ANTS XVII ▪ Seventeenth Algorithmic Number Theory Symposium

On the Logarithmic Density of Rank ≥ 1 and Rank ≥ 2 Genus-2 Jacobians

and applications to hyperelliptic curve cryptography

(tentative slides - actual presentation might differ)

George C. Țurcaș

Babeș–Bolyai University & certSIGN

the two lower bounds ▪ 15 min

Răzvan Bărbulescu

CNRS, IMB & Inria Bordeaux

cryptographic applications ▪ 15 min

joint work with Mugurel Barcău & Vicențiu Pașol

Bernoulli Institute, University of Groningen • 6–10 July 2026

arXiv:2601.17142 • ePrint 2026/110

The problem and the ordering

For a genus-2 curve $C/\mathbb{Q}$, the Mordell–Weil group decomposes as $\mathrm{Jac}(C)(\mathbb{Q}) \simeq \mathbb{Z}^r \oplus T$ with T finite; the invariant of interest is the rank $r = \operatorname{rank} \mathrm{Jac}(C)(\mathbb{Q})$.

Ordering by height. Following Booker et al., we enumerate integral Weierstrass models

$$C : \; y^2 + h(x)\, y = f(x), \qquad \deg f \le 6, \;\; \deg h \le 3, \;\; h_i \in \{0, 1\},$$

by the naive height $\mathrm{H}(C) = \max_i |a_i|$. (The simplified model $y^2 = f(x)$ is the case $h = 0$.) The box of such models of height $\le X$ satisfies

$$\#S_1(X) = \#\{(f, h) : |a_i| \le X, \;\; h_i \in \{0, 1\}\} \asymp \; X^7.$$

The difficulty. Curves of positive rank are expected to have *proportion zero*, so ordinary density is blind to them. We measure them on a logarithmic scale.

Logarithmic density and the main theorem

Definition

A set $S \subseteq S_1$ has *logarithmic density* $\alpha = \liminf_{X\to\infty} \dfrac{\log \#S(X)}{\log \#S_1(X)}$. Since $\#S_1(X) \asymp X^7$, this is the growth exponent of S divided by 7.

Theorem (Bărbulescu–Barcău–Pașol–Ț., 2026)

With respect to this height ordering, **unconditionally**:

$$\operatorname{rank} \mathrm{Jac}(C)(\mathbb{Q}) \geq 1 : \text{logarithmic density} \;\geq\; \tfrac{13}{14} \qquad \text{(Corollary 2.7)},$$
$$\operatorname{rank} \mathrm{Jac}(C)(\mathbb{Q}) \geq 2 : \text{logarithmic density} \;\geq\; \tfrac{5}{7} \qquad \text{(Corollary 2.11)}.$$

Two parallel arguments

The proofs for the two lower bounds follow the *same* four-step pattern.

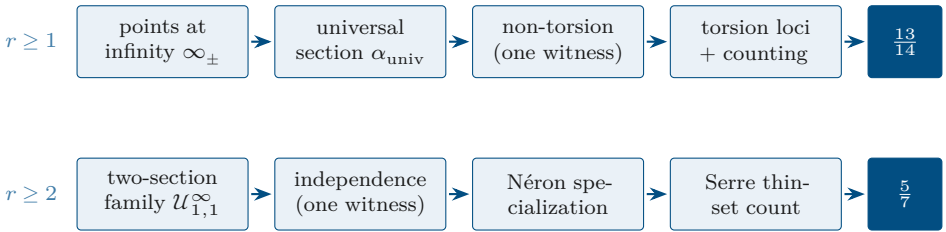

Both rest on a divisor class that is rational *for free*, a *single* explicit witness fixing the generic behaviour, and a count of the exceptional locus.

Part I ▪ George C. Țurcaș

Rank ≥ 1: logarithmic density $\frac{13}{14}$

A divisor class supported at infinity is, generically, non-torsion.

Rank ≥ 1: the points at infinity

Complete the square on $y^2 + h(x)y = f(x)$:

$$g(x) = 4f(x) + h(x)^2, \qquad c = \mathrm{coeff}_{x^6}(g).$$

If $\deg g = 6$, the two points at infinity $\infty_\pm$ are defined over $\mathbb{Q}(\sqrt{c})$, and

$$\infty_+, \infty_- \in C(\mathbb{Q}) \quad \Longleftrightarrow \quad c \in (\mathbb{Q}^\times)^2.$$

In that case we obtain a canonical rational class

$$\alpha = [\infty_+ - \infty_-] \in \mathrm{Jac}(C)(\mathbb{Q}).$$

Observation

If α is *non-torsion*, then $\operatorname{rank} \mathrm{Jac}(C)(\mathbb{Q}) \geq 1$ — a rational *Jacobian point* (a divisor class), obtained for free from the model.

The universal section

Let (f, h) vary. Over the smooth (discriminant $\neq 0$) sextic locus, adjoin $\sqrt{c}$ to form the double cover

$$\mathcal{U}^\infty = \{u^2 = c\} \longrightarrow \{\deg g = 6\}.$$

There the Jacobian is an abelian scheme $\mathcal{J} \to \mathcal{U}^\infty$ carrying the *universal section*

$$\alpha_{\text{univ}} = [\infty_+ - \infty_-] \in \mathcal{J}(\mathcal{U}^\infty).$$

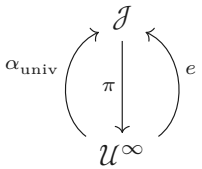

$\pi \circ \alpha_{\text{univ}} = \text{id}$, e the zero section.

Reduction

Each specialization at which α_{univ} stays non-torsion has $\operatorname{rank} \mathrm{Jac} \geq 1$. So the question becomes: how often is α_{univ} *torsion* after specialization? — which we now show is rare, in two steps.

Fact 1: a single witness suffices

To prove α_{univ} is non-torsion it is enough to exhibit *one* specialization where α has infinite order: a relation $n\,\alpha_{\mathrm{univ}} = 0$ would specialize identically.

Witness (Tengely)

For

$$C:\ y^2 = x^6 + 18x^5 + 75x^4 + 120x^3 + 120x^2 + 72x + 28,$$

the Jacobian has $\operatorname{rank} \mathrm{Jac}(C)(\mathbb{Q}) = 1$, generated by $[\infty_+ - \infty_-]$.

Hence α_{univ} is non-torsion. (The parameter space splits into 16 sheets indexed by $h \in \{0,1\}^4$; one LMFDB witness per sheet covers all of them.)

Fact 2a: only finitely many torsion loci can occur

The n-torsion locus

$$\mathcal{V}_n = \{\, \underline{a} \in \mathcal{U}^\infty : \; n\,\alpha_{\mathrm{univ}}(\underline{a}) = 0 \,\}$$

is a *proper* Zariski-closed subset, defined over $\mathbb{Q}$ (proper by Fact 1).

Uniform torsion bound

$|\operatorname{Jac}(C)(\mathbb{Q})_{\mathrm{tors}}| \ll (\log \mathrm{H}(C))^2$ (Kieffer θ-heights + Pazuki + Gaudron–Rémond).

Hence a torsion $\alpha_{\underline{a}}$ has order $n \ll (\log X)^2$, so

$$T(X) \;\subseteq \bigcup_{n \ll (\log X)^2} \mathcal{V}_n$$ — only get contribution from $O((\log X)^2)$ torsion

Fact 2b: counting the torsion loci

Degree. For a symmetric, relatively very ample L,

$$[n]^*L \;\simeq\; L^{\otimes n^2},$$

so in the induced projective coordinates $\mathcal{V}_n$ is cut out by equations of degree $O(n^2)$ — not n^4.[1]

Count & sum. A hypersurface estimate gives

$$\#(\mathcal{V}_n(\mathbb{Z}) \cap \{\mathrm{H} \le X\}) \ll n^2 X^6 \implies \#T(X) \ll X^6 \sum_{n \ll (\log X)^2} n^2 \;\ll\; X^6 (\log X)^6.$$

[1] Multiplication by n on the abelian surface has degree $n^{2g} = n^4$; the input here is the *polarization scaling* $[n]^*L \simeq L^{\otimes n^2}$, which controls the degree in projective coordinates.

Counting gives $\frac{13}{14}$

On the subfamily $S_1^{\square}(X)$ where $\infty_{\pm}$ are rational (c a square):

$$\#S_1^{\square}(X) \asymp \underbrace{X^{1/2}}_{c=\square} \cdot \underbrace{X^{6}}_{a_5,\dots,a_0} = X^{13/2}, \qquad \#T(X) \ll X^6(\log X)^6 = o(X^{13/2}).$$

Thus all but a negligible subset have α non-torsion, hence $\operatorname{rank}\operatorname{Jac}(C)(\mathbb{Q}) \geq 1$. Inside the box $\#S_1(X) \asymp X^7$, the logarithmic density is

$$\frac{13/2}{7} = \frac{13}{14}.$$

Corollary 2.7 (unconditional)

$$\liminf_{X\to\infty} \frac{\log \#\{(f,h) \in S_1(X) : \operatorname{rank}\operatorname{Jac}(C)(\mathbb{Q}) \geq 1\}}{\log \#S_1(X)} \geq \frac{13}{14}.$$

Part II ▪ George C. Țurcaș

Rank ≥ 2: logarithmic density $\frac{5}{7}$

Same machine — now we need two independent sections.

Rank ≥ 2: a two-section family

We need a second independent class, so restrict to the 5-parameter family ($h = 0$, $a_6 = a_0 = 1$):

$$C_{\underline{a}}: \; y^2 = x^6 + a_5x^5 + a_4x^4 + a_3x^3 + a_2x^2 + a_1x + 1.$$

- $a_6 = 1 \Rightarrow c = 4$ is a square, so $\infty_\pm$ are rational;
- $a_0 = 1 \Rightarrow P = (0, -1) \in C_{\underline{a}}(\mathbb{Q})$.

This yields *two* universal sections on $\mathcal{U}^{\infty}_{1,1}$:

$$\alpha_{\text{univ}} = [\infty_+ - \infty_-], \qquad \beta_{\text{univ}} = [P - \infty_+].$$

Observation

If the *specialized* classes $\alpha_{\underline{a}}, \beta_{\underline{a}}$ are $\mathbb{Z}$-independent, then $\operatorname{rank} \operatorname{Jac}(C_{\underline{a}})(\mathbb{Q}) \geq 2$.

Generic independence of the sections is not enough by itself — the next two slides upgrade it to almost every specialization.

Independence from one witness

Again a single specialization controls the generic fibre.

Witness (LMFDB `15625.a.15625.1`)

For $C_0 : y^2 = x^6 + 8x^5 + 10x^4 + 10x^3 + 5x^2 + 2x + 1$, one has $\mathrm{Jac}(C_0)(\mathbb{Q}) \simeq \mathbb{Z}^2$, with generators G_1, G_2 satisfying

$$G_2 = \beta_{\mathrm{univ}}, \qquad G_1 - 2G_2 = \alpha_{\mathrm{univ}}.$$

Hence α_{univ} and β_{univ} generate a rank-2 subgroup at $\underline{a}_0$, so they are $\mathbb{Z}$-independent in $\mathcal{J}(\mathcal{U}^{\infty}_{1,1})$.

Néron specialization and Serre's count

Néron specialization. The map

$$\mathrm{sp}_{\underline{a}} : \Gamma = \langle \alpha_{\mathrm{univ}}, \beta_{\mathrm{univ}} \rangle \longrightarrow \mathrm{Jac}(C_{\underline{a}})(\mathbb{Q})$$

is injective for all $\underline{a}$ outside a *thin* set $\Omega \subset \mathcal{U}^{\infty}_{1,1}(\mathbb{Q})$; for $\underline{a} \notin \Omega$, $\mathrm{rank}\,\mathrm{Jac}(C_{\underline{a}})(\mathbb{Q}) \geq 2$.

Serre's estimate for affine thin sets ($n = 5,\ d = 1$):

$$\#(\Omega \cap [-X, X]^5) \ \ll \ X^{9/2}(\log X)^{\gamma}.$$

The exceptional locus is a vanishing fraction of the $\asymp X^5$ family.

exponent comparison	
family	$\asymp X^5$
bad set Ω	$\ll X^{9/2+o(1)}$
ratio	$\ll X^{-1/2+o(1)} \to 0$

Counting gives $\frac{5}{7}$

Exactly as in the rank-1 count, with the five free coefficients $a_1, \dots, a_5$:

$$\#\,\mathcal{U}^{\infty}_{1,1}(\mathbb{Z}; X) \;\asymp\; X^5, \qquad \#\{\text{rank} < 2\} \;\ll\; X^{9/2}(\log X)^{\gamma} = o(X^5).$$

All but a negligible subset have $\operatorname{rank} \operatorname{Jac}(C_{\underline{a}})(\mathbb{Q}) \geq 2$. This X^5 family sits inside the box $\#S_1(X) \asymp X^7$, giving logarithmic density $\frac{5}{7}$.

Corollary 2.11 (unconditional)

$$\liminf_{X \to \infty} \frac{\log \#\{(f,h) \in S_1(X) : \operatorname{rank} \operatorname{Jac}(C)(\mathbb{Q}) \geq 2\}}{\log \#S_1(X)} \;\geq\; \frac{5}{7}.$$

Context: comparison with elliptic curves

Family / rank	log. density	status
Elliptic curves, $r \geq 2$	$\frac{19}{24} \approx 0.79$	heuristic (random matrix)
Genus 2, $r \geq 1$	$\frac{13}{14} \approx 0.93$	unconditional
Genus 2, $r \geq 2$	$\frac{5}{7} \approx 0.71$	unconditional

Orientation only: the elliptic-curve entry is heuristic (random-matrix) and uses a *different* (H_2, Bhargava–Gross) height ordering — not a theorem-to-theorem comparison. The genus-2 bounds are unconditional.

Heuristic. If rank *parities* equidistribute in the two-points-at-infinity subfamily, the same construction would suggest $\frac{13}{14}$ for $r \geq 2$ as well — but this is only heuristic.

→ *Over to Răzvan: what this abundance buys for quantum cryptanalysis.*

Part III ▪ Răzvan Bărbulescu

Cryptographic applications

Why high Mordell–Weil rank is the resource for Regev-style quantum DLP.

Hyperelliptic curve cryptography and the Regev advantage

HECC (Köblitz, 1989) is the genus-2 analogue of ECC: the DLP lives in $G = \mathrm{Jac}(C)(\mathbb{F}_q)$, $n = \log_2 |G|$.

- **Shor:** $\approx n$ group operations per run.
- **Regev** (2023) + Ekerå–Gärtner (2024) for DLP: $\approx d + \frac{n}{d}$ per run — an extra parameter d.
- Advantage over Shor: $\min\!\left(d, \frac{n}{d}\right)$.
- Increasing d = finding rational Mordell–Weil generators on a lift / twist.

cost per run (group ops)	
Shor	$\approx n$
Regev	$\approx d + \frac{n}{d}$
advantage	$\min\!\left(d, \frac{n}{d}\right)$
best d, fields	$\sqrt{n}\,(\log n)^{O(1)}$
best d, curves	$(\log n)^{1/2}$

fields: Pilatte 2026 • curves: BBP 2025.

The link to our theorems. A rank-r lift gives $d = r + 2$, so large d wants large rank — and the density bounds say high-rank lifts are *abundant*.

Regev in a nutshell: rank supplies the parameter

For the chosen d, take auxiliary $g_1, \dots, g_{d-2}$, set $g_{d-1} = h$ (the DLP target), $g_d = g$ (the base), and form

$$L = \left\{ z \in \mathbb{Z}^d : \sum_{i=1}^{d} [z_i] g_i = 0 \right\}.$$

A superposition over multi-scalar products with $\frac{n}{d}$-bit coefficients $\to$ a quantum basis of $L \to$ a linear solve returns $(0, \dots, 0, -1, *)$ whose last entry is $\log_g h$.

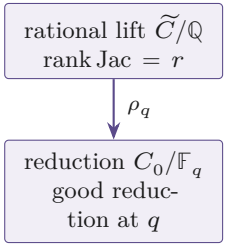

$d = r + 2$

$P_1, \dots, P_r \in \mathrm{Jac}(\widetilde{C})(\mathbb{Q})$ indep. mod torsion, reduce to $\overline{P}_i = \rho_q(P_i)$: the r auxiliaries.

So $d - 2 = r$. Reducing r independent Mordell–Weil generators supplies the r auxiliaries, hence $d = r + 2$ — under the usual independence heuristic (it can fail in degenerate Ekerå–Gärtner cases).

Two search problems — and what the density buys

High Rank Curve Problem

Sample a curve of rank $\mathrm{Jac} \geq r$ from the height-$\leq X$ family — the input to preprocessing / experiments.

High Rank Twist Problem

For a *fixed* $C/\mathbb{F}_q$: find a twist / lift of higher rank whose reduction is compatible with C.

The bridge. Brute-force sampling from $\mathcal{C}_1(X)$ finds a curve of

$$\begin{aligned} \text{rank} \geq 1: &\quad O\big(X^{1/2+o(1)}\big) \text{ trials,} \\ \text{rank} \geq 2: &\quad O\big(X^{2+o(1)}\big) \text{ trials,} \end{aligned}$$

since $\#\text{trials} \approx |\mathcal{C}_1(X)|^{1-\alpha}$ and $1 - \frac{13}{14}$, $1 - \frac{5}{7}$ give exponents $\frac{1}{2}, 2$ over the X^7 box.

The density theorems *are* the guarantee. $\frac{13}{14}$ and $\frac{5}{7}$ say exactly that the high-rank preprocessing input is easy to find — polynomially few random trials.

Experiments: curves from the cryptographic literature

Family (source)	search / outcome	best d
$y^2 = x^5 + 17$ (Buhler–Koblitz), CM, simple Jac	twists 17δ, $\delta \le 128{\cdot}10^4$ ($\approx$125 core-h) $\to$ rank 7	8
$y^2 = x^5 + a$, $a \le 10^5$ (CM search)	15/11/2/2 twists of rank 5/6/7/8	10
Split Jac $\sim E \times E$ (Freeman–Satoh; Dryło)	rank-4 elliptic twist $\to$ genus-2 rank 8 (splitting field)	10

- A rank-6 twist of $y^2 = x^5 + 17$ gives $d = 8$: an eightfold speedup over Shor on the matching $\mathbb{F}_q$ instance.
- The comparable Booker et al. (2016) set contains *no* genus-2 curve of rank ≥ 5.

Caveat: "suitable for Regev" = suitable as an attack / preprocessing input — not a secure-curve recommendation; a fixed $\mathbb{F}_q$-curve still needs a compatible lift or twist.

Summary

- Rank ≥ 1: logarithmic density $\frac{13}{14}$ — via the free class $[\infty_+ - \infty_-]$ and counting of torsion loci.
- Rank ≥ 2: logarithmic density $\frac{5}{7}$ — via the two-section family and Néron–Serre specialization.
- **One method behind both:** a universal section, a single explicit witness, specialization, and a count of a thin or closed exceptional set.
- Cryptographic pay-off: high-rank genus-2 Jacobians are abundant in the height ordering, so the Regev preprocessing input ($d = r + 2$) is easy to find — with explicit rank-7/8 twists on curves from the literature.

Thank you!

Bărbulescu, Barcău, Pașol & Ț. · arXiv:2601.17142 • ePrint 2026/110