Characterizing Gibbs states for area-tilted Brownian lines

Mriganka Basu Roy Chowdhury

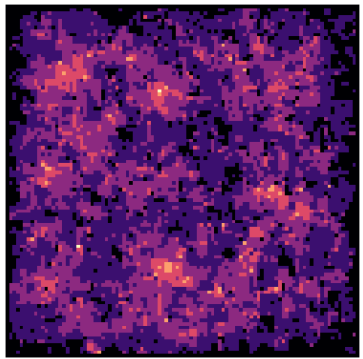

Figure: Height map for an entropically repulsed Ising interface in (2+1)-dimensions.

Introduction

- Systems of random curves, also referred to as **line ensembles**, appear throughout probability and statistical physics.
- Canonical examples of these processes are provided by eigenvalues of random matrices, or by level curves in $(2+1)$-dimensional random surface models such as the six-vertex model or interfaces in 3-dimensional spin systems.

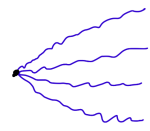

Dyson Brownian motion

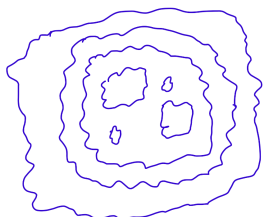

Level curves of height functions.

- A phenomenon of interest is the behavior of models with a hard substrate. This leads to what is known as **entropic repulsion**, where the floor propels the surface up.

- A phenomenon of interest is the behavior of models with a hard substrate. This leads to what is known as **entropic repulsion**, where the floor propels the surface up.
- An instance of this can be seen in the 3D Ising model on a box of side-length L, with positive boundary conditions on every face except the bottom, which is negative.
- This produces a random interface between the thermodynamically stable positive and negative phases, which is pushed up by the hard floor.

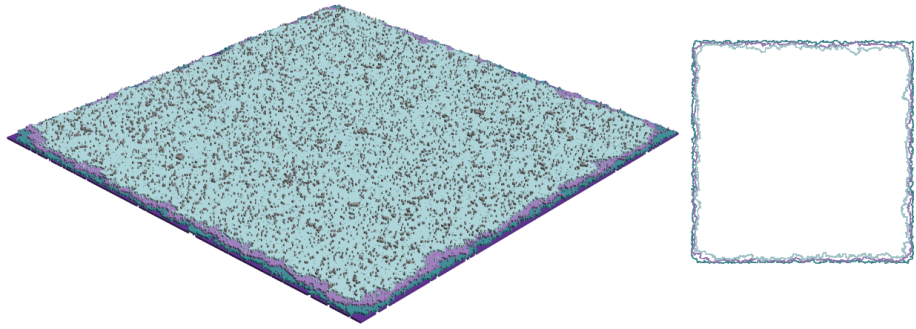

Figure: Entropic repulsion due to a hard floor constraint in (2+1)-dimensional Ising interfaces. The figure on the right illustrates the level curves of this interface. Picture due to Gheissari and Lubetzky (2021).

- Without the constraint of a hard floor, in low temperature, the surface would typically be rigid around height zero.

- Without the constraint of a hard floor, in low temperature, the surface would typically be rigid around height zero.
- But in this case, such a behavior incurs significant *entropic cost* — i.e., **a severe reduction in the number of possible configurations** since the surface cannot move downwards.

- Without the constraint of a hard floor, in low temperature, the surface would typically be rigid around height zero.
- But in this case, such a behavior incurs significant *entropic cost* — i.e., **a severe reduction in the number of possible configurations** since the surface cannot move downwards.
- Therefore, the surface chooses to quickly rise up to a height of $\Theta(\log L)$, ensuring **greater entropy in the interior.**

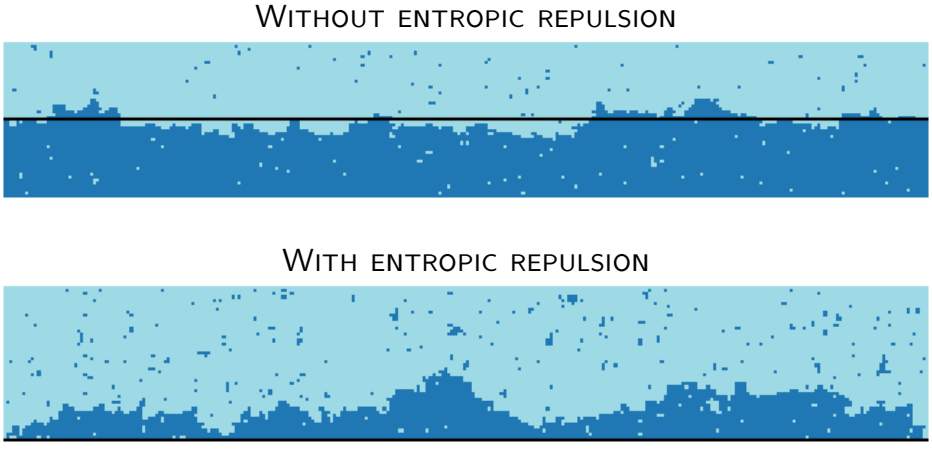

Figure: Effect of entropic repulsion in $(1 + 1)$ dimensional Ising interfaces. Observe the difference in fluctuations about the black lines.

- Entropic repulsion has been extensively studied in the literature, in particular via the **Solid-on-Solid model** (among others), a random height function aiming to capture the low-temperature fluctuations of Ising interfaces.

- Entropic repulsion has been extensively studied in the literature, in particular via the **Solid-on-Solid model** (among others), a random height function aiming to capture the low-temperature fluctuations of Ising interfaces.
- But as curves are often easier to study than surfaces, it is natural to analyze this height function via its *ensemble of level curves*, i.e., curves which pass through sites at the same height. Observe that by definition, **level curves are nonintersecting**.

- Entropic repulsion has been extensively studied in the literature, in particular via the **Solid-on-Solid model** (among others), a random height function aiming to capture the low-temperature fluctuations of Ising interfaces.
- But as curves are often easier to study than surfaces, it is natural to analyze this height function via its *ensemble of level curves*, i.e., curves which pass through sites at the same height. Observe that by definition, **level curves are nonintersecting**.
- Caputo, Ioffe and Wachtel (2019) proposed a putative scaling limit of this ensemble of curves called the **area-tilted line ensemble** — the primary object of interest in this talk. As the name suggests, these level curves exhibit a particular area dependent interaction which we will discuss shortly.

Gibbs property

- The seminal work of Corwin and Hammond (2014) on constructing the Airy line ensemble was primarily based on the observation that its prelimiting model satisfies a Gibbs property.
- The Gibbs property is essentially a **resampling invariance property**. More specifically, it prescribes the conditional distributions on **finite domains given boundary data**.

- The seminal work of Corwin and Hammond (2014) on constructing the Airy line ensemble was primarily based on the observation that its prelimiting model satisfies a Gibbs property.
- The Gibbs property is essentially a **resampling invariance property**. More specifically, it prescribes the conditional distributions on **finite domains given boundary data**.
- Gibbs properties are common in statistical physics, like in the study of spin systems.
- For example, the Gibbs property for the Ising model states that on any domain, the conditional distribution inside the domain (given the exterior) is that of an Ising model on that domain (with the appropriate boundary conditions).

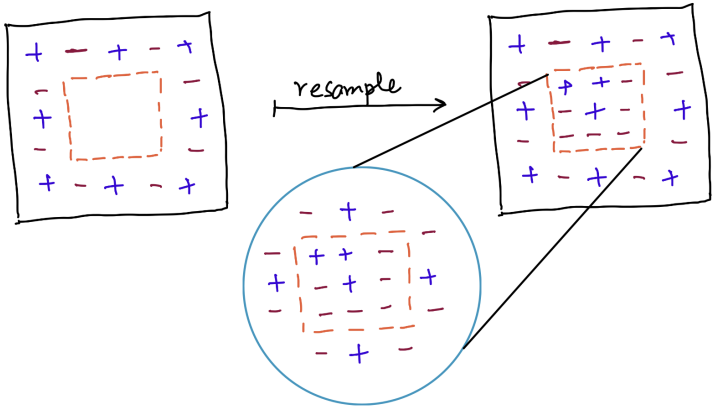

Figure: Gibbs property in the 2D Ising model. The dashed orange domain is resampled using an Ising model with the induced boundary conditions.

- The particular invariance property satisfied by the Airy line ensemble was termed the **"Brownian Gibbs property"**, which we describe next.

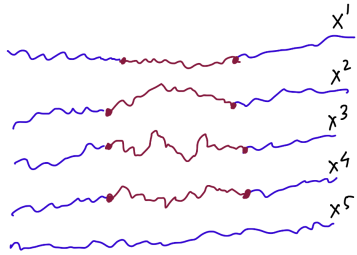

Figure: The Brownian Gibbs property with $k = 4$. We resample to get the red lines given the blue data.

- Fix any domain $[\ell, r]$ and an index k. Then the conditional distribution of $X^1, X^2, \ldots, X^k$ on $[\ell, r]$ given the blue data in the picture is that of independent Brownian bridges connecting $X^i(\ell)$ and $X^i(r)$ conditioned to *not intersect* and *avoid* X^{k+1} (illustrated in red).

- Such a description enables the use of probabilistic and geometric techniques via **monotonicity**, **coupling arguments** and **correlation decay estimates**.
- Since these tools are key inputs in our proofs as well, I will come back to these later in this talk.

- However, a Gibbs property is essentially just a **local specification**, prescribing conditional distributions only on finite domains.

- However, a Gibbs property is essentially just a **local specification**, prescribing conditional distributions only on finite domains.
- Consequently, a central question in statistical physics is that of determining all possible **infinite-volume Gibbs states**. These are infinite-volume measures (i.e., on all of $\mathbb{R}$ or $\mathbb{Z}^d$) satisfying the relevant Gibbs property.

- However, a Gibbs property is essentially just a **local specification**, prescribing conditional distributions only on finite domains.
- Consequently, a central question in statistical physics is that of determining all possible **infinite-volume Gibbs states**. These are infinite-volume measures (i.e., on all of $\mathbb{R}$ or $\mathbb{Z}^d$) satisfying the relevant Gibbs property.
- The existence and uniqueness of these Gibbs states are often highly nontrivial problems, and have been investigated for many models of interest.

- However, a Gibbs property is essentially just a **local specification**, prescribing conditional distributions only on finite domains.
- Consequently, a central question in statistical physics is that of determining all possible **infinite-volume Gibbs states**. These are infinite-volume measures (i.e., on all of $\mathbb{R}$ or $\mathbb{Z}^d$) satisfying the relevant Gibbs property.
- The existence and uniqueness of these Gibbs states are often highly nontrivial problems, and have been investigated for many models of interest.
- Under very general conditions, the collection of all infinite-volume Gibbs states form a **simplex**. Therefore is sufficient to determine all **extremal Gibbs states.**

- That the existence is not guaranteed can already be seen from a simple example:

- That the existence is not guaranteed can already be seen from a simple example:
- Consider a process X on $\mathbb{R}$ such that for any interval $[\ell, r]$, the conditional distribution of $X_{[\ell,r]}$ given everything outside is that of a **Brownian bridge** connecting $X(\ell)$ and $X(r)$.

- That the existence is not guaranteed can already be seen from a simple example:
- Consider a process X on $\mathbb{R}$ such that for any interval $[\ell, r]$, the conditional distribution of $X_{[\ell,r]}$ given everything outside is that of a **Brownian bridge** connecting $X(\ell)$ and $X(r)$.
- Such a process **does not exist.** To see this, fix T and condition on $X(-T) = x, X(+T) = y$. Then

$$X(0) \sim \mathcal{N}\left(0, \Theta\left(T\right)\right) + \frac{x+y}{2}.$$

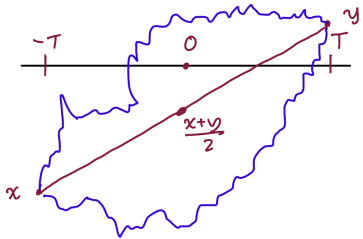

Figure: Two typical trajectories of X conditioned on $X(-T) = x, X(T) = y$.

- Observe that regardless of the value of $(x + y)/2$,

$$\mathbb{P}\left(X(0) \notin [-\sqrt{T}, \sqrt{T}]\right) \geq \Theta(1), \quad \text{uniformly in } T, x, y.$$

- Sending $T \to \infty$ shows that $X(0)$ **is not tight**. Therefore X does not exist.

- The natural analog of the Brownian bridge in higher dimensions is given by **the Gaussian free field (GFF)**.
- In the GFF, infinite-volume limits do not exist for $d = 2$, but exist for all $d \geq 3$, owing to the behavior of the Green's function. In fact, in $d \geq 3$, it has **infinitely** many Gibbs states, corresponding to harmonic functions on $\mathbb{Z}^d$.

- The natural analog of the Brownian bridge in higher dimensions is given by **the Gaussian free field (GFF)**.
- In the GFF, infinite-volume limits do not exist for $d = 2$, but exist for all $d \geq 3$, owing to the behavior of the Green's function. In fact, in $d \geq 3$, it has **infinitely** many Gibbs states, corresponding to harmonic functions on $\mathbb{Z}^d$.
- Although the dimension is a crucial parameter, many statistical physics models also possess a **temperature parameter** that plays a significant role in deciding the uniqueness of Gibbs states.

- The natural analog of the Brownian bridge in higher dimensions is given by **the Gaussian free field (GFF)**.
- In the GFF, infinite-volume limits do not exist for $d = 2$, but exist for all $d \geq 3$, owing to the behavior of the Green's function. In fact, in $d \geq 3$, it has **infinitely** many Gibbs states, corresponding to harmonic functions on $\mathbb{Z}^d$.
- Although the dimension is a crucial parameter, many statistical physics models also possess a **temperature parameter** that plays a significant role in deciding the uniqueness of Gibbs states.
- For example, the **high temperature** Ising model has **exactly one infinite-volume Gibbs state**, but **in low temperature** there are usually **at least two extremal states** – the positive and negative phases.

The area-tilted Brownian Gibbs property

- Returning to our primary object of interest, the area-tilted line ensembles also satisfy a certain *modified* Brownian Gibbs property. This is referred to as the **area-tilted Brownian Gibbs property**.

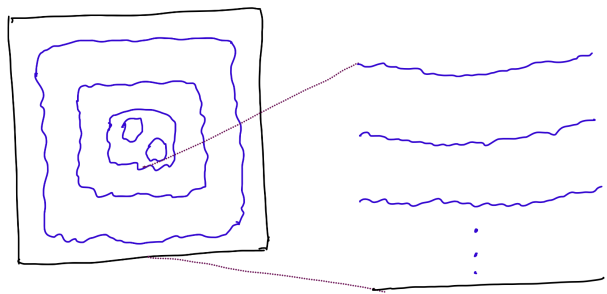

Figure: The scaling process to obtain the line ensemble.

- In order to describe the area-tilted Brownian Gibbs property, I first need to introduce the **area-tilted Brownian bridge.**
- An area-tilted Brownian bridge is simply a Brownian bridge with path $X(t)$
  - conditioned to be positive, and,
  - tilted exponentially by *area enclosed between it and the floor*, i.e., by

$$\exp\left(-(\text{tilt strength}) \cdot \int X(t)dt\right).$$

- In order to describe the area-tilted Brownian Gibbs property, I first need to introduce the **area-tilted Brownian bridge.**
- An area-tilted Brownian bridge is simply a Brownian bridge with path $X(t)$
  - conditioned to be positive, and,
  - tilted exponentially by *area enclosed between it and the floor*, i.e., by

$$\exp\left(-(\text{tilt strength}) \cdot \int X(t)dt\right).$$

- Given this, the area-tilted Brownian Gibbs property is the same as the Brownian Gibbs property, with **Brownian bridges replaced by area-tilted Brownian bridges with geometrically increasing tilt-strength.**

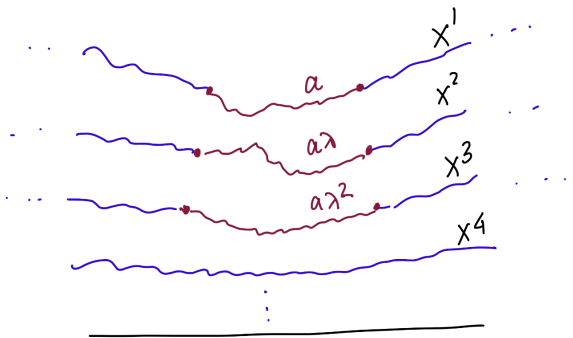

Figure: The area-tilted Brownian Gibbs property with $k = 3$ lines resampled on the interval $[\ell, r]$, resulting in the red curves. The red numbers on the lines are the area-tilt strengths of the lines, geometrically increasing at a rate of λ.

- That the level curves of entropically repulsed Ising interfaces face *geometrically increasing area-tilt* results from the fact that the entropy cost of having a large area A at height h scales like

$$(1 - \exp(-ch))^A \approx \exp\left(-e^{-ch}A\right)$$

Therefore, as h **decreases** the cost **increases geometrically**.

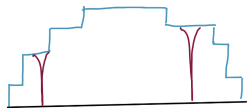

Figure: The entropy cost is decided by the probability that single site spikes (in red) hit the floor.

- That the level curves of entropically repulsed Ising interfaces face *geometrically increasing area-tilt* results from the fact that the entropy cost of having a large area A at height h scales like

$$(1 - \exp(-ch))^A \approx \exp\left(-e^{-ch}A\right)$$

Therefore, as h **decreases** the cost **increases geometrically**.

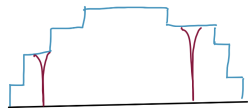

Figure: The entropy cost is decided by the probability that single site spikes (in red) hit the floor.

- In the scaling limit, the kth curve will have a penalty of $\exp\left(-a_k \int X^k(t)dt\right)$ with the a_k increasing geometrically.

- Given this Gibbs property, one then naturally asks the corresponding classification question: *What are all possible line ensembles with the area-tilted Brownian Gibbs property?*

- Given this Gibbs property, one then naturally asks the corresponding classification question: *What are all possible line ensembles with the area-tilted Brownian Gibbs property?*
- Indeed, this is our main result, which is stated next.

Theorem (B.R.C., Caputo and Ganguly, 2023 – informal)

Each possible extremal Gibbs state is characterized by the **coefficient of linear correction from parabolic growth**, *at $-\infty$ and $+\infty$. That is, if $X \sim \mu$ is a candidate process, the limits*

$$L = L(\mu) \stackrel{def}{=} \lim_{t \to -\infty} \frac{X^1(t) - t^2}{|t|}, \quad R = R(\mu) \stackrel{def}{=} \lim_{t \to \infty} \frac{X^1(t) - t^2}{|t|}.$$

exist as numbers and determine the measure. These limits can be $-\infty$, but must satisfy $L + R < 0$. Conversely, every choice of these limits corresponds to a unique extremal Gibbs state, called $\mu_{L,R}$.

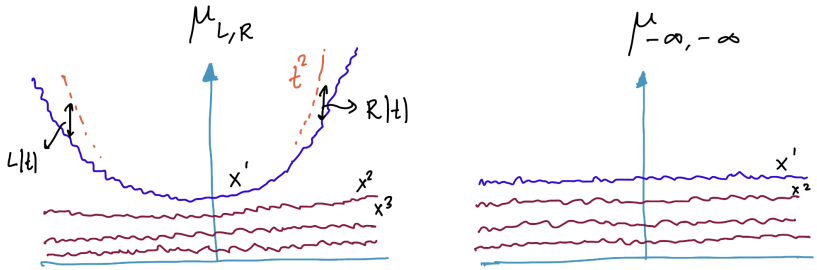

Figure: Various kinds of area-tilted BG measures.

History and known results

- As mentioned in the introduction, the area-tilted BG property appears in the context of entropically repulsed Ising interfaces, in particular, the entropically repulsed Solid-on-Solid (SOS) model, a model on height-functions ϕ, constructed as a **proxy for low-temperature** Ising interfaces.

- As mentioned in the introduction, the area-tilted BG property appears in the context of entropically repulsed Ising interfaces, in particular, the entropically repulsed Solid-on-Solid (SOS) model, a model on height-functions ϕ, constructed as a **proxy for low-temperature** Ising interfaces.
- To model the hard floor, one introduces the constraint of $\phi \geq 0$.

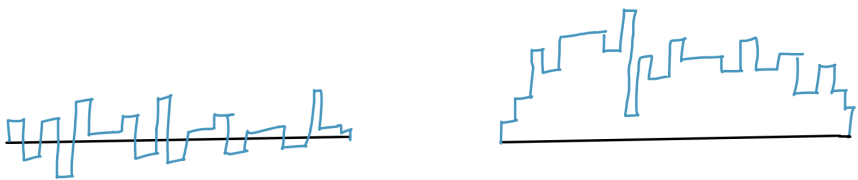

Figure: The effect of introducing the constraint $\phi \geq 0$. The surface on the left is rigid around zero, but the one on the right is entropically repulsed.

- As described previously, this causes an entropic repulsion effect, pushing the surface upto a height of $\Theta(\log(\text{system size}))$ in the bulk.

- As described previously, this causes an entropic repulsion effect, pushing the surface upto a height of $\Theta(\log(\text{system size}))$ in the bulk.
- This phenomenon was first studied by Bricmont, El Mellouki and Fröhlich (1986), where the picture predicted above was established.

- As described previously, this causes an entropic repulsion effect, pushing the surface upto a height of $\Theta(\log(\text{system size}))$ in the bulk.
- This phenomenon was first studied by Bricmont, El Mellouki and Fröhlich (1986), where the picture predicted above was established.
- On account of this quick rise of the height function, one observes **macroscopic level curves**, i.e., constant-height curves each encompassing a **constant fraction** of the domain.

- Caputo, Lubetzky, Martinelli, Sly and Toninelli (2013) performed a detailed analysis of these level curves, proving macroscopic limit shapes as well as fluctuation estimates.

- Caputo, Lubetzky, Martinelli, Sly and Toninelli (2013) performed a detailed analysis of these level curves, proving macroscopic limit shapes as well as fluctuation estimates.
- As a putative scaling limit of this system of level curves, Caputo, Ioffe and Wachtel (2019) consequently introduced the **area-tilted line ensemble**.

- Caputo, Lubetzky, Martinelli, Sly and Toninelli (2013) performed a detailed analysis of these level curves, proving macroscopic limit shapes as well as fluctuation estimates.
- As a putative scaling limit of this system of level curves, Caputo, Ioffe and Wachtel (2019) consequently introduced the **area-tilted line ensemble**.
- The associated area-tilted Brownian Gibbs property was also introduced, leading to the questions of **existence, uniqueness, and in general, classification of all Gibbs states with this property**.

- At this point, it will be convenient to introduce some notation.

- At this point, it will be convenient to introduce some notation.
- As explained earlier, the area-tilted Brownian bridge is defined as a Brownian bridge **tilted exponentially by area and conditioned to stay positive**.

- At this point, it will be convenient to introduce some notation.
- As explained earlier, the area-tilted Brownian bridge is defined as a Brownian bridge **tilted exponentially by area and conditioned to stay positive**.
- We will use the notation $\mathcal{L}_{\ell,r}^{a;x,y}$ to denote a strength-a area-tilted Brownian bridge X on $[\ell, r]$ satisfying $X(\ell) = x, X(r) = y$. Similarly, $\mathbf{B}_{\ell,r}^{x,y}$ will denote a Brownian bridge in the same setting.

- At this point, it will be convenient to introduce some notation.
- As explained earlier, the area-tilted Brownian bridge is defined as a Brownian bridge **tilted exponentially by area and conditioned to stay positive**.
- We will use the notation $\mathcal{L}_{\ell,r}^{a;x,y}$ to denote a strength-a area-tilted Brownian bridge X on $[\ell, r]$ satisfying $X(\ell) = x, X(r) = y$. Similarly, $\mathbf{B}_{\ell,r}^{x,y}$ will denote a Brownian bridge in the same setting.
- Then $\mathcal{L}$ is defined via the Radon-Nikodym derivative

$$\frac{d\mathcal{L}_{\ell,r}^{a;x,y}}{d\mathbf{B}_{\ell,r}^{x,y}}(X) = \underbrace{\frac{1}{Z_{\ell,r}^{a;x,y}}}_{\text{partition function}} \underbrace{\mathbf{1}_{X\geq 0}}_{\text{avoiding the floor}} \underbrace{\exp\left(-a\int_{\ell}^{r} X(t)dt\right)}_{\text{exponential tilting by area}}$$

- In the multi-line case, for $k > 1$, the definition of $\mathcal{L}^{a,\lambda;\overline{x},\overline{y}}_{k;\ell,r}$ is that of k **independent area-tilted Brownian bridges conditioned to avoid each other**, with the ith of them:
  - Pinned at $\overline{x}_i$ at ℓ, and $\overline{y}_i$ at r.
  - Facing an area-tilt strength of $a\lambda^{i-1}$.

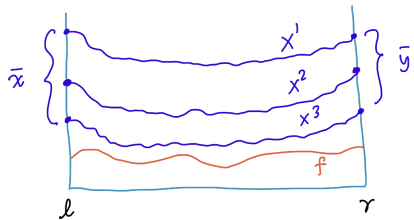

Figure: Illustration of $\mathcal{L}^{a,\lambda;\overline{x},\overline{y}}_{k;\ell,r}[f]$ with $k = 3$.

- In the multi-line case, for $k > 1$, the definition of $\mathcal{L}^{a,\lambda;\overline{x},\overline{y}}_{k;\ell,r}$ is that of k **independent area-tilted Brownian bridges conditioned to avoid each other**, with the ith of them:
  - Pinned at $\overline{x}_i$ at ℓ, and $\overline{y}_i$ at r.
  - Facing an area-tilt strength of $a\lambda^{i-1}$.

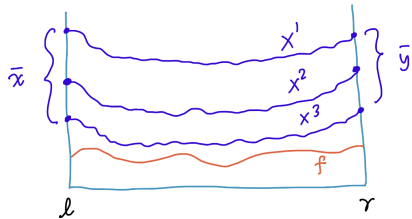

Figure: Illustration of $\mathcal{L}^{a,\lambda;\overline{x},\overline{y}}_{k;\ell,r}[f]$ with $k = 3$.

- We will also use $\mathcal{L}^{\cdots}_{\cdots}[f]$ to indicate a sample from $\mathcal{L}^{\cdots}_{\cdots}$ conditioned to stay above the function f.

- In this notation, the resampling distribution for the **area-tilted Brownian Gibbs property** is given by $\mathcal{L}^{a,\lambda;\overline{x},\overline{y}}_{k;\ell,r}$ where $[\ell, r]$ is the domain, k is the number of lines being resampled, and $\overline{x}_i = X^i(\ell)$ and $\overline{y}_i = X^i(r)$ for $1 \leq i \leq k$.

- In this notation, the resampling distribution for the **area-tilted Brownian Gibbs property** is given by $\mathcal{L}^{a,\lambda;\overline{x},\overline{y}}_{k;\ell,r}$ where $[\ell, r]$ is the domain, k is the number of lines being resampled, and $\overline{x}_i = X^i(\ell)$ and $\overline{y}_i = X^i(r)$ for $1 \leq i \leq k$.
- As alluded to before, the study of this Gibbs property benefits from several crucial monotonicity properties satisfied by the measures $\mathcal{L}^{\cdots;\overline{x},\overline{y}}_{\cdots}[f]$, discussed next.

- If the **endpoints and floor** for the blue lines are higher than those for the red lines, then **the law of the blue lines stochastically dominates that for the red lines.**

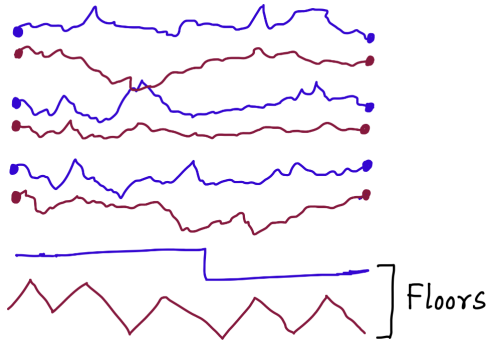

Figure: The blue lines stochastically dominate the red lines.

- Caputo, Ioffe, and Wachtel (2019) demonstrated the existence of measures with this Gibbs property by constructing the **zero-boundary** ensemble μ^0 as follows:

- Caputo, Ioffe, and Wachtel (2019) demonstrated the existence of measures with this Gibbs property by constructing the **zero-boundary** ensemble μ^0 as follows:
- Denote by $\mu^0_{k;T} = \mathcal{L}^{0,0}_{k;-T,T}$ the k-line zero-boundary ensemble, i.e., every line X^i (for $i = 1, 2, \ldots, k$) satisfies $X^i(\pm T) = 0$. The authors showed that the **measures $\mu^0_{k;T}$ are tight**. Then by **monotonicity, the local weak limit of these measures exists** as $k, T \to \infty$. This is called μ^0.

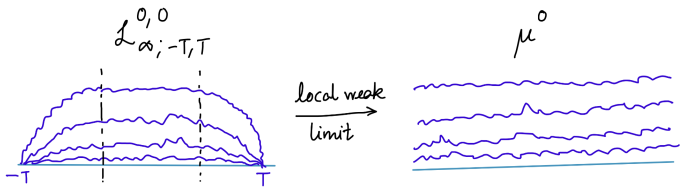

Figure: The limiting procedure to obtain μ^0.

- Caputo, Ioffe, and Wachtel (2019) demonstrated the existence of measures with this Gibbs property by constructing the **zero-boundary** ensemble μ^0 as follows:
- Denote by $\mu^0_{k;T} = \mathcal{L}^{0,0}_{k;-T,T}$ the k-line zero-boundary ensemble, i.e., every line X^i (for $i = 1, 2, \ldots, k$) satisfies $X^i(\pm T) = 0$. The authors showed that the **measures $\mu^0_{k;T}$ are tight**. Then by **monotonicity, the local weak limit of these measures exists** as $k, T \to \infty$. This is called μ^0.

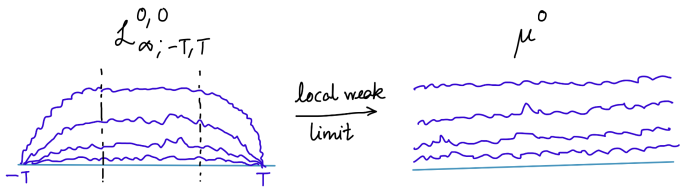

Figure: The limiting procedure to obtain μ^0.

- μ^0 is translation invariant.

- Another ensemble of interest arises when the **boundary is allowed to vary freely**, akin to **free boundary conditions** in spin models.

- Another ensemble of interest arises when the **boundary is allowed to vary freely**, akin to **free boundary conditions** in spin models.
- In our setup, this ensemble is called the **free boundary variant** $\mu^f_{k;T}$. *These are measures on* $[-T, T]$.

- Another ensemble of interest arises when the **boundary is allowed to vary freely**, akin to **free boundary conditions** in spin models.
- In our setup, this ensemble is called the **free boundary variant** $\mu^f_{k;T}$. *These are measures on* $[-T, T]$.
- The well-definedness of this measure for any k, T, as well as the tightness of this collection of measures was established by Caputo, Ioffe and Wachtel (2019).

- Another ensemble of interest arises when the **boundary is allowed to vary freely**, akin to **free boundary conditions** in spin models.
- In our setup, this ensemble is called the **free boundary variant** $\mu^f_{k;T}$. *These are measures on* $[-T, T]$.
- The well-definedness of this measure for any k, T, as well as the tightness of this collection of measures was established by Caputo, Ioffe and Wachtel (2019).
- Tightness guarantees the existence of some limit, and hence it is natural to ask if this collection has multiple limits, and how these limits behave.

- Interpreting $\mathcal{L}$ as a Markov process, Dembo, Lubetzky and Zeitouni (2022) established that **for every finite k, $\mu^f_{k;T}$, converges to the infinite-volume zero boundary variant** (for k lines).

- Interpreting $\mathcal{L}$ as a Markov process, Dembo, Lubetzky and Zeitouni (2022) established that **for every finite k, $\mu^f_{k;T}$, converges to the infinite-volume zero boundary variant** (for k lines). Consequently, one may take limits as the number of lines $k \to \infty$ to obtain

$$\lim_{k\to\infty} \lim_{T\to\infty} \mu^f_{k;T} = \mu^0$$

- Interpreting $\mathcal{L}$ as a Markov process, Dembo, Lubetzky and Zeitouni (2022) established that **for every finite k, $\mu^f_{k;T}$, converges to the infinite-volume zero boundary variant** (for k lines). Consequently, one may take limits as the number of lines $k \to \infty$ to obtain

$$\lim_{k\to\infty} \lim_{T\to\infty} \mu^f_{k;T} = \mu^0$$

- However the full limit remained open, and was established by Caputo and Ganguly (2023), who among other results concerning the tail behavior of the top line, correlation decay and ergodicity, also establish the following **uniqueness** result:

Theorem (Caputo and Ganguly, 2023)

Call an ensemble $X = (X^1, X^2, \dots,)$ uniformly tight if for every $\varepsilon > 0$, there is a C such that

$$\sup_t \mathbb{P}(X^1(t) \geq C) \leq \varepsilon.$$

If X satisfies this property and is asymptotically pinned to zero, then $X \sim \mu^0$.

Theorem (Caputo and Ganguly, 2023)

Call an ensemble $X = (X^1, X^2, \dots,)$ uniformly tight if for every $\varepsilon > 0$, there is a C such that

$$\sup_t \mathbb{P}(X^1(t) \geq C) \leq \varepsilon.$$

If X satisfies this property and is asymptotically pinned to zero, then $X \sim \mu^0$.

- The condition of **asymptotic pinning to zero** essentially says that the line ensemble has lines arbitrarily close to the floor at zero.
- This technical condition is needed because if X is an infinite ensemble satisfying the area-tilted Brownian Gibbs property, so is $X + h$ for any $h > 0$. *We will implicitly assume this condition throughout this presentation.*

- As a consequence of the result of Caputo and Ganguly, **every possible limit of $\mu^f_{k;T}$ is μ^0.**
- This is because **each such limit is uniformly tight via prelimiting tail bounds.**

- As a consequence of the result of Caputo and Ganguly, **every possible limit of $\mu^f_{k;T}$ is μ^0.**
- This is because **each such limit is uniformly tight via prelimiting tail bounds.**
- Also, every **translation-invariant ensemble is uniformly tight** implying that μ^0 is, in particular, **the unique translation invariant ensemble** (that is asymptotically pinned to zero).

- As a consequence of the result of Caputo and Ganguly, **every possible limit of $\mu^f_{k;T}$ is μ^0.**
- This is because **each such limit is uniformly tight via prelimiting tail bounds.**
- Also, every **translation-invariant ensemble is uniformly tight** implying that μ^0 is, in particular, **the unique translation invariant ensemble** (that is asymptotically pinned to zero).
- **But are there non-translation-invariant states?** This is answered in the positive by our main result.

Theorem (Main result formally stated)

Fix $\lambda > 1$. Let $\mathcal{G}_{\text{ext}}$ be the collection of all **extremal Gibbs measures** *satisfying the area-tilted Brownian Gibbs property (with strength given by $a = 2$ and λ). If $\mu \in \mathcal{G}_{\text{ext}}$ and $X \sim \mu$, then the following limits are* **deterministic** *and hold almost surely:*

$$L = L(\mu) \stackrel{def}{=} \lim_{t \to -\infty} \frac{X^1(t) - t^2}{|t|}, \quad R = R(\mu) \stackrel{def}{=} \lim_{t \to \infty} \frac{X^1(t) - t^2}{|t|}.$$

These limits satisfy $L, R \in [-\infty, \infty)$ and $L + R < 0$.Conversely, for every L, R satisfying these properties, there is exactly one measure $\mu_{L,R} \in \mathcal{G}_{\text{ext}}$ with these limits. As a consequence, $\mu_{-\infty,-\infty} = \mu^0$.

For every such μ, $X^2(t)$ (and hence $X^k(t)$ for all $k \geq 2$) is uniformly tight. Further all these measures are related by the following property: if $X \sim \mu_{L,R}$, then $X(\bullet - h) \sim \mu_{L+2h,R-2h}$.

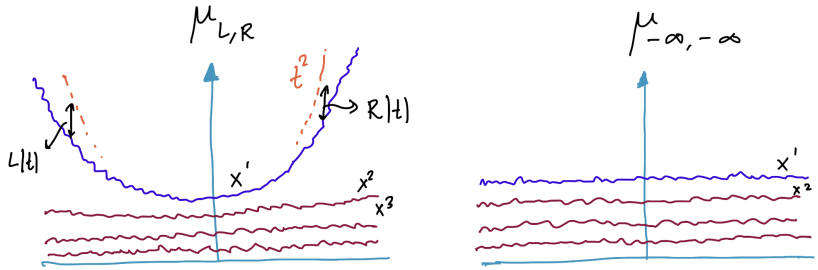

Figure: Only the top line can have parabolic growth, and only at rate t^2 with linear corrections.

- Although our results focus on the infinite-line ensemble, it is clear that the top line plays a special role as the **only line allowed to have parabolic growth.**

- Although our results focus on the infinite-line ensemble, it is clear that the top line plays a special role as the **only line allowed to have parabolic growth.**
- In the absence of the other lines, the top line satisfies a simplified Gibbs property where the resampling strategy is given by just **an area-tilted Brownian bridge conditioned to avoid the floor**.

- Although our results focus on the infinite-line ensemble, it is clear that the top line plays a special role as the **only line allowed to have parabolic growth.**
- In the absence of the other lines, the top line satisfies a simplified Gibbs property where the resampling strategy is given by just **an area-tilted Brownian bridge conditioned to avoid the floor**.
- This property was first studied by Ferrari and Spohn (2003) in the context of polynuclear growth. They considered a one-layer approximation to the multi-layer polynuclear growth model, replacing every layer except the top by a semicircle. This reduced the question to that of a **Brownian bridge conditioned to avoid a circle.**

- The gap process between this circle-avoiding Brownian bridge and the (semi)-circular floor has a local limit, called the **Ferrari-Spohn diffusion** μ^{FS} — a key member of the KPZ universality class.

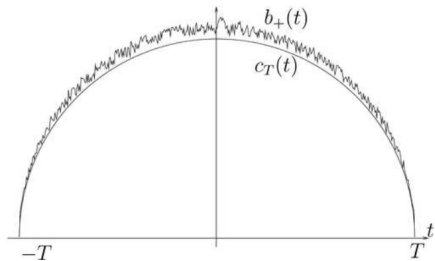

Figure: Brownian bridge conditioned to avoid a circle, from Ferrari and Spohn (2003). The local limit of the gap process $b_+ - c_T$ is the Ferrari-Spohn diffusion.

- In fact, a straightforward application of the Cameron-Martin / Girsanov theorem explains the relationship between area-tilting and circle/parabolic avoidance:

- In fact, a straightforward application of the Cameron-Martin / Girsanov theorem explains the relationship between area-tilting and circle/parabolic avoidance:

Lemma (Parabolic avoidance)

Fix $x, y, T > 0$. Let $X \sim \mathcal{L}_{-T,T}^{a;x,y}$, and let Y be a Brownian bridge from $x - \mathbf{p}(-T)$ to $y - \mathbf{p}(T)$, conditioned to avoid $-\mathbf{p}(\cdot)$ where $-\mathbf{p}$ is the parabola $t \mapsto -\frac{a}{2}t^2$. Then, $X \stackrel{d}{=} Y + \mathbf{p}$.

- In fact, a straightforward application of the Cameron-Martin / Girsanov theorem explains the relationship between area-tilting and circle/parabolic avoidance:

Lemma (Parabolic avoidance)

Fix $x, y, T > 0$. Let $X \sim \mathcal{L}_{-T,T}^{a;x,y}$, and let Y be a Brownian bridge from $x - \mathbf{p}(-T)$ to $y - \mathbf{p}(T)$, conditioned to avoid $-\mathbf{p}(\cdot)$ where $-\mathbf{p}$ is the parabola $t \mapsto -\frac{a}{2}t^2$. Then, $X \stackrel{d}{=} Y + \mathbf{p}$.

- That is, **Brownian bridges conditioned to avoid a parabola have a gap process identical in law to area-tilted Brownian bridges**.

- In fact, a straightforward application of the Cameron-Martin / Girsanov theorem explains the relationship between area-tilting and circle/parabolic avoidance:

Lemma (Parabolic avoidance)

Fix $x, y, T > 0$. Let $X \sim \mathcal{L}_{-T,T}^{a;x,y}$, and let Y be a Brownian bridge from $x - \mathbf{p}(-T)$ to $y - \mathbf{p}(T)$, conditioned to avoid $-\mathbf{p}(\cdot)$ where $-\mathbf{p}$ is the parabola $t \mapsto -\frac{a}{2}t^2$. Then, $X \stackrel{d}{=} Y + \mathbf{p}$.

- That is, **Brownian bridges conditioned to avoid a parabola have a gap process identical in law to area-tilted Brownian bridges**.
- All our results for the one-line case implicitly assume $a = 2$ since the corresponding parabola is $-t^2$. Scaling allows us to do this without loss of generality.

- In fact, a straightforward application of the Cameron-Martin / Girsanov theorem explains the relationship between area-tilting and circle/parabolic avoidance:

Lemma (Parabolic avoidance)

Fix $x, y, T > 0$. Let $X \sim \mathcal{L}_{-T,T}^{a;x,y}$, and let Y be a Brownian bridge from $x - \mathbf{p}(-T)$ to $y - \mathbf{p}(T)$, conditioned to avoid $-\mathbf{p}(\cdot)$ where $-\mathbf{p}$ is the parabola $t \mapsto -\frac{a}{2}t^2$. Then, $X \stackrel{d}{=} Y + \mathbf{p}$.

- That is, **Brownian bridges conditioned to avoid a parabola have a gap process identical in law to area-tilted Brownian bridges**.
- All our results for the one-line case implicitly assume $a = 2$ since the corresponding parabola is $-t^2$. Scaling allows us to do this without loss of generality.
- Since a circle and a parabola behave **similarly near the origin**, the local weak limit of the circle-avoiding process is the same as that when avoiding a parabola instead (both are the FS diffusion).

- Naturally, there is a corresponding classification question for the one-line area-tilted Brownian Gibbs property.

- Naturally, there is a corresponding classification question for the one-line area-tilted Brownian Gibbs property.
- A complete answer to this question serves as a key input to our main result.

Theorem (B.R.C., Caputo and Ganguly, 2023 – one line case)

Each possible extremal Gibbs state for the one-line case is characterized by the **coefficient of linear correction from parabolic growth**, *at $-\infty$ and $+\infty$. That is, if $X \sim \mu$ is a candidate process, the limits*

$$L = L(\mu) \stackrel{def}{=} \lim_{t \to -\infty} \frac{X(t) - t^2}{|t|}, \quad R = R(\mu) \stackrel{def}{=} \lim_{t \to \infty} \frac{X(t) - t^2}{|t|}.$$

exist as numbers and determine the measure. These limits can be $-\infty$, but must satisfy $L + R < 0$. Conversely, each choice of limits corresponds to a unique extremal Gibbs state. Further, μ^{FS} corresponds to $L = R = -\infty$, the unique translation-invariant state.

Observe that in terms of the "dual" $Y = X - t^2$, the limits considered above are $L = \lim_{t \to -\infty} Y(t)/|t|$ and $R = \lim_{t \to \infty} Y(t)/|t|$.

- The key lesson from the one-line analysis is the following:
- **The only possible rate of parabolic growth is given by $at^2/2$ (upto a linear correction) where a is the tilt strength. Otherwise, the process is uniformly tight.**

- The key lesson from the one-line analysis is the following:
- **The only possible rate of parabolic growth is given by $at^2/2$ (upto a linear correction) where a is the tilt strength. Otherwise, the process is uniformly tight.**
- In fact, this heuristically explains why only the top line can grow parabolically in the multi-line case.
- This is because the top line can at most grow as t^2.

- The key lesson from the one-line analysis is the following:
- **The only possible rate of parabolic growth is given by $at^2/2$ (upto a linear correction) where a is the tilt strength. Otherwise, the process is uniformly tight.**
- In fact, this heuristically explains why only the top line can grow parabolically in the multi-line case.
- This is because the top line can at most grow as t^2. But it **must also stay above all the other lines**, which have **higher tilts $2\lambda^{i-1}$. Therefore none of them can grow parabolically** because if they do, they **need to grow like $\lambda^{i-1}t^2$ and $\lambda^{i-1}t^2 - t^2 \gg t$.**

Overview of the proof

- We will now review some of the key ingredients in the proofs of the one-line and multi-line statements.

- We will now review some of the key ingredients in the proofs of the one-line and multi-line statements.
- Existence in both cases involves defining **finite domain versions** and then taking suitable (local weak) limits.

- We will now review some of the key ingredients in the proofs of the one-line and multi-line statements.
- Existence in both cases involves defining **finite domain versions** and then taking suitable (local weak) limits.
- The uniqueness question on the other hand involves taking **two measures with a given limit** and **coupling** them on (arbitrary) compact domains.

- We will now review some of the key ingredients in the proofs of the one-line and multi-line statements.
- Existence in both cases involves defining **finite domain versions** and then taking suitable (local weak) limits.
- The uniqueness question on the other hand involves taking **two measures with a given limit** and **coupling** them on (arbitrary) compact domains.
- In the case of infinitely many lines, this coupling is quite delicate, since one has to **couple infinitely many lines**.

- To get around this difficulty, we do not couple the lines exactly, but only upto a small error, via a **reverse coupling argument**.
- Without getting into details, essentially this considers two ordered boundary data but produces with high probability a coupling which reverses, for finitely many lines at the top, the ordering obtained from the usual monotone coupling.

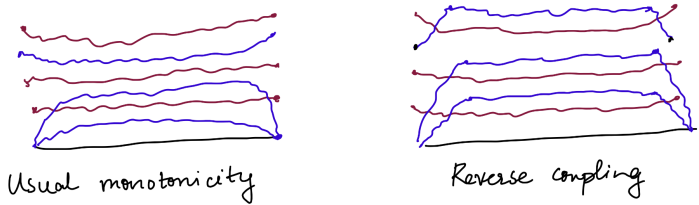

Figure: Coupling via usual monotonicity, vs using reverse coupling. Such "reversed" times exist with high probability in the independent coupling, allowing the reversal.

- Since the one-line result is much simpler, I will now describe some of the geometric ideas involved in its proof.

- Since the one-line result is much simpler, I will now describe some of the geometric ideas involved in its proof. Throughout this discussion, let X be a candidate process, and define $Y = X - t^2$, the dual representation.

- Since the one-line result is much simpler, I will now describe some of the geometric ideas involved in its proof. Throughout this discussion, let X be a candidate process, and define $Y = X - t^2$, the dual representation.
- The intuition here mostly relies on the **tangent method** (coined by Colomo and Sportiello in 2016).

- In our case, it states that the **macroscopic path of a Brownian bridge conditioned to avoid a parabola is that of the convex hull of the endpoints and the parabola.**

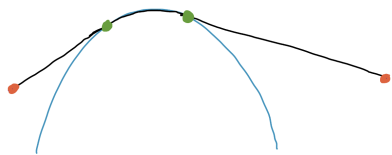

Figure: Brownian bridges starting and ending at the orange points, conditioned to avoid the parabola, macroscopically follow the black trajectory. The effect of the parabola only becomes pronounced near and between the green tangency points.

- In particular, this implies that if the endpoints are **far from the parabola**, the **avoidance constraint is not felt** until it is close to the parabola, and **therefore follows a Brownian trajectory**.

- In particular, this implies that if the endpoints are **far from the parabola**, the **avoidance constraint is not felt** until it is close to the parabola, and **therefore follows a Brownian trajectory**. For example, in the setup below, the portion of the path outside the red lines can be coupled with Brownian bridges connecting the endpoints to the tangency points in red.

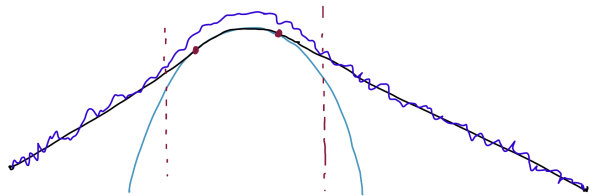

Figure: The portion outside the red lines can be coupled with Brownian bridges.

- A first illustration of this principle is manifested in the proof that the limits L and R exist. For concreteness, let us consider the existence of R.

- A first illustration of this principle is manifested in the proof that the limits L and R exist. For concreteness, let us consider the existence of R.
- Since the existence of the limit is the same as $\limsup = \liminf$, let us fix $\delta > 0$ and see why $\limsup_{t\to\infty} Y(t)/t > -h$ implies $\liminf_{t\to\infty} Y(t)/t \geq -h - \delta$ as well.

- A first illustration of this principle is manifested in the proof that the limits L and R exist. For concreteness, let us consider the existence of R.
- Since the existence of the limit is the same as $\limsup = \liminf$, let us fix $\delta > 0$ and see why $\limsup_{t\to\infty} Y(t)/t > -h$ implies $\liminf_{t\to\infty} Y(t)/t \geq -h - \delta$ as well.
- $\limsup_{t\to\infty} Y(t)/t > -h$ means that there are arbitrarily large times T such that $Y(T)/T > -h$.

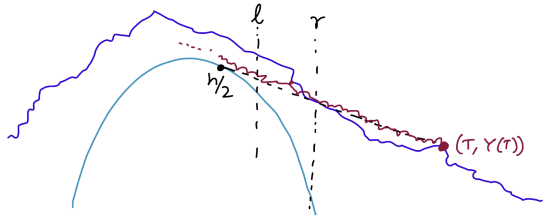

Figure: Upgrading a witness time T to claim uniform lower bounds on intervals via the tangent method.

- The tangent method then says that the red trajectory from a witness time T to the tangent location $h/2$ is Brownian, and hence **linear upto diffusive fluctuations.**
- This shows that $Y(t)/t > -h - \delta$ essentially throughout.

- As a second illustration, consider the existence part of the theorem with $L = R = -2K$.

- As a second illustration, consider the existence part of the theorem with $L = R = -2K$.

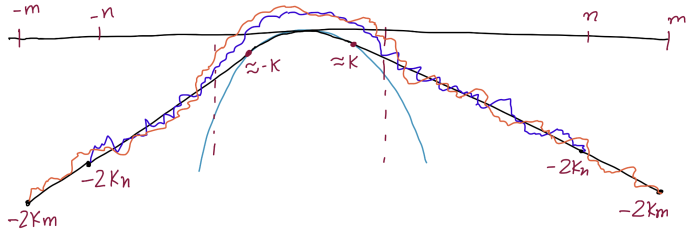

- Letting μ_T^K be the law of a parabola avoiding Brownian bridge with endpoints at $(\pm T, -2KT)$, we choose $n < m$ large and try to couple μ_n^K and μ_m^K in the bulk.

- As a second illustration, consider the existence part of the theorem with $L = R = -2K$.

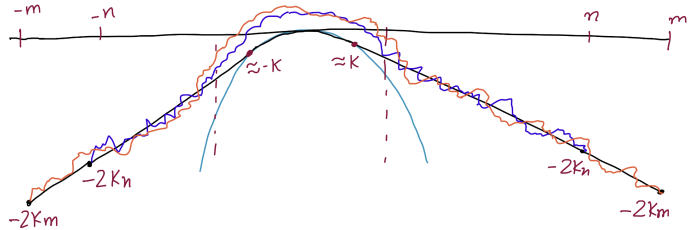

- Letting μ_T^K be the law of a parabola avoiding Brownian bridge with endpoints at $(\pm T, -2KT)$, we choose $n < m$ large and try to couple μ_n^K and μ_m^K in the bulk.
- Appealing to the tangent method, outside the red lines, the orange and blue curves are **close to Brownian bridges** (from their respective endpoints to $(\pm K, -K^2)$), and hence are **close in TV to each other** slightly away from K, allowing us to couple them.

- While these are some of the key ideas that go into the proof of the one-line statement, the analysis of the multi-line case is significantly more delicate.

- While these are some of the key ideas that go into the proof of the one-line statement, the analysis of the multi-line case is significantly more delicate.
- As already indicated, the key steps involve
  1. developing almost optimal "coming down" estimates,
  2. setting up an inductive framework to decouple the top line from the rest, and,
  3. implementing a reverse coupling strategy relying on stretched-exponential tail estimates.

- While these are some of the key ideas that go into the proof of the one-line statement, the analysis of the multi-line case is significantly more delicate.
- As already indicated, the key steps involve
  1. developing almost optimal "coming down" estimates,
  2. setting up an inductive framework to decouple the top line from the rest, and,
  3. implementing a reverse coupling strategy relying on stretched-exponential tail estimates.

$\sim$

Thank You!