Counting on the variety of modules over the quantum plane

Yifeng Huang

University of Michigan $\to$ UBC Vancouver

May. 2, 2022

arXiv: 2110.15570, to appear on Algebr. Comb.

Background

Given a field $\mathbb{F}$ and $n \geq 0$, define the n-th **commuting variety** over F as

$$K_{1,n}(\mathbb{F}) := \{(A, B) \in \mathrm{Mat}_n(\mathbb{F}) \times \mathrm{Mat}_n(\mathbb{F}) : AB = BA\}.$$

(The meaning of the notation will be clear later.)

What's known:

- When $\mathbb{F} = \mathbb{C}$, the commuting variety $K_{1,n}(\mathbb{C})$ is a complex algebraic variety. Motzkin and Taussky (1955) and Gerstenhaber (1961) showed that $K_{1,n}(\mathbb{C})$ is irreducible.
- When $\mathbb{F} = \mathbb{F}_q$, the finite field of q elements, the set $K_{1,n}(\mathbb{F}_q)$ is a finite set. Feit and Fine (1960) gave its cardinality by the formula:

$$\sum_{n=0}^{\infty} \frac{|K_{1,n}(\mathbb{F}_q)|}{(q^n - 1)(q^n - q) \ldots (q^n - q^{n-1})} x^n = \prod_{i,j \geq 1} \frac{1}{1 - x^i q^{2-j}}. \qquad (1)$$

Background

We now consider a quantum deformation of the commuting variety. Let ζ be a nonzero element of $\mathbb{F}$, define the n-th ζ-**commuting variety** as

$$K_{\zeta,n}(\mathbb{F}) := \{(A, B) \in \mathrm{Mat}_n(\mathbb{F}) \times \mathrm{Mat}_n(\mathbb{F}) : AB = \zeta BA\}.$$

If $\zeta = 1$, then it simply becomes the commuting variety, hence the notation $K_{1,n}$ for the commuting variety.

Efforts have been spent to extend the work of Motzkin, Taussky and Gerstenhaber to the ζ-commuting variety:

- Chen and Wang (2018) described the irreducible components of the anti-commuting variety $K_{-1,n}(\mathbb{C})$. There are more than one, unlike the $\zeta = 1$ case.
- Chen and Lu (2019) further extended the above result to general ζ.

Main result

We give a direct generalization of Feit–Fine's formula.

Main Theorem (H., 2021).

Let ζ be a nonzero element of $\mathbb{F}_q$, and let m be the smallest positive integer such that $\zeta^m = 1$. Then

$$\sum_{n=0}^{\infty} \frac{|K_{\zeta,n}(\mathbb{F}_q)|}{(q^n - 1)(q^n - q)\dots(q^n - q^{n-1})} x^n = \prod_{i=1}^{\infty} F_m(x^i; q),$$

where

$$F_m(x;q) := \frac{1 - x^m}{(1-x)(1-x^m q)} \cdot \frac{1}{(1-x)(1-xq^{-1})(1-xq^{-2})\dots}.$$

When $\zeta = 1$, we have $m = 1$, so $F_1(x^i; q) = \prod_{j \geq 1} \dfrac{1}{1 - x^i q^{2-j}}$ and we recover Feit–Fine.

Variety of modules

The commuting variety $K_{1,n}(\mathbb{F})$ parametrizes and classifies finite-$\mathbb{F}$-dimensional modules over the polynomial ring $\mathbb{F}[X,Y]$. So $K_{1,n}(\mathbb{F})$ is also called the **variety of modules** over $\mathbb{F}[X,Y]$. To specify an $\mathbb{F}[X,Y]$-module with underlying space $\mathbb{F}^n$, it suffices to specify the x-action $A:\mathbb{F}^n \to \mathbb{F}^n$ and the y-action $B:\mathbb{F}^n \to \mathbb{F}^n$ under the constraint $AB = BA$. This constraint is because x and y commmute in $\mathbb{F}[X,Y]$.

Similarly, the ζ-commuting variety parametrizes finite-$\mathbb{F}$-dimensional modules over the associative algebra $\mathbb{F}\{X,Y\}/(XY-\zeta YX)$. This algebra is called the **quantum plane**, and is considered as a quantum deformation of $\mathbb{F}[X,Y]$.

Remarks on Main Theorem

$$\sum_{n=0}^{\infty} \frac{|K_{\zeta,n}(\mathbb{F}_q)|}{(q^n-1)(q^n-q)\dots(q^n-q^{n-1})} x^n = \prod_{i=1}^{\infty} F_m(x^i;q),$$

$$F_m(x;q) := \frac{1-x^m}{(1-x)(1-x^m q)} \cdot \frac{1}{(1-x)(1-xq^{-1})(1-xq^{-2})\dots}.$$

- The cardinality of $K_{\zeta,n}(\mathbb{F}_q)$ depends only on the order of ζ as a root of unity of $\mathbb{F}_q$. This is expected.
- The denominator $(q^n-1)(q^n-q)\dots(q^n-q^{n-1})$ is precisely the size of $\mathrm{GL}_n(\mathbb{F}_q)$. This is the natural denominator in this type of generating function. In fact, the coefficient $|K_{\zeta,n}(\mathbb{F}_q)|/|\mathrm{GL}_n(\mathbb{F}_q)|$ is the number of n-dimensional modules over the quantum plane up to isomorphism, each measured with a weight of 1/(size of automorphism group).
- Bavula (1997) classified simple modules over the quantum plane; Main Theorem should encode some statistical information about this classification.

Further breakdown

We now state a refinement of Main Theorem. Let

$$U_{\zeta,n}(\mathbb{F}_q) := \{(A,B) \in \mathrm{GL}_n(\mathbb{F}_q) \times \mathrm{Mat}_n(\mathbb{F}_q) : AB = \zeta BA\},$$

and

$$N_{\zeta,n}(\mathbb{F}_q) := \{(A,B) \in \mathrm{Nilp}_n(\mathbb{F}_q) \times \mathrm{Mat}_n(\mathbb{F}_q) : AB = \zeta BA\}.$$

It turns out that the varieties $U_{\zeta,n}(\mathbb{F}_q)$ and $N_{\zeta,n}(\mathbb{F}_q)$ are building blocks of $K_{\zeta,n}(\mathbb{F}_q)$, in the sense that

$$\sum_{n=0}^{\infty} \frac{|K_{\zeta,n}(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n = \left(\sum_{n=0}^{\infty} \frac{|U_{\zeta,n}(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n \right) \left(\sum_{n=0}^{\infty} \frac{|N_{\zeta,n}(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n \right)$$

Recall that the left-hand side is the content of Main Theorem.

Further breakdown

Refined Theorem (H., 2021)

Let m be the order of ζ. Then

- $$\sum_{n=0}^{\infty} \frac{|U_{\zeta,n}(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n = \prod_{i=1}^{\infty} G_m(x^i; q),$$

  where

  $$G_m(x;q) := \frac{1 - x^m}{(1-x)(1-x^m q)}.$$

- $$\sum_{n=0}^{\infty} \frac{|N_{\zeta,n}(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n = \prod_{i=1}^{\infty} H(x^i; q),$$

  where

  $$H(x;q) := \frac{1}{(1-x)(1-xq^{-1})(1-xq^{-2})\dots}.$$

Remarks on Refined Theorem

Refined Theorem can be interpreted as that in the formula

$$F_m(x;q) := \frac{1-x^m}{(1-x)(1-x^m q)} \cdot \frac{1}{(1-x)(1-xq^{-1})(1-xq^{-2})\dots}$$

related to the count of $\{(A,B) : AB = \zeta BA\}$, the factor $\frac{1-x^m}{(1-x)(1-x^m q)}$ is the contribution of invertible A, while the factor $\frac{1}{(1-x)(1-xq^{-1})(1-xq^{-2})\dots}$ is the contribution of nilpotent A.

Note that the latter does not depend on m, so $|N_{\zeta,n}(\mathbb{F}_q)|$ does not depend on ζ.

Ideas of proof: decomposition

- Given $A, B \in \mathrm{Mat}_n(\mathbb{F}_q)$ such that $AB = \zeta BA$, by Fitting's lemma, there is a unique direct sum decomposition ${\mathbb{F}_q}^n = V \oplus W$ such that $A(V) \subseteq V, A(W) \subseteq W$, $A|_V$ is invertible, and $A|_W$ is nilpotent.
- It turns out that B must satisfy $B(V) \subseteq V, B(W) \subseteq W$. All we need in the proof is that $\zeta \neq 0$.
- This allows $K_{\zeta,n}(\mathbb{F}_q)$ to be "decomposed" into $U_{\zeta,n}(\mathbb{F}_q)$ (requiring invertible A) and $N_{\zeta,n}(\mathbb{F}_q)$ (requiring nilpotent A), in the sense of

$$\sum_{n=0}^{\infty} \frac{|K_{\zeta,n}(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n = \left(\sum_{n=0}^{\infty} \frac{|U_{\zeta,n}(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n \right) \left(\sum_{n=0}^{\infty} \frac{|N_{\zeta,n}(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n \right)$$

Ideas of proof: nilpotent part

- To compute $|N_{\zeta,n}(\mathbb{F}_q)| = |\{(A,B) : AB = \zeta BA, A \text{ nilp}\}|$, we first fix A and count the number of B.
- The number of B only depends on the similarity class of A, so we may assume A is in the Jordan canonical form.
- The general form of B can then be determined entry-wise.
- In particular, the number of B does not depend on ζ (this works even for $\zeta = 0$).

Ideas of proof: invertible part

- To compute $|U_{\zeta,n}(\mathbb{F}_q)| = |\{(A,B) : AB = \zeta BA, A \text{ invertible}\}|$, we first fix B and count the number of A. (Opposite to the nilpotent case!!)
- Not every B contributes. In order for the number of A to be nonzero, we must have that B is similar to ζB (by the definition of similarity).
- Using the standard orbit-stabilizer argument, it suffices to count the number of similarity classes of B such that B is similar to ζB.
- **This is where m, the order of ζ, matters.** The similarity class corresponds to a finite sequence $(g_1, g_2, \dots)$ of monic polynomials over $\mathbb{F}_q$ such that g_i divides g_{i+1}. Requiring B to be similar to ζB is equivalent to requiring every g_i in the sequence of polynomials associated to B to be of the following form: $t^d + c_1 t^{d-m} + c_2 t^{d-2m} + \dots$

Pause

Second half: Cohen–Lenstra series

A broader framework

The result can be put in a general framework which I call the Cohen–Lenstra series. Let R be an $\mathbb{F}_q$-algebra with some reasonable finiteness assumption (e.g., $\mathbb{F}_q[X]$, $\mathbb{F}_q[[X,Y]]/(XY)$, or the quantum plane). Define the Cohen–Lenstra series of R as

$$\widehat{Z}_R(x) = \sum_{M/R} \frac{1}{|\mathrm{Aut}\, M|} x^{\dim_{\mathbb{F}_q} M},$$

where M runs over all isomorphism classes of finite-cardinality modules over R. This framework has the advantages of:

- can unify many matrix enumeration problems of very distinct flavors, by varying R;
- can easily tell new problems from old.
- has local-global if R is commutative: $\widehat{Z}_R(x) = \prod_{p \in \mathrm{Specm}\, R} \widehat{Z}_{R_p}(x)$.

  (follows naturally from the definition)

Cohen–Lenstra series: History

The concept of Cohen–Lenstra series has been considered before in various formulations.

- Cohen and Lenstra, 1984: considered the series for Dedekind domain R, in the same formulation.
- Feit and Fine, 1960: The generating function they gave for counting commuting matrices matches the series for $\mathbb{F}_q[X, Y]$.
- Bryan and Morrison, 2015: reinterpreted Feit–Fine in motivic Donaldson–Thomas theory, where they considered a generating series for motivic classes of the stack of coherent sheaves over R, which is a refined version of $\widehat{Z}_R(x)$.

Matrix formulation

If $R = \mathbb{F}_q[X_1, \ldots, X_m]/(f_1, \ldots, f_r)$, consider the variety of modules over R:

$$K_n(R)(\mathbb{F}_q) = \{A_1, \ldots, A_m \in \mathrm{Mat}_n(\mathbb{F}_q) : [A_i, A_j] = 0, f_k(\underline{A}) = 0\}.$$

Then $K_n(R)$ parametrizes (with some nonuniqueness) modules over R that are n-dimensional over $\mathbb{F}_q$, and

$$\widehat{Z}_R(x) = \sum_{n \geq 0} \frac{|K_n(R)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n.$$

Punctual version: If $R = \mathbb{F}_q[[X_1, \ldots, X_m]]/(f_1, \ldots, f_r)$ where $f_k(\underline{0}) = 0$, the variety of modules over R is

$$K_n(R)(\mathbb{F}_q) = \{A_1, \ldots, A_m \in \mathrm{Nilp}_n(\mathbb{F}_q) : [A_i, A_j] = 0, f_k(\underline{A}) = 0\}.$$

Matrix enumeration interpreted as Cohen–Lenstra

By varying R, many generating series of matrix enumeration problems can be viewed as $\widehat{Z}_R(x)$:

- $R = \mathbb{F}_q[[X]]$: $\widehat{Z}_R(x)$ is a generating series for $|\mathrm{Nilp}_n(\mathbb{F}_q)|$ ($= q^{n^2-n}$ by Fine–Herstein 1958)
- $R = \mathbb{F}_q[X,Y]$: $\widehat{Z}_R(x)$ is the series counting commuting matrices computed in Feit–Fine 1960.
- $R = \mathbb{F}_q[X,Y]/(XY)$: $\widehat{Z}_R(x)$ counts mutually annihilating matrices $AB = BA = 0$ (computed in H., 2021).
- R is the quantum plane $\mathbb{F}_q\{X,Y\}/(XY - \zeta YX)$: $\widehat{Z}_R(x)$ is the series in Main Theorem.

Behaviors of the Cohen–Lenstra series

The status of knowledge of $\widehat{Z}_R(x)$ is best summarized in terms of the algebraic geometry of R:

- R noncommutative: mostly unknown, except Main Theorem.

For the below, R is commutative and $X = \operatorname{Spec} R$:

- X a smooth curve: $\widehat{Z}_R(x) = \prod_{j\geq 1} Z_X(xq^{-j})$, where Z_X is the Hasse–Weil zeta.
- X a smooth surface: $\widehat{Z}_R(x) = \prod_{i,j\geq 1} Z_X(x^i q^{-j})$.

The two above follow from classical formulas for $R = \mathbb{F}_q[X], \mathbb{F}_q[X,Y]$ and local-global.

- X a nodal singular curve: has an explicit formula with mysterious combinatorics (H., 2021)
- X other singular curve: has conjectural patterns (H., 2021)
- Fun fact: $R = \mathbb{F}_q$ (a point) is not a trivial case; in fact $\widehat{Z}_R(x)$ is the Rogers–Ramanujan series! $(1, 4 \mod 5 ...)$
- All other cases: wide open.

A dream

Since the combinatorial behavior of $\widehat{Z}_R(x)$ depends heavily on the algebraic geometry of R, **any observation about $\widehat{Z}_R(x)$ may suggest interesting geometry of R.**

In the rest of the talk, we will discuss an elementary observation about our Refined Theorem that seems to owe a higher-level explanation. It might inspire interesting noncommutative geometry of the quantum plane.

A cut-and-paste

Let's consider the following series:

$$\widehat{Z}_{\mathbb{F}_q[X]}(x) = \sum_{n \geq 0} \frac{|\mathrm{Mat}_n(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n$$

$$\widehat{Z}_{\mathbb{F}_q[X,X^{-1}]}(x) = \sum_{n \geq 0} \frac{|\mathrm{GL}_n(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n = \frac{1}{1-x}$$

$$\widehat{Z}_{\mathbb{F}_q[[X]]}(x) = \sum_{n \geq 0} \frac{|\mathrm{Nilp}_n(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n$$

Local-to-global $\implies$ $\widehat{Z}_{\mathbb{F}_q[X]}(x) = \widehat{Z}_{\mathbb{F}_q[X,X^{-1}]}(x)\widehat{Z}_{\mathbb{F}_q[[X]]}(x)$.
Essentially the cut-and-paste $\mathbb{A}^1 = (\mathbb{A}^1 \setminus \{0\}) \sqcup \{0\}$.
Since $\mathrm{Mat}_n(\mathbb{F}_q)$ is easy to count (there are q^{n^2} matrices), this computes $|\mathrm{Nilp}_n(\mathbb{F}_q)|$ (an alternative proof of Fine–Herstein).

Global-to-local

The decomposition above is essentially the same phenomenon as

$$\sum_{n=0}^{\infty} \frac{|K_{\zeta,n}(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n = \left(\sum_{n=0}^{\infty} \frac{|U_{\zeta,n}(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n \right) \left(\sum_{n=0}^{\infty} \frac{|N_{\zeta,n}(\mathbb{F}_q)|}{|\mathrm{GL}_n(\mathbb{F}_q)|} x^n \right).$$

It appears that $\widehat{Z}_{\mathbb{F}_q[X,X^{-1}]}(x)$ and $\widehat{Z}_{\mathbb{F}_q[[X]]}(x)$ are "disjoint and independent" building blocks, just like $|U_{\zeta,n}(\mathbb{F}_q)|$ vs $|N_{\zeta,n}(\mathbb{F}_q)|$. But the story doesn't end here – $\widehat{Z}_{\mathbb{F}_q[X,X^{-1}]}(x)$ and $\widehat{Z}_{\mathbb{F}_q[[X]]}(x)$ are amazingly interconnected!
Key: $\mathbb{A}^1 \setminus \{0\}$ is composed of closed points, each of which "looks like" the origin.

Global-to-local

Recall local-to-global for $R = \mathbb{F}_q[X, X^{-1}]$:

$$\widehat{Z}_{\mathbb{F}_q[X,X^{-1}]}(x) = \prod_p \widehat{Z}_{R_p}(x).$$

Now, each R_p is smooth of dimension one, so its completion is isomorphic to a power series ring over some finite field. Therefore, all $\widehat{Z}_{R_p}(x)$ is essentially $\widehat{Z}_{\mathbb{F}_q[[X]]}(x)$ (though a substitution is needed), so $\widehat{Z}_R(x)$ is determined by $\widehat{Z}_{\mathbb{F}_q[[X]]}(x)$.

Moreover, because $\widehat{Z}_R(x)$ is determined by $\widehat{Z}_{\mathbb{F}_q[[X]]}(x)$ alone, we can reverse this process ("global-to-local") and write down a formula that recovers $\widehat{Z}_{\mathbb{F}_q[[X]]}(x)$ from $\widehat{Z}_R(x) = 1/(1-x)$. (Yet another proof of Fine–Herstein!)

This idea is due to Bryan and Morrison (2015) in a more refined language of motivic classes. They actually did the case for $R = \mathbb{F}_q[X, Y]$ (namely, $K_{1,n}$).

How about the quantum plane?

Noncommutative Question.

Can you recover $|U_{\zeta,n}(\mathbb{F}_q)|$ from $|N_{\zeta,n}(\mathbb{F}_q)|$ (or vice versa), using the geometry of the quantum plane?

- The commutative analogue requires the notion of "localization at a prime ideal". Does this notion exist for the quantum plane?
- Recall that $|U_{\zeta,n}(\mathbb{F}_q)|$ depends on ζ (or the order of ζ), while $|N_{\zeta,n}(\mathbb{F}_q)|$ doesn't. So if $|N_{\zeta,n}(\mathbb{F}_q)|$ recovers $|U_{\zeta,n}(\mathbb{F}_q)|$, it must be because it also takes account of some geometry of the quantum plane that depends on ζ.

Final takeaway

- We extend a formula that counts matrix pairs $AB = BA$ to the case $AB = \zeta BA$ where ζ is nonzero. The answer depends on the order of ζ as a root of unity.
- The count of $AB = \zeta BA$ encodes statistical information about modules over the quantum plane.
- The count in question has two seemingly independent building blocks that turn out to be interdependent in the $\zeta = 1$ case, using ingredients from (commutative) algebraic geometry. I hope that the study of a possible interdependence in the case of general ζ will inspire interesting noncommutative geometry.