ECC RESEARCH PROGRAM · P2.2

Relation Rank, Not Label Count

A relation-rank meta-reduction for q-SDH in structured representations

math.st [1] **@math__street** [1]

with GPT 5.6 Sol · Claude Fable 5

[1]Independent research collective, math.st · July 2026

ABSTRACT

A q-type assumption such as q-SDH hands the adversary a power ladder $g, g^x, \ldots, g^{x^q}$ whose length grows with a scheme's usage parameter. Existing meta-reductions rule out fully black-box reductions from prime-order q-SDH to fixed-size assumptions in generic representations, but do not directly cover one named concrete encoding with native label operations. We identify the missing invariant. At every adversary-call prefix, collect all efficiently available, group-verifiable affine relations among source-label discrete logarithms and define the **prefix affine trace dimension** as the dimension of their solution space. Our main theorem shows that a fully-black-box fixed-representation reduction yields an attack on its fixed-size assumption whenever this dimension is bounded by $\kappa < q - 1$. The proof online-reparametrizes every label by κ unknowns and imports only the column-space root-list lemma of Lu–Zhandry. The theorem strictly strengthens a raw fresh-label bound: arbitrarily many labels derived from k native seeds through certified affine operations cost only k dimensions. It also gives a precise conditional bridge to structured generic groups with available factor-base exponent vectors; label density alone remains insufficient. For target-valued bilinear games we derive the safe bound $\binom{\kappa+2}{2} + \tau < q$. We audit equality branches, collisions, late relations, rewinding, and nonlinear native operations, and validate the rank bookkeeping with exact finite-field tests. The residual fully-black-box class must cross dimension $q - 1$ at an actual call prefix or rely on native structure outside the certified affine trace; non-black-box code access remains separate.

Keywords. q-SDH · q-type assumptions · meta-reductions · relation rank · structured generic groups · black-box separations

1 Introduction

Pairing-based cryptography leans, in large parts of its deployed and standardized corpus, on assumption **families** parametrized by an integer q that grows with usage. The prototype is the strong Diffie–Hellman assumption of Boneh and Boyen [1]: in a group of prime order r, given the power ladder $g, g^x, g^{x^2}, \ldots, g^{x^q}$, it must be infeasible to output any pair $\left(c, g^{1/(x+c)}\right)$. The parameter is not cosmetic. In the Boneh–Boyen short-signature proof, q equals the adversary's signing-query budget [1]; in the BBS family it tracks the number of issued signatures, and recent concrete-security work shows the dependence is real in both directions [2], [3]; and Cheon's algorithms show that the assumption **quantitatively degrades** as the ladder lengthens, by roughly $\sqrt{d}$ bits of security

for favorable divisors d of $r \pm 1$ [4]. A q-type assumption is therefore an awkward foundation: a family whose strength deteriorates in exactly the parameter that deployments increase.

Problem P2.2 poses the natural structural question. **Either** exhibit a black-box reduction from q-SDH to a constant-size assumption — one whose instance size does not grow with q — **or** prove via meta-reduction that no such reduction exists. The required deliverables are a uniform-notation catalogue, an implication graph with cited edges and loss factors, a written account of where an attempted constant-size reduction breaks, and a precisely scoped meta-reduction statement, proved or stated as a target.

Our answer is a **relation-aware scoped impossibility**: it goes beyond the published generic-representation boundary while making the remaining representation-specific escape routes explicit.

Main finding (scoped). The strongest justified deliverable is a scoped impossibility, not a reduction and not an all-black-box impossibility. CITED On the published side, no fully black-box generic-representation reduction bases prime-order q-SDH on a true fixed-size assumption whose challenger emits n group elements, once $n < q - 1$ in a single group or $\binom{n+2}{2} < q$ with a bilinear map [5]. PROVED Our main theorem extends the separation to **one named concrete representation**. If the maximum affine dimension κ of the verified source-label trace at any q-SDH call satisfies $\kappa < q - 1$, a fully-black-box reduction yields an attack on its fixed-size premise. Derived labels cost no new dimension when their affine exponent relations are available and group-verifiable. PROVED What remains open is exactly the residual class: reductions whose prefix affine trace reaches $q - 1$, reductions whose useful native relations are nonlinear or unavailable, and reductions that use the q-SDH adversary's code non-black-box.

1.1 Contributions and honest scope

We contribute (i) the typed catalogue and audited implication graph with per-edge losses (§3–4, Figure 1); (ii) the catalogue of constant-size prime-order candidates and what each does **not** supply (§5); (iii) Cheon's quantitative obstruction restated in typed notation (§6); (iv) the exact break point of the direct SXDH embedding — an exponent-span argument that kills the attempt already at x^2 (§7); (v) an audit of the Lu–Zhandry generic-representation separation and a new boundary corollary: every **representation-uniform** fully black-box standard-oracle reduction is already a generic-representation reduction and inherits the published thresholds (§8); (vi) the prefix relation-rank meta-reduction, our main new theorem, with a strict fresh-label corollary, structured-operation corollaries, and bilinear safe overcount (§9, Figure 2, Figure 3); (vii) post-mortems of three dead extensions — random relabeling, efficient sparse encodings, and density-only transfer — each with its exact failure point (§10); and (viii) a final scope classification separating what is ruled out from what is open (§11–12).

We state the scope plainly, as the problem's ground rules demand. Nothing here reduces q-SDH to a constant-size prime-order assumption, and nothing here is an unrestricted black-box impossibility. The status of P2.2 is **partial**: complete for the catalogue, the published generic boundary, and the bounded-relation-rank representation-dependent class; open once the prefix rank reaches the threshold, for unmodelled nonlinear native structure, and for non-black-box use of the adversary.

2 Setting and typed notation

CITED Throughout, r is prime, $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ is a non-degenerate bilinear map on groups of order r, and g_1, g_2 are fixed generators. We write $[z]_1 = g_1^z$, $[z]_2 = g_2^z$, and $[z]_T = e(g_1, g_2)^z$ for $z \in \mathbb{Z}_r$ [1].

Definition 1 (typed power ladder). PROVED For $I \subset \mathbb{Z}_{\geq 0}$, the typed power ladder is

$$P_i(x; I) = \left(\left[x^j \right]_i \right)_{j \in I}, \tag{1}$$

where $\left[x^0 \right]_i = [1]_i$ is the named base. "Search" means output the named value; "decision" means distinguish the named value from an independent uniform element of the same typed group. The two tasks are never silently interchanged.

The typing discipline matters more than it may appear: the names DHE, strong DH, BDHE, and their index shifts are not canonical across the literature, so every claim below refers to a displayed tuple, never to a bare acronym [4], [5]. Near-name variants are not identified without tuple equality.

3 The assumption catalogue

The first deliverable is a uniform-notation catalogue (sub-goal SG-01, complete). Table 1 lists the inversion and strong-DH family; Table 2 lists the exponent and gap family. Each row records the exact instance and target in typed notation, with the epistemic tag under which the research log holds it.

Table 1: The inversion and strong-DH family in typed notation (SG-01).

Name	Exact instance and target
CITED q-SDH (source core)	Input $P_1(x; \{0, \ldots, q\})$; output any $(c, [1/(x+c)]_1)$ with $c \in \mathbb{Z}_r$, $x + c \neq 0$ [1], [5].
CITED q-SDH (typed Boneh–Boyen)	The source core plus $[1]_2, [x]_2$; the extra $\mathbb{G}_2$ terms permit public verification but do not alter the search relation [1].
CITED q-DHI	Input $P_1(x; \{0, \ldots, q\})$; output $[1/x]_1$ — the prescribed choice $c = 0$ of the q-SDH relation [1].
PROVED q-DDHI	Input $P_1(x; \{0, \ldots, q\})$ and $T \in \mathbb{G}_1$; decide whether $T = [1/x]_1$ or $T \leftarrow \mathbb{G}_1$ uniform.
CITED q-BDHI	Symmetric original: input $(g, g^x, \ldots, g^{x^q})$, output $e(g, g)^{1/x}$. Typed lift: input $P_1(x; \{0, \ldots, q\})$ and $[1]_2$; output $[1/x]_T$ [6], [7].
PROVED q-DBDHI	Decision version of the typed target: distinguish $T = [1/x]_T$ from uniform $T \leftarrow \mathbb{G}_T$.
CITED q-wBDHI	Input independent generators g_1, h_2 and $P_1(x; \{1, \ldots, q\})$; output $e(g_1, h_2)^{1/x}$ [7].
CITED q-wBDHI*	Same input as q-wBDHI; output $e(g_1, h_2)^{x^{q+1}}$ [7].

Table 2: The exponent and gap family in typed notation (SG-01). The gap tuple omits the power x^q that its target exponentiates.

Name	Exact instance and target
CITED q-DHE	Input $P_1(x; \{0, \ldots, q\})$; output $\left[x^{q+1} \right]_1$ [5].

Name	Exact instance and target
CITED q-DDHE	Input $P_1(x;\{0,\dots,q\})$ and $T \in \mathbb{G}_1$; decide whether $T = [x^{q+1}]_1$ or uniform. Also called q-strong DDH in some sources [5].
CITED q-BDHE (gap form)	Input independent g_1, h_2 and $[x^j]_1$ for $j \in \{0,\dots,q-1,q+1,\dots,2q\}$; output $e(g_1,h_2)^{x^q}$ [4], [7].
PROVED q-DBDHE	Decision version of the preceding gap target in $\mathbb{G}_T$.
CITED q-aBDH	Input $P_1(x;\{0,\dots,q\})$, $[1]_2, [x]_2$, an independent h_2, and $h_2^{x^{q+2}}$; output $e(g_1,h_2)^{x^{q+1}}$ [1].

4 The implication graph and its loss audit

PROVED An arrow $H(A) \to H(B)$ means "hardness of A implies hardness of B": operationally, a breaker for B is transformed into a breaker for A. Every edge in Figure 1 was audited against the primary reduction text — not a survey summary — and carries its transformation, call count, and loss in Table 3.

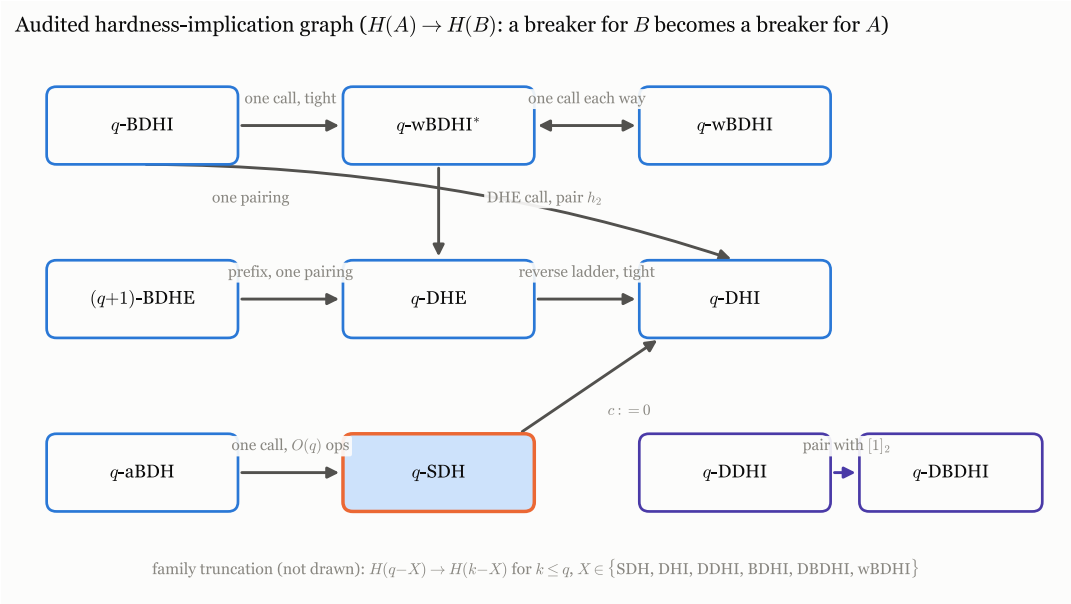

Figure 1: The audited implication graph (SG-02), drawn from the edge table below; schematic rendering of proved/cited reductions, not data. Solid computational family in blue; the decisional pair at lower right; q-SDH highlighted as the problem's anchor. Every edge is one oracle call and success-preserving. The truncation family $H(q\text{-}X) \to H(k\text{-}X)$, $k \le q$, applies to $X \in \{\text{SDH, DHI, DDHI, BDHI, DBDHI, wBDHI}\}$ but not to targets whose exponent depends on q (DHE, wBDHI*, gap-BDHE).

Table 3: Audited edge table (SG-02). ε denotes the invoked breaker's success probability or advantage.

Edge	Breaker transformation	Loss
CITED $H(q\text{-aBDH}) \to H(q\text{-SDH})$	Run the q-SDH breaker to get $(c, [1/(x+c)]_1)$. Divide $X^{q+2} - (-c)^{q+2}$ by $X + c$ to get $X^{q+1} + w(X)$ with $\deg w \le q$; pairing and the supplied $h_2^{x^{q+2}}$ isolate the wanted x^{q+1} term [1].	1 call; ε preserved; $O(q)$ ops

Edge	Breaker transformation	Loss
PROVED $H(q\text{-SDH}) \to H(q\text{-DHI})$	Feed the same ladder to the DHI breaker; return its output as the SDH answer with $c = 0$ [1].	1 call; ε; $O(1)$
PROVED $H(q\text{-BDHI}) \to H(q\text{-DHI})$	Pair the DHI output $[1/x]_1$ with $[1]_2$ [1], [6].	1 call; ε; 1 pairing
CITED $H(q\text{-BDHI}) \to H(q\text{-wBDHI}^*)$	From $(w, w^{\beta}, ..., w^{\beta^q})$, give the wBDHI* breaker base w^{β^q}, the reversed ladder, and $h = w^s$; raise its answer to $1/s$ [7].	1 call; ε; $O(q)$ ops
CITED $H(q\text{-wBDHI}^*) \leftrightarrow H(q\text{-wBDHI})$	Reverse the ladder with new base $g_1^{x^q}$ and hidden exponent $1/x$; the two targets transform into one another [7].	1 call each way; ε; $O(q)$
PROVED $H(q\text{-wBDHI}^*) \to H(q\text{-DHE})$	Run the DHE breaker on $(g_1, g_1^x, ..., g_1^{x^q})$ and pair its output $g_1^{x^{q+1}}$ with h_2 [4], [7].	1 call; ε; 1 pairing
PROVED $H(q\text{-DHE}) \to H(q\text{-DHI})$	To use a DHI breaker, set base $g_1^{x^q}$, hidden exponent $1/x$, and reverse the ladder; its output is $g_1^{x^{q+1}}$ [4].	1 call; ε; $O(q)$ reversal
PROVED $H((q+1)\text{-BDHE}) \to H(q\text{-DHE})$	On the gap tuple, run the q-DHE breaker on the prefix through x^q and pair $g_1^{x^{q+1}}$ with h_2; ignore higher powers [4].	1 call; ε; 1 pairing
PROVED $H(q\text{-DDHI}) \to H(q\text{-DBDHI})$	Pair every source challenge with $[1]_2$ and invoke the DBDHI distinguisher [6].	1 call; advantage preserved; 1 pairing
PROVED $H(q\text{-}X) \to H(k\text{-}X)$, $k \le q$	Delete the unused suffix of the ladder before calling the k-breaker; valid for $X \in \{\text{SDH}, \text{DHI}, \text{DDHI}, \text{BDHI}, \text{DBDHI}, \text{wBDHI}\}$ [1].	1 call; preserved; $O(q)$ handling

PROVED The truncation edge formalizes why the assumptions strengthen as q grows: hardness at the larger parameter implies hardness at the smaller, while the converse is not supplied by tuple truncation. CITED Boneh and Boyen further prove q-SDH random self-reducible by rescaling the hidden exponent and bases — one solver call, success preserved, $O(q)$ exponentiations [1].

5 Constant-size prime-order candidates

Which fixed-size assumptions could serve as a target? Table 4 catalogues the standard candidates (SG-03), separating source-group, cross-source, and target-group consequences.

Table 4: Constant-size prime-order candidates (SG-03). For every true fixed-size candidate here, the generic separation of §8 rules out the fully black-box generic reduction from q-SDH once q crosses the dimension threshold; this is **not** a claim that any of these assumptions is false [5].

Assumption	Fixed-size typed experiment	What it supplies — and does not
CITED DDH in $\mathbb{G}_i$	Distinguish $([1]_i, [a]_i, [b]_i, [ab]_i)$ from $([1]_i, [a]_i, [b]_i, [z]_i)$.	A decisional degree-two correlation in one source group; no source power ladder [5].
CITED XDH / SXDH	XDH: DDH holds in the designated source group despite the pairing; SXDH: DDH in both $\mathbb{G}_1$ and $\mathbb{G}_2$.	Constant-size decisional assumptions for asymmetric pairings; no map from $\mathbb{G}_T$ back to a source group [8], [9].

Assumption	Fixed-size typed experiment	What it supplies — and does not
CITED DLIN in $\mathbb{G}_i$	For independent u, v, h: distinguish $(u, v, h, u^a, v^b, h^{a+b})$ from the tuple with random last element.	A constant-size linear-relation challenge; still a fixed-dimensional source exponent span [10].
CITED co-CDH	Given $([1]_1, [a]_1, [1]_2)$, output $[a]_2$.	A constant-size cross-source search problem; pairing reveals $[a]_T$ but not $[a]_2$ in the typed interface [9], [11].
CITED DBDH	Given $([1]_1, [a]_1, [b]_1, [1]_2, [c]_2, T)$, distinguish $T = [abc]_T$ from uniform.	A constant-size target-group decision problem; it does not manufacture $[x^j]_1$ for unboundedly many j [6].

6 The quantitative obstruction: Cheon's exponent recovery

Even before any separation, the ladder itself leaks. The following two statements are Cheon's, restated in typed notation.

Proposition 2 (Cheon, $d \mid r-1$). CITED If a divisor $d \le q$ of $r-1$ is available, the ladder contains g^x and g^{x^d}, and x is recoverable in $O\left(\log r\left(\sqrt{(r-1)/d} + \sqrt{d}\right)\right)$ group operations [4].

Proposition 3 (Cheon, $d \mid r+1$). CITED If $d \mid r+1$ and the ladder reaches degree $2d$, then x is recoverable in $O\left(\log r\left(\sqrt{(r+1)/d} + d\right)\right)$ group operations [4].

Corollary 4. PROVED Recovering x immediately solves q-SDH: choose any $c \neq -x$ and compute $[1/(x+c)]_1$ by known-scalar exponentiation.

CITED These attacks reduce generic square-root security by roughly a factor $\sqrt{d}$ for favorable divisors, so concrete group sizes must account for both q and the factorization of $r \pm 1$ [4]. This is the quantitative face of the same phenomenon the rest of the paper treats structurally: the ladder is a high-dimensional correlated object, and nothing constant-size manufactures it.

7 Where a direct reduction dies: the exponent-span barrier

The problem asks for a written account of where an attempted constant-size-assumption reduction breaks. Attempt A001 tried the most natural route: use one source-group DDH challenge — hence one component of an SXDH challenge — to synthesize the ladder a q-SDH adversary expects, and convert its answer into a DDH decision. The reduction class tested is straight-line, type-preserving, and algebraic: it combines challenge elements by group operations and known-scalar exponentiation, invokes one q-SDH adversary, and has no source-group encoding of a target-group pairing value.

Proposition 5 (exponent-span barrier). PROVED Write the source DDH challenge exponents as $(1, a, b, z)$, with $z = ab$ in the real branch and independent z in the random branch. Every $\mathbb{G}_1$ element such a reduction creates before its oracle call has exponent in $L = \mathrm{span}_{\mathbb{Z}_r}\{1, a, b, z\}$. In the random branch, no nonconstant x satisfies both $x \in L$ and $x^2 \in L$; hence even the prefix $\left([1]_1, [x]_1, [x^2]_1\right)$ of a nontrivial q-SDH instance cannot be embedded.

Proof. In the random branch, treat a, b, z as algebraically independent. Put $x = A + Ba + Cb + Dz \in L$. If $x^2 \in L$, comparing the coefficients of $a^2, b^2, z^2, ab, az, bz$ — which do not occur in L

— forces, over an odd prime field, $B = C = D = 0$. So x is a known constant, statistically independent of the challenge bit, and an oracle query built from it cannot transfer any distinguishing advantage. Pairing does not repair the embedding: it produces degree-two exponents in $\mathbb{G}_T$ only, while q-SDH needs the missing powers in $\mathbb{G}_1$, and the typed interface has no operation $\mathbb{G}_T \to \mathbb{G}_1$. □

> **Remark 6.** PROVED This kills the stated straight-line algebraic class **at** x^2 — before the first oracle call — but is not by itself a proof against rewinding, adaptive multiple calls, or representation-dependent and non-algebraic reductions. Those classes are the subject of the next two sections.

The failure is dimensional, and it is the germ of everything that follows: a constant-size challenge exposes a fixed-dimensional algebraic span, while the q-SDH ladder demands $q + 1$ correlated powers.

8 The generic-representation boundary

8.1 The Lu–Zhandry separation, audited

The fixed-dimensional-span failure of A001 was turned, by Lu and Zhandry, into a separation for a substantially broader class [5]. Attempt A002 audited the primary proof bodies; we restate the result in the repository's terms. Methodologically the result descends from Coron's meta-reduction technique for signature-scheme optimality bounds [12], which is ancestry, not itself a q -SDH separation.

> **Definition 7** (GR-BB reduction). CITED A generic-representation black-box (GR-BB) reduction is polynomial-time, fully black-box in a possibly **inefficient** adversary whose generic group-operation queries it may observe and answer, and representation-independent: it receives concrete bit-string labels and may apply arbitrary bit gates to them, but its advantage is defined as the infimum over all — possibly inefficient — implementations of the group. The formulation permits interactive assumptions and adaptive oracle calls [5].

> **Definition 8** (fixed-size and q-type assumptions). CITED A **fixed-size** assumption may be interactive and may contain non-group auxiliary data, but across the experiment its challenger outputs at most n group elements. A ***q*-type** assumption requires that a transcript expose at least q linearly independent bounded-degree polynomial functions tied to a hidden x_0, such that knowing x_0 lets an efficient algorithm break the assumption with noticeable probability [5].

> **Theorem 9** (Lu–Zhandry). CITED In a prime-order generic group, if a GR-BB reduction maps such a q-type assumption to a fixed-size assumption of size $n < q - 1$, then a generic polynomial-time algorithm breaks the fixed-size assumption. In the generic bilinear extension the same conclusion holds when $\binom{n+2}{2} < q$ [5].

Proof. (Audited skeleton; the proof is Lu–Zhandry's.) The meta-reduction selects an inefficient perfect q-adversary that brute-forces the hidden exponents. It traces every reduction-created source element as a known linear combination of the n fixed challenge elements, so the exponent vectors submitted to the adversary lie in a subspace of dimension at most $n + 1$. A Vandermonde rank argument and finite-field root finding leave only polynomially many candidates for x_0; generic equality tests identify the valid ones, so the otherwise inefficient adversary is simulated efficiently,

and the reduction with its oracle replaced by this simulation is a generic attack on the fixed-size assumption. With a pairing, recorded exponents may be quadratic in the n input exponents; linearizing in all degree-at-most-two monomials gives dimension $\binom{n+2}{2}$ and the bilinear threshold [5]. □

Corollary 10 (q-SDH instantiation). CITED q-SDH meets the q-type definition: its ladder supplies the required independent polynomials, and knowledge of x yields a valid answer by selecting any $c \neq -x$ after a consistency check. Hence no generic reduction from q-SDH to a true fixed-size assumption exists at the stated thresholds [5].

Remark 11 (why composite order escapes). CITED The prime-order proof step fails in unknown-composite-order groups: finding the required polynomial roots can be as hard as factoring, and independently rerandomizable hidden subgroups supply the "shadow" dimensions that the positive Déjà Q reductions exploit — computational q-SDH and broad polynomial q-type classes **do** reduce to constant-size subgroup hiding in composite-order pairing groups, with $\Theta(q)$ hybrid loss in the original framework and logarithmic loss in the follow-up [5], [13], [14]. The positive and negative results are consistent; the barrier is specifically prime-order.

CITED The theorem's tracing argument also covers algebraic reductions whose outputs carry algebraic explanations, but it does not rule out every black-box reduction in every formalization: representation-specific standard-model reductions, and reductions non-black-box in the q-adversary, are outside its stated class [5]. Section 3.5 of the paper likewise leaves open derived groups whose law depends on auxiliary bit strings, equality branches, or concrete encodings [5]. These gaps are exactly where the remainder of this paper operates.

8.2 Representation-uniform reductions are already generic (A003)

One might hope that a **standard-oracle** reduction — a single machine that may inspect encodings as bit strings — escapes the GR-BB class. It does not, as long as its guarantee is uniform over representations.

Definition 12 (UR-FBB reduction). PROVED For single-stage prime-order group games F and Q, a reduction is **universally representation-uniform fully black-box** (UR-FBB) if a single PPT oracle machine R (i) treats the Q-adversary as an oracle, though it may run multiple copies and rewind; (ii) uses the group through labeling, group-operation, equality, and optional pairing interfaces, and otherwise performs arbitrary bit computation on labels; and (iii) satisfies pointwise representation uniformity: for **every** admissible implementation G of the group, including possibly inefficient ones, and every possibly inefficient adversary A, if $\mathrm{Adv}_{Q^G}(A) \geq \varepsilon$ for non-negligible ε, then $\mathrm{Adv}_{F^G}\left(R^{A,G}\right) \geq \delta$ for some non-negligible δ independent of G and A. Conditions (i)–(ii) match the classical fully-black-box interface [15].

Theorem 13 (boundary theorem). PROVED Every UR-FBB reduction is a GR-BB reduction. Consequently no UR-FBB reduction bases prime-order q-SDH on a true fixed-size assumption of size n when $n < q - 1$, nor, with a generic bilinear map, when $\binom{n+2}{2} < q$.

Proof. A GR adversary with non-negligible infimum advantage breaks Q^G for every implementation G with a common non-negligible lower bound ε. Applying uniformity pointwise gives one non-negligible δ with

$$\inf_{G} \mathrm{Adv}_{F^G}\left(R^{A,G}\right) \geq \delta. \tag{2}$$

Conditions (i)–(ii) make R a GR oracle circuit, so this is a GR-BB reduction — indeed stronger than required, since the source definition may assign a reduction per adversary. The thresholds then follow from Theorem 9 and its q-SDH instantiation. For single-stage games the statement transfers to the type-safe model, since TS-BB and GR-BB reductions exist if and only if one another [5]. Label inspection does not invalidate the wrapper: for label length $\ell > \log_2 r + \omega(\log \lambda)$, a polynomially bounded interaction guesses a valid unseen label only with negligible probability [5]. □

Remark 14. PROVED The qualifier **universal** is load-bearing: it includes possibly inefficient random sparse implementations. A guarantee promised only over efficiently computable or standardized representations need not include them and is **not** promoted to GR-BB by this argument. That weaker quantifier is precisely the residual case attacked next.

9 Main theorem: prefix affine relation rank

The remaining gap after §8 is a reduction promised for **one named concrete representation** G_* — an elliptic-curve group with a standardized encoding, say — where the machine may branch on encoding bits and invoke native validation, decompression, hash-to-group, tables, or structured label operations. Random relabeling cannot reach this class (§10.1). Merely counting native labels is also too crude: a public operation may generate millions of labels while exposing an exact algebraic relation for every output. The correct resource is the number of relation-independent logarithmic degrees of freedom present when the adversary is actually called.

9.1 Certified relations and the online invariant

Definition 15 (certified affine source relation). PROVED Fix a concrete cyclic source group $G_* = \langle g \rangle$ of prime order r. At an execution prefix τ, enumerate the distinct typed source labels $L_1, \ldots, L_m$ and write $z_i = \log_g L_i$. A **certified affine relation** is a public row

$$a_1 z_1 + \ldots + a_m z_m = b \bmod r \tag{3}$$

with known coefficients whose validity is checked by the group identity

$$\prod_{i=1}^{m} L_i^{a_i} = g^b. \tag{4}$$

The verified rows form a consistent affine system $A_\tau z = b_\tau$ over $\mathbb{F}_r$.

Definition 16 (prefix affine trace dimension). PROVED The **affine trace dimension** at prefix τ is

$$d(\tau) = \dim_{\mathbb{F}_r}\{z \in \mathbb{F}_r^m : A_\tau z = b_\tau\} = m - \mathrm{rank}_{\mathbb{F}_r}(A_\tau). \tag{5}$$

The execution-wise relation-rank budget is the maximum of $d(\tau)$ over prefixes at which the reduction invokes its q-type adversary. Taking the maximum over every prefix is a safe, possibly looser, sufficient bound.

PROVED The definition is operational. A raw native label adds a variable and no row, increasing d by one. A group-operation output, known-scalar map, or public structured operation whose exponent is an affine combination of earlier exponents adds a variable and a row with nonzero

coefficient on that variable, leaving d unchanged. A true equality or collision adds a row and can only decrease d; a false equality provides only a disequality and is not used for compression. All rows used by the meta-reduction are therefore efficiently checkable without computing a discrete logarithm.

Lemma 17 (online affine reparametrization). PROVED At every prefix with $d(\tau) = d$, polynomial-time Gaussian elimination computes a particular solution z^0 and a nullspace basis $u^1, \ldots, u^d$ such that every consistent logarithm vector is

$$z = z^0 + \theta_1 u^1 + \ldots + \theta_d u^d. \tag{6}$$

Consequently, every current source label has a known affine coefficient vector in $\mathbb{F}_r^{d+1}$, although neither the parameters θ_i nor any individual discrete logarithm is known.

Proof. Row-reduce the augmented matrix $[A_\tau \texttt{ mid } b_\tau]$. The actual logarithm vector witnesses consistency. Set all free variables to zero for z^0 and once to one for each nullspace vector. The number of field operations is polynomial in the transcript size. After a label or relation event, recompute the presentation; on rewinding, restore the matrix and reduction state from the same branch snapshot. No later relation is used at an earlier prefix. □

9.2 Fixed-representation separation

Theorem 18 (bounded prefix relation-rank separation). PROVED Let R be a PPT fully-black-box reduction from a fixed-size assumption F^{G_*} to q-SDH over the fixed prime-order representation G_*. Suppose that every source label is entered into an online typed trace, every relation used for compression is efficiently available and certified as above, and every prefix at which R calls its q-SDH adversary satisfies

$$d(\tau) \le \kappa < q - 1. \tag{7}$$

Then R yields a PPT attack on F^{G_*} with its success probability against a selected perfect q-SDH adversary. Consequently, if F^{G_*} is hard, no such reduction exists.

Proof. The meta-reduction runs R on the real F^{G_*} challenge and executes the same concrete representation code on the same bit strings. At a q-SDH call prefix, the online affine reparametrization lemma represents every queried source label by a known row in $\mathbb{F}_r^{d+1}$. Hence the matrix M of the q requested power functions has at most $d + 1 \le \kappa + 1 < q$ columns.

The low-dimensional root-list argument of Lu–Zhandry uses only the column space of M, not the logarithms assigned to its coordinates [5]. Their Lemma 5.1 constructs a polynomial-size list containing the hidden exponent of every valid nonidentity ladder, and Lemma 5.3 checks the candidates and returns a response consistent with a fixed perfect q-SDH adversary. The affine constant is simply the extra column. The unknown parameters θ_i never have to be evaluated.

After each new certified relation the simulator reparametrizes the current trace. Equality branches and actual collisions preserve consistency because the real logarithm vector satisfies the added row; a false equality adds no row. Multiple calls and rewinding are branch-local, so the presentation used at each call is exactly the one available at that prefix. The root-list simulator, field elimination, and real group checks are polynomial time and reproduce the selected perfect adversary's transcript. Thus the reduction's success probability becomes the meta-reduction's success probability against F^{G_*}. □

Remark 19. PROVED Prefix sensitivity is load-bearing. A relation learned after a q-SDH call may lower the final dimension but cannot retroactively lower the dimension governing that earlier call. The theorem therefore bounds a maximum over actual call prefixes, not a final matrix rank.

Figure 2 shows the quotient-trace architecture. Relative to the published proof, the only new step is the online affine reparametrization; the imported root-list lemma receives fewer columns, not a different query distribution [5].

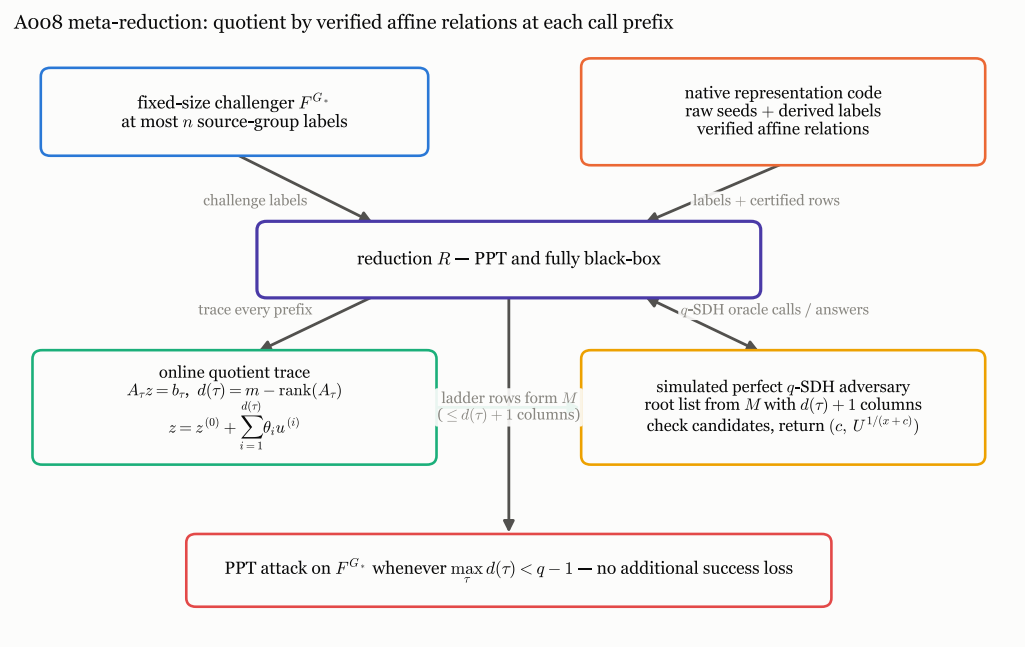

Figure 2: Architecture of the relation-rank meta-reduction (A008); a proof schematic, not measured data. Verified rows $A_\tau z = b_\tau$ quotient all source labels to $d(\tau)$ unknown parameters plus one affine constant. At an adversary-call prefix, the Lu–Zhandry root-list simulator applies whenever $d(\tau) + 1 < q$.

9.3 Strict corollaries: raw labels and structured operations

Corollary 20 (bounded fresh labels, recovered). PROVED If a fixed-assumption challenger supplies at most n source labels and the native representation introduces at most s unexplained source labels, then $d(\tau) \le n + s$. Therefore the bounded-fresh-label separation of A006 follows from the bounded prefix relation-rank theorem whenever $n + s < q - 1$.

Proof. Give each challenger or unexplained label a free coordinate. Every recorded group operation contributes a dependent label. This is exactly the empty- relation upper bound underlying A006. The relation-rank theorem is strictly sharper whenever nontrivial certified rows are available. □

Corollary 21 (structured derivation DAG). PROVED Suppose the fixed-assumption challenge contributes at most n relation-independent source coordinates and the native interface contributes k raw source seeds. If every other native label is produced by a recorded group operation, known-scalar map, or public operation whose exponent is a certified affine combination of earlier exponents, then $d(\tau) \le n + k$ at every prefix, independently of the total number of derived labels. The separation holds when $n + k < q - 1$.

Proof. Each derived label adds one variable together with an independent defining row whose new-label coefficient is nonzero. Variable count and relation rank both rise by one. A collision may add another row and can only lower the dimension. Unique factorization is not required when the derivation row itself is recorded. □

Corollary 22 (available factor-base presentation). PROVED If a structured representation exposes a factor base of size k and an efficiently available, verifiable exponent vector for every source label relative to that base, then its native contribution to $d(\tau)$ is at most k. Thus a challenge contributing n dimensions is separated whenever $n + k < q - 1$.

CITED This corollary identifies a concrete bridge to the structured generic-group literature. Its partial operation agrees with the hidden group operation wherever defined, and its index-calculus algorithm explicitly records factor-base exponent vectors and linear relations [16]. PROVED The corollary is conditional on those vectors being available to this transcript; the model's density parameter by itself does not provide them for every adaptively chosen label.

9.4 Bilinear lift and exact boundary

Proposition 23 (source-valued pairing case). PROVED For typed q-SDH whose oracle instance and answer lie in one source group, target-group pairing outputs do not enlarge the relevant source trace. The sharp sufficient threshold remains $\kappa < q - 1$, provided there is no unrecorded target-to-source conversion.

Proposition 24 (broader bilinear games). PROVED For a bilinear q-type game whose independent polynomial family may occur in the target group, a relation-aware source dimension κ spans at most the degree-at-most-two monomial space of dimension $\binom{\kappa+2}{2}$. With τ additional relation-independent target coordinates, replaying the bilinear root-list theorem gives the safe separation threshold

$$\binom{\kappa + 2}{2} + \tau < q. \tag{8}$$

This deliberately overcounts monomials forbidden by typing and is a sufficient, not necessarily tight, condition.

Corollary 25 (what a surviving reduction must do). PROVED Contrapositively, a fully-black-box fixed-representation reduction consistent with hardness of F^{G_*} must have some q-SDH-call prefix with $d(\tau) \geq q - 1$, use native source structure not captured by certified affine relations, or leave the fully-black-box class.

PROVED Nonlinear relations such as $z_{\text{out}} = z_{\text{in}}^2$ are outside the bounded prefix relation-rank theorem: one must either count the output as a new degree of freedom or lift the transcript to an explicitly bounded monomial space. An unverifiable relation is ignored. These are declared boundaries, not hidden assumptions.

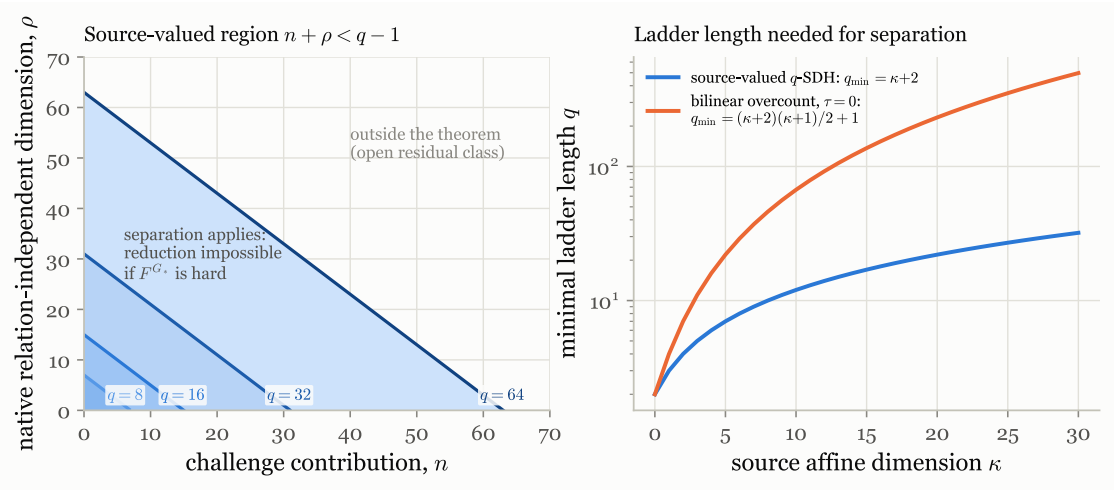

Figure 3: Parameter regions of the relation-rank theorem; direct plots of its inequalities, not measured data. Left: the source-valued separation region $n + \rho < q - 1$, where ρ is the native relation-independent dimension. Right: the minimal ladder length for the linear source theorem versus the conservative degree-two bilinear lift with $\tau = 0$.

EMPIRICAL: exact finite-field audit, 2026-07-25 Exact Gaussian-elimination fixtures over $\mathbb{F}_{101}$ validate the bookkeeping: three challenge coordinates and two native seeds remain five-dimensional after 24 certified derived labels; one late raw label raises the prefix maximum to six; its later defining relation lowers only the current dimension to five; and a collision row lowers it to four. The independent-label control has dimension 26. The scripts, tests, and JSON certificate are identified in the reproducibility statement.

10 Dead ends and their exact failure points

Three attempted extensions failed. Each post-mortem narrows the open residual class, so we record them as results, not anecdotes.

10.1 Random relabeling cannot settle a fixed representation (A004)

Proposition 26 (the representation quantifier). PROVED A reduction guaranteed only for one family G_* has quantifier form

$$\exists G_* \; \exists R_* \; \forall A_* : \quad \mathrm{Break}_{Q^{G_*}}(A_*) \Longrightarrow \mathrm{Break}_{F^{G_*}}\left(R_*^{A_*}\right). \tag{9}$$

The random-relabeling simulation needs this implication for a freshly sampled random sparse implementation G_L, which the premise does not entail — and cannot be made to entail. Given any R_* working for a publicly specified G_*, modify it to check the canonical encoding of the public generator (or a public representation identifier), abort on mismatch, and otherwise run R_*: the modified machine behaves identically on G_* and can fail on every random relabeling. Abstract group isomorphism does not transfer the behavior of a machine that reads encodings.

Proposition 27 (the trace invariant fails syntactically). PROVED Independently, concrete encodings break the fixed-dimensional trace at the syntax level: native validation, decompression, or hash-to-curve code may produce a valid point without an observed group-operation query, so the simulator receives no algebraic explanation of that point in the span of the fixed-size challenge.

CITED The published structural discussion likewise leaves open derived group laws using auxiliary strings, equality branches, or concrete representations; the classification covers natural affine algebraic constructions only [5]. HEURISTIC Fresh representation-derived points do not visibly generate the correlated ladder $g^x, \ldots, g^{x^q}$, so this trace failure is not positive evidence that a useful reduction exists; it only prevents the relabeling route from certifying impossibility. A006 repaired this partially by charging every unexplained point; A008 is sharper, charging only the logarithmic degrees of freedom that remain after all available certified affine relations.

10.2 Efficient sparse encodings cannot replace random injections (A005)

Could the GR wrapper's information-theoretic random injection be replaced by an **efficiently computable** sparse encoding, extending the separation to reductions promised over all efficient representations? No:

> **Theorem 28** (finite-seed counting obstruction). PROVED Let an efficient encoding family use an $m(\lambda)$-bit seed, m polynomial, to choose injections from an r-element group into an N-element label space, $N \geq r$. A possibly inefficient adversary distinguishes the family from a uniform random injection with advantage at least $1 - 2^m/(N)_r$, where $(N)_r = N(N-1)\cdots(N-r+1)$ — an overwhelming bound, since $\log (N)_r \geq \log(r!) = \Omega(r \log r)$ while m is polynomial and r is exponential in λ.

Proof. The family's support contains at most 2^m injective tables; the uniform distribution has support $(N)_r$. The distinguisher queries the encoding on all r inputs, reconstructs the full table, enumerates all 2^m seeds, and accepts exactly when the table belongs to the family: it accepts a family member with probability one and a random injection with probability at most $2^m/(N)_r$.□

PROVED Neither proposed weakening survives. A t-wise independent encoding matches only a bounded-query view, and the fully-black-box adversary quantifier supplies no query bound — the distinguisher above uses $r > t$ queries. A PRP protects only computationally bounded distinguishers, so it buys a computational premise, not the needed information-theoretic wrapper. CITED This boundary accords with Zhandry's generic-query convention, in which bit computation is unbounded and no efficient keyed permutation internal to the type-safe model is a secure PRP [17], and with the later sparse-GGM separations for admissibly encoded, dense, elliptic-curve, and bilinear settings [18].

PROVED A second, independent obstruction: breaking a static assumption in an artificial efficient encoding would refute hardness only **for that encoding**. If the assumption is asserted for a named curve family G_*, an attack on a different artificial family does not contradict its premise. PROVED What does transfer is the query-bounded version: if the reduction class quantifies only over adversaries making at most t encoding or group queries, a t-wise or lazy-table simulation becomes plausible — a strictly narrower class that must state its query bound explicitly.

10.3 Density does not bound rank; relations can (A007–A008)

The structured generic-group model gives algorithms free access to a partial label operation $\star$ agreeing with the hidden group law, and charges only generic group-oracle queries. CITED For a prime-order-r group whose constrained-label fraction is δ, its hard labeling distribution bounds T -query discrete-log advantage by

$$\delta(3T+2) + (3T+1)^2/r + 1/r \tag{10}$$

[16]. Could density δ alone bound the prefix affine trace?

Proposition 29 (density is not a label count). PROVED The δT term bounds a **random-hybrid probability** — the chance a random challenge interaction touches the structured region — whereas the trace dimension concerns labels **adaptively selected** along an actual execution and the rank of their available relations. A publicly recognizable structured subset of size $\delta r \geq q$ lets a reduction deliberately address q distinct relation-independent labels even when δ is negligible. The query parameter T cannot substitute either: $\star$ evaluations and all other label computation are free in the model's cost measure, so T does not bound the labels enumerated, parsed, or produced through $\star$ and then used as group elements.

PROVED Instantiating the model's hard labeling distribution instead repeats A004′s quantifier error: the theorem is existential in a labeling distribution, and a reduction promised for one named labeling need not retain its guarantee there [16].

PROVED The bounded prefix relation-rank theorem supplies the positive statement that A007 lacked. If each $\star$ output comes with the verified affine exponent relation induced by its inputs, arbitrarily many such outputs add no dimension beyond the raw structured seeds. More generally, available factor-base exponent vectors bound the native contribution by the factor-base rank. Thus density and relation rank are distinct resources: density alone does not bound the trace, while explicit relations can bound it without any label-count bound.

11 Scope classification and scheme evidence

Table 5 is the final classification: which reduction classes are ruled out, by what, and what survives.

Table 5: Final scope classification. The strongest justified deliverable is the relation-aware scoped impossibility of rows 1–4.

Reduction class	Result
CITED TS-BB / GR-BB, fully black-box in possibly inefficient adversaries	Ruled out at the published single-group ($n < q - 1$) and bilinear ($\binom{n+2}{2} < q$) thresholds [5].
CITED Algebraic reduction with explanations	The same tracing separation applies [5].
PROVED UR-FBB standard-oracle reduction, universal over representations	Implies GR-BB (Theorem 13, A003); ruled out at the same thresholds.
PROVED One fixed concrete representation, prefix affine source dimension at most κ	Direct quotient-trace separation at $\kappa < q - 1$ (A008), without random relabeling. This includes $n_1 + s_1$ fresh-label and $n + k$ structured-seed bounds.

Reduction class	Result
PROVED Representation class not closed under random sparse relabeling; prefix affine rank reaches $q-1$ or relations are not available and affine	A004 blocks relabeling and the relation-rank theorem no longer applies; existence or impossibility remains open.
CITED Non-black-box use of the adversary's code	Outside the fully-black-box taxonomy and outside every theorem here [5], [15].

Scheme-level evidence reinforces the same picture from the outside: Table 6 summarizes the positive results and why none crosses the prime-order barrier.

Table 6: Positive results and scheme taxonomy (NOTES §6). Current BBS-family results retain q-dependent premises in both directions.

Setting	What is established	Why it does not close P2.2
CITED Déjà Q, 2014	Computational q-SDH and broad polynomial q-type classes reduce to constant-size subgroup hiding in composite-order pairings; the q-SDH proof uses $q+2$ shadow copies with $\Theta(q)$ hybrid loss [13].	Not prime-order; asymmetric, computational-source only.
CITED Déjà Q follow-up, 2016	Broader classes, symmetric and asymmetric composite order, logarithmic tightness loss [14].	Still does not cross the prime-order barrier.
CITED Boneh–Boyen signatures	Standard-model proof under q-SDH with q the signing-query bound, roughly factor-two advantage loss [1].	A use of q-SDH, not a reduction of it.
CITED Dual-form exponent inversion	Modified Boneh–Boyen-like signatures and Gentry-IBE variants proved under SXDH [8].	Altered dual-form schemes — not the original assumption or scheme.
CITED BBS/BBS+ concrete security, 2025	After q signatures, attacks recover the key at $\Theta(q)$-DL cost; scheme security **implies** $\Theta(q)$-SDH [2].	A reverse, scheme-to-q-SDH comparison; reinforces usage-dependent security.
CITED BBS tightness, 2026	Tight q-SDH proof when each message is signed once; with repeats, no tight algebraic reduction to the considered q-SDH variants [3].	Tightens a q-type proof; neither removes q nor adds a static basis.

EMPIRICAL: bounded primary-source search, 2026-07-25 Primary-title, author, proceedings, and ePrint-oriented searches found the Lu–Zhandry column-space root-list proof and the structured-GGM

factor-base relation machinery, but no primary paper stating the prefix relation-rank meta-reduction proved in A008. This is evidence for a candidate contribution, not an unconditional priority claim. The same search found no unrestricted prime-order reduction of q-SDH to a fixed-size assumption or impossibility theorem covering every representation-specific black-box reduction.

12 Open problems

The residual gap is logged as P2.2/Q003 and now has three sharply separated branches.

PROVED **Branch 1 — threshold-crossing affine rank.** No result here rules out a representation-dependent reduction with an actual adversary-call prefix satisfying $d(\tau) \geq q - 1$. Closure requires either an upper bound on prefix rank for a precisely named representation class or a positive constant-size reduction that crosses the threshold. A large final label set is not enough: the independent dimension must be present at the call prefix.

PROVED **Branch 2 — nonlinear or unavailable native structure.** A concrete encoding may expose useful relations that are nonlinear in source logarithms, or whose coefficients cannot be efficiently obtained and verified by group identities. Closure requires a bounded-degree extension of the trace, a representation-specific extraction theorem, or a construction that exploits this structure.

PROVED **Branch 3 — code access.** A reduction that reads the code of the q-SDH adversary is outside the fully-black-box taxonomy [15] and outside every theorem above. Closure requires an explicitly non-black-box construction or a meta-reduction whose scope includes code access.

PROVED **Do not repeat.** Random relabeling (A004), finite-seed substitution (A005), and density-to-rank transfer (A007) each fail for the recorded, distinct reasons; the structural conclusion of the whole record is that the common obstruction is **dimensional** — a fixed-size prime-order challenge exposes a fixed-dimensional algebraic span, bilinearity raises the degree but leaves an $O(n^2)$ monomial space, and only composite-order hidden subgroups have so far manufactured the missing dimensions [5], [13]. The relevant resource for any future attack on Branch 1 is the prefix dimension after quotienting by all available verified relations, not label density, efficiency, raw label count, or an expectation over executions.

13 Conclusion

P2.2 asked for a reduction or a separation. The answer supported by the record is a **relation-aware layered separation**. The catalogue and audited implication graph fix the objects (§3–5); Cheon's attacks quantify the ladder's intrinsic leak (§6); and the direct SXDH embedding already dies at x^2 (§7). Lu–Zhandry turn the dimension barrier into a generic- representation impossibility, and representation-uniform standard-oracle reductions inherit it (§8). The bounded prefix relation-rank theorem then crosses the fixed-representation boundary: after quotienting all source labels by the verified affine relations available at an actual call prefix, fewer than $q - 1$ degrees of freedom suffice for the same root-list meta-reduction. This strictly replaces raw label count by relation rank, covers unbounded affine derivation DAGs over a bounded seed set, and isolates the exact conditional bridge to factor-base structured representations (§9–10).

The remaining doors are correspondingly narrower: reach rank $q - 1$ at a call prefix, exploit native structure outside an efficiently certified affine trace, or use the adversary's code non-black-

box. The result is not an unrestricted separation, but it identifies a new invariant, proves the strongest theorem supported by it, and leaves falsifiable targets rather than an undifferentiated representation gap.

Reproducibility

Exact relation-rank bookkeeping is implemented in `problems/P2.2-q-type-assumptions/code/relation_rank.py`; the deterministic audit driver is `code/audit_relation_rank.py`, and six known-answer tests are under `code/tests/`. The reported full fixture uses $\mathbb{F}_{101}$, seed 2208, and 24 derived labels; its machine-readable certificate is `data/audit_relation_rank_p101_g24_s2208_20260725.json`. It validates the linear-algebra accounting, not the cryptographic hardness theorem.

The three figures are schematics generated by `papers/figures/P2.2/make.py` — the implication graph from the audited edge table, the relation-rank quotient architecture from A008, and direct plots of $n + \rho < q - 1$ and $\binom{\kappa+2}{2} + \tau < q$; none depicts measured data. Primary proof bodies checked during the research include four implication reductions, Lu–Zhandry Lemmas 5.1 and 5.3 and Theorems 5.2 and 5.10, the Reingold–Trevisan–Vadhan fully-black-box definition, Zhandry's labeling model, and the structured-GGM definitions and factor-base relation algorithm. Every mathematical claim above carries one of the epistemic tags PROVED, CITED, HEURISTIC, EMPIRICAL: scope stated exactly as in the research log; untagged sentences are exposition, not claims.

References

[1] D. Boneh and X. Boyen, "Short Signatures Without Random Oracles and the SDH Assumption in Bilinear Groups," *Journal of Cryptology*, vol. 21, no. 2, pp. 149–177, 2008.

[2] R. Chairattana-Apirom and S. Tessaro, "On the Concrete Security of BBS/BBS+ Signatures," in *ASIACRYPT 2025*, 2025.

[3] R. Chairattana-Apirom, D. Hofheinz, and S. Tessaro, "Tight Security for BBS Signatures," in *EUROCRYPT 2026*, 2026.

[4] J. H. Cheon, "Security Analysis of the Strong Diffie–Hellman Problem," in *EUROCRYPT 2006*, in LNCS, vol. 4004. 2006, pp. 1–11.

[5] G. Lu and M. Zhandry, "Limits on the Power of Prime-Order Groups: Separating Q-Type from Static Assumptions," in *CRYPTO 2024, Part V*, in LNCS, vol. 14924. 2024, pp. 46–74.

[6] D. Boneh and X. Boyen, "Efficient Selective-ID Secure Identity-Based Encryption Without Random Oracles," in *EUROCRYPT 2004*, in LNCS, vol. 3027. 2004, pp. 223–238.

[7] D. Boneh, X. Boyen, and E.-J. Goh, "Hierarchical Identity Based Encryption with Constant Size Ciphertext," in *EUROCRYPT 2005*, in LNCS, vol. 3494. 2005, pp. 440–456.

[8] T. H. Yuen, S. S. M. Chow, H. Wu, C. Zhang, and S.-M. Yiu, "Exponent-Inversion P-Signatures and Accountable Identity-Based Encryption from SXDH," *IACR Communications in Cryptology*, vol. 1, no. 3, 2024.

[9] S. D. Galbraith, K. G. Paterson, and N. P. Smart, "Pairings for Cryptographers," *Discrete Applied Mathematics*, vol. 156, no. 16, pp. 3113–3121, 2008.

[10] D. Boneh, X. Boyen, and H. Shacham, "Short Group Signatures," in *CRYPTO 2004*, in LNCS, vol. 3152. 2004, pp. 41–55.

[11] D. Boneh, C. Gentry, B. Lynn, and H. Shacham, “Aggregate and Verifiably Encrypted Signatures from Bilinear Maps,” in *EUROCRYPT 2003*, in LNCS, vol. 2656. 2003, pp. 416–432.

[12] J.-S. Coron, “Optimal Security Proofs for PSS and Other Signature Schemes,” in *EUROCRYPT 2002*, in LNCS, vol. 2332. 2002, pp. 272–287.

[13] M. Chase and S. Meiklejohn, “Déjà Q: Using Dual Systems to Revisit q-Type Assumptions,” in *EUROCRYPT 2014*, in LNCS, vol. 8441. 2014, pp. 622–639.

[14] M. Chase, M. Maller, and S. Meiklejohn, “Déjà Q All Over Again: Tighter and Broader Reductions of q-Type Assumptions,” in *ASIACRYPT 2016, Part II*, in LNCS, vol. 10032. 2016, pp. 655–681.

[15] O. Reingold, L. Trevisan, and S. Vadhan, “Notions of Reducibility between Cryptographic Primitives,” in *TCC 2004*, in LNCS, vol. 2951. 2004, pp. 1–20.

[16] H. Corrigan-Gibbs, A. Henzinger, and D. J. Wu, “The Structured Generic-Group Model,” in *EUROCRYPT 2026*, 2026.

[17] M. Zhandry, “To Label, or Not To Label (in Generic Groups),” in *CRYPTO 2022*, 2022.

[18] T. Wang *et al.*, “Attention is Still What You Need: Another Round of Exploring Shoup's GGM,” in *ASIACRYPT 2025*, 2025.