

From Yang–Baxter to Robinson–Schensted–Knuth

Overview talk: Probabilistic methods (notes in progress)

Leonid Petrov

Conference “Solvable lattice models, representation theory of quantum groups and algebraic combinatorics” at Mittag-Leffler, July 2026

Abstract

These notes accompany two lectures on the probabilistic point of view on solvable lattice models and algebraic combinatorics. The goal is to present the Yang–Baxter equation as a source of Markov operators that transport probability measures attached to vertex models, via local Yang–Baxter moves. We derive the classical Robinson–Schensted–Knuth correspondence directly from the Yang–Baxter equation for a monochrome vertex model connected to the symmetric Schur polynomials, and discuss generalizations related to (spin) q -Whittaker and (spin) Hall–Littlewood polynomials.

Contents

1	Vertex models for Schur polynomials	3
1.1	Conventions	3
1.2	Schur vertex model I: Horizontal strips	4
1.3	Skew Young diagrams and vertex model rows	7
1.4	Schur vertex model II: Vertical strips	8
1.5	Yang–Baxter equation and Cauchy identity	10
1.6	Summary: Three kinds of Schur vertex models	12
2	q-deformations: Hall–Littlewood, q-Whittaker, and fusion	13
2.1	Vertex model for Hall–Littlewood polynomials	13
2.2	Vertex model for q -Whittaker polynomials	15
2.3	Macdonald duality	15
2.4	Fusion	16
3	Bijectivization: From identities to Markov transitions	19
4	Deriving RSK from the Yang–Baxter equation	20
4.1	RSK via toggles	20
4.2	Yang–Baxter bijectivization in the Schur case	21
5	Deformed weights and particle systems	24
5.1	The q -Whittaker case and q -PushTASEP	24
5.2	Hall–Littlewood: a row-indexed RSK and the six-vertex model	27

6	Further directions	30
6.1	Stationary measures via the forgetting property	30
6.2	Symmetry in the spectral parameters: parameter permutation and time reversal . .	31
6.3	Shift-invariance and hidden symmetries	32
6.4	Markov duality	32

Introduction

The Yang–Baxter equation is, on its face, an algebraic identity, which expresses that the partition functions of two configurations differing by a local transformation are equal. When can such an identity be upgraded into an algorithm, or into a random transformation? When both sides are sums of nonnegative weights, $\sum_a \text{wt}(a) = \sum_b \text{wt}(b)$, we *randomize* the identity by introducing a transformation — a Markov step — which turns a configuration on the left into a randomly chosen configuration on the right. This randomization must be consistent with the weights, but its choice need not be unique. Carried across the lattice, these random steps are interpreted as a random growth process. We call this passage from an identity to a randomized transformation a *bijectivization* or a *probabilistic bijection*. It is an application, in a solvable lattice setting, of the general principle of *coupling* in probability theory [Lin02], [Tho00].

For the Yang–Baxter equation, the bijectivization was introduced in [BP19], after two waves of randomized insertion algorithms generalizing the Robinson–Schensted–Knuth (RSK) correspondence: the q -Whittaker randomized RSK of [OP13], [BP16], and [MP17], and the Hall–Littlewood randomized RSK of [BP15], [BBW16], and [BM18]. The bijectivization can stay genuinely random, or collapse to a deterministic one.

These notes are expository, written from the probabilistic vantage point; we assume basic familiarity with RSK and Schur symmetric polynomials.




Figure 1: A single vertex has an *input* (the incoming occupations (i_1, j_1) on the bottom and left edges) and an *output* (the outgoing occupations (i_2, j_2) on the top and right). The Boltzmann weight of the vertex $w(i_1, j_1; i_2, j_2)$ is called *stochastic* when these weights are the conditional probabilities of the output given the input.

This circle of ideas around probabilistic applications of the Yang–Baxter equation grew out of three older subjects. From algebraic combinatorics it takes the Robinson–Schensted–Knuth (RSK) correspondence [Sch61], [Knu73] and its growth-diagram form [Fom88] and the theory of symmetric functions [Mac95]. From quantum integrable systems it takes the Yang–Baxter equation itself, the six-vertex [Lie67] and higher-spin models, the quantum inverse scattering method and algebraic Bethe ansatz of the Leningrad school [Fad96], and the quantum groups [Jim86] whose R -matrices produce the vertex weights. In particular, the presentation of Schur polynomials as free-fermionic lattice partition functions is classical [ZJ09]; its deformation (a Schur polynomial times a deformed

Weyl denominator) first appeared as a generating function over strict Gelfand–Tsetlin patterns [Tok88], with classical-group analogues [HK05]. These identities were recast as solvable six-vertex partition functions obeying the Yang–Baxter equation [BBF11], connecting the subject to p -adic representation theory. Finally, from probability it takes its founding solvable models: Ulam’s longest-increasing-subsequence problem and its Tracy–Widom asymptotic behavior due to Baik–Deift–Johansson [BDJ99], Okounkov’s Schur measures [Oko01], and exclusion processes [Spi70], [Lig75], [Ros81].

The Yang–Baxter bijectivization serves as a dictionary between symmetric functions and interacting particle systems. In the Schur case, the maximal part of the random partition, λ_1 , is a last-passage time [Joh00]. The structure of the Schur measure leads, in the edge scaling, to the GUE Tracy–Widom law. Deformed models of random partitions (obtained by replacing the Schur polynomials by their generalizations from the Macdonald symmetric function hierarchy [BC14] and beyond) display connections to further particle systems generalizing last-passage percolation and TASEP (Totally Asymmetric Simple Exclusion Process). In particular, one can directly utilize the Yang–Baxter bijectivization to construct these particle systems: the q -TASEP [BC14] and q -PushTASEP [MP17], the stochastic six-vertex model and ASEP [BCG16], and integrable models of directed random polymers, including the log-gamma [Sep12] and O’Connell–Yor [OY01] polymers.

Let us also mention that in rigorous two-dimensional statistical mechanics, the random star-triangle transformation (itself a bijectivization of the Yang–Baxter equation) underlies the proof of rotational invariance of the critical Ising, FK, and six-vertex models [DCKK⁺20].

Outline

In Section 1 we recall the three vertex models whose partition functions are the Schur symmetric polynomials. We also write out the Yang–Baxter equation for one of the models, and utilize it to prove the Cauchy summation identity. Section 2 describes the t - and q -deformations of the Schur vertex models (leading, respectively, to the Hall–Littlewood and q -Whittaker polynomials). The Hall–Littlewood and q -Whittaker polynomials are connected by the Macdonald duality, which at the level of vertex models can be realized as the fusion procedure. This section may be skipped if the reader is only interested in deriving the Robinson–Schensted–Knuth correspondence from the Yang–Baxter equation. Section 3 introduces the general bijectivization procedure. In Section 4 we apply it to the Schur vertex model and derive the Robinson–Schensted–Knuth correspondence. In the optional Section 5 we indicate how the deformed vertex weights of Section 2 and the bijectivization yield stochastic particle systems generalizing the Schur-level objects, the last-passage percolation and TASEP; this part is only sketched. Finally, Section 6 points to further directions in which a probabilistic reading of the Yang–Baxter equation has been fruitful in recent research.

1 Vertex models for Schur polynomials

1.1 Conventions

A *signature* of length N is a weakly decreasing integer vector $\lambda = (\lambda_1 \geq \dots \geq \lambda_N)$, $\lambda_i \in \mathbb{Z}$; the set of these is Sign_N . A *partition* is a signature with $\lambda_N \geq 0$, identified with its Young diagram (the left-justified array of λ_i boxes in row i , drawn top to bottom in the English convention); $|\lambda| = \sum_i \lambda_i$ is its size, $\ell(\lambda)$ its number of nonzero parts, and $s_\lambda(x_1, \dots, x_n)$ is the Schur polynomial. We write $\mathbb{Z}_{\geq 0} = \{0, 1, 2, \dots\}$. The *conjugate* (or transpose) partition λ' is the reflection of the diagram across the main diagonal,

$$\lambda'_i = \#\{j : \lambda_j \geq i\} \quad (\text{the height of the } i\text{th column of } \lambda),$$

so that $(\lambda')' = \lambda$ and $|\lambda'| = |\lambda|$. Our running example is

$$\lambda = (4, 4, 1), \quad \lambda' = (3, 2, 2, 2),$$

displayed in Figure 2.

Strips and interlacing. Two partitions $\mu \subseteq \lambda$ (diagram inclusion) differ by a *horizontal strip* if their skew shape λ/μ has at most one box in each *column*, and by a *vertical strip* if λ/μ has at most one box in each *row*. These two conditions are exchanged by conjugation, and each is encoded by an interlacing relation on signatures. We write $\mu \prec \lambda$ (equivalently $\lambda \succ \mu$) for the *horizontal* interlacing

$$\lambda_1 \geq \mu_1 \geq \lambda_2 \geq \mu_2 \geq \cdots,$$

which for partitions says exactly that λ/μ is a horizontal strip. The *transposed* (vertical) interlacing is its conjugate,

$$\mu \prec' \lambda \stackrel{\text{def}}{\iff} \mu' \prec \lambda', \quad \text{i.e.} \quad \lambda'_1 \geq \mu'_1 \geq \lambda'_2 \geq \mu'_2 \geq \cdots.$$

See Figure 3 for an illustration.




Figure 2: The partition $\lambda = (4, 4, 1)$ (left), its conjugate $\lambda' = (3, 2, 2, 2)$ obtained by reflecting across the main diagonal (middle), and a semistandard Young tableau T of shape λ with entries in $\{1, \dots, n\}$, $n = 3$ (right): rows weakly increase left to right, columns strictly increase top to bottom. The cells filled with a given entry k form a horizontal strip (at most one per column), so T records the horizontal-strip interlacing sequence $\emptyset \prec (2) \prec (4, 1) \prec (4, 4, 1)$. The weight of T is $x^{\text{wt}(T)} = x_1^2 x_2^3 x_3^4$, and $s_\lambda = \sum_T x^{\text{wt}(T)}$ is the sum of weights over all semistandard tableaux.



Figure 3: Left: a horizontal strip $\mu \prec \lambda$ (shaded), at most one box per column. Right: a vertical strip $\mu \prec' \lambda$ (shaded), at most one box per row. Conjugation exchanges the two.

Weights. Vertex weights are taken *nonnegative* — and that is all the Yang–Baxter bijectivization of Section 3 will require of them. A model is moreover called *stochastic* when its weights, read as conditional probabilities of the outgoing edges given the incoming ones, sum to one, so that the row-to-row transfer matrix is a Markov operator; this is a convenient extra structure used in connection with interacting particle systems.

1.2 Schur vertex model I: Horizontal strips

A solvable lattice model assigns Boltzmann weights to the local configurations on the square grid, with paths (arrows) conserved at each vertex, and the weight of a whole configuration is the product

of its vertex weights. The most classical instance is the *free-fermion five-vertex model*, whose partition function is a Schur polynomial. We describe it concretely, following Aggarwal–Borodin–Petrov–Wheeler [ABPW23] and Mucciconi–Petrov [MP22]; this identity is classical [ZJ09].

Fix n and work on the grid $\mathbb{Z}_{\geq 1} \times \{1, \dots, n\}$: n horizontal rows indexed from the bottom by $i = 1, \dots, n$, and columns indexed by a half-line coordinate. Each edge carries an occupation in $\{0, 1\}$, and we encode a local vertex by its four edge occupations

$$(i_1, j_1; i_2, j_2) = (\text{bottom in, left in; top out, right out}), \quad i_1, j_1, i_2, j_2 \in \{0, 1\}. \quad (1.1)$$

Path conservation forces the weight to vanish unless $i_1 + j_1 = i_2 + j_2$, leaving six admissible vertices; the *five-vertex* condition discards the doubly-occupied vertex $(1, 1; 1, 1)$. The remaining five vertices, and the Boltzmann weights attached to a vertex sitting in row i , are shown in Figure 4.




Figure 4: Top: the edge-occupation convention of (1.1) — a vertex in row i with vertical-in i_1 (bottom), horizontal-in j_1 (left), vertical-out i_2 (top), horizontal-out j_2 (right); paths run up and to the right. Bottom: the five admissible vertices with their Boltzmann weights in row i .

Remark 1.1 (Free fermion condition). The six Boltzmann weights (the five admissible ones together with the discarded doubly-occupied vertex of weight 0) satisfy the *free fermion condition*

$$w\left(\begin{array}{c} | \\ | \\ | \end{array}\right) \cdot w\left(\begin{array}{c} | \\ + \\ | \end{array}\right) + w\left(\begin{array}{c} | \\ \rightarrow \\ | \end{array}\right) \cdot w\left(\begin{array}{c} \uparrow \\ | \\ | \end{array}\right) = w\left(\begin{array}{c} \uparrow \\ | \\ | \end{array}\right) \cdot w\left(\begin{array}{c} | \\ + \\ | \end{array}\right). \quad (1.2)$$

A signature $\lambda = (\lambda_1 \geq \dots \geq \lambda_N)$ enters as boundary data through the “parts + staircase” point set

$$\mathcal{S}(\lambda) = \{\lambda_j + N - j + 1\}_{1 \leq j \leq N} \subset \mathbb{Z}_{\geq 1}, \quad (1.3)$$

a strictly increasing set of occupied sites; here N is the number of parts of λ . For the partition function Z_λ we take a single block of n rows on which λ has $N = n$ parts, feed in n paths from the bottom at the positions $\mathcal{S}(\lambda)$, require them to exit through the right boundary, and keep the top and left boundaries empty. Summing the product of weights from Figure 4 over all admissible five-vertex configurations with this boundary, by definition, is the partition function $Z_\lambda(x_1, \dots, x_n)$.

Proposition 1.2 (Schur polynomial as a five-vertex partition function). *We have*

$$s_\lambda(x_1, \dots, x_n) = Z_\lambda(x_1, \dots, x_n).$$

Sketch of proof. One can use a straightforward bijection between five-vertex configurations and semistandard Young tableaux which preserves the weights $x^{\text{wt}(T)}$ up to reversing the variables

$x_i \mapsto x_{n+1-i}$ (the reversal does not affect the partition function, which is symmetric in the x_i 's). See Figure 5 for an illustration of a vertex model configuration corresponding to the semistandard Young tableau from Figure 2. Note that each row of the vertex model corresponds to a horizontal strip in the tableau. $\square$




Figure 5: The five-vertex configuration corresponding to the semistandard tableau T of Figure 2 of shape $\lambda = (4, 4, 1)$ with $n = 3$.

Remark 1.3. Fixing an additional signature μ and reading the paths between the bottom boundary $\mathcal{S}(\lambda)$ and the top boundary $\mathcal{S}(\mu)$ yields the skew Schur polynomial $s_{\lambda/\mu}$.

Remark 1.4 (Higher spin version). The free-fermion five-vertex model described above can be bijectively mapped to a free-fermion instance of the higher spin six-vertex model studied in [Bor17], [BP18]. The allowed configurations in this model have $i_1, i_2 \in \mathbb{Z}_{\geq 0}$ and $j_1, j_2 \in \{0, 1\}$. The vertex weights (in the Schur case) in the i -th row are given by

$$w(i_1, j_1; i_2, j_2) = \mathbf{1}_{i_1+j_1=i_2+j_2} x_i^{j_2}.$$

The configuration which bijectively corresponds to the one in Figure 5 is shown in Figure 6. Now the top boundary is given by the multiset $\{\lambda_1, \lambda_2, \dots, \lambda_n\}$, and the paths enter everywhere from the left boundary, while the bottom and the right boundaries are empty. In contrast with the five-vertex model, the columns are now indexed by $\mathbb{Z}_{\geq 0}$ (zero parts of λ occupy column 0); this indexing is needed for the weights $x_i^{j_2}$ to produce the correct total power of x_i in each row.




Figure 6: Left: the four vertex types $(j_1, j_2) \in \{0, 1\}^2$, labelled $(i_1, j_1; i_2, j_2)$, with weights $w = x_i^{j_2}$. The vertical occupation $i_1, i_2 \geq 0$ is arbitrary, subject to $i_1 + j_1 = i_2 + j_2$; here we draw the representatives with $i_1 = 3$. Right: the higher-spin six-vertex configuration corresponding to the tableau T (equivalently, to the five-vertex configuration of Figure 5); the columns are indexed from 0.

Remark 1.5 (The free-fermion six-vertex model, and the Aztec diamond). Allowing the vertex $(1, 1; 1, 1)$ turns the five-vertex model into the *six-vertex model*. Keeping the free fermion condition (1.2), one can define free-fermion Schur functions [ABPW23], [Nap24], which generalize both horizontal- and vertical-strip Young tableaux (the vertical strip setting is detailed in Section 1.4 below). This family of functions includes as particular cases the factorial and shifted Schur polynomials, as well as the Frobenius–Schur functions.

1.3 Skew Young diagrams and vertex model rows

Drop the condition $j_1, j_2 \leq 1$ on the horizontal edges of the higher-spin model (Remark 1.4), keeping $i_1, i_2 \in \mathbb{Z}_{\geq 0}$ and the arrow conservation $i_1 + j_1 = i_2 + j_2$ at each vertex. Consider a single row of vertices on this more general vertex model. We read configurations of vertical arrows as multisets, and interpret them as partitions λ (at the top) and μ (at the bottom). In particular, the vertical occupation i_2 at the top of column x is the multiplicity $m_x(\lambda) = \#\{i : \lambda_i = x\}$. Writing h_x for the occupation of the horizontal edge entering column x from the left, we have

$$h_x = \sum_{y \geq x} (m_y(\lambda) - m_y(\mu)) = \lambda'_x - \mu'_x.$$

The horizontal edges are thus nonnegative if and only if $\lambda'_x \geq \mu'_x$ for all x , which is equivalent to the Young diagram inclusion $\mu \subseteq \lambda$.




Figure 7: The dictionary between general vertex models and skew Young diagrams.

Restricting λ/μ to horizontal strips recovers the higher-spin model of Remark 1.4. Restricting to vertical strips, we obtain another combinatorial condition on the allowed vertex types:

Lemma 1.6. *For a vertex row configuration corresponding to a skew Young diagram λ/μ , the inequality $j_2 \leq i_1$ holds at every vertex if and only if λ/μ is a vertical strip.*

Proof. At column x the bottom vertical edge carries $i_1 = m_x(\mu) = \mu'_x - \mu'_{x+1}$ and the outgoing horizontal edge carries $j_2 = h_{x+1} = \lambda'_{x+1} - \mu'_{x+1}$, so

$$j_2 \leq i_1 \iff \lambda'_{x+1} - \mu'_{x+1} \leq \mu'_x - \mu'_{x+1} \iff \lambda'_{x+1} \leq \mu'_x.$$

This completes the proof. $\square$

The vertical-strip model is discussed in detail in Section 1.4 below.

Remark 1.7. The five-vertex model is not a direct specialization of the model in Remark 1.4, but it is a bijective change of coordinates from the horizontal-strip case, which corresponds to the shifted coordinates $\mathcal{S}(\lambda)$ (1.3).

The rest of this subsection is a digression, not used later in the notes. In the shifted coordinates (1.3), the six-vertex model of Remark 1.5 (that is, with the doubly occupied vertex $(1, 1; 1, 1)$ allowed) captures one more distinguished class of skew Young diagrams — *ribbons* (border strips): connected skew diagrams containing no 2×2 block of boxes.

Lemma 1.8 (Ribbon condition). *Let λ, μ be partitions with at most N parts, both encoded as in (1.3) with this common N . Consider a single row of the six-vertex model (all edges carry occupations in $\{0, 1\}$, and all six vertices of Figure 4, including $(1, 1; 1, 1)$, are admissible) with bottom boundary $\mathcal{S}(\mu)$, top boundary $\mathcal{S}(\lambda)$, and empty left and right horizontal edges. An admissible configuration exists if and only if $\mu \subseteq \lambda$ and every connected component of λ/μ is a ribbon (equivalently, λ/μ contains no 2×2 block of boxes). The configuration is then unique; see Figure 8.*

Proof. The configuration is forced by the boundary data: the occupation h_x of the horizontal edge between columns x and $x + 1$ must equal

$$h_x = \#(\mathcal{S}(\mu) \cap \{1, \dots, x\}) - \#(\mathcal{S}(\lambda) \cap \{1, \dots, x\}) = \#\{i: \lambda_i - i \geq x - N\} - \#\{i: \mu_i - i \geq x - N\}, \quad (1.4)$$

and any assignment with all $h_x \in \{0, 1\}$ is admissible. Row i of λ/μ contains a box of content $x - N$ (content = column - row) if and only if $\mu_i - i < x - N \leq \lambda_i - i$. Hence $h_x \geq 0$ for all x is equivalent to $\mu \subseteq \lambda$, and then h_x counts the boxes of λ/μ of content $x - N$. The condition $h_x \leq 1$ states that every diagonal of λ/μ carries at most one box, which is equivalent to the absence of a 2×2 block: two boxes of equal content force boxes of that content in all intermediate rows, and two of them in adjacent rows produce a 2×2 block. $\square$




Figure 8: Left: the unique single-row six-vertex configuration with bottom $\mathcal{S}(\mu) = \{1, 4, 6, 7, 10\}$ and top $\mathcal{S}(\lambda) = \{2, 6, 7, 8, 12\}$, where $\mu = (5, 3, 3, 2)$, $\lambda = (7, 4, 4, 4, 1)$, and $N = 5$. Right: the disconnected skew shape λ/μ , a disjoint union of three ribbons of heights 1, 3, 1, with the content of each box indicated; the components do not even share corners, being separated by the empty diagonals of contents $-3, -2$ and $3, 4$. The horizontal edge between columns x and $x + 1$ is occupied precisely when λ/μ has a box of content $x - N$. The path emitted at column 4 jumps over the two stationary particles at columns 6 and 7 — the doubly occupied vertices $(1, 1; 1, 1)$ — which match the two boxes with a bottom neighbor (contents 1 and 2) in the height-3 ribbon, in agreement with Remark 1.9.

Compare (1.4) with the identity $h_x = \lambda'_x - \mu'_x$ from the beginning of this subsection: in the multiplicity coordinates the horizontal edges count the columns of λ/μ , while in the shifted coordinates they count the diagonals.

Remark 1.9 (Murnaghan–Nakayama structure). In the configuration of Lemma 1.8, the number of vertices of type $(1, 0; 0, 1)$ equals the number of ribbon components of λ/μ , and the number of doubly occupied vertices $(1, 1; 1, 1)$ equals $\sum_r (\text{ht}(r) - 1)$, where $\text{ht}(r)$ is the number of rows of the ribbon r (the doubly occupied columns correspond to the boxes of λ/μ having a bottom neighbor). Thus, assigning weight -1 to the vertex $(1, 1; 1, 1)$ produces the signs of the Murnaghan–Nakayama rule [Mac95, Ch. I, §7]. Matrix elements of the free-fermion six-vertex transfer matrices, with weights factorized over the ribbons of λ/μ , are computed in [Nap24, §4.2]; for the Yang–Baxter algebra of the general asymmetric six-vertex model in terms of ribbons, leading to cylindric Murnaghan–Nakayama rules for Hecke characters, see [Kor21, §4.3].

1.4 Schur vertex model II: Vertical strips

We now describe the third Schur vertex model which operates with the transposed Young diagram, and corresponds to vertical-strip interlacing sequences as described in Section 1.3 above. We follow [MP22]. This vertex model may be viewed as a “dual” to the free-fermion five-vertex model of Section 1.2 (and its higher-spin equivalent version from Remark 1.4).

We encode top and bottom boundary conditions of this new vertex model by Young diagrams. Namely, to a Young diagram λ we associate the multiset of its column lengths via

$$\lambda_i - \lambda_{i+1} = \#\{\text{paths at the } i\text{-th vertical edge}\}, \quad i = 1, 2, \dots \quad (1.5)$$

For example, the Young diagram $\lambda = (4, 4, 1)$ corresponds to the configuration with three vertical edges in column 2, one vertical edge in column 3, and no vertical edges in the other columns.

Remark 1.10. Note that the convention (1.5) used for our second Schur vertex model differs from the one of Section 1.3 by conjugation. Therefore, in the second model, we still get *horizontal-strip differences* between consecutive rows of the vertex model. In the Schur case this change of convention is immaterial, since Schur and skew Schur functions are invariant under conjugation, $s_{\lambda'/\mu'} = s_{\lambda/\mu}$. However, we follow the existing literature [BW21], [MP22] which deals with deformations of the models where this conjugation is meaningful. We discuss some of the deformations in Sections 2 and 5 below.

We now define two vertex models, exchanged by the reflection $i_1 \leftrightarrow i_2$ of the vertical axis. The *up-right* (blue) weight is

$$W(i_1, j_1; i_2, j_2) = \mathbf{1}_{i_1+j_1=i_2+j_2} \mathbf{1}_{i_1 \geq j_2} x^{j_2}, \quad (1.6)$$

with x the spectral parameter of the row, and the *down-right* (red) weight is its reflection,

$$W^*(i_1, j_1; i_2, j_2) = \mathbf{1}_{i_2+j_1=i_1+j_2} \mathbf{1}_{i_2 \geq j_2} x^{j_2}. \quad (1.7)$$

Here $i_1, i_2, j_1, j_2 \in \mathbb{Z}_{\geq 0}$.

Encode the bottom and top boundaries of a single row by the column-length coordinates (1.5) of two partitions μ and λ . The condition $i_1 \geq j_2$ in (1.6) (respectively $i_2 \geq j_2$ in (1.7)) is equivalent to the requirement that λ/μ be a *horizontal strip*, see Lemma 1.6 and Remark 1.10. The skew functions $G_{\lambda/\mu}(x)$ and $G_{\lambda/\mu}^*(x)$ are defined as the weights of the single rows with these boundary conditions (for a single row, there is only one admissible configuration in both vertex models).

Gluing k rows with spectral parameters $y_1, \dots, y_k$ produces $G_{\lambda/\mu}(y_1, \dots, y_k)$ and $G_{\lambda/\mu}^*(y_1, \dots, y_k)$ in the similar manner.

Remark 1.11. The horizontal arrows on the left boundary in the models for $G_{\lambda/\mu}$ and $G_{\lambda/\mu}^*$ may be arbitrary. This may be modeled by placing an infinite number of vertical edges at column 0 with occupation $i_1 = i_2 = \infty$. Then, the weights become $W(\infty, j_1; \infty, j_2) = x^{j_2}$, which are independent of j_1 (and similarly for W^*). These boundary weights are the correct choice for our partition functions. This convention is also equivalent to setting $\lambda_0 = +\infty$, see (1.5).

Configurations contributing a nonzero weight to $G_{\lambda/\mu}$ or $G_{\lambda/\mu}^*$ are in bijection with interlacing chains of partitions $\mu = \lambda^{(0)} \prec \lambda^{(1)} \prec \dots \prec \lambda^{(k)} = \lambda$, with the i -th row creating the horizontal strip $\lambda^{(i)}/\lambda^{(i-1)}$. Unlike for the models in Section 1.2, both G and G^* are defined for any number of variables and are stable under $y_k \rightarrow 0$. On the other hand, the number of variables k must be at least the length of the longest column of λ/μ for the weight to be nonzero.

Proposition 1.12. *For partitions $\mu \subseteq \lambda$ and any number of variables, we have*

$$G_{\lambda/\mu}(y_1, \dots, y_k) = G_{\lambda/\mu}^*(y_1, \dots, y_k) = s_{\lambda/\mu}(y_1, \dots, y_k), \quad G_\lambda := G_{\lambda/\emptyset}.$$

Proof. Follows by identifying vertex model configurations with chains of interlacing partitions (differing by horizontal strips), equivalently, with semistandard Young tableaux of shape λ/μ . An illustration of this identification is given in Figure 9. The vertex weights (1.6) and (1.7) produce the correct weight $y^{\text{wt}(T)}$ for each tableau T entering the (skew) Schur polynomial. $\square$




Figure 9: A configuration of the vertex model W (1.6) and the corresponding semistandard Young tableau. The shape is $\lambda = (5, 3, 2, 1, 0)$, and the weight is $y_1 y_2^2 y_3^3 y_4^3 y_5^2 = y^{\text{wt}(T)}$. No path in the vertex model can go below the diagonal thanks to the condition $i_1 \geq j_2$ in (1.6). The thick edge at column 0 carries the infinite vertical occupation (see Remark 1.11), from which the paths peel off to the right.

1.5 Yang–Baxter equation and Cauchy identity

Let us now formulate the *Yang–Baxter equation* connecting the models W and W^* . Define the cross-vertex weights by

$$\mathbb{R}_{x,y}(i_1, j_1; i_2, j_2) = \mathbb{R}_{x,y} \left(\begin{array}{ccc} & j_1 & \\ & \searrow & \nearrow \\ i_1 & & i_2 \\ & \nearrow & \searrow \\ & j_2 & \end{array} \right) = \mathbf{1}_{i_1+j_2=j_1+i_2} (xy)^{\min(i_2, j_2)}, \quad (1.8)$$

with (i_1, j_1) the incoming and (i_2, j_2) the outgoing arrow counts.

Proposition 1.13 (Yang–Baxter equation). *For all boundary occupations $i_1, i_2, i_3, j_1, j_2, j_3 \in \mathbb{Z}_{\geq 0}$ and spectral parameters x, y , we have*

$$\begin{aligned} \sum_{k_1, k_2, k_3 \geq 0} \mathbb{R}_{x,y}(i_2, i_1; k_2, k_1) W(k_3, k_2; j_3, j_2) W^*(i_3, k_1; k_3, j_1) \\ = \sum_{k'_1, k'_2, k'_3 \geq 0} W(i_3, i_2; k'_3, k'_2) W^*(k'_3, i_1; j_3, k'_1) \mathbb{R}_{x,y}(k'_2, k'_1; j_2, j_1), \end{aligned} \quad (1.9)$$

where $W = W_x$ (1.6) and $W^* = W_y^*$ (1.7) carry the row parameters x and y . Graphically, the two local vertex configurations of Figure 10 have equal partition functions: the $\mathbb{R}$ -cross may be dragged across the W/W^* pair while preserving the six boundary edges.

Proof. Follows by a direct computation with explicit weights. □

Let us now use the weights W and W^* to construct a vertex model for the sum

$$\sum_{\lambda} s_{\lambda}(x_1, \dots, x_N) s_{\lambda}(y_1, \dots, y_N), \quad (1.10)$$




Figure 10: The Yang–Baxter equation (1.9) of Proposition 1.13. The boundary arrow counts $i_1, i_2, i_3, j_1, j_2, j_3 \in \mathbb{Z}_{\geq 0}$ are fixed, and the internal arrow counts k_1, k_2, k_3 and k'_1, k'_2, k'_3 are summed over. The vertical edge on the left or on the right serves as a source or a sink, respectively.

where λ runs over all partitions with at most N parts. That is, we realize the Schur polynomial $s_\lambda(x_1, \dots, x_N)$ as the partition function of the up-right model with N rows, weights W , and top boundary λ ; and we realize $s_\lambda(y_1, \dots, y_N)$ as the partition function of the down-right model with N rows, weights W^* , and bottom boundary λ . See Figure 11 for an illustration.

Remark 1.14. The sums in the Yang–Baxter equation (1.9) are finite for any fixed finite boundary $i_1, i_2, i_3, j_1, j_2, j_3 \in \mathbb{Z}_{\geq 0}$ (see Figure 10), so the Yang–Baxter equation is a polynomial identity in x, y . However, for the Cauchy sum (1.10) we employ the infinite boundary condition in the zeroth column, so the sum runs over all partitions (with at most N parts), and is a power series in the x_i, y_j , convergent for $|x_i|, |y_j| < 1$. Alternatively, it may be treated algebraically as a formal series. The two viewpoints are essentially equivalent, but since we are dealing with probabilistic connections, we adopt the analytic one.




Figure 11: The vertex model for the Cauchy sum $\sum_\lambda s_\lambda(x_1, \dots, x_N) s_\lambda(y_1, \dots, y_N)$ with $N = 4$.

We now explain how to use the Yang–Baxter equation to evaluate the sum (1.10) as a finite product; the three steps are illustrated in Figure 12. Insert an auxiliary $N \times N$ square of the cross-vertices $\mathbb{R}_{x_i, y_j}$ to the right of the lattice in Figure 11. The crosses are initially empty thanks to the

boundary conditions $\emptyset$ on the right, so this inserted square must be empty, which contributes a factor of 1 to the partition function.

We then drag this square leftward through the lattice, one cross at a time, by the sequence of Yang–Baxter moves of Proposition 1.13 (see Figure 10). Each move preserves the partition function while altering the lattice by exchanging one blue line with one red. After N^2 such moves, one for each pair x_i, y_j , every cross has reached the empty left boundary, where its incoming occupations are $(0, 0)$. Thanks to the boundary conditions ∞ in the zeroth column, the outgoing occupations of each cross can be arbitrary (but equal), so each cross contributes

$$\sum_{a,b \geq 0} \mathbb{R}_{x_i, y_j}(0, 0; a, b) = \sum_{a,b \geq 0} \mathbf{1}_{a=b} (x_i y_j)^a = \frac{1}{1 - x_i y_j}.$$

See Figure 12 for an illustration. Meanwhile, the lattice to the right of the crosses, with the blue and red blocks interchanged, carries no paths at all (and thus has weight 1): an up-right path there would have to exit through the empty top boundary, and a down-right path through the empty bottom one. Thus, we obtain

Proposition 1.15. *For any N and $x_1, \dots, x_N, y_1, \dots, y_N$ with $|x_i|, |y_j| < 1$, we have*

$$\sum_{\lambda} s_{\lambda}(x_1, \dots, x_N) s_{\lambda}(y_1, \dots, y_N) = \prod_{i,j=1}^N \frac{1}{1 - x_i y_j}, \quad (1.11)$$

where the sum is over all partitions λ with at most N parts.

1.6 Summary: Three kinds of Schur vertex models

The Schur polynomial has now appeared as the partition function of three vertex models: the five-vertex model (Section 1.2), the higher spin vertex model with capacity one of the horizontal edges (Remark 1.4), and the vertical-strip model (Section 1.4) with arbitrary capacities at the edges, but with the restriction like $j_2 \leq i_1$. Each of the three classes of models is a member of a hierarchy of solvable lattice models:

- The five-vertex model (Section 1.2) sits inside the *free-fermion six-vertex model* (Remark 1.5, [ABPW23], [Nap24]). Furthermore, dropping the free fermion condition, we arrive at partition functions of general six-vertex models. Notably, this family also includes symmetric Grothendieck polynomials [MS13].
- The higher spin vertex model (Remark 1.4) was studied in connection with symmetric functions [Bor17], and then in connection with particle systems [CP16], [BP18]. This lifts the Schur polynomials to the Hall–Littlewood polynomials and further to the spin Hall–Littlewood symmetric rational functions.
- The vertical-strip model (Section 1.4) lifts the Schur polynomials to the q -Whittaker and spin q -Whittaker polynomials [BW21], [MP22], [BK24], [Kor24].

Despite the differences, these deformations are essentially instances of one general fully fused higher spin six-vertex model (connected to the quantum group $U_q(\widehat{\mathfrak{sl}}_2)$). We will not describe this general model in full generality, but in Section 2 below we will discuss some of the deformations of the second and the third models, and connect them together via the fusion procedure.




Figure 12: Illustration of the proof of the Cauchy identity (Proposition 1.15) by the Yang–Baxter equation. (a) Attaching the empty $N \times N$ square of $\mathbb{R}_{x_i, y_j}$ vertices to the right of the lattice in Figure 11. (b) Dragging one cross through the lattice by the Yang–Baxter equation. (c) After N^2 moves, all crosses have reached the left boundary (see Figure 13 for a smaller but more detailed example). Each cross contributes a geometric series to the partition function.

2 q -deformations: Hall–Littlewood, q -Whittaker, and fusion

This section is optional: it treats the t - and q -deformations and is not needed for the derivation of RSK in Section 4.

Putting the five-vertex model aside, we now focus on the two other Schur models of Section 1. Each of them carries deformations that preserve the Yang–Baxter equation and hence the Cauchy identity (proven exactly in the same way as in Section 1.5). We do not explore the full generality of the deformations (detailed in [Bor17], [BP18] and [BW21], [MP22], [BK24], [Kor24]), but focus on the symmetric-function cases of Hall–Littlewood and q -Whittaker polynomials.

2.1 Vertex model for Hall–Littlewood polynomials

The Hall–Littlewood weights are the t -deformation of the higher-spin six-vertex model of Remark 1.4: horizontal edges carry $j_1, j_2 \in \{0, 1\}$, the vertical edges have unbounded capacity, and arrows are conserved, $i_1 + j_1 = i_2 + j_2$. The four nonzero weights, with spectral parameter x , are

$$\begin{array}{cccc}
 \begin{array}{c} \uparrow \uparrow \uparrow \\ \text{---} \end{array} & \begin{array}{c} \uparrow \uparrow \uparrow \rightarrow \\ \text{---} \end{array} & \begin{array}{c} \uparrow \uparrow \uparrow \\ \text{---} \end{array} & \begin{array}{c} \uparrow \uparrow \uparrow \rightarrow \\ \text{---} \end{array} \\
 w = 1 & w = x & w = 1 & w = x(1 - t^g)
 \end{array} \tag{2.1}$$

where g denotes the occupation of the incoming (bottom) vertical edge ($g = 3$ in the pictures), and all other configurations are weighted 0. At $t = 0$, we recover the Schur weights from Figure 6.




Figure 13: An example of a configuration in the final step of the Yang–Baxter moves in the proof of the Cauchy identity (Proposition 1.15). The numbers along the edges indicate occupations. The framed numbers on the vertices indicate $\min(i_2, j_2)$, which enter the weights $\mathbb{R}_{x,y}$ (1.8). The edge occupations are uniquely determined by the framed numbers, thanks to the arrow preservation. The framed integers form a nonnegative integer matrix (the Robinson–Schensted–Knuth input matrix of Proposition 4.4), and the arrows recorded by the edge occupations are its Viennot shadow lines [Vie77], in the generalization from permutation matrices to arbitrary $\mathbb{Z}_{\geq 0}$ -matrices [Pra15].

Proposition 2.1 ([Tsi06], [FW09], [Bor17]). *For a partition λ , the partition function of the higher-spin model with N rows of spectral parameters $x_1, \dots, x_N$, weights (2.1), top boundary λ (given by the multiset $\{\lambda_1, \dots, \lambda_N\}$), fully packed left boundary, and empty bottom and right boundaries (see Figure 6) equals the Hall–Littlewood polynomial $P_\lambda(x_1, \dots, x_N; 0, t)$, up to a normalization independent of the x_i ’s.*

A standard reference on Hall–Littlewood and Macdonald polynomials is the book [Mac95]. Here and below, we use the notation $P_\lambda(\cdots; q, t)$ for the Macdonald polynomials. For $q = 0$, they reduce to the Hall–Littlewood polynomials. For $t = 0$, Macdonald polynomials become the q -Whittaker polynomials, which we discuss in Section 2.2 below. Both families become Schur polynomials at $q = t$ (which, in applications to vertex models, is taken to be $q = t = 0$).

Remark 2.2 (Spin Hall–Littlewood). Adjoining a *spin* parameter s deforms the four weights (2.1), in the same order, into

$$\frac{1 - sxt^g}{1 - sx}, \quad \frac{x - st^g}{1 - sx}, \quad \frac{1 - s^2t^g}{1 - sx}, \quad \frac{x(1 - t^g)}{1 - sx},$$

which return (2.1) at $s = 0$. Their partition functions are the *spin Hall–Littlewood* rational functions [Bor17], [BP18]. One may even let $s \rightsquigarrow s_j$ vary column by column without breaking the Yang–Baxter equation.

2.2 Vertex model for q -Whittaker polynomials

Now consider a q -deformation of the up-right weight W (1.6), whose edges carry arbitrary multiplicities $i_1, j_1, i_2, j_2 \in \mathbb{Z}_{\geq 0}$, subject to the conservation law $i_1 + j_1 = i_2 + j_2$ and the restriction $j_2 \leq i_1$. The q -deformed weight is defined as

$$W(i_1, j_1; i_2, j_2) = \mathbf{1}_{i_1+j_1=i_2+j_2} \mathbf{1}_{i_1 \geq j_2} x^{j_2} \frac{(q; q)_{i_1}}{(q; q)_{j_2} (q; q)_{i_1-j_2}}, \quad (2.2)$$

where $(a; q)_m = \prod_{k=0}^{m-1} (1 - aq^k)$ is the q -Pochhammer symbol (with $(a; q)_0 = 1$, and $(a; q)_m = \prod_{k=0}^{-m-1} (1 - aq^{m+k})^{-1}$ for $m < 0$). As $q \rightarrow 0$, the weight reduces to (1.6).

Remark 2.3. The ratio $(q; q)_{i_1} / (q; q)_{j_2} (q; q)_{i_1-j_2}$ is the q -binomial coefficient $\begin{bmatrix} i_1 \\ j_2 \end{bmatrix}_q$, a standard q -deformation of the binomial coefficient $\binom{i_1}{j_2}$. Note that the “classical” limit for us is $q \rightarrow 0$, not $q \rightarrow 1$.

Proposition 2.4. For a partition λ , the partition function of the up-right model with N rows of spectral parameters $x_1, \dots, x_N$, weights (2.2), top boundary λ and empty bottom boundary (see Figure 9) equals the q -Whittaker polynomial $P_\lambda(x_1, \dots, x_N; q, 0)$, up to a normalization independent of the x_i ’s.

Remark 2.5 (Spin and diagonal parameters). One can add a spin parameter s further deforming (2.2) into the spin q -Whittaker weight of [BW21], [MP22]:

$$W(i_1, j_1; i_2, j_2) = \mathbf{1}_{i_1+j_1=i_2+j_2} \mathbf{1}_{i_1 \geq j_2} x^{j_2} \frac{(-s/x; q)_{j_2} (-sx; q)_{i_1-j_2} (q; q)_{i_2}}{(q; q)_{j_2} (q; q)_{i_1-j_2} (s^2; q)_{i_2}}.$$

The corresponding partition functions are the spin q -Whittaker polynomials (unlike the spin Hall–Littlewood symmetric rational functions, these are polynomials in the x_i ’s). Going further, [BK24], [Kor24] enriched the spin q -Whittaker weights with extra *diagonal* parameters carried by the vertices; these have a representation-theoretic meaning and interpolation properties.

2.3 Macdonald duality

Here we illustrate how one can turn Hall–Littlewood polynomials into the q -Whittaker ones, and connect this construction to vertex models. First, consider the Cauchy identity for Hall–Littlewood polynomials [Mac95, Ch. III, §4]:¹

$$\sum_{\lambda} P_{\lambda}(x_1, \dots, x_N; 0, t) Q_{\lambda}(y_1, \dots, y_M; 0, t) = \prod_{i=1}^N \prod_{j=1}^M \frac{1 - t x_i y_j}{1 - x_i y_j}. \quad (2.3)$$

Here, Q_{λ} is a multiple of P_{λ} , where the coefficient depends on t and λ but not on the y_j ’s.

We now perform the following steps with (2.3):

- Replace each y_i by a finite geometric sequence $y_i, t y_i, \dots, t^{J-1} y_i$ of some length $J \in \mathbb{Z}_{\geq 1}$ (denote this finite geometric sequence by $y_i^{(J)}$). This is a finite version of the *plethystic substitution* appearing in symmetric function theory. On the vertex model side, this corresponds

¹This Cauchy identity also admits a Yang–Baxter proof, e.g., see [Bor17].

to *fusion* of J rows of the Hall–Littlewood vertex model, and dates back to [KRS81]. This turns the Cauchy identity (2.3) into

$$\sum_{\lambda} P_{\lambda}(x_1, \dots, x_N; 0, t) Q_{\lambda}(y_1^{(J)}, \dots, y_M^{(J)}; 0, t) = \prod_{i=1}^N \prod_{j=1}^M \frac{1 - t^J x_i y_j}{1 - x_i y_j}. \quad (2.4)$$

- Without proof, we claim that both sides of (2.4) depend on the parameters y_i and t^J analytically (this assertion would follow from the vertex model discussion in Section 2.4 below). In particular, we may perform the substitutions

$$y_i \rightsquigarrow 0, \quad t^J y_i \rightsquigarrow -z_i, \quad i = 1, \dots, M,$$

where z_i are new indeterminates. The Cauchy identity (2.4) becomes

$$\sum_{\lambda} P_{\lambda}(x_1, \dots, x_N; 0, t) F_{\lambda}(z_1, \dots, z_M) = \prod_{i=1}^N \prod_{j=1}^M (1 + x_i z_j). \quad (2.5)$$

- It remains to identify the symmetric functions $F_{\lambda}(z_1, \dots, z_M)$ in (2.5). Thanks to the Macdonald duality $\omega_{q,t}$ which interchanges q and t in the Macdonald polynomials (and transposes the partitions), the Hall–Littlewood and q -Whittaker polynomials satisfy the following dual Cauchy identity:

$$\sum_{\lambda} P_{\lambda}(x_1, \dots, x_N; 0, t) P_{\lambda'}(z_1, \dots, z_M; t, 0) = \prod_{i=1}^N \prod_{j=1}^M (1 + x_i z_j). \quad (2.6)$$

Since the identity (2.6) uniquely determines the functions $P_{\lambda'}(\dots; t, 0)$ (due to the orthogonality of the Hall–Littlewood polynomials), we conclude that

$$F_{\lambda}(z_1, \dots, z_M) = P_{\lambda'}(z_1, \dots, z_M; t, 0),$$

which are exactly the q -Whittaker polynomials (where the role of q is taken by the parameter t).

2.4 Fusion

The same steps as in Section 2.3 may be carried out directly at the level of the vertex weights. Let us briefly outline this procedure, which is known as *fusion* [KRS81], [KR87], [Bor17], [CP16], [BP18].

The fundamental building block is the Hall–Littlewood vertex (2.1), of horizontal capacity one, so $j_1, j_2 \in \{0, 1\}$. Fix $J \in \mathbb{Z}_{\geq 1}$ and consider J copies of the vertex, stacked vertically. Assign to the r -th row the spectral parameter $t^{r-1}x$, so that the column carries the geometric string of parameters. This mirrors the substitution $y_i \mapsto y_i^{(J)}$ from Section 2.3 above.

The fused weight is defined as the sum over the J incoming and outgoing horizontal occupations $\vec{h}_1, \vec{h}_2 \in \{0, 1\}^J$:

$$w_x^{(J)}(i_1, j_1; i_2, j_2) = \frac{1}{Z_{j_1}(J)} \sum_{\substack{|\vec{h}_1|=j_1 \\ |\vec{h}_2|=j_2}} t^{\sum_{r=1}^J (r-1) h_1^{(r)}} \prod_{r=1}^J w_{t^{r-1}x}(\ell_{r-1}, h_1^{(r)}; \ell_r, h_2^{(r)}), \quad (2.7)$$




Figure 14: The fusion procedure (2.7) for $J = 3$. Summing over the incoming and outgoing strings $\vec{h}_1, \vec{h}_2 \in \{0, 1\}^J$ with totals $j_1 = |\vec{h}_1|$ and $j_2 = |\vec{h}_2|$, with the t -exchangeable weight on the incoming arrows, collapses the column to a single fused vertex of horizontal capacity J .

where w_u are the weights (2.1) with spectral parameter u , and the internal vertical occupations are determined by arrow conservation, $\ell_0 = i_1$ and $\ell_r = \ell_{r-1} + h_1^{(r)} - h_2^{(r)}$ (so that $\ell_J = i_2$). See Figure 14 for an illustration of the fusion procedure.

The normalization in (2.7) is the t -deformed binomial

$$Z_j(J) = \sum_{\vec{h}_1: |\vec{h}_1|=j} t^{\sum_{r=1}^J (r-1) h_1^{(r)}} = t^{\binom{j}{2}} \begin{bmatrix} J \\ j \end{bmatrix}_t.$$

The weight $t^{\sum_{r=1}^J (r-1) h_1^{(r)}} / Z_{j_1}(J)$ placed on the incoming arrows is the t -exchangeable normalization (e.g., see [GO10] for a probabilistic perspective). Fusion is built so as to preserve the Yang–Baxter equation [KRS81], [KR87]; the fused weights (2.7) therefore satisfy it as well, as do their stochastic counterparts [Bor17], [BP18]. By construction (2.7) is a genuine vertex weight of horizontal capacity J — a function of the four occupations i_1, j_1, i_2, j_2 alone.

In fact, the outer sum over the outgoing string in (2.7) is redundant: a fused column maps t -exchangeable inputs to t -exchangeable outputs [CP16, Prop. 3.6], so that for *every* fixed $\vec{h}_2$ with $|\vec{h}_2| = j_2$, we have the following refined identity:

$$\sum_{\vec{h}_1: |\vec{h}_1|=j_1} t^{\sum_{r=1}^J (r-1) h_1^{(r)}} \prod_{r=1}^J w_{t^{r-1}x}(\ell_{r-1}, h_1^{(r)}; \ell_r, h_2^{(r)}) = \frac{Z_{j_1}(J)}{Z_{j_2}(J)} t^{\sum_{r=1}^J (r-1) h_2^{(r)}} w_x^{(J)}(i_1, j_1; i_2, j_2). \quad (2.8)$$

Remark 2.6. Let us illustrate (2.8) for $J = 2$. The general- J case of (2.8) follows from the same computation applied to adjacent rows, since the symmetrization underlying t -exchangeability is generated by swaps of neighboring rows.

Take $i_1 = 0, j_1 = 1$, so the incoming string is $(1, 0)$ with weight 1 and $(0, 1)$ with weight t . For the outgoing string $\vec{h}_2 = (0, 1)$ the two incoming terms give

$$\underbrace{t w_x(0, 0; 0, 0) w_{tx}(0, 1; 0, 1)}_{= t^2 x} + \underbrace{w_x(0, 1; 1, 0) w_{tx}(1, 0; 0, 1)}_{= tx(1-t)} = tx = x t^1,$$

while $\vec{h}_2 = (1, 0)$ gives $x = x t^0$. Since $j_1 = j_2 = 1$ makes $Z_{j_1}/Z_{j_2} = 1$ in (2.8), either outgoing string already yields $w_x^{(2)}(0, 1; 0, 1) = x$, with no outer sum needed.

Carrying out the fusion sum (2.7) — a direct if lengthy computation — gives the fused weight in closed form:²

$$w_x^{(J)}(i_1, j_1; i_2, j_2) = \mathbf{1}_{i_1+j_1=i_2+j_2} t^{\frac{1}{2}i_1(i_1+2j_1-1)+\frac{1}{2}(j_1-i_2)(j_1-i_2-1)} x^{j_2} \frac{(t; t)_{j_1}}{(t; t)_{j_2} (t; t)_{i_2}} \\ \times \sum_{k=0}^{i_1} \frac{(t^{-i_1}; t)_k (t^{-i_2}; t)_k}{(t; t)_k} (t^{1+j_1-i_2+k}; t)_{i_1-k} (t^{1+J-i_1-j_1+k}; t)_{i_1-k} t^k. \quad (2.9)$$

Since the prefactor multiplying x^{j_2} is free of both x and J , we may write $w_x^{(J)}(i_1, j_1; i_2, j_2) = x^{j_2} \sum_{p=0}^{i_1} c_p(t) (t^J)^p$. In particular t^J may be set to any complex value.

Expanding the one t^J -dependent factor by the t -binomial theorem and evaluating the inner sum by the terminating t -Chu–Vandermonde identity [GR04, (II.6)] gives $c_p(t)$ explicitly:

$$c_p(t) = (-1)^p t^{(j_2-p)} \frac{(t; t)_{j_1}}{(t; t)_{j_2} (t; t)_{i_2}} \begin{bmatrix} i_1 \\ p \end{bmatrix}_t (t^{1+j_2-p}; t)_p (t^{1+j_1}; t)_{i_1-p}, \quad 0 \leq p \leq i_1.$$

The factor $(t^{1+j_2-p}; t)_p$ acquires the vanishing term $1 - t^0$ as soon as $p > j_2$, so we see that the degree of (2.9) in t^J is equal to $\min(i_1, j_2)$.

Put $z = -t^J x$. Then

$$w_x^{(J)}(i_1, j_1; i_2, j_2) = x^{j_2} \sum_{p=0}^{\min(i_1, j_2)} c_p(t) (-z/x)^p = \sum_{p=0}^{\min(i_1, j_2)} c_p(t) (-1)^p x^{j_2-p} z^p. \quad (2.10)$$

Send $x \rightarrow 0$ while keeping z fixed. If $i_1 < j_2$, then (2.10) vanishes, which leads to the indicator $\mathbf{1}_{i_1 \geq j_2}$ in (2.2). If $i_1 \geq j_2$, then the limit picks only the term with $p = j_2$, giving the q -Whittaker weight:

$$\mathbf{1}_{i_1 \geq j_2} z^{j_2} \begin{bmatrix} i_1 \\ j_2 \end{bmatrix}_t = \mathbf{1}_{i_1 \geq j_2} z^{j_2} \frac{(t; t)_{i_1}}{(t; t)_{j_2} (t; t)_{i_1-j_2}},$$

where the parameter t plays the role of q in (2.2).

Remark 2.7 (Spin parameter). Carrying the spin parameter s of Remark 2.2 through the fusion works equally well, and produces the *spin* q -Whittaker weights of Remark 2.5. The fused J -dependent weights one has to go through take the following q -hypergeometric form [Man14], [Bor17], [CP16]:

$$w_{x,s}^{(J)}(i_1, j_1; i_2, j_2) = \mathbf{1}_{i_1+j_1=i_2+j_2} \frac{(-1)^{i_1+j_2} t^{\frac{1}{2}i_1(i_1+2j_1-1)} s^{j_2-i_1} x^{i_1} (t; t)_{j_1} (s^2; t)_{i_2} (xs^{-1}; t)_{j_1-i_2}}{(t; t)_{j_2} (t; t)_{i_2} (s^2; t)_{i_1} (xs; t)_{i_1+j_1}} \\ \times {}_4\bar{\varphi}_3 \left(\begin{matrix} t^{-i_1}, t^{-i_2}, t^J s x, t s x^{-1} \\ s^2, t^{1+j_1-i_2}, t^{1+J-i_1-j_1} \end{matrix} \middle| t, t \right).$$

Here ${}_{r+1}\bar{\varphi}_r$ is the regularized terminating basic hypergeometric series:

$${}_{r+1}\bar{\varphi}_r \left(\begin{matrix} t^{-n}, a_1, \dots, a_r \\ b_1, \dots, b_r \end{matrix} \middle| t, z \right) = \sum_{k=0}^n z^k \frac{(t^{-n}; t)_k}{(t; t)_k} \prod_{i=1}^r (a_i; t)_k (b_i t^k; t)_{n-k} = \prod_{i=1}^r (b_i; t)_n \cdot {}_{r+1}\varphi_r. \quad (2.11)$$

The fused weight (2.10) reduces to the spin Hall–Littlewood weight of Remark 2.2 at $J = 1$, and to (2.9) at $s = 0$. Here s is the parameter of the vertical (generic Verma) representation of $U_t(\widehat{\mathfrak{sl}}_2)$. Setting $s^2 = t^{-I}$ caps the vertical occupations at I .

²This computation may be reduced to checking certain recurrences for $w_x^{(J)}$. We refer to [CP16, Prop. 3.6] for a detailed derivation via these recurrences in the presence of the additional spin parameter.

3 Bijectivization: From identities to Markov transitions

We outline the idea of *bijectivization* (probabilistic bijection), the device that converts the algebraic Yang–Baxter identities of the previous sections into stochastic (Markov) steps.

Bijectivization of the Yang–Baxter equation was introduced by Bufetov and Petrov [BP19], extending the earlier works which constructed randomized (Markov) RSK-like insertion processes for Hall–Littlewood and q -Whittaker polynomials [OP13], [BP16], [MP17], [BP15], [BBW16]. In integrable probability, Yang–Baxter bijectivization has been applied further to connect spin Hall–Littlewood and spin q -Whittaker polynomials [BMP21], [MP22] to various stochastic vertex models; to construct Markov maps preserving the TASEP measure [PS22], [PS24] or the six-vertex measure on the full plane [NP22]. Furthermore, bijectivization has been extended to half-space models [CD21] and to the full two-parameter Macdonald setting [AF21], [FSA24].

Strikingly, the same mechanism appears independently in two-dimensional critical statistical mechanics in the proof of rotational invariance of the Ising, random-cluster, and six-vertex models. Namely, the star-triangle (Yang–Baxter) relation is realized as a move that “uses extra randomness ... and is not a deterministic matching of the configurations” [DCKK⁺20]. This is the general mechanism which we are about to describe.

Consider any identity with positive terms,

$$\sum_{a \in A} \text{wt}(a) = \sum_{b \in B} \text{wt}(b), \quad (3.1)$$

where A, B are finite or countable index sets and $\text{wt} > 0$. A *bijectivization* (or probabilistic bijection) of (3.1) is a pair of nonnegative matrices $\mathbf{p}^{\text{fwd}}(a \rightarrow b)$ and $\mathbf{p}^{\text{bwd}}(b \rightarrow a)$, indexed by $a \in A$ and $b \in B$, satisfying the stochasticity

$$\sum_{b \in B} \mathbf{p}^{\text{fwd}}(a \rightarrow b) = 1, \quad \sum_{a \in A} \mathbf{p}^{\text{bwd}}(b \rightarrow a) = 1,$$

and the reversibility (detailed balance) relation

$$\mathbf{p}^{\text{fwd}}(a \rightarrow b) \text{wt}(a) = \mathbf{p}^{\text{bwd}}(b \rightarrow a) \text{wt}(b) \quad \text{for all } a \in A, b \in B. \quad (3.2)$$

Summing (3.2) over b recovers the left side of (3.1) from $\mathbf{p}^{\text{fwd}}$ being stochastic, and summing over a recovers the right side; thus a bijectivization is exactly a *coupling* [Lin02], [Tho00] of the two sides of the identity. That is, the joint distribution on $A \times B$ with weights $\mathbf{p}^{\text{fwd}}(a \rightarrow b) \text{wt}(a)$ has marginals $\text{wt}(a)$ and $\text{wt}(b)$ on A and B , respectively.

Let us highlight two extreme cases of bijectivization:

- If $\mathbf{p}^{\text{fwd}}$ is equal to zero or one, i.e., supported on the graph of a map $a \mapsto b$, then the coupling is an honest *bijection* (and $\mathbf{p}^{\text{bwd}}$ is its inverse). One can say that this bijection provides a combinatorial proof of (3.1).
- The opposite extreme is the *independent* (“maximally random”) bijectivization

$$\mathbf{p}^{\text{fwd}}(a \rightarrow b) = \frac{\text{wt}(b)}{\sum_{b' \in B} \text{wt}(b')},$$

which makes the output independent of the input. The coupling in this case is just the product of the two marginals.

A bijectivization always exists, as the second example shows. It is unique only when one of A, B is a singleton. When both A and B have exactly two elements, the bijectivization is a one-parameter family.

When bijectivization is applied to the Yang–Baxter equation (1.9), the sets A and B are the sets of all possible arrow configurations of the internal edges (indexed by k_1, k_2, k_3 and k'_1, k'_2, k'_3 , respectively), and the weights are the products of the three vertex weights on each side. The full power of bijectivization is realized when it is applied sequentially on the lattice, by dragging the cross vertex through several columns, one at a time. In particular, if the cross is dragged from left to right, then the Markov property implies that the distribution of k'_1, k'_2, k'_3 in the first (leftmost) column does not depend on the subsequent columns.


4 Deriving RSK from the Yang–Baxter equation

We now specialize bijectivization to the Yang–Baxter equation for the Schur vertex model that we used in Section 1.5. We will see that a natural bijectivization in this case may be taken to be deterministic, and identify it with the classical Robinson–Schensted–Knuth (RSK) row insertion algorithm. This material first appeared in [MP22, §5]. First, we recall the RSK in the toggle form.

4.1 RSK via toggles

Classical RSK can be built entirely from local moves [Pak01], [Hop14], which we now explain. First, recall Fomin’s *growth-diagram* form of RSK [Fom88], [Fom95a]. Instead of building tableaux by insertion, one records the algorithm as a labeling of the vertices of an $N \times N$ grid by partitions: the partitions at adjacent vertices interlace (that is, differ by a horizontal strip), and the nonnegative integer entry of the input matrix is placed inside the corresponding unit cell. The two output tableaux P and Q are then read off the chains of partitions along the top and the right boundaries of the grid. See Figure 15 for an illustration of the growth diagram. A complete worked example — the insertion algorithm running in parallel with the corresponding vertex model states — is presented in Figure 20 at the end of this section.

The entire diagram is generated by a single *local (toggle) rule* that updates the top-right corner of a cell (denoted by ν) from the other three corners (denoted by $\lambda, \mu, \varkappa$) and the integer inside the cell (denoted by k):



The local rule assumes that μ, λ are fixed, and it is a bijection between the following two sets:

$$\{\varkappa : \varkappa \prec \lambda, \varkappa \prec \mu\} \times \mathbb{Z}_{\geq 0} \longleftrightarrow \{\nu : \nu \succ \lambda, \nu \succ \mu\}, \quad (4.1)$$

which is defined by the “toggle” relations

$$\nu_1 = k + \max(\lambda_1, \mu_1), \quad \nu_{c+1} = \max(\lambda_{c+1}, \mu_{c+1}) + \min(\lambda_c, \mu_c) - \varkappa_c \quad (c \geq 1). \quad (4.2)$$

In Figure 19 at the end of this section, the roles of $\varkappa, \lambda, \mu, \nu$ are marked in one cell of the growth diagram of the running example of Figure 20.

Remark 4.1. Note that other bijections between the two sets in (4.1) would correspond to other versions of the RSK correspondence, for example, the column insertion version. We refer to [Fom95a], [Fom95b] for further discussion.




Figure 15: Fomin's growth diagram. Each vertex carries a partition, drawn as its Young diagram; the left and bottom boundaries are empty ($\emptyset$); adjacent partitions interlace by a horizontal strip. Each cell carries a nonnegative integer (the input matrix entry), and the top-right partition is determined by the other three corners and this integer via the local (toggle) rule (4.2). The top and right boundaries give the output Young tableaux P and Q , respectively.

4.2 Yang–Baxter bijectivization in the Schur case

The Schur Yang–Baxter equation (1.9) for the weights (1.6), (1.7) and the cross-vertex weights (1.8) possesses a *deterministic* bijectivization. First, note that, as in any Yang–Baxter equation of the form (1.9), for fixed boundary conditions occupations $i_1, i_2, i_3, j_1, j_2, j_3 \in \mathbb{Z}_{\geq 0}$, both sides of (1.9) vanish unless the global balance $i_1 + j_2 + j_3 = i_2 + i_3 + j_1$ holds. Under this balance condition, path conservation at the three vertices leaves a single free internal coordinate on each side,

$$(k_1, k_2, k_3) = (k_1, i_2 + k_1 - i_1, i_3 + j_1 - k_1), \quad (k'_1, k'_2, k'_3) = (k'_1, k'_1 + j_2 - j_1, j_3 + i_1 - k'_1), \quad (4.3)$$

respectively.

Lemma 4.2. *In the Yang–Baxter equation (1.9) for the Schur vertex model, the summands in the left- and right-hand sides are, respectively,*

$$x^{j_2} y^{j_1} (xy)^{\min(k_1, k_2)} \quad \text{and} \quad x^{k'_2} y^{k'_1} (xy)^{\min(j_1, j_2)}. \quad (4.4)$$

These are nonzero exactly on the integer intervals

$$\max(0, i_1 - i_2) \leq k_1 \leq \min(i_3, i_3 + j_1 - j_2), \quad \max(0, j_1 - j_2) \leq k'_1 \leq \min(j_3, j_3 + i_1 - i_2), \quad (4.5)$$

which contain the same number of points.

Moreover, the map

$$k'_1 = j_1 - \min(j_1, j_2) + \min(k_1, k_2) = k_1 + \max(0, j_1 - j_2) - \max(0, i_1 - i_2) \quad (4.6)$$

is a weight-preserving bijection between the weights (4.4), so (1.9) admits a deterministic bijectivization — one with $\mathbf{p}^{\text{fwd}}, \mathbf{p}^{\text{bwd}} \in \{0, 1\}$ in (3.2).

Proof. The formulas (4.4) are obtained by substituting (4.3) into the vertex weights (1.6), (1.7), and (1.8). The occupation bounds $k_2 \geq 0$ and $k_3 \geq \max(j_1, j_2)$ in (1.6), (1.7) (and their primed counterparts) cut out the two desired intervals (4.5). Finally, note that since $k'_1 + \min(j_1, j_2) = j_1 + \min(k_1, k_2)$ and $k'_2 = k'_1 + j_2 - j_1$, the map (4.6) turns the right-hand summand $x^{k'_2} y^{k'_1} (xy)^{\min(j_1, j_2)}$ into the left-hand one, $x^{j_2} y^{j_1} (xy)^{\min(k_1, k_2)}$. Moreover, it is a bijection between the two intervals (4.5). This completes the proof. $\square$

Remark 4.3. The bijection (4.6) takes a cleaner form in terms of the vertical arrow counts $k_3 = i_3 + j_1 - k_1$, $k'_3 = j_3 + i_1 - k'_1$:

$$k'_3 = k_3 + \max(i_1, i_2) - \max(j_1, j_2). \quad (4.7)$$




Figure 16: Dragging the Yang–Baxter cross through two rows μ/λ (top) and λ/ν (bottom) of a growth diagram. Each vertical edge carries a column multiplicity $\square_c - \square_{c+1}$. The middle row is rewritten $\lambda \mapsto \nu$ (the shaded columns are already updated). The only free input is $k \in \mathbb{Z}_{\geq 0}$ at the first column, which corresponds to the state of the cross vertex $\mathbb{R}_{x,y}$ to the left of the zeroth column (the simplest example is in Figure 12, (c)). We discuss this free input in detail at the end of the proof of Proposition 4.4.

Proposition 4.4. *Dragging the cross (1.8) rightward across the two rows λ/ν and μ/λ of Figure 16, and rewriting the middle rail $\lambda \mapsto \nu$ one column at a time by the deterministic move of Lemma 4.2, reproduces the local rule (4.2). Consequently, the Yang–Baxter moves over the $N \times N$ grid (Figure 15) produce the Robinson–Schensted–Knuth correspondence.³*

Proof. Fix a column $c \geq 1$. The three rails give the vertical occupations

$$i_3 = \lambda_c - \lambda_{c+1}, \quad j_3 = \mu_c - \mu_{c+1}, \quad k_3 = \lambda_c - \lambda_{c+1}, \quad k'_3 = \nu_c - \nu_{c+1}, \quad (4.8)$$

the middle edge reading λ ahead of the cross and ν behind it. Each horizontal leg carries the number of paths crossing it, a partial sum of the strip multiplicities: the two legs entering the cross carry $i_1 = \nu_c - \mu_c$ and $i_2 = \nu_c - \lambda_c$, and the two leaving it carry $j_1 = \lambda_{c+1} - \lambda_{c+1}$ and $j_2 = \mu_{c+1} - \lambda_{c+1}$, all nonnegative by $\lambda \prec \mu, \mu \prec \nu$. Hence

$$\max(i_1, i_2) = \nu_c - \min(\lambda_c, \mu_c), \quad \max(j_1, j_2) = \max(\lambda_{c+1}, \mu_{c+1}) - \lambda_{c+1}. \quad (4.9)$$

³This RSK correspondence is the most classical one, based on row insertions. It would be interesting to see if the other three bijections (column insertion, dual row insertion, dual column insertion) can be derived from the Yang–Baxter equation in a similar way.

Substituting (4.8)–(4.9) into the k_3 -identity (4.7) and cancelling ν_c and $\varkappa_{c+1}$ gives

$$\nu_{c+1} = \max(\lambda_{c+1}, \mu_{c+1}) + \min(\lambda_c, \mu_c) - \varkappa_c,$$

which is precisely the toggle (4.2) for $c \geq 1$.

The case $c = 0$ is slightly different, because here each Yang–Baxter cross carries the free input (an element of the integer input matrix for the RSK). This input matrix is encoded by the $N \times N$ grid as in Figure 15. Assume that we need to insert the cross $\mathbb{R}_{x_a, y_b}$, $1 \leq a, b \leq N$, through the zeroth column of the grid which carries infinitely many vertical paths. Let this cross carry the parameter $k_{ab} \in \mathbb{Z}_{\geq 0}$. Assume that the input arrows to $\mathbb{R}_{x_a, y_b}$ are i_1 and j_1 , see Figure 17, and $k_{ab} = \min(k_1, k_2)$. This means that

$$k_1 = k_{ab} + \max(0, i_1 - i_2), \quad k_2 = k_{ab} + \max(0, i_2 - i_1).$$




Figure 17: Dragging the cross through the zeroth column.

After dragging the cross, we have $k'_1 = \nu_1 - \mu_1$ and $k'_2 = \nu_1 - \lambda_1$. Comparing the weights on the left and on the right in Figure 17, we see that

$$\nu_1 = k_{ab} + \max(\lambda_1, \mu_1),$$

which does not depend on the incoming occupations i_1, i_2 .

Thus, the data $k_{ab} = \min(k_1, k_2)$ at each cross $\mathbb{R}_{x_a, y_b}$ in Figure 13 is exactly the RSK input matrix. This completes the proof of Proposition 4.4. $\square$

Figure 18 fixes the input matrix M of the running example and its biword. Figure 20 illustrates the resulting correspondence step by step, with the bumping cascades on the RSK side matched to the path hops on the seam of the vertex model. Figure 19 displays the growth diagram of the same input matrix.

$$M = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 2 & 0 \\ 1 & 1 & 0 \end{pmatrix} \longrightarrow \left(\begin{array}{cc|cc|cc} 1 & 1 & 2 & 2 & 3 & 3 \\ 0 & 2 & 2 & 2 & 1 & 2 \\ 1 & 1 & 0 & 0 & 0 & 0 \end{array} \right) \begin{array}{l} b \text{ --- recorded in } Q \\ a \text{ --- inserted into } P \end{array}$$

entry m_{ba} : row b , column a biword: m_{ba} copies of $\binom{b}{a}$, in lexicographic order

Figure 18: The input matrix M of the running example and its biword. The top row of the biword (the b 's, red) is recorded in Q , and the bottom row (the a 's, blue) is inserted into P .




Figure 19: The growth diagram for the 3×3 input matrix of Figure 18; the top and right boundary chains are the P - and Q -chains. In the shaded cell, the corners play the roles $\kappa = (1)$, $\lambda = (1)$, $\mu = (1)$ in the local rule (4.2) with $k = 2$, which yields $\nu = (3)$.

5 Deformed weights and particle systems

This section is optional and may be skipped on a first reading.

5.1 The q -Whittaker case and q -PushTASEP

Replace the Schur weights of Section 4 by their q -Whittaker deformation: the up-right weight W (2.2) and the corresponding down-right dual W^* .⁴ The Yang–Baxter equation (1.9) persists, with a q -deformed cross $\mathbb{R}_{x,y}^{(q)}$ in place of (1.8). The cross is the unique (up to scale) solution of (1.9); it obeys the same arrow conservation $i_1 + j_2 = j_1 + i_2$ as (1.8), depends on x, y only through the product xy , and degenerates to (1.8) as $q \rightarrow 0$. Its entries include

$$\begin{aligned} \mathbb{R}_{x,y}^{(q)}(0, 0; m, m) &= \frac{(xy)^m}{(q; q)_m}, & \mathbb{R}_{x,y}^{(q)}(m, m; 0, 0) &= (q; q)_m, \\ \mathbb{R}_{x,y}^{(q)}(1, 1; 1, 1) &= xy + q, & \mathbb{R}_{x,y}^{(q)}(1, 2; 1, 2) &= \mathbb{R}_{x,y}^{(q)}(2, 1; 2, 1) = xy(1 + q) + q^2, \\ \mathbb{R}_{x,y}^{(q)}(2, 2; 1, 1) &= (1 - q^2)(xy + q + q^2), & \mathbb{R}_{x,y}^{(q)}(2, 2; 2, 2) &= (xy)^2 + xyq(1 + q)^2 + q^4, \end{aligned} \quad (5.1)$$

and in general it is the regularized terminating ${}_2\bar{\varphi}_1$ series

$$\mathbb{R}_{x,y}^{(q)}(i_1, j_1; i_2, j_2) = \mathbf{1}_{i_1+j_2=j_1+i_2} \frac{(q; q)_{j_1}}{(q; q)_{j_2} (q; q)_{i_2}} (xy)^{i_2} {}_2\bar{\varphi}_1 \left(\begin{matrix} q^{-i_2}; q^{-i_1} \\ q^{1+j_2-i_2} \end{matrix} \middle| q, \frac{q^{i_1+j_2+1}}{xy} \right), \quad (5.2)$$

where ${}_2\bar{\varphi}_1$ is the regularized terminating series (2.11). Written out, the q -deformed cross is the finite sum

$$\begin{aligned} \mathbb{R}_{x,y}^{(q)}(i_1, j_1; i_2, j_2) &= \mathbf{1}_{i_1+j_2=j_1+i_2} \frac{(q; q)_{j_1}}{(q; q)_{j_2} (q; q)_{i_2}} \\ &\quad \times \sum_{k=0}^{i_2} \frac{(q^{-i_2}; q)_k (q^{-i_1}; q)_k}{(q; q)_k} (q^{1+j_2-i_2+k}; q)_{i_2-k} q^{(i_1+j_2+1)k} (xy)^{i_2-k}. \end{aligned} \quad (5.3)$$

⁴The dual is the q -analogue of (1.7), $W^*(i_1, j_1; i_2, j_2) = \frac{(q; q)_{i_1}}{(q; q)_{i_2}} \mathbf{1}_{i_2+j_1=i_1+j_2} \mathbf{1}_{i_2 \geq j_2} y^{j_2} \begin{bmatrix} i_2 \\ j_2 \end{bmatrix}_q$. The scalar prefactor $(q; q)_{i_1}/(q; q)_{i_2}$ is a gauge on the vertical edges.

Remark 5.1. This q -deformation of the benign expression $(xy)^{\min(i_2, j_2)}$ in (1.8) may appear contrived, but it is required to satisfy the Yang–Baxter equation. There also exists a further spin deformation, see [BMP21], [MP22].

Remark 5.2. The sum over the outputs of the cross is independent of (i_1, j_1) :

$$\sum_{i_2, j_2} \mathbb{R}_{x, y}^{(q)}(i_1, j_1; i_2, j_2) = \frac{1}{(xy; q)_\infty} \quad \text{for every } (i_1, j_1). \quad (5.4)$$

For $(i_1, j_1) = (0, 0)$ this is the (infinite) q -binomial theorem

$$\sum_{m \geq 0} \frac{(xy)^m}{(q; q)_m} = \frac{1}{(xy; q)_\infty}.$$

This identity is the q -Whittaker Cauchy kernel. For $q = 0$, it reduces to the geometric series $\sum_{a \geq 0} (xy)^a = 1/(1 - xy)$ which appeared throughout the Schur case.

Let us apply the bijectivization in the same setting as the one which produced the RSK correspondence in Section 4.2.

In the bulk the q -deformation turns the deterministic move of Lemma 4.2 into a random coupling: k_1, k'_1 still range over the equal-length intervals (4.5), but the summands (4.4) now carry q -binomial factors, so only the two *sums* (and not the individual summands) match, and so $\mathbf{p}^{\text{fwd}}(k_1 \rightarrow k'_1)$ must spread the probabilistic mass.

However, at the leftmost column the bijectivization is unique. There the cross pushes its arrows into the system through the reservoir of infinitely many vertical paths, see Remark 1.11 and Figure 13. The vertex weights with infinitely many vertical arrows do not depend on the incoming horizontal occupation: $W(\infty, j_1; \infty, j_2) = x^{j_2}/(q; q)_{j_2}$, the q -deformation of Remark 1.11. Hence, in the Yang–Baxter equation with i_1, i_2 fixed, we may sum over the internal occupations k_1, k_2 in the left-hand side of (1.9), and by (5.4) that side collapses to a *single* term:

$$\frac{1}{(xy; q)_\infty} \frac{x^{j_2}}{(q; q)_{j_2}} \frac{y^{j_1}}{(q; q)_{j_1}} = \sum_{k'_1, k'_2} \frac{x^{k'_2}}{(q; q)_{k'_2}} \frac{y^{k'_1}}{(q; q)_{k'_1}} \mathbb{R}_{x, y}^{(q)}(k'_2, k'_1; j_2, j_1). \quad (5.5)$$

Note that both sides are independent of the fixed i_1, i_2 . A bijectivization of an identity with a singleton on one side is unique (Section 3). Under it, the pair (k'_1, k'_2) is drawn with probability proportional to its weight in the right-hand side of (5.5).

Let us use the symmetry

$$\frac{x^{k'_2} y^{k'_1}}{(q; q)_{k'_2} (q; q)_{k'_1}} \mathbb{R}_{x, y}^{(q)}(k'_2, k'_1; j_2, j_1) = \frac{x^{j_2} y^{j_1}}{(q; q)_{j_2} (q; q)_{j_1}} \mathbb{R}_{x, y}^{(q)}(j_2, j_1; k'_2, k'_1)$$

of the q -deformed cross, and denote these probabilities to pick (k'_1, k'_2) by

$$\mathbb{L}_{x, y}(i_1, j_1; i_2, j_2) := \frac{\mathbb{R}_{x, y}^{(q)}(i_1, j_1; i_2, j_2)}{\sum_{i'_2, j'_2} \mathbb{R}_{x, y}^{(q)}(i_1, j_1; i'_2, j'_2)} = (xy; q)_\infty \mathbb{R}_{x, y}^{(q)}(i_1, j_1; i_2, j_2). \quad (5.6)$$

We have $\mathbb{L}_{x, y} \geq 0$ and $\sum_{i_2, j_2} \mathbb{L}_{x, y}(i_1, j_1; i_2, j_2) = 1$.

Let us explain a connection to particle systems. Applied across the lattice, the moves $\mathbb{L}_{x, y}$ grow a field $\{\lambda^{(x, y)}\}$ of Young diagrams on the quadrant $\mathbb{Z}_{\geq 0}^2$ — the q -Whittaker process — whose *first*

parts $\lambda_1^{(x,y)}$ encode the particles of the (discrete-time) q -PushTASEP of Matveev–Petrov [MP17, §6.3] (the $s = 0$ case of the q -Hahn PushTASEP of Corwin–Matveev–Petrov [CMP19]). The four diagrams at the corners of a unit cell,

$$\varkappa = \lambda^{(x,y)}, \quad \mu = \lambda^{(x+1,y)}, \quad \lambda = \lambda^{(x,y+1)}, \quad \nu = \lambda^{(x+1,y+1)},$$

obey the local rule which randomly picks ν based on $\mu, \varkappa, \lambda$ and an integer k (the analogue of the RSK input k_{ab} from the proof of Proposition 4.4). Thanks to the unique bijectivization at the zeroth column, the evolution of the first parts of the partitions is *marginally Markovian*, governed by the stochasticized cross $\mathbb{L}_{x,y}(i_1, j_1; i_2, j_2)$ with

$$(i_1, j_1) = (\lambda_1 - \varkappa_1, \mu_1 - \varkappa_1), \quad (i_2, j_2) = (\nu_1 - \mu_1, \nu_1 - \lambda_1), \quad i_1 + j_2 = j_1 + i_2 = \nu_1 - \varkappa_1. \quad (5.7)$$

The first coordinate of the field plays the role of time, $t = x$. The update (5.7) involves only the first parts, so the array $(\lambda_1^{(t,j)})_{j \geq 1}$ evolves as a Markov chain in t , which we describe as a particle system. Set

$$P_j(t) := \lambda_1^{(t,j)} + j, \quad j \geq 1, \quad P_0 \equiv 0. \quad (5.8)$$

By the interlacing, $P_0 < P_1(t) < P_2(t) < \dots$, and the particles move to the right, pushing one another. One step $t \rightarrow t+1$ updates them in order $j = 1, 2, \dots$: given the move ℓ of P_{j-1} and the gap $g = P_j - P_{j-1} - 1$, the particle P_j moves right by L with probability $\mathbb{L}_{x,y}(g, \ell; g + L - \ell, L)$. Indeed, the four counts inside $\mathbb{L}_{x,y}$ in the latter expression can be matched to (5.7) for the cell with the lower left corner $\varkappa = \lambda^{(t,j-1)}$, with the outgoing pair $(g + L - \ell, L)$ recording the new gap and the jump. In particular, the leading particle P_1 receives “push” $\ell = 0$ from the frozen P_0 and jumps q -geometrically,

$$\mathbb{L}_{x,y}(g, 0; g + \ell, \ell) = (xy; q)_\infty \frac{(xy)^\ell}{(q; q)_\ell}, \quad (5.9)$$

independently of the gap g . The dynamics preserves the order: $\mathbb{L}_{x,y} = 0$ unless $L \geq \ell - g$, so $\ell > g$ forces a push.

A continuous-time limit of the q -PushTASEP (as $xy \rightarrow 0$ and the time step is rescaled proportionally to $(xy)^{-1}$) produces a simpler dynamics introduced in [CP15]. Under the continuous-time q -PushTASEP, each particle carries an independent exponential clock; when it rings the particle steps forward by one and, with probability q^g (with g the gap to the next particle ahead), instantaneously pushes that neighbor forward by one as well, the push cascading down the line. At $q = 0$ the push probability is 1 across a zero gap and 0 otherwise, so a jump shoves the whole block of adjacent particles ahead by one. That is, for $q = 0$ we get back the classical PushTASEP.

5.2 Hall–Littlewood: a row-indexed RSK and the six-vertex model

The Hall–Littlewood weights (2.1) — the t -deformation of the higher-spin model of Remark 1.4 — carry a randomized bijectivization parallel to the q -Whittaker one of Section 5.1. The *horizontal* edges in the Hall–Littlewood case have capacity one ($j_1, j_2 \in \{0, 1\}$), while the vertical edges stay unbounded ($i_1, i_2 \in \mathbb{Z}_{\geq 0}$). Read along the *rows* of the growth diagram rather than its columns, the insertion is indexed by rows, and, as we now explain, its scalar marginal is the *stochastic six-vertex model* rather than the q -PushTASEP.

Recall the vertex weights entering the Hall–Littlewood Cauchy identity (2.3). The up-right weight is (2.1), in formulas

$$W(i_1, j_1; i_2, j_2) = \mathbf{1}_{i_1+j_1=i_2+j_2} \times \begin{cases} 1, & (j_1, j_2) = (0, 0), \\ x, & (j_1, j_2) = (1, 1), \\ 1, & (j_1, j_2) = (1, 0), \\ x(1-t^{i_1}), & (j_1, j_2) = (0, 1), \end{cases} \quad (5.10)$$

whose N -row partition function is $P_\lambda(x_1, \dots, x_N; 0, t)$ (Proposition 2.1). The y -side of (2.3) carries the functions Q_λ , whose branching rule differs from that of P_λ : the factor $1-t^m$ is attached to a path *entering* a column of multiplicity m rather than leaving it [Mac95, Ch. III, §5]. Accordingly, the down-right weight (the analogue of the reflection (1.7), with the conservation law $i_2 + j_1 = i_1 + j_2$) is

$$W^*(i_1, j_1; i_2, j_2) = \mathbf{1}_{i_2+j_1=i_1+j_2} \times \begin{cases} 1, & (j_1, j_2) = (0, 0), \\ y, & (j_1, j_2) = (1, 1), \\ 1-t^{i_1}, & (j_1, j_2) = (1, 0), \\ y, & (j_1, j_2) = (0, 1). \end{cases} \quad (5.11)$$

Gluing N rows of each model along a common boundary λ , as in Figure 11, produces the vertex model for the left-hand side of the Cauchy identity (2.3).

The Yang–Baxter equation (1.9) with the weights (5.10)–(5.11) has the unique (up to scale) solution

$$\begin{aligned} \mathbb{R}_{x,y}^{(t)}(0, 0; 0, 0) &= 1, & \mathbb{R}_{x,y}^{(t)}(1, 0; 1, 0) &= \mathbb{R}_{x,y}^{(t)}(0, 1; 0, 1) = \frac{1-xy}{1-xy}, & \mathbb{R}_{x,y}^{(t)}(1, 1; 1, 1) &= t, \\ \mathbb{R}_{x,y}^{(t)}(0, 0; 1, 1) &= \frac{xy(1-t)}{1-xy}, & \mathbb{R}_{x,y}^{(t)}(1, 1; 0, 0) &= \frac{1-t}{1-xy}. \end{aligned} \quad (5.12)$$

Like the q -Whittaker cross (5.2), it depends on x, y only through the product xy .

Remark 5.3. All row sums of the cross coincide (cf. Remark 5.2):

$$\sum_{i_2, j_2 \in \{0, 1\}} \mathbb{R}_{x,y}^{(t)}(i_1, j_1; i_2, j_2) = \frac{1-xy}{1-xy} \quad \text{for every } (i_1, j_1), \quad (5.13)$$

which is the building block for the Hall–Littlewood Cauchy kernel.

Dragging the N^2 crosses $\mathbb{R}_{x_i, y_j}^{(t)}$ through the lattice exactly as in Section 1.5 yields the Yang–Baxter proof of the Cauchy identity (2.3). This Yang–Baxter proof is essentially the same as the one in [WZJ16], [Bor17], and can be traced back to [Tsi06].

Stochasticize the cross as in (5.6),

$$\mathbb{L}_{x,y}^{(t)}(i_1, j_1; i_2, j_2) := \frac{1-xy}{1-xy} \mathbb{R}_{x,y}^{(t)}(i_1, j_1; i_2, j_2), \quad (5.14)$$

and abbreviate

$$b_1 := \frac{1-xy}{1-xy}, \quad b_2 := t \frac{1-xy}{1-xy} = t b_1, \quad 0 < b_2 < b_1 < 1. \quad (5.15)$$

The six transition probabilities become

$$\begin{aligned}\mathbb{L}_{x,y}^{(t)}(1, 0; 1, 0) &= \mathbb{L}_{x,y}^{(t)}(0, 1; 0, 1) = 1, \\ \mathbb{L}_{x,y}^{(t)}(0, 0; 0, 0) &= b_1, \quad \mathbb{L}_{x,y}^{(t)}(0, 0; 1, 1) = 1 - b_1, \\ \mathbb{L}_{x,y}^{(t)}(1, 1; 1, 1) &= b_2, \quad \mathbb{L}_{x,y}^{(t)}(1, 1; 0, 0) = 1 - b_2.\end{aligned}\tag{5.16}$$

Figure 21, top row.

As in the q -Whittaker case, at the leftmost column the bijectivization is unique. With the reservoir weights $W(\infty, j_1; \infty, j_2) = x^{j_2}$ and $W^*(\infty, j_1; \infty, j_2) = y^{j_2}$ (Remark 1.11; the factors $1 - t^\infty = 1$ disappear), summing the Yang–Baxter equation over the internal occupations collapses its left-hand side to a single term, and the analogue of (5.5) reads

$$\frac{1 - txy}{1 - xy} x^{j_2} y^{j_1} = \sum_{k'_1, k'_2 \in \{0,1\}} x^{k'_2} y^{k'_1} \mathbb{R}_{x,y}^{(t)}(k'_2, k'_1; j_2, j_1).\tag{5.17}$$

The same symmetry of the cross as in the q -Whittaker case identifies the resulting unique coupling with $\mathbb{L}_{x,y}^{(t)}$ (5.14).

Dragging the crosses through the Hall–Littlewood Cauchy lattice — the construction of [BP19] — grows a field $\{\lambda^{(m,n)}\}$ of partitions on the quadrant, with $\lambda^{(m,n)}$ distributed as the Hall–Littlewood measure

$$\text{Prob}(\lambda) \propto P_\lambda(x_1, \dots, x_n; 0, t) Q_\lambda(y_1, \dots, y_m; 0, t),\tag{5.18}$$

the t -analogue of the Schur measure. Along any down-right cut the field is the *Hall–Littlewood process*, the t -deformation of the Schur process [OR03]. In the bulk the coupling is randomized (as in Section 5.1), while at the leftmost column it is the unique (5.17). Around a unit cell with corners $\varkappa = \lambda^{(m,n)}$, $\mu = \lambda^{(m+1,n)}$, $\lambda = \lambda^{(m,n+1)}$, $\nu = \lambda^{(m+1,n+1)}$, the stochasticized cross $\mathbb{L}_{x_{n+1}, y_{m+1}}^{(t)}$ governs the *lengths* of the partitions:

$$(i_1, j_1) = (\ell(\lambda) - \ell(\varkappa), \ell(\mu) - \ell(\varkappa)), \quad (i_2, j_2) = (\ell(\nu) - \ell(\mu), \ell(\nu) - \ell(\lambda)),\tag{5.19}$$

with $i_1 + j_2 = j_1 + i_2 = \ell(\nu) - \ell(\varkappa)$. This is the row-indexed counterpart of (5.7), with the first parts λ_1 replaced by the lengths $\ell(\lambda) = \lambda'_1$ (cf. Remark 1.10). In particular, the array of lengths — equivalently, of the numbers of zero parts $m_0(\lambda^{(m,n)}) = n - \ell(\lambda^{(m,n)})$ — evolves in a marginally Markovian way.

To recognize this marginal, recall the stochastic six-vertex model [GS92], [BCG16]. The general six-vertex model on $\mathbb{Z}^2$ has six admissible arrow configurations around a vertex, weighted $a_1, a_2, b_1, b_2, c_1, c_2 \geq 0$ (see Figure 21, bottom row, and [Bax07]). The weights are called *stochastic* [GS92] if

$$a_1 = a_2 = 1, \quad b_1 + c_1 = b_2 + c_2 = 1, \quad b_1, b_2 \in [0, 1].\tag{5.20}$$

Then the weight of each vertex is the conditional probability of its outgoing arrows (through the top and the right) given the incoming ones (from below and from the left), cf. Figure 1, and the model in the quadrant $\mathbb{Z}_{\geq 1}^2$ is sampled by a Markov chain, vertex by vertex along the successive diagonals $\{m + n = \text{const}\}$. Impose the *infinite domain wall* boundary conditions: a path enters from the left at every row, and no paths enter through the bottom. The *height function* is

$$H(m, n) := \#\{\text{up-right paths passing strictly to the right of } (m, n)\}, \quad 0 \leq H(m, n) \leq n,\tag{5.21}$$

see Figure 22 for an illustration.

The next statement is proven in the same way as its q -Whittaker counterpart from Section 5.1:




Figure 21: Top row: the six states $(i_1, j_1; i_2, j_2)$ of the stochasticized cross $\mathbb{L}_{x,y}^{(t)}$ (5.16) (blue = the P -line carrying x , red = the Q -line carrying y ; dotted = empty). Bottom row: the corresponding vertices of the stochastic six-vertex model, in the standard order $a_1, a_2, b_1, c_1, b_2, c_2$.




Figure 22: A configuration of the stochastic six-vertex model in the quadrant with the infinite domain wall boundary conditions, and several values of the height function (5.21).

Proposition 5.4 ([Bor18], [BBW16], [BM18], [BP19]). *Under the Yang–Baxter field $\{\lambda^{(m,n)}\}$ constructed from the bijectivization of the Hall–Littlewood Yang–Baxter equation, the array $m_0(\lambda^{(m,n)})$ evolves as the height function of the stochastic six-vertex model in the quadrant with the infinite domain wall boundary conditions, with b_1, b_2 given by (5.15) under $xy \rightsquigarrow x_n y_m$ at the coordinate (m, n) .*

In the Poisson-type limit around the diagonal ($xy \rightarrow 1$, so that $b_1, b_2 \rightarrow 0$ with the ratio $b_2/b_1 = t$ fixed by (5.15)) the six-vertex model degenerates to the ASEP with left and right jump rates t and 1 [BCG16] — the row-indexed counterpart of the q -PushTASEP produced from columns in Section 5.1. At $t = 0$ the Hall–Littlewood weights (2.1) become the Schur weights (Remark 1.4), and the height field reduces to the corresponding Schur-measure observable. The particle system along the diagonal becomes the continuous-time TASEP.

6 Further directions

As we saw in the previous sections, it is fruitful to read the Yang–Baxter equation as a *probabilistic normalization* identity between two random configurations, not only an algebraic one. Dragging a cross through a lattice gave the Cauchy identity (Section 1.5), RSK and its deformations (Sections 4 and 5), and their particle-system marginals. Below, we outline a few more ways in which a probabilistic reading of the Yang–Baxter equation has been fruitful in recent research.

6.1 Stationary measures via the forgetting property

Sending the number of arrows carried by one line of a stochastic vertex model to infinity turns that line into an inexhaustible reservoir that “forgets” how many arrows entered, so its output is independent of its input (we saw this first in Remark 1.11). Then, the Yang–Baxter equation

allows one to exchange this infinite line with another, finite stochastic line (encoding the evolution of a particle system). Since the distribution before and after applying the stochastic line is the same, such “reservoir measures” are *stationary* for particle systems. This approach was used to construct stationary measures for the colored ASEP, the colored q -Boson process, and the colored q -PushTASEP, on the ring and in the quadrant [ANP25]. For a single color the quadrant construction is due to [Agg18].

The reservoir weights are a lattice form of the multiline queues [FM07], [Mar20], and the underlying single-color Burke property appeared earlier for directed polymers [OY01], [Sep12]. Stationary measures for multiclass particle systems were described in terms of the Matrix Product Ansatz [FM05], [FM09]. The stationary measures tie to nonsymmetric Macdonald polynomials [CdGW15], [CMW22], [BW22b], [BW22a] and, for the colored q -Boson process, to modified Macdonald polynomials [AMM22], [AMM23].

6.2 Symmetry in the spectral parameters: parameter permutation and time reversal

The partition functions of Sections 1 and 2 build symmetric polynomials one row at a time, feeding in the spectral parameters in a fixed order; that they come out symmetric is a corollary of the Yang–Baxter equation, since an R -cross (of a different form than the ones we use here) dragged through two adjacent rows interchanges their parameters while fixing the partition function. Bijectivizing this single swap turns the symmetry into a local Markov move on semistandard Young tableaux of a fixed shape: the move preserves the partition function but carries the tableau distribution for one order of the parameters to the one for the swapped order. It is a probabilistic version of the Bender–Knuth involution [BK72].

Iterating this single swap up an infinite interlacing array drags the bottom parameter off to infinity, and homogeneity turns the residual rescaling into a shift $t \mapsto qt$ of the process time. At $q \rightarrow 1$ this is a time-homogeneous *backward* process: [PS22] build a backward Hammersley-type process with $\mu_t \mathbf{L}_\tau = \mu_{e^{-\tau}t}$, mapping the fixed-time law μ_t of TASEP with the step initial condition to the law at an earlier time $e^{-\tau}t$.

In the multiparameter q -Hahn TASEP each particle x_n carries its own jump parameter ν_n (coming from the vertex model spin parameter) [Pet21]. For $\nu_{n+1} < \nu_n$ the swap $\nu_n \leftrightarrow \nu_{n+1}$ is a q -deformed beta-binomial redistribution of x_n , reproducing the process with the two parameters swapped at every fixed time. The existence of this swap is due to the fact that the law of $x_n(t)$ is symmetric in $\nu_1, \dots, \nu_n$. The same elementary swap degenerates to q -TASEP and to the directed beta polymer, where it becomes a Beta-splitting of two adjacent partition functions. A random matrix shadow of this probabilistic symmetry was explored in [PT20].

The symmetry lifts from fixed-time marginals to whole trajectories. Indeed, [PS24] constructed “rewriting history” Markov operators on the space of trajectories of a particle system such as TASEP. These operators act by resampling a particle’s entire past trajectory while carrying the joint law onto that of the parameter-swapped process. Dragging to infinity gives the shift intertwiner, whose continuous limit is a Lax equation $t \dot{\mathbf{T}} = [\mathbf{B}, \mathbf{T}]$ for the TASEP semigroup $\mathbf{T} = \mathbf{T}(t)$, with $\mathbf{B}$ the generator of the backward process. The Lax equation has the same form for q -TASEP and TASEP, and differs only in the form of the generator $\mathbf{B}$ of the backward process.

6.3 Shift-invariance and hidden symmetries

The transposition of spectral parameters also surfaces in identities between *distant* observables, though the phenomenon is less a transposition than a hidden symmetry that the transposition merely records. [BGW22] prove that the joint law of colored height functions of the inhomogeneous colored stochastic six-vertex model is invariant under moving one observation point up by one lattice step and raising its color cutoff by one, at the cost of transposing two row spectral parameters. The identity is nonlocal, as the swapped rows may be arbitrarily far apart. Their proof is analytic: read the partition function as a polynomial in the spectral parameters and fix it by Lagrange interpolation at the nodes $x_i = y_j$, where the R -matrix at spectral parameter 1 degenerates to a permutation (vertex splitting); Yang–Baxter and unitarity then re-route the split partition function onto a smaller domain closed under induction. Fusion and analytic continuation in the spin erase the swap in the homogeneous limit and push the invariance down to the Beta, Gamma, and O’Connell–Yor polymers, the KPZ equation, and the Airy sheet. The degenerate specialization of this invariance is the color–position symmetry [BW22a], [BB21].

The cleanest formulation [Gal21] is a single *flip* invariance of the partition function (a 180° rotation of the boundary data with reversed row spectral parameters), of which the shift is a composition of two flips. Even this proof is algebraic, identifying the model with the Yang–Baxter basis of the Hecke algebra; independent proofs are given in [Dau22], [Zha23]. The two sides of a flip carry *equal numbers* of configurations, yet the weights match only in aggregate. It would be interesting to find useful couplings of the two sides, as we have done for the Cauchy identities in these notes.

6.4 Markov duality

Markov duality extends the integrable structure to the *dynamics* of observables of particle systems coming from stochastic vertex models. The symmetry at work is now the co-product rather than the spectral parameter: the stochastic Markov generator of a particle system commutes with the U_q -action, so a *duality function* $D(\eta, \xi)$ intertwines the generator with its space-reversed transpose,

$$\mathbb{E}_\eta[D(\eta_t, \xi)] = \mathbb{E}_\xi[D(\eta, \xi_t)].$$

The whole hierarchy of integrable particle systems (ASEP, the stochastic six-vertex model [Lin19], q -TASEP [BCS14], the higher-spin vertex models [CP16], and their dynamic/elliptic deformations) admits dualities falling out of the Yang–Baxter equation. Namely, D can be built as a matrix element of the universal R -matrix: a row of vertex weights for an auxiliary line adjoined to the lattice [KZ23].

Two regimes should be distinguished. *Triangular* dualities build D from raising operators on the reversible measure [Sch97], [GKRV09], [BS18], [Kua18]; they collapse a k -point q -moment of the height function to a closed evolution of k dual particles, solved by nested contour integrals (the engine behind the exact one-point laws) but they return moments rather than the law, so recovering a distribution still needs a (frequently ill-posed) moment problem. *Orthogonal-polynomial* dualities [KZ23] instead take D orthogonal with respect to the reversible measure; the orthogonality gives direct probabilistic access — Boltzmann–Gibbs principles, correlation identities, even Tracy–Widom asymptotics — without routing through moments.

Acknowledgements

The author is grateful to Alexei Borodin, Ben Brubaker, Matteo Mucciconi, and Greta Panova for helpful suggestions. The author was partially supported by the NSF grant DMS-2153869 and the Simons Collaboration Grant for Mathematicians 709055.

Claude Code, helped draft portions of the exposition, generate several of the TikZ figures, and carry out coding and numerical consistency checks. All definitions, statements, and proofs were verified by the author, who bears sole responsibility for their correctness.

References

- [ABPW23] A. Aggarwal, A. Borodin, L. Petrov, and M. Wheeler. Free Fermion Six Vertex Model: Symmetric Functions and Random Domino Tilings. *Selecta Math.*, 29:article 36, 2023. arXiv:2109.06718 [math.PR].
- [AF21] F. Aigner and G. Frieden. qRSt: A probabilistic Robinson–Schensted correspondence for Macdonald polynomials. *Intern. Math. Research Notices*, page rnab083, 2021. arXiv:2009.03526 [math.CO].
- [Agg18] A. Aggarwal. Current Fluctuations of the Stationary ASEP and Six-Vertex Model. *Duke Math J.*, 167(2):269–384, 2018. arXiv:1608.04726 [math.PR].
- [AMM22] A. Ayer, O. Mandelshtam, and J.B. Martin. Modified Macdonald polynomials and the multispecies zero range process: II. *arXiv preprint*, 2022. arXiv:2209.09859 [math.CO].
- [AMM23] A. Ayer, O. Mandelshtam, and J.B. Martin. Modified Macdonald polynomials and the multispecies zero-range process: I. *Algebr. Comb.*, 6(1):243–284, 2023. arXiv:2011.06117 [math.CO].
- [ANP25] A. Aggarwal, M. Nicoletti, and L. Petrov. Colored Interacting Particle Systems on the Ring: Stationary Measures from Yang–Baxter Equation. *Compositio Math.*, 161(8):1855–1922, 2025. arXiv:2309.11865 [math.PR].
- [Bax07] R. Baxter. *Exactly solved models in statistical mechanics*. Courier Dover Publications, 2007.
- [BB21] A. Borodin and A. Bufetov. Color-position symmetry in interacting particle systems. *Ann. Probab.*, 49(4):1607–1632, 2021. arXiv:1905.04692 [math.PR].
- [BBF11] B. Brubaker, D. Bump, and S. Friedberg. Schur polynomials and the Yang-Baxter equation. *Comm. Math. Phys.*, 308(2):281, 2011.
- [BBW16] A. Borodin, A. Bufetov, and M. Wheeler. Between the stochastic six vertex model and hall-littlewood processes. *arXiv preprint*, 2016. arXiv:1611.09486 [math.PR]. To appear in Jour. Comb. Th. A.
- [BC14] A. Borodin and I. Corwin. Macdonald processes. *Probab. Theory Relat. Fields*, 158:225–400, 2014. arXiv:1111.4408 [math.PR].
- [BCG16] A. Borodin, I. Corwin, and V. Gorin. Stochastic six-vertex model. *Duke J. Math.*, 165(3):563–624, 2016. arXiv:1407.6729 [math.PR].

- [BCS14] A. Borodin, I. Corwin, and T. Sasamoto. From duality to determinants for q -TASEP and ASEP. *Ann. Probab.*, 42(6):2314–2382, 2014. arXiv:1207.5035 [math.PR].
- [BDJ99] J. Baik, P. Deift, and K. Johansson. On the distribution of the length of the longest increasing subsequence of random permutations. *Jour. AMS*, 12(4):1119–1178, 1999. arXiv:math/9810105 [math.CO].
- [BGW22] A. Borodin, V. Gorin, and M. Wheeler. Shift-invariance for vertex models and polymers. *Proc. Lond. Math. Soc.*, 124:182–299, 2022. arXiv:1912.02957 [math.PR].
- [BK72] E.A. Bender and D.E. Knuth. Enumeration of plane partitions. *J. Comb. Theo. A*, 13(1):40–54, 1972.
- [BK24] A. Borodin and S. Korotkikh. Inhomogeneous spin q -Whittaker polynomials. *Ann. Fac. Sci. Toulouse Math. (6)*, 33(1):1–68, 2024. arXiv:2104.01415 [math.CO].
- [BM18] A. Bufetov and K. Matveev. Hall-Littlewood RSK field. *Selecta Math.*, 24(5):4839–4884, 2018. arXiv:1705.07169 [math.PR].
- [BMP21] A. Bufetov, M. Mucciconi, and L. Petrov. Yang-Baxter random fields and stochastic vertex models. *Adv. Math.*, 388:107865, 2021. arXiv:1905.06815 [math.PR].
- [Bor17] A. Borodin. On a family of symmetric rational functions. *Adv. Math.*, 306:973–1018, 2017. arXiv:1410.0976 [math.CO].
- [Bor18] A. Borodin. Stochastic higher spin six vertex model and macdonald measures. *Jour. Math. Phys.*, 59(2):023301, 2018. arXiv:1608.01553 [math-ph].
- [BP15] A. Bufetov and L. Petrov. Law of Large Numbers for Infinite Random Matrices over a Finite Field. *Selecta Math.*, 21(4):1271–1338, 2015. arXiv:1402.1772 [math.PR].
- [BP16] A. Borodin and L. Petrov. Nearest neighbor Markov dynamics on Macdonald processes. *Adv. Math.*, 300:71–155, 2016. arXiv:1305.5501 [math.PR].
- [BP18] A. Borodin and L. Petrov. Higher spin six vertex model and symmetric rational functions. *Selecta Math.*, 24(2):751–874, 2018. arXiv:1601.05770 [math.PR].
- [BP19] A. Bufetov and L. Petrov. Yang-Baxter field for spin Hall-Littlewood symmetric functions. *Forum Math. Sigma*, 7:e39, 2019. arXiv:1712.04584 [math.PR].
- [BS18] V. Belitsky and G. Schütz. Self-duality and shock dynamics in the n -component priority ASEP. *Stochast. Proc. Appl.*, 128(4):1165–1207, 2018. arXiv:1606.04587 [math.PR].
- [BW21] A. Borodin and M. Wheeler. Spin q -Whittaker polynomials. *Adv. Math.*, 376:107449, 2021. arXiv:1701.06292 [math.CO].
- [BW22a] A. Borodin and M. Wheeler. Colored stochastic vertex models and their spectral theory. *Astérisque*, 437, 2022. arXiv:1808.01866 [math.PR].
- [BW22b] A. Borodin and M. Wheeler. Nonsymmetric Macdonald polynomials via integrable vertex models. *Trans. Amer. Math. Soc.*, 375:8353–8397, 2022. arXiv:1904.06804 [math-ph].

- [CD21] K. Chen and X. Ding. Stable spin hall-littlewood symmetric functions, combinatorial identities, and half-space yang-baxter random field. *arXiv preprint*, 2021. arXiv:2106.12557 [math-ph].
- [CdGW15] L. Cantini, J. de Gier, and M. Wheeler. Matrix product formula for Macdonald polynomials. *Jour. Phys. A*, 48(38):384001, 2015. arXiv:1505.00287 [math-ph].
- [CMP19] I. Corwin, K. Matveev, and L. Petrov. The q-Hahn PushTASEP. *Intern. Math. Research Notices*, rnz106, 2019. arXiv:1811.06475 [math.PR].
- [CMW22] S. Corteel, O. Mandelshtam, and L. Williams. From multiline queues to Macdonald polynomials via the exclusion process. *American Journal of Mathematics*, 144(2):395–436, 2022. arXiv:1811.01024 [math.CO].
- [CP15] I. Corwin and L. Petrov. The q-PushASEP: A New Integrable Model for Traffic in 1+1 Dimension. *J. Stat. Phys.*, 160(4):1005–1026, 2015. arXiv:1308.3124 [math.PR].
- [CP16] I. Corwin and L. Petrov. Stochastic higher spin vertex models on the line. *Commun. Math. Phys.*, 343(2):651–700, 2016. arXiv:1502.07374 [math.PR].
- [Dau22] D. Dauvergne. Hidden invariance of last passage percolation and directed polymers. *Ann. Probab.*, 50(1):18–60, 2022. arXiv:2002.09459 [math.PR].
- [DCKK⁺20] H. Duminil-Copin, K. K. Kozłowski, D. Krachun, I. Manolescu, and M. Oulamarara. Rotational invariance in critical planar lattice models. *arXiv preprint*, 2020. arXiv:2012.11672 [math.PR].
- [Fad96] L.D. Faddeev. How algebraic Bethe ansatz works for integrable model. 1996. Les-Houches lectures (1996), arXiv:hep-th/9605187.
- [FM05] P.A. Ferrari and J.B. Martin. Multiclass processes, dual points and M/M/1 queues. *arXiv preprint*, 2005. arXiv:math-ph/0509045.
- [FM07] P.A. Ferrari and J.B. Martin. Stationary distributions of multi-type totally asymmetric exclusion processes. *Ann. Probab.*, 35(3):807–832, 2007.
- [FM09] P. A. Ferrari and J. B. Martin. Multiclass Hammersley-Aldous-Diaconis process and multiclass-customer queues. *Annales de l’Institut Henri Poincaré (B) Probability and Statistics*, 45(1):250–265, 2009. arXiv:0707.4202 [math.PR].
- [Fom88] S. V. Fomin. Generalized Robinson-Schensted-Knuth correspondence. *J. Math. Sci.*, 41(2):979–991, 1988. Translated from Zap. Nauchn. Semin. LOMI, vol. 155, pp. 156–175, 1986.
- [Fom95a] S. Fomin. Schensted algorithms for dual graded graphs. *J. Algebr. Comb.*, 4(1):5–45, 1995.
- [Fom95b] S. Fomin. Schur Operators and Knuth Correspondences. *Journal of combinatorial theory. Series A*, 72(2):277–292, 1995.
- [FSA24] G. Frieden and F. Schreier-Aigner. qtRSK*: A probabilistic dual RSK correspondence for Macdonald polynomials. *arXiv preprint*, 2024. arXiv:2403.16243 [math.CO].

- [FW09] O. Foda and M. Wheeler. Hall-Littlewood Plane Partitions and KP. *Int. Math. Res. Notices*, 2009(14):2597, 2009. arXiv:0809.2138 [math-ph].
- [Gal21] P. Galashin. Symmetries of stochastic colored vertex models. *Ann. Probab.*, 49(5):2175–2219, 2021. arXiv:2003.06330 [math.PR].
- [GKRV09] C. Giardinà, J. Kurchan, F. Redig, and K. Vafayi. Duality and hidden symmetries in interacting particle systems. *J. Stat. Phys.*, 135(1):25–55, 2009. arXiv:0810.1202 [math-ph].
- [GO10] A. Gnedin and G. Olshanski. q-exchangeability via quasi-invariance. *Ann. Probab.*, 38(6):2103–2135, 2010. arXiv:0907.3275 [math.PR].
- [GR04] G. Gasper and M. Rahman. *Basic hypergeometric series*. Cambridge University Press, 2004.
- [GS92] L.-H. Gwa and H. Spohn. Six-vertex model, roughened surfaces, and an asymmetric spin Hamiltonian. *Phys. Rev. Lett.*, 68(6):725–728, 1992.
- [HK05] A.M. Hamel and R.C. King. U-turn alternating sign matrices, symplectic shifted tableaux and their weighted enumeration. *J. Algebraic Combin.*, 21(4):395–421, 2005. arXiv:math/0312169 [math.CO].
- [Hop14] S. Hopkins. RSK via local transformations. *Unpublished notes*, 2014. <https://www.samuelfhopkins.com/docs/rsk.pdf>.
- [Jim86] M. Jimbo. Quantum R matrix for the generalized Toda system. *Commun. Math. Phys.*, 102:537–547, 1986.
- [Joh00] K. Johansson. Shape fluctuations and random matrices. *Commun. Math. Phys.*, 209(2):437–476, 2000. arXiv:math/9903134 [math.CO].
- [Knu73] D.E. Knuth. *The Art of Computer Programming, Vol. 3: Sorting and Searching*. Addison–Wesley, London, 1973.
- [Kor21] C. Korff. Cylindric Hecke Characters and Gromov–Witten Invariants via the Asymmetric Six-Vertex Model. *Commun. Math. Phys.*, 381(2):591–640, 2021. arXiv:1906.02565 [math-ph].
- [Kor24] S. Korotkiikh. Representation theoretic interpretation and interpolation properties of inhomogeneous spin q -Whittaker polynomials. *Selecta Math.*, 30(40), 2024. arXiv:2204.06166 [math.CO].
- [KR87] A.N. Kirillov and N. Reshetikhin. Exact solution of the integrable XXZ Heisenberg model with arbitrary spin. I. The ground state and the excitation spectrum. *J. Phys. A*, 20(6):1565–1585, 1987.
- [KRS81] P. Kulish, N. Reshetikhin, and E. Sklyanin. Yang-Baxter equation and representation theory: I. *Letters in Mathematical Physics*, 5(5):393–403, 1981.
- [Kua18] J. Kuan. An Algebraic Construction of Duality Functions for the Stochastic $U_q(A_n^{(1)})$ Vertex Model and Its Degenerations. *Comm. Math. Phys.*, 359:121–187, 2018. arXiv:1701.04468 [math.PR].

- [KZ23] J. Kuan and Z. Zhou. Orthogonal dualities of dynamic stochastic higher spin vertex models, using the drinfeld twister. *arXiv preprint*, 2023. arXiv:2305.17602 [math.PR].
- [Lie67] E.H. Lieb. Residual entropy of square ice. *Physical Review*, 162(1):162–172, 1967.
- [Lig75] Thomas M Liggett. Ergodic theorems for the asymmetric simple exclusion process. *Transactions of the American Mathematical Society*, 213:237–261, 1975.
- [Lin02] T. Lindvall. *Lectures on the Coupling Method*. Dover Publications, Mineola, NY, 2002. Corrected reprint of the 1992 Wiley original.
- [Lin19] Y. Lin. Markov Duality for Stochastic Six Vertex Model. *arXiv preprint*, 2019. arXiv:1901.00764 [math.PR].
- [Mac95] I.G. Macdonald. *Symmetric functions and Hall polynomials*. Oxford University Press, 2nd edition, 1995.
- [Man14] V. Mangazeev. On the Yang–Baxter equation for the six-vertex model. *Nuclear Physics B*, 882:70–96, 2014. arXiv:1401.6494 [math-ph].
- [Mar20] J.B. Martin. Stationary distributions of the multi-type ASEP. *Electronic Journal of Probability*, 25(43):1–41, 2020. arXiv:1810.10650 [math.PR].
- [MP17] K. Matveev and L. Petrov. q -randomized Robinson–Schensted–Knuth correspondences and random polymers. *Annales de l’IHP D*, 4(1):1–123, 2017. arXiv:1504.00666 [math.PR].
- [MP22] M. Mucciconi and L. Petrov. Spin q -Whittaker polynomials and deformed quantum Toda. *Commun. Math. Phys.*, 389(3):1331–1416, 2022. arXiv:2003.14260 [math.PR].
- [MS13] K. Motegi and K. Sakai. Vertex models, TASEP and Grothendieck polynomials. *J. Phys. A: Math. Theor.*, 46(35):355201, 2013. arXiv:1305.3030 [math-ph].
- [Nap24] S. Naprienko. Free Fermionic Schur Functions. *Adv. Math.*, 436:109413, 2024. arXiv:2301.12110.
- [NP22] M. Nicoletti and L. Petrov. Irreversible Markov Dynamics and Hydrodynamics for KPZ States in the Stochastic Six Vertex Model. *Electron. J. Probab.*, to appear, 2022. arXiv:2201.12497 [math.PR].
- [Oko01] A. Okounkov. Infinite wedge and random partitions. *Selecta Math.*, 7(1):57–81, 2001. arXiv:math/9907127 [math.RT].
- [OP13] N. O’Connell and Y. Pei. A q -weighted version of the Robinson-Schensted algorithm. *Electron. J. Probab.*, 18(95):1–25, 2013. arXiv:1212.6716 [math.CO].
- [OR03] A. Okounkov and N. Reshetikhin. Correlation function of Schur process with application to local geometry of a random 3-dimensional Young diagram. *Jour. AMS*, 16(3):581–603, 2003. arXiv:math/0107056 [math.CO].
- [OY01] N. O’Connell and M. Yor. Brownian analogues of Burke’s theorem. *Stochastic Processes and their Applications*, 96(2):285–304, 2001.

- [Pak01] I. Pak. Hook length formula and geometric combinatorics. *Sém. Lothar. Combin.*, 46(B46f):13, 2001.
- [Pet21] L. Petrov. Parameter permutation symmetry in particle systems and random polymers. *SIGMA*, 17(021):34, 2021. arXiv:1912.06067 [math.PR].
- [Pra15] A. Prasad. *Representation Theory: A Combinatorial Viewpoint*, volume 147 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2015.
- [PS22] L. Petrov and A. Saenz. Mapping TASEP back in time. *Prob. Theory Related Fields*, 182:481–530, 2022. arXiv:1907.09155 [math.PR].
- [PS24] L. Petrov and A. Saenz. Rewriting History in Integrable Stochastic Particle Systems. *Commun. Math. Phys.*, 405(300):1–75, 2024. arXiv:2212.01643 [math.PR].
- [PT20] L. Petrov and M. Tikhonov. Parameter symmetry in perturbed GUE corners process and reflected drifted Brownian motions. *Jour. Stat. Phys.*, 181:1996–2010, 2020. arXiv:1912.08671 [math.PR].
- [Ros81] H. Rost. Nonequilibrium behaviour of a many particle process: density profile and local equilibria. *Z. Wahrsch. Verw. Gebiete*, 58(1):41–53, 1981.
- [Sch61] C. Schensted. Longest increasing and decreasing subsequences. *Canad. J. Math.*, 13:179–191, 1961.
- [Sch97] G. Schütz. Duality relations for asymmetric exclusion processes. *J. Stat. Phys.*, 86(5-6):1265–1287, 1997.
- [Sep12] T. Seppäläinen. Scaling for a one-dimensional directed polymer with boundary conditions. *Ann. Probab.*, 40(1):19–73, 2012. arXiv:0911.2446 [math.PR].
- [Spi70] F. Spitzer. Interaction of Markov processes. *Adv. Math.*, 5(2):246–290, 1970.
- [Tho00] H. Thorisson. *Coupling, Stationarity, and Regeneration*. Probability and its Applications. Springer-Verlag, New York, 2000.
- [Tok88] T. Tokuyama. A generating function of strict Gelfand patterns and some formulas on characters of general linear groups. *J. Math. Soc. Japan*, 40(4):671–685, 1988.
- [Tsi06] N. Tsilevich. Quantum inverse scattering method for the q-boson model and symmetric functions. *Functional Analysis and Its Applications*, 40(3):207–217, 2006. arXiv:math-ph/0510073.
- [Vie77] G. Viennot. Une forme géométrique de la correspondance de Robinson–Schensted. In D. Foata, editor, *Combinatoire et Représentation du Groupe Symétrique*, volume 579 of *Lecture Notes in Mathematics*, pages 29–58. Springer, Berlin, 1977.
- [WZJ16] M. Wheeler and P. Zinn-Justin. Refined Cauchy/Littlewood identities and six-vertex model partition functions: III. Deformed bosons. *Adv. Math.*, 299:543–600, 2016. arXiv:1508.02236 [math-ph].
- [Zha23] L. Zhang. Shift-invariance of the colored TASEP and finishing times of the oriented swap process. *Adv. Math.*, 415:108884, 2023. arXiv:2107.06350 [math.PR].

- [ZJ09] P. Zinn-Justin. Six-Vertex, Loop and Tiling models: Integrability and Combinatorics. *arXiv preprint*, 2009. arXiv:0901.0665 [math-ph].