The lower tail of q-pushTASEP

Milind Hegde
(based on joint work with Ivan Corwin)

Columbia University

MIT Probability Seminar
March 6, 2023

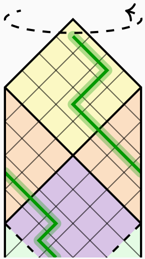

The KPZ universality class

The Kardar-Parisi-Zhang universality class contains a very broad class of "stochastic growth" models, including

- polymer models,
- interacting particle systems,
- metric/last passage percolation models, and more.

These each have a relevant observable, often called a height function which exhibits universal behavior, eg. in limiting distributions.

There are two subclasses: positive and zero temperature models.

Integrable models

Currently, only a small subset of models are tractable for deeper analysis, known as integrable models.

They are connected to algebraic and representation-theoretic structures which yield exact formulas. These can be analysed directly, though this is not trivial!

Integrable models

Currently, only a small subset of models are tractable for deeper analysis, known as integrable models.

They are connected to algebraic and representation-theoretic structures which yield exact formulas. These can be analysed directly, though this is not trivial!

In the past decade, integrable tools have started to be combined with probabilistic arguments to deepen the study of these models.

Integrable models

Currently, only a small subset of models are tractable for deeper analysis, known as integrable models.

They are connected to algebraic and representation-theoretic structures which yield exact formulas. These can be analysed directly, though this is not trivial!

In the past decade, integrable tools have started to be combined with probabilistic arguments to deepen the study of these models.

In these studies, a crucial integrable input has often been needed: estimates on the upper and lower tails of the relevant observables, and on the fluctuation scale.

The most prominent examples are probably last passage percolation and Gibbs line ensembles.

The upper vs. lower tail

It turns out that the upper tail is usually much easier to analyse than the lower tail: often there is a determinantal formula and bounds on the kernel suffice.

Techniques for the lower tail are considerably more involved, but a toolbox has been developed for zero-temperature models, including Riemann-Hilbert methods and connections to random matrix theory.

The upper vs. lower tail

It turns out that the upper tail is usually much easier to analyse than the lower tail: often there is a determinantal formula and bounds on the kernel suffice.

Techniques for the lower tail are considerably more involved, but a toolbox has been developed for zero-temperature models, including Riemann-Hilbert methods and connections to random matrix theory.

Techniques for positive temperature are not as extensive. A few approaches are

- using determinantal formulas for Laplace transforms
- Riemann-Hilbert representations for Laplace transforms
- geometric methods for polymer models.

These do not seem able to address all models, including ours.

The model of q-pushTASEP

q-pushTASEP is a discrete time interacting particle system on $\mathbb{Z}$ introduced by Matveev-Petrov—related to pushTASEP, q-TASEP, and q-pushTASEP (continuous time).

There are N particles whose positions at time T are denoted $x_1(T) < \ldots < x_N(T)$, and whose initial positions are $x_i(0) = i$ (step initial condition).

At each time step, particle positions are updated from left to right. At time T, the k^{th} attempts to move $J_{k,T} + P_{k,T}$ positions to the right, where

$$\text{(jump)} \qquad \mathbb{P}(J_{k,T} = s) = q^s \cdot \frac{(q;q)_\infty}{(q;q)_s}$$

$$\text{(push)} \qquad \mathbb{P}(P_{k,T} = s) = q^{s \cdot \mathrm{gap}_k(T)} \cdot (q^{\mathrm{gap}_k(T)}; q^{-1})_{\Delta_{k-1,T} - s} \cdot \frac{(q^{-1};q^{-1})_{\Delta_{k-1,T}}}{(q^{-1};q^{-1})_s \cdot (q^{-1};q^{-1})_{\Delta_{k-1,T} - s}},$$

where

$$(z;x)_n = \prod_{i=0}^{n-1}(1 - zx^i), \qquad \mathrm{gap}_k(T) = x_k(T) - x_{k-1}(T),$$

$$\Delta_{k-1,T} = x_{k-1}(T+1) - x_{k-1}(T).$$

The model of q-pushTASEP

At each time step, particle positions are updated from left to right. At time T, the k^{th} attempts to move $J_{k,T} + P_{k,T}$ positions to the right, where

$$\text{(jump)} \qquad \mathbb{P}(J_{k,T} = s) = q^s \cdot \frac{(q;q)_\infty}{(q;q)_s}$$

$$\text{(push)} \qquad \mathbb{P}(P_{k,T} = s) = q^{s \cdot \mathrm{gap}_k(T)} \cdot (q^{\mathrm{gap}_k(T)}; q^{-1})_{\Delta_{k-1,T} - s} \cdot \frac{(q^{-1};q^{-1})_{\Delta_{k-1,T}}}{(q^{-1};q^{-1})_s \cdot (q^{-1};q^{-1})_{\Delta_{k-1,T} - s}},$$

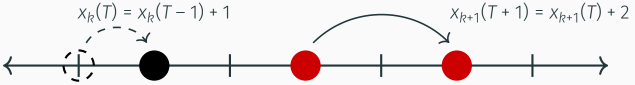

The model of q-pushTASEP

The definition of the model may not seem very natural, but it is integrable; indeed, using this structure, Vető computed the law of large numbers of the model and showed limiting Tracy-Widom GUE fluctuations.

The model of q-pushTASEP

The definition of the model may not seem very natural, but it is integrable; indeed, using this structure, Vető computed the law of large numbers of the model and showed limiting Tracy-Widom GUE fluctuations.

It is known that $x_N(T)$ has the same distribution as the length of the top row of a partition from the q-Whittaker measure—a crucial connection for our arguments.

This is a measure on partitions defined in terms of the q-Whittaker polynomials, though its precise definition won't be needed for us here.

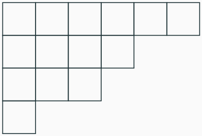

The model of q-pushTASEP

The definition of the model may not seem very natural, but it is integrable; indeed, using this structure, Vető computed the law of large numbers of the model and showed limiting Tracy-Widom GUE fluctuations.

It is known that $x_N(T)$ has the same distribution as the length of the top row of a partition from the q-Whittaker measure—a crucial connection for our arguments.

This is a measure on partitions defined in terms of the q-Whittaker polynomials, though its precise definition won't be needed for us here.

Another point of interest of the model is that, taking $q \to 1$ and rescaling appropriately, $x_N(T)$ converges to the free energy of the log-gamma polymer.

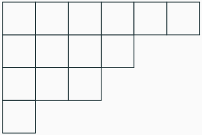

Main result: q-pushTASEP lower tail

Theorem

Let $q \in (0,1)$. There exist absolute constants c and θ_0 such that, for large N and $\theta > \theta_0$,

$$\mathbb{P}\left(x_N(N) - f_q N < -\theta N^{1/3}\right) \leq \exp\left(-c\theta^{3/2}\right).$$

We expect the true tail decay to be $\exp(-c\theta^3)$, at least for $\theta \ll N^{2/3}$ (i.e., before the large deviations regime), but we do not prove this.

A connection to last passage percolation

In geometric last passage percolation, we have an i.i.d. environment $\{\xi_v\}$ of Geo(z) random variables, one for each of the N^2 small squares in a $N \times N$ big square.

(Here, $X \sim \text{Geo}(z)$ means $\mathbb{P}(X \geq k) = z^k$ for $k = 0, 1, 2, \ldots$)

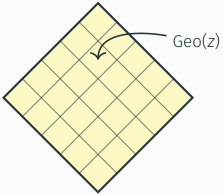

A connection to last passage percolation

In geometric last passage percolation, we have an i.i.d. environment $\{\xi_v\}$ of Geo(z) random variables, one for each of the N^2 small squares in a $N \times N$ big square.

(Here, $X \sim \text{Geo}(z)$ means $\mathbb{P}(X \geq k) = z^k$ for $k = 0, 1, 2, \ldots$)

We consider downward paths. Each path γ has a weight $w(\gamma)$ given by

$$w(\gamma) = \sum_{v \in \gamma} \xi_v.$$

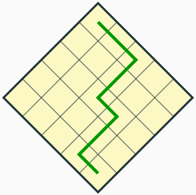

A connection to last passage percolation

In geometric last passage percolation, we have an i.i.d. environment $\{\xi_v\}$ of Geo(z) random variables, one for each of the N^2 small squares in a $N \times N$ big square.

(Here, $X \sim \text{Geo}(z)$ means $\mathbb{P}(X \geq k) = z^k$ for $k = 0, 1, 2, \ldots$)

We consider downward paths. Each path γ has a weight $w(\gamma)$ given by

$$w(\gamma) = \sum_{v \in \gamma} \xi_v.$$

The LPP value $L = \max_\gamma w(\gamma)$.

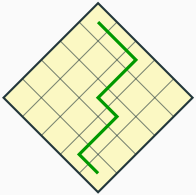

A connection to last passage percolation

Surprisingly, there is a connection between q-pushTASEP and a certain LPP problem.

Consider an infinite strip with independent geometric RVs as shown; the parameter is the same in a single big square, and decays geometrically as we move down.

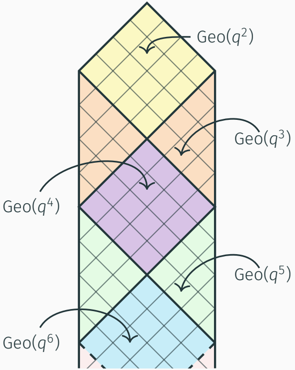

A connection to last passage percolation

Surprisingly, there is a connection between q-pushTASEP and a certain LPP problem.

Consider an infinite strip with independent geometric RVs as shown; the parameter is the same in a single big square, and decays geometrically as we move down.

Because of this decay, by the Borel-Cantelli lemma, only finitely many small squares are non-zero. (Recall $\mathbb{P}(\mathrm{Geo}(q^k) \neq 0) = q^k$.)

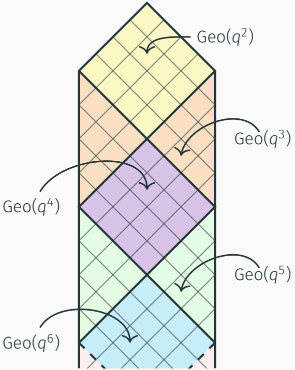

A connection to last passage percolation

We regard this strip as being periodic, i.e., a cylinder. We define the LPP value L as being from the top to infinitely down, but allow paths to wrap around.

The maximizing path is called a geodesic.

Note that L is well-defined and finite almost surely!

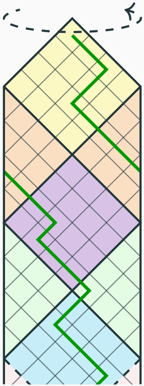

A connection to last passage percolation

We regard this strip as being periodic, i.e., a cylinder. We define the LPP value L as being from the top to infinitely down, but allow paths to wrap around.

The maximizing path is called a geodesic.

Note that L is well-defined and finite almost surely!

Remarkably, it was shown by Imamura-Mucciconi-Sasamoto that

$$L \stackrel{d}{=} \lambda_1,$$

where λ_1 is the first row of a partition $\sim q$-Whittaker measure.

So, it follows that $L \stackrel{d}{=} x_N(N)$!

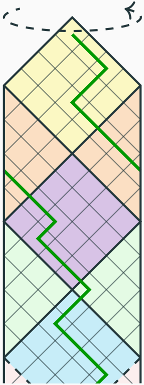

A fairly sharp lower bound using the variational problem

This identity is very useful to bound the lower tail of q-pushTASEP. Indeed, for *any* path γ,

$$x_N(N) \stackrel{d}{=} L \geq w(\gamma).$$

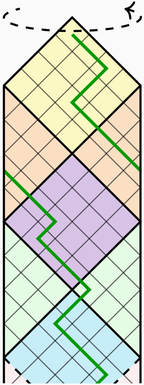

This identity is very useful to bound the lower tail of q-pushTASEP. Indeed, for *any* path γ,

$$x_N(N) \overset{d}{=} L \geq w(\gamma).$$

As a consequence,

$$\mathbb{P}(x_N(N) \leq t) \leq \mathbb{P}(w(\gamma) \leq t).$$

So if we find a γ which is close in weight to L and whose weight can be analyzed, we get a tail bound on $x_N(N)$.

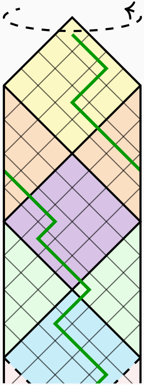

A fairly sharp lower bound using the variational problem

This identity is very useful to bound the lower tail of q-pushTASEP. Indeed, for *any* path γ,

$$x_N(N) \stackrel{d}{=} L \geq w(\gamma).$$

As a consequence,

$$\mathbb{P}(x_N(N) \leq t) \leq \mathbb{P}(w(\gamma) \leq t).$$

So if we find a γ which is close in weight to L and whose weight can be analyzed, we get a tail bound on $x_N(N)$.

We don't have tail bounds for LPP values in periodic or inhomogeneous environments. So we consider an easier path instead of the geodesic itself.

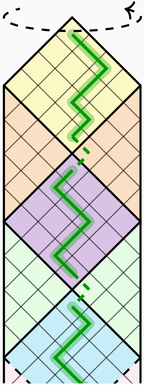

A fairly sharp lower bound using the variational problem

The easier path can be written as a concatenation of LPP paths in homogeneous environments, and the smaller paths cannot wrap around the strip.

Perhaps surprisingly, this path's weight and $x_N(N)$ match to the first order law of large numbers term!

The path's weight is a infinite sum of independent RVs. If we can control its fluctuations, we will be done.

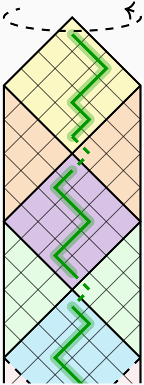

An issue...

Sharp lower tail bounds for geometric LPP are available, but for fixed q; the constants may blow up as $q \to 0$ or 1.

But the parameter of the geometric LPP problem in the i^{th} big square from the top is $q^{2i} \to 0$ as $i \to \infty$.

We need to develop new lower tail estimates for geometric LPP that hold uniformly for $q \in (0, 1)$.

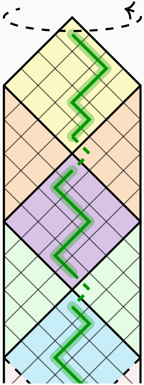

Main result: Uniform lower tail for geometric LPP

Let T_N be the LPP value from top to bottom of an $N \times N$ square in an environment given by i.i.d. Geo(q) random variables, and let $\mu_q = (1 + q^{1/2})^2/(1 - q)$.

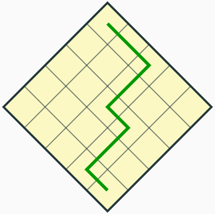

Theorem

There exist positive constants c, x_0, and N_0 such that, for $q \in (0, 1)$, $N \geq N_0$, and $x > x_0$,

$$\mathbb{P}\left(T_N \leq (\mu_q - 1)N - x \cdot \frac{q^{1/6}}{1-q} N^{1/3}\right) \leq \exp(-cx^{3/2}).$$

Main result: Uniform lower tail for geometric LPP

Theorem

There exist positive constants c, x_0, and N_0 such that, for $q \in (0,1)$, $N \geq N_0$, and $x > x_0$,

$$\mathbb{P}\left(T_N \leq (\mu_q - 1)N - x \cdot \frac{q^{1/6}}{1-q} N^{1/3}\right) \leq \exp(-cx^{3/2}).$$

Note that the constant c and x_0 are independent of q, which is allowed to range over all of $(0,1)$. As we saw, this is crucial for the application.

Effectively, $q \geq N^{-2}$. If $q = cN^{-2}$ and $N \to \infty$, geometric LPP becomes Poissonian LPP.

Main result: Uniform lower tail for geometric LPP

Theorem

There exist positive constants c, x_0, and N_0 such that, for $q \in (0,1)$, $N \geq N_0$, and $x > x_0$,

$$\mathbb{P}\left(T_N \leq (\mu_q - 1)N - x \cdot \frac{q^{1/6}}{1-q} N^{1/3}\right) \leq \exp(-cx^{3/2}).$$

Note that the constant c and x_0 are independent of q, which is allowed to range over all of $(0,1)$. As we saw, this is crucial for the application.

Effectively, $q \geq N^{-2}$. If $q = cN^{-2}$ and $N \to \infty$, geometric LPP becomes Poissonian LPP.

The true decay behaviour is $\exp(-cx^3)$, but we do not achieve this.

Main result: Uniform lower tail for geometric LPP

Theorem

There exist positive constants c, x_0, and N_0 such that, for $q \in (0,1)$, $N \geq N_0$, and $x > x_0$,

$$\mathbb{P}\left(T_N \leq (\mu_q - 1)N - x \cdot \frac{q^{1/6}}{1-q} N^{1/3}\right) \leq \exp(-cx^{3/2}).$$

Note that the constant c and x_0 are independent of q, which is allowed to range over all of $(0,1)$. As we saw, this is crucial for the application.

Effectively, $q \geq N^{-2}$. If $q = cN^{-2}$ and $N \to \infty$, geometric LPP becomes Poissonian LPP.

The true decay behaviour is $\exp(-cx^3)$, but we do not achieve this.

With this tail estimate, we obtain the lower tail of $x_N(N)$ by using a concentration inequality to control the contribution from all the big squares.

The concentration inequality has to take into account the variables' decaying scale.

Proof ideas

Determinantal point processes

Determinantal point processes are a class of point processes with several features making them easier to analyse.

Distributions of point processes can be described by their n-point correlation functions $\rho_n(x_1, \ldots, x_n)$; for DPPs, ρ_n is given as an $n \times n$ determinant of a kernel.

Many properties of the point process can be understood by studying the asymptotics of the kernel only.

Independent

DPP

Determinantal point processes

Determinantal point processes are a class of point processes with several features making them easier to analyse.

Distributions of point processes can be described by their n-point correlation functions $\rho_n(x_1, \ldots, x_n)$; for DPPs, ρ_n is given as an $n \times n$ determinant of a kernel.

Many properties of the point process can be understood by studying the asymptotics of the kernel only.

Independent

DPP

T_N from geometric LPP is the top particle in a DPP (Meixner ensemble)!

Widom's trick

Widom's trick provides a simpler way to obtain lower tail estimates for determinantal point processes, at the cost of a worse tail exponent ($\frac{3}{2}$ instead of 3).

Heuristically, it allows one to treat the process as if the points are independent, ignoring the repulsion that DPPs actually exhibit.

For a DPP ($\lambda_1 > \lambda_2 > \ldots > \lambda_n$) with kernel K_n,

$$\mathbb{P}(\lambda_1 \leq t) = \det\left(I_n - K_n^t\right)$$

.

Widom's trick

Widom's trick provides a simpler way to obtain lower tail estimates for determinantal point processes, at the cost of a worse tail exponent ($\frac{3}{2}$ instead of 3).

Heuristically, it allows one to treat the process as if the points are independent, ignoring the repulsion that DPPs actually exhibit.

For a DPP ($\lambda_1 > \lambda_2 > \ldots > \lambda_n$) with kernel K_n,

$$\mathbb{P}(\lambda_1 \leq t) = \det\left(I_n - K_n^t\right) = \prod_{i=1}^{n}(1 - \rho_i^t)$$

.

Widom's trick

Widom's trick provides a simpler way to obtain lower tail estimates for determinantal point processes, at the cost of a worse tail exponent ($\frac{3}{2}$ instead of 3).

Heuristically, it allows one to treat the process as if the points are independent, ignoring the repulsion that DPPs actually exhibit.

For a DPP ($\lambda_1 > \lambda_2 > \ldots > \lambda_n$) with kernel K_n,

$$\mathbb{P}(\lambda_1 \leq t) = \det\left(I_n - K_n^t\right) = \prod_{i=1}^{n}(1 - \rho_i^t)$$

$$\leq \exp\left(-\sum_{i=1}^{n} \rho_i^t\right)$$

.

Widom's trick

Widom's trick provides a simpler way to obtain lower tail estimates for determinantal point processes, at the cost of a worse tail exponent ($\frac{3}{2}$ instead of 3).

Heuristically, it allows one to treat the process as if the points are independent, ignoring the repulsion that DPPs actually exhibit.

For a DPP ($\lambda_1 > \lambda_2 > \ldots > \lambda_n$) with kernel K_n,

$$\begin{aligned}\mathbb{P}(\lambda_1 \leq t) = \det\left(I_n - K_n^t\right) &= \prod_{i=1}^{n}(1 - \rho_i^t) \\ &\leq \exp\left(-\sum_{i=1}^{n} \rho_i^t\right) \\ &= \exp\left(-\mathrm{Tr}\, K_n^t\right).\end{aligned}$$

Analysing the trace

There are many representation for the trace of K_N^t, eg. in terms of the kernel K itself, which leads to contour integral representations.

Because of the uniformity in q that we require, contour integral representations are hard to extract asymptotics from.

Instead, a different representation will be useful for us. It involves the mean empirical spectral distribution $\nu_{q,N}$ of the DPP:

$$\nu_{q,N}(A) = \mathbb{E}\left[\frac{1}{N}\sum_{i=1}^{N} \mathbb{1}_{\lambda_i \in A}\right].$$

Analysing the trace

There are many representation for the trace of K_N^t, eg. in terms of the kernel K itself, which leads to contour integral representations.

Because of the uniformity in q that we require, contour integral representations are hard to extract asymptotics from.

Instead, a different representation will be useful for us. It involves the mean empirical spectral distribution $\nu_{q,N}$ of the DPP:

$$\nu_{q,N}(A) = \mathbb{E}\left[\frac{1}{N}\sum_{i=1}^{N} \mathbb{1}_{\lambda_i \in A}\right].$$

It turns out that

$$\operatorname{Tr} K_N^t = \nu_{q,N}\left([t,\infty)\right).$$

Since $\mathbb{P}(\lambda_1 \leq t) \leq \exp(-\operatorname{Tr} K_N^t)$, we want a lower bound on the upper tail of $\nu_{q,N}$.

An argument of Ledoux

Ledoux has an argument, introduced for GUE, that converts sharp asymptotics on moments of $\nu_{q,N}$ to such a lower bound. Let $X \sim \nu_{q,N}$.

An argument of Ledoux

Ledoux has an argument, introduced for GUE, that converts sharp asymptotics on moments of $\nu_{q,N}$ to such a lower bound. Let $X \sim \nu_{q,N}$.

First, by Cauchy-Schwarz,

$$\mathbb{E}\left[X^{2k}\mathbb{1}_{X\geq t}\right] \leq \mathbb{E}\left[X^{4k}\right]^{1/2} \cdot \mathbb{P}(X \geq t)^{1/2}.$$

An argument of Ledoux

Ledoux has an argument, introduced for GUE, that converts sharp asymptotics on moments of $\nu_{q,N}$ to such a lower bound. Let $X \sim \nu_{q,N}$.

First, by Cauchy-Schwarz,

$$\mathbb{E}\left[X^{2k}\mathbb{1}_{X\geq t}\right] \leq \mathbb{E}\left[X^{4k}\right]^{1/2} \cdot \mathbb{P}(X \geq t)^{1/2}.$$

Also,

$$\begin{aligned}\mathbb{E}\left[X^{2k}\mathbb{1}_{X\geq t}\right] &= \mathbb{E}\left[X^{2k}\right] - \mathbb{E}\left[X^{2k}\mathbb{1}_{X<t}\right] \\ &\geq \mathbb{E}\left[X^{2k}\right] - \mathbb{E}\left[X^{k}\right] \cdot t^{k}.\end{aligned}$$

An argument of Ledoux

Ledoux has an argument, introduced for GUE, that converts sharp asymptotics on moments of $\nu_{q,N}$ to such a lower bound. Let $X \sim \nu_{q,N}$.

First, by Cauchy-Schwarz,

$$\mathbb{E}\left[X^{2k}\mathbb{1}_{X \geq t}\right] \leq \mathbb{E}\left[X^{4k}\right]^{1/2} \cdot \mathbb{P}(X \geq t)^{1/2}.$$

Also,

$$\begin{aligned} \mathbb{E}\left[X^{2k}\mathbb{1}_{X \geq t}\right] &= \mathbb{E}\left[X^{2k}\right] - \mathbb{E}\left[X^{2k}\mathbb{1}_{X<t}\right] \\ &\geq \mathbb{E}\left[X^{2k}\right] - \mathbb{E}\left[X^{k}\right] \cdot t^{k}. \end{aligned}$$

So overall,

$$\mathbb{P}\left(X \geq t\right) \geq \mathbb{E}\left[X^{4k}\right]^{-1}\left(\mathbb{E}\left[X^{2k}\right] - \mathbb{E}\left[X^{k}\right] \cdot t^{k}\right)^{2}.$$

An argument of Ledoux

Ledoux has an argument, introduced for GUE, that converts sharp asymptotics on moments of $\nu_{q,N}$ to such a lower bound. Let $X \sim \nu_{q,N}$.

First, by Cauchy-Schwarz,

$$\mathbb{E}\left[X^{2k}\mathbb{1}_{X\geq t}\right] \leq \mathbb{E}\left[X^{4k}\right]^{1/2} \cdot \mathbb{P}(X \geq t)^{1/2}.$$

Also,

$$\begin{aligned}\mathbb{E}\left[X^{2k}\mathbb{1}_{X\geq t}\right] &= \mathbb{E}\left[X^{2k}\right] - \mathbb{E}\left[X^{2k}\mathbb{1}_{X<t}\right] \\ &\geq \mathbb{E}\left[X^{2k}\right] - \mathbb{E}\left[X^{k}\right] \cdot t^{k}.\end{aligned}$$

So overall,

$$\mathbb{P}\left(X \geq t\right) \geq \mathbb{E}\left[X^{4k}\right]^{-1}\left(\mathbb{E}\left[X^{2k}\right] - \mathbb{E}\left[X^{k}\right] \cdot t^{k}\right)^{2}.$$

So, with sharp upper and lower bounds on high moments of X, we can optimize over k and obtain a lower bound.

$$\mathbb{P}\left(X \geq t\right) \geq \mathbb{E}\left[X^{4k}\right]^{-1}\left(\mathbb{E}\left[X^{2k}\right] - \mathbb{E}\left[X^{k}\right] \cdot t^{k}\right)^{2}$$

For GUE, one can get sharp estimates on $\mathbb{E}[X^k]$ using some simple recursions. Not so for Meixner!

$$\mathbb{P}\left(X \geq t\right) \geq \mathbb{E}\left[X^{4k}\right]^{-1}\left(\mathbb{E}\left[X^{2k}\right] - \mathbb{E}\left[X^{k}\right] \cdot t^{k}\right)^{2}$$

For GUE, one can get sharp estimates on $\mathbb{E}[X^k]$ using some simple recursions. Not so for Meixner!

However, with such estimates, taking $t = \mu_q N(1 - q^{1/6}\varepsilon)$, one gets

$$\mathbb{P}\left(X \geq \mu_q N(1 - q^{1/6}\varepsilon)\right) \geq \varepsilon^{3/2}.$$

Note that we need the $q^{1/6}$ factor in front of ε in the probability to get a bound independent of q!

This bound is sharp, and we also prove an upper bound of the same order. The $\frac{3}{2}$ here is the source of the $\frac{3}{2}$ tail exponents in our main theorems.

Sharp moment asymptotics for $\nu_{q,N}$

Theorem

Let $X \sim \nu_{q,N}$. For all large N and uniformly over essentially the whole range of k (up to δN) and $q \in (0,1)$,

$$\mathbb{E}[X^k] = (q^{1/6}k)^{-3/2}(\mu_q N)^k \exp\left(\pm\Theta\left(q^{1/2}\frac{k^3}{N^2}\right)\right).$$

The bulk of the technical work in our argument is obtaining the above estimate.

Sharp moment asymptotics for $\nu_{q,N}$

Theorem

Let $X \sim \nu_{q,N}$. For all large N and uniformly over essentially the whole range of k (up to δN) and $q \in (0,1)$,

$$\mathbb{E}[X^k] = (q^{1/6}k)^{-3/2}(\mu_q N)^k \exp\left(\pm\Theta\left(q^{1/2}\frac{k^3}{N^2}\right)\right).$$

The bulk of the technical work in our argument is obtaining the above estimate.

The starting point is a formula of Ledoux for the *factorial* moments of X:

$$\mathbb{E}[(X)_k] := \mathbb{E}\left[X(X-1)\cdots(X-k+1)\right] = \frac{q^k}{(1-q)^k}\frac{1}{N}\cdot\frac{1}{k+1}\sum_{i=0}^{k} q^{-i}\binom{k}{i}^2\cdot\frac{(N+k-i)!}{(N-i-1)!}.$$

We obtain asymptotics for the factorial moments from this using a Laplace method argument (though there is some delicacy because we need to allow k up to $\Theta(N)$).

Factorial to polynomial moments?

But going from factorial moments to polynomial moments is very non-trivial.

One reason is that when $k = \Theta(N)$, $(X)_k$ and X^k differ by an exponential-in-N factor:

$$\frac{X(X-1)\cdots(X-k+1)}{X^k} = \prod_{i=1}^{k-1}(1 - iX^{-1}) = \exp\left(-\sum_{i=1}^{k-1}\log(1 - iX^{-1})\right)$$

Factorial to polynomial moments?

But going from factorial moments to polynomial moments is very non-trivial.

One reason is that when $k = \Theta(N)$, $(X)_k$ and X^k differ by an exponential-in-N factor:

$$\frac{X(X-1)\cdots(X-k+1)}{X^k} = \prod_{i=1}^{k-1}(1 - iX^{-1}) = \exp\left(-\sum_{i=1}^{k-1}\log(1-iX^{-1})\right)$$

If $k = \alpha N$ and $X \approx \mu_q N$, we can evaluate this as a Riemann sum to obtain

$$\exp\left((\alpha-\mu_q)\log\left(1-\frac{\alpha}{\mu_q}\right)-\alpha\right).$$

Factorial to polynomial moments?

But going from factorial moments to polynomial moments is very non-trivial.

One reason is that when $k = \Theta(N)$, $(X)_k$ and X^k differ by an exponential-in-N factor:

$$\frac{X(X-1)\cdots(X-k+1)}{X^k} = \prod_{i=1}^{k-1}(1-iX^{-1}) = \exp\left(-\sum_{i=1}^{k-1}\log(1-iX^{-1})\right)$$

If $k = \alpha N$ and $X \approx \mu_q N$, we can evaluate this as a Riemann sum to obtain

$$\exp\left((\alpha-\mu_q)\log\left(1-\frac{\alpha}{\mu_q}\right)-\alpha\right).$$

This factor has to be exactly tracked and shown to cancel with an expression arising from asymptotics of $\mathbb{E}[(X)_k]$.

It would be interesting to find a more direct approach, perhaps by finding a formula for $\mathbb{E}[X^k]$ directly.

Summary and future directions

- We focused on $x_N(N)$, and not $x_N(T)$. In the infinite LPP problem, general T would correspond to the homogeneous domains being rectangular of dimension $N \times T$.
    - Requires developing uniform geometric LPP estimates in directions other than $(1, 1)$.
- Improving the tail exponent to 3: this just requires getting the same for geometric LPP (uniformly in q).
    - Possibly Riemann-Hilbert approaches could yield it.
    - Another approach is to use the geometric bootstrapping approach developed in Ganguly-H. This would also require uniform-in-q estimates with exponent $3/2$ for LPP in other directions.
- A LPP representation for the log-gamma polymer free energy?

Summary and future directions

- We focused on $x_N(N)$, and not $x_N(T)$. In the infinite LPP problem, general T would correspond to the homogeneous domains being rectangular of dimension $N \times T$.
    - Requires developing uniform geometric LPP estimates in directions other than $(1, 1)$.
- Improving the tail exponent to 3: this just requires getting the same for geometric LPP (uniformly in q).
    - Possibly Riemann-Hilbert approaches could yield it.
    - Another approach is to use the geometric bootstrapping approach developed in Ganguly-H. This would also require uniform-in-q estimates with exponent $3/2$ for LPP in other directions.
- A LPP representation for the log-gamma polymer free energy?

Thank you!