow \mathbb{R}$.

Proof of Theorem 2.12. This is in essence a consequence of \mathcal{L} being a Brownian Gibbs ensemble that is ordered, since this forces the acceptance probability on $\llbracket 1, k \rrbracket \times I$ to be almost surely positive. \square

The conclusion offered by Theorem 2.12 was recently strengthened and quantified by an extensive analysis using the Brownian Gibbs property.

A little further notation is needed to describe these more recent results.

Let $f : [a, b] \rightarrow \mathbb{R}$ be a continuous function. Such a function will be called a *standard bridge* if $f(a) = f(b) = 0$.

For a general continuous function $f : [a, b] \rightarrow \mathbb{R}$, we define $f^{a,b} : [a, b] \rightarrow \mathbb{R}$ to be the unique standard bridge obtained by affine translation of f . That is,

$$f^{[a,b]}(x) = f(x) - \frac{b-x}{b-a}f(a) - \frac{x-a}{b-a}f(b) \text{ for } x \in [a, b].$$

The notation extends to line ensembles: let $n \in \mathbb{N}$ and $a, b \in \mathbb{R}$ satisfy $-z_n \leq a \leq b$. Consistently with the bridge notation, we specify the standard bridge ensemble $\mathcal{L}_n^{[a,b]} : \llbracket 1, n \rrbracket \times [a, b] \rightarrow \mathbb{R}$ induced on $[a, b]$ by \mathcal{L}_n to be

$$\mathcal{L}_n^{[a,b]}(i, x) = \mathcal{L}_n(i, x) - \ell_n^{[a,b]}(i, x) \text{ for } (i, x) \in \llbracket 1, n \rrbracket \times [a, b],$$

where $\ell_n^{[a,b]}(i, \cdot)$ denotes the affine function whose values at a and b are $\mathcal{L}_n(i, a)$ and $\mathcal{L}_n(i, b)$.

Write $\mathcal{C}_{0,0}([a, b], \mathbb{R}) = \{f : [a, b] \rightarrow \mathbb{R} : f \text{ is a standard bridge}\}$. If $B : [a, b] \rightarrow \mathbb{R}$ is Brownian motion, then $B^{[a,b]}$ is $\mathcal{B}_{1,0,0}^{[a,b]}$ -distributed. This law is called *standard Brownian bridge*.

To say that $\mathcal{L}_n^{\text{sc}}(1, \cdot) : [-1, 1] \rightarrow \mathbb{R}$ is *uniformly absolutely continuous* with respect to Brownian motion on $[-1, 1]$ is to assert that for all $\varepsilon > 0$, there exists $\delta > 0$ such that, for all $n \in \mathbb{N}$, the condition that

$$\mathcal{B}_{1,0,0}^{[a,b]}(B \in A) < \delta$$

for any given measurable subset $A \subseteq \mathcal{C}_{0,0}([a, b], \mathbb{R})$ (and where $B = B(1, \cdot)$) implies that

$$\mathbb{P}(\mathcal{L}^{[a,b]}(1, \cdot) \in A) < \varepsilon,$$

where here we write $\mathcal{L} = \mathcal{L}_n^{\text{sc}}$. This is in essence the state of affairs achieved by [CH14] – uniformity of comparison between ensemble curves and Brownian motion is achieved in the curve index n , but without any quantitative relation being demonstrated between the parameters ε and δ . (The comparison is in fact achieved in [CH14] between ensemble curves and Brownian *motion*. We have stated the comparison in terms of affinely shifted ensemble curves and Brownian *bridge* because it is in these terms that we will succeed in quantifying such estimates.)

When a quantitative relation is known between these parameters that takes a power law form, a moment bound results on the Radon-Nikodym derivative of the two concerned measures on continuous curves.

Indeed, suppose that μ and ν are two probability measures on a common measurable space (Ω, \mathcal{F}) . To make the assertion concerning a *Radon-Nikodym moment bound* that

$$\frac{d\nu}{d\mu} \in L^{\beta-}(d\mu) \text{ for a given } \beta \in [1, \infty)$$

is the same as claiming about the *deformation in probability of rare events* that

$$\forall \eta \in (0, 1 - \beta^{-1}), \exists C = C_\eta \text{ such that } \forall A \in \mathcal{F}, \nu(A) \leq C\mu(A)^\eta.$$

We quantify the comparison of the curves in $\mathcal{L}_n^{\text{sc}}$ or in its high n limit \mathcal{L} with Brownian bridge by using the latter language, of deformation in probability of rare events.

Recall that each $\mathcal{L}_n^{\text{sc}}$ is a (c, C) -regular ensemble – and so is \mathcal{L} , with a tiny abuse of notation.

Here is the quantified comparison – one of the main theorems of [H17a].

THEOREM 2.13 (Brownian bridge regularity). *Let $n \in \mathbb{N}$. Suppose that \mathcal{L}_n is an n -curve regular ensemble. Let $K \in \mathbb{R}$ satisfy $[K, K + 1] \subset c/2 \cdot [-n^{1/9}, n^{1/9}]$. Let $k \in \mathbb{N}$ denote a curve index. Let $A \subseteq \mathcal{C}_{0,0}([K, K + 1], \mathbb{R})$ be measurable, and set $a = \mathcal{B}_{1,0,0}^{[K,K+1]}(A)$.*

If $n \geq n_0(k)$, $a \leq a_0(k)$ and

$$a \geq \exp \left\{ -c_1(k)n^{1/12} \right\}, \quad (5)$$

then

$$\mathbb{P} \left(\mathcal{L}_n^{[K,K+1]}(k, \cdot) \in A \right) \leq a \cdot \exp \left\{ (\log a^{-1})^{5/6} O_k(1) \right\},$$

where $O_k(1)$ is a k -dependent term which has no dependence on a or n .

That is, deformation of rare events is controlled strongly, with for example an event whose Brownian bridge probability is a having ensemble probability at most $a^{1-o(1)}$ where the term $a^{-o(1)}$ is controlled uniformly in high ensemble curve number n .

There are two caveats about the theorem:

1. Comparison is made of the standard bridge-valued $\mathcal{L}_n^{[K,K+1]}$ to Brownian bridge. The theorem does not attempt comparison of $\mathcal{L}_n(1, \cdot) - \mathcal{L}_n(1, K) : [K, K + 1] \rightarrow \mathbb{R}$ to standard Brownian motion on $[K, K + 1]$.

2. Imposing condition (5) entails that extremely small probability events are not considered – in a strongly n -dependent sense of ‘small’. Note however that when \mathcal{L} is considered, so that $n = \infty$, this condition is vacuously satisfied. In the $n = \infty$ case, then, the theorem implies that $\frac{d\nu}{d\mu} \in L^{\infty-}(d\mu)$ where $\mu = \mathcal{B}_{1,0,0}^{[-1,1]}$ and ν is the law of $\mathcal{L}^{[-1,1]}(1, \cdot) - \mathcal{L}^{[-1,1]}(1, -1) : [-1, 1] \rightarrow \mathbb{R}$.

To give a practical sense of the strength of the comparison, we provide a corollary. Recall first that

$$\mathcal{B}_{1,0,0}^{[0,1]} \left(\sup_{x \in [0,1]} |B(1, x)| > s \right) \in [1, 2] \cdot e^{-2s^2}.$$

COROLLARY 2.14.

$$\mathbb{P} \left(\sup_{x \in [0,1]} |\mathcal{L}^{[0,1]}(1, x)| \geq s \right) \leq C_1 \exp \left\{ -2s^2(1 - c_1 s^{-1/3}) \right\}.$$

Now the second main theorem.

THEOREM 2.15 (k -curve closeness at one point). *Let \mathcal{L}_n be a regular ensemble. Then the probability that $\mathcal{L}_n(k, 0) \geq \mathcal{L}_n(1, 0) - \eta$ behaves as $\eta^{k^2-1+o(1)}$ for small $\eta > 0$.*

The event that the top k curves in a regular ensemble come within η of each other has particular significance when that ensemble is $\mathcal{L}_n^{\text{sc}}$. As Figure 2.6 illustrates, this circumstance occurs when there exists a collection of mutually disjoint staircases with shared endpoints each of whose members comes close to attaining the maximum energy among staircases with these endpoints – where by ‘close’ is meant within order η if the measurement is undertaken in scaled units. This phenomenon is closely related to the mutual coexistence of geodesics with endpoints that are at a small positive distance when this distance is also measured in scaled units – so that Theorem 2.15, in asserting the rarity of near touch among the top several curves in regular ensembles, also has much to say in an effort to prove that such mutual coexistence of geodesics is unlikely in Brownian LPP. The last two theorems are the principal conclusions of the Brownian Gibbs analysis in [Ham17a]. Along

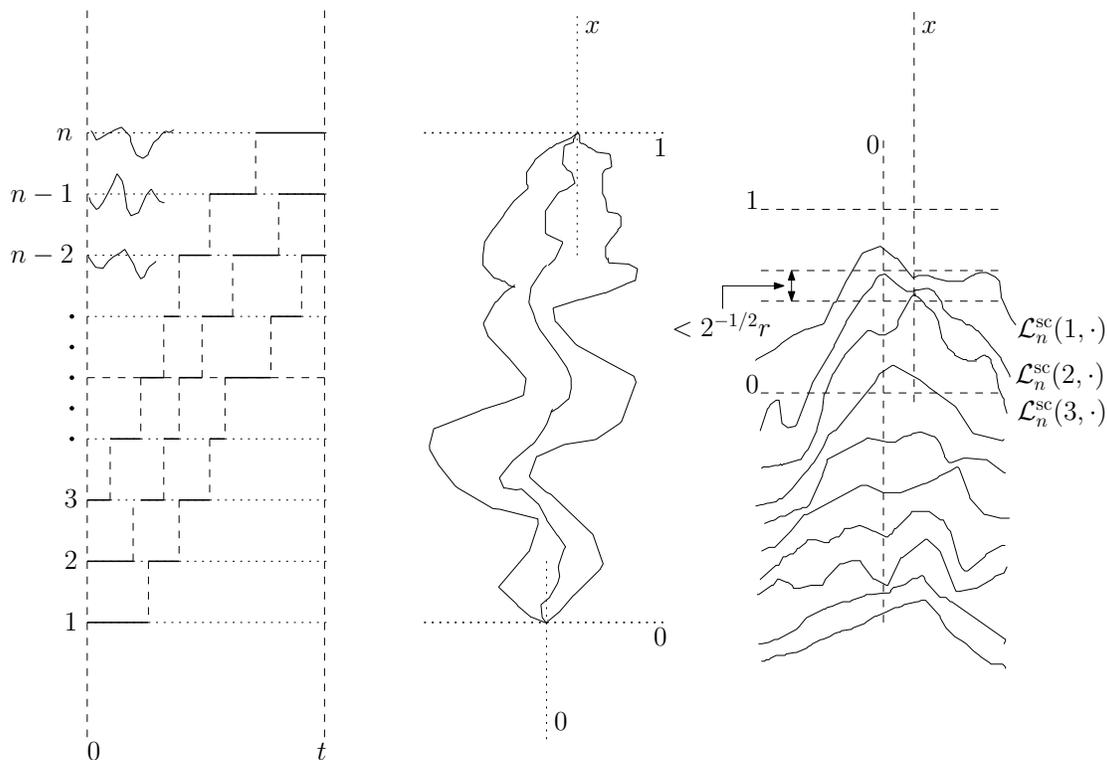


FIGURE 2.6. *Left:* In the formation of the Brownian last passage percolation line ensemble, the maximum triple energy $M_n^3(t) = \sum_{i=1}^3 L_n(i, t)$, for $t > 0$ given, is formed by considering the sum of the increments on the intervals indicated by horizontal solid black lines of the depicted independent Brownian motions and finding the maximum possible such value. *Middle:* Taking n large and setting $t = n + 2n^{2/3}x$ for a given $x \in \mathbb{R}$, we may consider the maximizing triple and depict it after the change of coordinates $(x_1, x_2) \rightarrow (\frac{1}{2}n^{-2/3}(x_1 - x_2), x_2n^{-1})$. If n is high enough, the semi-discrete structure will be indiscernible in the new sketch, and the triple of paths – a multi-polymer watermelon – will appear to share the endpoints $(0, 0)$ and $(1, x)$. *Right:* Suppose that, in the middle sketch, the elements in the path triple have very similar energies, with a collective deficit of $rn^{1/3}$ over the total available in principle, where $r > 0$ is a given small value. Measuring the deficit in units of $2^{1/2}n^{1/3}$, a $2^{-1/2}r$ -near touch will arise between the top three curves in the scaled ensemble $\mathcal{L}_n^{\text{sc}}$ over location x .

the way to their derivation, a result depending on some of the apparatus needed to derive them concerns local fluctuation of the curves in regular ensembles. For convenience, we state this result merely for the $n = \infty$ case.

THEOREM 2.16 (Local Fluctuation). *Let $\mathcal{A} : \mathbb{R} \rightarrow \mathbb{R}$ denote the Airy₂ process. There exist $M, m > 0$ such that, for all $x \in \mathbb{R}$, $\varepsilon \in (0, 1]$ and $K > 0$,*

$$\mathbb{P}\left(\sup_{h \in (0, \varepsilon)} |\mathcal{A}(x+h) - \mathcal{A}(x)| \geq K\varepsilon^{1/2}\right) \leq Me^{-mK^{3/2}};$$

here \mathcal{A} could equally be replaced by \mathcal{L} .

The route to our two main theorems can be presented as a journey of three stages, labelled ‘basic’, ‘intermediate’ and ‘advanced’. Figure 2.7 summarises this.

Level	Method	Results
BASIC	Simple BG resamples	Monotonicity; No Big Max ...
INTERMEDIATE	Missing closed middle reconstruction and the Wiener candidate	Construction of the AiryLE Control on local fluctuation of its curves
ADVANCED	The jump ensemble	L^∞ Brownian bridge regularity k-curve closeness above a point

FIGURE 2.7. The three-stage journey to our two main theorems.

As we turn to attempting to explain the proof of the principal result Theorem 2.13 on Brownian bridge regularity, we will follow this three-stage journey. This second lecture ends by discussing the first step – the basic elements of the Brownian Gibbs approach.

2.4. Stage One: Brownian Gibbs basics

Five basic general properties of Brownian Gibbs ensembles are needed.

- A: Monotonicity lemmas
- B: Stopping domain Brownian Gibbs resampling
- C: No Big Max
- D: Near parabolic invariance of regular ensembles
- E: Control on lower curves

We discuss these briefly in turn.

A: Monotonicity lemmas. Consider the basic law $\mathcal{B}_{k;\bar{x},\bar{y}}^{[a,b]}$ which appears in the BG resampling associated to $\llbracket 1, k \rrbracket \times [a, b]$. We want to record two forms of stochastic monotonicity regarding this law. The first concerns the replacement of the left or right boundary vectors \bar{x} or \bar{y} by a new vector whose components are all at least as high as their counterparts in the replaced vector. The effect is to send the law upwards, in the sense that the original law and the new one may be coupled together so that the curves in the new law are always at least as high as their counterparts in the original.

In its second form, the monotonicity result asserts that the same response occurs when we instead replace the lower boundary condition by a function which is pointwise at least as large.

These results were proved in [CH14] by a technique which is a simple illustration of the principle that, in order to understand a given random system, it is sometimes useful to embed that system in a higher dimensional random object, and analyse the latter instead. In the present case, we first move to a discrete system, with fine mesh mutually avoiding random walks approximating the curves in the law $\mathcal{B}_{k;\bar{x},\bar{y}}^{[a,b]}$. Then we introduce a Markov dynamic that at a given step attempts a local change in the form of one of these walks. We design the dynamic so that it converges in the limit of many updates to the mutually avoiding system of walks that interests us – our approximation of $\mathcal{B}_{k;\bar{x},\bar{y}}^{[a,b]}$. Then we run the dynamic simultaneously on two configurations, one with boundary data approximating $\mathcal{B}_{k;\bar{x},\bar{y}}^{[a,b]}$ and one with data approximately the modified form of this law in which either the boundary vectors or the lower boundary condition have been increased. We find an explicit initial condition for the two sets of walkers that respect the sought monotonicity – with the curves in the altered system being at least as high as in the original one. Then we check that any given update for the dynamic – an update which is shared in the two systems – never creates a violation of monotonicity. Since the two systems converge in the limit of many updates to the approximations of the law $\mathcal{B}_{k;\bar{x},\bar{y}}^{[a,b]}$ and its upward perturbed counterpart, this monotonicity transmits as desired to this law and its perturbation.

B: Stopping domain Brownian Gibbs resampling.

The basic BG resample indexed by $\llbracket 1, k \rrbracket \times I$ involves a deterministic number of curves k and a deterministic interval I . We strengthen this, keeping k fixed, but permitting I to be random. The random interval I cannot have an arbitrary form of randomness if the BG rule governing the resampling is to remain valid – one must not peek inside the curve data in $\llbracket 1, k \rrbracket \times I$ if the validity of the rule is to remain valid. The random interval I is permitted to be a *stopping domain* – in essence, we may examine the top k curves in a rightward sweep, stopping whenever we please as the new data comes in, and declare our location of stopping to be the left endpoint of I . And similarly, we may sweep to the left to find the right endpoint. This done so that I is non-empty, naturally. Such a stopping domain has been found without the observer ever peeking into $\llbracket 1, k \rrbracket \times I$ – so that the original BG resample rule works in this setting.

C: No Big Max

The axioms of a regular ensemble ensure that the upper tail of the top curve above one point – a random variable such as $\mathcal{L}_n^{\text{sc}}(1, 0)$ – has a rapidly decaying tail. But what of the maximum value attained by the ensemble's top curve on a compact interval – a quantity such as $\sup \{ \mathcal{L}_n^{\text{sc}}(1, x) : x \in [-1, 1] \}$? In principle, this could be much higher, even if the finite dimensional distributions of the ensemble are controlled. But the Brownian Gibbs property keeps the maximum on the same scale as the one-point statistic.

PROPOSITION 2.17 (No Big Max). *Let \mathcal{L}_n be an n -curve regular ensemble. For $r \in [0, c/2 \cdot n^{1/9}]$, $t \in [2^{7/2}, 2n^{1/3}]$ and $n \geq (2c)^{-18}$,*

$$\mathbb{P}\left(\sup_{x \in [-r, r]} (\mathcal{L}_n(1, x) + 2^{-1/2}x^2) \geq t\right) \leq (r + 1) \cdot 6C \exp\{-2^{-9/2}ct^{3/2}\}.$$

Proof. This amounts to showing that

$$\mathbb{P}(\text{High}(t)) \leq Ce^{-ct^{3/2}}, \tag{6}$$

where

$$\text{High}(t) = \left\{ \sup_{x \in [-1, 1]} \mathcal{L}_n(1, x) \geq t \right\}.$$

We prove this by using a stopping domain BG resampling. We search in $[-1, 1]$ from the left until we encounter a random X in this interval for which $\mathcal{L}_n(1, X) \geq t$. Then we consider the random interval $[X, 2]$ – it may be the interval $[1, 2]$ if such X is never encountered; but no matter. A little control is needed at a point, chosen to be 2, beyond the right-hand endpoint of the interval $[-1, 1]$ – we know that $\mathcal{L}_n(1, 2)$ is at least $-t/3$ except on an event of probability $Ce^{-ct^{3/2}}$, because this one-point lower tail bound is addressed in the definition of a regular ensemble.

Consider what happens when the BG resampling associated to $\{1\} \times [X, 1]$ occurs in the presence of the event $\text{High}(t) \cap \{\mathcal{L}_n(1, 2) \geq -t/3\}$. In this case, $X < 1$. Consider the value of the resampled curve at the point 1. The resample has the law of Brownian bridge on $[X, 1]$ with left endpoint at least t ; with right endpoint at least $-t/3$; and with lower boundary condition provided by the second ensemble curve on $[X, 2]$. If we remove this second curve, and replace the left endpoint by the value t , and the right endpoint by the value $-t/3$, we only decrease the probability that the resampled curve at 1 exceeds any given level – this is the content of the two monotonicity results treated in *Basics A*. But the law resulting from these alterations could not be simpler – it is Brownian bridge on the interval $[X, 1]$ with left value t and right value $-t/3$. This process has probability at least one-half to assume a value greater than the affine interpolation of those end values; and, since $X \geq -1$, the latter value is at least $t/9$.

In summary, the resampled process' value at 1 exceeds $t/9$ with probability at least

$$\frac{1}{2} \cdot \mathbb{P}\left(\text{High}(t) \cap \{\mathcal{L}_n(1, 2) \geq -t/3\}\right).$$

But the random resampling maintains the law of the ensemble – and since the ensemble is regular, we have a bound on how rare this outcome is:

$$\mathbb{P}\left(\mathcal{L}_n(1, 1) \geq t/9\right) \leq C \exp\{-c(t/9)^{3/2}\}.$$

Thus we find that

$$\mathbb{P}\left(\text{High}(t)\right) \leq \mathbb{P}\left(\mathcal{L}_n(1, 2) < -t/3\right) + 2C \exp\{-c(t/9)^{3/2}\} \leq \Theta(1) \exp\{-\Theta(1)(t/9)^{3/2}\}.$$

This is a consequence of the sought form (6). □

The last proof is a simple demonstration of BG resampling – it shows how exceptional behaviour, such as a very high value for the top curve in a BG ensemble, can become typical, with in this case the high value spreading under resampling to occur on a unit-order interval.

D: Near parabolic invariance of regular ensembles

This is a simple property about which we comment only briefly. A regular ensemble's top curves on a compact interval about the origin are locally Brownian, and, on a slightly larger scale, they are globally parabolic – they hew to the parabola $-Q$, where $Q(x) = 2^{-1/2}x^2$. Suppose instead that we view the top curves on a compact interval centred instead at a point x which satisfies $|x| \leq c/2 \cdot n^{1/9}$. The top curves are rapidly ascending if x is negative (and large, subject to its constraint); or rapidly descending if x is positive and large. Change coordinates by adding a suitable linear function to flatten the curves in this neighbourhood – and then shift the picture so that the top curves in this neighbourhood are now above the origin at roughly unit height. It is a basic consequence of the definition of a regular ensemble that the resulting ensemble remains regular – with a small

adjustment, in the values of the parameters c and C . That is, the top curves are locally Brownian, and hew again to the parabola $-Q(x)$ on large scales, up to horizontal scale $n^{1/9}$.

The role of this near parabolic invariance is to permit us to take a result proved near the origin, such as the No Big Max Proposition 2.17 – and to extend its domain of validity to a much wider region of x -coordinates – to these x at or below scale $n^{1/9}$.

E: Control on lower curves

This is our final basic BG input. The definition of regular ensemble insists on the BG property and one-point control on the top curve – nothing else. Using the BG property, we work to propagate understanding of one-point behaviour to lower curves in the ensemble. By induction, for increasing $k \in \mathbb{N}$, we may thus prove:

$$\mathbb{P}\left(\mathcal{L}_n(k, 0) \leq -s\right) \leq C_k e^{-c_k s^{3/2}}.$$

We will not explain how this induction works – but it is a BG resample argument. We will need this lower curve control in the intermediate and advanced steps, as we progress towards the proof of the Brownian bridge regularity Theorem 2.13. The recorded bound could well be also available from determinantal or other integrable techniques but we have not found it in the literature.

CHAPTER 3

Lecture Three: Local fluctuation results via the Wiener candidate

This lecture is devoted to proving the Local Fluctuation Theorem 2.16. This is the second, *intermediate*, stage in the three-stage journey to the principal Brownian bridge regularity Theorem 2.13. The intermediate step involves introducing a new resampling of an ensemble, one that is a little more subtle than the basic BG resampling – the objects we introduce, which will be used again in the final lecture, are *missing closed middle reconstruction* and the associated *Wiener candidate*.

First a note about the construction of the Brownian bridge law $\mathcal{B}_{1;x,y}^{[a,b]}$, whose sample $B : [a, b] \rightarrow \mathbb{R}$, $B(a) = x$, $B(b) = y$, we here call simply B , rather than $B(1, \cdot)$.

Let $[u, v] \subseteq [a, b]$. The process B may be constructed by

- (1) sampling its pair value $(B(u), B(v))$ according to the density on \mathbb{R}^2 whose value at (z_1, z_2) is, up to a factor of normalization, equal to

$$\exp \left\{ -\frac{(x-z_1)^2}{2(u-a)} - \frac{(z_2-z_1)^2}{2(v-u)} - \frac{(z_2-y)^2}{2(b-v)} \right\};$$

- (2) constructing three independent standard Brownian bridges on the intervals $[a, u]$, $[u, v]$ and $[v, b]$;
- (3) and then forming $B : [a, b] \rightarrow \mathbb{R}$ by adding the function that interpolates the values of B on the four-point set $\{a, u, v, b\}$ on which it is already constructed and the sum of the three sampled standard bridges.

Very similarly, we may represent the process B by a list of data

- its values $(B(u), B(v))$;
- and the three standard bridges $B^{[a,u]}$, $B^{[u,v]}$ and $B^{[v,b]}$.

We now split this data in different way, into two pieces of data, which – for reasons to be explained momentarily – we call the *retained* and the *lost* data. The retained data is

- the bridges $B^{[a,u]}$ and $B^{[v,b]}$.

The lost data is then presented in a *two-piece list* in the form:

- the values $(B(u), B(v))$;
- and the bridge $B^{[u,v]}$.

We may equally present the lost data in a *one-piece list*:

- the marginal $B : [u, v] \rightarrow \mathbb{R}$.

Why these names? Suppose that an experimenter realizes the process B according to the law $\mathcal{B}_{1;x,y}^{[a,b]}$, and records the outcome in the form of retained and lost data. Suppose then that the experimenter

discards the lost data, and hands only the retained data to an observer, who is aware that B has been sampled according to $\mathcal{B}_{1;x,y}^{[a,b]}$. The observer has all information except the lost data.

Let \mathcal{F} denote the σ -algebra of information to which the observer is privy – the σ -algebra generated by the retained data. Let $\mathbb{P}_{\mathcal{F}}$ denote conditional probability given \mathcal{F} . That is, $\mathbb{P}_{\mathcal{F}}(A) = \mathbb{E}[\mathbf{1}_A | \mathcal{F}]$.

We will view the observer's perspective in view of an attempt that the observer may make to reconstruct the process B given the available, retained, data. The law \mathcal{F} represents this perspective. The observer may choose to describe the conditional distribution of what is unknown by recording a random process of lost data in either one-piece or two-piece list form.

If two-piece list form is adopted, then the unknown data has the form of a point-pair value at u and v ; and a standard bridge on $[u, v]$. The sampling of these data may be viewed by placing a pair of beads on vertical rods at x -coordinates u and v at heights respectively dictated by the sampled point-pair value; adding the standard bridge on $[u, v]$ to the affine interpolation on $[u, v]$ of the bead locations; and affinely translating the two side bridges to meet their fixed endpoints at a and at b and the bead locations.

In the one-piece list perspective, the observer reconstructs the lost data – the marginal process $B : [u, v] \rightarrow \mathbb{R}$ – by sampling $W : [u, v] \rightarrow \mathbb{R}$ as the marginal on $[u, v]$ of a realization of the bridge law $\mathcal{B}_{1;x,y}^{[a,b]}$ that is sampled independently of the retained data; then the observer sets B on the side intervals $[a, u]$ and $[v, b]$ by affinely translating the retained side bridges in accordance with their fixed endpoints and the variable ones specified by the already determined values $B(u)$ and $B(v)$.

3.1. Reprising this for real – with a regular ensemble

We will work with the ensemble $\mathcal{L} : \mathbb{N} \times \mathbb{R} \rightarrow \mathbb{R}$ for simplicity. Missing closed middle reconstruction is the reconstruction of \mathcal{L} which is counterpart to the Brownian bridge case just discussed.

The reconstruction has three parameters $T > 0$, $\ell \in [-T, 0]$ and $r \in [0, T]$. The interval $[-2T, 2T]$ is partitioned

$$[-2T, 2T] = [-2T, \ell] \cup [\ell, r] \cup [r, 2T]$$

The left interval, counterpart to $[a, b]$ is partitioned into a *middle* interval $[\ell, r]$, counterpart to $[u, v]$, and two *side* intervals: $[-2T, \ell]$ on the left and $[r, 2T]$ on the right.

We now record a presentation of the ensemble $\mathcal{L} : \mathbb{N} \times \mathbb{R} \rightarrow \mathbb{R}$ into retained and lost data, in which the lost data is recorded in either one-piece or two-piece list form.

The retained data is:

- the top curve $\mathcal{L}(1, \cdot)$ on $\mathbb{R} \setminus [-2T, 2T]$;
- all lower curve data $\mathcal{L}(k, x)$ for $k \geq 2$ and $x \in \mathbb{R}$;
- and the side bridges $\mathcal{L}^{[-2T, \ell]}(1, \cdot)$ and $\mathcal{L}^{[r, 2T]}(1, \cdot)$.

In its two-piece list form, the lost data is

- the middle bridge $\mathcal{L}^{[\ell, r]}(1, \cdot)$;
- and the pair values $(\mathcal{L}(1, \ell), \mathcal{L}(1, r))$.

In its one-piece list form, the same data is

- the process $\mathcal{L}(1, \cdot) : [\ell, r] \rightarrow \mathbb{R}$.

Consider then the perspective of the observer who is presented merely with the retained data. We again write \mathcal{F} for the σ -algebra of the retained data to which the observer has access; and write $\mathbb{P}_{\mathcal{F}}$ for conditional probability given \mathcal{F} .

Viewing \mathcal{L} from the observer's perspective. The observer's point of view is depicted in Figure 3.1. Here, the observer realizes the law $\mathcal{B}_{1;u,v}^{[-2T,2T]}$ with $u = \mathcal{L}(1, -2T)$ and $v = \mathcal{L}(1, 2T)$, where note that these (u, v) values are indeed known to the observer. The Wiener candidate W is defined to be the marginal on $[\ell, r]$ of the realized process. Using the bead locations thus dictated, the two side interval process values

$$\mathcal{L}(1, \cdot) : [-2T, \ell] \cup [r, 2T] \rightarrow \mathbb{R}$$

are obtained by affine translation as discussed before.

However, the proposed outcome must be checked for avoidance constraints between the top and the second curve on the side intervals and on the middle interval. In this way, there are *three tests* that must be passed: these are associated to, and named after, the *left side* $[-2T, \ell]$, the *middle interval* $[\ell, r]$, and the *right side* $[r, 2T]$.

This then is *missing closed middle reconstruction*, so named because the observer lacks knowledge of $\mathcal{L}(1, \cdot) : [\ell, r] \rightarrow \mathbb{R}$ throughout the middle interval $[\ell, r]$; so that the bead locations $\mathcal{L}(1, \ell)$ and $\mathcal{L}(1, r)$ are *random* for the observer.

We now use this technique to prove the Local Fluctuation Theorem 2.16: namely,

$$\mathbb{P}\left(\sup_{x \in (0, \varepsilon)} |\mathcal{L}(1, x+h) - \mathcal{L}(1, x)| \geq K\varepsilon^{1/2}\right) \leq Ce^{-cK^{3/2}}.$$

First consult Figure 3.2 for an explanation of the form of the left and right side tests, including the role of two \mathcal{F} -measurable real random variables `LeftCorner` and `RightCorner`.

We make use a *favourable* event F . This is an \mathcal{F} -measurable event – whose occurrence or otherwise is known to the observer – on which the observer's perspective is sufficiently pleasant to admit analysis of the application in question – Theorem 2.16 at present.

For this application of missing closed middle reconstruction, we make the parameter settings $T = 1$, $\ell = -1$ and $r = 1$. For now, our specification of the event $F = F_t$ will depend on a parameter $t > 0$ – to be specified shortly. Indeed, for $t > 0$, we specify $F = F_t$ to be the event that

- (1) $\mathcal{L}(1, -2)$ and $\mathcal{L}(1, 2)$ belong to $[-t, t]$;
- (2) `LeftCorner` and `RightCorner` belong to $[-t, t]$;
- (3) and $\mathcal{L}(2, x) \leq t$ for all $x \in [-1, 1]$.

LEMMA 3.1 (The favourable event is typical).

$$\mathbb{P}(F_t^c) \leq C_1 e^{-c_1 t^{3/2}}.$$

Proof. Property (1) is one-point upper and lower tail bounds. Regarding property (2), note that

$$\mathcal{L}(2, -1) \leq \text{LeftCorner} \leq \mathcal{L}(1, -1)$$

and

$$\mathcal{L}(2, 1) \leq \text{RightCorner} \leq \mathcal{L}(1, 1)$$

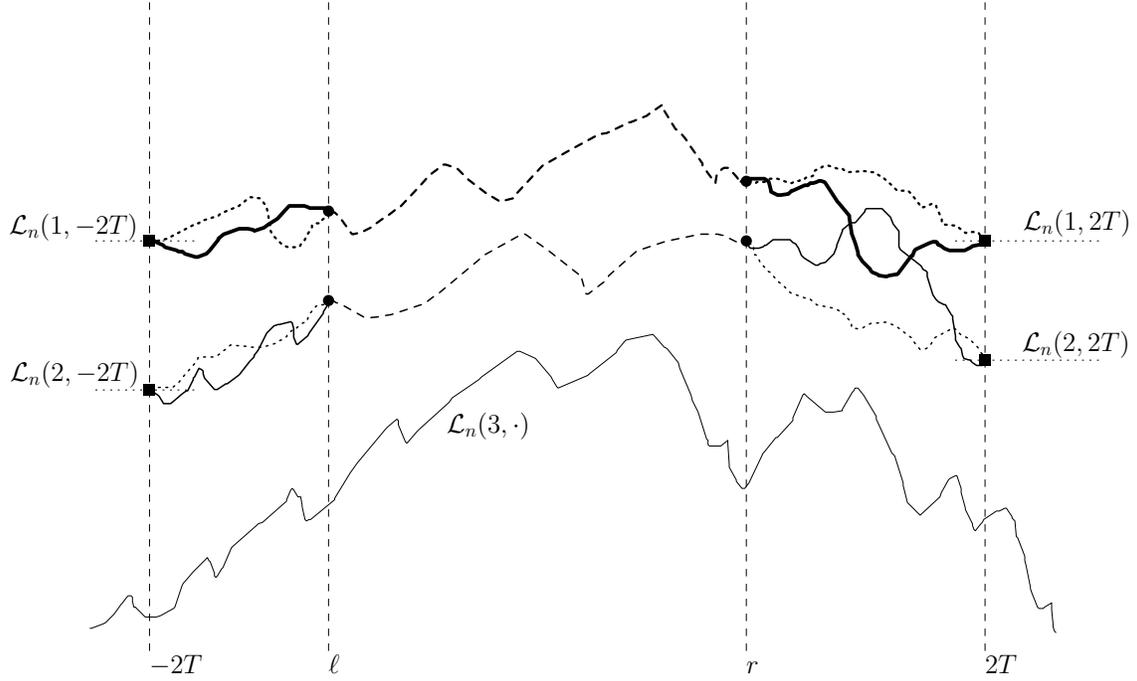


FIGURE 3.1. The perspective of the observer of \mathcal{F} and the construction of the Wiener candidate are depicted for the case of two curves – recall however that we merely consider the top curve in these lectures for reasons of convenience. The two curves on $[-2T, 2T]$ that are dotted, dashed and then dotted again are samples of the law $\mathcal{B}_{2; \bar{u}, \bar{v}}^{[-2T, 2T]}$, with $(u_1, u_2) = (\mathcal{L}_n(1, -2T), \mathcal{L}_n(2, -2T))$ and $(v_1, v_2) = (\mathcal{L}_n(1, 2T), \mathcal{L}_n(2, 2T))$. The marginal of these curves on $[\ell, r]$, i.e., their dashed sections, forms the Wiener candidate W . The thicker solid black curve on $[-2T, \ell]$ is the affine translation of $\mathcal{L}_n^{[-2T, \ell]}$ with left endpoint $\mathcal{L}_n(1, -2T)$ and right endpoint $W(1, \ell)$. Similarly on the right, and for the thinner solid black second curve. The black beads on the vertical line with x -coordinate ℓ , located at heights $(W(1, \ell), W(2, \ell))$, thus dictate the form of the black curves to the left via affine translation subject to the fixed left endpoints; and similarly of course on the right. The left side and middle interval tests are passed but the right side test is failed by the Wiener candidate, so that the candidate is unsuccessful in this instance.

The right-hand terms here are not \mathcal{F} -measurable – but that doesn't matter. We see from the displays that (2) reduces to one-point upper and lower tail bounds.

Property (3) is handled via ensemble ordering and No Big Max. □

Now for the proof of Theorem 2.16 – using the favourable event F_t , and setting t in terms of K .

Recall that, under $\mathbb{P}_{\mathcal{F}}$, the Wiener candidate $W : [-1, 1] = [\ell, r] \rightarrow \mathbb{R}$ has the marginal law on $[-1, 1]$ of $\mathcal{B}_{1; u, v}^{[-2, 2]}$ with $u = \mathcal{L}(1, -2)$ and $v = \mathcal{L}(1, 2)$.

What are the prospects for success of the Wiener candidate?

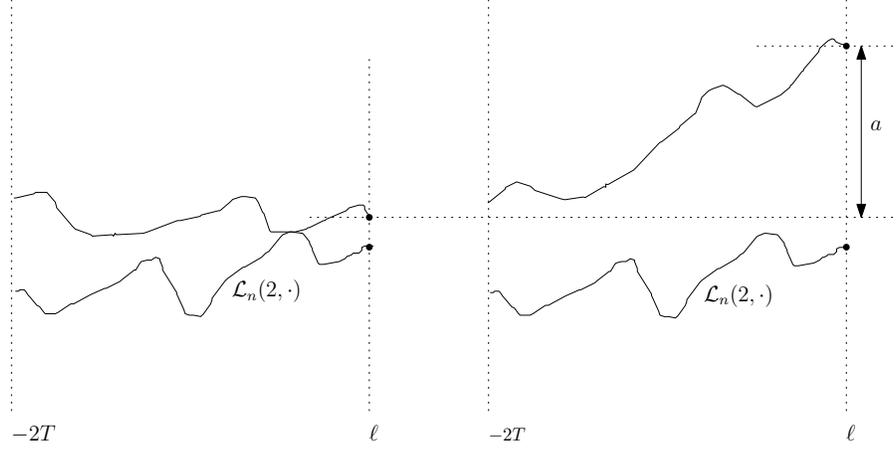


FIGURE 3.2. In the left sketch, the observer has brought the upper bead over ℓ down as far as possible compatibly with passing the left side test. The curve which ends at this bead touches the second curve on $[-2T, \ell]$. This curve is being affinely translated in accordance with the bead location and with the curve's fixed left endpoint – see the right sketch for a depiction of the outcome when the bead is raised by distance $a > 0$. The minimum location for the bead – the level of the horizontal dotted line – is defined to be **LeftCorner**. Admissible locations for the Wiener candidate W at ℓ are those at or above **LeftCorner**, and the conditional distribution $W(\ell)$ given that W passes the left-side test is given by the Gaussian law of $W(\ell)$, conditioned on $W(\ell) \geq \text{LeftCorner}$. There is a similar story for the right-side test – and a counterpart minimum level **RightCorner** for $W(r)$ so that the Wiener candidate does not violate the right-side test.

LEMMA 3.2 (Wiener candidate success lower bound).

$$\mathbb{P}_{\mathcal{F}}(\text{the Wiener candidate succeeds}) \geq \exp\{-O(1)t^2\} \cdot \mathbf{1}_{F_t}.$$

Proof. Consider the event that, under $\mathbb{P}_{\mathcal{F}}$,

$$t \leq W(x) \leq 2t \quad \forall x \in [-1, 1]. \quad (7)$$

Any such W is successful (provided that F_t occurs) because

$$W(x) \geq t > \mathcal{L}(2, x) \quad \forall x \in [-1, 1]$$

– so the middle interval test is passed;

$$W(-1) \geq t \geq \text{LeftCorner}$$

– so the left side test is passed; and

$$W(1) \geq t \geq \text{RightCorner}$$

– so the right side test is passed.

Q1: What is the $\mathbb{P}_{\mathcal{F}}$ -probability of the event in (7), given any \mathcal{F} -data that verifies F_t ?

A1: At least $e^{-C_1 t^2}$ with $C_1 \approx 20$.

Why this answer? See Figure 3.3 and its caption.

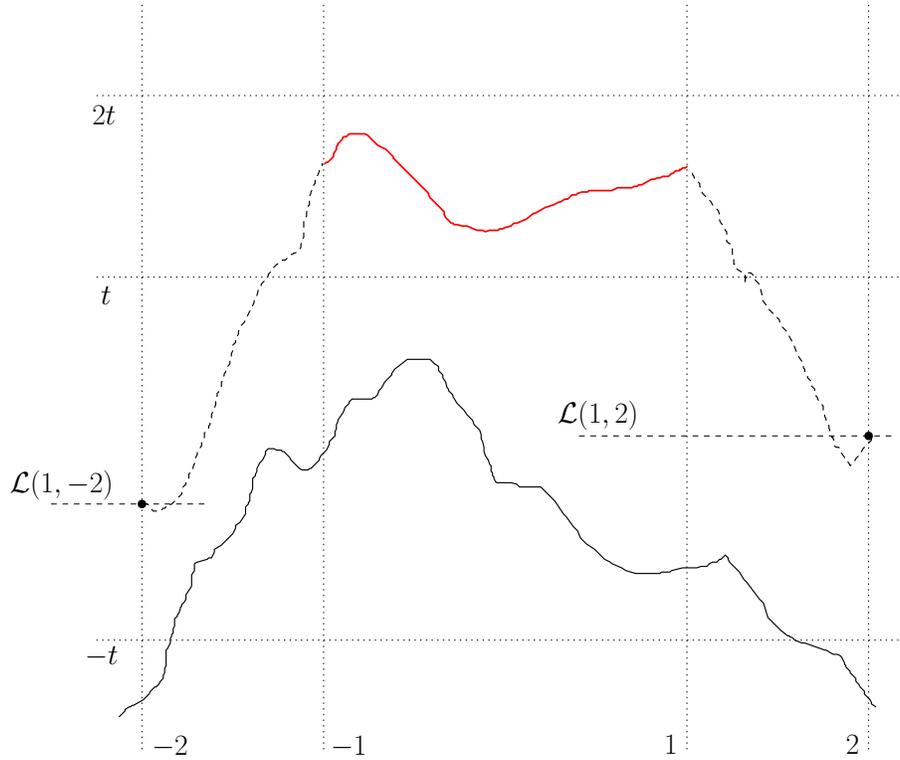


FIGURE 3.3. The red Wiener candidate realizes $W(x) \in [t, 2t]$ throughout $x \in [-1, 1]$ in the depiction. The Gaussian probability of its endpoint values $W(-1)$ and $W(1)$ lying in the middle-third of $[t, 2t]$ are at least of the form $e^{-(3t)^2}$ due to the lower bounds on $\mathcal{L}(1, -2)$ and $\mathcal{L}(1, 2)$. The need for these endpoints to exceed `LeftCorner` and `RightCorner` only helps in achieving this outcome for $W(-1)$ and $W(1)$. If it is achieved, then with positive conditional probability the red Wiener candidate stays in the channel $[t, 2t]$ during $x \in [-1, 1]$. Thus the candidate realizes $W(x) \in [t, 2t]$ for $x \in [-1, 1]$ with probability at least $e^{-O(1)T^2}$ whenever F_t occurs.

Q2: What is the Wiener candidate's chance of making a big fluctuation on $[0, \varepsilon]$?

A2: The answer is the next bound:

$$\mathbb{P}_{\mathcal{F}}\left(\sup_{x \in (0, \varepsilon)} |W(x + \varepsilon) - W(x)| \geq K\varepsilon^{1/2}\right) \cdot \mathbf{1}_{F_t} \leq C_2 e^{-c_1 K^2}$$

when $t \leq O(k)$.

Why does the bound hold?

See Figure 3.4. The movement of W on $[0, \varepsilon]$ takes the form

$$O(t)\varepsilon + \varepsilon^{1/2}G,$$

where G is a standard Gaussian random variable. With $t \leq K$ and $K\varepsilon \leq \varepsilon^{1/2}$ (which is equivalent to the harmlessly supposed upper bound on ε , $\varepsilon \leq K^{-2}$), the movement in question is at most $2\varepsilon^{1/2}G$. Thus we confirm the given answer.

We are ready to put these elements together.

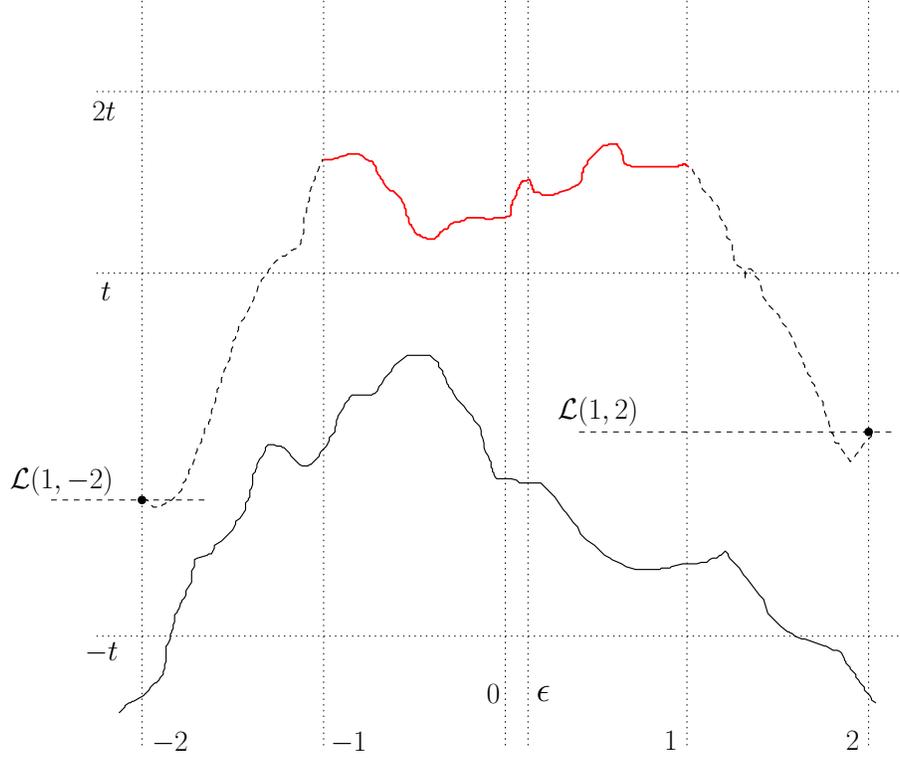


FIGURE 3.4. The Wiener candidate is in essence Brownian motion of drift of order at most t – and this permits us to bound above the probability of that it experiences a big fluctuation on the interval $[0, \varepsilon]$.

We choose $t = O(1)K$. For a random process X defined on an interval that includes $[0, \varepsilon]$, we write

$$\text{BigFluc}(X) = \left\{ \sup_{x \in (0, \varepsilon)} |W(x + \varepsilon) - W(x)| \geq K\varepsilon^{1/2} \right\}.$$

Note then that

$$\mathbb{P}(\text{BigFluc}(\mathcal{L})) \leq \mathbb{P}(\text{BigFluc}(\mathcal{L}) \cap F_t) + \mathbb{P}(F_t^c).$$

Now,

$$\mathbb{P}(\text{BigFluc}(\mathcal{L}) \cap F_t)$$

is equal to

$$\mathbb{E} \left[\mathbb{P}_{\mathcal{F}}(\text{BigFluc}(\mathcal{L})) \cdot \mathbf{1}_{F_t} \right].$$

Since

$$\mathbb{P}_{\mathcal{F}}(\text{BigFluc}(\mathcal{L})) = \mathbb{P}_{\mathcal{F}}(\text{BigFluc}(W) | W \text{ succeeds}),$$

the last expression is at most

$$\mathbb{E} \left[\frac{\mathbb{P}_{\mathcal{F}}(\text{BigFluc}(W))}{\mathbb{P}_{\mathcal{F}}(W \text{ succeeds})} \cdot \mathbf{1}_{F_t} \right].$$

and thus also at most

$$\mathbb{E} \left[\frac{C_2 e^{-c_1 K^2}}{e^{-C_1 O(K^2)}} \cdot \mathbf{1}_{F_t} \right].$$

in view of Answers A1 and A2 as well as the choice $t = O(K)$. Indeed, selecting $t = O(1)K$ with the $O(1)$ term of unit order but small enough, we find that

$$\mathbb{P}(\text{BigFluc}(\mathcal{L}) \cap F_t) \leq e^{-c_3 K^2}$$

with $c_3 > 0$. We insist that $t \geq cK$ because we further use Lemma 3.1:

$$\mathbb{P}(F_t^c) \leq C_1 e^{-c_1 t^{3/2}}.$$

Thus,

$$\mathbb{P}(\text{BigFluc}(\mathcal{L})) \leq O(1)e^{-O(1)t^2} + O(1)e^{-O(1)t^{3/2}};$$

and Theorem 2.16, i.e.,

$$\mathbb{P}\left(\sup_{x \in (0, \varepsilon)} |\mathcal{L}(1, x + \varepsilon) - \mathcal{L}(1, x)| \geq K\varepsilon^{1/2}\right) \leq C e^{-cK^{3/2}}$$

is proved. □

These intermediate tools – missing closed middle reconstruction and the Wiener candidate – may also be used to give a fairly direct proof of the key step in the construction of the Airy line ensemble – Proposition 2.11, the uniform lower bound on the acceptance probability; this is forthcoming work with Jacob Calvert and Milind Hegde.

CHAPTER 4

**Lecture Four: proving Brownian bridge regularity
via the jump ensemble**

Here we prove Theorem 2.13 in the special case of the ensemble $\mathcal{L} : \mathbb{N} \times \mathbb{R} \rightarrow \mathbb{R}$. For the purpose of recall, and because the assertion slightly simplifies in this special case, we state the relevant result.

THEOREM 4.1 (Brownian bridge regularity). *Let $a \in (0, 1)$ denote $\mathcal{B}_{1;0,0}^{[0,1]}(A)$ where $A \subseteq \mathcal{C}_{0,0}([0, 1], \mathbb{R})$ is measurable. Then*

$$\mathbb{P}\left(\mathcal{L}^{[0,1]}(1, \cdot) \in A\right) \leq a \cdot \exp\left\{\left(\log a^{-1}\right)^{5/6} O(1)\right\}.$$

Note that we are attempting a proof of this result only for curve index $k = 1$ – working in this case is not much more than a notational convenience. We work merely with the spatial interval $[0, 1]$ because the stationarity of the Airy line ensemble permits this reduction – in the general case of a regular ensemble, it is near parabolic invariance (Basics *D*) which performs the corresponding role.

In crude summary of this statement, then, we are aiming to show a result of the form:

$$\begin{aligned} &\text{an event whose Brownian bridge probability is } \varepsilon && (8) \\ &\text{has probability for the top curve in } \mathcal{L} \text{ which is at most } \varepsilon^{1-o(1)}. \end{aligned}$$

We now reserve the symbol $\varepsilon > 0$ to denote the small probability of a Brownian bridge $\mathcal{B}_{1;0,0}^{[0,1]}$ event A .

First, we review our prospects of proving an implication of the form (8) using our present, intermediate, apparatus: MCM reconstruction and the Wiener candidate.

We have specified a favourable event F_t , depending on a parameter $t > 0$. We will choose t in terms of $\varepsilon > 0$.

The basic inequality of the intermediate method asserts that, for an event A with $\mathcal{B}_{1;0,0}^{[0,1]}(A) = \varepsilon$,

$$\mathbb{P}\left(\mathcal{L}^{[0,1]}(1, \cdot) \in A\right) \leq \mathbb{E}\left[\mathbb{P}_{\mathcal{F}}\left(W^{[0,1]}(\cdot) \in A \mid W \text{ succeeds}\right) \cdot \mathbf{1}_{F_t}\right] + \mathbb{P}(F_t^c);$$

or

$$\mathbb{P}\left(\mathcal{L}^{[0,1]}(1, \cdot) \in A\right) \leq \mathbb{E}\left[\frac{\mathbb{P}_{\mathcal{F}}\left(W^{[0,1]}(\cdot) \in A\right)}{\mathbb{P}_{\mathcal{F}}\left(W \text{ succeeds}\right)} \cdot \mathbf{1}_{F_t}\right] + \mathbb{P}(F_t^c).$$

The right-hand numerator $\mathbb{P}_{\mathcal{F}}(W^{[0,1]}(\cdot) \in A)$ equals $\mathcal{B}_{1;0,0}^{[0,1]}(A) = \varepsilon$. In order to achieve

$$\mathbb{P}\left(\mathcal{L}^{[0,1]}(1, \cdot) \in A\right) \leq \varepsilon^{1-o(1)},$$

we thus need two things:

(1) *plausible success of the Wiener candidate*, i.e., on F_t that

$$\mathbb{P}_{\mathcal{F}}(W \text{ succeeds}) \geq \varepsilon^{o(1)};$$

(2) and *rarity of failure of the favourable event*, i.e.,

$$\mathbb{P}(F_t^c) \leq \varepsilon^{1-o(1)}.$$

Note then that (1) is an upper bound on t , since $\mathbb{P}_{\mathcal{F}}(W \text{ succeeds}) \geq e^{-O(1)t^2}$ on F_t ; and (2) is a lower bound on t , because $\mathbb{P}(F_t^c) \leq e^{-O(1)t^{3/2}}$.

Indeed, for these reasons, (1) forces $t \leq o(1)(\log \varepsilon^{-1})^{1/2}$, while (2) forces $t \geq \Theta(1)(\log \varepsilon^{-1})^{3/2}$.

These incompatible choices show that the present method cannot work – at least not with its present parameter choices $T = 1$, $\ell = -1$ and $r = 1$.

Our first recourse to remedy this trouble is to vary the parameter settings. Suppose that we try

$$T = (\log \varepsilon^{-1})^\alpha > 0$$

with ℓ and $-r$ equal to, or in any case of the order of, T . Here, $\alpha > 0$ is a parameter which we are at liberty to vary in seeking a better outcome for the method.

The new advantage is that the Wiener candidate, who needs to jump over a hill of height $O((\log \varepsilon^{-1})^{2/3})$, – a height whose order is dictated by the consideration (2) above, where the favourable event F_t is specified by a natural variation now that the parameter choices (T, ℓ, r) may no longer equal $(1, -1, 1)$ – now has a duration $(\log \varepsilon^{-1})^\alpha$ in which to make the jump, rather than merely unit-order time. However, the height of the hill has increased, because the attempted high jump begins at lower locations than before, in view of parabolic curvature. See Figure 4.1. Indeed, $\mathcal{L}(1, \pm 2T) \asymp -(\log \varepsilon^{-1})^{2\alpha}$.

Consideration (2) continues to dictate a choice of t that satisfies

$$t \geq \Theta(1)(\log \varepsilon^{-1})^{2/3};$$

and, in fact, we may choose equality here.

Regarding (1), the Wiener success probability on F_t has the form

$$\exp \left\{ - \frac{1}{(\log \varepsilon^{-1})^\alpha} \left((\log \varepsilon^{-1})^{2/3} + (\log \varepsilon^{-1})^{2\alpha} \right)^2 O(1) \right\} = \exp \left\{ - (\log \varepsilon^{-1})^{(4/3-\alpha)\vee 3\alpha} O(1) \right\},$$

where inside the first exponential, the term in parentheses in the numerator is the hill height and the denominator is the duration in which the jump may be completed. This expression is maximized when $4/3 - \alpha = 3\alpha$, in which case, it takes the form

$$\exp \left\{ - O(1)(\log \varepsilon^{-1}) \right\} = \varepsilon^{O(1)}.$$

We have come close to, *but cannot achieve*, the desired form $\varepsilon^{o(1)}$. What we instead want – what would clearly suffice – is an expression of the form $\exp \{ (\log \varepsilon^{-1})^{1-\zeta} \}$ for some $\zeta > 0$. We refer to our failure to obtain it – we merely obtained $\zeta = 0$ – as the *high jump* difficulty.

Set back by this difficulty we may seem to be, we have nonetheless learnt a sensible way MCM parameter $T > 0$: henceforth, we set

$$T = D(\log \varepsilon^{-1})^{1/3},$$

where $D > 0$ is a constant to be chosen suitably in applications.

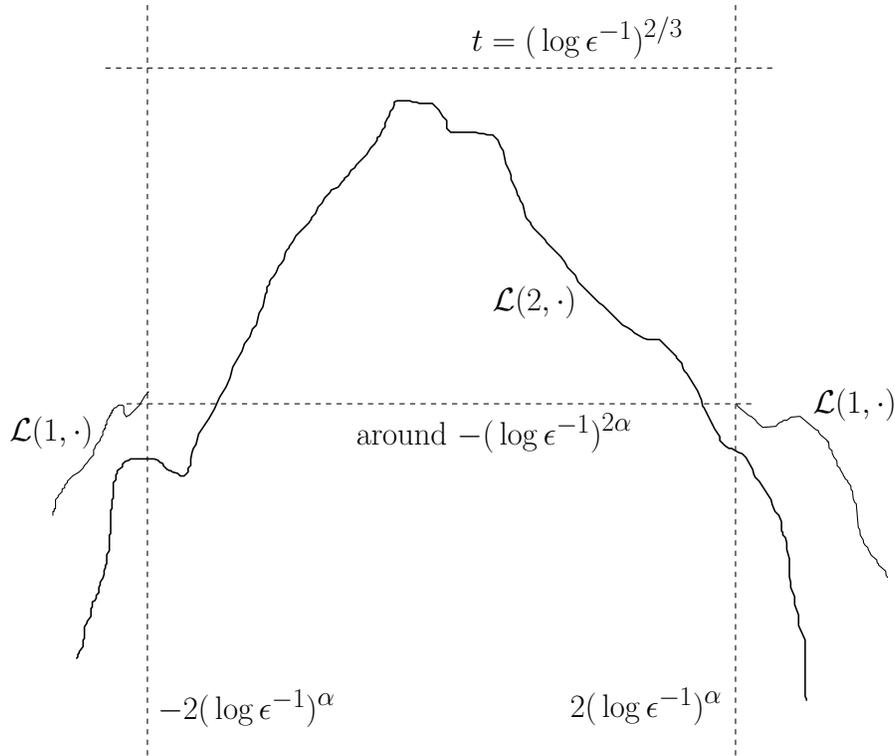


FIGURE 4.1. With $T = -(\log \varepsilon^{-1})^\alpha$, the Wiener candidate has duration of order $(\log \varepsilon^{-1})^\alpha$ to jump over a hill of height $(\log \varepsilon^{-1})^{2/3}$, but must begin and end its journey in a valley below the base of the hill by order $(\log \varepsilon^{-1})^{2\alpha}$.

We continue to work with a favourable event $\mathbf{Fav} = F_t$, with $t = \Theta(1)(\log \varepsilon^{-1})^{2/3}$, with F_t now specified analogously to before for the new choice of T . See Figure 4.2. The upper bound on the probability of failure of the favourable event now takes the form:

LEMMA 4.2.

$$\mathbb{P}(\mathbf{Fav}^c) \leq \varepsilon^{O(1)D^3}.$$

Next, we set the parameters ℓ and r . The obvious respective choices are $-T$ and T , but we will make a choice $\ell \in [-T, 0]$ and $r \in [0, T]$ that renders the lower boundary condition $\mathcal{L}(2, \cdot)$ slightly more regular on $[\ell, r]$.

To specify our choice of (ℓ, r) , we introduce the *least concave majorant* $\mathfrak{c}_+ : [-T, T] \rightarrow \mathbb{R}$ of the curve $\mathcal{L}_n(2, \cdot) : [-T, T] \rightarrow \mathbb{R}$.

Define a random variable pair (ℓ, r) according to

$$\begin{aligned} \ell &= \inf \{x \in [-T, T] : \mathfrak{c}'_+(x) \leq 4T\} \\ \text{and } r &= \sup \{x \in [-T, T] : \mathfrak{c}'_+(x) \geq -4T\}, \end{aligned}$$

where the convention that $\inf \emptyset = T$ and $\sup \emptyset = -T$ is adopted. That is, $\ell \in [-T, T]$ is the leftmost location at which \mathfrak{c}_+ has slope at most $4T$, and r is the rightmost at which this slope is at least $-4T$.

It is easily seen that the occurrence of \mathbf{Fav} forces $\ell \in [-T, -T/2]$ and $r \in [T/2, T]$.

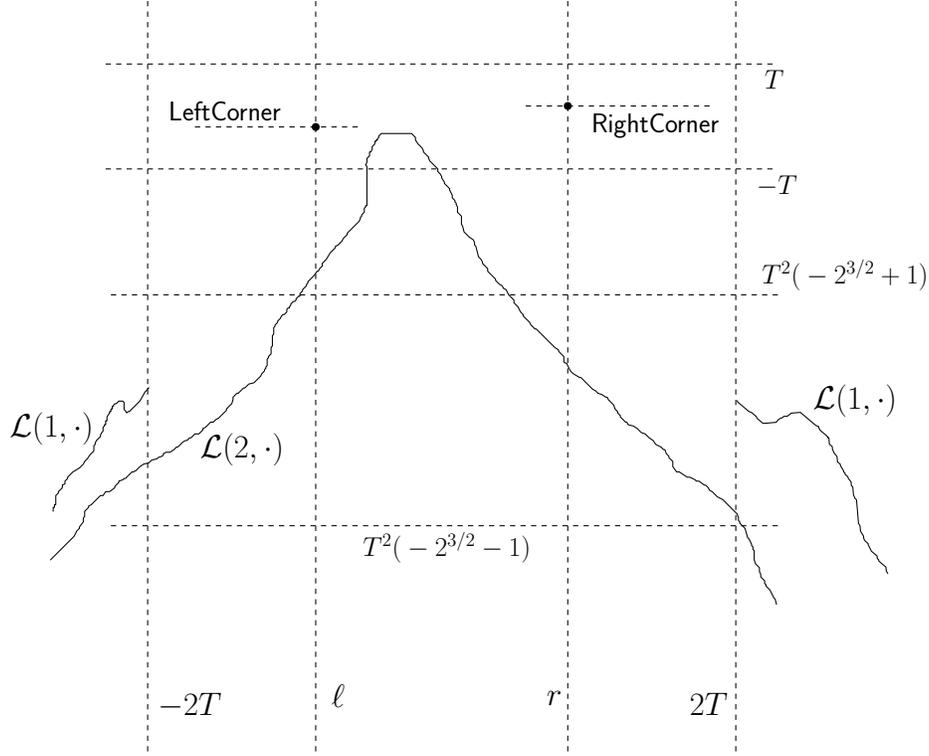


FIGURE 4.2. Depicted is an example of a realization of the new favourable event Fav. The conditions that collectively specify this event are that LeftCorner and RightCorner belong to $[-T, T]$; that $\mathcal{L}(2, x) \leq T$ for all $x \in [\ell, r]$; and that $\mathcal{L}(1, -2T)$ and $\mathcal{L}(1, 2T)$ differ from the parabolically determined level $-2^{3/2}T^2$ by at most T^2 .

4.1. More promising than the Wiener candidate – the jump ensemble

Consider the law $\mathbb{P}_{\mathcal{F}}$ for \mathcal{F} -data that realizes Fav. We have seen that Wiener candidate success occurs with probability $\varepsilon^{O(1)}$ – but we need $\varepsilon^{o(1)}$. That is, we need to solve the high jump difficulty.

We will vary the Wiener candidate $W : [\ell, r] \rightarrow \mathbb{R}$ to obtain the jump curve $J : [\ell, r] \rightarrow \mathbb{R}$. The success probability of J will be at least $\varepsilon^{o(1)}$ for \mathcal{F} -data verifying Fav as desired.

The Wiener candidate finds it difficult to pass the middle interval test – that it exceeds $\mathcal{L}(2, \cdot)$ on $[\ell, r]$.

Idea: make this easier by first conditioning the Wiener candidate on exceeding a coarse-grained caricature of $\mathcal{L}(2, \cdot)$ on $[\ell, r]$.

Recall that $\mathbf{c}_+ : [-T, T] \rightarrow \mathbb{R}$ is the least concave majorant of the curve $\mathcal{L}_n(k+1, \cdot) : [-T, T] \rightarrow \mathbb{R}$. Let $\text{xExt}(\mathbf{c}_+) \subset [\ell, r]$ denote the set of x -coordinates of extreme points of the closed set $\{(x, y) : \ell \leq x \leq r, y \leq \mathbf{c}_+(x)\}$. Note that $\text{xExt}(\mathbf{c}_+)$ consists of the intersection with $[\ell, r]$ of the set of points of local non-constancy of \mathbf{c}'_+ ; necessarily, $\{\ell, r\} \in \text{xExt}(\mathbf{c}_+)$. Let P denote a subset of $\text{xExt}(\mathbf{c}_+)$ with the properties that

- $\{\ell, r\} \in P$;
- any distinct elements $p_1, p_2 \in P$ satisfy $|p_1 - p_2| > 1$;

- and, if $x \in \text{xExt}(\mathbf{c}_+) \setminus P$, then some element $p \in P$ satisfies $|p - x| \leq 1$.

We have $|P| \leq 2T$ since $[\ell, r] \subseteq [-T, T]$.

Why the ‘pole’ set? The upcoming Figure 4.4 explains.

Recall that, under $\mathbb{P}_{\mathcal{F}}$, the Wiener candidate $W : [\ell, r] \rightarrow \mathbb{R}$ is the marginal on $[\ell, r]$ of $\mathcal{B}_{k;u,v}^{[-2T,2T]}$, where $u = \mathcal{L}(1, -2T)$ and $v = \mathcal{L}(1, 2T)$. We say that the Wiener candidate passes the *jump* test if it clears all the pole tops, namely if

$$W(x) > \mathcal{L}_n(2, x) \quad \text{for all } x \in P.$$

This test is weaker than the middle interval test which entails the above bound for all $x \in [\ell, r]$ (and we have $P \subseteq [\ell, r]$).

Indeed, the examination of the Wiener candidate may be represented in three steps:

- Test 1 is the side intervals test;
- Test 2 is the jump test;
- and Test 3 is the middle interval test.

See Figure 4.3.

Under $\mathbb{P}_{\mathcal{F}}$, the jump curve $J : [\ell, r] \rightarrow \mathbb{R}$ is constructed so that it has the conditional distribution of $W : [\ell, r] \rightarrow \mathbb{R}$ given that W passes Test 1 and 2 – see Figure 4.4. We may picture the jump curve as being a halfway house between the unadulterated Brownian randomness of the Wiener candidate W and the object of interest $\mathcal{L}(1, \cdot)$ whose Brownianity is to be established.

Next: the jump curve is a serious candidate to pass the third and final test; in a sense, it solves the high jump difficulty, with $\zeta = 1/3$.

PROPOSITION 4.3 (Realized promise of the jump curve). *We have that*

$$\mathbb{P}_{\mathcal{F}}\left(J \text{ passes Test 3}\right) \geq \exp\left\{-O(1)(\log \varepsilon^{-1})^{2/3}\right\} \cdot \mathbf{1}_{\text{Fav}}.$$

This is the key proposition in the general apparatus of the jump ensemble method. We will see to what we have reduced the Brownian bridge regularity Theorem 2.13 by assuming the proposition – then we will indicate something of its proof.

4.2. Proving Brownian bridge regularity via the realized promise of the jump curve

Recall that we want to show that for $A \subseteq \mathcal{C}_{0,0}([0, 1], \mathbb{R})$ with $\mathcal{B}_{1,0,0}^{[0,1]}(A) = \varepsilon$, we have that

$$\mathbb{P}\left(\mathcal{L}^{[0,1]}(1, \cdot) \in A\right) \leq \varepsilon \cdot C \exp\left\{C(\log \varepsilon^{-1})^{5/6}\right\},$$

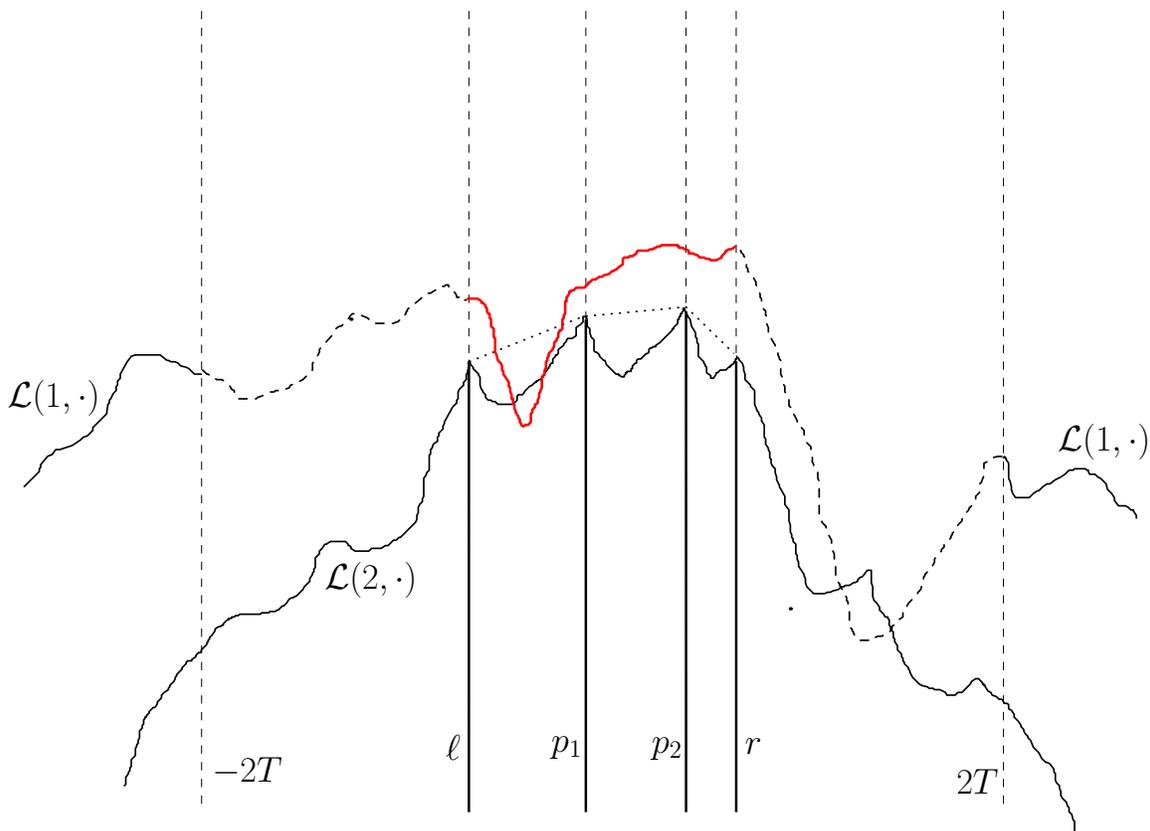


FIGURE 4.3. Here, the pole set P has four elements, with $P = \{\ell, p_1, p_2, r\}$. The dotted map Tent straddles the four bold poles indexed by P . The red Wiener candidate $W : [\ell, r] \rightarrow \mathbb{R}$ has transcript: test 1 is failed, with the left side test passed and the right side test failed; test 2 is passed, because W jumps over the four poles; test 3, the middle interval test, is failed, because W touches $\mathcal{L}(2, \cdot)$ on $[\ell, p_1]$.

Proof. The basic inequality of the jump method is

$$\begin{aligned}
& \mathbb{P}\left(\mathcal{L}^{[0,1]}(1, \cdot) \in A\right) \\
& \leq \mathbb{P}\left(\mathcal{L}^{[0,1]}(1, \cdot) \in A, \text{Fav}\right) + \mathbb{P}(\text{Fav}^c) \\
& = \mathbb{E}\left[\mathbb{P}_{\mathcal{F}}\left(\mathcal{L}^{[0,1]}(1, \cdot) \in A\right) \mathbf{1}_{\text{Fav}}\right] + \mathbb{P}(\text{Fav}^c) \\
& = \mathbb{E}\left[\mathbb{P}_{\mathcal{F}}\left(J^{[0,1]}(\cdot) \in A \mid J \text{ passes Test 3}\right) \mathbf{1}_{\text{Fav}}\right] + \mathbb{P}(\text{Fav}^c) \\
& \leq \mathbb{E}\left[\frac{\mathbb{P}_{\mathcal{F}}\left(J^{[0,1]}(\cdot) \in A\right)}{\mathbb{P}_{\mathcal{F}}\left(J \text{ passes Test 3}\right)} \mathbf{1}_{\text{Fav}}\right] + \mathbb{P}(\text{Fav}^c).
\end{aligned}$$

By the Realized Promise Proposition 4.3 and the Favourable Event Probability Lemma 4.2, we find that

$$\mathbb{P}\left(\mathcal{L}^{[0,1]}(1, \cdot) \in A\right) \leq \exp\left\{O(1)(\log \varepsilon^{-1})^{2/3}\right\} \mathbb{P}_{\mathcal{F}}\left(J^{[0,1]}(\cdot) \in A\right) \mathbf{1}_{\text{Fav}} + \varepsilon^{O(1)D^3}.$$

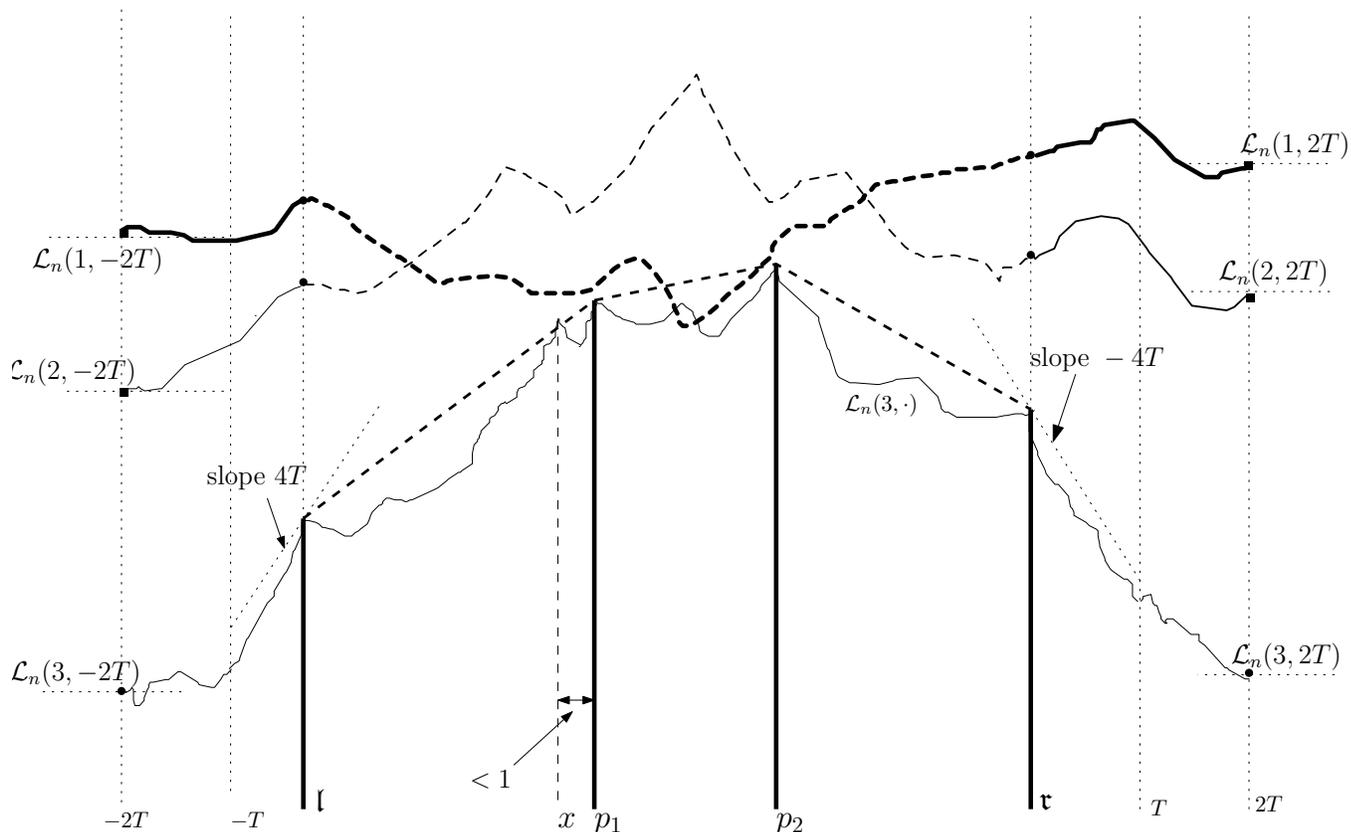


FIGURE 4.4. The jump ensemble depicted with $k = 2$ curves – we work with $k = 1$, so that we may call this object simply the jump curve. The depiction is not to scale: the intervals $[-2T, -T]$ and $[T, 2T]$ are too short. The pole set P in this example equals $\{\ell, p_1, p_2, r\}$. The vertical poles are depicted in thick solid lines. The function $\text{Tent} : [\ell, r] \rightarrow \mathbb{R}$ is defined to be the affine function that interpolates the values $\mathcal{L}(k+1, p)$ for $p \in P$. We picture vertical poles rising above the locations $p \in P$, with the tent map stretched over their tops. In the figure, the point x is an element of $\text{xExt}(\mathbf{c}_+)$ but not of P , because $|x - p_1| < 1$. (In fact, there are almost surely infinitely many elements of $\text{xExt}(\mathbf{c}_+) \setminus P$.) The piecewise affine dashed curve defined on $[\ell, r]$ is Tent . The rougher dashed curves are the jump ensemble $J : [1, 2] \times [\ell, r] \rightarrow \mathbb{R}$. The jump ensemble fails the criterion of passing Test 3 due to the meeting of the two J -curves and contact between $J(1, \cdot)$ and $\mathcal{L}_n(3, \cdot)$.

Since we may choose the jump method parameter $D > 0$ high enough that $O(1)D^3 \geq 1$, we see that the input needed to complete the proof of Theorem 2.13 – given the jump ensemble method, including Proposition 4.3 – is

$$\mathbb{P}_{\mathcal{F}}(J^{[0,1]}(\cdot) \in A) \mathbf{1}_{\text{Fav}} \leq \varepsilon \cdot \exp \left\{ C(\log \varepsilon^{-1})^{5/6} \right\},$$

where $\mathcal{B}_{1;0,0}^{[0,1]}(A) = \varepsilon$. The next proposition furnishes this input.

PROPOSITION 4.4 (Brownian regularity for J). *Let $A \subseteq \mathcal{C}_{0,0}([0, 1], \mathbb{R})$ satisfy the non-smallness condition that*

$$\mathcal{B}_{1;0,0}^{[0,1]}(A) \geq \varepsilon^{D^2/2}. \quad (9)$$

Then

$$\mathbb{P}_{\mathcal{F}}\left(J^{[0,1]}(\cdot) \in A\right) \mathbf{1}_{\text{Fav}} \leq \mathcal{B}_{1;0,0}^{[0,1]}(A) \cdot O(1) \exp\left\{O(1)(\log \varepsilon^{-1})^{5/6}\right\}.$$

Note that, in the application, $\mathcal{B}_{1;0,0}^{[0,1]}(A) = \varepsilon$, so that condition (9) is satisfied provided that we choose, as we may, $D^2 \geq 2$.

Two steps remain then to complete the proof of Theorem 2.13:

- (1) verifying the general jump method component Proposition 4.3;
- (2) and verifying the problem-determined Proposition 4.4.

We now sketch arguments for these two results in turn.

Sketch of proof of Proposition 4.3. Suppose for simplicity that $\ell = -T$, $r = T$, and that the pole set has three elements in the form $P = \{-T, 0, T\}$; and that

$$\mathcal{L}(2, 0) = T^2, \quad \mathcal{L}(2, -T) = \mathcal{L}(2, T) = -T^2, \quad \mathcal{L}(1, -2T) = \mathcal{L}(1, 2T) = -2T^2.$$

This is in essence an extreme case.

What hope is there for the curve J to succeed, i.e., to pass the middle interval test: $J(x) \geq \mathcal{L}(2, x)$ for all $x \in [-T, T]$? See Figure 4.5.

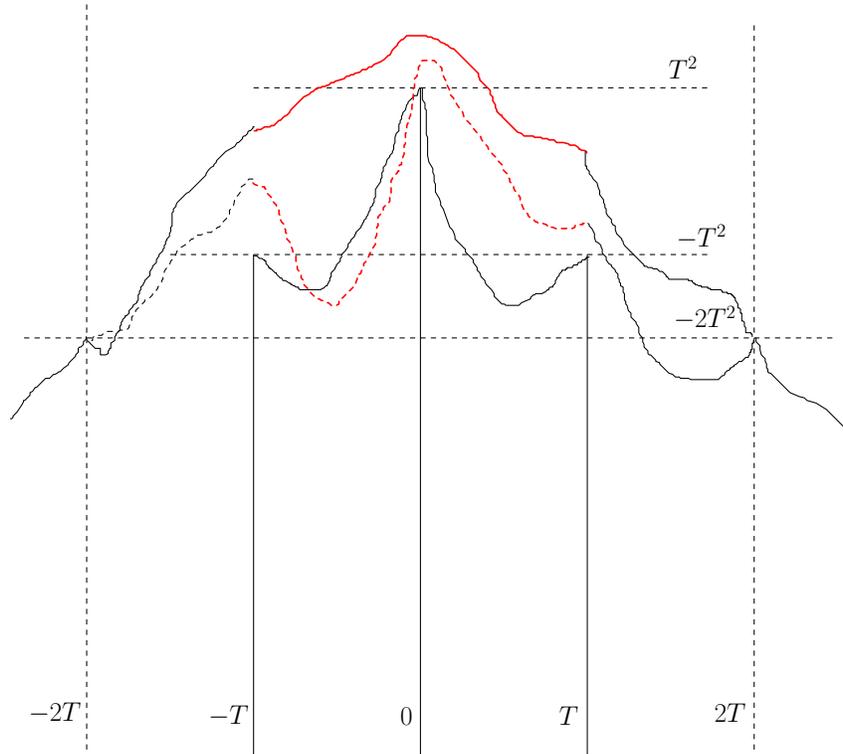


FIGURE 4.5. Two jump curves – the unbroken curve passes the middle interval test, and the dashed curve fails it. The random variable $J(-T)$ may be caricatured as a normal random variable – whose mean $-2T^2$ and whose variance has order T – conditioned to exceed the pole top at $-T$.

Crudely, $J(-T)$ is a normally distributed random variable of mean $-2T^2$ and variance of order T *conditioned* to be above $\mathcal{L}(2, -T) = -T^2$. By how much does $J(-T)$ typically exceed this minimum level? The density at x for this random excess equals

$$\frac{\exp\left\{-\frac{(T^2+x)^2}{T}\right\}}{\exp\left\{-\frac{(T^2)^2}{T}\right\}}$$

which when $x = o(T^2)$ is approximately $\exp\left\{-\frac{T^2x}{T}\right\} = \exp\{-Tx\}$.

This is an example of *cancellation of first-order kinetic costs*: the first-order term in the numerator and the denominator are both

$$\exp\{-T^3O(1)\} = \varepsilon^{O(1)D^3},$$

where recall that $T = D(\log \varepsilon^{-1})^{1/3}$. This cost is the unacceptably high term that we saw in the high jump difficulty – in the explanation of why the intermediate, Wiener candidate, approach was inadequate. Here, however, these high terms cancel, reflecting the conditioning to which the jump curve J is subject by its definition.

In the application, we choose x to be of order $(\log \varepsilon^{-1})^{1/3}$. We find that the probability under $\mathbb{P}_{\mathcal{F}}$ that $J(-T)$ (which is $J(\ell)$ in our present simplification) exceeds its minimum possible value by an order of $(\log \varepsilon^{-1})^{1/2}$ behaves as $\exp\{-\Theta(1)(\log \varepsilon^{-1})^{2/3}\}$ – this is because the argument of the exponential $-Tx$ has this form in view of $T = D(\log \varepsilon^{-1})^{1/3}$ and $x = (\log \varepsilon^{-1})^{1/3}$.

What happens in the circumstance that this excess indeed has this order? See Figure 4.6. Note that J is above Tent at $-T$, 0 and T by order $(\log \varepsilon^{-1})^{1/3}$. It deviates from affine interpolation of these endpoints by order $T^{1/2} = (\log \varepsilon^{-1})^{1/6}$ – that is, typically, J exceeds Tent on $[-T, T]$ by order $(\log \varepsilon^{-1})^{1/3}$ consistently.

Now $\mathcal{L}(2, \cdot)$ on $[-T, T]$ may rise above Tent sometimes – because not every extreme point of $\{(x, y) : -T \leq x \leq T, y \leq \mathfrak{c}_+(x)\}$ is in the pole set P . But by how much may it rise above Tent ?

LEMMA 4.5. For $x \in [-T, T] = [\ell, r]$,

$$\mathcal{L}(2, x) - \text{Tent}(x) \leq 8T.$$

See Figure 4.7 for a proof sketch.

Since $T \asymp (\log \varepsilon^{-1})^{1/3}$, we see that, with suitable selection of constant factors, J has in the circumstance – depicted in Figure 4.6 – of exceeding the pole tops by the indicated margin positive conditional probability to exceed $\mathcal{L}(2, \cdot)$ on $[-T, T]$ – and thus to pass Test 3, the middle interval test, and be successful. That is,

$$\mathbb{P}_{\mathcal{F}}(J \text{ succeeds}) \geq \exp\left\{-O(1)(\log \varepsilon^{-1})^{2/3}\right\} \mathbf{1}_{\text{Fav}},$$

as we sought to show in deriving Proposition 4.3. □

4.3. Sketch of proof of Proposition 4.4

Recall that, in essence, we seek to show that if $\varepsilon = \mathcal{B}_{1;0,0}^{[-1,1]}(A)$, then

$$\mathbb{P}_{\mathcal{F}}\left(J^{[-1,1]}(\cdot) \in A\right) \mathbf{1}_{\text{Fav}} \leq \varepsilon \cdot O(1) \exp\left\{O(1)(\log \varepsilon^{-1})^{5/6}\right\}.$$

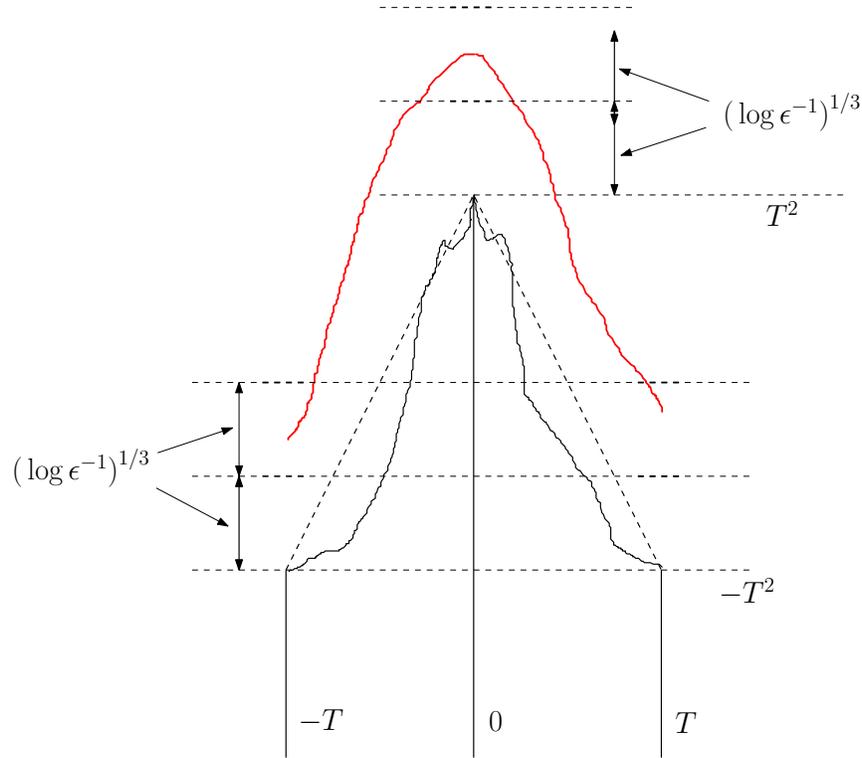


FIGURE 4.6. The jump curve on $[\ell, r] = [-T, T]$ is depicted as exceeding each of the three poles by at least $(\log \varepsilon^{-1})^{1/2}$; it thus typically exceeds the dotted map Tent by the same order throughout $[-T, T]$.

Here, we arbitrarily but harmlessly switched the interval of attention from $[0, 1]$ to $[-1, 1]$.

Again, we work with a special case to focus our attention: $P = \{-T, 0, T\}$, $\mathcal{L}(2, -T) = \mathcal{L}(2, T) = -T^2$ and $\mathcal{L}(2, 0) = T^2$.

How does J on $[-1, 1]$ differ from ordinary Brownian motion on this interval? More to the point: how does $J^{[-1, 1]}$ differ from standard Brownian bridge?

Let's make this comparison after fixing $J(-T)$ and $J(T)$ at the lowest possible locations that these values may adopt: $J(-T) = J(T) = -T^2$. This in essence gives a worst case scenario, in which the jump that J faces is as high as it ever may be – see Figure 4.8.

Now, $J^{[-1, 1]}$ can be caricatured as $B(x) + \text{SimpleTent}(x)$, where $\text{SimpleTent} : [-1, 1] \rightarrow \mathbb{R}$ is the affine function that interpolates the values zero at -1 ; T at zero; and zero at T . Similarly, $J^{[-1, 1]}$ is roughly a Brownian bridge that experiences drift T during $[-1, 0]$ and drift $-T$ during $[0, 1]$ – the reason for the rough validity of these models being that the stated drifts are dictated by the need to make a movement of order T^2 in a duration T .

We are trying to show that an ε -probability event for standard Brownian bridge has probability at most $\varepsilon \cdot O(1) \exp \left\{ O(1) (\log \varepsilon^{-1})^{5/6} \right\}$ for $J^{[-1, 1]}$. Given our caricature of $J^{[-1, 1]}$, which event is liable to maximize the $J^{[-1, 1]}$ -probability among those events whose probability under $\mathcal{B}_{1,0,0}^{[-1, 1]}$ equals ε ?

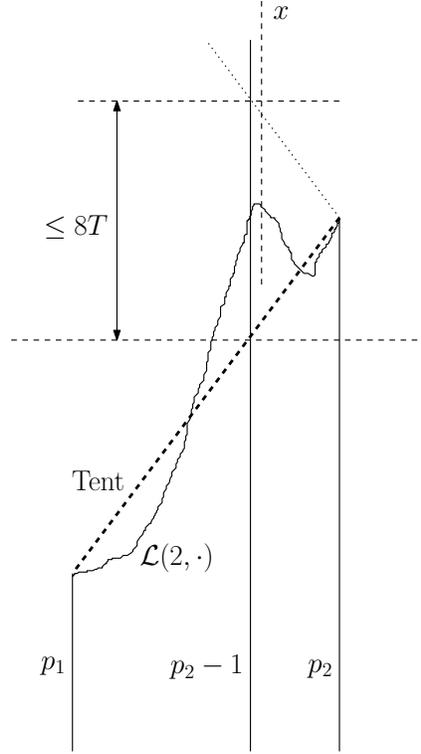


FIGURE 4.7. This sketch illustrates a proof of Lemma 4.5. An extreme point of the graphs of the concave majorant \mathfrak{c}_+ has horizontal coordinate x lying between consecutive pole set elements p_1 and p_2 . Thus x lies within unit distance of either p_1 or p_2 – here, we suppose the latter. In the sketch, the tent map, of slope at most $4T$, and a downward sloping dotted line, of slope $-4T$, touch the top of the pole at p_2 . The planar point $(x, \mathcal{L}(2, x))$ is bounded below by Tent and above by the downward sloping line. The indicated arrow distance, which is an upper bound on how much $\mathcal{L}(2, \cdot)$ may exceed Tent, is thus at most $8T$.

The event is clear enough: it is

$$\{J^{[-1,1]}(0) \geq R\},$$

where R satisfies

$$\mathcal{B}_{1;0,0}^{[-1,1]}(B(0) \geq R) = \varepsilon,$$

and thus also satisfies $\exp\{-R^2 O(1)\} = \varepsilon$ or equivalently $R = O(1)(\log \varepsilon^{-1})^{1/2}$.

What then is $\mathbb{P}(J^{[-1,1]}(0) \geq R)$ when R is so chosen? It is (and see Figure 4.9!)

$$\exp\{-(R-T)^2\} = \mathcal{B}_{1;0,0}^{[-1,1]}(B(0) \geq R) \times \exp\{RO(T)\}.$$

Here, the term T appearing in the left-hand expression is the roughly the mean of $J^{[-1,1]}(0)$. The term $\mathcal{B}_{1;0,0}^{[-1,1]}(B(0) \geq R)$ is a first-order term – of type “ e^{-R^2} ” – which in effect equals the dominant term when the bracket inside the exponential on the left-hand side is expanded. What we have seeing in the displayed equation is in effect a cancellation of first-order kinetic costs. That the remaining product on the right-hand side indeed has the form $\exp\{RO(T)\}$ depends on $T = o(R)$ – a valid bound given our choices.

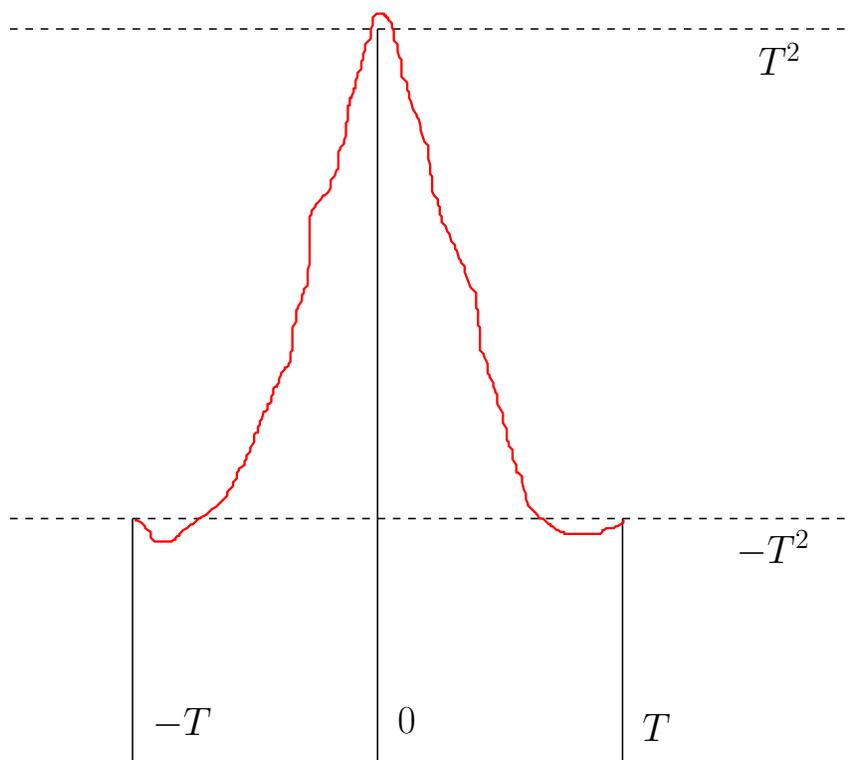


FIGURE 4.8. The jump curve illustrated in the special case under review and with boundary conditions making its effective drift before and after zero as pronounced as possible.

Thus,

$$\mathbb{P}(J^{[-1,1]}(0) \geq R) = \mathcal{B}_{1;0,0}^{[-1,1]}(B(0) \geq R) \exp \{O(1)(\log \varepsilon^{-1})^{5/6}\};$$

the first term on the right-hand side equals ε , and the second has its form because it equals $\exp \{RO(T)\}$ where $T = (\log \varepsilon^{-1})^{1/3}$ and $R = O(1)(\log \varepsilon^{-1})^{1/2}$.

This case is supposed to be the worst – this is merely a sketch of a proof! – and assuming that it is, our conclusion is

$$\mathbb{P}(J^{[-1,1]} \in A) \leq \varepsilon \cdot \exp \{O(1)(\log \varepsilon^{-1})^{5/6}\}$$

whenever $\mathcal{B}_{1;0,0}^{[-1,1]}(B \in A) = \varepsilon$. This completes the proof sketch of Proposition 4.4. \square

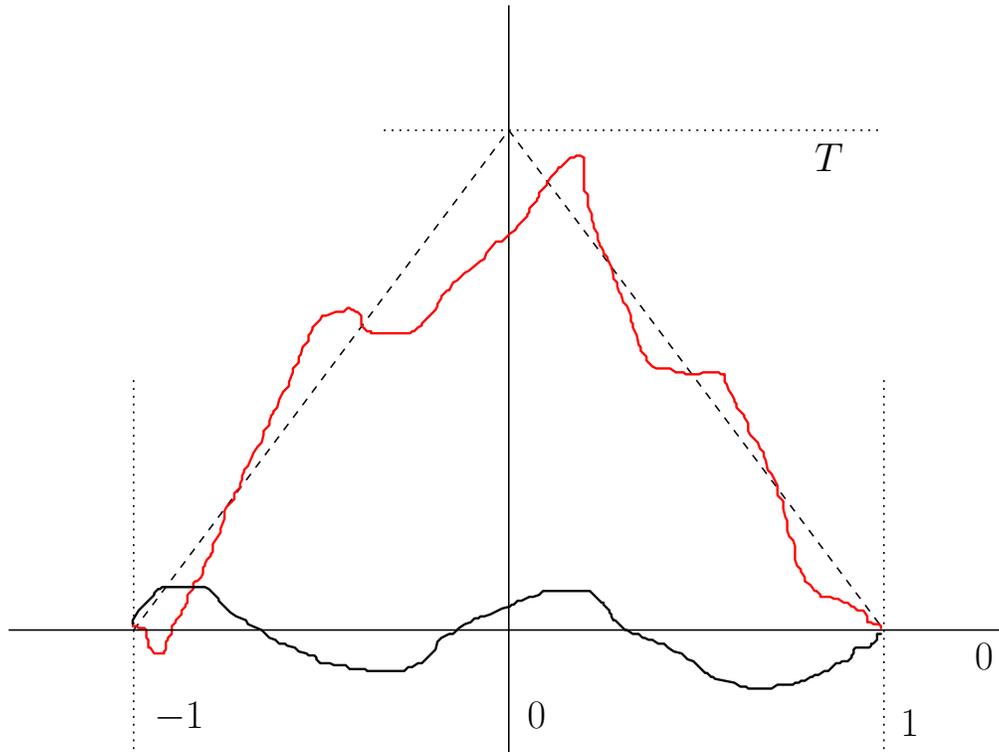


FIGURE 4.9. The dotted affine curve is SimpleTent. The red curve, our caricature of $J^{[-1,1]}$, is a sample of standard Brownian bridge on $[-1, 1]$ to which SimpleTent has been added. The red curve is to be compared in law to the black curve, which is standard Brownian bridge on the same interval. The event which is liable to expose the contrast between the curves most vividly is that either exceeds a given high level at zero.

Bibliography

- [CH14] Ivan Corwin and Alan Hammond. Brownian Gibbs property for Airy line ensembles. *Invent. Math.*, 195(2):441–508, 2014.
- [GW91] Peter W. Glynn and Ward Whitt. Departures from many queues in series. *Ann. Appl. Probab.*, 1(4):546–572, 1991.
- [Ham17] Alan Hammond. Brownian regularity for the Airy line ensemble, and multi-polymer watermelons in Brownian last passage percolation. *arXiv:1609.02971. Mem. Amer. Math. Soc., to appear*, 2017.
- [OY02] Neil O’Connell and Marc Yor. A representation for non-colliding random walks. *Electron. Comm. Probab.*, 7:1–12 (electronic), 2002.
- [PS02] Michael Prähofer and Herbert Spohn. Scale invariance of the PNG droplet and the Airy process. *J. Statist. Phys.*, 108(5-6):1071–1106, 2002. Dedicated to David Ruelle and Yasha Sinai on the occasion of their 65th birthdays.