J. Eur. Math. Soc. 13, 1079–1112

DOI 10.4171/JEMS/276

Patrick Brosnan · Zinovy Reichstein · Angelo Vistoli

Essential dimension of moduli of curves and other algebraic stacks

(with an appendix by Najmuddin Fakhruddin)

Received July 4, 2009 and in revised form October 4, 2009

Abstract. In this paper we consider questions of the following type. Let k be a base field and K/k be a field extension. Given a geometric object X over a field K (e.g. a smooth curve of genus g), what is the least transcendence degree of a field of definition of X over the base field k? In other words, how many independent parameters are needed to define X? To study these questions we introduce a notion of essential dimension for an algebraic stack. Using the resulting theory, we give a complete answer to the question above when the geometric objects X are smooth, stable or hyperelliptic curves. The appendix, written by Najmuddin Fakhruddin, answers this question in the case of abelian varieties.

Keywords. Essential dimension, stack, gerbe, moduli of curves, moduli of abelian varieties

Contents

1. Introduction

This paper was motivated by the following question.

P. Brosnan, Z. Reichstein: Department of Mathematics, The University of British Columbia, 1984 Mathematics Road, Vancouver, B.C., Canada V6T 1Z2;
e-mail: brosnan@math.ubc.ca, reichst@math.ubc.ca

A. Vistoli: Scuola Normale Superiore, Piazza dei Cavalieri 7, 56126 Pisa, Italy;
e-mail: angelo.vistoli@sns.it

N. Fakhruddin: School of Mathematics, Tata Institue of Fundamental Research, Homi Bhabha Road, Mumbai 400005, India; e-mail: naf@math.tifr.res.in

Mathematics Subject Classification (2010): Primary 14A20, 14H10, 14K10

Question 1.1. Let k be a field and $g \geq 0$ be an integer. What is the smallest integer d such that for every field K/k, every smooth curve X of genus g defined over K descends to a subfield $k \subset K_0 \subset K$ with $\operatorname{tr\,deg}_k K_0 \leq d$?

Here by "X descends to K_0" we mean that there exists a curve X_0 over K_0 such that X is K-isomorphic to $X_0 \times_{\operatorname{Spec} K_0} \operatorname{Spec} K$.

In order to address this and related questions, we will introduce and study the notion of essential dimension for algebraic stacks; see §2. The essential dimension $\operatorname{ed} \mathcal{X}$ of a scheme $\mathcal{X}$ is simply the dimension of $\mathcal{X}$; on the other hand, the essential dimension of the classifying stack $\mathcal{B}_k G$ of an algebraic group G is the essential dimension of G in the usual sense; see [Rei00] or [BF03]. The notion of essential dimension of a stack is meant to bridge these two examples. The minimal integer d in Question 1.1 is the essential dimension of the moduli stack of smooth curves $\mathcal{M}_g$. We show that $\operatorname{ed} \mathcal{X}$ is finite for a broad class of algebraic stacks of finite type over a field; see Corollary 3.4. This class includes all Deligne–Mumford stacks and all quotient stacks of the form $\mathcal{X} = [X/G]$, where G a linear algebraic group.

Our main result is the following theorem.

Theorem 1.2. *Let $\mathcal{M}_{g,n}$ (respectively, $\overline{\mathcal{M}}_{g,n}$) be the stacks of n-pointed smooth (respectively, stable) algebraic curves of genus g over a field k of characteristic 0. Then*

$$\operatorname{ed} \mathcal{M}_{g,n} = \begin{cases} 2 & \textit{if } (g,n) = (0,0) \textit{ or } (1,1), \\ 0 & \textit{if } (g,n) = (0,1) \textit{ or } (0,2), \\ +\infty & \textit{if } (g,n) = (1,0), \\ 5 & \textit{if } (g,n) = (2,0), \\ 3g-3+n & \textit{otherwise}. \end{cases}$$

Moreover for $2g - 2 + n > 0$ we have $\operatorname{ed} \overline{\mathcal{M}}_{g,n} = \operatorname{ed} \mathcal{M}_{g,n}$.

In particular, the values of $\operatorname{ed} \mathcal{M}_{g,0} = \operatorname{ed} \mathcal{M}_g$ give a complete answer to Question 1.1.

Note that $3g - 3 + n$ is the dimension of the moduli space $\mathbf{M}_{g,n}$ in the stable range $2g - 2 + n > 0$ (and the dimension of the stack in all cases); the dimension of the moduli space represents an obvious lower bound for the essential dimension of a stack. The first four cases are precisely the ones where a generic object in $\mathcal{M}_{g,n}$ has non-trivial automorphisms, and $(g,n) = (1,0)$ is the only case where the automorphism group scheme of an object of $\mathcal{M}_{g,n}$ is not affine.

Our proof of Theorem 1.2 for $(g,n) \neq (1,0)$ relies on two results of independent interest. One is the "Genericity Theorem" 6.1 which says that the essential dimension of a smooth integral Deligne–Mumford stack satisfying an appropriate separation hypothesis is the sum of its dimension and the essential dimension of its generic gerbe. This somewhat surprising result implies that the essential dimension of a non-empty open substack equals the essential dimension of the stack. In particular, it proves Theorem 1.2 in the cases where a general curve in $\mathcal{M}_{g,n}$ has no non-trivial automorphisms. It also brings into relief the important role played by gerbes in this theory.

The second main ingredient in our proof of Theorem 1.2 is the following formula, which we use to compute the essential dimension of the generic gerbe.

Theorem 1.3. *Let $\mathcal{X}$ be a gerbe over a field K banded by a group G. Let $[\mathcal{X}] \in \mathrm{H}^2(K, G)$ be the Brauer class of $\mathcal{X}$.*

(a) *If $G = \mathbb{G}_\mathrm{m}$ and $\operatorname{ind}[\mathcal{X}]$ is a prime power then $\operatorname{ed}\mathcal{X} = \operatorname{ind}[\mathcal{X}] - 1$.*
(b) *If $G = \boldsymbol{\mu}_{p^r}$, where p is a prime and $r \geq 1$, then $\operatorname{ed}\mathcal{X} = \operatorname{ind}[\mathcal{X}]$.*

Our proof of this theorem can be found in the preprint [BRV07, Section 7]. A similar argument was used by N. Karpenko and A. Merkurjev in the proof of [KM08, Theorem 3.1], which generalizes Theorem 1.3(b). For the sake of completeness, we include an alternative proof of Theorem 1.3 in §4.

Theorem 1.3 has a number of applications beyond Theorem 1.2. Some of these have already appeared in print. In particular, we used Theorem 1.3 to study the essential dimension of spinor groups in [BRV10], N. Karpenko and A. Merkurjev [KM08] used it to study the essential dimension of finite p-groups, and A. Dhillon and N. Lemire [DL] used it, in combination with the Genericity Theorem 6.1, to give an upper bound for the essential dimension of the moduli stack of SL_n-bundles over a projective curve. In this paper Theorem 1.3 (in combination with Theorem 6.1) is also used to study the essential dimension of the stacks of hyperelliptic curves (Theorem 7.2) and, in the appendix written by Najmuddin Fakhruddin, of principally polarized abelian varieties.

In the case where $(g, n) = (1, 0)$ Theorem 1.2 requires a separate argument, which is carried out in §8. In this case Theorem 1.2 is a consequence of the fact that the group scheme of l^n-torsion points on a Tate curve has essential dimension l^n, where l is a prime.

2. The essential dimension of a stack

Let k be a field. We will write Fields_k for the category of field extensions K/k. Let $F\colon \mathrm{Fields}_k \to \mathrm{Sets}$ be a covariant functor.

Definition 2.1. Let $a \in F(L)$, where L is an object of Fields_k. We say that *a descends* to an intermediate field $k \subset K \subset L$ or equivalently that K is a *field of definition* for a if a is in the image of the induced map $F(K) \to F(L)$.

The *essential dimension* $\operatorname{ed} a$ of $a \in F(L)$ is the minimum of the transcendence degrees $\operatorname{tr deg}_k K$ taken over all intermediate fields $k \subseteq K \subseteq L$ such that a descends to K.

The essential dimension $\operatorname{ed} F$ of the functor F is the supremum of $\operatorname{ed} a$ taken over all $a \in F(L)$ with L in Fields_k. We will write $\operatorname{ed} F = -\infty$ if F is the empty functor.

These notions are relative to the base field k. To emphasize this, we will sometimes write $\operatorname{ed}_k a$ or $\operatorname{ed}_k F$ instead of $\operatorname{ed} a$ or $\operatorname{ed} F$, respectively.

The following definition singles out a class of functors that is sufficiently broad to include most interesting examples, yet "geometric" enough to allow one to get a handle on their essential dimension.

Definition 2.2. Suppose $\mathcal{X}$ is an algebraic stack over k. The *essential dimension* $\operatorname{ed}\mathcal{X}$ of $\mathcal{X}$ is defined to be the essential dimension of the functor $F_{\mathcal{X}}\colon \text{Fields}_k \to \text{Sets}$ which sends a field L/k to the set of isomorphism classes of objects in the groupoid $\mathcal{X}(L)$.[1]

As in Definition 2.1, we will write $\operatorname{ed}_k \mathcal{X}$ when we need to be specific about the dependence on the base field k. Similarly for $\operatorname{ed}_k \xi$, where ξ is an object of $F_{\mathcal{X}}$.

Example 2.3. Let G be an algebraic group defined over k and $\mathcal{X} = \mathcal{B}_k G$ be the classifying stack of G. Then $F_{\mathcal{X}}$ is the Galois cohomology functor sending K to the set $\mathrm{H}^1(K, G)$ of isomorphism classes of G-torsors over $\operatorname{Spec}(K)$, in the fppf topology. The essential dimension of this functor is a numerical invariant of G, which, roughly speaking, measures the complexity of G-torsors over fields. This number is usually denoted by $\operatorname{ed}_k G$ or (if k is fixed throughout) simply by $\operatorname{ed} G$; following this convention, we will often write $\operatorname{ed} G$ in place of $\operatorname{ed}\mathcal{B}_k G$. Essential dimension was originally introduced and has since been extensively studied in this context; see e.g., [BR97, Rei00, RY00, Kor00, Led02, JLY02, BF03, Lem04, CS06, Gar09]. The more general Definition 2.1 is due to A. Merkurjev; see [BF03, Proposition 1.17].

Example 2.4. Let $\mathcal{X} = X$ be a scheme of finite type over a field k, and let $F_X\colon \text{Fields}_k \to \text{Sets}$ denote the functor given by $K \mapsto X(K)$. Then an easy argument due to Merkurjev shows that $\operatorname{ed} F_X = \dim X$; see [BF03, Proposition 1.17].

In fact, this equality remains true for any algebraic space X. Indeed, an algebraic space X has a stratification by schemes X_i. Any K-point $\eta\colon \operatorname{Spec} K \to X$ must land in one of the X_i. Thus $\operatorname{ed} X = \max \operatorname{ed} X_i = \dim X$.

Example 2.5. Let $\mathcal{X} = \mathcal{M}_{g,n}$ be the stack of smooth algebraic curves of genus g. Then the functor $F_{\mathcal{X}}$ sends K to the set of isomorphism classes of n-pointed smooth algebraic curves of genus g over K. Question 1.1 asks about the essential dimension of this functor in the case where $n = 0$.

Example 2.6. Suppose a linear algebraic group G is acting on an algebraic space X over a field k. We shall write $[X/G]$ for the quotient stack $[X/G]$. Recall that K-points of $[X/G]$ are by definition diagrams of the form

$$\begin{array}{ccc} T & \xrightarrow{\psi} & X \\ \big\downarrow{\scriptstyle \pi} & & \\ \operatorname{Spec} K & & \end{array} \tag{2.1}$$

where π is a G-torsor and ψ is a G-equivariant map. The functor $F_{[X/G]}$ associates with a field K/k the set of isomorphism classes of such diagrams.

In the case where G is a *special* group (recall that this means that every G-torsor over $\operatorname{Spec} K$ is split, for every field K/k) the essential dimension of $F_{[X/G]}$ has been

[1] In the literature the functor $F_{\mathcal{X}}$ is sometimes denoted by $\widehat{\mathcal{X}}$ or $\overline{\mathcal{X}}$.

previously studied in connection with the so-called "functor of orbits" $\mathbf{Orb}_{X,G}$ given by the formula

$$\mathbf{Orb}_{X,G}(K) \stackrel{\text{def}}{=} \text{set of } G(K)\text{-orbits in } X(K).$$

Indeed, if G is special, the functors $F_{[X/G]}$ and $\mathbf{Orb}_{X,G}$ are isomorphic; an isomorphism between them is given by sending an object (2.1) of $F_{[X/G]}$ to the $G(K)$-orbit of the point $\psi s\colon \operatorname{Spec} K \to X$, where $s\colon \operatorname{Spec} K \to T$ is a section of $\pi\colon T \to \operatorname{Spec} K$.

Of particular interest are the natural GL_n-actions on $\mathbb{A}^N$ = affine space of homogeneous polynomials of degree d in n variables and on $\mathbb{P}^{N-1}$ = projective space of degree d hypersurfaces in $\mathbb{P}^{n-1}$, where $N = \binom{n+d-1}{d}$ is the number of degree d monomials in n variables. For general n and d the essential dimension of the functor of orbits in these cases is not known. The study of this problem was initiated in [BF04] and [BR05, Sections 14–15]; stronger results can be found in the recent preprint [RV11].

Remark 2.7. If the functor F in Definition 2.1 is limit-preserving, a condition satisfied in all cases of interest to us, then every element $a \in F(L)$ descends to a field $K \subset L$ that is finitely generated over k. Thus in this case $\operatorname{ed} a$ is finite. In particular, if $\mathcal{X}$ is an algebraic stack over k, $\operatorname{ed} \xi$ is finite for every object $\xi \in \mathcal{X}(K)$ and every field extension K/k; the limit-preserving property in this case is proved in [LMB00, Proposition 4.18],

In §3 we will show that, in fact, $\operatorname{ed} \mathcal{X} < \infty$ for a broad class of algebraic stacks $\mathcal{X}$; cf. Corollary 3.4. On the other hand, there are interesting examples where $\operatorname{ed} \mathcal{X} = \infty$; see Theorem 1.2 or [BS08].

The following observation is a variant of [BF03, Proposition 1.5].

Proposition 2.8. *Let $\mathcal{X}$ be an algebraic stack over k, and let K be a field extension of k. Then* $\operatorname{ed}_K \mathcal{X}_K \leq \operatorname{ed}_k \mathcal{X}$.

Here, as in what follows, we denote by $\mathcal{X}_K$ the stack $\operatorname{Spec} K \times_{\operatorname{Spec} k} \mathcal{X}$.

Proof. If L/K is a field extension, then the natural morphism $\mathcal{X}_K(L) \to \mathcal{X}(L)$ is an equivalence. Suppose that M/k is a field of definition for an object ξ in $\mathcal{X}(L)$. Let N be a composite of M and K over k. Then N is a field of definition for ξ, $\operatorname{tr deg}_K N \leq \operatorname{tr deg}_k M$, and the proposition follows. □

3. A fiber dimension theorem

We now recall Definitions (3.9) and (3.10) from [LMB00]. A morphism $f\colon \mathcal{X} \to \mathcal{Y}$ of algebraic stacks (over k) is said to be *representable* if, for every k-morphism $T \to \mathcal{Y}$, where T is an affine k-scheme, the fiber product $\mathcal{X} \times_{\mathcal{Y}} T$ is representable by an algebraic space over T. A representable morphism $f\colon \mathcal{X} \to \mathcal{Y}$ is said to be *locally of finite type and of fiber dimension* $\leq d$ if for every $T \to \mathcal{Y}$ as above, the projection $\mathcal{X} \times_{\mathcal{Y}} T \to T$ is locally of finite type over T and every fiber has dimension $\leq d$.

Example 3.1. Let G be an algebraic group defined over k, and let $X \to Y$ be a G-equivariant morphism of k-algebraic spaces, locally of finite type and of relative dimension $\leq d$. Then the induced map of quotient stacks $[X/G] \to [Y/G]$ is representable, locally of finite type and of relative dimension $\leq d$.

The following result may be viewed as a partial generalization of the fiber dimension theorem (see [Har77, Exercise II.3.22 or Proposition III.9.5]) to the setting where schemes are replaced by stacks and dimension by essential dimension.

Theorem 3.2. *Let d be an integer, and $f: \mathcal{X} \to \mathcal{Y}$ be a representable k-morphism of algebraic stacks which is locally of finite type and of fiber dimension at most d. Let L/k be a field, and $\xi \in \mathcal{X}(L)$. Then*

(a) $\mathrm{ed}_k\, \xi \leq \mathrm{ed}_k\, f(\xi) + d$,
(b) $\mathrm{ed}_k\, \mathcal{X} \leq \mathrm{ed}_k\, \mathcal{Y} + d$.

In particular, if $\mathrm{ed}_k\, \mathcal{Y}$ *is finite, then so is* $\mathrm{ed}_k\, \mathcal{X}$.

Proof. (a) By the definition of $\mathrm{ed}_k\, f(\xi)$ we can find an intermediate field $k \subset K \subset L$ and a morphism η: $\operatorname{Spec} K \to \mathcal{Y}$ such that $\operatorname{tr\,deg}_k K \leq \mathrm{ed}\, f(\xi)$ and the following diagram commutes:

$$\begin{array}{ccc} \operatorname{Spec} L & \xrightarrow{\ \xi\ } & \mathcal{X} \\ \downarrow & & \downarrow{\scriptstyle f} \\ \operatorname{Spec} K & \xrightarrow{\ \eta\ } & \mathcal{Y} \end{array}$$

Let $\mathcal{X}_K \stackrel{\text{def}}{=} \mathcal{X} \times_{\mathcal{Y}} \operatorname{Spec} K$. By the hypothesis, $\mathcal{X}_K$ is an algebraic space, locally of finite type over K and of relative dimension at most d. By the commutativity of the diagram above, the morphism ξ: $\operatorname{Spec} L \to \mathcal{X}$ factors through $\mathcal{X}_K$:

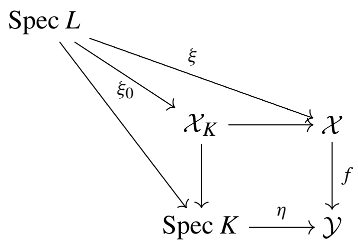

Moreover, ξ factors through $K(p)$, where p denotes the image of ξ_0 in $\mathcal{X}_K$. Since $\mathcal{X}_K$ has dimension at most d over K, we have $\operatorname{tr\,deg}_K K(p) \leq d$. Therefore,

$$\operatorname{tr\,deg}_k K(p) = \operatorname{tr\,deg}_k K + \operatorname{tr\,deg}_K K(p) \leq \mathrm{ed}\, f(\xi) + d$$

and part (a) follows.

Part (b) follows from (a) by taking the maximum on both sides over all L/k and all $\xi \in \mathcal{X}(L)$. □

Corollary 3.3. *Consider an action of an algebraic group G on an algebraic space X, defined over a field k. Assume X is locally of finite type over k. Then*

$$\mathrm{ed}_k\, G \geq \mathrm{ed}_k\, [X/G] - \dim X.$$

Proof. The natural G-equivariant map $X \to \operatorname{Spec} k$ gives rise to a map $[X/G] \to \mathcal{B}_k G$ of quotient stacks. This latter map is locally of finite type and of relative dimension $\leq \dim X$; see Example 3.1. Theorem 3.2(b) applied to this map yields the desired inequality. □

Corollary 3.4 (Finiteness of essential dimension). *Let $\mathcal{X}$ be an algebraic stack of finite type over k. Suppose that for any algebraically closed extension Ω of k and any object ξ of $\mathcal{X}(\Omega)$ the group scheme $\underline{\operatorname{Aut}}_{\Omega}(\xi) \to \operatorname{Spec} \Omega$ is affine. Then* $\operatorname{ed}_k \mathcal{X} < \infty$.

Note that Corollary 3.4 fails without the assumption that all the $\underline{\operatorname{Aut}}_{\Omega}(\xi)$ are affine. For example, by Theorem 1.2, $\operatorname{ed} \mathcal{M}_{1,0} = +\infty$.

Proof. We may assume without loss of generality that $\mathcal{X} = [X/G]$ is a quotient stack for some affine algebraic group G acting on an algebraic space X. Indeed, by a theorem of Kresch [Kre99, Proposition 3.5.9], $\mathcal{X}$ is covered by quotient stacks $[X_i/G_i]$ of this form, and hence $\operatorname{ed} \mathcal{X} = \max_i \operatorname{ed} [X_i/G_i]$.

If $\mathcal{X} = [X/G]$ then by Corollary 3.3,

$$\operatorname{ed} [X/G] \leq \operatorname{ed}_k G + \dim X.$$

The desired conclusion now follows from the well-known fact that $\operatorname{ed}_k G < \infty$ for any affine algebraic group G; see [Rei00, Theorem 3.4] or [BF03, Proposition 4.11]. □

4. The essential dimension of a gerbe over a field

The goal of this section is to prove Theorem 1.3 stated in the Introduction. We proceed by briefly recalling some background material on gerbes from [Mil80, p. 144] and [Gir71, IV.3.1.1], and on canonical dimension from [KM06] and [BR05].

Gerbes. Let $\mathcal{X}$ be a gerbe defined over a field K *banded* by an abelian K-group scheme G. In particular, X is a stack over K which becomes isomorphic to $\mathcal{B}_K G$ over the algebraic closure of K.

There is a notion of equivalence of gerbes banded by G; the set of equivalence classes is in a natural bijective correspondence with the group $\mathrm{H}^2(K, G)$. The identity element of $\mathrm{H}^2(K, G)$ corresponds to the class of the neutral gerbe $\mathcal{B}_K G$. Recall that the group $\mathrm{H}^2(K, \mathbb{G}_{\mathrm{m}})$ is canonically isomorphic to the Brauer group $\operatorname{Br} K$ of Brauer equivalence classes of central simple algebras over K. Here, as usual, $\mathbb{G}_{\mathrm{m}}$ denotes the multiplicative group scheme over K.

Canonical dimension. Let X be a smooth projective variety defined over a field K. We say that L/K is a *splitting field* for X if $X(L) \neq \emptyset$. A splitting field L/K is called *generic* if for every splitting field L_0/K there exists a K-place $L \to L_0$. The *canonical dimension* $\operatorname{cd} X$ of X is defined as the minimal value of $\operatorname{tr deg}_K L$, where L/K ranges over all generic splitting fields. Note that the function field $L = K(X)$ is a generic splitting field of X; see [KM06, Lemma 4.1]. In particular, generic splitting fields exist and $\operatorname{cd} X$ is finite. If X is a smooth complete projective variety over K then $\operatorname{cd} X$ has the following simple geometric interpretation: $\operatorname{cd} X$ is the minimal value of $\dim(Y)$ as Y ranges over the closed

K-subvarieties of X which admit a rational map $X \dashrightarrow Y$ defined over K; see [KM06, Corollary 4.6].

The *determination functor* $D_X\colon$ Fields$_K \to$ Sets is defined as follows. For any field extension L/K, $D_X(L)$ is the empty set if $X(L) = \emptyset$, and a set consisting of one element if $X(L) \neq \emptyset$. The natural map $D(L_1) \to D(L_2)$ is then uniquely determined for any $K \subset L_1 \subset L_2$. It is shown in [KM06] that if X is a complete regular K-variety then

$$\operatorname{cd} X = \operatorname{ed} D_X. \tag{4.1}$$

Of particular interest to us will be the case where X is a Brauer–Severi variety over K. Let m be the index of X. If $m = p^a$ is a prime power then

$$\operatorname{cd} X = p^a - 1; \tag{4.2}$$

see [KM06, Example 3.10] or [BR05, Theorem 11.4].

If $m = p_1^{a_1} \dots p_r^{a_r}$ is the prime decomposition of m then the class of X in $\operatorname{Br} L$ is the sum of classes $\alpha_1, \dots, \alpha_r$ whose indices are $p_1^{a_1}, \dots, p_r^{a_r}$. Denote by $X_1, \dots, X_r$ the Brauer–Severi varieties associated with $\alpha_1, \dots, \alpha_r$. It is easy to see that $K(X_1 \times \dots \times X_r)$ is a generic splitting field for X. Hence,

$$\operatorname{cd} X \leq \dim(X_1 \times \dots \times X_r) = p_1^{a_1} + \dots + p_r^{a_r} - r.$$

J.-L. Colliot-Thélène, N. Karpenko and A. Merkurjev [CTKM07] conjectured that equality holds, i.e.,

$$\operatorname{cd} X = p_1^{a_1} + \dots + p_r^{a_r} - r. \tag{4.3}$$

As we mentioned above, this is known to be true if m is a prime power (i.e., $r = 1$). Colliot-Thélène, Karpenko and Merkurjev also proved (4.3) for $m = 6$; see [CTKM07, Theorem 1.3]. Their conjecture remains open for all other m.

Theorem 4.1. *Let d be an integer with $d > 1$. Let K be a field and $x \in \mathrm{H}^2(K, \boldsymbol{\mu}_d)$. Denote the image of x in $\mathrm{H}^2(K, \mathbb{G}_\mathrm{m})$ by y, the $\boldsymbol{\mu}_d$-gerbe associated with x by $\mathcal{X} \to \operatorname{Spec} K$, the $\mathbb{G}_\mathrm{m}$-gerbe associated with y by $\mathcal{Y} \to \operatorname{Spec} K$, and the Brauer–Severi variety associated with y by P. Then*

(a) $\operatorname{ed} \mathcal{Y} = \operatorname{cd} P$,
(b) $\operatorname{ed} \mathcal{X} = \operatorname{cd} P + 1$.

In particular, if the index of x is a prime power p^r then $\operatorname{ed} \mathcal{Y} = p^r - 1$ and $\operatorname{ed} \mathcal{X} = p^r$.

Proof. The last assertion follows from (a) and (b) by (4.2).

(a) The functor $F_\mathcal{Y}\colon$ Fields$_K \to$ Sets sends a field L/K to the empty set if $P(L) = \emptyset$, and to a set consisting of one point if $P(L) \neq \emptyset$. In other words, F_Y is the determination functor D_P introduced above. The essential dimension of this functor is $\operatorname{cd} P$; see (4.1).

(b) First note that the natural map $\mathcal{X} \to \mathcal{Y}$ is of finite type and representable of relative dimension ≤ 1. By Theorem 3.2(b) we conclude that $\operatorname{ed} \mathcal{X} \leq \operatorname{ed} \mathcal{Y} + 1$. By part (a) it remains to prove the opposite inequality, $\operatorname{ed} \mathcal{X} \geq \operatorname{ed} \mathcal{Y} + 1$. We will do this by constructing an object α of $\mathcal{X}$ whose essential dimension is $\geq \operatorname{ed} \mathcal{Y} + 1$.

We will view $\mathcal{X}$ as a torsor for $\mathcal{B}_K\boldsymbol{\mu}_d$ in the following sense. There exist maps

$$\mathcal{X} \times \mathcal{B}_K\boldsymbol{\mu}_d \to \mathcal{X}, \quad \mathcal{X} \times \mathcal{X} \to \mathcal{B}_K\boldsymbol{\mu}_d$$

satisfying various compatibilities, where the first map is the "action" of $\mathcal{B}_K\boldsymbol{\mu}_d$ on $\mathcal{X}$ and the second map is the "difference" of two objects of $\mathcal{X}$. For the definition and a discussion of the properties of these maps, see [Gir71, Chapter IV, Sections 2.3, 2.4 and 3.3]. (Note that, in the notation of Giraud's book, $\mathcal{X} \wedge \mathcal{B}_K\boldsymbol{\mu}_d \cong \mathcal{X}$ and the action operation above arises from the map $\mathcal{X} \times \mathcal{B}_K\boldsymbol{\mu}_d \to \mathcal{X} \wedge \mathcal{B}_K\boldsymbol{\mu}_d$ given in Chapter IV, Proposition 2.4.1. The difference operation, which we will not use here, arises similarly from the fact that, in Giraud's notation, $\mathrm{HOM}(\mathcal{X}, \mathcal{X}) \cong \mathcal{B}_K\boldsymbol{\mu}_d$.)

Let $L = K(P)$ be the function field of P. Since L splits P, we have a natural map $a\colon \operatorname{Spec} L \to \mathcal{Y}$. Moreover since L is a generic splitting field for P,

$$\operatorname{ed} a = \operatorname{cd} P = \operatorname{ed} \mathcal{Y}, \tag{4.4}$$

where we view a as an object in $\mathcal{Y}$. Non-canonically lift $a\colon \operatorname{Spec} L \to \mathcal{Y}$ to a map $\operatorname{Spec} L \to \mathcal{X}$ (this can be done, because $\mathcal{X} \to \mathcal{Y}$ is a $\mathbb{G}_{\mathrm{m}}$-torsor). Let $\operatorname{Spec} L(t) \to \mathcal{B}_L\boldsymbol{\mu}_d$ denote the map classified by $(t) \in \mathrm{H}^1(L(t), \boldsymbol{\mu}_d) = L(t)^{\times}/L(t)^{\times d}$. Composing these two maps, we obtain an object

$$\alpha\colon \operatorname{Spec} L(t) \to \mathcal{X} \times \mathcal{B}_L\boldsymbol{\mu}_d \to \mathcal{X}.$$

in $\mathcal{X}(L(t))$. Our goal is to prove that $\operatorname{ed} \alpha \geq \operatorname{ed} \mathcal{Y} + 1$. In other words, given a diagram of the form

$$\begin{array}{ccc} \operatorname{Spec} L(t) & \xrightarrow{\alpha} & \mathcal{X} \\ \downarrow & \nearrow_{\beta} & \\ \operatorname{Spec} M & & \end{array} \tag{4.5}$$

where $K \subset M \subset L$ is an intermediate field, we want to prove the inequality $\operatorname{tr\,deg}_K M \geq \operatorname{ed} \mathcal{Y} + 1$. Assume the contrary: there is a diagram as above with $\operatorname{tr\,deg}_K M \leq \operatorname{ed} \mathcal{Y}$. Let $\nu\colon L(t)^* \to \mathbb{Z}$ be the usual discrete valuation corresponding to t and consider two cases.

Case 1. Suppose the restriction $\nu_{|M}$ of ν to M is non-trivial. Let M_0 denote the residue field of ν and $M_{\geq 0}$ denote the valuation ring. Since $\operatorname{Spec} M \to \mathcal{X} \to \mathcal{Y}$, there exists an M-point of P. Then by the valuative criterion of properness for P, there exists an $M_{\geq 0}$-point and thus an M_0-point of P. Passing to residue fields, we obtain the diagram

$$\begin{array}{ccc} \operatorname{Spec} L & \xrightarrow{a} & \mathcal{Y} \\ \downarrow & \nearrow & \\ \operatorname{Spec} M_0 & & \end{array}$$

which shows that $\operatorname{ed} a \leq \operatorname{tr\,deg}_K M_0 = \operatorname{tr\,deg}_K M - 1 \leq \operatorname{ed} \mathcal{Y} - 1$, contradicting (4.4).

Case 2. Now suppose the restriction of ν to M is trivial. The map $\operatorname{Spec} L \to \mathcal{X}$ sets up an isomorphism $\mathcal{X}_L \cong \mathcal{B}_L\boldsymbol{\mu}_d$. The map $\operatorname{Spec} L(t) \to \mathcal{X}$ factors through $\mathcal{X}_L$ and thus induces a class in $\mathcal{B}_L\boldsymbol{\mu}_d(L(t)) = \mathrm{H}^1(L(t), \boldsymbol{\mu}_d)$. This class is (t). Tensoring the diagram (4.5) with L over K, we obtain

$$\begin{array}{ccc} \operatorname{Spec} L(t) \otimes L & \xrightarrow{\alpha} & \mathcal{X}_L \cong \mathcal{B}_L\boldsymbol{\mu}_d \\ \downarrow & \nearrow_{\beta} & \\ \operatorname{Spec} M \otimes L & & \end{array}$$

Recall that $L = K(P)$ is the function field of P. Since P is absolutely irreducible, the tensor products $L(t)\otimes L$ and $M\otimes L$ are fields. The map $\operatorname{Spec} M\otimes L \to \mathcal{B}_L\boldsymbol{\mu}_d$ is classified by some $m \in (M \otimes L)^\times/(M \otimes L)^{\times d} = \mathrm{H}^1(M \otimes L, \boldsymbol{\mu}_d)$. The image of m in $L(t) \otimes L$ is equal to t modulo d-th powers. We will now derive a contradiction by comparing the valuations of m and t.

To apply the valuation to m, we lift ν from $L(t)$ to $L(t) \otimes L$. That is, we define ν_L as the valuation on $L(t) \otimes L = (L \otimes L)(t)$ corresponding to t. Since $\nu_L(t) = \nu(t) = 1$, we conclude that $\nu_L(m) \equiv 1 \pmod d$. This shows that ν_L is not trivial on $M \otimes L$ and thus ν is not trivial on M, contradicting our assumption. This contradiction completes the proof of part (b). □

Corollary 4.2. *Let* $1 \to Z \to G \to Q \to 1$ *denote an extension of group schemes over a field k with Z central and isomorphic to* (a) $\mathbb{G}_{\mathrm{m}}$ *or* (b) $\boldsymbol{\mu}_{p^r}$ *for some prime p and some $r \geq 1$. Let* $\operatorname{ind}(G, Z)$ *be the maximal value of* $\operatorname{ind}(\partial_K(t))$ *as K ranges over all field extensions of k and t ranges over all torsors in* $\mathrm{H}^1(K, Q)$. *If* $\operatorname{ind}(G, Z)$ *is a prime power* (*which is automatic in case* (b)) *then*

$$\operatorname{ed}_k G \geq \operatorname{ind}(G, Z) - \dim G.$$

Proof. Choose $t \in \mathrm{H}^1(K, Z)$ so that $\operatorname{ind}(\partial_K(t))$ attains its maximal value, $\operatorname{ind}(G, Z)$. Let $X \to \operatorname{Spec} K$ be the Q-torsor representing t. Then G acts on X via the projection $G \to Q$, and $[X/G]$ is the Z-gerbe over $\operatorname{Spec} K$ corresponding to the class $\partial_K(t) \in \mathrm{H}^2(K, Z)$. By Theorem 1.3,

$$\operatorname{ed}[X/G] = \begin{cases} \operatorname{ind}(\partial_K(t)) - 1 & \text{in case (a),} \\ \operatorname{ind}(\partial_K(t)) & \text{in case (b).} \end{cases}$$

Since $\dim X = \dim Q$, applying Corollary 3.3 to the G-action on X, we obtain

$$\operatorname{ed}_K G_K \geq \begin{cases} (\operatorname{ind}(G, Z) - 1) - \dim Q = \operatorname{ind}(G, Z) - \dim G & \text{in case (a),} \\ \operatorname{ind}(G, Z) - \dim Q = \operatorname{ind}(G, Z) - \dim G & \text{in case (b).} \end{cases}$$

Since $\operatorname{ed}_k G \geq \operatorname{ed}_K G_K$ (see [BF03, Proposition 1.5] or our Proposition 2.8), the corollary follows. □

5. Gerbes over complete discrete valuation rings

In this section we prove two results on the structure of étale gerbes over complete discrete valuation rings that will be used in the proof of Theorem 6.1.

5.1. Big and small étale sites

Let S be a scheme. We let Sch/S denote the category of all schemes T equipped with a morphism to S. As in [SGA72], we equip Sch/S with the étale topology. Let $\text{ét}/S$ denote the full subcategory of Sch/S consisting of all schemes étale over S (also with the étale topology). The site Sch/S is the *big étale site* and the category $\text{ét}/S$ is the *small étale site*. We let $S_{\text{Ét}}$ denote the category of sheaves on Sch/S, and $S_{\text{ét}}$ the category of sheaves on $\text{ét}/S$. Since the obvious inclusion functor from the small to the big étale site is continuous, it induces a continuous morphism of sites $u\colon \text{ét}/S \to \mathrm{Sch}/S$ and thus a morphism $u\colon S_{\text{Ét}} \to S_{\text{ét}}$. Moreover, the adjunction morphism $F \to u_* u^* F$ is an isomorphism for F a sheaf in $S_{\text{ét}}$ [SGA72, VII.4.1]. We can therefore regard $S_{\text{ét}}$ as a full subcategory of $S_{\text{Ét}}$.

Definition 5.1. Let S be a scheme. An *étale gerbe* over S is a separated locally finitely presented Deligne–Mumford stack over S that is a gerbe in the étale topology.

Let $\mathcal{X} \to S$ be an étale gerbe over a scheme S. Then, by definition, there is an étale atlas, i.e., a morphism $U_0 \to \mathcal{X}$, where $U_0 \to S$ is surjective, étale and finitely presented over S. This atlas gives rise to a groupoid $\mathcal{G} \overset{\text{def}}{=} [U_1 \overset{\text{def}}{=} U_0 \times_{\mathcal{X}} U_0 \rightrightarrows U_0]$ in which each term is étale over S. Since $\mathcal{X}$ is the stackification of $\mathcal{G}$ which is a groupoid on the small étale site $S_{\text{ét}}$, it follows that $\mathcal{X} = u^*\mathcal{X}'$ for a gerbe $\mathcal{X}'$ on $S_{\text{ét}}$. In other words, we have the following proposition.

Proposition 5.2. *Let $\mathcal{X} \to S$ be an étale gerbe over a scheme S. Then there is a gerbe $\mathcal{X}'$ on $S_{\text{ét}}$ such that $\mathcal{X} = u^*\mathcal{X}'$.*

If S is a *henselian trait* (i.e., the spectrum of a henselian discrete valuation ring) we can do better:

Proposition 5.3. *Let S be a henselian trait and $f\colon T \to S$ be a surjective étale morphism. Then there is an open component T' of T such that $f_{|T'}\colon T' \to S$ is a finite étale morphism.*

Proof. Let s denote the closed point of S. Since f is surjective, there exists a $t \in T$ such that $f(t) = s$. Since f is étale, f is quasi-finite at t by [Gro67, 17.6.1]. Now, it follows from [Gro67, 18.5.11] that $T' \overset{\text{def}}{=} \operatorname{Spec} \mathcal{O}_{T,t}$ is an open component of T which is finite and étale. □

Now for a scheme S, let $\text{fét}/S$ denote the category of finite étale covers $T \to S$. We can consider $\text{fét}/S$ as a site in the obvious way. Then the inclusion morphism induces a continuous morphism of sites $v\colon \text{ét}/S \to \text{fét}/S$. If S is a henselian trait with closed point s, then the inclusion morphism $i\colon s \to S$ induces an equivalence of categories

$i^*\colon \text{fét}/S \to \text{fét}/s$. Since the site $\text{fét}/s$ is equivalent to $s_{\text{ét}}$, this induces the *specialization morphism* $\text{sp}\colon S_{\text{ét}} \to s_{\text{ét}}$, which is inverse to the inclusion morphism $i\colon s_{\text{ét}} \to S$; cf. [SGA73, p. 89]. Let $\tau = \text{sp} \circ u\colon S_{\text{Ét}} \to s_{\text{ét}}$.

Corollary 5.4. *Let $\mathcal{X} \to S$ be an étale gerbe over a henselian trait S with closed point s. Then there is a gerbe $\mathcal{X}''$ over $s_{\text{ét}}$ such that $\mathcal{X} = \tau^*\mathcal{X}''$.*

Proof. Since $\mathcal{X} \to S$ is an étale gerbe, there is an étale atlas $X_0 \to S$ of $\mathcal{X}$. By Proposition 5.3 we may assume that X_0 is finite over S. Then $X_1 \stackrel{\text{def}}{=} X_0 \times_{\mathcal{X}} X_0$ is also finite, because $\mathcal{X}$ is separated, by hypothesis. Now the equivalence of categories $i^*\colon \text{fét}/S \to \text{fét}/s$ produces a gerbe $\mathcal{X}''$ over $s_{\text{ét}}$ such that $\mathcal{X} = \tau^*\mathcal{X}''$. □

5.2. Group extensions and gerbes

Let k be a field with separable closure $\bar{k}$ and absolute Galois group $G = \text{Gal}(\bar{k}/k)$. Let

$$1 \to F \stackrel{i}{\to} E \stackrel{p}{\to} G \to 1 \tag{5.1}$$

be an extension of profinite groups with F finite and all maps continuous. From this data, we can construct a gerbe $\mathcal{X}_E$ over $(\text{Spec}\, k)_{\text{ét}}$. To determine the gerbe it is enough to give its category of sections over $\text{Spec}\, L$ where L/k is a finite separable extension. Let $K = \{g \in G \mid g(\alpha) = \alpha,\ \alpha \in L\}$. Then the objects of the category $\mathcal{X}_E(L)$ are the solutions of the embedding problem given by (5.1). That is, an object of $\mathcal{X}_E(L)$ is a continuous homomorphism $\sigma\colon K \to E$ such that $p \circ \sigma(k) = k$ for $k \in K$. If $s_i\colon K \to E$, $i = 1, 2$, are two objects in $\mathcal{X}_E(L)$ then a morphism from s_1 to s_2 is an element $f \in F$ such that $f s_1 f^{-1} = s_2$; cf. [DD99, p. 581].

By the results of Giraud [Gir71, Chapter VIII], it is easy to see that any gerbe $\mathcal{X} \to \text{Spec}\, k$ with finite inertia arises from a sequence (5.1) as above. We explain how to get the extension: Given $\mathcal{X}$, we can find a separable Galois extension L/k and an object $\xi \in \mathcal{X}(L)$. This gives an extension of groups $\text{Aut}_{\mathcal{X}}(\xi) \to \text{Aut}_{\text{Spec}\, k}(\text{Spec}\, L) = \text{Gal}(L/k)$. Pulling back this extension via the map $G = \text{Gal}(k) \to \text{Gal}(L/k)$ gives the desired sequence (5.1).

Now, suppose that E is as in (5.1). Let L/k be a field extension, which is separable but not necessarily finite. Let $\bar{L}$ denote a fixed separable closure of L and let $\bar{k}$ denote the separable closure of k in $\bar{L}$. Then there is an obvious map $r\colon \text{Gal}(\bar{L}/L) \to \text{Gal}(\bar{k}/k)$. Let $u\colon (\text{Spec}\, k)_{\text{Ét}} \to (\text{Spec}\, k)_{\text{ét}}$ denote the functor of §5.1. Then $u^*\mathcal{X}(L)$ has the same description as in the case where L is a finite extension of k. In other words, we have the following proposition.

Proposition 5.5. *Let L/k be a separable extension and let $\mathcal{X}_E$ be the gerbe defined above. Then the objects of the category $u^*\mathcal{X}_E(L)$ are the morphisms $s\colon \text{Gal}(L) \to E$ making the following diagram commute:*

$$\begin{array}{ccccccccc} & & & & & & \text{Gal}(L) & & \\ & & & & & \stackrel{s}{\swarrow} & \downarrow r & & \\ 1 & \longrightarrow & F & \longrightarrow & E & \longrightarrow & \text{Gal}(k) & \longrightarrow & 1 \end{array}$$

Moreover, if s_i, $i = 1, 2$, *are two objects in* $u^*\mathcal{X}_E(L)$, *then the morphisms from* s_1 *to* s_2 *are the elements* $f \in F$ *such that* $f s_1 f^{-1} = s_2$. □

5.3. Splitting the inertia sequence

We begin by recalling some results and notation from Serre's chapter in [GMS03].

Let A be a discrete valuation ring. Write $S = S_A$ for $\operatorname{Spec} A$, $s = s_A$ for the closed point in S and $\eta = \eta_A$ for the generic point. When A is the only discrete valuation ring under consideration, we suppress the subscripts. If A is henselian, then the choice of a separable closure $k(\overline{\eta})$ of $k(\eta)$ induces a separable closure of $k(s)$ and a map $\operatorname{Gal}(k(\overline{\eta})) \to \operatorname{Gal}(k(\overline{s}))$ between the absolute Galois groups. The kernel of this map is called the *inertia*, written as $I = I_A$. If $\operatorname{char} k(s) = p > 0$, then we set $I^{\mathrm{w}} = I_A^{\mathrm{w}}$ equal to the unique p-Sylow subgroup of I; otherwise we set $I^{\mathrm{w}} = \{1\}$. The group I^{w} is called the *wild inertia*. The group $I_{\mathrm{t}} = I_{A,t} \stackrel{\mathrm{def}}{=} I/I^{\mathrm{w}}$ is called the *tame inertia* and the group $\operatorname{Gal}(k(\eta))_{\mathrm{t}} \stackrel{\mathrm{def}}{=} \operatorname{Gal}(k(\eta))/I^{\mathrm{w}}$ is called the *tame Galois group*. We therefore have the following exact sequences:

$$1 \to I \to \operatorname{Gal}(k(\eta)) \to \operatorname{Gal}(k(s)) \to 1, \tag{5.2}$$

$$1 \to I_{\mathrm{t}} \to \operatorname{Gal}(k(\eta))_{\mathrm{t}} \to \operatorname{Gal}(k(s)) \to 1. \tag{5.3}$$

The sequence (5.2) is called the *inertia exact sequence* and (5.3) the *tame inertia exact sequence*.

For each prime l, set $\mathbb{Z}_l(1) = \varprojlim \boldsymbol{\mu}_{l^n}$ so that

$$\prod_{l \neq p} \mathbb{Z}_l(1) = \varprojlim_{p \nmid n} \boldsymbol{\mu}_n.$$

Then there is a canonical isomorphism $c\colon I_{\mathrm{t}} \to \prod_{l \neq p} \mathbb{Z}_l(1)$ [GMS03, p. 17]. To explain this isomorphism, let $g \in I_{\mathrm{t}}$ and let $\pi^{1/n}$ be an n-th root of a uniformizing parameter $\pi \in A$ with n not divisible by p. Then the image of $c(g)$ in $\boldsymbol{\mu}_n$ is $g(\pi^{1/n})/\pi^{1/n}$.

Proposition 5.6. *Let A be a henselian discrete valuation ring. Then the sequence* (5.2) *is split.*

The proposition extends Lemma 7.6 of [GMS03], where A is assumed to be complete.

Proof. Because we need the ideas from the proof, we will repeat Serre's argument. Set $K = k(\eta)$ and $\overline{K} = k(\overline{\eta})$. Set $K_{\mathrm{t}} = \overline{K}^{I_{\mathrm{t}}}$: the maximal tamely ramified extension of K. Let π be a uniformizing parameter in A. Then, for each non-negative integer n not divisible by p, choose an n-th root π_n of π in $\overline{K}$ such that $\pi_{nm}^m = \pi_n$. Set $K_{\mathrm{ram}} \stackrel{\mathrm{def}}{=} K[\pi_n]_{(p \nmid n)}$. Then K_{ram} is totally and tamely ramified over K. Moreover any $K_{\mathrm{t}} = K_{\mathrm{ram}} K_{\mathrm{unr}}$. It follows that $\operatorname{Gal}(k(s))$ may be identified with the subgroup of elements $g \in \operatorname{Gal}(K)_{\mathrm{t}}$ fixing each of the π_n; cf. [Del80]. This splits the sequence (5.3).

Now, in [GMS03], Serre extends this splitting non-canonically to a splitting of (5.2) as follows. Since $k(s)$ has characteristic p, the p-cohomological dimension of $\operatorname{Gal}(k(s))$

is ≤ 1; see [Ser02]. Consequently, any homomorphism $\mathrm{Gal}(k(s)) \to \mathrm{Gal}(K)_{\mathrm{t}}$ can be lifted to $\mathrm{Gal}(K)$. ◻

While the splitting of (5.3) is not canonical, we need to know that it is possible to split two such sequences, associated with henselian discrete valuation rings $A \subseteq B$, in a compatible way.

Proposition 5.7. *Let $A \subseteq B$ be an extension of henselian discrete valuation rings, such that a uniformizing parameter for A is also a uniformizing parameter for B. Then there exist maps $\sigma_B\colon \mathrm{Gal}(k(s_B)) \to \mathrm{Gal}(k(\eta_B))_{\mathrm{t}}$ (resp. $\sigma_A\colon \mathrm{Gal}(k(s_A)) \to \mathrm{Gal}(k(\eta_A))_{\mathrm{t}}$) splitting the tame inertia exact sequence* (5.3) *for B* (*resp. A*) *and such that the diagram*

$$\begin{array}{ccc} \mathrm{Gal}(k(s_B)) & \xrightarrow{\sigma_B} & \mathrm{Gal}(k(\eta_B))_{\mathrm{t}} \\ \downarrow & & \downarrow \\ \mathrm{Gal}(k(s_A)) & \xrightarrow{\sigma_A} & \mathrm{Gal}(k(\eta_A))_{\mathrm{t}} \end{array}$$

with vertical morphisms given by restriction commutes.

Proof. Let $\pi \in A$ be a uniformizing parameter for A, and hence for B. For each n not divisible by $p = \mathrm{char}(k(s_A))$, choose an n-th root π_n of π in $k(\overline{\eta}_B)$. Now, set $\sigma_B(k(s_B)) = \{g \in \mathrm{Gal}(k(\eta_B))_{\mathrm{t}} \mid g(\pi_n) = \pi_n \text{ for all } n\}$ and similarly for A. By the proof of Proposition 5.6, this defines splitting of the tame inertia sequences. Moreover, these splittings lift to splittings of the inertia exact sequence. ◻

Remark 5.8. By the proof of Proposition 5.6, the splittings σ_B and σ_A in Proposition 5.7 can be lifted to maps $\tilde{\sigma}_B\colon \mathrm{Gal}(k(s_B)) \to \mathrm{Gal}(k(\eta_B))$ (resp. $\tilde{\sigma}_A\colon \mathrm{Gal}(k(s_A)) \to \mathrm{Gal}(k(\eta_a)))$. However, since these liftings are non-canonical it is not clear that $\tilde{\sigma}_B$ and $\tilde{\sigma}_A$ can be chosen compatibly.

5.4. Tame gerbes and splittings

The following result is certainly well known; for the sake of completeness we supply a short proof.

Proposition 5.9. *Let $\mathcal{X} \to S$ be an étale gerbe over a henselian trait, with closed point s. Denote by $i\colon s \to S$ the inclusion of the closed point and by $\mathrm{sp}\colon S \to s$ the specialization map. Then the restriction map*

$$i^*\colon \mathcal{X}(S) \to \mathcal{X}(s)$$

induces an equivalence of categories with quasi-inverse given by

$$\mathrm{sp}^*\colon \mathcal{X}(s) \to \mathcal{X}(S).$$

Proof. Since the composite $s \to S \overset{\mathrm{sp}}{\to} s$ is an auto-equivalence and $\mathcal{X}$ is obtained by pullback from $\mathcal{X}_s$, it suffices to show that the functor $i^*\colon \mathcal{X}(S) \to \mathcal{X}(s)$ is faithful. For this, suppose $\xi_i\colon S \to \mathcal{X}$, $i = 1, 2$, are two objects of $\mathcal{X}(S)$. Then the sheaf $\mathrm{Hom}(\xi_1, \xi_2)$ is étale over S. Since S is henselian, it follows that the sections of $\mathrm{Hom}(\xi_1, \xi_2)$ over S are isomorphic (via restriction) to the sections over s. Thus $i^*\colon \mathcal{X}(S) \to \mathcal{X}(s)$ is fully faithful. □

A Deligne–Mumford stack $\mathcal{X} \to S$ is *tame* if, for every geometric point $\xi\colon \operatorname{Spec}\Omega \to \mathcal{X}$, the order of the automorphism group $\mathrm{Aut}_{\operatorname{Spec}\Omega}(\xi)$ is prime to the characteristic of Ω. For tame gerbes over a henselian discrete valuation ring, we have the following analogue of the splitting in Proposition 5.7.

Theorem 5.10. *Let $h\colon \operatorname{Spec} B \to \operatorname{Spec} A$ be the morphism of henselian traits induced by an inclusion $A \hookrightarrow B$ of henselian discrete valuation rings (here we assume that a uniformizing parameter for A is sent to a uniformizing parameter for B). Let $\mathcal{X}$ be a tame étale gerbe over* Spec A*. Write $j_B\colon \{\eta_B\} \to \operatorname{Spec} B$ (resp. $j_A\colon \{\eta_A\} \to \operatorname{Spec} A$) for the inclusion of the generic points. Then there exist functors*

$$\tau_A\colon \mathcal{X}(k(\eta_A)) \to \mathcal{X}(A) \quad \textit{and} \quad \tau_B\colon \mathcal{X}(k(\eta_B)) \to \mathcal{X}(B)$$

such that the diagram

$$\begin{array}{ccccc} \mathcal{X}(A) & \xrightarrow{j_A^*} & \mathcal{X}(k(\eta_A)) & \xrightarrow{\tau_A} & \mathcal{X}(A) \\ \downarrow{\scriptstyle h^*} & & \downarrow{\scriptstyle h^*} & & \downarrow{\scriptstyle h^*} \\ \mathcal{X}(B) & \xrightarrow{j_B^*} & \mathcal{X}(k(\eta_B)) & \xrightarrow{\tau_B} & \mathcal{X}(B) \end{array}$$

commutes (up to natural isomorphism) and the horizontal composites are isomorphic to the identity.

Proof. Since $\mathcal{X}$ is an étale gerbe, there is an extension E as in (5.1) with $G = \mathrm{Gal}(k(s_A))$ such that $\mathcal{X}$ is the pullback of $\mathcal{X}_E$ to the big étale site over S_A. Since $\mathcal{X}$ is tame, the band, i.e., the group F in (5.1), has order prime to $\operatorname{char} k(s_A)$.

Now, pick splittings σ_B and σ_A compatibly, as in Proposition 5.7.

We define a functor $\tau_B\colon \mathcal{X}(k(\eta_B)) \to \mathcal{X}(B)$ as follows. Using Proposition 5.5 we can identify $\mathcal{X}(k(\eta_B))$ with the category of sections $s\colon \mathrm{Gal}(k(\eta_B)) \to E$. Given such a section s, the tameness of E implies that $s(I^{\mathrm{w}}) = 1$. Therefore, s induces a map $\mathrm{Gal}(k(\eta_B))_{\mathrm{t}} \to \mathrm{Gal}(k(s_B))$, which we will also denote by the symbol s. Let $\tau_B(s)$ denote the section $s \circ \sigma_B\colon \mathrm{Gal}(k(s_B)) \to E$. This defines τ_B on the objects in $\mathcal{X}(k(\eta_B))$. If we define τ_A in the same way, it is clear that the diagram above commutes on objects. We define τ_B (resp. σ_A) on morphisms by setting $\tau_B(f) = f$ (and similarly for A). We leave the rest of the verification to the reader. □

5.5. *Genericity*

Theorem 5.11. *Let R be a discrete valuation ring, $S = \operatorname{Spec} R$ and $\mathcal{X} \to S$ a tame étale gerbe. Then $\mathrm{ed}_{k(s)} \mathcal{X}_s \le \mathrm{ed}_{k(\eta)} \mathcal{X}_\eta$, where s is the closed point of S and η is the generic point.*

Proof. We may assume without loss of generality that R is complete. Indeed, otherwise replace R with its completion at s. The field $k(s)$ does not change, but $k(\eta)$ is replaced by a field extension. By Proposition 2.8, the essential dimension of $\mathcal{X}_{k(\eta)}$ does not increase.

If R is equicharacteristic, then by Cohen's structure theorem, $R = k[[t]]$ with $k = k(s)$. If not, denote by $W(k(s))$ the unique complete discrete valuation ring with residue field $k(s)$ and uniformizing parameter p. This is called a Cohen ring of $k(s)$ in [Gro64, 19.8]. If $k(s)$ is perfect then $W(k(s))$ is the ring of Witt vectors of $k(s)$, but this is not true in general, and $W(k(s))$ is only determined up to a non-canonical isomorphism. By [Gro64, Théorème 19.8.6], there is a homomorphism $W(k(s)) \to R$ inducing the identity on $k(s)$. Since $\mathcal{X}$ is pulled back from $k(s)$ via the specialization map, we can replace R by $W(k(s))$.

Now suppose $b\colon \operatorname{Spec} L \to \mathcal{X}_s$ is a morphism from the spectrum of a field with $\operatorname{ed}_{k(s)} b = \operatorname{tr deg}_{k(s)} L = \operatorname{ed} \mathcal{X}_s$. (Such a morphism exists because $\operatorname{ed} \mathcal{X}_s$ is finite.) Set $B := L[[t]]$ if R is equicharacteristic and $B := W(L)$ otherwise. In either case, B is a complete discrete valuation ring with residue field L. In the first case we have a canonical embedding $R = k[[t]] \subseteq L[[t]] = B$; in the second case, again by [Gro64, Théorème 19.8.6] (due to Cohen), we have a lifting $R = W(k(s)) \to W(L) = B$ of the embedding $k(s) \subseteq L$, which is easily seen to be injective. Therefore there is a unique morphism $\beta = b \circ \operatorname{sp}\colon S_B \to \mathcal{X}$ whose specialization to the closed point of B coincides with ξ.

Suppose there is a subfield M of $k(\eta_B)$ containing $k(\eta_R)$ such that the following conditions hold:

(1) the restriction $j_B^*\beta$ of β to $k(\eta_B)$ factors through M,
(2) $\operatorname{tr deg}_{k(\eta_R)} M < \operatorname{ed}_{k(s)} b$.

Complete M with respect to the discrete valuation induced from $k(\eta_B)$ and call the resulting complete discrete valuation ring A. It follows that there is a class α in $\mathcal{X}(k(\eta_A))$ whose restriction to $k(\eta_B)$ coincides with $j_B^*\beta$. But then, by Theorem 5.10, we have $\beta = h_*\sigma_A(\alpha)$. This implies that $b\colon \operatorname{Spec} L \to \mathcal{X}_s$ factors through the special fiber of A. Since the transcendence degree of $k(s_A)$ over $k(s)$ is less than $\operatorname{ed}_{k(s)} b$, this is a contradiction. □

Corollary 5.12. *Let R be an equicharacteristic complete discrete valuation ring and $\mathcal{X} \to \operatorname{Spec} R$ be a tame étale gerbe. Then*

$$\operatorname{ed}_{k(s)} \mathcal{X}_s = \operatorname{ed}_{k(\eta)} \mathcal{X}_\eta,$$

where s denotes the closed point of Spec R *and η denotes the generic point.*

Proof. Set $k = k(s)$. Since R is equicharacteristic, we have $R = k[[t]]$ and $\mathcal{X}_{k(\eta)}$ is the pullback to $k(\eta)$ of $\mathcal{X}_{k(s)}$ via the inclusion of k in $k((t))$. Therefore $\operatorname{ed}_{k(s)} \mathcal{X}_{k(s)} \geq \operatorname{ed}_{k(\eta)} \mathcal{X}_{k(\eta)}$. The opposite inequality is given by Theorem 5.11. □

Theorem 5.13. *Suppose that $\mathcal{X}$ is an étale gerbe over a smooth scheme $\mathbf{X}$ locally of finite type over a perfect field k. Let K be an extension of k, and $\xi \in \mathcal{X}(\operatorname{Spec} K)$. Then*

$$\operatorname{ed} \xi \leq \operatorname{ed}_{k(\mathbf{X})} \mathcal{X}_{k(\mathbf{X})} + \dim \mathbf{X} - \operatorname{codim}_{\mathcal{X}} \xi.$$

Proof. We proceed by induction on $\operatorname{codim}\xi$. If $\operatorname{codim}_{\mathcal{X}}\xi = 0$, then the morphism $\xi\colon \operatorname{Spec} K \to \mathcal{X}$ is dominant. Hence ξ factors through $\mathcal{X}_{k(\mathbf{X})}$, and the result is obvious.

Assume $\operatorname{codim}_{\mathcal{X}}\xi > 0$. Let $\mathbf{Y}$ be the closure of the image of $\operatorname{Spec} K$ in $\mathbf{X}$. Since we are assuming that k is perfect, $\mathbf{Y}$ is generically smooth over $\operatorname{Spec} k$. By restricting to a neighborhood of the generic point of $\mathbf{Y}$, we may assume that $\mathbf{Y}$ is contained in a smooth hypersurface $\mathbf{X}'$ of $\mathbf{X}$. Denote by $\mathcal{Y}$ and $\mathcal{X}'$ the inverse images in $\mathcal{X}$ of $\mathbf{Y}$ and $\mathcal{X}'$ respectively. Set $R = \mathcal{O}_{\mathbf{X},\mathbf{Y}}$ and denote the pullback of $\mathcal{X}$ to R also by $\mathcal{X}$. Then we can apply Theorem 5.11 to the gerbe $\mathcal{X}_R \to \operatorname{Spec} R$ and conclude that

$$\operatorname{ed}_{k(\mathbf{X}')}\mathcal{X}'_{k(\mathbf{X}')} \le \operatorname{ed}_{k(\mathbf{X})}\mathcal{X}_{k(\mathbf{X})}.$$

Using the inductive hypothesis we have

$$\begin{aligned}\operatorname{ed}\xi &\le \operatorname{ed}_{k(\mathbf{X}')}\mathcal{X}_{k(\mathbf{X}')} + \dim\mathbf{X}' - \operatorname{codim}_{\mathcal{X}'}\xi \le \operatorname{ed}_{k(\mathbf{X})}\mathcal{X}_{k(\mathbf{X})} + \dim\mathbf{X} - 1 - \operatorname{codim}_{\mathcal{X}'}\xi\\ &\le \operatorname{ed}_{k(\mathbf{X})}\mathcal{X}_{k(\mathbf{X})} + \dim\mathbf{X} - \operatorname{codim}_{\mathcal{X}}\xi. \qquad\square\end{aligned}$$

6. A genericity theorem for a smooth Deligne–Mumford stack

It is easy to see that Theorem 5.13 fails if $\mathcal{X}$ is not assumed to be a gerbe. In this section we will use Theorem 5.13 to prove the following weaker result for a wider class of Deligne–Mumford stacks.

Recall that a Deligne–Mumford stack $\mathcal{X}$ over a field k is *tame* if the order of the automorphism group of any object of $\mathcal{X}$ over an algebraically closed field is prime to the characteristic of k.

Theorem 6.1. *Let $\mathcal{X}$ be a smooth integral tame Deligne–Mumford stack locally of finite type over a perfect field k. Then*

$$\operatorname{ed}\mathcal{X} = \operatorname{ed}_{k(\mathbf{X})}\mathcal{X}_{k(\mathbf{X})} + \dim\mathcal{X}.$$

Here the dimension of $\mathcal{X}$ is the dimension of the moduli space of any non-empty open substack of $\mathcal{X}$ with finite inertia.

Before proceeding with the proof, we record two immediate corollaries.

Corollary 6.2. *For $\mathcal{X}$ as above, if $\mathcal{U}$ is an open dense substack, then* $\operatorname{ed}_k\mathcal{M} = \operatorname{ed}_k\mathcal{U}$. $\square$

Corollary 6.3. *If the conditions of Theorem* 6.1 *are satisfied, and the generic object of $\mathcal{X}$ has no non-trivial automorphisms (i.e., $\mathcal{X}$ is an orbifold, in the topologists' terminology), then* $\operatorname{ed}_k\mathcal{X} = \dim\mathcal{X}$.

Proof. Here the generic gerbe $\mathcal{X}_K$ is a scheme, so $\mathcal{X}_K$ is the 0-dimensional scheme $\operatorname{Spec}(K)$, and $\operatorname{ed}_K\mathcal{X}_K = \dim_K \operatorname{Spec}(K) = 0$. $\square$

Proof of Theorem 6.1. The inequality $\operatorname{ed}\mathcal{X} \ge \operatorname{ed}_{k(\mathbf{X})}\mathcal{X}_{k(\mathbf{X})} + \dim\mathcal{X}$ is obvious: so we only need to show that

$$\operatorname{ed}\xi \le \operatorname{ed}_{k(\mathbf{X})}\mathcal{X}_{k(\mathbf{X})} + \dim\mathcal{X} \tag{6.1}$$

for any field extension L of k and any object ξ of $\mathcal{X}(L)$.

First of all, let us reduce the general result to the case that $\mathcal{X}$ has finite inertia. The reduction is immediate from the following lemma, essentially due to Keel and Mori.

Lemma 6.4 (Keel–Mori). *There exists an integral Deligne–Mumford stack with finite inertia $\mathcal{X}'$, together with an étale representable morphism of finite type $\mathcal{X}' \to \mathcal{X}$, and a factorization* $\operatorname{Spec} L \to \mathcal{X}' \to \mathcal{X}$ *of the morphism* $\operatorname{Spec} L \to \mathcal{X}$ *corresponding to* ξ.

Proof. We follow an argument due to B. Conrad. By [Con, Lemma 2.2] there exist

(i) an étale representable morphism $\mathcal{W} \to \mathcal{X}$ such that every morphism $\operatorname{Spec} L \to \mathcal{X}$, where L is a field, lifts to $\operatorname{Spec} L \to \mathcal{W}$, and
(ii) a finite flat representable map $Z \to \mathcal{W}$, where Z is a scheme.

Condition (ii) implies that $\mathcal{W}$ is a quotient of Z by a finite flat equivalence relation $Z \times_{\mathcal{W}} Z \rightrightarrows Z$, which in particular tells us that $\mathcal{W}$ has finite inertia. We can now take $\mathcal{X}'$ to be a connected component of $\mathcal{W}$ containing a lifting $\operatorname{Spec} L \to \mathcal{W}$ of $\operatorname{Spec} L \to \mathcal{X}$. □

Suppose that we have proved the inequality (6.1) whenever $\mathcal{X}$ has finite inertia. If we denote by ξ' the object of $\mathcal{X}'$ corresponding to a lifting $\operatorname{Spec} L \to \mathcal{X}'$, we have

$$\operatorname{ed} \xi \leq \operatorname{ed} \xi' \leq \operatorname{ed}_{k(\mathbf{X}')} \mathcal{X}'_{k(\mathbf{X}')}.$$

On the other hand, the morphism $\mathcal{X}'_{k(\mathbf{X}')} \to \mathcal{X}_{k(\mathbf{X})}$ induced by the étale representable morphism $\mathcal{X}' \to \mathcal{X}$ is representable with fibers of dimension 0, hence

$$\operatorname{ed}_{k(\mathbf{X}')} \mathcal{X}'_{k(\mathbf{X}')} = \operatorname{ed}_{k(\mathbf{X})} \mathcal{X}'_{k(\mathbf{X}')} \leq \operatorname{ed}_{k(\mathbf{X})} \mathcal{X}_{k(\mathbf{X})}$$

by Theorem 3.2 (the first equality follows immediately from the fact that the extension $k(\mathbf{X}) \subseteq k(\mathbf{X}')$ is finite).

So, in order to prove the inequality (6.1) we may assume that $\mathcal{X}$ has finite inertia. Denote by $\mathbf{Y} \subseteq \mathbf{X}$ the closure of the image of the composite $\operatorname{Spec} L \to \mathcal{X} \to \mathbf{X}$, where $\operatorname{Spec} L \to \mathcal{X}$ corresponds to ξ, and denote by $\mathcal{Y}$ the reduced inverse image of $\mathbf{Y}$ in $\mathcal{X}$. Since k is perfect, $\mathcal{Y}$ is generically smooth; by restricting to a neighborhood of the generic point of $\mathbf{Y}$ we may assume that $\mathcal{Y}$ is smooth.

Denote by $\mathcal{N} \to \mathcal{Y}$ the normal bundle of $\mathcal{Y}$ in $\mathcal{X}$. Consider the deformation to the normal bundle $\phi\colon \mathcal{M} \to \mathbb{P}^1_k$ for the embedding $\mathcal{Y} \subseteq \mathcal{X}$. This is a smooth morphism such that $\phi^{-1}\mathbb{A}^1_k = \mathcal{X} \times_{\operatorname{Spec} k} \mathbb{A}^1_k$ and $\phi^{-1}(\infty) = \mathcal{N}$, obtained as an open substack of the blow-up of $\mathcal{X} \times_{\operatorname{Spec} k} \mathbb{P}^1_k$ along $\mathcal{Y} \times \{\infty\}$ (the well-known construction, explained for example in [Ful98, Chapter 5], generalizes immediately to algebraic stacks). Denote by $\mathcal{M}^0$ the open substack whose geometric points are the geometric points of $\mathcal{M}$ with stabilizer of minimal order (this is well defined because $\mathcal{M}$ has finite inertia).

We claim that $\mathcal{M}^0 \cap \mathcal{N} \neq \emptyset$. This would be evident if $\mathcal{X}$ were a quotient stack $[V/G]$, where G is a finite group of order not divisible by the characteristic of k, acting linearly on a vector space V, and $\mathcal{Y}$ were of the form $[X/G]$, where W is a G-invariant linear subspace of V. However, étale-locally on $\mathbf{X}$ every tame Deligne–Mumford stack is a quotient $[X/G]$, where G is a finite group of order not divisible by the characteristic of k (see, e.g., [AV02, Lemma 2.2.3]). Since G is tame and X is smooth, it is well known that étale-locally on $\mathbf{X}$, the stack $\mathcal{X}$ has the desired form, and this is enough to prove the claim.

Set $\mathcal{N}^0 \overset{\text{def}}{=} \mathcal{M}^0 \cap \mathcal{N}$. The object ξ corresponds to a dominant morphism $\operatorname{Spec} L \to \mathcal{Y}$. The pullback $\mathcal{N} \times_{\mathcal{Y}} \operatorname{Spec} K$ is a vector bundle V over $\operatorname{Spec} L$, and the inverse image $\mathcal{N}^0 \times_{\mathcal{Y}} \operatorname{Spec} L$ of $\mathcal{N}^0$ is not empty. We may assume that L is infinite; otherwise $\operatorname{ed} \xi = 0$ and there is nothing to prove. Assuming that L is infinite, $\mathcal{N}^0 \times_{\mathcal{Y}} \operatorname{Spec} L$ has an L-rational point, so there is a lifting $\operatorname{Spec} L \to \mathcal{N}^0$ of $\operatorname{Spec} L \to \mathcal{Y}$, corresponding to an object η of $\mathcal{N}^0(\operatorname{Spec} L)$. Clearly the essential dimension of ξ as an object of $\mathcal{X}$ is the same as its essential dimension as an object of $\mathcal{Y}$, and $\operatorname{ed} \xi \le \operatorname{ed} \eta$. Let us apply Theorem 5.13 to the gerbe $\mathcal{M}^0$. The function field of the moduli space $\mathbf{M}$ of $\mathcal{M}$ is $k(\mathbf{X})(t)$, and its generic gerbe is $\mathcal{X}_{k(\mathbf{X})(t)}$; by Proposition 2.8, we have $\operatorname{ed}_{k(\mathbf{X})(t)} \mathcal{X}_{k(\mathbf{X})(t)} \le \operatorname{ed}_{k(\mathbf{X})} \mathcal{X}_{k(\mathbf{X})}$. The composite $\operatorname{Spec} L \to \mathcal{N}^0 \subseteq \mathcal{M}^0$ has codimension at least 1, hence we obtain

$$\operatorname{ed} \xi < \operatorname{ed}_{k(\mathbf{X})(t)} \mathcal{X}_{k(\mathbf{X})(t)} + \dim \mathbf{M} \le \operatorname{ed}_{k(\mathbf{X})} \mathcal{X}_{k(\mathbf{X})} + \dim \mathbf{X} + 1.$$

This concludes the proof. □

Example 6.5. The following examples show that Corollary 6.3 (and thus Corollary 6.2 and Theorem 6.1) fail for more general algebraic stacks, such as (a) singular Deligne–Mumford stacks, (b) non-Deligne–Mumford stacks, including quotient stacks of the form $[W/G]$, where W is a smooth complex affine variety with an action of a connected complex reductive linear algebraic group G acting on W.

(a) Let $r, n \ge 2$ be integers. Assume that the characteristic of k is prime to r. Let $W \subseteq \mathbb{A}^n$ be the Fermat hypersurface defined by the equation $x_1^r + \cdots + x_n^r = 0$ and $\Delta \subset \mathbb{A}^n$ be the union of the coordinate hyperplanes defined by $x_i = 0$ for $i = 1, \dots, n$. The group $G := \boldsymbol{\mu}_r^n$ acts on $\mathbb{A}^n$ via the formula

$$(s_1, \dots, s_n)(x_1, \dots, x_n) = (s_1 x_1, \dots, s_n x_n),$$

leaving W and Δ invariant. Let $\mathcal{X} := [W/G]$. Since the G-action on $W \smallsetminus \Delta$ is free, $\mathcal{X}$ is generically an affine scheme of dimension $n - 1$. On the other hand, $[\{0\}/G] \simeq \mathcal{B}_k \boldsymbol{\mu}_r^n$ is a closed substack of $\mathcal{X}$ of essential dimension n; hence, $\operatorname{ed}(\mathcal{X}) \ge n$.

(b) Consider the action of $G = \mathrm{GL}_n$ on the affine space M of all $n \times n$-matrices by multiplication on the left. Since G has a dense orbit, and the stabilizer of a non-singular matrix in M is trivial, we see that $[M/G]$ is generically a scheme of dimension 0. On the other hand, let Y be the locus of matrices of rank $n - 1$, which is a locally closed subscheme of M. There is a surjective GL_n-equivariant morphism $Y \to \mathbb{P}^{n-1}$, sending each matrix of rank $n - 1$ to its kernel, which induces a morphism $[Y/G] \to \mathbb{P}^{n-1}$. If L is an extension of $\mathbb{C}$, every L-valued point of $\mathbb{P}^{n-1}$ lifts to an L-valued point of Y. Hence,

$$\operatorname{ed} [M/G] \ge \operatorname{ed} [Y/G] \ge n - 1.$$

As an aside, we remark that a similar argument with Y replaced by the locus of matrices of rank r, shows that the essential dimension of $[M/G]$ is in fact the maximum of the dimensions of the Grassmannians of r-planes in $\mathbb{C}^n$, as r ranges between 1 and $n - 1$, which is $n^2/4$ if n is even, and $(n^2 - 1)/4$ if n is odd.

Question 6.6. Under what hypotheses does the genericity theorem hold? Let $\mathcal{X} \to \operatorname{Spec} k$ be an integral algebraic stack. Using the results of [LMB00, Chapter 11], one can define the generic gerbe $\mathcal{X}_K \to \operatorname{Spec} K$ of $\mathcal{X}$, which is an fppf gerbe over a field of finite transcendence degree over k. What conditions on $\mathcal{X}$ ensure the equality

$$\operatorname{ed}_k \mathcal{X} = \operatorname{ed}_K \mathcal{X}_K + \operatorname{tr\,deg}_k K\ ?$$

Smoothness seems necessary, as there are counterexamples even for Deligne–Mumford stacks with very mild singularities; see Example 6.5(a). We think that the best result that one can hope for is the following. Suppose that $\mathcal{X}$ is smooth with quasi-affine diagonal, and let $\xi \in \mathcal{X}(\operatorname{Spec} L)$ be a point. Assume that the automorphism group scheme of ξ over L is linearly reductive. Then $\operatorname{ed} \xi \leq \operatorname{ed}_K \mathcal{X}_K + \operatorname{tr\,deg}_k K$. In particular, if all the automorphism groups are linearly reductive, then $\operatorname{ed} \mathcal{X} = \operatorname{ed}_K \mathcal{X}_K + \operatorname{tr\,deg}_k K$. (Added in proof: This has recently been established; see [RV11, Theorem 1.2].)

7. The essential dimension of $\mathcal{M}_{g,n}$ for $(g, n) \neq (1, 0)$

Recall that the base field k is assumed to be of characteristic 0.

The assertion that $\operatorname{ed} \overline{\mathcal{M}}_{g,n} = \operatorname{ed} \mathcal{M}_{g,n}$ whenever $2g - 2 + n > 0$ is an immediate consequence of Corollary 6.2. Moreover, if $g \geq 3$, or $g = 2$ and $n \geq 1$, or $g = 1$ and $n \geq 2$, then

$$\operatorname{ed}_k \mathcal{M}_{g,n} = \operatorname{ed}_k \overline{\mathcal{M}}_{g,n} = 3g - 3 + n.$$

Indeed, in all these cases the automorphism group of a generic object of $\mathcal{M}_{g,n}$ is trivial, so the generic gerbe is trivial, and $\operatorname{ed} \mathcal{M}_{g,n} = \dim \mathcal{M}_{g,n}$ by Corollary 6.3.

The remaining cases of Theorem 1.2, with the exception of $(g, n) = (1, 0)$, are covered by the following proposition. The case where $(g, n) = (1, 0)$ requires a separate argument which will be carried out in the next section.

Proposition 7.1.

(a) $\operatorname{ed} \mathcal{M}_{0,1} = 2$,
(b) $\operatorname{ed} \mathcal{M}_{0,1} = \operatorname{ed} \mathcal{M}_{0,2} = 0$,
(c) $\operatorname{ed} \mathcal{M}_{1,1} = 2$,
(d) $\operatorname{ed} \mathcal{M}_{2,0} = 5$.

Proof. (a) Since $\mathcal{M}_{0,0} \simeq \mathcal{B}_k\mathrm{PGL}_2$, we have $\operatorname{ed} \mathcal{M}_{0,0} = \operatorname{ed} \mathrm{PGL}_2 = 2$, where the last inequality is proved in [Rei00, Lemma 9.4(c)] (the argument there is valid for any field k of characteristic $\neq 2$).

Alternative proof of (a): The inequality $\operatorname{ed} \mathcal{M}_{0,0} \leq 2$ holds because every smooth curve of genus 0 over a field K is a conic C in $\mathbb{P}^2_K$. After a change of coordinates in $\mathbb{P}^2_K$ we may assume that C is given by an equation of the form $ax^2 + by^2 + z^2 = 0$ for some $a, b \in K$, and hence descends to the field $k(a, b)$ of transcendence degree ≤ 2 over k. The opposite inequality follows from Tsen's theorem.

(b) A smooth curve C of genus 0 with one or two rational points over an extension K of k is isomorphic to $(\mathbb{P}^1_k, 0)$ or $(\mathbb{P}^1_k, 0, \infty)$. Hence, it is defined over k.

Alternative proof of (b): $\mathcal{M}_{0,2} = \mathcal{B}_k\mathbb{G}_\mathrm{m}$ and $\mathcal{M}_{0,1} = \mathcal{B}_k(\mathbb{G}_\mathrm{m} \ltimes \mathbb{G}_\mathrm{a})$, and the groups $\mathbb{G}_\mathrm{m}$ and $\mathbb{G}_\mathrm{m} \ltimes \mathbb{G}_\mathrm{a}$ are special (and hence have essential dimension 0).

(c) Let $\mathcal{M}_{1,1} \to \mathbb{A}^1_k$ denote the map given by the j-invariant and let $\mathcal{X}$ denote the pullback of $\mathcal{M}_{1,1}$ to the generic point $\operatorname{Spec} k(j)$ of $\mathbb{A}^1$. Then $\mathcal{X}$ is banded by $\boldsymbol{\mu}_2$ and is neutral by [Sil86, Proposition 1.4 (c)], and so $\operatorname{ed}_k \mathcal{X} = \operatorname{ed}_{k(j)} \mathcal{X}+1 = \operatorname{ed} \mathcal{B}_{k(j)}\boldsymbol{\mu}_2+1 = 2$.

(d) is a special case of Theorem 7.2 below, since $\mathcal{H}_2 = \mathcal{M}_{2,0}$. □

Let $\mathcal{H}_g$ denote the stack of hyperelliptic curves of genus $g > 1$ over a field k of characteristic 0. This must be defined with some care; defining a family of hyperelliptic curves as a family $C \to S$ in $\mathcal{M}_{g,0}$ whose fiber are hyperelliptic curves will not yield an algebraic stack. There are two possibilities.

(a) One can define $\mathcal{H}_g$ as the closed reduced substack of $\mathcal{M}_g$ whose geometric points corresponds to hyperelliptic curves.
(b) As in [AV04], an object of $\mathcal{H}_g$ can be defined as two morphisms of schemes $C \to P \to S$, where $P \to S$ is a Brauer–Severi, $C \to P$ is a flat finite finitely presented morphism of constant degree 2, and the composite $C \to S$ is a smooth morphism whose fibers are connected curves of constant genus g.

We adopt the second definition; $\mathcal{H}_g$ is then a smooth algebraic stack of finite type over k (this is shown in [AV04]). Furthermore, there is a natural morphism $\mathcal{H}_g \to \mathcal{M}_{g,0}$, which sends $C \to P \to S$ to the composite $C \to S$. This morphism is easily seen to be a closed embedding. Hence the two stacks defined above are in fact isomorphic.

Theorem 7.2. $\operatorname{ed} \mathcal{H}_g = \begin{cases} 2g & \text{if } g \geq 3 \text{ is odd,} \\ 2g+1 & \text{if } g \geq 2 \text{ is even.} \end{cases}$

Proof. Denote by $\mathbf{H}_g$ the moduli space of $\mathcal{H}_g$; the dimension of $\mathcal{H}_g$ is $2g-1$. Let K be the field of rational functions on $\mathbf{H}_g$, and denote by $(\mathcal{H}_g)_K \overset{\text{def}}{=} \operatorname{Spec} K \times_{\mathbf{H}_g} \mathcal{H}_g$ the generic gerbe of $\mathcal{H}_g$. From Theorem 6.1 we have

$$\operatorname{ed} \mathcal{H}_g = 2g - 1 + \operatorname{ed}_K (\mathcal{H}_g)_K,$$

so we need to show that $\operatorname{ed}_K (\mathcal{H}_g)_K$ is 1 if g is odd, and 2 if g is even. For this we need some standard facts about stacks of hyperelliptic curves, which we will now recall.

Let $\mathcal{D}_g$ be the stack over k whose objects over a k-scheme S are pairs $(P \to S, \Delta)$, where $P \to S$ is a conic bundle (that is, a Brauer–Severi scheme of relative dimension 1), and $\Delta \subseteq P$ is a Cartier divisor that is étale of degree $2g+2$ over S. Let $C \xrightarrow{\pi} P \to S$ be an object of $\mathcal{H}_g$; denote by $\Delta \subseteq P$ the ramification locus of π. Sending $C \xrightarrow{\pi} P \to S$ to $(P \to S, \Delta)$ gives a morphism $\mathcal{H}_g \to \mathcal{D}_g$. Recall the usual description of ramified double covers: if we split $\pi_* \mathcal{O}_C$ as $\mathcal{O}_P \oplus L$, where L is the part of trace 0, then multiplication yields an isomorphism $L^{\otimes 2} \simeq \mathcal{O}_P(-\Delta)$. Conversely, given an object $(P \to S, \Delta)$ of $\mathcal{D}_g(S)$ and a line bundle L on P, with an isomorphism $L^{\otimes 2} \simeq \mathcal{O}_P(-\Delta)$, the direct sum $\mathcal{O}_P \oplus L$ has an algebra structure, whose relative spectrum is a smooth curve $C \to S$ with a flat map $C \to P$ of degree 2.

The morphism $\mathcal{H}_g \to \mathbf{H}_g$ factors through $\mathcal{D}_g$, and the morphism $\mathcal{D}_g \to \mathbf{H}_g$ is an isomorphism over the non-empty locus of divisors on a curve of genus 0 with no non-trivial automorphisms (this is non-empty because $g \geq 2$, hence $2g + 2 \geq 5$). Denote by $(P \to \operatorname{Spec} K, \Delta)$ the object of $\mathcal{D}_g(\operatorname{Spec} K)$ corresponding to the generic point $\operatorname{Spec} K \to \mathbf{H}_g$. It is well known that $P(K) = \emptyset$; we give a proof for lack of a suitable reference.

Let C be a conic without rational points defined over some extension L of k. Let V be the L-vector space $\mathrm{H}^0(C, \omega_{C/L}^{-(g+1)})$; denote the function field of V by $F = L(V)$. Then there is a tautological section σ of $\mathrm{H}^0(C_F, \omega_{C_F/F}^{-(g+1)}) = \mathrm{H}^0(C, \omega_{C/L}^{-(g+1)}) \otimes_L F$. Note that $C_F(F) = \emptyset$, because the extension $L \subseteq F$ is purely transcendental. The zero scheme of σ is a divisor on C_F that is étale over $\operatorname{Spec} F$, and defines a morphism $C_F \to \mathcal{D}_g$. This morphism is clearly dominant: so $K \subseteq F$, and $C_F = P \times_{\operatorname{Spec} L} \operatorname{Spec} F$. Since $C_F(F) = \emptyset$ we have $P(K) = \emptyset$, as claimed.

By the description above, the gerbe $(\mathcal{H}_g)_K$ is the stack of square roots of $\mathcal{O}_P(-\Delta)$, which is banded by $\boldsymbol{\mu}_2$. When g is odd then there exists a line bundle of degree $g + 1$ on P, whose square is isomorphic to $\mathcal{O}_P(-\Delta)$; this gives a section of $(\mathcal{H}_g)_K$, which is therefore isomorphic to $\mathcal{B}_K \boldsymbol{\mu}_2$, whose essential dimension over K is 1. If g is even then such a section does not exist, and the stack is isomorphic to the stack of square roots of the relative dualizing sheaf $\omega_{P/K}$ (since $\mathcal{O}_{P/K}(-\Delta) \simeq \omega_{P/K}^{g+1}$, and $g + 1$ is odd), whose class in $\mathrm{H}^2(K, \boldsymbol{\mu}_2)$ represents the image in $\mathrm{H}^2(K, \boldsymbol{\mu}_2)$ of the class $[P]$ in $\mathrm{H}^1(K, \mathrm{PGL}_2)$ under the non-abelian boundary map $\mathrm{H}^1(K, \mathrm{PGL}_2) \to \mathrm{H}^2(K, \boldsymbol{\mu}_2)$. According to Theorem 1.3 its essential dimension is the index of $[P]$, which equals 2. ◻

The results above apply to more than stable curves. Assume that we are in the stable range $2g - 2 + n > 0$. Denote by $\mathfrak{M}_{g,n}$ the stack of all reduced n-pointed local complete intersection curves of genus g. This is the algebraic stack over $\operatorname{Spec} k$ whose objects over a k-scheme T are finitely presented proper flat morphisms $C \to T$, where C is an algebraic space, whose geometric fibers are connected reduced local complete intersection curves of genus g, together with n sections $T \to C$ whose images are contained in the smooth locus of $C \to T$. We do not require the sections to be disjoint.

The stack $\mathfrak{M}_{g,n}$ contains $\mathcal{M}_{g,n}$ as an open substack. By standard results in deformation theory, every reduced local complete intersection curve is unobstructed, and is a limit of smooth curves. Furthermore there is no obstruction to extending the sections, since these map into the smooth locus. Therefore $\mathfrak{M}_{g,n}$ is smooth and connected, and $\mathcal{M}_{g,n}$ is dense in $\mathfrak{M}_{g,n}$. However, the stack $\mathfrak{M}_{g,n}$ is very large (it is certainly not of finite type), and in fact it is very easy to see that its essential dimension is infinite. However, consider the open substack $\mathfrak{M}_{g,n}^{\mathrm{fin}}$ consisting of objects whose automorphism group is finite. Then $\mathfrak{M}_{g,n}^{\mathrm{fin}}$ is a Deligne–Mumford stack, and Theorem 6.1 applies to it. Thus we get the following strengthened form of Theorem 1.2 (under the assumption that $2g - 2 + n > 0$).

Theorem 7.3. *If $2g - 2 + n > 0$ and the characteristic of k is 0, then*

$$\operatorname{ed} \mathfrak{M}_{g,n}^{\mathrm{fin}} = \begin{cases} 2 & \text{if } (g, n) = (1, 1), \\ 5 & \text{if } (g, n) = (2, 0), \\ 3g - 3 + n & \text{otherwise.} \end{cases}$$

It is not hard to show that $\mathfrak{M}_{g,n}^{\mathrm{fin}}$ does not have finite inertia.

8. Tate curves and the essential dimension of $\mathcal{M}_{1,0}$

In this section we will finish the proof of Theorem 1.2 by showing that $\operatorname{ed}\mathcal{M}_{1,0} = +\infty$.

We remark that the moduli stack $\mathcal{M}_{1,0}$ of genus 1 curves should not be confused with the moduli stack $\mathcal{M}_{1,1}$ of elliptic curves. The objects of $\mathcal{M}_{1,0}$ are torsors for elliptic curves, whereas the objects of $\mathcal{M}_{1,1}$ are elliptic curves themselves. The stack $\mathcal{M}_{1,1}$ is Deligne–Mumford and, as we saw in the last section, its essential dimension is 2. The stack $\mathcal{M}_{1,0}$ is not Deligne–Mumford, and we will now show that its essential dimension is ∞.

Let R be a complete discrete valuation ring with function field K and uniformizing parameter q. For simplicity, we will assume that $\operatorname{char} K = 0$. Let $E = E_q/K$ denote the Tate curve over K [Sil86, §4]. This is an elliptic curve over K with the property that, for every finite field extension L/K, $E(L) \cong L^*/q^{\mathbb{Z}}$. It follows that the kernel $E[n]$ of multiplication by an integer $n > 0$ fits canonically into a short exact sequence

$$0 \to \boldsymbol{\mu}_n \to E[n] \to \mathbb{Z}/n \to 0. \tag{8.1}$$

Let $\partial\colon \mathrm{H}^0(K, \mathbb{Z}/n) \to \mathrm{H}^1(K, \boldsymbol{\mu}_n)$ denote the connecting homomorphism. Then it is well known (and easy to see) that $\partial(1) = q \in \mathrm{H}^1(K, \boldsymbol{\mu}_n) \cong K^*/(K^*)^n$.

Lemma 8.1. *Let $E = E_q$ be a Tate curve as above and let l be a prime integer not equal to* $\operatorname{char} R/q$*. Then, for any integer $n > 0$,*

$$\operatorname{ed} E[l^n] = l^n.$$

Proof. First observe that $E[l^n]$ admits an l^n-dimensional generically free representation $V = \operatorname{Ind}_{\boldsymbol{\mu}_{l^n}}^{E[l^n]} \chi$, over K, where $\chi : \boldsymbol{\mu}_{l^n} \to \mathbb{G}_{\mathrm{m}}$ is the tautological character. Thus,

$$\operatorname{ed} \mathcal{B}E[l^n] \le \dim V = l^n;$$

see [BR97, Theorem 3.1] or [BF03, Proposition 4.11].

It remains to show that

$$\operatorname{ed} E[l^n] \ge l^n. \tag{8.2}$$

Let $R' \stackrel{\text{def}}{=} R[1^{1/l^n}]$ with fraction field $K' = K[1^{1/l^n}]$. Since l is prime to the residue characteristic, R' is a complete discrete valuation ring, and the Tate curve E_q/K' is the pullback to K' of E_q/K. Since $\operatorname{ed}(E_q/K') \le \operatorname{ed}(E_q/K)$, it suffices to prove the lemma with K' replacing K. In other words, it suffices to prove the inequality (8.2) under the assumption that K contains the l^n-th roots of unity.

In that case, we can pick a primitive l^n-th root of unity ζ and write $\boldsymbol{\mu}_{l^n} = \mathbb{Z}/l^n$. Let $L = K(t)$ and consider the class $(t) \in \mathrm{H}^1(L, \boldsymbol{\mu}_{l^n}) = L^*/(L^*)^n$. It is not difficult to see that

$$\partial(t) = q \cup (t).$$

Since the map $\alpha \mapsto \alpha \cup (t)$ is injective by cohomological purity, the exponent of $q \cup (t)$ is l^n. Therefore $\operatorname{ind}(q \cup (t)) = l^n$. Then, since $\dim \mathbb{Z}/l^n = 0$, Corollary 4.2 applied to the sequence (8.1) implies that $\operatorname{ed}\mathcal{B}E[l^n] \ge l^n$, as claimed. □

Theorem 8.2. *Let $E = E_q$ denote the Tate curve over a field K as above. Then* $\mathrm{ed}_K\, E = +\infty$.

Proof. For each prime power l^n, the morphism $\mathcal{B}E[l^n] \to \mathcal{B}E$ is representable of fiber dimension 1. By Theorem 3.2,

$$\mathrm{ed}\, E \geq \mathrm{ed}\, \mathcal{B}E[l^n] = l^n - 1$$

for every $n \geq 1$. ◻

Remark 8.3. It is shown in [BS08] that if A is an abelian variety over k and k is a number field then $\mathrm{ed}_k\, A = +\infty$. On the other hand, if $k = \mathbb{C}$ is the field of complex numbers then $\mathrm{ed}_{\mathbb{C}}(A) = 2\dim(A)$; see [Bro07].

Now we can complete the proof of Theorem 1.2.

Theorem 8.4. *Let k be a field. Then* $\mathrm{ed}_k\, \mathcal{M}_{1,0} = +\infty$.

Proof. Set $F = k((t))$. By Proposition 2.8, $\mathrm{ed}_F(\mathcal{M}_{1,0} \otimes_k F) \leq \mathrm{ed}_k\, \mathcal{M}_{1,0}$, so it suffices to show that $\mathrm{ed}_F(\mathcal{M}_{1,0} \otimes_k F)$ is infinite. Consider the morphism $\mathcal{M}_{1,0} \to \mathcal{M}_{1,1}$ which sends a genus 1 curve to its Jacobian. Let E denote the Tate elliptic curve over F, which is classified by a morphism $\mathrm{Spec}\, F \to \mathcal{M}_{1,1}$. We have a Cartesian diagram:

$$\begin{array}{ccc} \mathcal{B}_k E & \longrightarrow & \mathcal{M}_{1,0} \otimes_k F \\ \downarrow & & \downarrow \\ \mathrm{Spec}\, F & \longrightarrow & \mathcal{M}_{1,1} \otimes_k F \end{array}$$

It follows that the morphism $\mathcal{B}_k E \to \mathcal{M}_{1,0}$ is representable, with fibers of dimension ≤ 0. Applying Theorem 3.2 once again, we see that

$$+\infty = \mathrm{ed}\, \mathcal{B}_F E \leq \mathrm{ed}_F(\mathcal{M}_{1,0} \otimes_k F) \leq \mathrm{ed}_k\, \mathcal{M}_{1,0},$$

as desired. ◻

Acknowledgments. We would like to thank the Banff International Research Station in Banff, Alberta (BIRS) for providing the inspiring meeting place where this work was started. We are grateful to J. Alper, K. Behrend, C.-L. Chai, D. Edidin, A. Merkurjev, B. Noohi, G. Pappas, M. Reid and B. Totaro for helpful conversations.

The research of P. Brosnan and Z. Reichstein was supported in part by NSERC Discovery grants.

The research of A. Vistoli was supported in part by the PRIN Project “Geometry of algebraic varieties and their moduli spaces”, financed by MIUR.

9. Appendix: Essential dimension of moduli of abelian varieties (by Najmuddin Fakhruddin)

In Theorem 1.2, Brosnan, Reichstein and Vistoli compute the essential dimension of various moduli stacks of curves as an application of their "genericity theorem" for the essential dimension of smooth and tame Deligne–Mumford stacks. Here we use this theorem to compute the essential dimension of some stacks of abelian varieties. Our main result is:

Theorem 9.1. *Let $g \geq 1$ be an integer, $\mathcal{A}_g$ the stack of g-dimensional principally polarised abelian varieties over a field K, and $\mathcal{B}_g$ the stack of all g-dimensional abelian varieties over K.*

(1) *If $\operatorname{char}(K) = 0$ then $\operatorname{ed}\mathcal{A}_g = g(g+1)/2 + 2^a = \operatorname{ed}\mathcal{B}_g$, where 2^a is the largest power of 2 dividing g.*
(2) *If $\operatorname{char}(K) = p > 0$ and $p \nmid |\mathrm{Sp}_{2g}(\mathbb{Z}/\ell\mathbb{Z})|$ for some prime $\ell > 2$ then $\operatorname{ed}\mathcal{A}_g = g(g+1)/2 + 2^a$ with a as above.*

For g odd this result is due to Miles Reid.

We do not know if the restriction on $\operatorname{char}(K)$ is really necessary; in Theorem 9.7 we show by elementary methods that for $g = 1$ it is not.

The main ingredient in the proof, aside from Theorem 6.1, is:

Theorem 9.2. *Let K be a field with $\operatorname{char}(K) \neq 2$ and let $\mathcal{R}_g$ be the moduli stack of (connected) étale double covers of smooth projective curves of genus g with $g > 2$ over K. Then the index of the generic gerbe of $\mathcal{R}_g$ is 2^b, where 2^b is the largest power of 2 dividing $g-1$. Furthermore, if $\mathcal{R}_g$ is tame then $\operatorname{ed}\mathcal{R}_g = 3g - 3 + 2^b$.*

The two theorems stated above are connected via the Prym map $\mathcal{R}_{g+1} \to \mathcal{A}_g$.

9.1.

It is easy to get an upper bound on the index of the generic gerbe of $\mathcal{A}_{g,d}$ over any field. This gives an upper bound on the essential dimension whenever $\mathcal{A}_{g,d}$ is smooth and tame.

Proposition 9.3. *Let $d > 0$ be an integer and $\mathcal{A}_{g,d}$ the moduli stack of abelian varieties with a polarisation of degree d over K.*

(1) *The index of the generic gerbe of each irreducible component of $\mathcal{A}_{g,d}$ is $\leq 2^a$ if $\operatorname{char}(K) \neq 2$. If $\operatorname{char}(K) = 2$ then the generic gerbes are all trivial.*
(2) *If $p = \operatorname{char}(K) > 0$ assume that $p \nmid d \cdot |\mathrm{GL}_{2g}(\mathbb{Z}/\ell\mathbb{Z})|$ (or if $d = 1$, $p \nmid |\mathrm{Sp}_{2g}(\mathbb{Z}/\ell\mathbb{Z})|$) for some prime $\ell > 2$. Then $\operatorname{ed}\mathcal{A}_{g,d} \leq g(g+1)/2 + 2^a$.*

Proof. For any g, d, $\mathcal{A}_{g,d}$ is a Deligne–Mumford stack over K with each irreducible component of dimension $g(g+1)/2$ (see [NO80] for the case $\operatorname{char}(K) \mid d$). It is a consequence of a theorem of Grothendieck [Oor71, Theorem 2.4.1], that if $p \nmid d$ then $\mathcal{A}_{g,d}$ is smooth. Furthermore, if $p \nmid |\mathrm{GL}_{2g}(\mathbb{Z}/\ell\mathbb{Z})|$ (or if $d = 1$, $p \nmid |\mathrm{Sp}_{2g}(\mathbb{Z}/\ell\mathbb{Z})|$) for ℓ as above then $\mathcal{A}_{g,d}$ is also tame. By Theorem 6.1 we see that (2) follows from (1).

Assume $\mathrm{char}(K) \neq 2$. The generic gerbe is a gerbe banded by $\mathbb{Z}/2\mathbb{Z} = \boldsymbol{\mu}_2$ so the index is a power of 2. The Lie algebra $\mathrm{Lie}_{g,d}$ of the universal family of abelian varieties over $\mathcal{A}_{g,d}$ is a vector bundle of rank g on which the automorphism $x \mapsto -x$ of the universal family induces multiplication by -1. So $\mathrm{Lie}_{g,d}$ gives rise to a *twisted* sheaf (see e.g. [Lie08, Section 3]) on the generic gerbe of each component, hence the index divides g. We conclude that the index divides the largest power of 2 dividing g, i.e. 2^a.

For any field L of characteristic 2, $\mathrm{H}^2(L, \mathbb{Z}/2\mathbb{Z}) = 0$ so the generic gerbes above are all trivial if $\mathrm{char}(K) = 2$. □

If g is odd then it follows that $\mathrm{ed}\,\mathcal{A}_{g,d} = g(g+1)/2$ whenever $\mathcal{A}_{g,d}$ is tame and smooth; this was first proved by Miles Reid using Kummer varieties. For even g we now use Theorem 9.2, which we will prove later, to complete the proof of Theorem 9.1.

Proof that Theorem 9.2 implies Theorem 9.1. We may assume that $g > 1$ since it is known that if $g = 1$ then $\mathrm{ed}\,\mathcal{A}_g (= \mathcal{B}_g) = 2$ (by Theorem 1.2 or Section 9.5).

We first recall the construction of the Prym map $P : \mathcal{R}_{g+1} \to \mathcal{A}_g$.

Let $f : X \to S$ be a family of smooth projective curves of genus $g+1$ and let $\pi : Y \to X$ be a finite étale double cover (so that the fibres of the composite morphism $f' : Y \to S$ are smooth projective curves of genus $2g+1$). Let $\mathrm{Pic}^0_{X/S}, \mathrm{Pic}^0_{Y/S}$ be the corresponding relative Jacobians and let $N : \mathrm{Pic}^0_{Y/S} \to \mathrm{Pic}^0_{X/S}$ be the norm map. The identity component of the kernel of N is an abelian scheme $\mathrm{Prym}(Y/X)$ over S of relative dimension g and the involution of Y over X induces an automorphism of $\mathrm{Pic}^0_{Y/S}$ which restricts to multiplication by -1 on $\mathrm{Prym}(Y/X)$. Furthermore, the canonical principal polarisation on $\mathrm{Pic}^0_{Y/S}$ restricts to 2λ, where λ is a principal polarisation on $\mathrm{Prym}(Y/X)$. Then P is given by sending $(f : X \to S, \pi : Y \to X)$ to $(\mathrm{Prym}(Y/X) \to S, \lambda)$. The coarse moduli space $\mathbf{R}_{g+1}$ of $\mathcal{R}_{g+1}$ is an irreducible variety and P induces a morphism, which we also denote by P, $\mathbf{R}_{g+1} \to \mathbf{A}_g$.

Let $\mathbf{A}'_g$ be the open subvariety of $\mathbf{A}_g$ corresponding to principally polarised abelian varieties A with $\mathrm{Aut}(A) = \{\pm \mathrm{Id}\}$. Then $\mathcal{A}_g|_{\mathbf{A}'_g} \to \mathbf{A}'_g$ is a μ_2 gerbe. Since $P(\mathbf{R}_{g+1}) \cap \mathbf{A}'_g \neq \emptyset$ it follows that the generic gerbe of $\mathcal{R}_{g+1}$ is isomorphic to $\mathcal{A}_g \times_{\mathbf{A}_g} \mathrm{Spec}\, K(\mathbf{R}_{g+1})$. Since the index at the generic point of an element of the Brauer groups of a smooth variety is greater than or equal to the index at any other point, it follows that the index of the generic gerbe of $\mathcal{A}_g$ is greater than or equal to the index of the generic gerbe of $\mathcal{R}_{g+1}$. By Theorem 9.2 the latter index is 2^a and then using Proposition 9.3 we deduce the first equality of Theorem 9.1(1) and also (2), since $\mathcal{A}_g$ is tame whenever $p \nmid |\mathrm{Sp}_{2g}(\mathbb{Z}/\ell\mathbb{Z})|$ for some prime $\ell \neq 2$.

Now suppose $\mathrm{char}(K) = 0$ and let A be any abelian variety of dimension g over an extension field L of K. Since A is projective, it follows that A has a polarisation of degree d for some $d > 0$ and hence corresponds to an object of $\mathcal{A}_{g,d}(L)$. By Proposition 9.3, it follows that A together with its polarisation can be defined over a field of transcendence degree $\leq g(g+1)/2 + 2^a$ over K, hence $\mathrm{ed}\,\mathcal{B}_g \leq g(g+1)/2 = 2^a$. A principally polarised abelian variety A over L such that the image of $\mathrm{Spec}\, L$ in $\mathbf{A}_g$ is the generic point has a unique polarisation which is defined whenever the abelian variety is defined. It then follows from the previous paragraph that there exists an abelian variety defined over an extension of transcendence degree $g(g+1) + 2^a$ over K which cannot be defined over a subextension of smaller transcendence degree. This proves the second equality of Theorem 9.1(1). □

9.2.

For any morphism $f : X \to S$, we denote by $\mathrm{Pic}_{X/S}$ the relative Picard functor [BLR90, Chapter 8]. If $\mathrm{Pic}_{X/S}$ is representable we use the same notation to denote the representing scheme and if $S = \mathrm{Spec}\, K$ is a field we drop it from the notation.

We recall from [BLR90, Chapter 8, Proposition 4] that if f is proper and cohomologically flat in dimension 0, then for any S-scheme T we have a canonical exact sequence

$$0 \to \mathrm{Pic}(T) \to \mathrm{Pic}(X \times_S T) \to \mathrm{Pic}_{X/S}(T) \xrightarrow{\delta} \mathrm{Br}(T) \to \mathrm{Br}(X \times T) \tag{9.1}$$

so $\delta(\tau) \in \mathrm{Br}(T)$, for $\tau \in \mathrm{Pic}_{X/S}(T)$, is the obstruction to the existence of a line bundle $\mathcal{L}$ on $X \times_S T$ representing τ.

If X is a smooth projective curve over a field, then using the morphisms $\mathrm{Sym}^d(X) \to \mathrm{Pic}^d_X$ for $d > 0$, the Riemann–Roch theorem and Serre duality one sees that the index of $\delta(\tau)$ divides $\chi(\tau) = \deg(\tau) + 1 - g$. Since δ is a homomorphism it follows that if τ is of order m then the order of $\delta(\tau)$ divides m. We deduce that in this case the index of $\delta(\tau)$ divides the largest integer dividing $g - 1$ all of whose prime divisors also divide m. Note that if $g = 1$ then we do not get any bound on the index.

9.3.

Let A be an abelian variety over a field K, let $\tau \in \mathrm{Pic}^0_A(K)$, let $\theta \in \mathrm{H}^1(K, A)$ and let P be the A-torsor corresponding to θ. Since Pic^0_P is canonically isomorphic to Pic^0_A, we may view τ as an element τ_P of $\mathrm{Pic}^0_P(K)$.

Lemma 9.4. *With the notation as above, the subgroups of* $\mathrm{Br}(K)$ *generated by* $\delta(\tau_P)$ *and* $\partial(\theta)$ *are equal, where* ∂ *is the boundary map in the long exact sequence of Galois cohomology corresponding to the extension of commutative group schemes*

$$1 \to \mathbb{G}_{\mathrm{m}} \to S \to A \to 0$$

associated to τ *via the isomorphism* $\mathrm{Pic}^0_A(K) = \mathrm{Ext}^1(A, \mathbb{G}_{\mathrm{m}})$.

Proof. We first remark that as a $\mathbb{G}_{\mathrm{m}}$-bundle on A, S is just the complement of the zero section of $\mathcal{L}$, where $\mathcal{L}$ is the line bundle on A corresponding to τ (see e.g. [Mum70, Theorem 1, p. 225]).

Now let L be any field extension of K. If $\delta(\tau_p) = 0$ in $\mathrm{Br}(L)$ then τ_P is represented by a line bundle $\mathcal{L}$ on P_L. Using the remark above, one sees that Q, the complement of the zero section in $\mathcal{L}$, is an S_L-torsor such that $Q \times_{S_L} A_L = P_L$. This implies that $\partial(\theta) = 0$ in $\mathrm{Br}(L)$. Conversely, if $\partial(\theta) = 0$ in $\mathrm{Br}(L)$ then there is a (unique) S_L-torsor Q such that $Q \times_{S_L} A_L = P_L$. This gives a $\mathbb{G}_{\mathrm{m}}$-bundle over P_L and hence a line bundle on P_L which represents τ_P, so we must have $\delta(\tau_p) = 0$ in $\mathrm{Br}(L)$.

It follows that the splitting fields of $\partial(\theta)$ and $\delta(\tau_P)$ are the same, hence the two elements must generate the same subgroup in $\mathrm{Br}(K)$. □

It is very likely that $\delta(\tau_P)$ and $\partial(\theta)$ are equal, at least up to sign, but we shall not need this.

9.4.

Given a smooth projective curve X over a field K and an element τ of $\mathrm{Pic}_X(K)$, one may ask how large the index of $\delta(\tau)$ can be. In the case that τ is torsion, the theorem below shows that the best upper bound on the index which is valid over all fields is the one given in Section 9.2.

Theorem 9.5. *Let $g > 0$ be an integer, $n > 0$ an integer such that n divides $g - 1$ and $\mathrm{char}(K) \nmid n$, and $m > 0$ such that $m \mid n$ and m, n have the same prime factors. Then there exists an extension L of K, a smooth projective curve X of genus g over L with $\mathrm{Aut}(X_{\overline{L}}) = \{\mathrm{Id}\}$ if $g > 2$, and an element τ of order m in $\mathrm{Pic}_X(L)$ such that the index of $\delta(\tau)$ is n.*

The theorem for all g will be deduced from the slightly stronger result below for $g = 1$.

Proposition 9.6. *Let $n > 0$ be an integer such that $\mathrm{char}(K) \nmid n$ and $m > 0$ such that $m \mid n$ and m, n have the same prime factors. Then there exists an extension M of K, a smooth projective geometrically irreducible curve P of genus 1 over M and an element σ of order m in $\mathrm{Pic}_P(M)$ such that the index of $\delta(\sigma)$ is n. Furthermore, there exists an extension M' of M of degree n such that $P(M')$ is infinite.*

Proof. We first replace K by $\overline{K}(s, t)$ where s, t are indeterminates. We fix an isomorphism $\boldsymbol{\mu}_n \cong \mathbb{Z}/n\mathbb{Z}$ which we use to identify all $\boldsymbol{\mu}_n^{\otimes i}$, $i \in \mathbb{Z}$. For the elements $(s), (t) \in \mathrm{H}^1(K, \boldsymbol{\mu}_n)$ consider $\alpha = (s) \cup (t) \in \mathrm{H}^2(K, \boldsymbol{\mu}_n^{\otimes 2}) \cong \mathrm{H}^2(K, \boldsymbol{\mu}_n) = {}_n\mathrm{Br}(K)$. It is well known and easy to see that this element of $\mathrm{Br}(K)$ has both order and index equal to n. Let K' be the function field of the Brauer–Severi variety corresponding to the division algebra over K representing $m\alpha$. By a theorem of Amitsur [Ami55, Theorem 9.3] the image of α in $\mathrm{Br}(K')$ has order m and by a theorem of Schofield and Van den Bergh [SVdB92, Theorem 2.1] its index is still n.

Let M be the field of Laurent series $K'((q))$ and let E be the Tate elliptic curve over M associated to the element $sq^n \in M^\times$. For any finite extension M' of M there is a canonical Galois equivariant isomorphism

$$E(M') \cong {M'}^\times / \langle sq^n \rangle.$$

From this we get a canonical exact sequence

$$1 \to \boldsymbol{\mu}_n \to E[n] \to \mathbb{Z}/n\mathbb{Z} \to 0$$

where $1 \in \mathbb{Z}/n\mathbb{Z}$ is the image of any n-th root of sq^n in $\overline{M'}$. For any $\phi \in \mathrm{H}^1(M, \mathbb{Z}/n\mathbb{Z})$, one easily checks using the definitions that $\partial(\phi) \in \mathrm{H}^2(M, \boldsymbol{\mu}_n) = \mathrm{H}^2(M, \boldsymbol{\mu}_n \otimes \mathbb{Z}/n\mathbb{Z})$

is equal to $(s) \cup \phi$, where ∂ denotes the boundary map in the long exact sequence of Galois cohomology associated to the above short exact sequence of Galois modules. It follows that if we identify $(t) \in \mathrm{H}^1(M, \boldsymbol{\mu}_n)$ with an element of $\mathrm{H}^1(M, \mathbb{Z}/n\mathbb{Z})$ using our chosen isomorphism $\mathbb{Z}/n\mathbb{Z} \cong \boldsymbol{\mu}_n$, then $\beta := \partial((t)) = (s) \cup (t) \in \mathrm{H}^2(K, \boldsymbol{\mu}_n) \subset \mathrm{Br}(M)$. Thus β also has order m and index n (since the index is the smallest dimension of a linear subvariety of the Brauer–Severi variety and such varieties are preserved by specialisation). In particular, it is in the image of the inclusion map $\mathrm{H}^1(M, \boldsymbol{\mu}_m) \to \mathrm{H}^1(M, \boldsymbol{\mu}_n)$.

Let E' be the quotient of E by $\boldsymbol{\mu}_m$, so E' is also an elliptic curve over M. Let $I_n \subset E'[n]$ be the image of $E[n]$, so we have exact sequences

$$1 \to \boldsymbol{\mu}_m \to E[n] \to I_n \to 0 \quad \text{and} \quad 1 \to \boldsymbol{\mu}_{n/m} \to I_n \to \mathbb{Z}/n\mathbb{Z} \to 0.$$

By construction, the boundary map of the second sequence maps the element $(t) \in \mathrm{H}^1(M, \mathbb{Z}/n\mathbb{Z})$ to 0 in $\mathrm{H}^1(M, \boldsymbol{\mu}_{n/m})$, hence (t) lifts to an element $\gamma \in \mathrm{H}^1(M, I_n)$. Clearly γ is mapped to $\beta \in \mathrm{H}^2(M, \boldsymbol{\mu}_m)$ by the boundary map of the first exact sequence.

Now let $M' = M(t^{1/n}) = K'(t^{1/n})((q))$. The restriction of γ in $\mathrm{H}^1(M', I_n)$ goes to 0 in $\mathrm{H}^1(M', \mathbb{Z}/n\mathbb{Z})$ by construction, hence it comes from an element of $\mathrm{H}^1(M', \boldsymbol{\mu}_{n/m})$. We have a commutative diagram

$$\begin{array}{ccc} \mathrm{H}^1(M', \boldsymbol{\mu}_n) & \longrightarrow & \mathrm{H}^1(M', E) \\ \downarrow & & \downarrow \\ \mathrm{H}^1(M', \boldsymbol{\mu}_{n/m}) & \longrightarrow & \mathrm{H}^1(M', E') \end{array}$$

where the vertical maps are induced by quotienting by $\boldsymbol{\mu}_m$. The first vertical map is surjective and the inclusion $\boldsymbol{\mu}_n \to E(\overline{M'})$ factors as $\boldsymbol{\mu}_n \to \overline{M'}^{\times} \to E(\overline{M'})$, so it follows from Hilbert's Theorem 90 that the bottom horizontal map is zero. Therefore θ, the image of γ in $\mathrm{H}^1(M, E')$, restricts to 0 in $\mathrm{H}^1(M', E')$.

Let P be the E'-torsor corresponding to θ, so Pic^0_P is canonically isomorphic to E'. The image of $E[m]$ in E' is naturally isomorphic to $\mathbb{Z}/m\mathbb{Z}$; let σ denote $1 \in \mathbb{Z}/m\mathbb{Z} \subset E'(M) = \mathrm{Pic}_P(M)$. Pushing out the exact sequence

$$1 \to \boldsymbol{\mu}_m \to E \to E' \to 0$$

via the inclusion $\boldsymbol{\mu}_m \to \mathbb{G}_\mathrm{m}$ we get an exact sequence

$$1 \to \mathbb{G}_\mathrm{m} \to S \to E' \to 0$$

whose class in $\mathrm{Ext}^1(E', \mathbb{G}_\mathrm{m})$ generates the kernel of the map $\mathrm{Ext}^1(E', \mathbb{G}_\mathrm{m}) \to \mathrm{Ext}^1(E, \mathbb{G}_\mathrm{m})$. Under the canonical isomorphisms $\mathrm{Ext}^1(E', \mathbb{G}_\mathrm{m}) \cong \mathrm{Pic}^0_{E'}(M) \cong E'(M)$, $1 \in \mathbb{Z}/m\mathbb{Z} \subset E'(M)$ is a generator of the above kernel, so it follows that the two elements generate the same subgroup of $\mathrm{Ext}^1(E', \mathbb{G}_\mathrm{m})$.

It now follows from Lemma 9.4 that $\delta(\sigma)$ and β generate the same subgroup of $\mathrm{Br}(M)$; in particular, $\delta(\sigma)$ has index n. Since θ becomes 0 in $\mathrm{H}^1(M', E')$, it follows that $P_{M'} \cong E'_{M'}$. Since $E'(M)$ is infinite, so is $E(M')$ and therefore also $P(M')$.

We conclude that M, P, σ and M' satisfy all the conditions of the proposition. □

Proof of Theorem 9.5. If $g = 1$ the result follows from Proposition 9.6 so we may assume that $g > 1$.

Let $r = (g-1)/n$ and let M, P, σ and M' be as in Proposition 9.6. Note that since the index of $\delta(\sigma)$ is n, any closed point of P must have degree divisible by n. Let $p_1, \dots, p_r$ be distinct closed points of P of degree n and let Y be the stable curve over M obtained by gluing two copies of P along all the p_i's, i.e. the p_i in one copy is identified with the p_i in the other copy using the identity map on residue fields. The arithmetic genus of Y is $1 + 1 + rn - 1 = 1 + rn = g$. Using the natural map $\pi : Y \to P$ which is the identity on both components, we get a morphism $\pi^* : \mathrm{Pic}_P \to \mathrm{Pic}_Y$ and we let $\sigma' = \pi^*(\sigma) \in \mathrm{Pic}_Y(M)$. Note that $\delta(\sigma) = \delta(\sigma') \in \mathrm{Br}(M)$.

Let $R = M[[x]]$ and let $f : \mathbf{Y} \to \operatorname{Spec} R$ be a generic smoothing of Y. So $\mathbf{Y}$ is a regular scheme and f is a flat proper morphism with closed fibre equal to Y (see for example [DM69, Section 1]). By a theorem of Raynaud [Ray70, Théorème 8.2.1], $\mathrm{Pic}^0_{\mathbf{Y}/R}$ is representable by a separated and smooth group scheme of finite type over R. Since $\mathrm{char}(M) \nmid m$, the endomorphism of $\mathrm{Pic}^0_{\mathbf{Y}/R}$ given by multiplication by m is étale. Since R is complete, it follows that σ' can be lifted to an element $\boldsymbol{\sigma}$ in $\mathrm{Pic}^0_{\mathbf{Y}/R}(R)$ of order m.

Consider $\delta(\boldsymbol{\sigma}) \in \mathrm{Br}(R)$. Since $\mathrm{Br}(R) = \mathrm{Br}(M)$, we see by the functoriality of the exact sequences in (9.1) that $\delta(\boldsymbol{\sigma}) = \delta(\sigma') = \delta(\sigma)$.

Now let $L = M(\!(x)\!)$, let X be the generic fibre of f and let τ be the restriction of $\boldsymbol{\sigma}$ in $\mathrm{Pic}^0_X(L)$; by the genericity of the deformation it follows that $\mathrm{Aut}(X_{\overline{L}}) = \{\mathrm{Id}\}$ if $g > 2$. Again by the functoriality of the exact sequences in (9.1) we see that $\delta(\tau)$ is the image of $\delta(\boldsymbol{\sigma}) = \delta(\sigma)$ in $\mathrm{Br}(L)$. Thus $\delta(\tau)$ has index n as required. □

Theorem 9.2 is a simple consequence of Theorem 9.5.

Proof of Theorem 9.2. Since $\mathcal{R}_g$ is a smooth irreducible Deligne–Mumford stack of dimension $3g - 3$, it follows from Theorem 6.1 that to compute $\operatorname{ed} \mathcal{R}_g$ when $\mathcal{R}_g$ is tame it suffices to compute the index of the generic gerbe.

The coarse moduli space $\mathbf{R}_g$ of $\mathcal{R}_g$ is generically a fine moduli space parametrizing smooth projective curves X of genus g over S with a non-trivial element of order 2 of $\mathrm{Pic}_{X/S}(S)$. Thus over the generic point $\operatorname{Spec} K(\mathbf{R}_g) \in \mathbf{R}_g$ we have a smooth projective curve C of genus g and an element $\sigma \in \mathrm{Pic}_C(K(\mathbf{R}_g))$ of order 2. It follows that the element of $\mathrm{Br}(K(\mathbf{R}_g))$ represented by the generic gerbe of $\mathcal{R}_g$ is the obstruction to the existence of a line bundle $\mathcal{L}$ over C whose class in $\mathrm{Pic}_C(K(\mathbf{R}_g))$ is equal to σ.

If $b = 0$, then $g - 1$ is odd hence the generic gerbe is trivial. So assume $b > 0$ and let X, L and σ be obtained by applying Theorem 9.5 with $m = 2$ and $n = 2^b$. Since $\mathrm{Aut}(X_{\overline{L}}) = \{\mathrm{Id}\}$ it follows that the image of the map $\operatorname{Spec} L \to \mathbf{R}_g$ lies in the smooth locus $\mathbf{R}'_g$ of $\mathcal{R}_g$. Since the restriction of the map $\mathcal{R}_g \to \mathbf{R}_g$ is a $\boldsymbol{\mu}_2$ gerbe, it follows that the index of the generic gerbe is $\geq 2^b$. Since the index must also divide $g - 1$ it follows that we must have equality as claimed. □

9.5. The essential dimension of $\mathcal{A}_1$ over arbitrary fields

We do not know the essential dimension of $\mathcal{A}_g$ over fields of small characteristic. However, it follows from classical formulae [Sil86, Appendix A, Proposition 1.1] that $\operatorname{ed} \mathcal{A}_1$

$= 2$ over any field of characteristic $\neq 2$ and $\operatorname{ed} \mathcal{A}_1 \leq 3$ over any field of characteristic 2. We prove here the following

Theorem 9.7. $\operatorname{ed} \mathcal{A}_1 = 2$ *over any field of characteristic* 2.

Proof. It suffices to prove the theorem over $\mathbb{F}_2$ since it is easy to see that $\operatorname{ed} \mathcal{A}_1 \geq 2$ over any field.

Any elliptic curve E over a field K of characteristic 2 with $j(E) \neq 0$ has an affine equation [Sil86, Appendix A]

$$y^2 + xy = x^3 + a_2x^2 + a_6, \quad a_2, 0 \neq a_6 \in K,$$

hence it suffices to compute the essential dimension of the residual gerbe corresponding to elliptic curves E with $j(E) = 0$. Any such curve has an affine equation

$$y^2 + a_3y = x^3 + a_4x + a_6, \quad a_3 \neq 0, a_4, a_6 \in K.$$

We let E be the curve corresponding to the equation $y^2 + y = x^3$ over $\mathbb{F}_2$ and denote by $\operatorname{Aut}(E)$ its automorphism group scheme.

By [Sil86, Appendix A, Proposition 1.2] and its proof, $\operatorname{Aut}(E)$ is an étale group scheme over $\mathbb{F}_2$ of order 24. As a scheme it is given by the equations $U^3 = 1$, $S^4 + S = 0$ and $T^2 + T = 0$, where U, S, T are coordinates on $\mathbb{A}^3$. Given a solution (u, s, t) of these equations, the corresponding automorphism $E \to E$ is given in the above coordinates by $(x, y) \mapsto (x', y')$ with $x = u^2x' + s^2$ and $y = y' + u^2sx' + t$. Thus, if $f_i : E \to E$, $i = 1, 2$, over a field K is given by a tuple (u_i, s_i, t_i) then $f_2 \circ f_1 : E \to E$ is given by the coordinate change

$$x = u_1^2x_1 + s_1^2 = u_1^2(u_2^2x_2 + s_2^2) + s_1^2 = (u_1u_2)^2x_2 + (u_1s_2 + s_1)^2$$

and

$$\begin{aligned} y &= y_1 + u_1^2s_1x_1 + t_1 = (y_2 + u_2^2s_2x_2 + t_2) + u_1^2s_1(u_2^2x_2 + s_2^2) + t_1 \\ &= y_2 + (u_1u_2)^2(u_1s_2 + s_1)x_2 + (t_1 + u_1^2s_1s_2^2 + t_2). \end{aligned}$$

Thus $f_2 \circ f_1$ corresponds to the triple $(u_1u_2, u_1s_2 + s_1, t_1 + t_2 + u_1^2s_1s_2^2)$.

Clearly $\operatorname{Aut}(E)$ becomes a constant group scheme over any field containing $\mathbb{F}_4$; one may see that this constant group scheme is isomorphic to $\mathrm{SL}_2(\mathbb{F}_3)$ by considering its action on $E[3]$. The centre of $\operatorname{Aut}(E)$ is the constant group scheme $\mathbb{Z}/2\mathbb{Z}$, the non-trivial element corresponds to the tuple $(1, 0, 1)$ and acts by multiplication by -1 on E. Let G be the quotient of $\operatorname{Aut}(E)$ by its centre. It is given by the equations $U^3 = 1$, $S^4 + S = 0$ and the quotient map corresponds to forgetting the last coordinate.

Let $B \subset \mathrm{SL}_2(\mathbb{F}_4)$ be the subgroup of upper triangular matrices, viewed as a closed subgroup scheme of $\mathrm{SL}_{2,\mathbb{F}_2}$ in the natural way. The formula for composition in $\operatorname{Aut}(E)$ given above shows that the map on points $G \to B$ given by $(u, s) \mapsto \left(\begin{smallmatrix} u & us \\ 0 & u^2 \end{smallmatrix}\right)$ induces an isomorphism of group schemes over $\mathbb{F}_2$. Thus G is a closed subgroup scheme of $\mathrm{GL}_{2,\mathbb{F}_2}$ which maps injectively into $\mathrm{PGL}_{2,\mathbb{F}_2}$, so $\operatorname{ed} G = 1$.

Now we have a central extension of group schemes over $\mathbb{F}_2$,

$$0 \to \mathbb{Z}/2\mathbb{Z} \to \mathrm{Aut}(E) \to G \to 1,$$

which for any extension field K of $\mathbb{F}_2$ gives rise to an exact sequence of pointed sets

$$\mathrm{H}^1(K, \mathbb{Z}/2\mathbb{Z}) \xrightarrow{\alpha} \mathrm{H}^1(K, \mathrm{Aut}(E)) \xrightarrow{\beta} \mathrm{H}^1(K, G) \xrightarrow{\partial} \mathrm{H}^2(K, \mathbb{Z}/2\mathbb{Z}).$$

Since $\mathrm{H}^2(K, \mathbb{Z}/2\mathbb{Z}) = 0$ it follows that β is surjective. Thus $\mathrm{H}^1(K, \mathbb{Z}/2\mathbb{Z})$ operates on $\mathrm{H}^1(K, \mathrm{Aut}(E))$ and the quotient is $\mathrm{H}^1(K, G)$ by [Gir71, III, Proposition 3.4.5(iv)]. Since both $\mathbb{Z}/2\mathbb{Z}$ and G have essential dimension 1, it follows that $\operatorname{ed} \mathrm{Aut}(E) \leq 2$.

The residual gerbe at the point E of $\mathcal{A}_1$ is neutral, so it is isomorphic to $\mathcal{B}\,\mathrm{Aut}(E)$, hence has ed ≤ 2. Since the generic gerbe is isomorphic to $\mathcal{B}\,\mathbb{Z}/2\mathbb{Z}$, we conclude that $\operatorname{ed} \mathcal{A}_1 = 2$. □

Acknowledgments. I thank Arvind Nair and Madhav Nori for useful conversations.

References

[AV02] Abramovich, D., Vistoli, A.: Compactifying the space of stable maps. J. Amer. Math. Soc. **15**, 27–75 (2002) Zbl 0991.14007 MR 1862797

[Ami55] Amitsur, S. A.: Generic splitting fields of central simple algebras. Ann. of Math. (2) **62**, 8–43 (1955) Zbl 0066.28604 MR 0070624

[AV04] Arsie, A., Vistoli, A.: Stacks of cyclic covers of projective spaces. Compos. Math. **140**, 647–666 (2004) Zbl 1169.14301 MR 2041774

[BF03] Berhuy, G., Favi, G.: Essential dimension: a functorial point of view (after A. Merkurjev). Doc. Math. **8**, 279–330 (2003) Zbl 1101.14324 MR 2029168

[BF04] Berhuy, G., Favi, G.: Essential dimension of cubics. J. Algebra **278**, 199–216 (2004) Zbl 1068.14035 MR 2068074 MR 1457337

[BR05] Berhuy, G., Reichstein, Z.: On the notion of canonical dimension for algebraic groups. Adv. Math. **198**, 128–171 (2005) Zbl 1097.11018 MR 2183253

[BLR90] Bosch, S., Lütkebohmert, W., Raynaud, M.: Néron Models. Ergeb. Math. Grenzgeb. (3) 21, Springer, Berlin (1990) Zbl 0705.14001 MR 1045822

[Bro07] Brosnan, P.: The essential dimension of a g-dimensional complex abelian variety is $2g$. Transform. Groups **12**, 437–441 (2007) Zbl 1127.14043 MR 2356317

[BRV10] Brosnan, P., Reichstein, Z., Vistoli, A.: Essential dimension, spinor groups and quadratic forms. Ann. of Math. **171**, 533–544 (2010) Zbl pre05712735 MR 2630047

[BRV07] Brosnan, P., Reichstein, Z., Vistoli, A.: Essential dimension and algebraic stacks. arXiv:math/0701903v1 [math.AG] (2007)

[BS08] Brosnan, P., Sreekantan, R.: Essential dimension of abelian varieties over number fields. C. R. Math. Acad. Sci. Paris **346**, 417–420 (2008) Zbl 1137.14034 MR 2417561

[BR97] Buhler, J., Reichstein, Z.: On the essential dimension of a finite group. Compos. Math. **106**, 159–179 (1997) Zbl 0905.12003 MR 1457337

[CS06] Chernousov, V., Serre, J.-P.: Lower bounds for essential dimensions via orthogonal representations. J. Algebra **305**, 1055–1070 (2006) Zbl 1181.20042 MR 2266867

[CTKM07] Colliot-Thélène, J.-L., Karpenko, N. A., Merkur'ev, A. S.: Rational surfaces and the canonical dimension of the group PGL_6. Algebra i Analiz **19**, 159–178 (2007) (in Russian) Zbl 1206.14070 MR 2381945

[Con] Conrad, B.: Keel–Mori theorem via stacks. http://www.math.lsa.umich.edu/˜bdconrad/papers/coarsespace.pdf

[DD99] Dèbes, P., Douai, J.-C.: Gerbes and covers. Comm. Algebra **27**, 577–594 (1999) Zbl 0917.18008 MR 1671938

[Del80] Deligne, P.: La conjecture de Weil. II. Inst. Hautes Études Sci. Publ. Math. **52**, 137–252 (1980) Zbl 0456.14014 MR 0601520

[DM69] Deligne, P., Mumford, D.: The irreducibility of the space of curves of given genus. Inst. Hautes Études Sci. Publ. Math. **36**, 75–109 (1969) Zbl 0181.48803 MR 0262240

[DL] Dhillon, A., Lemire, N.: Upper bounds for the essential dimension of the moduli stack of SL_n-bundles over a curve. Transform. Groups **14**, 747–770 (2009) Zbl 1190.14004 MR 2577196

[Ful98] Fulton, W.: Intersection Theory. 2nd ed., Ergeb. Math. Grenzgeb. 2, Springer, Berlin (1998) Zbl 0885.14002 MR 1644323

[Gar09] Garibaldi, S.: Cohomological invariants: exceptional groups and spin groups. Mem. Amer. Math. Soc. **200**, no. 937 (2009) Zbl 1191.11009 MR 2528487

[GMS03] Garibaldi, S., Merkurjev, A., Serre, J.-P.: Cohomological Invariants in Galois Cohomology. Univ. Lecture Ser. 28, Amer. Math. Soc., Providence, RI (2003) Zbl 1159.12311 MR 1999383

[Gir71] Giraud, J.: Cohomologie non abélienne. Grundlehren Math. Wiss. 179, Springer, Berlin (1971) Zbl 0226.14011 MR 0344253

[Gro64] Grothendieck, A.: Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas. I. Inst. Hautes Études Sci. Publ. Math. **20**, 259 pp. (1964) Zbl 0136.15901 MR 0173675

[Gro67] Grothendieck, A.: Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas. IV. Inst. Hautes Études Sci. Publ. Math. **32**, 361 pp. (1967) Zbl 0153.22301 MR 0238860

[Har77] Hartshorne, R.: Algebraic Geometry. Grad. Texts in Math. 52, Springer, New York (1977) Zbl 0367.14001 MR 0463157

[JLY02] Jensen, C. U., Ledet, A., Yui, N.: Generic Polynomials. Math. Sci. Res. Inst. Publ. 45, Cambridge Univ. Press, Cambridge (2002) Zbl 1042.12001 MR 1969648

[KM06] Karpenko, N. A., Merkurjev, A. S.: Canonical p-dimension of algebraic groups, Adv. Math. **205**, 410–433 (2006) Zbl 1119.14041 MR 2258262

[KM08] Karpenko, N. A., Merkurjev, A. S.: Essential dimension of finite p-groups. Invent. Math. **172**, 491–508 (2008) Zbl 1200.12002 MR 2393078

[Kor00] Kordonskiĭ, V. È.: On the essential dimension and Serre's conjecture II for exceptional groups. Mat. Zametki **68**, 539–547 (2000) (in Russian) Zbl 1009.2055 MR 1823140

[Kre99] Kresch, A.: Cycle groups for Artin stacks. Invent. Math. **138**, 495–536 (1999) Zbl 0938.14003 MR 1719823

[LMB00] Laumon, G.: Moret-Bailly, L.: Champs algébriques. Ergeb. Math. Grenzgeb. 39, Springer, Berlin (2000) Zbl 0945.14005 MR 1771927

[Led02] Ledet, A.: On the essential dimension of some semi-direct products. Canad. Math. Bull. **45**, 422–427 (2002) Zbl 1024.12004 MR 1935274

[Lem04] Lemire, N.: Essential dimension of algebraic groups and integral representations of Weyl groups. Transform. Groups **9**, 337–379 (2004) Zbl 1076.14060 MR 2105732

[Lie08] Lieblich, M.: Twisted sheaves and the period-index problem. Compos. Math. **144**, 1–31 (2008) Zbl 1133.14018 MR 2388554

[Mil80] Milne, J. S.: Étale Cohomology. Princeton Math. Ser. 33, Princeton Univ. Press, Princeton, NJ (1980) Zbl 0433.14012 MR 0559531

[Mum70] Mumford, D.: Abelian varieties. Tata Inst. Fundam. Res. Stud. Math. 5, Oxford Univ. Press (1970) Zbl 0223.14022 MR 0282985

[NO80] Norman, P., Oort, F.: Moduli of abelian varieties. Ann. of Math. (2) **112**, 413–439 (1980) Zbl 0483.14010 MR 0595202

[Oor71] Oort, F.: Finite group schemes, local moduli for abelian varieties, and lifting problems. Compos. Math. **23**, 265–296 (1971) Zbl 0223.14024 MR 0301026

[Ray70] Raynaud, M.: Spécialisation du foncteur de Picard. Inst. Hautes Études Sci. Publ. Math. **38**, 27–76 (1970) Zbl 0207.51602 MR 0282993

[Rei00] Reichstein, Z.: On the notion of essential dimension for algebraic groups. Transform. Groups **5**, 265–304 (2000) Zbl 0981.20033 MR 1780933

[RV11] Reichstein, Z., Vistoli, A.: A genericity theorem for algebraic stacks and essential dimension of hypersurfaces. arXiv:1103.1611

[RY00] Reichstein, Z., Youssin, B.: Essential dimensions of algebraic groups and a resolution theorem for G-varieties. Canad. J. Math. **52**, 1018–1056 (2000) Zbl 1044.14023 MR 1782331

[SVdB92] Schofield, A., Van den Bergh, M.: The index of a Brauer class on a Brauer–Severi variety. Trans. Amer. Math. Soc. **333**, 729–739 (1992) Zbl 0778.12004 MR 1061778

[Ser02] Serre, J.-P.: Galois Cohomology. Springer Monogr. Math., Springer, Berlin (2002) Zbl 1004.12003 MR 1867431

[SGA72] Théorie des topos et cohomologie étale des schémas. Tome 2. Séminaire de Géométrie Algébrique du Bois-Marie 1963–1964 (SGA 4). Lecture Notes in Math. 270, Springer, Berlin (1972) Zbl 0237.00012 MR 0354653

[SGA73] Groupes de monodromie en géométrie algébrique. II. Séminaire de Géométrie Algébrique du Bois-Marie 1967–1969 (SGA 7 II). Lecture Notes in Math. 340, Springer, Berlin (1973) Zbl 0258.00005 MR 0354657

[Sil86] Silverman, J. H.: The Arithmetic of Elliptic Curves. Grad. Texts in Math. 106, Springer, New York (1986) Zbl 0585.14026 MR 0817210