Comparing Three Presentations of the Absolute Galois Group of $\mathbb{Q}_2$

July 26, 2026

Abstract

We compare three proposed profinite presentations of $G_{\mathbb{Q}_2}$: the presentation proved in the main paper, David Roe's shorter presentation, and a new hybrid presentation. The comparison uses a five-component complexity vector measuring straight-line word length, the marked maximal pro-2 normalization, linear Fox calculus, the quadratic determinant calculation, and the integral unramified marking. The purpose is not to assign a canonical scalar complexity, but to make the tradeoffs between the three presentations explicit and reproducible. The paper's presentation is marking-transparent but word-heavy; Roe's is shortest and cohomologically transparent but has a harder integral normalization; the hybrid preserves the paper's explicit marked pro-2 boundary while substantially shortening its class-two collector.

1 Common setting

We use the convention

$$x^g = g^{-1}xg, \qquad [x,y] = x^{-1}y^{-1}xy.$$

Let $\omega_2 \in \widehat{\mathbb{Z}}$ be the idempotent which projects to 1 in $\mathbb{Z}_2$ and to 0 in every odd-primary component. Each presentation has generators

$$\sigma, \quad \tau, \quad x_0, \quad x_1,$$

the closed normal subgroup generated by x_0, x_1 is required to be pro-2, and the first relation is

$$\tau^\sigma = \tau^2. \tag{1}$$

Only the second, or wild, relation varies.

1.1 The presentation in the paper

Put

$$\begin{aligned}
\sigma_2 &= \sigma^{\omega_2}, & u_i &= (x_i\tau)^{\omega_2} \quad (i=0,1), & d_0 &= u_0 x_0^{-1}, \\
z_0 &= x_0^{\sigma_2}, & c_0 &= [d_0, z_0], & g_0 &= \sigma_2^2, \\
d_g &= d_0^{g_0}, & h_c &= [d_g, d_0], & h_0 &= x_0^{g_0} x_0\, d_g d_0 d_0^2 h_c.
\end{aligned}$$

The wild relation is

$$r_{\text{paper}} = h_0 u_1^{-1} x_1^\sigma c_0 = 1. \tag{2}$$

1.2 Roe's presentation

Put

$$\sigma_2 = \sigma^{\omega_2}, \qquad a = (x_0^{-3}\tau)^{\omega_2}, \qquad y_1 = x_1^{\sigma_2}, \qquad c = [x_1, y_1].$$

The wild relation is

$$r_{\mathrm{Roe}} = (x_0^{\sigma})^{-1} a\, x_1^2 c = 1. \tag{3}$$

Equivalently,

$$r_{\mathrm{Roe}} = (x_0^{\sigma})^{-1} (x_0^{-3}\tau)^{\omega_2} x_1^2 [x_1, x_1^{\sigma_2}].$$

1.3 The hybrid presentation

Use

$$\sigma_2 = \sigma^{\omega_2}, \qquad u_i = (x_i\tau)^{\omega_2}, \qquad d = u_0 x_0^{-1}.$$

The proposed hybrid wild relation is

$$r_* = (x_0 d^{-1})^{\sigma_2^2} x_0 d\, u_1^{-1} x_1^{\sigma} [d, x_0^{\sigma_2}] = 1. \tag{4}$$

If

$$y = x_0 d^{-1}, \qquad g = \sigma_2^2,$$

then its first block has the equivalent twisted-square form

$$(x_0 d^{-1})^g x_0 d = y^g y d^2. \tag{5}$$

2 The comparison methodology

2.1 Why a vector rather than a scalar

Tietze and Nielsen transformations can trade short relators for complicated generators, and auxiliary definitions can make a printed relator appear artificially short. Moreover, the proof in the paper does not use the word uniformly: different quotients of the word control the maximal pro-2 boundary, the linear deformation complex, the mixed cup product, and the quadratic determinant obstruction. We therefore use

$$C(P) = \big(L(P),\, C_{\text{pro-2}}(P),\, C_{\mathrm{Fox}}(P),\, C_{\mathrm{quad}}(P),\, C_{\mathrm{mark}}(P)\big). \tag{6}$$

The entries should not be added without first choosing application dependent weights.

2.2 Straight-line word complexity

To make $L(P)$ reproducible, we fix the following grammar.

- Generators and previously defined auxiliary words have cost 0.
- Each multiplication, inversion, conjugation, commutator, integer power, and ω_2-power has cost 1.
- Multiplication of n displayed factors costs $n-1$.
- Each auxiliary definition is charged once; later reuse is free.

- The common tame relation (1) is omitted, since it contributes equally to all three presentations.

Thus L is straight-line-program size rather than the number of letters in the fully expanded free word. In particular, an ω_2-power is charged as one primitive profinite operation. This convention is not intrinsic, but it prevents free abbreviations from distorting the comparison.

2.3 Proof-complexity coordinates

The remaining four entries are ordinal scores:

1	direct calculation or explicit one-line normalization,
2	one short structural lemma or a small case split,
3	a substantial word ledger, orientation calculation, or auxiliary normalization.

More specifically:

$C_{\text{pro-2}}$ The work needed to identify the maximal pro-2 quotient abstractly with the rank-three dyadic Demushkin group.

C_{Fox} The candidate-specific work needed to obtain the evaluated simple-factor Fox Jacobian.

C_{quad} The work needed to identify the mixed Heisenberg pairing and the extraspecial determinant quadratic form.

C_{mark} The additional work needed to identify the full integral unramified character, not merely its reduction modulo 2.

These scores measure the proof route furnished by the methods of the paper. They do not measure the computational speed of evaluating the word in a finite group, and they are not evidence of correctness by themselves.

3 Straight-line operation counts

The detailed count is as follows.

component	paper	Roe	hybrid
definitions before the principal block	12	6	7
principal square or collector block	7	1	6
remaining operations in the final relator	5	5	7
$L(P)$	24	12	20

Here the hybrid's principal-block entry is the cost of forming

$$(x_0 d^{-1})^{\sigma_2^2} x_0 d$$

as a single reusable block. The remaining entry includes u_1^{-1}, x_1^{σ}, the conjugated word $x_0^{\sigma_2}$, the commutator, and the three multiplications joining the four top-level factors. Equivalently, counting directly from (4), the definitions have cost 7 and the final relation has cost 13.

For clarity, the totals can also be recovered directly:

$$\begin{aligned} L(P_{\text{paper}}) &= 24, \\ L(P_{\text{Roe}}) &= 12, \\ L(P_*) &= 20. \end{aligned}$$

Roe's relation is therefore decisively shortest in this grammar. The hybrid improves the paper's word, but its main advantage is the simplicity of its collector identity rather than a dramatic reduction in raw operation count.

4 The maximal pro-2 quotient

In a finite 2-group quotient, $\tau = 1$ and ω_2-powering is the identity.

4.1 The paper and the hybrid

For the paper's presentation,

$$u_i = x_i, \quad d_0 = 1, \quad c_0 = 1, \quad g_0 = \sigma^2, \quad d_g = h_c = 1,$$

so (2) becomes

$$x_0^{\sigma^2} x_0 [x_1, \sigma] = 1. \tag{7}$$

For the hybrid,

$$u_i = x_i, \qquad d = 1, \qquad \sigma_2 = \sigma,$$

and (4) becomes the same relation (7).

The relation (7) admits the explicit Nielsen normalization

$$\sigma = s, \qquad x_1 = y, \qquad x_0 = s^{-2} a^{-1} \tag{8}$$

from the standard marked Demushkin presentation

$$\langle a, s, y \mid a^2 s^4 [s, y] = 1 \rangle_{\text{pro-2}}.$$

It follows immediately from

$$\nu_{\text{ur}}(a, s, y) = (-2, 1, 0)$$

that

$$\nu_{\text{ur}}(\sigma) = 1, \qquad \nu_{\text{ur}}(x_0) = \nu_{\text{ur}}(x_1) = 0.$$

Thus both the abstract pro-2 identification and its integral marking have score 1.

4.2 Roe's pro-2 quotient

Roe's relation specializes to

$$(x_0^{\sigma})^{-1} x_0^{-3} x_1^2 [x_1, x_1^{\sigma}] = 1. \tag{9}$$

Its degree-two initial form is nondegenerate and has the correct dyadic Demushkin type, so the abstract group can be recognized. The fully marked identification is less transparent. In the orientation calculation described in the accompanying verification of Roe's candidate, if

$$S = \chi(\sigma), \qquad X = \chi(x_0), \qquad Y = \chi(x_1),$$

then one obtains

$$Y = -X^2, \qquad X^3 + 2X^2 + 1 = 0, \qquad S = -\frac{X^3}{X^2 + X + 1}. \tag{10}$$

One must choose the unique root

$$X \equiv 5 \pmod{16},$$

verify that $S \equiv 13 \pmod{16}$, and then apply an orientation-preserving abelian correction. This is a genuine 2-adic normalization rather than a direct Nielsen substitution, so both the pro-2 and marking coordinates receive score 3.

5 Linear Fox calculus

At a tame lower map, write variations of (σ, τ, x_0, x_1) as (a, b, c, d). Let S, T be the lower actions of σ, τ, put

$$\mathsf{U} = \mathsf{S}^{\omega_2},$$

and let

$$P = 1 + \mathsf{T} + \cdots + \mathsf{T}^{e-1}$$

be the norm projector for the odd order e of T. The tame row is common to all three presentations:

$$L_{\mathrm{t}} = \mathsf{S}^{-1}(1 + \mathsf{T})a + (\mathsf{S}^{-1} + 1 + \mathsf{T})b.$$

5.1 The paper

The paper first proves that the long word h_0 has the elementary shadow of x_0^2, hence zero first derivative in characteristic 2. The remaining factors give

$$L_{\mathrm{w}} = Pb + (P + \mathsf{S}^{-1})d.$$

Thus

$$J_{\mathrm{paper}} = \begin{pmatrix} \mathsf{S}^{-1}(1+\mathsf{T}) & \mathsf{S}^{-1} + 1 + \mathsf{T} & 0 & 0 \\ 0 & P & 0 & P + \mathsf{S}^{-1} \end{pmatrix}. \tag{11}$$

The result is sparse, but reaching it requires the separate h_0-shadow lemma and the full finite-word ledger. We assign score 3.

5.2 Roe

For Roe's relation:

- $(x_0^\sigma)^{-1}$ contributes $\mathsf{S}^{-1}c$;
- $(x_0^{-3}\tau)^{\omega_2}$ contributes $P(c + b)$;
- x_1^2 and $[x_1, x_1^{\sigma_2}]$ have zero first derivative.

Therefore

$$J_{\mathrm{Roe}} = \begin{pmatrix} \mathsf{S}^{-1}(1+\mathsf{T}) & \mathsf{S}^{-1} + 1 + \mathsf{T} & 0 & 0 \\ 0 & P & P + \mathsf{S}^{-1} & 0 \end{pmatrix}. \tag{12}$$

This is (11) with the two wild columns interchanged. No auxiliary shadow lemma is needed, so the score is 1.

5.3 The hybrid

The hybrid retains the paper's remaining factors, but replaces its long h_0 by the block $y^g y d^2$. Its vanishing first derivative follows from a short case split:

- If $P = 0$, then $Dd = Dx_0 = c$, so $Dy = 0$, and d^2 has zero derivative.
- If $P = 1$, then $\mathsf{T} = 1$ and $g = \sigma_2^2$ acts trivially on the simple characteristic-2 quotient. The two y-terms cancel and d^2 again has zero derivative.

Consequently

$$J_* = J_{\text{paper}}.$$

The calculation requires a small structural observation but not the paper's collector ledger, so the score is 2.

6 Class-two and quadratic calculations

The distinction between the three presentations is sharpest at second order.

6.1 The hybrid square identity

The hybrid is governed by the following elementary lemma.

Lemma 1. *Let*

$$1 \longrightarrow Z \longrightarrow E \longrightarrow V \longrightarrow 1$$

be a central extension, and let ϕ be an automorphism of E which fixes Z pointwise. For $X, D \in E$, put

$$H_\phi(X, D) = (XD^{-1})^\phi (XD^{-1}) D^2.$$

Then:

(i) if X and D have the same image in V, then $H_\phi(X, D) = X^2$;

(ii) if $D \in Z$ and $\phi = 1$, then $H_\phi(X, D) = X^2$.

Proof. For (i), write $D = Xz$ with $z \in Z$. Then

$$XD^{-1} = z^{-1},$$

and therefore

$$H_\phi(X, D) = z^{-1} z^{-1} (Xz)^2 = X^2.$$

For (ii), centrality of D gives

$$(XD^{-1})^2 D^2 = X^2.$$

□

In the $P = 0$ case, the images of d and x_0 agree, so Lemma 1(i) applies. In the $P = 1$ case, d is central on the normalized representatives and σ_2^2 acts trivially on the full class-two coefficient group, so part (ii) applies. Thus the hybrid block has exactly the same Heisenberg and extraspecial coordinates as x_0^2.

6.2 The resulting forms

For the paper and the hybrid, normalized degree-one representatives are supported on x_0. The square block contributes

$$\lambda(c) \quad \text{or} \quad q(c),$$

while

$$[d_0, x_0^{\sigma_2}]$$

or its hybrid counterpart contributes the polar term. The normalized quadratic form is

$$Q^0_{\text{paper}}(c) = Q^0_*(c) = q(c) + b_q\big((P+1)c, \mathsf{U}^{-1}c\big). \tag{13}$$

Equivalently,

$$Q^0_{\text{paper}}(c) = Q^0_*(c) = \begin{cases} q(c), & P = 1, \\ q(c) + b_q(c, \mathsf{U}^{-1}c), & P = 0. \end{cases}$$

For Roe, normalized representatives are supported on x_1. The literal square and commutator give directly

$$Q^0_{\text{Roe}}(d) = q(d) + b_q(d, \mathsf{U}^{-1}d). \tag{14}$$

When $P = 1$, $\mathsf{U} = 1$ and $b_q(d, d) = 0$, so (14) agrees with (13). When $P = 0$, the two formulas are already identical after exchanging the wild coordinates.

Roe therefore receives quadratic score 1. The hybrid receives score 2, because Lemma 1 and the two cases must be invoked. The paper receives score 3, because the same conclusion is extracted from the longer h_0 collector and its full contribution ledger.

7 The integral marking

The integral marking is logically independent of the mod-2 cup form. Two one-relator pro-2 groups may have the same quadratic initial form while placing the unramified $\mathbb{Z}_2$-coordinate differently in their abelianizations.

For the paper and the hybrid, the exact reduction to (7) and the Nielsen transformation (8) settle the issue at once. Thus

$$C_{\text{mark}}(P_{\text{paper}}) = C_{\text{mark}}(P_*) = 1.$$

For Roe, the marking is not visible from (9). The orientation equations (10), the Hensel choice of X, and the final orientation-preserving correction are essential. Hence

$$C_{\text{mark}}(P_{\text{Roe}}) = 3.$$

8 Summary and Pareto comparison

The final table is

presentation	L	$C_{\text{pro-2}}$	C_{Fox}	C_{quad}	C_{mark}
paper	24	1	3	3	1
Roe	12	3	1	1	3
hybrid	20	1	2	2	1

(15)

In vector notation,

$$C(P_{\text{paper}}) = (24, 1, 3, 3, 1),$$
$$C(P_{\text{Roe}}) = (12, 3, 1, 1, 3),$$
$$C(P_*) = (20, 1, 2, 2, 1).$$

The hybrid strictly improves the paper's presentation in this complexity model: it has lower straight-line cost, the same direct marked pro-2 normalization, and shorter Fox and quadratic arguments. Roe and the hybrid do not dominate one another. Roe is shorter and has the most transparent low-order cohomological structure, while the hybrid has the simplest full integral marked normalization.

Thus the Pareto frontier consists of

$$\boxed{\text{Roe: shortest and cohomologically simplest}}$$

and

$$\boxed{\text{hybrid: simplest balanced marked proof}}.$$

If the primary goal is the shortest theorem statement or the smallest candidate-specific Fox calculation, Roe's presentation is preferable. If the goal is to minimize the complete proof while keeping the full unramified marking explicit, the hybrid is preferable.

9 Limitations of the comparison

The numerical values after L are ordinal assessments rather than intrinsic invariants. They depend on using the proof architecture of the main paper and on treating the integral marking as a separate load-bearing input. A new explicit Nielsen transformation for Roe's pro-2 relator would lower two of Roe's coordinates. Conversely, a different proof strategy might make the full finite word, rather than its first two jets, the dominant cost.

The straight-line count also depends on the chosen grammar. For example, charging an ω_2-power according to the cost of evaluating it in finite quotients would change all three length scores. The value of the present convention is not canonicity but reproducibility: every auxiliary word and every primitive profinite operation is charged exactly once.

References

[1] David Roe and David Turturean, *A Presentation of the Absolute Galois Group of* $\mathbb{Q}_2$, https://roed314.github.io/gq2/paper/paper.html.

[2] David Roe, *The Presentations*, https://roed314.github.io/gq2/presentations/.