

Patching and Galois theory

David Harbater*

Dept. of Mathematics, University of Pennsylvania

Abstract: Galois theory over $\mathbb{C}(x)$ is well-understood as a consequence of Riemann's Existence Theorem, which classifies the algebraic branched covers of the complex projective line. The proof of that theorem uses analytic and topological methods, including the ability to construct covers locally and to patch them together on the overlaps. To study the Galois extensions of $k(x)$ for other fields k , one would like to have an analog of Riemann's Existence Theorem for curves over k . Such a result remains out of reach, but partial results in this direction can be proven using patching methods that are analogous to complex patching, and which apply in more general contexts. One such method is formal patching, in which formal completions of schemes play the role of small open sets. Another such method is rigid patching, in which non-archimedean discs are used. Both methods yield the realization of arbitrary finite groups as Galois groups over $k(x)$ for various classes of fields k , as well as more precise results concerning embedding problems and fundamental groups. This manuscript describes such patching methods and their relationships to the classical approach over \mathbb{C} , and shows how these methods provide results about Galois groups and fundamental groups.

* Supported in part by NSF Grants DMS9970481 and DMS0200045. (Version 3.4, Aug. 6, 2002.)

2000 Mathematics Subject Classification. Primary 14H30, 12F12, 14D15; Secondary 13B05, 13J05, 12E30.

Key words and phrases: fundamental group, Galois covers, patching, formal scheme, rigid analytic space, affine curves, deformations, embedding problems.

Contents

Section 1: Introduction

- 1.1 Overview
- 1.2 Galois groups and fundamental groups
- 1.3 Structure of this manuscript

Section 2: Complex patching

- 2.1 Riemann's Existence Theorem
- 2.2 GAGA
- 2.3 Complex patching and constructing covers

Section 3: Formal patching

- 3.1 Zariski's formal holomorphic functions
- 3.2 Grothendieck's formal schemes
- 3.3 Formal patching and constructing covers

Section 4: Rigid patching

- 4.1 Tate's rigid analytic spaces
- 4.2 Rigid geometry via formal geometry
- 4.3 Rigid patching and constructing covers

Section 5: Toward Riemann's Existence Theorem

- 5.1 Embedding problems and the geometric case of Shafarevich's Conjecture
- 5.2 Arithmetic lifting, embedding problems, and patching
- 5.3 Abhyankar's Conjecture and embedding problems

References

Section 1: Introduction

This manuscript discusses patching methods and their use in the study of Galois groups and fundamental groups. There is a particular focus on Riemann’s Existence Theorem and the inverse Galois problem, and their generalizations (both known and conjectured). This first section provides an introduction, beginning with a brief overview of the topic in Section 1.1. More background about Galois groups and fundamental groups is provided in Section 1.2. Section 1.3 then discusses the overall structure of the paper, briefly indicating the content of each later section.

Section 1.1. Overview.

Galois theory is algebraic in its origins, arising from the study of polynomial equations and their solvability. But it has always had intimate connections to geometry. This is evidenced, for example, when one speaks of an “icosahedral Galois extension” — meaning a field extension whose Galois group is A_5 , the symmetry group of an icosahedron.

Much progress in Galois theory relies on connections to geometry, particularly on the parallel between Galois groups and the theory of covering spaces and fundamental groups in topology. This parallel is more than an analogy, with the group-theoretic and topological approaches being brought together in the context of algebraic geometry. The realization of all finite groups as Galois groups over $\mathbb{C}(x)$ is an early example of this approach.

In recent years, this approach has drawn heavily on the notion of *patching*, i.e. on “cut-and-paste” constructions that build covers locally and then combine them to form a global cover with desired symmetries. Classically this could be performed only for spaces defined over \mathbb{C} , in order to study Galois groups over fields like $\mathbb{C}(x)$. But by carrying complex analytic methods over to other settings — most notably via formal and rigid geometry — results in Galois theory have now been proven for a broad array of rings and fields by means of patching methods.

This paper provides an overview of this approach to Galois theory via patching, both in classical and non-classical contexts. A key theme in both contexts is *Riemann’s Existence Theorem*. In the complex case, that result provides a classification of the finite Galois extensions of the field $\mathbb{C}(x)$ and more generally of the function field K of any Riemann surface X . (In the case $K = \mathbb{C}(x)$, X is the Riemann sphere, i.e. the complex projective line $\mathbb{P}_{\mathbb{C}}^1$.) This classification relies on the correspondence between these field extensions and the branched covers of X , and on the classification of the branched covers of X with given branch locus B . This correspondence between field extensions and covers in turn proceeds by proving the equivalence of covers in the algebraic, analytic, and topological senses, and then relying on topology to classify the covering spaces of the complement of a finite set $B \subset X$. In demonstrating this equivalence, one regards branched covers as being given locally over discs, where the cover breaks up into a union of cyclic components, and where

agreement on the overlaps is given in order to define the cover globally. Using “complex patching” (specifically, Serre’s result GAGA), such an analytic cover in fact arises from a cover of complex algebraic curves, given by polynomial equations.

This patching approach proves that every finite group is a Galois group over $\mathbb{C}(x)$ and more generally over K as above, and it provides the structure of the absolute Galois group of the field. Moreover, if one fixes a finite set of points B , the approach shows which finite groups are Galois groups of covers with that branch locus, and how those groups fit together as quotients of the fundamental group of $X - B$.

The success of this approach made it desirable to carry it over to other settings, in order to study the Galois theory of other fields K — e.g. $K = k(x)$ where k is a field other than \mathbb{C} , or even arithmetic fields K . In order to do this, one needs to carry over the notion of patching. This is, if K is the function field of a scheme X (e.g. $X = \text{Spec } R$, where R is an integral domain whose fraction field is K), then one would like to construct covers of X locally, with agreement on the overlaps, and then be able to assert the existence of a global cover that induces this local data. The difficulty is that one needs an appropriate topology on X . Of course, there is the Zariski topology, but that is too coarse. Namely, if U is a Zariski open subset of an irreducible scheme X , then giving a branched cover $V \rightarrow U$ is already tantamount to giving a cover over all of X , since $X - U$ is just a closed subset of lower dimension (and one can take the normalization of X in V). Instead, one needs a finer notion, which behaves more like the complex metric topology in the classical setting, and where one can speak of the ring of “holomorphic functions” on any open set in this topology.

In this manuscript, after discussing the classical form of Riemann’s Existence Theorem via GAGA for complex curves, we present two refinements of the Zariski topology that allow patching constructions to take place in many (but not all) more general settings. These approaches of *formal* and *rigid patching* are roughly equivalent to each other, but they developed separately. Each relies on an analog of GAGA, whose proof parallels the proof of Serre’s original GAGA. These approaches are then used to realize finite groups as Galois groups over various function fields, and to show how these groups fit together in the tower of all extensions of the field (corresponding to information about the structure of the absolute Galois group — or of a fundamental group, if the branch locus is fixed).

Underlying this entire approach is the ability to pass back and forth between algebra and geometry. This ability is based on the relationship between field extensions and covers, with Galois groups of field extensions corresponding to groups of deck transformations of covers, and with absolute Galois groups of fields playing a role analogous to fundamental groups of spaces. This relationship is reviewed in Section 1.2 below, where basic terminology is also introduced. (Readers who are familiar with this material may wish to skip §1.2.) Section 1.3 then describes the structure of this paper as a whole.

I wish to thank Claus Lehr, Rachel Pries, Kate Stevenson, Jakob Stix, and Oliver Watson for their comments and suggestions on an earlier version of this manuscript, and Florian Pop for his collaboration on Theorem 5.2.3.

Section 1.2. Galois groups and fundamental groups.

Traditionally, Galois theory studies field extensions by means of symmetry groups (viz. their Galois groups). Covering spaces can also be studied using symmetry groups (viz. their groups of deck transformations). In fact, the two situations are quite parallel.

In the algebraic situation, the basic objects of study are field extensions $L \supset K$. To such an extension, one associates its symmetry group, viz. the Galois group $\text{Gal}(L/K)$, consisting of automorphisms of L that fix all the elements of K . If L is a finite extension of K of degree $[L : K] = n$, then the order of the Galois group is at most n ; and the extension is *Galois* if the order is exactly n , i.e. if the extension is as symmetric as possible. (For finite extensions, this is equivalent to the usual definition in terms of being normal and separable.)

In the geometric situation, one considers topological covering spaces $f : Y \rightarrow X$. There is the associated symmetry group $\text{Aut}(Y/X)$ of deck transformations, consisting of self-homeomorphisms ϕ of Y that preserve the map f (i.e. such that $f \circ \phi = f$). If $Y \rightarrow X$ is a finite cover of degree n , then the order of the covering group is at most n ; and the extension is “*regular*” (in the terminology of topology) if the order is exactly n , i.e. if the cover is as symmetric as possible. To emphasize the parallel, we will refer to the covering group $\text{Aut}(Y/X)$ as the *Galois group* of the cover, and will call the symmetric covers *Galois* rather than “regular” (and will instead reserve the latter word for another meaning that is used in connection with covers in arithmetic algebraic geometry).

The parallel extends further: If L is a finite Galois extension of K with Galois group G , then the intermediate extensions M of K are in bijection with the subgroups H of G ; namely a subgroup H corresponds to its fixed field $M = L^H$, and an intermediate field M corresponds to the Galois group $H = \text{Gal}(L/M) \subset G$. Moreover, M is Galois over K if and only if H is normal in G ; and in that case $\text{Gal}(M/K) = G/H$. Similarly, if $Y \rightarrow X$ is a finite Galois cover with group G , then the intermediate covers $Z \rightarrow X$ are in bijection with the subgroups H of G ; namely a subgroup H corresponds to the quotient space $Z = Y/H$, and an intermediate cover $Z \rightarrow X$ corresponds to the Galois group $H = \text{Gal}(Y/Z) \subset G$. Moreover, $Z \rightarrow X$ is Galois if and only if H is normal in G ; and in that case $\text{Gal}(Z/X) = G/H$. Thus in both the algebraic and geometric contexts, there is a “Fundamental Theorem of Galois Theory”.

The reason behind this parallel can be illustrated by a simple example. Let X and Y be two copies of $\mathbb{C} - \{0\}$, with complex parameters x and y respectively, and let $f : Y \rightarrow X$ be given by $x = y^n$, for some integer $n > 1$. This is a degree n Galois cover whose Galois group is the cyclic group C_n of order n , where the generator of the Galois group takes

$y \mapsto \zeta_n y$ (with $\zeta_n \in \mathbb{C}$ being a primitive n th root of unity). If one views X and Y not just as topological spaces, but as copies of the affine variety $\mathbb{A}_{\mathbb{C}}^1 - \{0\}$, then f is the morphism corresponding to the inclusion of function fields, $\mathbb{C}(x) \hookrightarrow \mathbb{C}(y)$, given by $x \mapsto y^n$. This inclusion is a Galois field extension of degree n whose Galois group is C_n , whose generator acts by $y \mapsto \zeta_n y$. (Strictly speaking, if the Galois group of covers acts on the left, then the Galois group of fields acts on the right.)

In this example, the Fundamental Theorem for covering spaces implies the Fundamental Theorem for the extension of function fields, since intermediate covers $Z \rightarrow X$ correspond to intermediate field extensions $M \supset \mathbb{C}(x)$ (where M is the function field of Z). More generally, one can consider Galois covers of schemes that are not necessarily defined over \mathbb{C} , and in that context have both algebraic and geometric forms of Galois theory.

In order to extend the idea of covering space to this setting, one needs to define a class of finite morphisms $f : Y \rightarrow X$ that generalizes the class of covering spaces (in the complex metric topology) for complex varieties. The condition of being a covering space in the Zariski topology does not do this, since an irreducible scheme X will not have any irreducible covers in this sense, other than the identity map. (Namely, if $Y \rightarrow X$ is evenly covered over a dense open set, then it is a disjoint union of copies of X , globally.) Instead one uses the notion of *finite étale covers*, i.e. finite morphisms $f : Y \rightarrow X$ such that locally at every point of Y , the scheme Y is given over X by m polynomials f_1, \dots, f_m in m variables y_1, \dots, y_m , and such that the Jacobian determinant $(\partial f_i / \partial y_j)$ is locally invertible. The point is that for spaces over \mathbb{C} , this condition is equivalent to f having a local section near every point (by the Inverse Function Theorem, where “local section” means in the complex metric); and for a finite morphism, satisfying this latter condition is equivalent to being a finite covering space (in the complex metric sense).

For finite étale covers of an irreducible scheme X , one then has a Fundamental Theorem of Galois Theory as above. If one restricts to complex varieties, one obtains the geometric situation discussed above. And if one restricts to X of the form $\text{Spec } K$ (for some field K), then one recovers Galois theory for field extensions (with $\text{Spec } L \rightarrow \text{Spec } K$ corresponding to a field extension $L \supset K$).

As in the classical situation, one can consider fundamental groups. Namely, let X be an irreducible normal scheme, and let K be the function field of X^{red} (or just of X , if X has no nilpotents). Also, let \bar{K} be the separable closure of K (so just the algebraic closure of K , if X has characteristic 0). Then the function fields of the (reduced) finite étale Galois covers of X form a direct system of extensions of K contained in \bar{K} , and so the covers form an inverse system — as do their Galois groups. In the complex case, this system of groups is precisely the one obtained by taking the finite quotients of the classical topological fundamental group of X . More generally, the *algebraic* (or *étale*) *fundamental*

group of the scheme X is defined to be the inverse limit of this inverse system of finite groups (or equivalently, the automorphism group of the inverse system of covers); this is a profinite group whose finite quotients precisely form the above inverse system. This group is denoted by $\pi_1^{\text{ét}}(X)$; it is the profinite completion of the topological fundamental group $\pi_1^{\text{top}}(X)$, in the special case of complex varieties X . (Thus, for $X = \mathbb{A}_{\mathbb{C}}^1 - \{0\}$, $\pi_1^{\text{ét}}(X)$ is $\hat{\mathbb{Z}}$ rather than just \mathbb{Z} .) When working in the algebraic context, one generally just writes $\pi_1(X)$ for $\pi_1^{\text{ét}}(X)$.

As in the classical situation, one may also consider *branched covers*. For Riemann surfaces X , giving a branched cover of X is equivalent to giving a covering space of $X - B$, where the finite set B is the branch locus of the cover (i.e. where it is not étale). More generally, we can define a finite *branched cover* (or for short, a *cover*) of a scheme X to be a finite morphism $Y \rightarrow X$ that is generically separable. Most often, one restricts to the case that X and Y are normal integral schemes. In this case, the finite branched covers of X are in natural bijection with the finite separable field extensions of the function field of K . The notions of “Galois” and “Galois group” carry over to this situation: The *Galois group* $\text{Gal}(Y/X)$ of a branched cover $f : Y \rightarrow X$ consists of the self-automorphisms of Y that preserve f . And a finite branched cover $Y \rightarrow X$ is *Galois* if X and Y are irreducible, and if the degree of the automorphism group is equal to the degree of f . Sometimes one wants to allow X or Y to be reducible, or even disconnected; and sometimes one wants to make explicit the identification of a given finite group G (e.g. the abstract group D_5) with the Galois group of a cover. In this situation, one speaks of a *G-Galois cover* $f : Y \rightarrow X$; this means a cover together with a homomorphism $\alpha : G \simeq \text{Gal}(Y/X)$ such that via α , the group G acts simply transitively on a generic geometric fibre.

Thus in order to understand the Galois theory of an integral scheme X , one would like to classify the finite étale covers of X in terms of their branch loci, ramification behavior, and Galois groups; and also to describe how they fit together in the tower of covers. In the classical case of complex curves (Riemann surfaces), this is accomplished by Riemann’s Existence Theorem (discussed in Section 2.1 below). A key goal is to carry this result over to more general contexts. Such a classification would in particular give an explicit description of the profinite group $\pi_1(X)$, and also of the set $\pi_A(X)$ of finite quotients of $\pi_1(X)$ (i.e. the Galois groups of finite étale covers of X). Similarly, on the generic level (where arbitrary branching is allowed), one would like to have an explicit description of the *absolute Galois group* $G_K = \text{Gal}(K^s/K)$ of the function field K of X (where K^s is the separable closure of K). This in turn would provide an explicit description of the finite quotients of G_K , i.e. the Galois groups of finite field extensions of K .

Beyond the above parallel between field extensions and covers, there is a second connection between those two theories, relating to *fields of definition* of covers. The issue can be illustrated by a variant on the simple example given earlier. As before, let X and

Y be two copies of $\mathbb{C} - \{0\}$, with complex parameters x and y respectively; let $n > 1$ be an integer; and this time let $f : Y \rightarrow X$ be given by $x = \pi y^n$ (where π is the usual transcendental constant). Again, the cover is Galois, with cyclic Galois group generated by $g : y \mapsto \zeta_n y$. This cover, along with its Galois action, is defined by polynomials over \mathbb{C} ; but after a change of variables $z = \pi^{1/n} y$, the cover is given by polynomials over $\bar{\mathbb{Q}}$ (viz. $z^n = x$, and $g : z \mapsto \zeta_n z$). In fact, if X is any curve that can be defined over $\bar{\mathbb{Q}}$ (e.g. if X is the complement of finitely many \mathbb{Q} -points in \mathbb{P}^1), then any finite étale cover of the induced complex curve $X_{\mathbb{C}}$ can in fact itself be defined over $\bar{\mathbb{Q}}$ (along with its Galois action, in the G -Galois case; see Remark 2.1.6 below). And since there are only finitely many polynomials involved, it can even be defined over some number field K .

The key question here is what this number field is, in terms of the topology of the cover. By Riemann's Existence Theorem, the Galois covers of a given base X are classified; e.g. those over $\mathbb{P}^1 - \{0, 1, \infty\}$ correspond to the finite quotients of the free group on two generators. So in that case, given a finite group G together with a pair of generators, what is the number field K over which the corresponding G -Galois cover of $\mathbb{P}^1 - \{0, 1, \infty\}$ is defined? Actually, this field of definition K is not uniquely determined, although there is an "ideal candidate" for K , motivated by Galois theory. Namely, if $\omega \in G_{\bar{\mathbb{Q}}} := \text{Gal}(\bar{\mathbb{Q}}/\mathbb{Q})$, then ω acts on the set of (isomorphism classes of) G -Galois covers, by acting on the coefficients. If a G -Galois cover is defined over a number field K , then any $\omega \in \text{Gal}(\bar{\mathbb{Q}}/K) \subset \text{Gal}(\bar{\mathbb{Q}}/\mathbb{Q})$ must carry this G -Galois cover to itself. So we may consider the *field of moduli* M for the G -Galois cover, defined to be the fixed field of all the ω 's in $G_{\bar{\mathbb{Q}}}$ that carry the G -Galois cover to itself. This is then contained in every field of definition of the G -Galois cover. Moreover it is the *intersection* of the fields of definition; and in key cases (e.g. if G is abelian or has trivial center) it is the unique minimal field of definition [CH]. One can then investigate the relationship between the (arithmetic) Galois theory of M and the (geometric) Galois theory of the given cover.

In particular, if X is a Zariski open subset of the Riemann sphere $\mathbb{P}_{\mathbb{C}}^1$, and if M is known to be the minimal field of definition of the given G -Galois cover, then G is a Galois group over $M(x)$ — and hence over the field M , by Hilbert's Irreducibility Theorem [FJ, Chapter 11]. The most important special case is that of finite simple groups G such that $M \subset \mathbb{Q}(\zeta_n)$ for some n , for some G -Galois cover of $X = \mathbb{P}^1 - \{0, 1, \infty\}$. Many examples of such simple groups and covers have been found, using the technique of *rigidity*, developed in work of Matzat, Thompson, Belyi, Fried, Feit, Shih, and others. In each such case, the group has thus been realized as a Galois group over some explicit $\mathbb{Q}(\zeta_n)$, the most dramatic example being Thompson's realization [Th], over \mathbb{Q} itself, of the *monster group*, the largest of the 26 sporadic finite simple groups (having order $\approx 8 \cdot 10^{53}$). See the books [Se7], [Vö], and [MM] for much more about rigidity. (Note that the name *rigidity* is related to the same term in the theory of local systems and differential equations [Ka], but is unrelated

to the notion of *rigid spaces* discussed elsewhere in this paper.)

Thus Galois theory appears in the theory of covering spaces in two ways — in a geometric form, coming from the parallel between covering groups and Galois groups of field extensions, and in an arithmetic form, coming from fields of definition. This situation can be expressed in another way, via the *fundamental exact sequence*

$$1 \rightarrow \pi_1(\bar{X}) \rightarrow \pi_1(X) \rightarrow G_k \rightarrow 1.$$

Here X is a geometrically connected variety over a field k ; G_k is the absolute Galois group $\text{Gal}(\bar{k}/k)$; and $\bar{X} = X \times_k \bar{k}$, the space obtained by regarding X over \bar{k} . The kernel $\pi_1(\bar{X})$ is the geometric part of the fundamental group, while the cokernel $G_k = \pi_1(\text{Spec } k)$ is the arithmetic part. There is a natural outer action of G_k on $\pi_1(\bar{X})$, obtained by lifting an element of G_k to $\pi_1(X)$ and conjugating; and this action corresponds to the action of the absolute Galois group on covers, discussed just above in connection with fields of moduli. Moreover, this exact sequence splits if X has a point ξ defined over k (by applying π_1 to $\text{Spec } k = \{\xi\} \hookrightarrow X$), and in that case there is a true action of G_k on $\pi_1(\bar{X})$.

For a given cover $Y \rightarrow X$ of irreducible varieties over a base field k , one can separate out the arithmetic and geometric parts by letting ℓ be the algebraic closure of the function field of X in the function field of Y . The given cover then factors as $Y \rightarrow X_\ell \rightarrow X$, where $X_\ell = X \times_k \ell$. Here $X_\ell \rightarrow X$ is a purely arithmetic extension, coming just from an extension of constants; one sometimes refers to it as a *constant* extension. The cover $Y \rightarrow X_\ell$ is purely geometric, in the sense that the base field ℓ of X_ℓ is algebraically closed in the function field of Y ; and so Y is in fact geometrically irreducible (i.e. even $Y \times_\ell \bar{k}$ is irreducible). One says that $Y \rightarrow X_\ell$ is a *regular* cover of ℓ -curves (or that Y is “regular over ℓ ”; note that this use of the word “regular” is unrelated to the topological notion of “regularity” mentioned at the beginning of this Section 1.2.) Algebraically, we say that the function field of Y is a *regular* extension of that of X_ℓ , as an ℓ -algebra (or that the function field of Y is “regular over ℓ ”). Of course if a field k is algebraically closed, then every cover of a k -variety (and every Galois extension of $k(x)$) is automatically regular over k .

This paper will discuss methods of patching to construct complex covers, and how those methods can be carried over to covers defined over various other classes of fields. One of those classes will be that of algebraically closed fields. But there will be other classes as well (particularly complete fields and “large” fields), for which the fields need not be algebraically closed. In those settings, we will be particularly interested in regular covers. In particular, in carrying over to those situations the fact that every finite group is a Galois group over $\mathbb{C}(x)$, it will be of interest to show that a field k has the property that every finite group is the Galois group of a *regular* extension of $k(x)$; this is the “regular inverse Galois problem” over k .

Section 1.3. Structure of this manuscript.

This paper is intended as an introduction to patching methods and their use in Galois theory. The main applications are to Riemann’s Existence Theorem and related problems, particularly finding fundamental groups and solving the inverse Galois problem over curves. This is first done in Section 2 in the classical situation of complex curves (Riemann surfaces), using patching in the complex topology. Guided by the presentation in Section 2, later sections describe non-classical patching methods that apply to curves over other fields, and use these methods to obtain analogs of results and proofs that were presented in Section 2. In particular, Sections 3 and 4 each parallel Section 2, with Section 3 discussing patching methods using formal schemes, and Section 4 discussing patching using rigid analytic spaces. In each of these cases, the context provides enough structure to carry over results and proofs from the classical situation of Section 2 to the new situation. Although a full analog of Riemann’s Existence Theorem remains unknown in the non-classical settings, the partial analogs that are obtained are sufficient to solve the geometric inverse Galois problem in these settings. Further results about Galois groups and fundamental groups are presented in Section 5, using both formal and rigid methods.

Section 2 begins in §2.1 with a presentation of Riemann’s Existence Theorem for complex curves (Theorem 2.1.1). There, an equivalence between algebraic, analytic, and topological notions of covers provides an explicit classification of the covers of a given base (Corollary 2.1.2). As a consequence, one solves the inverse Galois problem over $\mathbb{C}(x)$ (Corollary 2.1.4). Section 2.2 shows how Riemann’s Existence Theorem follows from Serre’s result GAGA (Theorem 2.1.1), which gives an equivalence between coherent sheaves in the algebraic and analytic settings (basically, between the set-up in Hartshorne [Hrt2] and the one in Griffiths-Harris [GH]). The bulk of §2.2 is devoted to proving GAGA, by showing that the two theories behave in analogous ways (e.g. that their cohomology theories agree, and that sufficient twisting provides a sheaf that is generated by its global sections). This proof follows that of Serre [Se3]. Specific examples of complex covers are considered in §2.3, including ones obtained by taking copies of the base and pasting along slits; these are designed to emphasize the “patching” nature of GAGA and Riemann’s Existence Theorem, and to motivate what comes after.

Section 3 treats *formal patching*, a method to extend complex patching to more general situations. The origins of this approach, going back to Zariski, are presented in §3.1, along with the original motivation of “analytic continuation” along subvarieties. Related results of Ferrand-Raynaud and Artin, which permit patching constructions consistent with Zariski’s original point of view, are also presented here. Grothendieck’s extension of Zariski’s viewpoint is presented in §3.2, where formal schemes are discussed. The key result presented here is Grothendieck’s Existence Theorem, or GFGA (Theorem 3.2.1), which is a formal analog of GAGA. We present a proof here which closely parallels Serre’s

proof of complex GAGA that appeared in Section 2. Afterwards, a strengthening of this result, due to the author, is shown, first for curves (Theorem 3.2.8) and then in higher dimensions (Theorem 3.2.12). Applications to covers and Galois theory are then given in §3.3. These include the author’s result that every finite group can be regularly realized over the fraction field of a complete local ring other than a field (Theorem 3.3.1); the corollary that the same is true for algebraically closed fields of arbitrary characteristic (Corollary 3.3.5); and Pop’s extension of this corollary to “large fields” (Theorem 3.3.6). There is also an example that illustrates the connection to complex “slit” covers that were considered in §2.3.

Section 4 considers a parallel approach, viz. *rigid patching*. Tate’s original view of this approach is presented in §4.1. Unlike formal patching, which is motivated by considerations of abstract varieties and schemes, this viewpoint uses an intuition that remains closer to the original analytic approach. On the other hand, there are technical difficulties to be overcome, relating to the non-uniqueness of analytic continuation with respect to a non-archimedean metric. Tate’s original method of dealing with this (via the introduction of “rigidifying data”) is given in §4.1, and the status of rigid GAGA from this point of view is discussed. Then §4.2 presents a reinterpretation of rigid geometry from the point of view of formal geometry, along the lines introduced by Raynaud and worked out later by him and by Bosch and Lütkebohmert. This point of view allows rigid GAGA to “come for free” as a consequence of the formal version. It also establishes a partial dictionary between the formal and rigid approaches, allowing one to use the formal machinery together with the rigid intuition. Applications to covers and Galois theory are then given in §4.3 — in particular the regular realization of groups over complete fields (Theorem 4.3.1, paralleling Theorem 3.3.1); and Pop’s “Half Riemann Existence Theorem” for henselian fields (Theorem 4.3.3), classifying “slit covers” in an arithmetic context.

Section 5 uses both formal and rigid methods to consider results that go further in the direction of a full Riemann’s Existence Theorem in general contexts. In order to go beyond the realization of individual Galois groups, §5.1 discusses *embedding problems* for the purpose of seeing how Galois groups “fit together” in the tower of all covers. In particular, a result of the author and Pop is presented, giving the structure of the absolute Galois group of the function field of a curve over an algebraically closed field (Theorem 5.1.1). This result relies on showing that finite embedding problems over such curves have proper solutions. That fact about embedding problems does not extend to more general fields, but we present Pop’s result that it holds for *split* embedding problems over large fields (Theorem 5.1.9). Section 5.2 presents Colliot-Thélène’s result on the existence of covers of the line with given Galois group and a given fibre, in the case of a large base field (Theorem 1). Both this result and Pop’s Theorem 5.1.9 can be subsumed by a single result, due to the author and Pop; this appears as Theorem 5.2.3. The classification of covers

with given branch locus is taken up in §5.3, where Abhyankar’s Conjecture (Theorem 5.3.1, proven by Raynaud and the author) is discussed, along with Pop’s strengthening in terms of embedding problems (Theorem 5.3.4). Possible generalizations to higher dimensional spaces are discussed, along with connections to embedding problems for such spaces and their function fields. As a higher dimensional local application, it is shown that every finite split embedding problem over $\mathbb{C}((x, y))$ has a proper solution. But as discussed there, most related problems in higher dimension, including the situation for the rational function field $\mathbb{C}(x, y)$, remain open.

This manuscript is adapted, in part, from lectures given by the author at workshops at MSRI during the fall 1999 semester program on Galois groups and fundamental groups. Like those lectures, this paper seeks to give an overall view of its subject to beginners and outsiders, as well as to researchers in Galois theory who would benefit from a general presentation, including new and recent results. It follows an approach that emphasizes the historical background and motivations, the geometric intuition, and the connections between various approaches to patching — in particular stressing the parallels between the proofs in the complex analytic and formal contexts, and between the frameworks in the formal and rigid situations. The manuscript ties together results that have appeared in disparate locations in the literature, and highlights key themes that have recurred in a variety of contexts. In doing so, the emphasis is on presenting the main themes first, and afterwards discussing the ingredients in the proofs (thus following, to some extent, the organization of a lecture series).

Certain results that have been stated in the literature in a number of special cases are given here in a more natural, or more general, setting (e.g. see Theorems 3.2.8, 3.2.12, 5.1.9, 5.1.10, and 5.2.3). Quite a number of remarks are given, describing open problems, difficulties, new directions, and alternative versions of results or proofs. In particular, there is a discussion in Section 5.3 of the higher dimensional situation, which is just beginning to be understood. A new result in the local case is shown there (Theorem 5.3.9), and the global analog is posed.

The only prerequisite for this paper is a general familiarity with concepts in algebraic geometry along the lines of Hartshorne [Hrt2], although some exposure to arithmetic notions would also be helpful. Extensive references are provided for further exploration, in particular the books on inverse Galois theory by Serre [Se7], Völklein [Vö], and Malle-Matzat [MM], and the book on fundamental groups in algebraic geometry edited by Bost, Loeser, and Raynaud [BLoR].

Section 2: Complex patching

This section presents the classical use of complex patching methods in studying Galois branched covers of Riemann surfaces, and it motivates the non-classical patching methods discussed in the later sections of this manuscript. Section 2.1 begins with the central result, Riemann’s Existence Theorem, which classifies covers. In its initial version, it shows the equivalence between algebraic covers and topological covers; but since topological covers can be classified group-theoretically, so can algebraic covers. It is the desire to classify algebraic covers (and correspondingly, field extensions) so concretely that provides much of the motivation in this manuscript.

The key ingredient in the proof of Riemann’s Existence Theorem is Serre’s result GAGA. This is proven in Section 2.2, using an argument that will itself motivate the proof of a key result in Section 3 (formal GAGA). Some readers may wish to skip Section 2.2 on first reading, and go directly to Section 2.3, where examples of Riemann’s Existence Theorem are given. These examples show how topological covers can be constructed by building them locally and then patching; and the “slit cover” example there will motivate constructions that will appear in analogous contexts later, in Sections 3 and 4.

Section 2.1. Riemann’s Existence Theorem

Algebraic varieties over the complex numbers can be studied topologically and analytically, as well as algebraically. This permits the use of tools that are not available for varieties over more general fields and rings. But in order to use these tools, one needs a link between the objects that exist in the algebraic, analytic, and topological categories. In the case of fundamental groups, this link is the correspondence between covers in the three settings. Specifically, in the case of complex algebraic curves, the key result is

Theorem 2.1.1. (Riemann’s Existence Theorem) *Let X be a smooth connected complex algebraic curve, which we also regard as a complex analytic space and as a topological space with respect to the classical topology. Then the following categories are equivalent:*

- (i) *Finite étale covers of the variety X ;*
- (ii) *Finite analytic covering maps of X ;*
- (iii) *Finite covering spaces of the topological space X .*

Using this theorem, results about topological fundamental groups, which can be obtained via loops or covering spaces, can be translated into results about étale covers and étale fundamental groups. In particular, there is the following corollary:

Corollary 2.1.2. (Explicit form of Riemann’s Existence Theorem) *Let $r \geq 0$, let $\xi_1, \dots, \xi_r \in \mathbb{P}_{\mathbb{C}}^1$, and let $X = \mathbb{P}_{\mathbb{C}}^1 - \{\xi_1, \dots, \xi_r\}$. Let G be a finite group, and let \mathcal{C} be the set of equivalence classes of r -tuples $\underline{g} = (g_1, \dots, g_r) \in G^r$ such that g_1, \dots, g_r generate G and satisfy $g_1 \cdots g_r = 1$. Here we declare two such r -tuples $\underline{g}, \underline{g}'$ to be equivalent if they are*

uniformly conjugate (i.e. if there is an $h \in G$ such that for $1 \leq i \leq r$ we have $g'_i = hg_i h^{-1}$). Then there is a bijection between the G -Galois connected finite étale covers of X and the elements of \mathcal{C} . Moreover this correspondence is functorial under the operation of taking quotients of G , and also under the operation of deleting more points from $\mathbb{P}_{\mathbb{C}}^1$.

Namely, the topological fundamental group of X is given by

$$\pi_1^{\text{top}}(X) = \langle x_1, \dots, x_r \mid x_1 \cdots x_r = 1 \rangle,$$

where the x_i 's correspond to loops around the ξ_i 's, from some base point $\xi_0 \in X$. The fundamental group can be identified with the Galois group (of deck transformations) of the universal cover, and the finite quotients of π_1 can be identified with pointed finite Galois covers of X . To give a quotient map $\pi_1 \rightarrow G$ is equivalent to giving the images of the x_i 's, i.e. giving g_i 's as above. Making a different choice of base point on the cover (still lying over ξ_0) uniformly conjugates the g_i 's. So the elements of \mathcal{C} are in natural bijection with G -Galois connected covering spaces of X ; and by Riemann's Existence Theorem these are in natural bijection with G -Galois connected finite étale covers of X .

In the situation of Corollary 2.1.2, the uniform conjugacy class of (g_1, \dots, g_r) is called the *branch cycle description* of the corresponding cover of X [Fr1]. It has the property that the corresponding *branched* cover of $\mathbb{P}_{\mathbb{C}}^1$ contains points η_1, \dots, η_r over ξ_1, \dots, ξ_r respectively, such that g_i generates the inertia group of η_i over ξ_i . (See [Fr1] and Section 2.3 below for a further discussion of this.)

The corollary can also be stated for more general complex algebraic curves. Namely if X is obtained by deleting r points from a smooth connected complex projective curve of genus γ , then the topological fundamental group of X is generated by elements $x_1, \dots, x_r, y_1, \dots, y_\gamma, z_1, \dots, z_\gamma$, subject to the relation $x_1 \cdots x_r [y_1, z_1] \cdots [y_\gamma, z_\gamma] = 1$, where the y 's and z 's correspond to loops around the “handles” of the topological surface X . The generalization of the corollary then replaces \mathcal{C} by the set of equivalence classes of $(r + 2\gamma)$ -tuples of generators that satisfy this longer relation.

Note that the above results are stated only for finite covers, whereas the topological results are a consequence of the fact that the fundamental group is isomorphic to the Galois group of the universal cover (which is of infinite degree, unless $X = \mathbb{P}_{\mathbb{C}}^1$ or $\mathbb{A}_{\mathbb{C}}^1$). Unfortunately, the universal cover is *not* algebraic — e.g. if E is a complex elliptic curve, then the universal covering map $\mathbb{C} \rightarrow E$ is not a morphism of varieties (only of topological spaces and of complex analytic spaces). As a result, in algebraic geometry there is no “universal étale cover”; only a “pro-universal cover”, consisting of the inverse system of finite covers. The *étale fundamental group* is thus defined to be the automorphism group of this inverse system; and for complex varieties, this is then the profinite completion of the topological fundamental group. By Corollary 2.1.2 and this definition, we have the following result, which some authors also refer to as “Riemann's Existence Theorem”:

Corollary 2.1.3. *Let $r \geq 1$, and let $S = \{\xi_1, \dots, \xi_r\}$ be a set of r distinct points in $\mathbb{P}_{\mathbb{C}}^1$. Then the étale fundamental group of $X = \mathbb{P}_{\mathbb{C}}^1 - S$ is the profinite group Π_r on generators x_1, \dots, x_r subject to the single relation $x_1 \cdots x_r = 1$. This is isomorphic to the free profinite group on $r - 1$ generators.*

Also note that there is a bijection between finite field extensions of $\mathbb{C}(x)$ and connected finite étale covers of (variable) Zariski open subsets of $\mathbb{P}_{\mathbb{C}}^1$. The reverse direction is obtained by taking function fields; and the forward direction is obtained by considering the integral closure of $\mathbb{C}[x]$ in the extension field, taking its spectrum and the corresponding morphism to the complex affine line, and then deleting points where the morphism is not étale. Under this bijection, Galois field extensions correspond to Galois finite étale covers. The corresponding statements remain true for general complex connected projective curves and their function fields.

Reinterpreting Corollary 2.1.2 via this bijection, we obtain a correspondence between field extensions and tuples of group elements (which is referred to as the “algebraic version of Riemann’s Existence Theorem” in [Vö, Thm.2.13]). From this point of view, we obtain as an easy consequence the following result in the Galois theory of field extensions:

Corollary 2.1.4. *The inverse Galois problem holds over $\mathbb{C}(x)$. That is, for every finite group G , there is a finite Galois field extension K of $\mathbb{C}(x)$ such that the Galois group $\text{Gal}(K/\mathbb{C}(x))$ is isomorphic to G .*

For short, we say that “every finite group is a Galois group over $\mathbb{C}(x)$ ”. Corollary 2.1.4 is immediate from Corollary 2.1.2, since for every finite group G we may pick a set of generators $g_1, \dots, g_r \in G$ whose product is 1, and a set of distinct points $\xi_1, \dots, \xi_r \in \mathbb{P}_{\mathbb{C}}^1$; and then obtain a connected G -Galois étale cover of $X = \mathbb{P}_{\mathbb{C}}^1 - \{\xi_1, \dots, \xi_r\}$. The corresponding extension of function fields is then the desired extension K of $\mathbb{C}(x)$. (Similarly, if K_0 is any field of transcendence degree 1 over \mathbb{C} , we may prove the inverse Galois problem over K_0 by applying the generalization of Corollary 2.1.2 to the complex projective curve with function field K_0 , minus r points.)

Even more is true:

Corollary 2.1.5. *The absolute Galois group of $\mathbb{C}(x)$ is a free profinite group, of rank equal to the cardinality of \mathbb{C} .*

This follows from the fact that the correspondences in Corollary 2.1.2 are compatible with quotient maps and with deleting more points. For then, one can deduce that the absolute Galois group is the inverse limit of the étale fundamental groups of $\mathbb{P}_{\mathbb{C}}^1 - S$, where S ranges over finite sets of points. The result then follows from Corollary 2.1.3, since $\pi_1^{\text{ét}}(\mathbb{P}_{\mathbb{C}}^1 - S)$ is free profinite on $\text{card } S - 1$ generators; see [Do] for details.

Remark 2.1.6. Corollaries 2.1.4 and 2.1.5 also hold for $\bar{\mathbb{Q}}(x)$, and this fact can be deduced from a refinement of Riemann’s Existence Theorem. Namely, consider a smooth curve V

defined over $\bar{\mathbb{Q}}$, and let $X = V_{\mathbb{C}}$ be the base change of V to \mathbb{C} (i.e. the “same” curve, viewed over the complex numbers). Then every finite étale cover of X is induced from a finite étale cover of V (along with its Galois action). In particular, take V to be an open subset of \mathbb{P}^1 . Then there is a bijection between topological covering spaces of $S^2 - (r \text{ points})$ and finite étale covers of $\mathbb{P}_{\bar{\mathbb{Q}}}^1 - (r \text{ points})$, where S^2 is the sphere. This yields the analogs of Corollaries 2.1.4 and 2.1.5 for $\bar{\mathbb{Q}}$.

This refinement of Riemann’s Existence Theorem can be proven by first observing that a finite étale Galois cover $f : Y \rightarrow X$ is defined over some subalgebra $A \subset \mathbb{C}$ that is of finite type over $\bar{\mathbb{Q}}$. That is, there are finitely many equations that define the cover (along with its Galois action, and the property of being étale); and their coefficients all lie in such an A , thereby defining a finite étale Galois cover $f_A : \mathcal{Y} \rightarrow \mathcal{X} := X \times_{\bar{\mathbb{Q}}} A$. This cover can be regarded as a *family* of covers of X , parametrized by $T := \text{Spec } A$. The inclusion $i : A \hookrightarrow \mathbb{C}$ defines a \mathbb{C} -point ξ of T , and the fibre over this point is (tautologically) the given cover $f : Y \rightarrow X$. Meanwhile, let κ be a $\bar{\mathbb{Q}}$ -point of T , and consider the corresponding fibre $g : W \rightarrow V$. Both ξ and κ define \mathbb{C} -points on $T_{\mathbb{C}} = T \times_{\bar{\mathbb{Q}}} \mathbb{C}$, corresponding to two fibres of a connected family of finite étale covers of X . But in any continuous connected family of covering spaces of a constant base, the fibres are the same (because $\pi_1(X_1 \times X_2) = \pi_1(X_1) \times \pi_1(X_2)$ in topology). Thus the complex cover induced by $g : W \rightarrow V$ agrees with $f : Y \rightarrow X$, as desired.

Using ideas of this type, Grothendieck proved a stronger result [Gr5, XIII, Cor. 2.12], showing that Riemann’s Existence Theorem carries over from \mathbb{C} to any algebraically closed field of characteristic 0. But in fact, Corollaries 2.1.4 and 2.1.5 even carry over to characteristic $p > 0$; see Sections 3.3 and 5.1 below. \square

The assertions in Corollaries 2.1.2-2.1.5 above (and the analogous results for $\bar{\mathbb{Q}}(x)$ mentioned in the above remark) are purely algebraic in nature. It would therefore be desirable to have purely algebraic proofs of these assertions — and this would also have the consequence of permitting generalizations of these results to a variety of other contexts beyond those considered in [Gr5, XIII]. Unfortunately, no purely algebraic proofs of these results are known. Instead, the only known proofs rely on Riemann’s Existence Theorem 2.1.1, which (because it states an equivalence involving algebraic, analytic, and topological objects) is inherently non-algebraic in nature.

Concerning the proof of Riemann’s Existence Theorem 2.1.1, the easy part is the equivalence of (ii) and (iii). Namely, there is a forgetful functor from the category in (ii) to the one in (iii). We wish to show that the functor induces a surjection on isomorphism classes of objects, and bijections on morphisms between corresponding pairs of objects. (Together these automatically guarantee injectivity on isomorphism classes of objects.) In the case of objects, consider a topological covering space $f : Y \rightarrow X$. The space X is evenly covered by Y ; i.e. X is a union of open discs D_i such that $f^{-1}(D_i)$ is a disjoint union of

finitely many connected open subsets D_{ij} of Y , each mapping homeomorphically onto D_i . By giving each D_{ij} the same analytic structure as D_i , and using the same identifications on the overlapping D_{ij} 's as on the overlapping D_i 's, we give Y the structure of a Riemann surface, i.e. a complex manifold of dimension 1; and $f : Y \rightarrow X$ is an object in (ii) whose underlying topological cover is the one we were given. This shows surjectivity on isomorphism classes. Injectivity on morphisms is trivial, and surjectivity on morphisms follows from surjectivity on objects, since if $f : Y \rightarrow X$ and $g : Z \rightarrow X$ are analytic covering spaces and if $\phi : Y \rightarrow Z$ is an morphism of topological covers (i.e. $g\phi = f$), then ϕ is itself a topological cover (of Z), hence a morphism of analytic curves and thus of analytic covers.

With regard to the equivalence of (i) and (ii), observe first of all that while the objects in (ii) and (iii) are covering spaces with respect to the metric topology, those in item (i) are finite étale covers rather than covering spaces with respect to the Zariski topology (since those are all trivial, because non-empty Zariski open subsets are dense). And indeed, if one forgets the algebraic structure and retains just the analytic (or topological) structure, then a finite étale cover of complex curves is a covering space in the metric topology, because of the Inverse Function Theorem. Thus every object in (i) yields an object in (ii). (Note also that finite étale covers can be regarded as “covering spaces in the étale topology”, making the parallels between (i), (ii), (iii) look a bit closer.)

The deeper and more difficult part of the proof of Riemann's Existence Theorem is going from (ii) to (i) — and in particular, showing that every finite analytic cover of an algebraic curve is itself algebraic. One approach to this is to show that every *compact* Riemann surface (i.e. compact one-dimensional complex manifold) is in fact a complex algebraic curve, with enough meromorphic functions to separate points. See [Vö, Chaps. 5,6] for a detailed treatment of this approach. Another approach is to use Serre's result GAGA (“géométrie algébrique et géométrie analytique”), from the paper [Se3] of the same name. That result permits the use of “analytic patching” in complex algebraic geometry; i.e. constructing analytic objects locally so as to agree on overlaps, and then concluding that there is a global algebraic object that induces the local structures compatibly. It is this approach that we describe next, and it is this approach that motivates much of the discussion in the later parts of this paper.

Section 2.2. GAGA

Serre's result GAGA [Se3] permits the construction of sheaves of modules over a complex projective algebraic curve, in the Zariski topology, by constructing the sheaf analytically, in the classical complex metric topology. From this assertion about sheaves of modules, the corresponding result follows for sheaves of algebras, and therefore for covers. This in turn leads to a proof of Riemann's Existence Theorem, as discussed below.

GAGA allows one to pass from an object whose definition is inherently infinite in

nature (viz. an analytic space, where functions are defined in terms of limits) to one whose definition is finite in nature (viz. an algebraic variety, based on polynomials). Intuitively, the idea is that the result is stated only for spaces that are projective, and hence compact (in the metric topology); and this compactness provides the finiteness condition that permits us to pass from the analytic to the algebraic.

To make this more precise, let X be a complex algebraic variety, with the Zariski topology, and let $\mathcal{O} = \mathcal{O}_X$ be its structure sheaf — so that (X, \mathcal{O}) is a locally ringed space. Meanwhile, let X^h be the space X with the complex metric topology, and let $\mathcal{H} = \mathcal{H}_X$ be the corresponding structure sheaf, which assigns to any metric open set $U \subset X$ the ring $\mathcal{H}(U)$ of holomorphic functions on U . So (X^h, \mathcal{H}) is also a locally ringed space, called the complex analytic space associated to (X, \mathcal{O}) ; this is a Riemann surface if X is a complex algebraic curve.

As usual, we say that a sheaf \mathcal{F} of \mathcal{O} -modules on X is *coherent* if it is locally (in the Zariski topology) of the form $\mathcal{O}^m \rightarrow \mathcal{O}^n \rightarrow \mathcal{F} \rightarrow 0$; i.e. locally it is finitely generated, with finitely many relations. Similarly, we say that a sheaf \mathcal{F} of \mathcal{H} -modules on X^h is *coherent* if it is locally (in the complex topology) of the form $\mathcal{H}^m \rightarrow \mathcal{H}^n \rightarrow \mathcal{F} \rightarrow 0$. The main point of GAGA is that every coherent sheaf of \mathcal{H} -modules on X^h comes from a (unique) coherent sheaf of \mathcal{O} -modules on X , via a natural passage from \mathcal{O} -modules to \mathcal{H} -modules.

More precisely, we may associate, to any sheaf \mathcal{F} of \mathcal{O} -modules on X , a sheaf \mathcal{F}^h of \mathcal{H} -modules on X^h . Following [Se3], this is done as follows: First, let \mathcal{O}' be the sheaf of rings on X^h given by $\mathcal{O}'(U) = \varinjlim_V \mathcal{O}(V)$, where V ranges over the set Z_U of Zariski open subsets $V \subset X$ such that $V \supset U$. (For example, if U is an open disc in $\mathbb{P}_{\mathbb{C}}^1$, then $\mathcal{O}'(U)$ is the ring of rational functions with no poles in U .) Similarly, for every sheaf \mathcal{F} of \mathcal{O} -modules on X , we can define a sheaf \mathcal{F}' of \mathcal{O}' -modules on X^h by $\mathcal{F}'(U) = \varinjlim_V \mathcal{F}(V)$, where again V ranges over Z_U . Then define \mathcal{F}^h , a sheaf of \mathcal{H} -modules on X^h , by $\mathcal{F}^h(U) = \mathcal{F}'(U) \otimes_{\mathcal{O}'} \mathcal{H}$. For example, $\mathcal{O}^h = \mathcal{H}$. The assignment $\mathcal{F} \mapsto \mathcal{F}^h$ is an exact functor; so if $\mathcal{O}^m \rightarrow \mathcal{O}^n \rightarrow \mathcal{F} \rightarrow 0$ on a Zariski open subset U , then we also have $\mathcal{H}^m \rightarrow \mathcal{H}^n \rightarrow \mathcal{F}^h \rightarrow 0$. Thus if \mathcal{F} is coherent, then so is \mathcal{F}^h .

Theorem 2.2.1. (GAGA) [Se3] *Let X be a complex projective variety. Then the functor $\mathcal{F} \mapsto \mathcal{F}^h$, from the category of coherent \mathcal{O}_X -modules to the category of coherent \mathcal{H}_X -modules, is an equivalence of categories.*

There are two main ingredients in proving GAGA. The first of these (Theorem 2.2.2 below) is that that functor $\mathcal{F} \mapsto \mathcal{F}^h$ preserves cohomology. This result, due to Serre [Se3, §12, Théorème 1], allows one to pass back and forth more freely between the algebraic and analytic settings. Namely, on (X, \mathcal{O}) and (X^h, \mathcal{H}) , as on any locally ringed space, we can consider Čech cohomology of sheaves. In fact, given any topological space X , a sheaf of abelian groups \mathcal{F} on X , and an open covering $\mathcal{U} = \{U_\alpha\}$ of X , we can define the i th

Čech cohomology group $\check{H}^i(\mathcal{U}, \mathcal{F})$ as in [Hrt2, Chap III, §4]. We then define $H^i(X, \mathcal{F}) = \varinjlim_{\mathcal{U}} \check{H}^i(\mathcal{U}, \mathcal{F})$, where \mathcal{U} ranges over all open coverings of X in the given topology. For schemes X and coherent (or quasi-coherent) sheaves \mathcal{F} , this Čech cohomology agrees with the (derived functor) cohomology considered in Hartshorne [Hrt2, Chap. III, §2], because of [Hrt2, Chap. III, Theorem 4.5]. Meanwhile, for analytic spaces, Čech cohomology is the cohomology considered in Griffiths-Harris [GH, p.39], and also agrees with analytic derived functor cohomology (cf. [Hrt2, p.211]).

Theorem 2.2.2. [Se3] *Let X be a complex projective variety, and \mathcal{F} a coherent sheaf on X . Then the natural map $\varepsilon : H^q(X, \mathcal{F}) \rightarrow H^q(X^h, \mathcal{F}^h)$ is an isomorphism for every $q \geq 0$.*

The second main ingredient in the proof of GAGA is the following result of Serre and Cartan:

Theorem 2.2.3. *Let $X = \mathbb{P}_{\mathbb{C}}^r$ or $(\mathbb{P}_{\mathbb{C}}^r)^h$, and let \mathcal{M} be a coherent sheaf on X . Then for $n \gg 0$, the twisted sheaf $\mathcal{M}(n)$ is generated by finitely many global sections.*

In the algebraic case (i.e. for $X = \mathbb{P}_{\mathbb{C}}^r$), this is due to Serre, and is from his paper “FAC” [Se2]; cf. [Hrt2, Chap. II, Theorem 5.17]. In the analytic case (i.e. for $X = (\mathbb{P}_{\mathbb{C}}^r)^h$), this is Cartan’s “Theorem A” [Ca, exp. XVIII]; cf. [GH, p.700]. Recall that the condition that a sheaf \mathcal{F} is generated by finitely many global sections means that it is a quotient of a free module of finite rank; i.e. that there is a surjection $\mathcal{O}^N \twoheadrightarrow \mathcal{F}$ in the algebraic case, and $\mathcal{H}^N \twoheadrightarrow \mathcal{F}$ in the analytic case, for some integer $N > 0$ (where the exponent indicates a direct sum of N copies).

Proof of Theorem 2.2.1 (GAGA). The proof will rely on Theorems 2.2.2 and 2.2.3 above, the proofs of which will be discussed afterwards.

Step 1: We show that the functor $\mathcal{F} \rightarrow \mathcal{F}^h$ induces bijections on morphisms. That is, if \mathcal{F}, \mathcal{G} are coherent \mathcal{O}_X -modules, then the natural map $\phi : \text{Hom}_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G}) \rightarrow \text{Hom}_{\mathcal{H}_X}(\mathcal{F}^h, \mathcal{G}^h)$ is an isomorphism of groups.

To accomplish this, let $\mathcal{S} = \underline{\text{Hom}}_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G})$; i.e. \mathcal{S} is the sheaf of \mathcal{O}_X -modules associated to the presheaf $U \mapsto \text{Hom}_{\mathcal{O}_X}(\mathcal{F}(U), \mathcal{G}(U))$. Similarly, let $\mathcal{T} = \underline{\text{Hom}}_{\mathcal{H}_X}(\mathcal{F}^h, \mathcal{G}^h)$. There is then a natural morphism $\iota : \mathcal{S}^h \rightarrow \mathcal{T}$ of (sheaves of) \mathcal{H} -modules, inducing $\iota_* : H^0(X^h, \mathcal{S}^h) \rightarrow H^0(X^h, \mathcal{T})$. Let $\varepsilon : H^0(X, \mathcal{S}) \rightarrow H^0(X^h, \mathcal{S}^h)$ be as in Theorem 2.2.2 above. With respect to the identifications $H^0(X, \mathcal{S}) = \text{Hom}_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G})$ and $H^0(X^h, \mathcal{T}) = \text{Hom}_{\mathcal{H}_X}(\mathcal{F}^h, \mathcal{G}^h)$, the composition $\iota_* \varepsilon : \text{Hom}_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G}) \rightarrow \text{Hom}_{\mathcal{H}_X}(\mathcal{F}^h, \mathcal{G}^h)$ is the natural map ϕ taking $f \in \text{Hom}_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G})$ to $f^h \in \text{Hom}_{\mathcal{H}_X}(\mathcal{F}^h, \mathcal{G}^h)$. We want to show that ϕ is an isomorphism. Since ε is an isomorphism (by Theorem 2.2.2), it suffices to show that ι_* is also — which will follow from showing that $\iota : \mathcal{S}^h \rightarrow \mathcal{T}$ is an isomorphism. That in turn can be checked on stalks. Here, the stalks of \mathcal{S} , \mathcal{S}^h , and \mathcal{T} at a point ξ are given by

$$\mathcal{S}_{\xi} = \text{Hom}_{\mathcal{O}_{X,\xi}}(\mathcal{F}_{\xi}, \mathcal{G}_{\xi}), \quad \mathcal{S}_{\xi}^h = \text{Hom}_{\mathcal{O}_{X,\xi}}(\mathcal{F}_{\xi}, \mathcal{G}_{\xi}) \otimes_{\mathcal{O}_{X,\xi}} \mathcal{H}_{X,\xi},$$

$$\mathcal{T}_\xi = \mathrm{Hom}_{\mathcal{H}_{X,\xi}}(\mathcal{F}_\xi \otimes_{\mathcal{O}_{X,\xi}} \mathcal{H}_\xi, \mathcal{G}_\xi \otimes_{\mathcal{O}_{X,\xi}} \mathcal{H}_\xi).$$

Now the local ring $\mathcal{H}_{X,\xi}$ is flat over $\mathcal{O}_{X,\xi}$, since the inclusion $\mathcal{O}_{X,\xi} \hookrightarrow \mathcal{H}_{X,\xi}$ becomes an isomorphism upon completion (with both rings having completion $\mathbb{C}[[x]]$, where x is a local uniformizer at ξ). So by [Bo, I, §2.10, Prop. 11], we may “pull the tensor across the Hom” here; i.e. $\mathcal{S}_\xi^{\mathrm{h}} \rightarrow \mathcal{T}_\xi$ is an isomorphism.

Step 2: We show that the functor $\mathcal{F} \rightarrow \mathcal{F}^{\mathrm{h}}$ is essentially surjective, i.e. is surjective on isomorphism classes (and together with Step 1, this implies that it is bijective on isomorphism classes).

First we reduce to the case $X = \mathbb{P}_{\mathbb{C}}^r$, by taking an embedding $j : X \hookrightarrow \mathbb{P}_{\mathbb{C}}^r$, considering the direct image sheaf $j_*\mathcal{F}$ on $\mathbb{P}_{\mathbb{C}}^r$, and using that $(j_*\mathcal{F})^{\mathrm{h}}$ is canonically isomorphic to $j_*(\mathcal{F}^{\mathrm{h}})$. Next, say that \mathcal{M} is a coherent sheaf on X^{h} , i.e. a coherent \mathcal{H} -module. By Theorem 2.2.3, there is a surjection $\mathcal{H}^M \rightarrow \mathcal{M}(m) \rightarrow 0$ for some integers m, M ; and so $\mathcal{H}(-m)^M \rightarrow \mathcal{M} \rightarrow 0$. Let the sheaf \mathcal{N} be the kernel of this latter surjection. Then \mathcal{N} is a coherent \mathcal{H} -module, and so there is similarly a surjection $\mathcal{H}(-n)^N \rightarrow \mathcal{N} \rightarrow 0$ for some n, N . Combining, we have an exact sequence $\mathcal{H}(-n)^N \xrightarrow{g} \mathcal{H}(-m)^M \rightarrow \mathcal{M} \rightarrow 0$. Now $\mathcal{H}(-n)^N = (\mathcal{O}(-n)^N)^{\mathrm{h}}$ and $\mathcal{H}(-m)^M = (\mathcal{O}(-m)^M)^{\mathrm{h}}$. So by Step 1, $g = f^{\mathrm{h}}$ for some $f \in \mathrm{Hom}(\mathcal{O}(-n)^N, \mathcal{O}(-m)^M)$. Let $\mathcal{F} = \mathrm{cok}f$. So $\mathcal{O}(-n)^N \xrightarrow{f} \mathcal{O}(-m)^M \rightarrow \mathcal{F} \rightarrow 0$ is exact, and hence so is $\mathcal{H}(-n)^N \xrightarrow{g} \mathcal{H}(-m)^M \rightarrow \mathcal{F}^{\mathrm{h}} \rightarrow 0$, using $g = f^{\mathrm{h}}$. Thus $\mathcal{M} \approx \mathcal{F}^{\mathrm{h}}$. \square

Having proven GAGA, we now use it to finish the proof of Riemann’s Existence Theorem for complex algebraic curves X . Two steps are needed. The first is to pass from an assertion about *modules* over a projective curve (or a projective variety) X to an assertion about *branched covers* of X . The second step is to pass from branched covers of a projective curve X to (unramified) covering spaces over a Zariski open subset of X .

For the first of these steps, observe that the equivalence between algebraic and analytic coherent modules, stated in GAGA, automatically implies the corresponding equivalence between algebraic and analytic coherent *algebras* (i.e. sheaves of algebras that are coherent as sheaves of modules). The reason is that an R -algebra A is an R -module together with some additional structure, given by module homomorphisms (viz. a product map $\mu : A \otimes_R A \rightarrow A$ and an identity $1 : R \rightarrow A$) and relations which can be given by commutative diagrams (corresponding to the associative, commutative, distributive, and identity properties); and the same holds locally for sheaves of algebras. The equivalence of categories $\mathcal{F} \mapsto \mathcal{F}^{\mathrm{h}}$ in GAGA is compatible with tensor product (i.e. it is an equivalence of tensor categories); so the additional algebra structure carries over under the equivalence — and thus the analog of GAGA for coherent algebras holds. Under this equivalence, generically separable \mathcal{O}_X -algebras correspond to generically separable \mathcal{H}_X -algebras (using that $\mathcal{H}(U)$ is faithfully flat over $\mathcal{O}(U)$ for a Zariski open subset $U \subset X$, because the inclusion of stalks becomes an isomorphism upon completion). So taking spectra, we also obtain an equivalence between *algebraic branched covers* and *analytic branched covers*.

This formal argument can be summarized informally in the following

General Principle 2.2.4. *An equivalence of tensor categories of modules induces a corresponding equivalence of categories of algebras, of branched covers, and of Galois branched covers for any given finite Galois group.*

The last point (about Galois covers) holds because an equivalence of categories between covers automatically preserves the Galois group.

In order to obtain Riemann’s Existence Theorem, one more step is needed, viz. passage from branched covers of a curve X to étale (or unramified) covers of an open subset of X . For this, recall that an algebraic branched cover is locally a covering space (in the metric topology) precisely where it is étale, by the Inverse Function Theorem. Conversely, an étale cover of a Zariski open subset of X extends to an algebraic branched cover of X (by taking the normalization in the function field of the cover). Such an extension also exists for analytic covers of curves, since it exists *locally* over curves. (Namely, a finite covering space of the punctured disc $0 < |z| < 1$ extends to an analytic branched cover of the disc $|z| < 1$, since the covering map — being bounded and holomorphic — has a removable singularity [Ru, Theorem 10.20].) Thus the above equivalence for branched covers induces an equivalence of categories between finite étale covers of a smooth complex algebraic curve X , and finite analytic covering maps to X^h . That is, the categories (i) and (ii) in Riemann’s Existence Theorem are equivalent; and this completes the proof of that theorem.

Apart from Riemann’s Existence Theorem, GAGA has a number of other applications, including several proven in [Se3]. Serre showed there that if V is a smooth projective variety over a number field K , and if X is the complex variety obtained from V via an embedding $j : K \hookrightarrow \mathbb{C}$, then the Betti numbers of X are independent of the choice of j [Se3, Cor. to Prop. 12]. Serre also used GAGA to obtain a proof of Chow’s Theorem [Ch] that every closed analytic subset of $\mathbb{P}_{\mathbb{C}}^n$ is algebraic [Se3, Prop. 13], as well as several corollaries of that result. In addition, he showed that if X is a projective algebraic variety, then the natural map $H^1(X, \mathrm{GL}_n(\mathbb{C})) \rightarrow H^1(X^h, \mathrm{GL}_n(\mathbb{C})^h)$ is bijective [Se3, Prop. 18]. As a consequence, the set of isomorphism classes of rank n algebraic vector bundles over X (in the Zariski topology) is in natural bijection with the set of isomorphism classes of rank n analytic vector bundles over X^h (in the metric topology). In a way, this is surprising, since the corresponding assertion for *covers* is false (because all covering spaces in the Zariski topology are trivial, over an irreducible complex variety).

Having completed the proofs of GAGA and Riemann’s Existence Theorem, we return to the proofs of Theorems 2.2.2 and 2.2.3.

Proof of Theorem 2.2.2. First we reduce to the case $X = \mathbb{P}_{\mathbb{C}}^r$ as in Step 2 of the proof of Theorem 2.2.1, using that $H^q(X, \mathcal{F}) = H^q(\mathbb{P}_{\mathbb{C}}^r, j_*\mathcal{F})$ if $j : X \hookrightarrow \mathbb{P}_{\mathbb{C}}^r$, and similarly for X^h .

Second, we verify the result directly for the case $\mathcal{F} = \mathcal{O}$ and $\mathcal{F}^h = \mathcal{H}$, for all $q \geq 0$. The case $q = 0$ is clear, since then both sides are just \mathbb{C} , because X is projective (and hence compact). On the other hand if $q > 0$, then $H^q(X, \mathcal{O}) = 0$ by the (algebraic) cohomology of projective space [Hrt2, Chap. III, Theorem 5.1], and $H^q(X^h, \mathcal{H}) = 0$ via Dolbeault's Theorem [GH, p.45].

Third, we verify the result for the sheaf $\mathcal{O}(n)$ on $X = \mathbb{P}_{\mathbb{C}}^r$. This step uses induction on the dimension r , where the case $r = 0$ is trivial. Assuming the result for $r - 1$, we need to show it for r . This is done by induction on $|n|$; for ease of presentation, assume $n > 0$ (the other case being similar). Let E be a hyperplane in $\mathbb{P}_{\mathbb{C}}^r$; thus $E \approx \mathbb{P}_{\mathbb{C}}^{r-1}$. Tensoring the exact sequence $0 \rightarrow \mathcal{O}(-1) \rightarrow \mathcal{O} \rightarrow \mathcal{O}_E \rightarrow 0$ with $\mathcal{O}(n)$, we obtain an associated long exact sequence $(*)$ which includes, in part:

$$H^{q-1}(E, \mathcal{O}_E(n)) \rightarrow H^q(X, \mathcal{O}(n-1)) \rightarrow H^q(X, \mathcal{O}(n)) \rightarrow H^q(E, \mathcal{O}_E(n)) \rightarrow H^q(X, \mathcal{O}(n-1))$$

Similarly, replacing \mathcal{O} by \mathcal{H} , we obtain an analogous long exact sequence $(*)^h$; and there are (commuting) maps ε from each term in $(*)$ to the corresponding term in $(*)^h$. By the inductive hypotheses, the map ε is an isomorphism on each of the outer four terms above. So by the Five Lemma, ε is an isomorphism on $H^q(X, \mathcal{O}(n))$.

Fourth, we handle the general case. By a vanishing theorem of Grothendieck ([Gr1]; see also [Hrt2, Chap. III, Theorem 2.7]), the q th cohomology vanishes for a Noetherian topological space of dimension n if $q > n$. (Cf. [Hrt2, p.5] for the definition of *dimension*.) So we can proceed by descending induction on q . Since \mathcal{F} is coherent, it is a quotient of a sheaf $\mathcal{E} = \bigoplus_i \mathcal{O}(n_i)$ [Hrt2, Chap. II, Cor. 5.18], say with kernel \mathcal{N} . The associated long exact sequence includes, in part:

$$H^q(X, \mathcal{N}) \rightarrow H^q(X, \mathcal{E}) \rightarrow H^q(X, \mathcal{F}) \rightarrow H^{q+1}(X, \mathcal{N}) \rightarrow H^{q+1}(X, \mathcal{E})$$

The (commuting) homomorphisms ε map from these terms to the corresponding terms of the analogous long exact sequence of coherent \mathcal{H} -modules on X^h . On the five terms above, the second map ε is an isomorphism by the previous step; and the fourth and fifth maps ε are isomorphisms by the descending inductive hypothesis. So by the Five Lemma, the middle ε map is surjective. This gives the surjectivity part of the result, for an arbitrary coherent sheaf \mathcal{F} . In particular, surjectivity holds with \mathcal{F} replaced by \mathcal{N} . That is, on the first of the five terms in the exact sequence above, the map ε is surjective. So by the Five Lemma, the middle ε is injective; so it is an isomorphism. \square

Concerning Theorem 2.2.3, that result is equivalent to the following assertion:

Theorem 2.2.5. *Let $X = \mathbb{P}_{\mathbb{C}}^r$ or $(\mathbb{P}_{\mathbb{C}}^r)^h$, and let \mathcal{M} be a coherent sheaf on X . Then there is an n_0 such that for all $n \geq n_0$ and all $q > 0$, we have $H^q(X, \mathcal{M}(n)) = 0$.*

In the algebraic setting, Theorem 2.2.5 is due to Serre; cf. [Hrt2, Chap. III, Theorem 5.2]. In the analytic setting, this is Cartan’s “Theorem B” ([Se1], exp. XVIII of [Ca2]); cf. also [GH, p.700].

The proof of Theorem 2.2.3 is easier in the algebraic situation than in the analytic one. In the former case, the proof proceeds by choosing generators of stalks \mathcal{M}_ξ ; multiplying each by an appropriate monomial to get a global section of some $\mathcal{M}(n)$; and using quasi-compactness to require only finitely many sections overall (also cf. [Hrt2, Chapter II, proof of Theorem 5.17]). But this strategy fails in the analytic case because the local sections are not rational, or even meromorphic; and so one cannot simply clear denominators to get a global section of a twisting of \mathcal{M} .

The proof in the analytic case proves Cartan’s Theorems A and B (i.e. 2.2.3 and 2.2.5) together, by induction on r . Denoting these assertions in dimension r by (A_r) and (B_r) , the proof in ([Se1], exp. XIX of [Ca3]) proceeds by showing that $(A_{r-1}) + (B_{r-1}) \Rightarrow (A_r)$ and that $(A_r) \Rightarrow (B_r)$. Since the results are trivial for $r = 0$, the two theorems then follow; and as a result, GAGA follows as well. Serre’s later argument in [Se3] is a variant on this inductive proof that simultaneously proves GAGA and Theorems A and B (i.e. Theorems 2.2.1, 2.2.3, and 2.2.5 above).

Theorems A and B were preceded by a non-projective version of those results, viz. for polydiscs in \mathbb{C}^r , and more generally for Stein spaces (exp. XVIII and XIX of [Ca2]; cf. also [GuR, pp. 207, 243]). There too, the two theorems are essentially equivalent. Also, no twisting is needed for Theorem B in the earlier version because the spaces were not projective there.

The proof of Theorem A in this earlier setting uses an “analytic patching” argument, applied to overlapping compact sets K', K'' on a Stein space X . In that situation, one considers metric neighborhoods U', U'' of K', K'' respectively, and one chooses generating sections $f'_1, \dots, f'_k \in \mathcal{M}(U')$ and $f''_1, \dots, f''_k \in \mathcal{M}(U'')$ for the given sheaf \mathcal{M} on U', U'' respectively. From this data, one produces generating sections $g_1, \dots, g_k \in \mathcal{M}(U)$, where U is an open neighborhood of $K = K' \cup K''$. This is done via Cartan’s Lemma on matrix factorization, which says (for appropriate choice of K', K'') that every element $A \in GL_n(K' \cap K'')$ can be factored as a product of an element $B \in GL_n(K')$ and an element $C \in GL_n(K'')$. That lemma, which can be viewed as a multiplicative matrix analog of Cousin’s Theorem [GuRo, p.32], had been proven earlier in [Ca1], with this application in mind; and a special case had been shown even earlier in [Bi]. See also [GuRo, Chap. VI, §E]. (Cartan’s Lemma is also sometimes called Cartan’s “attaching theorem”, where *attaching* is used in essentially the same sense as *patching* here.)

Cartan’s Lemma can be used to prove these earlier versions of Theorems A and B by taking bases f'_i and f''_i over U' and U'' , and letting A be the transition matrix between them (i.e. $\vec{f}' = A\vec{f}''$, where \vec{f}' and \vec{f}'' are the column vectors with entries f'_i and f''_i

respectively). The generators g_i as above can then be defined as the sections that differ from the f' 's by B^{-1} and from the f'' 's by C (i.e. $\vec{g} = B^{-1}\vec{f} = C\vec{f}''$). The g_i 's are then well-defined wherever either the f' 's or f'' 's are — and hence in a neighborhood of $K = K' \cup K''$. This matrix factorization strategy also appears elsewhere, e.g. classically, concerning the Riemann-Hilbert problem, in which one attempts to find a system of linear differential equations whose monodromy representation of the fundamental group is a given representation (this being a differential analog of the inverse Galois problem for covers).

This use of Cartan's Lemma also suggests another way of restating GAGA, in the case where a projective variety X is covered by two open subsets X_1, X_2 that are strictly contained in X . The point is that if one gives coherent analytic (sheaves of) modules over X_1 and over X_2 together with an isomorphism on the overlap, then there is a unique coherent *algebraic* module over X that induces the given data compatibly. Of course by definition of coherent sheaves, there is such an *analytic* module over X (and similarly, we can always reduce to the case of two metric open subsets X_1, X_2); but the assertion is that it is algebraic.

To state this compactly, we introduce some categorical terminology. If $\mathcal{A}, \mathcal{B}, \mathcal{C}$ are categories, with functors $f : \mathcal{A} \rightarrow \mathcal{C}$ and $g : \mathcal{B} \rightarrow \mathcal{C}$, then the *2-fibre product* of \mathcal{A} and \mathcal{B} over \mathcal{C} (with respect to f, g) is the category $\mathcal{A} \times_{\mathcal{C}} \mathcal{B}$ in which an *object* is a pair $(A, B) \in \mathcal{A} \times \mathcal{B}$ together with an isomorphism $\iota : f(A) \simeq g(B)$ in \mathcal{C} ; and in which a *morphism* $(A, B; \iota) \rightarrow (A', B'; \iota')$ is a pair of morphisms $A \rightarrow A'$ and $B \rightarrow B'$ that are compatible with the ι 's. For any variety [resp. analytic space] X , let $\mathfrak{M}(X)$ denote the category of algebraic [resp. analytic] coherent modules on X . (Similarly, for any ring R , we write $\mathfrak{M}(R)$ for the category of finitely presented R -modules. This is the same as $\mathfrak{M}(\text{Spec } R)$.) In this language, GAGA and its generalizations to algebras and covers can be restated as:

Theorem 2.2.6. *Let X be a complex projective algebraic variety, with metric open subsets X_1, X_2 such that $X = X_1 \cup X_2$; let X_0 be their intersection. Then the natural base change functor*

$$\mathfrak{M}(X) \rightarrow \mathfrak{M}(X_1) \times_{\mathfrak{M}(X_0)} \mathfrak{M}(X_2)$$

is an equivalence of categories. Moreover the same holds if \mathfrak{M} is replaced by the category of finite algebras, or of finite branched covers, or of Galois covers with a given Galois group.

Here X is regarded as an algebraic variety, and the X_i 's as analytic spaces (so that the left hand side of the equivalence consists of algebraic modules, and the objects on the right hand side consist of analytic modules). In the case of curves, each X_i is contained in an affine open subset U_i , so coherent sheaves of modules on X_i can be identified with coherent modules over the ring $\mathcal{H}(X_i)$; thus we may identify the categories $\mathfrak{M}(X_i)$ and $\mathfrak{M}(\mathcal{H}(X_i))$.

The approach in Theorem 2.2.6 will be useful in considering analogs of GAGA in Sections 3 and 4 below.

Section 2.3. Complex patching and constructing covers.

Consider a Zariski open subset of the Riemann sphere, say $U = \mathbb{P}_{\mathbb{C}}^1 - \{\xi_1, \dots, \xi_r\}$. By Riemann's Existence Theorem, every finite covering space of U is given by an étale morphism of complex algebraic curves. Equivalently, every finite branched cover of $\mathbb{P}_{\mathbb{C}}^1$, branched only at $S = \{\xi_1, \dots, \xi_r\}$, is given by a finite dominating morphism from a smooth complex projective curve Y to $\mathbb{P}_{\mathbb{C}}^1$. As discussed in Section 2.1, passage from topological to analytic covers is the easier step, but it requires knowledge of what topological covering spaces exist (essentially via knowledge of the fundamental group, which is understood via loops). Passage from analytic covers to algebraic covers is deeper, and can be achieved using GAGA, as discussed in Section 2.2.

Here we consider how covers can be constructed from this point of view using complex analytic patching, keeping an eye on possible generalizations. In particular, we raise the question of how to use these ideas to understand covers of curves that are not defined over the complex numbers.

We begin by elaborating on the bijection described in Corollary 2.1.2.

Taking $U = \mathbb{P}_{\mathbb{C}}^1 - \{\xi_1, \dots, \xi_r\}$ as above, choose a base point $\xi_0 \in U$. The fundamental group $\pi_1(U, \xi_0)$ is then the group $\langle x_1, \dots, x_r \mid x_1 \cdots x_r = 1 \rangle$, as discussed in Section 2.1. Up to isomorphism, the fundamental group is independent of the choice of ξ_0 , and so the mention of the base point is often suppressed; but fixing a base point allows us to analyze the fundamental group more carefully. Namely, we may choose a “bouquet of loops” at ξ_0 (in M. Fried's terminology [Fr1]), consisting of a set of counterclockwise loops $\sigma_1, \dots, \sigma_r$ at ξ_0 , where σ_j winds once around ξ_j and winds around no other ξ_k ; where the support of the σ_j 's are disjoint except at ξ_0 ; where $\sigma_1 \cdots \sigma_r$ is homotopic to the identity; and where the homotopy classes of the σ_j 's (viz. the x_j 's) generate $\pi_1(U, \xi_0)$. In particular, we can choose σ_j to consist of a path ϕ_j from ξ_0 to a point ξ_j' near ξ_j , followed by a counterclockwise loop λ_j around ξ_j , followed by ϕ_j^{-1} . The term “bouquet” is natural with this choice of loops (e.g. in the case that $\xi_0 = 0$ and $\xi_j = e^{j\pi i/r}$, with $j = 1, \dots, r$, and where each ϕ_j is a line segment from ξ_0 to $(1 - \varepsilon)\xi_j$ for some small positive value of ε).

Let $f : V \rightarrow U$ be a finite Galois covering space, say with Galois group G . Then $\pi_1(V)$ is a subgroup N of finite index in $\pi_1(U)$, and $G = \pi_1(U)/N$. Let $g_1, \dots, g_r \in G$ be the images of $x_1, \dots, x_r \in \pi_1(U)$, and let m_j be the order of g_j . Thus (g_1, \dots, g_r) is the branch cycle description of $V \rightarrow U$; i.e. the G -Galois cover $V \rightarrow U$ corresponds to the uniform conjugacy class of (g_1, \dots, g_r) in Corollary 2.1.2. By Riemann's Existence Theorem 2.1.1, the cover $V \rightarrow U$ can be given by polynomial equations and regarded as a finite étale cover. Taking the normalization of $\mathbb{P}_{\mathbb{C}}^1$ in V , we obtain a smooth projective curve Y containing V as a Zariski open subset, and a G -Galois connected branched covering map $f : Y \rightarrow \mathbb{P}_{\mathbb{C}}^1$ which is branched only over $S = \{\xi_1, \dots, \xi_r\}$.

In the above notation, with $\sigma_j = \phi_j \lambda_j \phi_j^{-1}$ (and multiplying paths from left to right),

we can extend ϕ_j to a path ψ_j from ξ_0 to ξ_j in $\mathbb{P}_{\mathbb{C}}^1$. The path ψ_j can be lifted to a path $\tilde{\psi}_j$ in Y from a base point $\eta_0 \in Y$ over ξ_0 , to a point $\eta_j \in Y$ over ξ_j . The element g_j generates the inertia group A_j of η_j (i.e. the stabilizer of η_j in the group G). If X_j is a simply connected open neighborhood of ξ_j that contains no other ξ_k , then the topological fundamental group of $X_j - \{\xi_j\}$ is isomorphic to \mathbb{Z} . So $f^{-1}(X_j)$ is a union of homeomorphic connected components, each of which is Galois and cyclic of order m_j over X_j , branched only at ξ_j . The component Y_j of $f^{-1}(X_j)$ that contains η_j has stabilizer $A_j = \langle g_j \rangle \subset G$, and by Kummer theory it is given by an equation of the form $s_j^{m_j} = t_j$, if t_j is a uniformizer on X_j at ξ_j . Moreover g_j acts by $g_j(s_j) = e^{2\pi i/m_j} s_j$. So $f^{-1}(X_j)$ is a (typically disconnected) G -Galois cover of X_j , consisting of a disjoint union of copies of the m_j -cyclic cover $Y_j \rightarrow X_j$, indexed by the left cosets of A_j in G . We say that $f^{-1}(X_j)$ is the G -Galois branched cover of X_j that is *induced* by the A_j -Galois cover $Y_j \rightarrow X_j$; and we write $f^{-1}(X_j) = \text{Ind}_{A_j}^G Y_j$. Similarly, if U' is a simply connected open subset of U (and so U' does not contain any branch points ξ_j), then $f^{-1}(U')$ is the trivial G -Galois cover of U' , consisting of $|G|$ copies of U' permuted simply transitively by the elements of G ; this cover is $\text{Ind}_1^G U'$.

Since the complex affine line is simply connected, the smallest example of the above situation is the case $r = 2$. By a projective linear change of variables, we may assume that the branch points are at $0, \infty$. The fundamental group of $U = \mathbb{P}_{\mathbb{C}}^1 - \{0, \infty\}$ is infinite cyclic, so a finite étale cover is cyclic, say with Galois group C_m ; and the cover has branch cycle description $(g, g^{-1}) = (g, g^{m-1})$ for some generator g of the cyclic group C_m . This cover is given over U by the single equation $y^m = x$. So no patching is needed in this case. (If we instead take two branch points $x = c_0, x = c_1$, with $c_0, c_1 \in \mathbb{C}$, then the equation is $y^m = (x - c_0)(x - c_1)^{m-1}$ over $\mathbb{P}_{\mathbb{C}}^1$ minus the two branch points.)

The next simplest case is that of $r = 3$. This is the first really interesting case, and in fact it is key to understanding cases with $r > 3$. By a projective linear transformation we may assume that the branch locus is $\{0, 1, \infty\}$. We consider this case next in more detail:

Example 2.3.1. We give a “recipe” for constructing Galois covers of $U = \mathbb{P}_{\mathbb{C}}^1 - \{0, 1, \infty\}$ via patching, in terms of the branch cycle description of the given cover.

The topological fundamental group of U is $\langle \alpha, \beta, \gamma \mid \alpha\beta\gamma = 1 \rangle$, and this is isomorphic to the free group on two generators, viz. α, β . If we take $z = 1/2$ as the base point for the fundamental group, then these generators can be taken to be counterclockwise loops at $1/2$ around $0, 1$, respectively. The paths ψ_0, ψ_1 as above can be taken to be the real line segments connecting the base point to $0, 1$ respectively, and ψ_{∞} can be taken to be the vertical path from $1/2$ to “ $1/2 + i\infty$ ”.

Let G be a finite group generated by two elements a, b . Let $c = (ab)^{-1}$, so that $abc = 1$. Consider the connected G -Galois covering space $f : V \rightarrow U$ with branch cycle description (a, b, c) , and the corresponding branched cover $Y \rightarrow P_{\mathbb{C}}^1$ branched at S . As above, after choosing a base point $\eta \in Y$ over $1/2 \in \mathbb{P}_{\mathbb{C}}^1$ and lifting the paths ψ_j , we obtain points $\eta_0,$

η_1, η_∞ over $0, 1, \infty$, with cyclic stabilizers $A_0 = \langle a \rangle, A_1 = \langle b \rangle, A_\infty = \langle c \rangle$ respectively. Let ι be the path in $\mathbb{P}_\mathbb{C}^1$ from 0 to 1 corresponding to the real interval $[0, 1]$, and let $\tilde{\iota}$ be the unique path in Y that lifts ι and passes through η . Observe that the initial point of $\tilde{\iota}$ is η_0 , and the final point is η_1 .

Consider the simply connected neighborhoods $X_0 = \{z \in \mathbb{C} \mid \operatorname{Re} z < 2/3\}$ of 0, and $X_1 = \{z \in \mathbb{C} \mid \operatorname{Re} z > 1/3\}$ of 1. We have that $X_0 \cup X_1 = \mathbb{C}$, and $U' := X_0 \cap X_1$ is contained in U . Also, $U = U_0 \cup U_1$, where $U_j = X_j - \{j\}$ for $j = 0, 1$. By the above discussion, $f^{-1}(X_0) = \operatorname{Ind}_{A_0}^G Y_0$, where $Y_0 \rightarrow X_0$ is a cyclic cover branched only at 0, and given by the equation $y_0^m = x$ (where m is the order of a). Similarly $f^{-1}(X_1) = \operatorname{Ind}_{A_1}^G Y_1$, where the branched cover $Y_1 \rightarrow X_1$ is given by $y_1^n = x - 1$ (where n is the order of b). Since the overlap $U' = X_0 \cap X_1$ does not meet the branch locus S , we have that $f^{-1}(U')$ is the trivial G -Galois cover $\operatorname{Ind}_1^G U'$. These induced covers have connected components that are respectively indexed by the left cosets of $A_0, A_1, 1$; and the identity coset corresponds to the component respectively containing η_0, η_1, η . Observe that the identity component of $\operatorname{Ind}_1^G U'$ is contained in the identity components of the other two induced covers, because $\tilde{\iota}$ passes through η_0, η_1, η .

Turning this around, we obtain the desired “patching recipe” for constructing the G -Galois cover of U with given branch cycle description (a, b, c) : Over the above open sets U_0 and U_1 , take the induced covers $\operatorname{Ind}_{A_0}^G V_0$ and $\operatorname{Ind}_{A_1}^G V_1$, where $V_0 \rightarrow U_0$ and $V_1 \rightarrow U_1$ are respectively given by $y_0^m = x$ and $y_1^n = x - 1$, and where $A_0 = \langle a \rangle, A_1 = \langle b \rangle$. Pick a point η over $1/2$ on the identity components of each of these two induced covers; thus $g(\eta)$ is a well-defined point on each of these induced covers, for any $g \in G$. The induced covers each restrict to the trivial G -Galois cover on the overlap $U' = U_0 \cap U_1$; now paste together the components of these trivial covers by identifying, for each $g \in G$, the component of $\operatorname{Ind}_{A_0}^G V_0$ containing $g(\eta)$ with the component of $\operatorname{Ind}_{A_1}^G V_1$ containing that point. The result is the desired cover $V \rightarrow U$. \square

The above example begins with a group G and a branch cycle description (a, b, c) , and constructs the cover $V \rightarrow U = \mathbb{P}_\mathbb{C}^1 - \{0, 1, \infty\}$ with that branch cycle description. In doing so, it gives the cover locally in terms of equations over two topological open discs U_0 and U_1 , and instructions for patching on the overlap. Thus it gives the cover analytically (not algebraically, since the U_i 's are not Zariski open subsets).

The simplest specific instance of the above example uses the cyclic group $C_3 = \langle g \rangle$ of order 3, and branch cycle description (g, g, g) . Over U_0 the cover is given by (one copy of) $y_0^3 = x$; and over U_1 it is given by $y_1^3 = x - 1$. Here, over U_i , the generator g acts by $g(y_i) = \zeta_3 y_i$, where $\zeta_3 = e^{2\pi i/3}$. By GAGA, the cover can be described algebraically, i.e. by polynomials over Zariski open sets. And in this particular example, this can even be done globally over U , by the single equation $z^3 = x(x - 1)$ (where $g(z) = \zeta_3 z$). Here $z = y_0 f_0(x)$ on U_0 , where $f_0(x)$ is the holomorphic function on U_0 such that $f_0(0) = -1$

and $f_0^3 = x - 1$; explicitly, $f_0(x) = -1 + \frac{1}{3}x + \frac{1}{9}x^2 + \dots$ in a neighborhood of $x = 0$. Similarly, $z = y_1 f_1(x)$ on U_1 , where $f_1(x)$ is the holomorphic function on U_1 such that $f_1(1) = 1$ and $f_1^3 = x$; here $f_1(x) = 1 + \frac{1}{3}(x - 1) - \frac{1}{9}(x - 1)^2 + \dots$ in a neighborhood of $x = 1$. (Note that for this very simple cover, the global equation can be written down by inspection. But in general, for non-abelian groups, the global polynomial equations are not at all obvious from the local ones, though by GAGA they must exist.)

Example 2.3.1 requires GAGA in order to pass from the analytic equations (locally, on metric open subsets) to algebraic equations that are valid on a Zariski open dense subset. In addition, it uses ideas of topology — in particular, knowledge of the fundamental group, and the existence of open sets that overlap and together cover the space U . In later sections of this paper, we will discuss the problem of performing analogous constructions over fields other than \mathbb{C} , in order to understand covers of algebraic curves over those fields. For that, we will see that often an analog of GAGA exists — and that analog will permit passage from “analytic” covers to algebraic ones. A difficulty that has not yet been overcome, however, is how to find analogs of the notions from topology — both regarding explicit descriptions of fundamental groups and regarding the need for having overlapping open sets (which in non-archimedean contexts do not exist in a non-trivial way). One way around this problem is to consider only certain types of covers, for which GAGA alone suffices (i.e. where the information from topology is not required). The next example illustrates this.

Example 2.3.2. Let G be a finite group, with generators g_1, \dots, g_r (whose product need not be 1). Let $S = \{\xi_1, \dots, \xi_{2r}\}$ be a set of $2r$ distinct points in $\mathbb{P}_{\mathbb{C}}^1$, and consider the G -Galois covering space $V \rightarrow U = \mathbb{P}_{\mathbb{C}}^1 - S$ with branch cycle description

$$(g_1, g_1^{-1}, g_2, g_2^{-1}, \dots, g_r, g_r^{-1}), \quad (*)$$

with respect to a bouquet of loops $\sigma_1, \dots, \sigma_{2r}$ at a base point $\xi_0 \in U$. Let $Y \rightarrow \mathbb{P}_{\mathbb{C}}^1$ be the corresponding branched cover. This cover is well defined since the product of the entries of $(*)$ is 1, and it is connected since the entries of $(*)$ generate G . The cover can be obtained by a “cut-and-paste” construction as follows: Choose disjoint simple (i.e. non-self-intersecting) paths s_1, \dots, s_r in $\mathbb{P}_{\mathbb{C}}^1$, where s_j begins at ξ_{2j-1} and ends at ξ_{2j} . Take $|G|$ distinct copies of $\mathbb{P}_{\mathbb{C}}^1$, indexed by the elements of G . Redefine the topology on the disjoint union of these copies by identifying the right hand edge of a “slit” along s_j on the g th copy of $\mathbb{P}_{\mathbb{C}}^1$ to the left hand edge of the “slit” along s_j on the gg_j th copy of $\mathbb{P}_{\mathbb{C}}^1$ (with the orientation as one proceeds along the slits). The resulting space maps to $\mathbb{P}_{\mathbb{C}}^1$ in the obvious way, and away from S it is the G -Galois covering space of $\mathbb{P}_{\mathbb{C}}^1 - S$ with branch cycle description $(*)$. Because of this construction, we will call covers of this type *slit covers* [Ha1, 2.4]. (The corresponding branch cycle descriptions $(*)$ have been referred to as “Harbater-Mumford representatives” [Fr3].)

Now choose disjoint simply connected open subsets $X_j \subset \mathbb{P}_{\mathbb{C}}^1$ for $j = 1, \dots, r$, such that $\xi_{2j-1}, \xi_{2j} \in X_j$. (If ξ_{2j-1} and ξ_{2j} are sufficiently close for all j , relative to their distances

to the other ξ_k 's, then the X_j 's can be taken to be discs.) In the above cut-and-paste construction, the paths s_1, \dots, s_r can be chosen so that the support of s_j is contained in X_j , for each j . Each X_j contains a strictly smaller simply connected open set X_j^* (e.g. a smaller disc) which also contains the support of s_j , and whose closure \bar{X}_j^* is contained in X_j . Let $U' = \mathbb{P}_{\mathbb{C}}^1 - \bigcup \bar{X}_j^*$. In the cut-and-paste construction of $V \rightarrow U$, we have that the topology of the disjoint union of the $|G|$ copies of $\mathbb{P}_{\mathbb{C}}^1$ is unaffected outside of the union of the X_j^* 's; and so the restriction of $V \rightarrow U$ to U' is a trivial cover, viz. $\text{Ind}_1^G U'$. Suppose that ξ_j is not the point $x = \infty$ on $\mathbb{P}_{\mathbb{C}}^1$; thus ξ_j corresponds to a point $x = c_j$, with $c_j \in \mathbb{C}$. Let m_j be the order of g_j , and let A_j be the subgroup of G generated by g_j . Then the restriction of $V \rightarrow U$ to $U_j = X_j \cap U$ is given by $\text{Ind}_{A_j}^G V_j$, where $V_j \rightarrow U_j$ is the A_j -Galois étale cover given by $y_j^{m_j} = (x - c_{2j-1})(x - c_{2j})^{m_j-1}$ (as in the two branch point case, discussed just before Example 2.3.1).




Figure 2.2.3: Base of a slit cover of $\mathbb{P}_{\mathbb{C}}^1$ with slits s_1 from ξ_1 to ξ_2 , and s_2 from ξ_3 to ξ_4 ; and with generators g_1, g_2 , corresponding to the loops σ_1, σ_3 , respectively. (Here the inverses g_1^{-1}, g_2^{-1} correspond to the loops σ_2, σ_4 .)

As a result, we obtain the following recipe for obtaining slit covers by analytic patching: Given G and generators g_1, \dots, g_r (whose product need not be 1), let $A_j = \langle g_j \rangle$, and let m_j be the order of g_j . Take r disjoint open discs X_j , choose smaller open discs $X_j^* \subset X_j$, and for each j pick two points $\xi_{2j-1}, \xi_{2j} \in X_j^*$. Over $U_j = X_j - \{\xi_{2j-1}, \xi_{2j}\}$, let V_j be the A_j -Galois cover given by $y_j^{m_j} = (x - c_{2j-1})(x - c_{2j})^{m_j-1}$. The restriction of V_j to $O_j := X_j - \bar{X}_j^*$ is trivial, and we identify it with the A_j -Galois cover $\text{Ind}_1^{A_j} O_j$. This identifies the restriction of $\text{Ind}_{A_j}^G V_j$ over O_j with $\text{Ind}_1^G O_j$ — which is also the restriction

of the trivial cover $\text{Ind}_1^G U'$ of $U' = \mathbb{P}_\mathbb{C}^1 - \bigcup \bar{X}_j^*$ to O_j . Taking the union of the trivial G -Galois cover $\text{Ind}_1^G U'$ of U' with the induced covers $\text{Ind}_{A_j}^G V_j$, with respect to these identifications, we obtain the slit G -Galois étale cover of $U = \mathbb{P}_\mathbb{C}^1 - \{\xi_1, \dots, \xi_{2r}\}$ with description $(g_1, g_1^{-1}, g_2, g_2^{-1}, \dots, g_r, g_r^{-1})$. \square

The slit covers that occur in Example 2.3.2 can also be understood in terms of degeneration of covers — and this point of view will be useful later on, in more general settings. Consider the G -Galois slit cover $V \rightarrow U = \mathbb{P}_\mathbb{C}^1 - S$ with branch cycle description $(*)$ as in Example 2.3.2; here $S = \{\xi_1, \dots, \xi_{2r}\}$ and s_j is a simple path connecting ξ_{2j-1} to ξ_{2j} , with the various s_j 's having disjoint support. This cover may be completed to a G -Galois branched cover $Y \rightarrow \mathbb{P}_\mathbb{C}^1$, with branch locus S , by taking the normalization of $\mathbb{P}_\mathbb{C}^1$ in (the function field of) V . Now deform this branched cover by allowing each point ξ_{2j} to move along the path s_j backwards toward ξ_{2j-1} . This yields a one (real) parameter family of irreducible G -Galois slit covers $Y_t \rightarrow \mathbb{P}_\mathbb{C}^1$, each of which is trivial outside of a union of (shrinking) simply connected open sets containing ξ_{2j-1} and (the moving) ξ_{2j} . In the limit, when ξ_{2j} collides with ξ_{2j-1} , we obtain a finite map $Y_0 \rightarrow \mathbb{P}_\mathbb{C}^1$ which is unramified away from $S' := \{\xi_1, \xi_3, \dots, \xi_{2r-1}\}$, such that Y_0 is connected; G acts on Y_0 over $\mathbb{P}_\mathbb{C}^1$, and acts simply transitively away from S' ; and the map is a trivial cover away from S' . In fact, Y_0 is a union of $|G|$ copies of $\mathbb{P}_\mathbb{C}^1$, indexed by the elements of G , such that the g th copy meets the gg_j th copy over ξ_{2j-1} . The map $Y_0 \rightarrow \mathbb{P}_\mathbb{C}^1$ is a *mock cover* [Ha1, §3], i.e. is finite and generically unramified, and such that each irreducible component of Y_0 maps isomorphically onto the base (here, $\mathbb{P}_\mathbb{C}^1$). This degeneration procedure can be reversed: starting with a connected G -Galois mock cover which is built in an essentially combinatorial manner in terms of the data g_1, \dots, g_r , one can then deform it near each branch point to obtain an irreducible G -Galois branched cover branched at $2r$ points with branch cycle description $(g_1, g_1^{-1}, g_2, g_2^{-1}, \dots, g_r, g_r^{-1})$. This is one perspective on the key construction in the next section, on formal patching.




Figure 2.3.4: A mock cover of the line, with Galois group S_3 , branched at two points η_1, η_2 . The sheets are labeled by the elements of S_3 . The cyclic subgroups $\langle (01) \rangle, \langle (012) \rangle$ are the stabilizers on the identity sheet over η_1, η_2 , respectively.

As discussed before Example 2.3.2, slit covers do not require topological input — i.e. knowledge of the explicit structure of topological fundamental groups, or the existence of overlapping open discs containing different branch points — unlike the general three-point cover in Example 2.3.1. Without this topological input, for general covers one obtains only the equivalence of algebraic and analytic covers in Riemann’s Existence Theorem — and in particular, we do not obtain the corollaries to Riemann’s Existence Theorem in Section 2.1. But one can obtain those corollaries as they relate to slit covers, without topological input. Since only half of the entries of the branch cycle description can be specified for a slit cover, such results can be regarded as a “Half Riemann Existence Theorem”; and can be used to motivate analogous results about fundamental groups for curves that are not defined over \mathbb{C} , where there are no “loops” or overlapping open discs. (Indeed, the term “half Riemann Existence Theorem” was first coined by F. Pop to refer to such an analogous result [Po2, Main Theorem]; cf. §4.3 below). In particular, we have the following variant on Corollary 2.1.3:

Theorem 2.3.5. (Analytic half Riemann Existence Theorem) *Let $r \geq 1$, let $S = \{\xi_1, \dots, \xi_{2r}\}$ be a set of $2r$ distinct points in $\mathbb{P}_{\mathbb{C}}^1$, and let $U = \mathbb{P}_{\mathbb{C}}^1 - S$. Let F_r be the free profinite group on generators x_1, \dots, x_r . Then F_r is a quotient of the étale fundamental group of U .*

Namely, let G be any finite quotient of F_r . That is, G is a finite group together with generators g_1, \dots, g_r . Consider the G -Galois slit cover with branch cycle description $(g_1, g_1^{-1}, g_2, g_2^{-1}, \dots, g_r, g_r^{-1})$. These form an inverse subsystem of the full inverse system of covers of U , as G varies; and the inverse limit of their Galois groups is F_r .

Here, in order for this inverse system to make sense, one can first fix a bouquet of loops around the points of S ; or one can fix a set of disjoint simple paths s_j from ξ_{2j-1} to ξ_{2j} and consider the corresponding set of slit covers. But to give a non-topological proof of this result (which of course is a special case of Corollary 2.1.3), one can instead give compatible local Kummer equations for the slit covers and then use GAGA; or one can use the deformation construction starting from mock covers, as sketched above. These approaches are in fact equivalent, and will be discussed in the next section in a more general setting.

Observe that the above “half Riemann Existence Theorem” is sufficient to prove the inverse Galois problem over $\mathbb{C}(x)$, which appeared above, as Corollary 2.1.4 of (the full) Riemann’s Existence Theorem. Namely, for any finite group G , pick a set of r generators of G (for some r), and pick a set S of $2r$ points in $\mathbb{P}_{\mathbb{C}}^1$. Then G is the Galois group of an unramified Galois cover of $\mathbb{P}_{\mathbb{C}}^1 - S$; and taking function fields yields a G -Galois field extension of $\mathbb{C}(x)$.

The above discussion relating to Example 2.3.2 brings up the question of constructing

covers of algebraic curves defined over fields other than \mathbb{C} , and of proving at least part of Riemann’s Existence Theorem for curves over more general fields. Even if the topological input can be eliminated (as discussed above), it is still necessary to have a form of GAGA to pass from “analytic” objects to algebraic ones. The “analytic” objects will be defined over a topology that is finer than the Zariski topology, and with respect to which modules and covers can be constructed locally and patched. It will also be necessary to have a structure sheaf of “analytic” functions on the space under this topology.

One initially tempting approach to this might be to use the étale topology; but unfortunately, this does not really help. One difficulty with this is that a direct analog of GAGA does not hold in the étale topology. Namely, in order to descend a module from the étale topology to the Zariski topology, one needs to satisfy a descent criterion [Gr5, Chap. VIII, §1]. In the language of Theorem 2.2.6, this says that one needs not just agreement on the overlap $X_1 \times_X X_2$ between the given étale open sets, but also on the “self-overlaps” $X_1 \times_X X_1$ and $X_2 \times_X X_2$, which together satisfy a compatibility condition. (See also [Gr3], in which descent is viewed as a special case of patching, or “recollement”.) A second difficulty is that in order to give étale open sets $X_i \rightarrow X$, one needs to understand covers of X ; and so this introduces an issue of circularity into the strategy for studying and constructing covers.

Two other approaches have proven quite useful, though, for large classes of base fields (though not for all fields). These are the Zariski-Grothendieck notion of formal geometry, and Tate’s notion of rigid geometry. Those approaches will be discussed in the following sections.

Section 3: Formal patching

This section and the next describe approaches to carrying over the ideas of Section 2 to algebraic curves that are defined over fields other than \mathbb{C} . The present section uses formal schemes rather than complex curves, in order to obtain analogs of complex analytic notions that can be used to obtain results in Galois theory. The idea goes back to Zariski; and his notion of a “formal holomorphic function”, which uses formal power series rather than convergent power series, is presented in Section 3.1. Grothendieck’s strengthening of this notion is presented in Section 3.2, including his formal analog of Serre’s result GAGA (and the proof presented here parallels that of GAGA, presented in Section 2.2). These ideas are used in Section 3.3 to solve the geometric inverse Galois problem over various fields, using ideas motivated by the slit cover construction of Section 2.3. Further applications of these ideas are presented later, in Section 5.

Section 3.1. Zariski’s formal holomorphic functions.

In order to generalize analytic notions to varieties over fields other than \mathbb{C} , one needs to have “small open neighborhoods”, and not just Zariski open sets. One also needs to have a notion of (“analytic”) functions on those neighborhoods.

Unfortunately, if there is no metric on the ground field, then one cannot consider discs around the origin in \mathbb{A}_k^1 , for example, or the rings of power series that converge on those discs. But one can consider the ring of *all* formal power series, regarded as analytic (or holomorphic) functions on the spectrum of the complete local ring at the origin (which we regard as a “very small neighborhood” of that point). And in general, given a variety V and a point $\nu \in V$, we can consider the elements of the complete local ring $\hat{\mathcal{O}}_{V,\nu}$ as holomorphic functions on $\text{Spec } \hat{\mathcal{O}}_{V,\nu}$.

While this point of view can be used to study local behaviors of varieties near a point, it does not suffice in order to study more global behaviors locally and then to “patch” (as one would want to do in analogs of GAGA and Riemann’s Existence Theorem), because these “neighborhoods” each contain only one closed point. The issue is that a notion of “analytic continuation” of holomorphic functions is necessary for that, so that holomorphic functions near one point can also be regarded as holomorphic functions near neighboring points.

This issue was Zariski’s main focus during the period of 1945-1950, and it grew out of ideas that arose from his previous work on resolution of singularities. The question was how to extend a holomorphic function from the complete local ring at a point $\nu \in V$ to points in a neighborhood. As he said later in a preface to his collected works [Za5, pp. xii-xiii], “I sensed the probable existence of such an extension provided the analytic continuation were carried out along an algebraic subvariety W of V .” That is, if W is a Zariski closed subset of V , then it should make sense to speak of “holomorphic functions”

in a “formal neighborhood” of W in V .

These formal holomorphic functions were defined as follows ([Za4, Part I]; see also [Ar5, p.3]): Let W be a Zariski closed subset of a variety V . First, suppose that V is affine, say with ring of functions R , so that $W \subset V$ is defined by an ideal $I \subset R$. Consider the ring of rational functions g on V that are regular along W ; this is a metric space with respect to the I -adic metric. The space of *strongly holomorphic functions f along W* (in V) is defined to be the metric completion of this space (viz. it is the space of equivalence classes of Cauchy sequences of such functions g). This space is also a ring, and can be identified with the inverse limit $\varprojlim R/I^n$.

More generally, whether or not V is affine, one can define a (*formal*) *holomorphic function along W* to be a function given locally in this manner. That is, it is defined to be an element $\{f_\omega\} \in \prod_{\omega \in W} \hat{\mathcal{O}}_{V,\omega}$ such that there is a Zariski affine open covering $\{V_i\}_{i \in I}$ of V together with a choice of a strongly holomorphic function $\{f_i\}_{i \in I}$ along $W_i := W \cap V_i$ in V_i (for each $i \in I$), such that f_ω is the image of f_i in $\hat{\mathcal{O}}_{V,\omega}$ whenever $\omega \in W_i$. These functions also form a ring, denoted $\hat{\mathcal{O}}_{V,W}$. Note that $\hat{\mathcal{O}}_{V,W}$ is the complete local ring $\hat{\mathcal{O}}_{V,\omega}$ if $W = \{\omega\}$. Also, if U is an affine open subset of W , and $U = \tilde{U} \cap W$ for some open subset $\tilde{U} \subset V$, then the ring $\hat{\mathcal{O}}_{\tilde{U},U}$ depends only on U , and not on the choice of \tilde{U} ; so we also denote this ring by $\hat{\mathcal{O}}_{V,U}$, and call it the ring of holomorphic functions along U in V .

Remark 3.1.1. Nowadays, if I is an ideal in a ring R , then the *I -adic completion* of R is defined to be the inverse limit $\varprojlim R/I^n$. This modern notion of formal completion is equivalent to Zariski’s above notion of metric completion via Cauchy sequences, which he first gave in [Za2, §5]. But Zariski’s approach more closely paralleled completions in analysis, and fit in with his view of formal holomorphic functions as being analogs of complex analytic functions. (Prior to his giving this definition, completions of rings were defined only with respect to *maximal* ideals.) In connection with his introduction of this definition, Zariski also introduced the class of rings we now know as *Zariski rings* (and which Zariski had called “semi-local rings”): viz. rings R together with a non-zero ideal I such that every element of $1 + I$ is a unit in R [Za2, Def. 1]. Equivalently [Za2, Theorem 5], these are the I -adic rings such that I is contained in (what we now call) the Jacobson radical of R . Moreover, every I -adically complete ring is a Zariski ring [Za2, Cor. to Thm. 4]; so the ring of strictly holomorphic functions on a closed subset of an affine variety is a Zariski ring. \square

A deep fact proven by Zariski [Za4, §9, Thm. 10] is that every holomorphic function along a closed subvariety of an affine scheme is strongly holomorphic. So those two rings of functions agree, in the affine case; and the ring of holomorphic functions along $W = \text{Spec } R/I$ in $V = \text{Spec } R$ can be identified with the formal completion $\varprojlim R/I^n$ of R with respect to I .

Example 3.1.2. Consider the x -axis $W \approx \mathbb{A}_k^1$ in the x, t -plane $V = \mathbb{A}^2$. Then W is defined by the ideal $I = (t)$, and the ring of holomorphic functions along W in V is $A_1 := k[x][[t]]$. Note that every element of A_1 can be regarded as an element in $\hat{\mathcal{O}}_{W, \nu}$ for every point $\nu \in W$; and in this way can be regarded as an analytic continuation of (local) functions along the x -axis. Intuitively, the spectrum S_1 of A_1 can be viewed as a thin tubular neighborhood of W in V , which “pinches down” as $x \rightarrow \infty$. For example, observe that the elements x and $x - t$ are non-units in A_1 , and so each defines a proper ideal of A_1 ; and correspondingly, their loci in $S_1 = \text{Spec } A_1$ are non-empty (and meet the x -axis at the origin). On the other hand, $1 - xt$ is a unit in A_1 , with inverse $1 + xt + x^2t^2 \dots$, so its locus in S_1 is empty; and geometrically, its locus in V (which is a hyperbola) approaches the x -axis only as $x \rightarrow \infty$, and so misses the (“pinched down”) spectrum of A_1 . One can similarly consider the ring $A_2 = k[x^{-1}][[t]]$; its spectrum S_2 is a thin neighborhood of $\mathbb{P}_k^1 - (x = 0)$ which “pinches down” near $x = t = 0$. \square

Example 3.1.3. Let V' be the complement of the t -axis ($x = 0$) in the x, t -plane \mathbb{A}^2 , and let $W' \subset V'$ be the locus of $t = 0$. Then the ring of holomorphic functions along W' in V' is $A_0 := k[x, x^{-1}][[t]]$. Geometrically, this is a thin tubular neighborhood of W' in V' , which “pinches down” in two places, viz. as x approaches either 0 or ∞ . Observe that $\text{Spec } A_0$ is *not* a Zariski open subset of $\text{Spec } A_1$, where A_1 is as in Example 3.1.2. In particular, A_0 is much larger than the ring $A_1[x^{-1}]$; e.g. $\sum_{n=1}^{\infty} x^{-n}t^n$ is an element of A_0 but not of $A_1[x^{-1}]$. Intuitively, $S_0 := \text{Spec } A_0$ can be viewed as an “analytic open subset” of $S_1 = \text{Spec } A_1$ but not a Zariski open subset — and similarly for S_0 and $S_2 = \text{Spec } A_2$ in Example 3.1.2. Moreover S_0 can be regarded as the “overlap” of S_1 and S_2 in $\mathbb{P}_k^1[[t]]$. This will be made more precise below. \square




Figure 3.1.4: A covering of $\mathbb{P}_k^1[[t]]$ (lower left) by two formal patches, $S_1 = \text{Spec } k[x][[t]]$ and $S_2 = \text{Spec } k[1/x][[t]]$. The “overlap” S_0 is $\text{Spec } k[x, 1/x][[t]]$. See Examples 3.1.2 and 3.1.3.

Remark 3.1.5. Just as the ring $A_0 = k[x, x^{-1}][[t]]$ in Example 3.1.3 is much larger than

$A_1[x^{-1}]$, where $A_1 = k[x][[t]]$ as in Example 3.1.2, it is similarly the case that the ring A_1 is much larger than the ring $T := k[[t]][x]$ (e.g. $\sum_{n=1}^{\infty} x^n t^n$ is in A_1 but not in T). The scheme $\text{Spec } T$ can be identified with the affine line over the complete local ring $k[[t]]$, and is a Zariski open subset of $\mathbb{P}_{k[[t]]}^1$ (given by $x \neq \infty$). This projective line over $k[[t]]$ can be viewed as a thin but uniformly wide tubular neighborhood of the projective x -line \mathbb{P}_k^1 , and its affine open subset $\text{Spec } T$ can correspondingly be viewed as a uniformly wide thin tubular neighborhood of the x -axis \mathbb{A}_k^1 (with no “pinching down” near infinity). As in Example 2, we have here that $\text{Spec } A_1$ is not a Zariski open subset of $\text{Spec } T$, and instead it can be viewed as an “analytic open subset” of $\text{Spec } T$. \square

Using these ideas, Zariski proved his Fundamental Theorem on formal holomorphic functions [Za4, §11, p.50]: If $f : V' \rightarrow V$ is a projective morphism of varieties, with V normal and with the function field of V algebraically closed in that of V' , and if $W' = f^{-1}(W)$ for some closed subset $W \subset V$, then the natural map $\hat{\mathcal{O}}_{V,W} \rightarrow \hat{\mathcal{O}}_{V',W'}$ is an isomorphism. (See [Ar5, pp.5-6] for a sketch of the proof.) This result in turn yielded Zariski’s Connectedness Theorem [Za4, §20, Thm. 14] (cf. also [Hrt2, III, Cor. 11.3]), and implied Zariski’s Main Theorem (cf. [Hrt2, III, Cor. 11.4]).

The above general discussion suggests that it should be possible to prove an analog of GAGA that would permit patching of modules using formal completions. And indeed, there is the following assertion, which is essentially a result of Ferrand and Raynaud (cf. [FR, Prop. 4.2]). Here the notation is as at the end of Section 2.2 above, and this result can be viewed as analogous to the version of GAGA given by Theorem 2.2.6.

Proposition 3.1.6. (Ferrand-Raynaud) *Let R be a Noetherian ring, let V be the affine scheme $\text{Spec } R$, let W be a closed subset of V , and let $V^\circ = V - W$. Let R^* be the ring of holomorphic functions along W in V , and let $W^* = \text{Spec } R^*$. Also let $W^\circ = W^* \times_V V^\circ$. Then the base change functor*

$$\mathfrak{M}(V) \rightarrow \mathfrak{M}(W^*) \times_{\mathfrak{M}(W^\circ)} \mathfrak{M}(V^\circ)$$

is an equivalence of categories.

Here R^* is the I -adic completion of R , where I is the ideal of W in V . Intuitively, we regard $W^* = \text{Spec } R^*$ as a “formal neighborhood” of W in V , and we regard W° as the “intersection” of W^* with V° (i.e. the “complement” of W in W^*).

Remark 3.1.7. The above result is essentially a special case of the assertion in [FR, Prop. 4.2]. That result was stated in terms of cartesian diagram of categories, which is equivalent to an assertion concerning 2-fibre products (i.e. the way Proposition 3.1.6 above is stated). The main difference between the above result and [FR, Prop. 4.2] is that the latter result allows W^* more generally to be any scheme for which there is a flat morphism

$f : W^* \rightarrow V$ such that the pullback $f_W : W^* \times_V W \rightarrow W$ is an isomorphism — which is the case in the situation of Proposition 3.1.6 above. Actually, though, [FR, Prop. 4.2] assumes that $f : W^* \rightarrow V$ is *faithfully* flat (unlike the situation in Proposition 1). But this extra faithfulness hypothesis is unnecessary for their proof; and in any event, given a flat morphism $f : W^* \rightarrow V$ such that f_W is an isomorphism, one can replace W^* by the disjoint union of W^* and V° , which is then faithfully flat — and applying [FR, Prop. 4.2] to that new W^* gives the desired conclusion for the original W^* . \square

The following result of Artin [Ar4, Theorem 2.6] generalizes Proposition 3.1.6:

Proposition 3.1.8. *In the situation of Proposition 3.1.6, let \tilde{V} be a scheme and let $f : \tilde{V} \rightarrow V$ be a morphism of finite type. Let \tilde{W}^* , \tilde{V}° , \tilde{W}° be the pullbacks of W^* , V° , W° with respect to f . Then the base change functor*

$$\mathfrak{M}(\tilde{V}) \rightarrow \mathfrak{M}(\tilde{W}^*) \times_{\mathfrak{M}(\tilde{W}^\circ)} \mathfrak{M}(\tilde{V}^\circ)$$

is an equivalence of categories.

Note that $\tilde{V}^\circ = \tilde{V} - \tilde{W}$ in Proposition 3.1.8, where $\tilde{W} = f^{-1}(W)$.

As an example of this result, let V be a smooth n -dimensional affine scheme over a field k , let W be a closed point of V , and let \tilde{V} be the blow-up of V at W . So \tilde{W} is a copy of \mathbb{P}_k^{n-1} , and \tilde{W}^* is the spectrum of a “uniformly wide tubular neighborhood” of \tilde{W} in \tilde{V} . Here \tilde{W}^* , which is irreducible, can be viewed as a “twisted version” of \mathbb{P}_k^{n-1} ; cf. [Hrt2, p.29, Figure 3] for the case $n = 2$. According to Proposition 3.1.8, giving a coherent module on \tilde{V} is equivalent to giving such modules on \tilde{W}^* and on the complement of \tilde{W} , with agreement on the “overlap” \tilde{W}° .

While the above two propositions required V to be affine, this hypothesis can be dropped if W is finite:

Corollary 3.1.9. *Let V be a Noetherian scheme, and let W be a finite set of closed points in V . Let R^* be the ring of holomorphic functions along W in V , let $W^* = \text{Spec } R^*$, let $V^\circ = V - W$, and let $W^\circ = W^* \times_V V^\circ$.*

a) Then the base change functor

$$\mathfrak{M}(V) \rightarrow \mathfrak{M}(W^*) \times_{\mathfrak{M}(W^\circ)} \mathfrak{M}(V^\circ)$$

is an equivalence of categories.

b) Let \tilde{V} be a scheme and let $f : \tilde{V} \rightarrow V$ be a morphism of finite type. Let \tilde{W}^ , \tilde{V}° , \tilde{W}° be the pullbacks of W^* , V° , W° with respect to f . Then the base change functor*

$$\mathfrak{M}(\tilde{V}) \rightarrow \mathfrak{M}(\tilde{W}^*) \times_{\mathfrak{M}(\tilde{W}^\circ)} \mathfrak{M}(\tilde{V}^\circ)$$

is an equivalence of categories.

Namely, we may cover V by finitely many affine open subsets $V_i = \text{Spec } R_i$, with R_i Noetherian. Applying Proposition 3.1.6 to each V_i and $W_i := V_i \cap W$, we obtain equivalences over each V_i . These equivalences agree on the overlaps $V_i \cap V_j$ (since each is given by base change), and so together they yield the desired equivalence over V , in part (a). Part (b) is similar, using Proposition 3.1.8.

Unfortunately, while the above results are a kind of GAGA, permitting the patching of modules, they do not directly help to construct covers (via the General Principle 2.2.4); and so they do not directly help prove an analog of Riemann’s Existence Theorem. The reason is that these results require that a module be given over a Zariski open subset V° (or \tilde{V}°), viz. the complement of the given closed subset W (or \tilde{W}). And a normal cover $Z \rightarrow V$ is determined by its restriction to a dense open subset V° (viz. it is the normalization of V in the function field of the cover — which is the same as the function field of the restriction). So these results provide a cover $Z \rightarrow V$ only in circumstances in which one already has the cover in hand.

Instead, in order to use Zariski’s approach to obtain results about covers, we will focus on spaces such as $\mathbb{P}_{k[[t]]}^1$ (and see the discussion in the Remark 3.1.5 above). In that situation, Grothendieck has proven a “formal GAGA”, which we discuss next. That result yields a version of Riemann’s Existence Theorem for many fields other than \mathbb{C} . Combining that approach with the above results of Ferrand-Raynaud and Artin yields even stronger versions of “formal GAGA”; and those formal patching results have been used to prove a number of results concerning covers and fundamental groups over various fields (as will be discussed later).

Section 3.2. Grothendieck’s formal schemes.

Drawing on Zariski’s notion of formal holomorphic functions, Grothendieck introduced the notion of *formal scheme*, and provided a framework for proving a “formal GAGA” that is sufficient for establishing formal analogs of (at least parts of) Riemann’s Existence Theorem. In his paper of the same name [Gr2], Grothendieck announced his result GFGA (“géométrie formelle et géométrie algébrique”), and sketched how it leads to results about covers and fundamental groups of curves. The details of this GFGA result appeared later in EGA [Gr4, III, Cor. 5.1.6], and the result in that form has become known as Grothendieck’s Existence Theorem. In SGA 1 [Gr5], the details about the results on covers and fundamental groups appeared.

To begin with, fix a Zariski closed subset W of a scheme V . Let $\mathcal{O}_{\mathfrak{W}} = \mathcal{O}_{\mathfrak{W},W}$ be the sheaf of holomorphic functions along W in V . That is, for every Zariski open subset $U \subset W$, let $\mathcal{O}_{\mathfrak{W}}(U)$ be the ring $\hat{\mathcal{O}}_{V,U}$ of holomorphic functions along U in V . Thus $\mathcal{O}_{\mathfrak{W}} = \varinjlim_n \mathcal{O}_V/\mathcal{I}^{n+1}$, where \mathcal{I} is the sheaf of ideals of \mathcal{O}_V defining W in V . The ringed

space $\mathfrak{W} := (W, \mathcal{O}_{\mathfrak{W}})$ is defined to be the *formal completion of V along W* .

The simplest example of this takes V to be the affine t -line over a field k , and W to be the point $t = 0$. Here we may identify $\mathcal{O}_{\mathfrak{W}}$ with the ring $k[[x]] = \varprojlim_n \mathcal{O}_n$, where $\mathcal{O}_n = k[t]/(t^{n+1})$. Here $n = 0$ corresponds to W , and $n > 0$ to infinitesimal thickenings of W . The kernel I_m of $\mathcal{O}_m \rightarrow \mathcal{O}_0$ is the ideal $t\mathcal{O}_m$, and the kernel of $\mathcal{O}_m \rightarrow \mathcal{O}_n$ is $t^{n+1}\mathcal{O}_m = I_m^{n+1}$.

As a somewhat more general example, let A be a ring that is complete with respect to an ideal I . Then $W = \text{Spec } A/I$ is a closed subscheme of $V = \text{Spec } A$, consisting of the prime ideals of A that are open in the I -adic topology. The formal completion $\mathfrak{W} = (W, \mathcal{O}_{\mathfrak{W}})$ of V along W consists of the underlying topological space W together with a structure sheaf whose ring of global sections is A . This formal completion is also called the *formal spectrum* of A , denoted $\text{Spf } A$. (For example, if $A = k[x][[t]]$ and $I = (t)$, then the underlying space of $\text{Spf } A$ is the affine x -line over k , and its global sections are $k[x][[t]]$.)

Note that the above definition of formal completion relies on the idea that the geometry of a space is captured by the structure sheaf on it, rather than on the underlying topological space. Indeed, the underlying topological space of \mathfrak{W} is the same as that of W ; but the structure sheaf $\mathcal{O}_{\mathfrak{W}}$ incorporates all of the information in the spectra of $\hat{\mathcal{O}}_{V,U}$ — and thus it reflects the local geometry of V near W .

More generally, suppose we are given a topological space X and a sheaf of topological rings $\mathcal{O}_{\mathfrak{X}}$ on X . Suppose also that $\mathcal{O}_{\mathfrak{X}} = \varprojlim_n \mathcal{O}_n$, where $\{\mathcal{O}_n\}_n$ is an inverse system of sheaves of rings on X such that (X, \mathcal{O}_n) is a scheme X_n for each n ; and that for $m \geq n$, the homomorphism $\mathcal{O}_m \rightarrow \mathcal{O}_n$ is surjective with kernel \mathcal{I}_m^{n+1} , where $\mathcal{I}_m = \ker(\mathcal{O}_m \rightarrow \mathcal{O}_0)$. Then the ringed space $\mathfrak{X} := (X, \mathcal{O}_{\mathfrak{X}})$ is a *formal scheme*. In particular, in the situation above for $W \subset V$ (taking $\mathcal{O}_n = \mathcal{O}_V/\mathcal{I}^{n+1}$), the formal completion $\mathfrak{W} = (W, \mathcal{O}_{\mathfrak{W}})$ of V along W is a formal scheme.

If W is a closed subset of a scheme V , with formal completion \mathfrak{W} , then to every sheaf \mathcal{F} of \mathcal{O}_V -modules on V we may canonically associate a sheaf $\hat{\mathcal{F}}$ of $\mathcal{O}_{\mathfrak{W}}$ modules on \mathfrak{W} . Namely, for every n let $\mathcal{F}_n = \mathcal{F} \otimes_{\mathcal{O}_V} \mathcal{O}_V/\mathcal{I}^{n+1}$, where \mathcal{I} is the sheaf of ideals defining W . Then let $\hat{\mathcal{F}} = \varprojlim_n \mathcal{F}_n$. Note that $\mathcal{O}_{\mathfrak{W}} = \hat{\mathcal{O}}_V$. Also observe that if \mathcal{F} is a coherent \mathcal{O}_V -module, then $\hat{\mathcal{F}}$ is a coherent $\mathcal{O}_{\mathfrak{W}}$ -module (i.e. it is locally of the form $\mathcal{O}_{\mathfrak{W}}^m \rightarrow \mathcal{O}_{\mathfrak{W}}^n \rightarrow \hat{\mathcal{F}} \rightarrow 0$).

Theorem 3.2.1. (GFGA, Grothendieck Existence Theorem) *Let A be a Noetherian ring that is complete with respect to a proper ideal I , let V be a proper A -scheme, and let $W \subset V$ be the inverse image of the locus of I . Let $\mathfrak{W} = (W, \mathcal{O}_{\mathfrak{W}})$ be the formal completion of V along W . Then the functor $\mathcal{F} \mapsto \hat{\mathcal{F}}$, from the category of coherent \mathcal{O}_V -modules to the category of coherent $\mathcal{O}_{\mathfrak{W}}$ -modules, is an equivalence of categories.*

Before turning to the proof of Theorem 3.2.1, we discuss its content and give some examples, beginning with

Corollary 3.2.2. [Gr2, Cor. 1 to Thm. 3] *In the situation of Theorem 3.2.1, the natural map from closed subschemes of V to closed formal subschemes of \mathfrak{V} is a bijection.*

Namely, such subschemes [resp. formal subschemes] correspond bijectively to coherent subsheaves of \mathcal{O}_V [resp. of $\mathcal{O}_{\mathfrak{V}}$]. So this is an immediate consequence of the theorem.

This corollary may seem odd, for example in the case where V is a curve over a complete local ring A , and W is thus a curve over the residue field of A — since then, the only reduced closed subsets of W (other than W itself) are finite sets of points. But while distinct closed subschemes of V can have the same intersection with the topological space W , the structure sheaves of their restrictions will be different, and so the induced formal schemes will be different.

Theorem 3.2.1 can be viewed in two ways: as a thickening result (emphasizing the inverse limit point of view), and as a patching result (emphasizing the analogy with the classical GAGA of Section 2.2).

From the point of view of thickening, given $W \subset V$ defined by a sheaf of ideals \mathcal{I} , we have a sequence of subschemes $V_n = \underline{\text{Spec}} \mathcal{O}_V/\mathcal{I}^{n+1}$. Each V_n has the same underlying topological space (viz. that of $W = V_0$), but has a different structure sheaf. The formal completion \mathfrak{V} of V along W can be regarded as the direct limit of the schemes V_n . What Theorem 3.2.1 says is that under the hypotheses of that result, to give a coherent sheaf \mathcal{F} on V is equivalent to giving a compatible set of coherent sheaves \mathcal{F}_n on the V_n 's (i.e. the restrictions of \mathcal{F} to the V_n 's). The hard part (cf. the proof below) is to show the existence of a coherent sheaf \mathcal{F} that restricts to a given compatible set of coherent sheaves \mathcal{F}_n . And later, the result will tell us that to give a branched cover of V is equivalent to giving a compatible system of covers of the V_n 's.

On the other hand, the point of view of patching is closer to that of Zariski's work on formal holomorphic functions. Given $W \subset V$, we can cover W by affine open subsets U_i . By definition, giving a coherent formal sheaf on W amounts to giving finitely presented modules over the rings $\hat{\mathcal{O}}_{V,U_i}$ that are compatible on the overlaps (i.e. over the rings $\hat{\mathcal{O}}_{V,U_{ij}}$, where $U_{ij} = U_i \cap U_j$). So Theorem 3.2.1 says that to give a coherent sheaf \mathcal{F} on V is equivalent to giving such modules locally (i.e. the pullbacks of \mathcal{F} to the “formal neighborhoods” $\text{Spec } \hat{\mathcal{O}}_{V,U_i}$ with agreements on the “formal overlaps” $\text{Spec } \hat{\mathcal{O}}_{V,U_{ij}}$). The same principle will be applied later to covers.

Example 3.2.3. Let k be a field, let $A = k[[t]]$, and let $V = \mathbb{P}_A^1$, the projective x -line over A . So W is the projective x -line over k . Let \mathfrak{V} be the formal completion of V at W . Theorem 3.2.1 says that giving a coherent \mathcal{O}_V -module is equivalent to giving a coherent $\mathcal{O}_{\mathfrak{V}}$ -module.

From the perspective of thickening, to give a coherent $\mathcal{O}_{\mathfrak{V}}$ -module \mathcal{F} amounts to giving an inverse system of coherent modules \mathcal{F}_n over the V_n 's, where V_n is the projective

x -line over $k[t]/(t^{n+1})$. Each finite-level thickening \mathcal{F}_n gives more and more information about the given module, and in the limit, the theorem says that the full \mathcal{O}_V -module \mathcal{F} is determined.

For the patching perspective, cover W by two open sets U_1 (where $x \neq \infty$) and U_2 (where $x \neq 0$), each isomorphic to the affine k -line. The corresponding rings of holomorphic functions are $k[x][[t]]$ and $k[x^{-1}][[t]]$, while the ring of holomorphic functions along the overlap $U_0 : (x \neq 0, \infty)$ is $k[x, x^{-1}][[t]]$. As in Examples 3.1.2 and 3.1.3, the spectra S_1, S_2 of the first two of these rings can be viewed as tubular neighborhoods of the two affine lines, pinching down near $x = \infty$ and near $x = 0$ respectively. The spectrum S_0 of the third ring (the “formal overlap”) can be viewed as a tubular neighborhood that pinches down near both 0 and ∞ . (See Figure 3.1.4.) These spectra can be viewed as “analytic open subsets” of V , which cover V (in the sense that the disjoint union $S_1 \cup S_2$ is faithfully flat over V) — and the theorem says that giving coherent modules over S_1 and S_2 , which agree over S_0 , is equivalent to giving a coherent module over V . \square

From the above patching perspective, Theorem 3.2.1 can be rephrased as follows, in a form that is useful in the case of relative dimension 1. In order to be able to apply it to Galois theory (in Section 3.3 below), we state it as well for algebras and covers.

Theorem 3.2.4. *In the situation of Theorem 3.2.1, suppose that U_1, U_2 are affine open subsets of W such that $U_1 \cup U_2 = W$, with intersection U_0 . For $i = 0, 1, 2$, let S_i be the spectrum of the ring of holomorphic functions along U_i in V . Then the base change functor*

$$\mathfrak{M}(V) \rightarrow \mathfrak{M}(S_1) \times_{\mathfrak{M}(S_0)} \mathfrak{M}(S_2)$$

is an equivalence of categories. Moreover the same holds if \mathfrak{M} is replaced by the category of finite algebras, or of finite branched covers, or of Galois covers with a given Galois group.

Compare this with the restatement of the classical GAGA at Theorem 2.2.6, and with the results of Ferrand-Raynaud and Artin (Propositions 3.1.6 and 3.1.8). See also Figure 3.1.4 for an illustration of this result in the situation of the above example. As in Theorem 2.2.6, the above assertions for algebras and covers follow formally from the result for modules, via the General Principle 2.2.4. (Cf. also [Ha2, Proposition 2.8].)

Remarks 3.2.5. (a) Theorem 3.2.1 does not hold if the properness hypothesis is dropped. For example, Corollary 3.2.2 is false in the case that $A = k[[t]]$ and $V = \mathbb{A}_A^1$ (since the subscheme $(1 - xt)$ in V induces the same formal subscheme of \mathcal{V} as the empty set). Similarly, Theorem 3.2.4 does not hold as stated if V is not proper over A (and note that $S_1 \cup S_2$ is not faithfully flat over V in this situation). But a variant of Theorem 3.2.4 does hold if V is affine: namely there is still an equivalence if $\mathfrak{M}(V)$ is replaced by $\mathfrak{M}(S)$, where S is the ring of holomorphic functions along W in V . This is essentially a restatement of Zariski’s result that holomorphic functions on an affine open subset of W are strongly

holomorphic. It is also analogous to the version of Cartan’s Theorem A for Stein spaces [Ca2] (cf. the discussion near the end of Section 2.2 above).

(b) The main content of Theorem 3.2.1 (or Theorem 3.2.4) can also be phrased in affine terms in the case of relative dimension 1. For instance, in the situation of the above example with $A = k[[t]]$ and $V = \mathbb{P}_A^1$, a coherent module \mathcal{M} over V is determined up to twisting by its restriction to $\mathbb{A}_A^1 = \text{Spec } A[x]$. Letting S_0, S_1, S_2 be as in the example, and restricting to the Zariski open subset \mathbb{A}_A^1 , we obtain an equivalence of categories

$$\mathfrak{M}(R) \rightarrow \mathfrak{M}(R_1) \times_{\mathfrak{M}(R_0)} \mathfrak{M}(R_2) \quad (*)$$

where $R = k[[t]][x]$; $R_1 = k[x][[t]]$; $R_2 = k[x^{-1}][[t]][x]$; and $R_0 = k[x, x^{-1}][[t]]$. (Here we adjoin x in the definition of R_2 because of the restriction to \mathbb{A}_A^1 .) In this situation, one can directly prove a formal version of Cartan’s Lemma, viz. that every element of $\text{GL}_n(R_0)$ can be written as the product of an element of $\text{GL}_n(R_1)$ and an element of $\text{GL}_n(R_2)$. This immediately gives the analog of (*) for the corresponding categories of finitely generated *free* modules, by applying this formal Cartan’s Lemma to the transition matrix between the bases over R_1 and R_2 . (Cf. the discussion in Section 2.2 above, and also [Ha2, Prop. 2.1] for a general result of this form.) Moreover, combining this formal Cartan’s Lemma with the fact that every element of R_0 is the sum of an element of R_1 and an element of R_2 , one can deduce all of (*), and thus essentially all of Theorem 3.2.1 in this situation. (See [Ha2, Proposition 2.6] for the general result, and see also Remark 1 after the proof of Corollary 2.7 there.)

(c) Using the approach of Remark (b), one can also prove analogous results where Theorem 3.2.1 does not apply. For example, let A and B be subrings of \mathbb{Q} , let $D = A \cap B$, and let C be the subring of \mathbb{Q} generated by A and B . (For instance, take $A = \mathbb{Z}[1/2]$ and $B = \mathbb{Z}[1/3]$, so $C = \mathbb{Z}[1/6]$ and $D = \mathbb{Z}$.) Then “Cartan’s Lemma” applies to the four rings $A[[t]]$, $B[[t]]$, $C[[t]]$, $D[[t]]$ (as can be proven by constructing the coefficients of the entries of the factorization, inductively). So by [Ha2, Proposition 2.6]), giving a finitely generated module over $D[[t]]$ is equivalent to giving such modules over $A[[t]]$ and $B[[t]]$ together with an isomorphism between the modules they induce over $C[[t]]$.

Another example involves the ring of convergent arithmetic power series $\mathbb{Z}\{t\}$, which consists of the formal power series $f(t) \in \mathbb{Z}[[t]]$ such that f converges on the complex disc $|t| < 1$. (Under the analogy between \mathbb{Z} and $k[x]$, the ring $\mathbb{Z}[[t]]$ is analogous to $k[x][[t]]$, and the ring $\mathbb{Z}\{t\}$ is analogous to $k[[t]][x]$.) Then with A, B, C, D as in the previous paragraph, “Cartan’s Lemma” applies to $A[[t]]$, $B\{t\}$, $C[[t]]$, $D\{t\}$ [Ha2, Prop. 2.3]. As a consequence, the analog of Theorem 3.2.4 holds for these rings: viz. giving a finitely presented module over $D\{t\}$ is equivalent to giving such modules over $A[[t]]$ and $B\{t\}$ together with an isomorphism between the modules they induce over $C[[t]]$ [Ha5, Theorem 3.6]. \square

The formal GAGA (Theorem 3.2.1) above can be proved in a way that is analogous to the proof of the classical GAGA (as presented in Section 2.2). In particular, there are two main ingredients in the proof. The first is:

Theorem 3.2.6. (Grothendieck) *In the situation of Theorem 3.2.1, if \mathcal{F} is a coherent sheaf on V , then the natural map $\varepsilon : H^q(V, \mathcal{F}) \rightarrow H^q(\mathfrak{B}, \hat{\mathcal{F}})$ is an isomorphism for every $q \geq 0$.*

This result was announced in [Gr2, Cor. 1 to Thm. 2] and proven in [Gr4, III, Prop. 5.1.2]. Here the formal H^q 's can (equivalently) be defined either via Čech cohomology or by derived functor cohomology. The above theorem is analogous to Theorem 2.2.2, concerning the classical case; and like that result, it is proven by descending induction on q (using that $H^q = 0$ for q sufficiently large). As in Section 2.2, it is the key case $q = 0$ that is used in proving GAGA. That case is known as Zariski's Theorem on Formal Functions [Hrt2, III, Thm. 11.1]; it generalizes the original version of Zariski's Fundamental Theorem on formal holomorphic functions [Za4, §11, p.50], which is the case $q = 0$ and $\mathcal{F} = \mathcal{O}_V$, and which was discussed in Section 3.1 above.

The second key ingredient in the proof of Theorem 3.2.1 is analogous to Theorem 2.2.3:

Theorem 3.2.7. *In the situation of Theorem 3.2.1 (with V assumed projective over A), let \mathcal{M} be a coherent \mathcal{O}_V -module or a coherent $\mathcal{O}_{\mathfrak{B}}$ -module. Then for $n \gg 0$ the twisted sheaf $\mathcal{M}(n)$ is generated by finitely many global sections.*

In the algebraic case, this is again Serre's result [Hrt2, Chap. II, Theorem 5.17]. In the formal case, this is a formal analog of Cartan's Theorem A, proven by Grothendieck [Gr4, III, Cor. 5.2.4], via a formal analog of Cartan's Theorem B [Gr4, III, Prop. 5.2.3].

Once one has Theorems 3.2.6 and 3.2.7 above, the projective case of Theorem 3.2.1 follows from them in exactly the same manner that Theorem 2.2.1 (classical GAGA) followed from Theorems 2.2.2 and 2.2.3 there. The proper case can then be deduced from the projective case using Chow's Lemma [Gr4, II, Thm. 5.6.1]; cf. [Gr4, III, 5.3.5] for details.

As for the proof of Theorem 3.2.7, the issue here, as in the analytic analog in Section 2, is to obtain a twist that will work for a given sheaf, even though the sheaf is not algebraic and we cannot simply clear denominators (as in the algebraic proof). In the situation here, a formal sheaf \mathcal{M} corresponds to an inverse system $\{\mathcal{M}_i\}$ of sheaves on the finite thickenings V_i . By the result in the algebraic case (applied to V_i), we have that for each i there is an n such that $\mathcal{M}_i(n)$ is generated by finitely many global sections. But we need to know that there is a single n that works for *all* i , and with compatible finite sets of global sections. The strategy is to pick a finite set of generating sections for $\mathcal{M}_0(n)$ for some n (and these will exist if n is chosen sufficiently large); and then inductively to lift them to sections of the $\mathcal{M}_i(n)$'s, in turn. If this is done, Theorem 3.2.7 follows, since the lifted sections automatically generate, by Nakayama's Lemma.

In order to carry out this inductive lifting, first reduce to the case $V = \mathbb{P}_A^m$ for some m , as in Section 2.2 (viz. embedding the given V in some \mathbb{P}_A^m and extending the module by 0). Now let $\text{gr } A$ be the associated graded ring to A and let $\text{gr } \mathcal{O} = (R/I)\mathcal{O} \oplus (I/I^2)\mathcal{O} \oplus \cdots$ (where $\mathcal{O} = \mathcal{O}_V$). Also write $\text{gr } \mathcal{M} = \mathcal{M}_0 \oplus (I/I^2)\mathcal{M}_1 \oplus \cdots$. Since \mathcal{M} is a coherent $\mathcal{O}_{\mathfrak{X}}$ -module, it follows that $\text{gr } \mathcal{M}$ is a coherent $\text{gr } \mathcal{O}$ -module on $\mathbb{P}_{\text{gr } \mathcal{O}}^m$. So by the algebraic analog of Cartan's Theorem B (i.e. by Serre's result [Hrt2, III, Theorem 5.2]), there is an integer n_0 such that for all $n \geq n_0$, $H^1(\mathbb{P}_{\text{gr } \mathcal{O}}^m, \text{gr } \mathcal{M}(n)) = 0$. But $\text{gr } \mathcal{M}(n) = \bigoplus_i (I^i/I^{i+1})\mathcal{M}_i(n)$, and so each $H^1(\mathbb{P}_{A/I^{i+1}}^m, (I^i/I^{i+1})\mathcal{M}_i(n)) = 0$. By the long exact sequence associated to the short exact sequence $0 \rightarrow (I^i/I^{i+1})\mathcal{M}_i(n) \rightarrow \mathcal{M}_i(n) \rightarrow \mathcal{M}_{i-1}(n) \rightarrow 0$, this H^1 is the obstruction to lifting sections of $\mathcal{M}_{i-1}(n)$ to sections of $\mathcal{M}_i(n)$. So choosing such an n which is also large enough so that $\mathcal{M}_0(n)$ is generated by its global sections, we can carry out the liftings inductively and thereby obtain Theorem 3.2.7.

As indicated above, Grothendieck's Existence Theorem is a strong enough form of "formal GAGA" to be useful in proving formal analogs of (at least parts of) the classical Riemann Existence Theorem. (This will be discussed further in Section 3.3.) But for certain purposes, it is useful to have a variant of Theorem 3.2.4 that allows U_1 and S_1 to be more local. Namely, rather than taking U_1 to be an affine open subset of the closed fibre, and S_1 its formal thickening, we would instead like to take U_1 to be the spectrum of the complete local ring in the closed fibre at some point ω , and S_1 its formal thickening (viz. the spectrum of the complete local ring at ω in V). In the relative dimension 1 case, the "overlap" U_0 of U_1 and U_2 is then the spectrum of the fraction field of the complete local ring at ω in the closed fibre, and S_0 is its formal thickening.

More precisely, in the case that V is of relative dimension 1 over A , there is the following formal patching result. First we introduce some notation and terminology. If ω is a closed point of a variety V_0 , then $\mathcal{K}_{V_0, \omega}$ denotes the total ring of fractions of the complete local ring $\hat{\mathcal{O}}_{V_0, \omega}$ (and thus the fraction field of $\hat{\mathcal{O}}_{V_0, \omega}$, if the latter is a domain). Let A be a complete local ring with maximal ideal \mathfrak{m} , let V be an A -scheme, and let V_n be the fibre of V over \mathfrak{m}^{n+1} (regarding $V_n \subset V_{n+1}$). Let $\omega \in V_0$, and let ω' denote $\text{Spec } \mathcal{K}_{V_0, \omega}$. Then *the ring of holomorphic functions in V at ω'* is defined to be $\hat{\mathcal{O}}_{V, \omega'} := \varprojlim \mathcal{K}_{V_n, \omega}$. (For example, if $A = k[[t]]$ and V is the affine x -line over A , and if ω is the point $x = t = 0$, then $\omega' = \text{Spec } k((x))$, $\mathcal{K}_{V_n, \omega} = k((x))[t]/(t^{n+1})$, and the ring of holomorphic functions at ω' is $\hat{\mathcal{O}}_{V, \omega'} = k((x))[[t]]$.)

Theorem 3.2.8. *Let V be a proper curve over a complete local ring A , let V_0 be the fibre over the closed point of $\text{Spec } A$, let W be a non-empty finite set of closed points of V_0 , and let $U = V_0 - W$. Let W^* be the union of the spectra of the complete local rings $\hat{\mathcal{O}}_{V, \omega}$ for $\omega \in W$. Let $U^* = \text{Spec } \hat{\mathcal{O}}_{V, U}$, and let $W'^* = \bigcup_{\omega \in W} \text{Spec } \hat{\mathcal{O}}_{V, \omega'}$, where $\omega' = \text{Spec } \mathcal{K}_{V_0, \omega}$ as*

above. Then the base change functor

$$\mathfrak{M}(V) \rightarrow \mathfrak{M}(W^*) \times_{\mathfrak{M}(W'^*)} \mathfrak{M}(U^*)$$

is an equivalence of categories. The same holds for finite algebras and for (Galois) covers.

This result appeared as [Ha6, Theorem 1], in the special case that V is regular, $A = k[[t_1, \dots, t_n]]$ for some field k and some $n \geq 0$, and where attention is restricted to projective modules. The proof involved showing that the appropriate form of Cartan’s Lemma is satisfied. In the form above, the result appeared at [Pr1, Theorem 3.4]. There, it was assumed that the complete local ring A is a discrete valuation ring, but that hypothesis was not necessary for the proof there. Namely, the proof there first showed the result for A/\mathfrak{m}^n , where \mathfrak{m} is the maximal ideal of A , using Corollary 3.1.9(a) (to the result of Ferrand and Raynaud [FR]); and afterwards used Grothendieck’s Existence Theorem (Theorem 3.2.1 above) to pass to A . (This use of [FR] was suggested by L. Moret-Bailly.)




Figure 3.2.9: Example 3.2.10 of Theorem 3.2.8, with $V = \mathbb{P}_k^1$, $W =$ one point. Here \mathbb{P}_k^1 is covered by the small patch $W^* = \text{Spec } k[[x, t]]$ and the larger patch $U^* = \text{Spec } k[1/x][[t]]$; their overlap is $W'^* = \text{Spec } k((x))[[t]]$ (upper right). Compare Fig. 3.1.4.

Example 3.2.10. Let k be a field, let $A = k[[t]]$, and let $V = \mathbb{P}_A^1$ (the projective x -line over $k[[t]]$), with closed fibre $V_0 = \mathbb{P}_k^1$ over $(t = 0)$. Let W consist of the single point ω where $x = t = 0$. In the notation of Theorem 3.2.8, $W^* = \text{Spec } k[[x, t]]$, which can be viewed as a “small neighborhood” of ω . The formal completion of V along $U := V_0 - W$ is $U^* = \text{Spec } k[1/x][[t]]$, whose “overlap” with W^* is $W'^* = \text{Spec } k((x))[[t]]$. (See Figure 3.2.9.) According to Theorem 3.2.8, giving a coherent module on V is equivalent to giving finite modules over W^* and over U^* together with an isomorphism on their

pullbacks (“restrictions”) to W'^* . The same holds for covers; and this permits modifying a branched cover of V near ω , e.g. by adding more inertia there (see Remarks 5.1.6(d,e)). \square

Example 3.2.11. Let k, A be as in Example 3.2.10, and let V be an irreducible normal curve over A , with closed fibre V_0 . Then V_0 is a k -curve which is connected (by Zariski’s Connectedness Theorem [Za4, §20, Thm. 14], [Hrt2, III, Cor. 11.3]) but not necessarily irreducible; let V_1, \dots, V_r be its irreducible components. The singular locus of V is a finite subset of V_0 , and it includes all the points where irreducible components V_i of V_0 intersect. Let W be a finite subset of V_0 that contains this singular locus, and contains at least one point on each irreducible component V_i of V_0 . For $i = 1, \dots, r$ let $W_i = V_i \cap W$, let $U_i = V_i - W_i$, and consider the ring $\hat{\mathcal{O}}_{V, U_i}$ of holomorphic functions along U_i . Also, for each point ω in W , we may consider its complete local ring $\hat{\mathcal{O}}_{V, \omega}$ in V . According to Theorem 4, giving a coherent module on V is equivalent to giving finite modules over each $\hat{\mathcal{O}}_{V, U_i}$ and over each $\hat{\mathcal{O}}_{V, \omega}$ together with isomorphisms on the “overlaps”. See [HS] for a formalization of this set-up. \square

Theorem 3.2.8 above can be generalized to allow V to be higher dimensional over the base ring A . In addition, by replacing the result of Ferrand-Raynaud (Proposition 3.1.6) by the related result of Artin (Proposition 3.1.8), one can take a proper morphism $\tilde{V} \rightarrow V$ and work over \tilde{V} rather than over V itself. Both of these generalizations are accomplished in the following result:

Theorem 3.2.12. *Let (A, \mathfrak{m}) be a complete local ring, let V be a proper A -scheme, and let $f : \tilde{V} \rightarrow V$ be a proper morphism. Let W be a finite set of closed points of V ; let $\tilde{W} = f^{-1}(W) \subset \tilde{V}$; let $W^* = \bigcup_{\omega \in W} \text{Spec } \hat{\mathcal{O}}_{V, \omega}$; and let $\tilde{W}^* = \tilde{V} \times_V W^*$. Let $\tilde{\mathcal{U}}$ [resp. $\tilde{\mathcal{U}}^*$] be the formal completion of $\tilde{V} - \tilde{W}$ [resp. of $\tilde{W}^* - \tilde{W}$] along its fibre over \mathfrak{m} . Then the base-change functor*

$$\mathfrak{M}(\tilde{V}) \rightarrow \mathfrak{M}(\tilde{W}^*) \times_{\mathfrak{M}(\tilde{\mathcal{U}}^*)} \mathfrak{M}(\tilde{\mathcal{U}})$$

is an equivalence of categories. The same holds for finite algebras and for (Galois) covers.

Note that the scheme $U^* = \text{Spec } \hat{\mathcal{O}}_{V, U}$ in the statement of Theorem 3.2.8 is replaced in Theorem 3.2.12 by a formal scheme, because the complement of W in the closed fibre of V will no longer be affine, if V is not a curve over its base ring (and so the ring $\hat{\mathcal{O}}_{V, U}$ of Theorem 4 would not be defined here). Similarly, the scheme W'^* in Theorem 3.2.8 is also replaced by a formal scheme in Theorem 3.2.12.

Proof. For $n \geq 0$ let \tilde{V}_n and \tilde{W}_n^* be the pullbacks of \tilde{V} and \tilde{W}^* , respectively, over $A_n := A/\mathfrak{m}^{n+1}$. Also, let $\tilde{U}_n = \tilde{V}_n - \tilde{W}_n$ and $\tilde{U}_n^* = \tilde{W}_n^* - \tilde{W}_n$; thus the formal schemes $\tilde{\mathcal{U}}, \tilde{\mathcal{U}}^*$ respectively correspond to the inverse systems $\{\tilde{U}_n\}_n, \{\tilde{U}_n^*\}_n$.

For every n , we have by Corollary 3.1.9(b) (to Artin's result, Proposition 3.1.8) that the base change functor

$$\mathfrak{M}(\tilde{V}_n) \rightarrow \mathfrak{M}(\tilde{W}_n^*) \times_{\mathfrak{M}(\tilde{U}_n^*)} \mathfrak{M}(\tilde{U}_n)$$

is an equivalence of categories. By definition of coherent modules over a formal scheme, we have that $\mathfrak{M}(\tilde{U}) = \varprojlim \mathfrak{M}(\tilde{U}_n)$ and $\mathfrak{M}(\tilde{U}^*) = \varprojlim \mathfrak{M}(\tilde{U}_n^*)$. Moreover, \tilde{V} is proper over A ; so Grothendieck's Existence Theorem (Theorem 3.2.1 above) implies that the functor $\mathfrak{M}(\tilde{V}) \rightarrow \varprojlim \mathfrak{M}(\tilde{V}_n)$ is an equivalence of categories. So it remains to show that the corresponding assertion holds for $\mathfrak{M}(\tilde{W}^*)$; i.e. that $\mathfrak{M}(\tilde{W}^*) \rightarrow \varprojlim \mathfrak{M}(\tilde{W}_n^*)$ is an equivalence of categories. In doing so, we immediately reduce to the case that W consists of just one point ω , which we now assume.

Let $T = \hat{\mathcal{O}}_{V,\omega}$, and let \mathfrak{m}_ω be the maximal ideal of T , corresponding to the closed point ω . Also, let $\mathfrak{n} = \mathfrak{m}T \subset T$. Thus $\mathfrak{n} \subset \mathfrak{m}_\omega$ (but unfortunately they are not equal — necessitating a more involved argument than otherwise). For every $n, m \geq 1$ let $T_n = T/\mathfrak{n}^n$, let $T^m = T/\mathfrak{m}_\omega^m$ and $T_n^m = T/\mathfrak{n}^n \mathfrak{m}_\omega^m$. So T^n is a quotient of T_n , and the quotient map $T^m \rightarrow T^{n+m}$ factors as the composition $T^m \rightarrow T_n^m \rightarrow T^{n+m}$. Let $\mathfrak{m}_{\omega,n} = \mathfrak{m}_\omega / (\mathfrak{m}_\omega \cap \mathfrak{n}^n)$; this is the maximal ideal in the complete local ring T_n . Finally, let \tilde{W}^{*m} and \tilde{W}_n^{*m} be the pullbacks of \tilde{W}^* to T^m and T_n^m respectively. Thus $\{\tilde{W}^{*m}\}_m$ is cofinal in the doubly indexed inverse system $\{\tilde{W}_n^{*m}\}_{n,m}$.

Since \tilde{W}^* is proper over the complete local ring T , and since \tilde{W}_n^* is proper over the complete local ring T_n for each n , Grothendieck's Existence Theorem asserts that the obvious functors $\mathfrak{M}(\tilde{W}^*) \rightarrow \varprojlim_{\bar{m}} \mathfrak{M}(\tilde{W}^{*m})$ and $\mathfrak{M}(\tilde{W}_n^*) \rightarrow \varprojlim_{\bar{m}} \mathfrak{M}(\tilde{W}_n^{*m})$ are (compatible) equivalences of categories. But $\varprojlim_{\bar{m}} \mathfrak{M}(\tilde{W}^{*m})$ can be identified with $\varprojlim_{\bar{n}} \varprojlim_{\bar{m}} \mathfrak{M}(\tilde{W}_n^{*m})$ because of the above cofinality. So the functor $\mathfrak{M}(\tilde{W}^*) \rightarrow \varprojlim_{\bar{n}} \mathfrak{M}(\tilde{W}_n^*)$ is an equivalence of categories, as desired. \square

Example 3.2.13. Let k, A be as in Examples 3.2.10 and 3.2.11, and let $V = \mathbb{P}_k^n$ for some $n \geq 1$, with homogeneous coordinates x_0, \dots, x_n . Let W consist of the closed point ω of V where $x_1 = \dots = x_n = t = 0$, and let $f : \tilde{V} \rightarrow V$ be the blow-up of V at ω . Let $V_0 = \mathbb{P}_k^n$ be the closed fibre of V over $(t = 0)$. For $i = 1, \dots, n$, let U_i be the affine open subset of V_0 given by $x_i \neq 0$, and consider the ring $\hat{\mathcal{O}}_{V,U_i}$ of holomorphic functions along U_i in V . Also consider the complete local ring $\hat{\mathcal{O}}_{V,\omega} = k[[x_1, \dots, x_n, t]]$ at ω in V , and consider the pullback \tilde{W}^* of \tilde{V} over $\hat{\mathcal{O}}_{V,\omega}$ (whose fibre over the closed point ω is a copy of \mathbb{P}_k^n). According to Theorem 3.2.12, giving a coherent module over V is equivalent to giving finite modules over the rings $\hat{\mathcal{O}}_{V,U_i}$, and a coherent module over \tilde{W}^* , together with compatible isomorphisms on the overlaps. (This uses that giving a coherent module on the formal completion of $V - W$ along its closed fibre is equivalent to giving compatible modules over the completions at the U_i 's; here we also identify $\tilde{V} - f^{-1}(W)$ with $V - W$.)

In particular, if $n = 1$, then \tilde{V} is an irreducible A -curve whose closed fibre consists of two projective lines meeting at one point (one being the proper transform of the given line V_0 , and the other being the exceptional divisor). This one-dimensional case is also within the context of Example 3.2.11, and so Theorem 3.2.8 could instead be used. (See also the end of Example 4.2.4 below.) \square

Remark 3.2.14. The above formal patching results (Theorems 3.2.4, 3.2.8, 3.2.12) look similar, though differing in terms of what types of “patches” are allowed. In each case, we are given a proper scheme V over a complete local ring A , and the assertion says that if a module is given over each of two patches (of a given form), with agreement on the “overlap”, then there is a unique coherent module over V that induces them compatibly. Theorem 3.2.4 (a reformulation of Grothendieck’s Existence Theorem) is the basic version of formal patching, modeled after the classical result GAGA in complex patching (see Theorem 2.2.6, where two metric open sets are used as patches). In Theorem 3.2.4, the patches correspond to thickenings along Zariski open subsets of the closed fibre of V ; see Example 3.2.3 above and see Figure 3.1.4 for an illustration. This basic type of formal patching will be sufficient for the results of Section 3.3 below, on the realization of Galois groups, via “slit covers”.

More difficult results about fundamental groups, discussed in Section 5 below, require Theorems 3.2.8 or 3.2.12 instead of Theorem 3.2.4 (e.g. Theorem 5.1.4 and Theorem 5.3.1 use Theorem 3.2.8, while Theorem 5.3.9 uses Theorem 3.2.12). In Theorem 3.2.8 above, one of the patches is allowed to be much smaller than in Theorem 3.2.4, viz. the spectrum of the complete local ring at a point, if the closed fibre is a curve; see Examples 3.2.10 and 3.2.11 above, and see Figure 3.2.9 above for an illustration. Theorem 3.2.12 is still more general, allowing the closed fibre to have higher dimension, and also allowing a more general choice of “small patch” because of the choice of a proper morphism $\tilde{V} \rightarrow V$; see Example 3.2.13 above. The advantage of these stronger results is that the overlap of the patches is “smaller” than in the situation of Theorem 3.2.4, and therefore less agreement is required between the given modules. This gives greater applicability to the patching method, in constructing modules or covers with given properties. (Recall that the similar-looking patching results at the end of Section 3.1, which allow the construction of modules by prescribing them along and away from a given closed set, do not directly give results for covers; but they were used, together with Grothendieck’s Existence Theorem, in proving Theorems 3.2.8 and 3.2.12 above.) \square

Section 3.3. Formal patching and constructing covers.

The methods of Section 3.2 allow one to construct covers of algebraic curves over various fields other than the complex numbers. The idea is to use the approach of Section 2.3, building “slit covers” using formal patching rather than analytic patching (as was used in Section 2). This will be done by relying on Grothendieck’s Existence Theorem,

in the form of Theorem 3.2.4. (As will be discussed in Section 5, by using variants of Theorem 3.2.4, in particular Theorems 3.2.8 and 3.2.12, it is possible to make more general constructions as well. See also [Ha6], [St1], [HS1], and [Pr2] for other applications of those stronger patching results, concerning covers with given inertia groups over certain points, or even unramified covers of projective curves.)

The first key result is

Theorem 3.3.1. [Ha4, Theorem 2.3, Corollary 2.4] *Let R be a normal complete local domain other than a field; let K be the fraction field of R ; and let G be a finite group. Then G is the Galois group of a Galois field extension L of $K(x)$, which corresponds to a Galois branched cover of \mathbb{P}_K^1 with Galois group G . Moreover L can be chosen to be regular, in the sense that K is algebraically closed in L .*

Before discussing the proof, we give several examples:

Example 3.3.2. a) Let $K = \mathbb{Q}_p$, or a finite extension of \mathbb{Q}_p , for some prime p . Then every finite group is a Galois group over \mathbb{P}_K^1 (i.e. of some Galois branched cover of the K -line), and so is a Galois group over $K(x)$.

b) Let k be a field, let n be a positive integer, and let $K = k((t_1, \dots, t_n))$, the fraction field of $k[[t_1, \dots, t_n]]$. Then every finite group is a Galois group over \mathbb{P}_K^1 , and so over $K(x)$.

c) If K is as in Example (b) above, and if $n > 1$, then every finite group is a Galois group over K (and not just over $K(x)$, as above). The reason is that K is separably Hilbertian, by Weissauer's Theorem [FJ, Theorem 14.17]. That is, every separable field extension of $K(x)$ specializes to a separable field extension of K , by setting $x = c$ for an appropriate choice of $c \in K$; such a specialization of a Galois field extension is then automatically Galois. (The condition of being separably Hilbertian is a bit weaker than being Hilbertian, but is sufficient for dealing with Galois extensions. See [FJ, Chapter 11], [Vö, Chapter 1], or [MM, Chapter IV, §1.1] for more about Hilbertian and separably Hilbertian fields.)

This example remains valid more generally, where the coefficient field k is replaced by any Noetherian normal domain A that is complete with respect to a prime ideal. Moreover if A is not a field, then the condition $n > 1$ can even be weakened to $n > 0$. In particular, if K is the fraction field of $\mathbb{Z}[[t]]$ (a field which is much smaller than $\mathbb{Q}((t))$), then every finite group is the Galois group of a regular cover of \mathbb{P}_K^1 , and is a Galois group over K itself. The proof of this generalization uses formal A -schemes, and parallels the proof of Theorem 1; see [Le].

d) Let K be the ring of algebraic p -adics (i.e. the algebraic closure of \mathbb{Q} in \mathbb{Q}_p), or alternatively the ring of algebraic Laurent series in n -variables over a field k (i.e. the algebraic closure of $k(t_1, \dots, t_n)$ in $k((t_1, \dots, t_n))$). Then every finite group is a Galois group over \mathbb{P}_K^1 . More generally this holds if K is the fraction field of R , a normal henselian

local domain other than a field. This follows by using Artin’s Algebraization Theorem ([Ar3], a consequence of Artin’s Approximation Theorem [Ar2]), in order to pass from formal elements to algebraic ones. See [Ha4, Corollary 2.11] for details. In the case of algebraic power series in $n > 1$ variables, Weissauer’s Theorem then implies that every finite group is a Galois group over K , as in Example (c). \square

Theorem 3.3.1 also implies that all finite groups are Galois groups over $K(x)$ for various other fields K , as discussed below (after the proof).

Theorem 3.3.1 can be proven by carrying over the slit cover construction of Section 2.3 to the context of formal schemes. Before doing so, it is first necessary to construct cyclic covers that can be patched together (as in Example 2.3.2). Rather than using complex discs as in §2.3, we will use “formal open subsets”, i.e. we will take the formal completions of \mathbb{P}_R^1 along Zariski open subsets of the closed fibre \mathbb{P}_k^1 (where k is the residue field of R). In order to be able to use Grothedieck’s Existence Theorem to patch these covers together, we will want the cyclic covers to agree on the “overlaps” of these formal completions — and this will be accomplished by having them be trivial on these overlaps (just as in Example 2.3.2).

In order to apply Grothedieck’s Existence Theorem, we will use it in the case of Galois branched covers (rather than for modules), as in Theorem 3.2.4. There, it was stated just for two patches U_1, U_2 and their overlap U_0 ; but by induction, it holds as well for finitely many patches, provided that compatible isomorphisms are given on overlaps (and cf. the statement of Theorem 3.2.1).

Grothedieck’s Existence Theorem will be applied to the following proposition, which yields the cyclic covers $Y \rightarrow \mathbb{P}^1$ that will be patched together in order to prove Theorem 3.3.1. The desired triviality on overlaps will be guaranteed by the requirement that the closed fibre $\phi_k : Y_k \rightarrow \mathbb{P}_k^1$ of the branched cover $\phi : Y \rightarrow \mathbb{P}_R^1$ be a *mock cover*; i.e. that the restriction of ϕ_k to each irreducible component of Y_k be an isomorphism. This condition guarantees that if $U \subset \mathbb{P}_k^1$ is the complement of the branch locus of ϕ_k , then the restriction of ϕ_k to U is trivial; i.e. $\phi_k^{-1}(U)$ just consists of a disjoint union of copies of U .

Proposition 3.3.3. [Ha4, Lemma 2.1] *Let (R, \mathfrak{m}) be a normal complete local domain other than a field, with fraction field K and residue field $k = R/\mathfrak{m}$. Let $S \subset \mathbb{P}_k^1$ be a finite set of closed points, and let $n > 1$. Then there is a cyclic field extension L of $K(x)$ of degree n , such that the normalization of \mathbb{P}_R^1 in L is an n -cyclic Galois branched cover $Y \rightarrow \mathbb{P}_R^1$ whose closed fibre $Y_k \rightarrow \mathbb{P}_k^1$ is a mock cover that is unramified over S .*

Proof. We follow the proof in [Ha4], first observing that we are reduced to the situation that n is a prime power p^r . (Namely, if $n = \prod p_i^{r_i}$, and if $Y_i \rightarrow \mathbb{P}_R^1$ are $p_i^{r_i}$ -cyclic covers, then we may take Y to be the fibre product of the Y_i ’s over \mathbb{P}_R^1 .)

The easiest case is if the field K contains a primitive n th root of unity ζ_n . Then we

may take L to be the field obtained by adjoining an n th root of $f(x)(f(x) - \alpha)^{n-1}$, where $f(x) \in R[x]$ does not vanish at any point of S , and where $\alpha \in \mathfrak{m} - \{0\}$. (For example, if k is infinite, we may choose $f(x) = x - c$ for some $c \in R$; compare Example 2.3.2.)

Next, suppose that K does not contain a primitive n th root of unity but that p is not equal to the characteristic of K . Then we can consider $K' = K[\zeta_n]$, and will construct an n -cyclic Kummer extension of $K'(x)$ which descends to a desired extension of $K(x)$. This will be done using constructions in [Slt] to find an element $g(x) \in R[\zeta_n, x]$ such that the extension $y^n - g(x)$ of $R[\zeta_n, x]$ descends to an n -cyclic extension of $R[x]$ whose closed fibre is a mock cover.

Specifically, first suppose that p is odd. Let s be the order of the cyclic group $\text{Gal}(K'/K)$, with generator $\tau : \zeta_n \mapsto \zeta_n^m$. Choose $\alpha \in \mathfrak{m} - \{0\}$ and let $b = f(x)^n - \zeta_n p^2 \alpha$, for some $f(x) \in R[x]$ which does not vanish on S . Let L' be the n -cyclic field extension of $K'(x)$ given by adjoining an n th root of $M(b) = b^{m^{s-1}} \tau(b)^{m^{s-2}} \dots \tau^{s-2}(b)^m \tau^{s-1}(b)$. Then $L' = L \otimes_K K'$ for some n -cyclic extension L of $K(x)$, by [Slt, Theorem 2.3]. (Note that the branch locus of the associated cover, which is given by $M(b) = 0$, is invariant under τ . Here the various powers of the factors of $M(b)$ are chosen so that τ will commute with the generator of $\text{Gal}(L'/K'(x))$, given by $y \mapsto \zeta_n y$. These two facts enable the Kummer cover of the K' -line to descend to a cyclic cover of the K -line.)

On the other hand, suppose $p = 2$. If K contains a square root of -1 then $\text{Gal}(K'/K)$ is again cyclic, so the same proof as in the odd case works. Otherwise, if $n = 2$ then take the extension of $K(x)$ given by adjoining a square root of $f(x)^2 - 4\alpha$. If $n = 4$, then adjoin a fourth root of $(f(x)^4 + 4i\alpha)^3 (f(x)^4 - 4i\alpha)$ to $K'(x)$; this descends to a 4-cyclic extension of $K(x)$ by [Slt, Theorem 2.4]. If $n = 2^r$ with $r \geq 3$, then $\text{Gal}(K'/K)$ is the product of a cyclic group of order 2 with generator $\kappa : \zeta_n \mapsto \zeta_n^{-1}$, and another of order $s \leq 2^{n-2}$ with generator $\zeta_n \mapsto \zeta_n^m$ for some $m \equiv 1 \pmod{4}$. Take $b = f(x)^n + 4\zeta_n \alpha$ and $a = b^{2^{n-1}+1} \kappa(b)^{2^{n-1}-1}$; and (in the notation of the odd case) consider the extension of $K'(x)$ given by adjoining an n th root of $M(a)$. By [Slt, Theorem 2.7], this descends to an n -cyclic extension of $K(x)$.

Finally, there is the case that p is equal to the characteristic of K . If $n = p$, we can adjoin a root of an Artin-Schreier polynomial $y^p - f(x)^{p-1}y - \alpha$, where $f(x) \in R[x]$ and $\alpha \in \mathfrak{m} - \{0\}$. More generally, with $n = p^r$, we can use Witt vectors, by adjoining the roots of the Witt coordinates of $\text{Fr}(y) - f(x)^{p-1}y - \alpha$, where $f(x)$ and y denote the elements of the truncated Witt ring $W_r(R[x, y_0, \dots, y_{r-1}])$ with Witt coordinates $(f(x), 0, \dots, 0)$ and (y_0, \dots, y_n) respectively, and where Fr denotes Frobenius.

In each of these cases, one checks that the extension L of $K(x)$ has the desired properties. (See [Ha4, Lemma 2.1] for details.) \square

Using this result together with Grothendieck's Existence Theorem (for covers), one easily obtains Theorem 3.3.1:

Proof of Theorem 3.3.1. Let G be a finite group, and let g_1, \dots, g_r be generators. Let H_i be the cyclic subgroup of G generated by g_i . By Proposition 3.3.3, for each i there is an irreducible normal H_i -Galois cover $Y_i \rightarrow \mathbb{P}_R^1$ whose closed fibre is a mock cover of \mathbb{P}_k^1 ; moreover these covers may be chosen inductively so as to have disjoint branch loci B_i (by choosing them so that the branch loci along the closed fibre are disjoint). For $i = 1, \dots, r$, let $U_i = \mathbb{P}_R^1 - \bigcup_{j \neq i} B_j$, let R_i be the ring of holomorphic functions on U_i along its closed fibre (i.e. the \mathfrak{m} -adic completion of the ring of functions on U_i), and let $\hat{U}_i = \text{Spec } R_i$. Also let $U_0 = \mathbb{P}_R^1 - \bigcup_{j=1}^r B_j$ (so that $U_0 = U_i \cap U_j$ for any $i \neq j$), let $H_0 = 1 \subset G$, and let $Y_0 = \mathbb{P}_R^1$. Then the restriction $\hat{Y}_i = Y_i \times_{\mathbb{P}_R^1} \hat{U}_i$ is an irreducible normal H_i -Galois cover, and we may identify the pullback $\hat{Y}_i \times_{\hat{U}_i} \hat{U}_0$ with the trivial cover $\hat{Y}_0 = \text{Ind}_1^{H_i} \hat{U}_0$. Finally, let $\hat{Z}_i = \text{Ind}_{H_i}^G \hat{Y}_i$; this is a (disconnected) G -Galois cover of \hat{U}_i , equipped with an isomorphism $\hat{Z}_i \times_{\hat{U}_i} \hat{U}_0 \simeq \hat{Z}_0$. By Grothendieck's Existence Theorem for covers (see Theorem 3.2.4), there is a unique G -Galois cover $Z \rightarrow \mathbb{P}_R^1$ whose restriction to \hat{U}_i is \hat{Z}_i , compatibly. This cover is connected since its closed fibre is (because H_1, \dots, H_r generate G); it is normal since each \hat{Z}_i is; and so it is irreducible (being connected and normal). The closed fibre of Z is a mock cover (and so reducible), since the same is true for each \hat{Z}_i ; and so K is algebraically closed in the function field L of Z . So L is as desired in Theorem 3.3.1. \square

Remark 3.3.4. A variant approach to Theorem 3.3.1 involves proving a modification of Proposition 3.3.3 — viz. requiring that Y_k contains a k -point that is not in the ramification locus of $Y_k \rightarrow \mathbb{P}_k^1$, rather than requiring that $Y_k \rightarrow \mathbb{P}_k^1$ is a mock cover. This turns out to be sufficient to obtain Theorem 3.3.1, e.g. by showing that after a birational change of variables on \mathbb{P}^1 , the cover Y is taken to a cover whose closed fibre is a mock cover (and thereby recapturing the original proposition above). This modified version of the proposition can be proven by first showing that there is *some* n -cyclic extension of $K(x)$, e.g. as in [FJ, Lemma 24.46]; and then adjusting the extension by a “twist” in order to obtain an unramified rational point [HV, Lemma 4.2(a)]. (In general, this twisting method works for abelian covers, and so in particular for cyclic covers.) This modified proposition first appeared in [Li], where it was used to provide a proof of Theorem 3.3.1 using rigid analytic spaces, rather than formal schemes. See Theorem 4.3.1 below for a further discussion of this. \square

As mentioned just after the statement of Theorem 3.3.1 above, that result can be used to deduce that many other fields K have the same inverse Galois property, even without being complete. In particular:

Corollary 3.3.5. [Ha3, Corollary 1.5] *Let k be an algebraically closed field. Then every finite group is a Galois group over $k(x)$; or equivalently, it is the Galois group of some branched cover of the k -line.*

In the case of $k = \mathbb{C}$, this result is classical, and was the subject of Section 2 above,

where the proof involved topology and analytic patching. For a more general algebraically closed field, the proof uses Theorem 3.3.1 above and a trick that relies on the fact that every finite extension is given by finitely many polynomials (also used in the remark after Corollary 2.1.5):

Proof of Corollary 3.3.5. Let $R = k[[t]]$ and $K = k((t))$. Applying Theorem 3.3.1 to R and a given finite group G , we obtain an irreducible G -Galois branched cover $Y \rightarrow \mathbb{P}_K^1$ such that K is algebraically closed in its function field. This cover is of finite type, and so it is defined (as a G -Galois cover) over a k -subalgebra A of K of finite type; i.e. there is an irreducible G -Galois branched cover $Y_A \rightarrow \mathbb{P}_A^1$ such that $Y_A \times_A K \approx Y$ as G -Galois branched covers of \mathbb{P}_K^1 . By the Bertini-Noether Theorem [FJ, Prop. 9.29], there is a non-zero element $\alpha \in A$ such that the specialization of Y_A to any k -point of $\text{Spec } A[\alpha^{-1}]$ is (geometrically) irreducible. Any such specialization gives an irreducible G -Galois branched cover of \mathbb{P}_k^1 . \square

In fact, as F. Pop later observed [Po4], the proof of the corollary relied on k being algebraically closed only to know that every $k((t))$ -point has a k -point. So for any field k with this more general property (a field k that is “existentially closed in $k((t))$ ”), the corollary holds as well. Moreover the resulting Galois extension of $k(x)$ can be chosen to be regular, i.e. with k algebraically closed in the extension, by the geometric irreducibility assertion in the Bertini-Noether Theorem. Pop proved [Po4, Proposition 1.1] that the fields k that are existentially closed in $k((t))$ can be characterized in another way: they are precisely those fields k with the property that every smooth k -curve with a k -rational point has infinitely many k -rational points. He called such fields “large”, because they are sufficiently large within their algebraic closures in order to recapture the finite-type argument used in the above corollary. (In particular, if k is large, then any extension field of k , contained in the algebraic closure of k , is also large [Po4, Proposition 1.2].) Thus we obtain the following strengthening of the corollary:

Theorem 3.3.6 [Po4] *Let k be a large field, and let G be a finite group.*

- a) *Then G is the Galois group of a Galois field extension L of $k(x)$, and the extension may be chosen to be regular.*
- b) *If k is (separably) Hilbertian, then G is a Galois group over k .*

Here part (b) follows from part (a) as in Example 3.3.2(c).

Example 3.3.7. a) Let K be a complete valuation field. Then K is large by [Po4, Proposition 3.1], the basic idea being that K satisfies an Implicit Function Theorem (and so one may move a K -rational point a bit to obtain other K -rational points). So every finite group is a Galois group over $K(x)$, by Theorem 3.3.6. In particular, this is true for the fraction field K of a complete discrete valuation ring R — as was already shown in

Theorem 3.3.1. On the other hand, Theorem 3.3.6 applies to complete valuation fields K that are not of that form.

b) More generally, a henselian valued field K (i.e. the fraction field of a henselian valuation ring) is large by [Po4, Proposition 3.1]. So again, every finite group is a Galois group over $K(x)$. If the valuation ring is a discrete valuation ring, then this conclusion can also be deduced using the Artin Algebraization Theorem, as in Example 3.3.2(d). But as in Example (a) above, K is large even if it is not discretely valued (in which case the earlier example does not apply).

c) It is immediate from the definition that a field k will be large if it is PAC (pseudo-algebraically closed); i.e. if every smooth geometrically integral k -variety has a k -point. Fields that are PRC (pseudo-real closed) or PpC (pseudo- p -adically closed) are also large. In particular, the field of all totally real algebraic numbers is large, and so is the field of totally p -adic algebraic numbers (i.e. algebraic numbers α such that $\mathbb{Q}(\alpha)$ splits completely over the prime p). Hence every finite group is a Galois group over $k(x)$, where k is any of the above fields. And if k is Hilbertian (as some PAC fields are), then every finite group is therefore a Galois group over k . See [Po4, Section 3] and [MB1, Thm. 1.3] for details.

d) Let K be a field that contains a large subfield K' . If K is algebraic over K' then K is automatically large [Po4, Proposition 1.2]; but otherwise K need not be large (e.g. $\mathbb{C}(t)$ is not large). Nevertheless, every finite group is the Galois group of a regular branched cover of $\mathbb{P}_{K'}^1$. The reason is that this property holds for K' ; and the function field F of the cover of $\mathbb{P}_{K'}^1$ is linearly disjoint from K over K' , because K' is algebraically closed in F (by regularity). In particular, we may use this approach to deduce Theorem 3.3.1 from Theorem 3.3.2, since every normal complete local domain R other than a field must contain a complete discrete valuation ring R_0 — whose fraction field is large. (Namely, if R contains a field k , then take $R_0 = k[[t]]$ for some non-zero element t in the maximal ideal of R ; otherwise, R contains \mathbb{Z}_p for some p .) Similarly, we may recover Example 3.3.2(d) in this way (taking the algebraic Laurent series in $k((t_1))$), even though it is not known whether $k((t_1, \dots, t_n))$ and its subfield of algebraic Laurent series are large. (Note that $k((t_1, \dots, t_n))$ is not a valuation field for $n > 1$, unlike the case of $n = 1$.) \square

Remark 3.3.8. a) An arithmetic analog of Example 3.3.7(b) holds for the ring $T = \mathbb{Z}\{t\}$ of power series over \mathbb{Z} convergent on the open unit disc. Namely, replacing Grothendieck's Existence Theorem by its arithmetic analog discussed in Remark 3.2.5(c) above, one obtains an analog of Theorem 3.3.1 above for $\mathbb{Z}\{t\}$ [Ha5, Theorem 3.7]; i.e. that every finite group is a Galois group over the fraction field of $\mathbb{Z}\{t\}$ (whose model over $\text{Spec } \mathbb{Z}\{t\}$ has a mock fibre modulo (t)). Moreover, the construction permits one to construct the desired Galois extension L of $\text{frac } T$ so that it remains a Galois field extension, with the same Galois group, even after tensoring with the fraction field of $T_r = \mathbb{Z}_{r+}[[t]]$, the ring of power series over \mathbb{Z} convergent on a neighborhood of the closed disc $|t| \leq r$. (Here $0 < r < 1$.)

Even more is true: Using an arithmetic analog of Artin’s Approximation Theorem (see [Ha5, Theorem 2.5]), it follows that these Galois extensions L_r of T_r can simultaneously be descended to a compatible system of Galois extensions L_r^h of $\text{frac } T_r^h$, where T_r^h is the ring of *algebraic* power series in T_r . Surprisingly, the intersection of the rings T_r^h has fraction field $\mathbb{Q}(t)$ [Ha2, Theorem 3.5] (i.e. every algebraic power series over \mathbb{Z} that converges on the open unit disc is *rational*). So since the Galois extensions L, L_r, L_r^h (for $0 < r < 1$) are all compatible, this suggests that it should be possible to descend the system $\{L_r^h\}$ to a Galois extension L^h of $\mathbb{Q}(t)$. If this could be done, it would follow that every finite group would be a Galois group over $\mathbb{Q}(t)$ and hence over \mathbb{Q} (since \mathbb{Q} is Hilbertian). See [Ha5, Section 4] for a further discussion of this (including examples that demonstrate pitfalls).

b) The field \mathbb{Q}^{ab} (the maximal abelian extension of \mathbb{Q}) is known to be Hilbertian [Vö, Corollary 1.28]. It is conjectured that this field is large; and if it is, then Theorem 3.3.6(b) above would imply that every finite group is a Galois group over \mathbb{Q}^{ab} . Much more is believed: The Shafarevich Conjecture asserts that the absolute Galois group of \mathbb{Q}^{ab} is a *free* profinite group on countably many generators. This conjecture has been posed more generally, to say that if K is a global field, then the absolute Galois group of K^{cycl} (the maximal cyclotomic extension of K) is a free profinite group on countably many generators. (Recall that $\mathbb{Q}^{\text{ab}} = \mathbb{Q}^{\text{cycl}}$, by the Kronecker-Weber Theorem in number theory.) The Shafarevich Conjecture (along with its generalization to arbitrary number fields) remains open — though it too would follow from knowing that \mathbb{Q}^{ab} is large (see Section 5). On the other hand, the generalized Shafarevich Conjecture has been proven in the geometric case, i.e. for function fields of curves [Ha10] [Po1] [Po3]; see Section 5 for a further discussion of this. □

As another example of the above ideas, consider covers of the line over *finite* fields. Not surprisingly (from the terminology), finite fields \mathbb{F}_q are not large. And it is unknown whether every finite group G is a Galois group over $k(x)$ for every finite field k . But it is known that every finite group G is a Galois group over $k(x)$ for *almost* every finite field k :

Proposition 3.3.9. (Fried-Völklein, Jarden, Pop) *Let G be a finite group. Then for all but finitely many finite fields k , there is a regular Galois field extension of $k(x)$ with Galois group G .*

Proof. First consider the case that k ranges just over prime fields \mathbb{F}_p . By Example 3.3.2(d) (or by Theorem 3.3.6 and Example 3.3.7(b) above), G is a regular Galois group over the field $\mathbb{Q}((t))^h(x)$, where $\mathbb{Q}((t))^h$ is the field of algebraic Laurent series over \mathbb{Q} (the t -adic henselization of $\mathbb{Q}(t)$). Such a G -Galois field extension is finite, so it descends to a G -Galois field extension of $K(x)$, where K is a finite extension of $\mathbb{Q}(t)$ (in which \mathbb{Q} is algebraically closed, since $K \subset \mathbb{Q}((t))$). This extension of $K(x)$ can be interpreted as the function field of a G -Galois branched cover $Z \rightarrow \mathbb{P}_V^1$; here V is a smooth projective curve over \mathbb{Q} with

function field K , viz. a finite branched cover of the t -line, say of genus g . For all points $\nu \in V$ outside some finite set Σ , the fibre of Z over ν is an irreducible G -Galois cover of $\mathbb{P}_{k(\nu)}^1$, where $k(\nu)$ is the residue field at ν . By taking a normal model $\mathcal{Z} \rightarrow \mathcal{V}$ of $Z \rightarrow V$ over \mathbb{Z} , we may consider the reductions V_p and Z_p for any prime p . For all primes p outside some finite set S , the reduction V_p is a smooth connected curve over \mathbb{F}_p of genus g ; the reduction Z_p is an irreducible G -Galois branched cover of $\mathbb{P}_{\mathbb{F}_p}^1$; and any specialization of this cover away from the reduction Σ_p of Σ is an irreducible G -Galois cover of the line. According to the Weil bound in the Riemann Hypothesis for curves over finite fields [FJ, Theorem 3.14], the number of k points on a k -curve of genus g is at least $|k| + 1 - 2g\sqrt{|k|}$. So for all $p \notin S$ with $p > (2g + \deg(\Sigma))^2$, the curve Z_p has an \mathbb{F}_p -point that does not lie in the reduction of Σ . The specialization at that point is a regular G -Galois cover of $\mathbb{P}_{\mathbb{F}_p}^1$, corresponding to a regular G -Galois field extension of $\mathbb{F}_p(x)$.




Figure 3.3.10: Base of the Galois cover $Z \rightarrow \mathbb{P}_V^1$ in the first case of the proof of Proposition 3.3.9. For most choices of ν in V , the restriction of Z over ν is an irreducible cover of the projective line; and for most primes p , the same is true for its reduction mod p .

For the general case, observe that if G is a regular Galois group over $\mathbb{F}_p(x)$, then it is also a regular Galois group over $\mathbb{F}_q(x)$ for every power q of p (by base change). Now

consider the finitely many primes p such that G is not known to be a regular Galois group over $\mathbb{F}_p(x)$. Arguing as above (but using $\mathbb{F}_p((t))$ instead of $\mathbb{Q}((t))$), we obtain a geometrically irreducible G -Galois cover $Y_p \rightarrow \mathbb{P}_{W_p}^1$, for some \mathbb{F}_p -curve W_p . Again using the Weil bound, there is a constant c_p such that if q is a power of p and $q > c_p$, then W_p has an \mathbb{F}_q -point at which Y_p specializes to a regular G -Galois cover of $\mathbb{P}_{\mathbb{F}_q}^1$. So if c is chosen larger than each of the finitely many c_p 's (as p ranges over the exceptional set of primes), then G is a regular Galois group over $k(x)$ for every finite field k of order $\geq c$. \square

Remark 3.3.11. a) The above result can also be proven via ultraproducts, viz. using that a non-principal ultraproduct of the \mathbb{F}_q 's is large (and even PAC); see [FV1, §2.3, Cor. 2]. In [FV1], just the case of prime fields was shown. But Pop showed that the conclusion holds for general finite fields (as in the statement of Proposition 3.3.9), using ultraproducts.

b) It is conjectured that in fact there are *no* exceptional primes in the above result, i.e. that every finite group is a Galois group over each $\mathbb{F}_q(x)$. But at least, it would be desirable to have a better understanding of the possible exceptional set of primes. For this, one could try to make more precise the sets S and Σ in the above proof, and also the bound on the exceptional primes. (The bound in the above proof is certainly not optimal.) \square

Remark 3.3.12. a) The class of large fields also goes under several other names in the literature. Following the introduction of this notion by Pop in [Po4] under the name “large”, D. Haran and M. Jarden referred to such fields as “ample” [HJ1]; P. Dèbes and B. Deschamps called them fields with “automatique multiplication des points lisses existants” (abbreviated AMPLE) [DD]; J.-L. Colliot-Thélène has referred to such fields as “epais” (thick); L. Moret-Bailly has called them “fertile” [MB2]; and the present author has even suggested that they be called “pop fields”, since the presence of a single smooth rational point on a curve over such a field implies that infinitely many rational points will “pop up”.

b) By whatever name, large fields form the natural context to generalize Corollary 3.3.5 above. As noted in Example 3.3.7(d), the class of fields K that *contain* large subfields also has the property that every finite group is a regular Galois group over $K(x)$; and this class is general enough to subsume Theorem 3.3.1, as well as Theorem 3.3.6. On the other hand, this Galois property holds for the fraction field of $\mathbb{Z}[[t]]$, as noted at the end of Example 3.3.2(c); but that field is not known to contain a large subfield. Conjecturally, *every* field K has the regular Galois realization property (see [Ha9, §4.5]; this conjecture has been referred to as the regular inverse Galois problem). But that degree of generality seems very far from being proved in the near future.

c) In addition to yielding regular Galois realizations, large fields have a stronger property: that *every finite split embedding problem is properly solvable* (Theorem 5.1.9 below). Conjecturally, all fields have this property (and this conjecture subsumes the one in Re-

mark (b) above). See Section 5 for more about embedding problems, and for other results in Galois theory that go beyond Galois realizations over fields. The results there can be proven using patching theorems from Section 3.2 (including those at the end of §3.2, which are stronger than Grothendieck’s Existence Theorem). \square

We conclude this section with a reinterpretation of the above patching construction in terms of thickening and deformation. Namely, as discussed after Theorem 3.2.1 (Grothendieck’s Existence Theorem), that earlier result can be interpreted either as a patching result or as a thickening result. Theorem 3.3.1 above, and its Corollary 3.3.5, relied on Grothendieck’s Existence Theorem, and were presented above in terms of patching. It is instructive to reinterpret these results in terms of thickening, and to compare these results from that viewpoint with the slit cover construction of complex covers, discussed in Section 2.3.

Specifically, the proof of Theorem 3.3.1 above yields an irreducible normal G -Galois cover $Z \rightarrow \mathbb{P}_R^1$ whose closed fibre is a connected mock cover $Z_0 \rightarrow \mathbb{P}_k^1$. Viewing $\text{Spec } R$ as a “small neighborhood” of $\text{Spec } k$, we can regard \mathbb{P}_R^1 as a “tubular neighborhood” of \mathbb{P}_k^1 ; and the construction of $Z \rightarrow \mathbb{P}_R^1$ can be viewed as a thickening (or deformation) of $Z_0 \rightarrow \mathbb{P}_k^1$, built in such a way that it becomes irreducible (by making it locally irreducible near each of the branch points). Regarding formal schemes as thickenings of their closed fibres (given by a compatible sequence of schemes over the R/\mathfrak{m}^i), this construction be viewed as the result of infinitesimal thickenings (over each R/\mathfrak{m}^i) which in the limit give the desired cover of \mathbb{P}_R^1 .

From this point of view, Corollary 3.3.5 above can be viewed as follows: As before, take $R = k[[t]]$ and as above obtain an irreducible normal G -Galois cover $Z \rightarrow \mathbb{P}_R^1$. Since this cover is of finite type, it is defined over a $k[t]$ -subalgebra E of R of finite type (i.e. there is a normal irreducible G -Galois cover $Z_E \rightarrow \mathbb{P}_E^1$ that induces $Z \rightarrow \mathbb{P}_R^1$), such that there is a maximal ideal \mathfrak{n} of E with the property that the fibre of $Z_E \rightarrow \mathbb{P}_E^1$ over the corresponding point $\xi_{\mathfrak{n}}$ is isomorphic to the closed fibre of $Z \rightarrow \mathbb{P}_R^1$ (viz. it is the mock cover $Z_0 \rightarrow \mathbb{P}_k^1$). The cover $Z_E \rightarrow \mathbb{P}_E^1$ can be viewed as a family of covers of \mathbb{P}_k^1 , parametrized by the variety $V = \text{Spec } E$, and which provides a deformation of $Z_0 \rightarrow \mathbb{P}_k^1$. A generically chosen member of this family will be an irreducible cover of \mathbb{P}_k^1 , and this G -Galois cover is then as desired.

In the case that $k = \mathbb{C}$, we can be even more explicit. There, we are in the easy case of Proposition 3.3.3 above, where the field contains the roots of unity, ramification is cyclic, and cyclic extensions are Kummer. So choosing generators g_1, \dots, g_r of G of orders n_1, \dots, n_r , and choosing corresponding branch points $x = a_1, \dots, a_r$ for the mock cover $Z_0 \rightarrow \mathbb{P}_{\mathbb{C}}^1$, we may choose $Z \rightarrow \mathbb{P}_R^1$ so that it is given locally by the (normalization of the) equation $z_i^{n_i} = (x - a_i)(x - a_i - t)^{n_i - 1}$ in a neighborhood of a point over $x = a_i, t = 0$ (and so the mock cover is given locally by $z_i^{n_i} = (x - a_i)^{n_i}$). By Artin’s Algebraization

Theorem [Ar3] (cf. Example 3.3.2(d) above), this cover descends to a cover $Z \rightarrow \mathbb{P}_{R^h}^1$, where $R^h \subset R = \mathbb{C}[[t]]$ is the ring of algebraic power series. Since that cover is of finite type, it can be defined over a $\mathbb{C}[t]$ -subalgebra of R^h of finite type; i.e. the cover further descends to a cover $Y_C \rightarrow \mathbb{P}_C^1$, where C is a complex curve together with a morphism $C \rightarrow \mathbb{A}_{\mathbb{C}}^1 = \text{Spec } \mathbb{C}[t]$, and where the fibre of Y_C over some point $\xi \in C$ over $t = 0$ is the given mock cover $Z_0 \rightarrow \mathbb{P}_{\mathbb{C}}^1$. This family $Y_C \rightarrow \mathbb{P}_C^1$ can be viewed as a family of covers of $\mathbb{P}_{\mathbb{C}}^1$ deforming the mock cover; and this deformation takes place by allowing the positions of the branch points to move. By the choice of local equations, if we take a typical point on C near ξ , the corresponding cover has $2r$ branch points $x = a_1, a'_1, \dots, a_r, a'_r$, with branch cycle description

$$(g_1, g_1^{-1}, \dots, g_r, g_r^{-1}) \quad (*)$$

(see Section 2.1 and the beginning of Section 2.3 for a discussion of branch cycle descriptions). So this is a slit cover, in the sense of Example 2.3.2. See also the discussion following that example, concerning the role of the mock cover as a degeneration of the typical member of this family (in which a'_i is allowed to coalesce with a_i).

For more general fields k , we may not be in the easy case of Proposition 3.3.3, and so may have to use more complicated branching configurations. As a result, the deformed covers may have more than $2r$ branch points, and they may come in clusters rather than in pairs. Moreover, while the tamely ramified branch points will move in \mathbb{P}^1 as one deforms the cover, wildly ramified branch points can stay at the same location (with just the Artin-Schreier polynomial changing; see the last case in the proof of Proposition 3.3.3).

Still, in the tame case, by following this construction with a further doubling of branch points, it is possible to pair up the points of the resulting branch locus so that the resulting cover has “branch cycle description” of the form $(h_1, h_1^{-1}, \dots, h_N, h_N^{-1})$, where each h_i is a power of some generator g_j . (Here, since we are not over \mathbb{C} , the notion of branch cycle description will be interpreted in the weak sense that the entries of the description are generators of inertia groups at some ramification points over the respective branch points.) This leads to a generalization of the “half Riemann Existence Theorem” (Theorem 2.3.5) from \mathbb{C} to other fields. Such a result (though obtained using the rigid approach rather than the formal approach) was proven by Pop [Po2]; see Section 4.3 below.

The construction in the tame case can be made a bit more general by allowing the r branch points $x = a_i$ of the mock cover to be deformed with respect to independent variables. For example, in the case $k = \mathbb{C}$, we can replace the ring R by $k[[t_1, t'_1, \dots, t_r, t'_r]]$ and use the (normalization of the) local equation $z_i^{n_i} = (x - a_i - t_i)(x - a_i - t'_i)^{n_i - 1}$ in a neighborhood of a point over $x = a_i$ on the closed fibre $\underline{t} = \underline{t}' = 0$. Using Artin’s Algebraization Theorem, we obtain a $2r$ -dimensional family of covers that deform the given mock cover, with each of the r mock branch points splitting in two, each moving independently. The resulting family $Z \rightarrow \mathbb{P}_V^1$ is essentially a component of a *Hurwitz*

family of covers (e.g. see [Fu1] and [Fr1]), which is by definition a total family $Y \rightarrow \mathbb{P}_H^1$ of covers of \mathbb{P}^1 over the moduli space H for branched covers with a given branch cycle description and variable branch points (the *Hurwitz space*). Here, however, a given cover is permitted to appear more than once in the family (though only finitely often), and part of the boundary of the Hurwitz space is included (in particular, the point of the parameter space V corresponding to the mock cover). That is, there is a finite-to-one morphism $V \rightarrow \bar{H}$, where \bar{H} is the compactification of H . From this point of view, the desirability of using branch cycle descriptions of the form $(*)$ is that one can begin with an easily constructed mock cover, and use it to construct algebraically a component of a Hurwitz space with this branch cycle description. See [Fr3] for more about this point of view.

As mentioned above, still more general formal patching constructions of covers can be performed if one replaces Grothendieck’s Existence Theorem by the variations at the end of Section 3.2. In particular, one can begin with a given irreducible cover, and then modify it near one point (e.g. by adding ramification there). Some constructions along these lines will be discussed in Section 5, in connection with the study of fundamental groups.

Section 4: Rigid patching

This section, like Section 3, discusses an approach to carrying over the ideas of Section 2 from complex curves to more general curves. The approach here is due to Tate, who introduced the notion of rigid analytic spaces. The idea here is to consider power series that converge on metric neighborhoods on curves over a valued field, and to “rigidify” the structure to obtain a notion of “analytic continuation”. Tate’s original point of view, which is presented in Section 4.1, is rather intuitive. But the details of carrying it out become somewhat complicated, as the reader will see (particularly with regard to the precise method of rigidifying “wobbly spaces”). A simplified approach, due to Grauert, Remmert, and Gerritzen, is discussed later in Section 4.1, including their approach to a rigid analog of GAGA. Section 4.2 then discusses a later reinterpretation of rigid geometry that is due to Raynaud, and which establishes a kind of “dictionary” between the formal and rigid set-ups (and allows rigid GAGA to be deduced from formal GAGA). Applications to the construction of Galois covers of curves are then presented in Section 4.3, including a version of the (geometric) regular inverse Galois problem, and Pop’s Half Riemann Existence Theorem. Additional applications of both rigid and formal geometry to Galois theory appear afterwards, in Section 5.

Section 4.1. Tate’s rigid analytic spaces.

Another approach to generalizing complex analytic notions to spaces over other fields is provided by Tate’s rigid analytic spaces. As in the formal approach discussed in Section 3, the rigid approach allows “small neighborhoods” of points, and permits objects (spaces,

maps, sheaves, covers) to be constructed by giving them locally and giving agreement on overlaps (i.e. “patching”). Here the small neighborhoods are metric discs, rather than formal neighborhoods of subvarieties, as in the formal patching approach.

This approach was introduced by Tate in [Ta], a 1962 manuscript which he never submitted for publication. The manuscript was circulated in the 1960’s by IHES, with the notation that it consisted of “private notes of J. Tate reproduced with(out) his permission”. Later, the paper was published in *Inventiones Mathematicae* on the initiative of the journal’s editors, who said in a footnote that they “believe that it is in the general interests of the mathematical community to make these notes available to everyone”.

Tate’s approach was motivated by the problem of studying bad reduction of elliptic curves (what we now know as the study of Tate curves; see e.g. [BGR, 9.7]). The idea is to work over a field K that is complete with respect to a non-trivial non-archimedean valuation — e.g. the p -adics, or the Laurent series over a coefficient field k . On spaces defined over such a field K , one can consider discs defined with respect to the metric on K ; and one can consider “holomorphic functions” on those discs, viz. functions given by power series that are convergent there. One then wants to work more globally by means of analytic continuation, and to carry over the classical results over \mathbb{C} (e.g. those of Section 2 above) to this context. As a result, one hopes to obtain a GAGA-type result, a version of Riemann Existence Theorem, the realization of all finite groups as Galois groups over $K(x)$, etc.

There are difficulties, however, that are caused by the fact that the topology on K is totally disconnected. For example, on the affine K -line, consider the characteristic function f_D of the open unit disc $|x| < 1$; i.e. $f(x) = 1$ for $|x| < 1$, and $f(x) = 0$ for $|x| \geq 1$. Then this function is continuous, and in a neighborhood of each point $x = x_0$ it is given by a power series. (Namely, on the open disc of radius 1 about $x = x_0$, it is identically 1 or identically 0, depending on whether or not $|x_0|$ is less than 1.) This is quite contrary to the situation over \mathbb{C} , where a holomorphic function is “rigid”, in the sense that it is determined by its values on any open disc. Thus, if one proceeds in the obvious way, objects will have a strictly local character, and there will be no meaningful “patching”.

Tate used two ideas to deal with this problem. The first of these is to consider functions that are locally given on *closed* discs, rather than on open discs, and to require agreement on overlapping boundaries. Note, though, that because the metric is non-archimedean, closed discs are in fact open sets. The second idea is to restrict the set of allowable maps between spaces, by choosing a class of maps that fulfills certain properties and creates a “rigid” situation.

Concerning the first of these ideas, let $K\{x\}$ denote the subring of $K[[x]]$ consisting of power series that converge on the closed unit disc $|x| \leq 1$. Because the metric is non-archimedean, this ring consists precisely of those series $\sum_{i=0}^{\infty} a_i x^i$ for which $a_i \rightarrow 0$ as

$i \rightarrow \infty$. Similarly, the power series in $K[[x_1, \dots, x_n]]$ that converge on the closed polydisc where each $|x_i| \leq 1$ form the ring $K\{x_1, \dots, x_n\}$ of series $\sum a_{\underline{i}} x^{\underline{i}}$, where \underline{i} ranges over n -tuples of non-negative integers, and where $a_{\underline{i}} \rightarrow 0$ as $|\underline{i}| \rightarrow \infty$. As an example, if $K = k((t))$ for some field k , then $K\{x\} = k[x][[t]][t^{-1}]$. (Verification of this equality is an exercise left to the reader.)

If $0 < r_1 \leq r_2$, then we may also consider the closed annulus $\{x \mid r_1 \leq |x| \leq r_2\}$. Since the metric is non-archimedean, this is an open subset, which we may consider even when $r_1 = r_2$. In particular, in the case $r_1 = r_2 = 1$, we may consider the ring $K\{x, x^{-1}\} = K\{x, y\}/(xy - 1)$ of functions converging on the annulus; this consists of doubly infinite series $\sum_{i=-\infty}^{\infty} a_i x^i$ such that $a_i \rightarrow 0$ as $|i| \rightarrow \infty$. Similarly, we may consider the ring $K\{x_1, \dots, x_n, x_1^{-1}, \dots, x_n^{-1}\} = K\{x_1, \dots, x_n, y_1, \dots, y_n\}/(x_i y_i - 1)$ of functions on the “poly-annulus” $|x_i| = 1$ (with $i = 1, \dots, n$). In the case that $K = k((t))$, we have that $K\{x, x^{-1}\} = k[x, x^{-1}][[t]][t^{-1}]$. (Verification of this is again left to the reader. In this situation, the one-dimensional rings $K\{x\}$ and $K\{x, x^{-1}\}$ are obtained by inverting t in the two-dimensional rings $k[x][[t]]$ and $k[x, x^{-1}][[t]]$; cf. Figure 3.1.4 above.)

In order to consider more general analytic “varieties” over K , Tate considered quotients of the rings $K\{x_1, \dots, x_n\}$ by ideals. He referred to such quotients by saying that they were of *topologically finite type*; these are also now referred to as *affinoid algebras* [BGR] or as *Tate algebras* [Ra1] (though the latter term is sometimes used only for the ring $K\{x_1, \dots, x_n\}$ itself [BGR]). Tate showed that a complete K -algebra A is an affinoid algebra if and only if it is a finite extension of some $K\{x_1, \dots, x_n\}$ [Ta, Theorem 4.4]; and in this case A is Noetherian, every ideal is closed, and the residue field of every maximal ideal is finite over K [Ta, Theorem 4.5]. The association $A \mapsto \text{Max } A$ is a contravariant functor from affinoid algebras to sets, where $\text{Max } A$ is the maximal spectrum of A . (The map $\text{Max } B \rightarrow \text{Max } A$ associated to $\phi : A \rightarrow B$ is denoted by ϕ° , and is called *rigid*.) Since A/ξ is a finite extension L of K for any $\xi \in \text{Max } A$, we may consider $f(\xi) \in L$ and $|f(\xi)| \in \mathbb{R}$ for any $f \in A$ (and thus regard A as a ring of functions on $\text{Max } A$). By an *affinoid variety*, we then mean a pair $\text{Sp } A := (\text{Max } A, A)$, where A is an affinoid algebra.

Tate defined an *affine subset* $Y \subset \text{Max } A$ to be a subset for which there is an affinoid algebra A_Y that represents the functor $h_Y : B \mapsto \{\phi : A \rightarrow B \mid \phi^\circ(\text{Max } B) \subset Y\}$; i.e. such that $h_Y(B) = \text{Hom}(A_Y, B)$. (This is called an *affinoid subdomain* in [BGR].) A *special affine subset* $Y \subset \text{Max } A$ is a subset of the form

$$Y = \{\xi \in \text{Max } A : |f_i(\xi)| \leq 1 (\forall i), |g_j(\xi)| \geq 1 (\forall j)\},$$

where $(f_i), (g_j)$ are finite families of elements of A . (These are called *Laurent domains* in [BGR].) Tate showed [Ta, Proposition 7.2] that every special affine subset is affine, viz. that if Y is given by $(f_i), (g_j)$ as above, then $A_Y = A\{f_i, g_j^{-1}\} := A\{x_i, y_j\}/(f_i - x_i, 1 - g_j y_j)$. Moreover if Y is an affine subset of $\text{Max } A$, then the canonical map $\text{Max } A_Y \rightarrow Y$ is a

bijection [Ta, Proposition 7.3]. In fact, it is a homeomorphism [Ta, Cor. 2 to Prop. 9.1], if we give $\text{Max } A$ the topology in which a fundamental system of neighborhoods of a point ξ_0 is given by sets of the form $U_\varepsilon(g_1, \dots, g_n) = \{\xi \in \text{Max } A : |g_i(\xi)| < \varepsilon \text{ for } 1 \leq i \leq n\}$, where $\varepsilon > 0$ and where $g_1, \dots, g_n \in A$ satisfy $g_i(\xi_0) = 0$.

Tate defined Čech cohomology for coverings of affinoid varieties $V = (\text{Max } A, A)$ by finitely many affine subsets, and proved his Acyclicity Theorem [Ta, Theorem 8.2], that $H^i(\mathfrak{V}, \mathcal{O}) = 0$ for $i > 0$; here \mathcal{O} is the presheaf that associates to any affine subset its affinoid algebra, and \mathfrak{V} is a finite covering of V by special affine subsets. (In fact, this holds even with a finite covering of V by affine subsets; see [BGR, §8.2, Theorem 1].) As a consequence, for such a covering \mathfrak{V} of V and any A -module M of finite type, $H^0(\mathfrak{V}, \tilde{M})$ is isomorphic to M , and $H^i(\mathfrak{V}, \tilde{M}) = 0$ for $i > 0$ [Ta, Theorem 8.7]; here \tilde{M} is the presheaf $Y \rightarrow M \otimes_A A_Y$ for Y an affine subset of V . These are analogs of the usual facts for the cohomology of affine varieties. Moreover, they imply that \mathcal{O} and \tilde{M} are sheaves. In particular [BGR, §8.2, Corollary 2], if $f, g \in A$ agree on each member U_i of a finite affine covering of V , then they are equal; and if for every i we are given a function f_i on U_i , with agreements on the overlaps, then they may be “patched” — i.e. there is a function $f \in A$ which restricts to each f_i .

As might be expected, if U is an affine open subset of an affinoid variety V , then the map $A_U \rightarrow \Gamma(U, \mathcal{O})$ is injective. Unfortunately, it is not surjective, e.g. because of characteristic functions like f_D , mentioned at the beginning of this section. Moreover, the functor $A \mapsto \text{Max } A$ is faithful, but not fully faithful [Ta, Corollary 2 to Proposition 9.3]; i.e. not every K -ringed space morphism between two affinoid varieties is induced by a homomorphism between the corresponding rings of functions. Because of this phenomenon, if one defines more global analytic K -spaces simply by considering ringed K -spaces that are locally isomorphic to affinoid varieties, then one instead obtains a theory of “wobbly analytic spaces”, rather than rigid ones.

In order to “rigidify” these wobbly spaces, Tate introduced the second of the two ideas mentioned earlier — viz. shrinking the class of allowable morphisms between such spaces, in such a way that in the case of affinoid varieties, the allowable morphisms are precisely the rigid ones (i.e. those induced by homomorphisms of the underlying algebras). He did this in a series of steps, which he said followed “fully and faithfully a plan furnished by Grothendieck” [Ta, §10]. First, he defined [Ta, Definition 10.1] an *h-structure* θ on a wobbly analytic space V to be a choice of a subset $V^\theta(A) \subset \text{Hom}(\text{Max } A, V)$ (of *structural* maps) for every affinoid K -algebra A , such that every point of V is in the image of some open structural immersion, and such that the composition of a rigid map of affinoids with a structural map is structural. An *h-space* is a wobbly analytic space together with an h-structure, and a *morphism* of h-spaces $(V, \theta) \rightarrow (V', \theta')$ is a ringed space morphism $V \rightarrow V'$ which pulls back structural maps to structural maps. If V, V' are affinoid, then a

morphism of h-spaces between them is the same as a rigid morphism between them [Ta, Corollary to Prop. 10.4].

Next, Tate defined a *special covering* of an h-space [Ta, Def. 10.9] to be one that is obtained by taking a finite covering by special affine subsets, then repeating this process on each of those subsets, a finite number of times. An h-space V is then said to be *special* [Ta, Def. 10.12] if it has the property that a ringed space morphism $\text{Max } B \rightarrow V$ is structural if and only if its restriction to each member of any special covering of $\text{Max } B$ is structural. An open covering of an h-space V is *admissible* if its pullback by any structural morphism has a refinement that is a special covering. A *semi-rigid* analytic space V over K is a special h-space that has an admissible covering by affine open h-spaces. Finally, a *rigid* analytic space is a semi-rigid space V such that the above admissible covering has the property that the intersection of any two members is semi-rigid [Ta, Definition 10.16].

This rather cumbersome approach to rigidifying “wobbly spaces” was simplified and extended in a number of papers in the 1960’s and 1970’s, particularly in [GrRe1], [GrRe2], [GG]. From this point of view, the key idea is that analytic continuation on rigid spaces is permitted only with respect to “admissible” coverings by affinoid varieties, and where the only morphisms permitted between affinoid varieties are the rigid ones (i.e. those induced by homomorphisms between the corresponding affinoid algebras). To make sense of “admissibility”, the notion of Grothendieck topology was used.

Recall (e.g. from [Ar1] or [Mi]) that a Grothendieck topology is a generalization of a classical topology on a space X , in which one replaces the collection of open sets $U \subset X$ by a collection of (admissible) maps $U \rightarrow X$, and in which certain families of such maps $\{V_i \rightarrow U\}_{i \in I}$ are declared to be (*admissible*) *coverings* (of U). This notion was originally introduced in order to provide a framework for the étale topology and for étale cohomology, which for algebraic varieties behaves much like classical singular cohomology in algebraic topology (unlike Zariski Čech cohomology).

In the case of rigid analytic spaces, a less general notion of Grothendieck topology is needed, in which the maps $U \rightarrow X$ are just inclusions of (certain) subsets of X , so that one speaks of “admissible subsets” of X [GuRo, §9.1]. According to the definition of a Grothendieck topology, the admissible subsets U and the admissible coverings of the U ’s satisfy several properties:

- the intersection of two admissible subsets is admissible;
- the singleton $\{U\}$ is an admissible covering of a set U ;
- choosing an admissible covering of each member of an admissible covering together gives an admissible covering; and
- the intersection of an admissible covering of U with an admissible subset $V \subset U$ is an admissible covering of V .

Here, though, several additional conditions are imposed [BGR, p.339]:

- the empty set and X are admissible subsets of X ;
- if V is a subset of an admissible $U \subset X$ and if the restriction to V of every member of some admissible covering of U is an admissible subset of X , then V is an admissible subset of X ; and
- a family of admissible subsets $\{U_i\}_{i \in I}$ whose union is an admissible subset U , and which admits a refinement that is an admissible covering of U , is itself an admissible covering.

In this framework, a *rigid analytic space* is a locally ringed space (V, \mathcal{O}_V) under a Grothendieck topology as above, with respect to which V has an admissible covering $\{V_i\}_{i \in I}$ where each $(V_i, \mathcal{O}_V|_{V_i})$ is an affinoid variety $\mathrm{Sp} A_i = (\mathrm{Max} A_i, A_i)$. (Here $A_i = \mathcal{O}_V|_{V_i}$.) A *morphism* of rigid analytic spaces $(V, \mathcal{O}_V) \rightarrow (W, \mathcal{O}_W)$ is a morphism (f, f^*) as locally ringed spaces. Thus morphisms between affinoid spaces are required to be rigid (i.e. of the form (ϕ°, ϕ) , for some algebra homomorphism ϕ), and global morphisms are locally rigid with respect to an admissible covering. Analogously to the classical and formal cases, a *coherent sheaf* \mathcal{F} (of \mathcal{O}_V -modules) is an \mathcal{O}_V -module that is locally (with respect to an admissible covering) of the form $\mathcal{O}_V^r \rightarrow \mathcal{O}_V^s \rightarrow \mathcal{F} \rightarrow 0$. In the case of an affinoid variety $\mathrm{Sp} A = (\mathrm{Max} A, A)$, coherent sheaves are precisely those of the form \tilde{M} , where M is a finite A -module [FP, III, 6.2].

Rigid analogs of key results in the classical and formal situations (cf. Sections 2.2 and 3.2 above) have been proven in this context. A rigid version of Cartan’s Lemma on matrix factorization [FP, III, 6.3] asserts that if $V = \mathrm{Sp} A$ is an affinoid variety and $f \in A$, and if we let V_1 [resp. V_2] be the set where $|f| \leq 1$ [resp. $|f| \geq 1$], then every invertible matrix in $\mathrm{GL}_n(\mathcal{O}(V_1 \cap V_2))$ that is sufficiently close to the identity can be factored as the product of invertible matrices over $\mathcal{O}(V_1)$ and $\mathcal{O}(V_2)$. There are also rigid analogs of Cartan’s Theorems A and B, proven by Kiehl [Ki2]; they assert that a coherent sheaf \mathcal{F} is generated by its global sections, and that $H^i(V, \mathcal{F}) = 0$ for $i > 0$, for “quasi-Stein” rigid analytic spaces V . (These are rigid spaces V that can be written as an increasing union of affinoid open subsets U_i that form an admissible covering of V , and such that $\mathcal{O}(U_{i+1})$ is dense in $\mathcal{O}(U_i)$. Compare Cartan’s original version for complex Stein spaces [Ca2] discussed in §2.2 above.) Kiehl also proved [Ki1] a rigid analog of Zariski’s Theorem on Formal Functions [Hrt2, III, Thm. 11.1], which together with Cartan’s Theorem B (or Theorem A) was used to obtain GAGA classically. And indeed, there is a rigid analog of GAGA (or in this case, a “GRGA”: géométrie rigide et géométrie algébrique) [Köp], asserting the equivalence between coherent rigid sheaves and coherent algebraic sheaves of modules over a projective algebraic K -variety. Thus, to give a coherent sheaf over such a variety, it suffices to give it over the members of an admissible covering (viewing the variety as a rigid analytic space), and giving the patching data on the overlaps.

As in Sections 2 and 3 above, it would be desirable to use these results in order to

obtain a version of Riemann’s Existence Theorem, which would classify covers. Ideally, this should be precise enough to give an explicit description of the tower of Galois groups of covers of a given space; and that description should be analogous to Corollary 2.1.2, the explicit form of the classical Riemann’s Existence Theorem given at the beginning of Section 2.1. Unfortunately, to give such an explicit description, one needs to have a notion of a “topological fundamental group”, and one needs to be able to compute that group explicitly. But unlike the complex case, one does not have such a notion, or computation, over more general fields K (in particular, because we cannot speak of “loops”). Thus, in this context, one does not have a full analog of Riemann’s Existence Theorem 2.1.1, because one cannot assert an equivalence between finite rigid analytic covering maps and finite topological covering spaces. Still, one can ask for an analog of the first part of Theorem, 2.1.1 viz. an equivalence between finite étale covers of an algebraic curve V over K , and finite analytic covering maps of V (viewed as a rigid analytic space).

Such a result has been obtained (with some restrictions) by Lütkebohmert [Lü2]. As in the proof of the complex version (see Section 2.2), the proof proceeds using GAGA (here, the rigid version discussed above). Namely, as in the complex case, once one has the equivalence of categories that GAGA provides for sheaves of modules, one also obtains an equivalence (as a purely formal consequence) for sheaves of algebras, and hence for branched covers. But as in the complex case, GAGA applies to projective curves, but not to affine curves. So GAGA shows that there is an equivalence between branched (algebraic) covers of a projective K -curve X , and rigid analytic branched covers of the curve. Then to prove the desired portion of Riemann’s Existence Theorem, it remains to show (both in the algebraic and rigid analytic settings) that covers of X branched only at a finite set B are equivalent to unramified covers of $V = X - B$ (i.e. that every unramified cover of V extends uniquely to a branched cover of X). In Section 2.2, we saw that this is immediate in the algebraic context, and follows easily from complex analysis in the analytic setting. But in the rigid analytic setting, this extension result for rigid analytic covers is harder, and moreover requires that the characteristic of K is 0.

Specifically, if $\text{char } K = 0$, then unramified rigid covers of an affine K -curve $V = X - B$ do extend (uniquely) to rigid branched covers of the projective curve X ; and so finite étale covers of V are equivalent to finite unramified rigid analytic covers of V . Moreover this generalizes to higher dimensions, where V is any K -scheme that is locally of finite type over K [Lü2, Theorem 3.1]. But there are counterexamples, even for curves, if $\text{char } K = p$. For example, let $K = k((t))$, let V be the affine x -line over K , and consider the rigid unramified covering map $W \rightarrow V$ given by $y^p - y = \sum_{i=1}^{\infty} t^{(p+1)^i} x^{p^i}$. Then this map does not extend to a finite (branched) cover of the projective line, and so is not induced by any algebraic cover of V [Lü2, Example 2.10]. On the other hand, if one restricts attention to *tamely* ramified covers, then the desired equivalence between rigid and algebraic unramified

covers does hold [Lü2, Theorem 4.1]. (Note that the above wildly ramified example does not contradict rigid GAGA, since that result applies in the projective case, whereas this example is affine.)

Still, we do not have an explicit description of the rigid analytic covers of a given curve (even apart from the difficulty with wildly ramified covers); so this result does not give explicit information about Galois groups and fundamental groups for K -curves (as a full rigid analog of Corollary 2.1.2 would). We return to this issue in Section 4.3, after considering another approach to rigid analytic spaces in Section 4.2.

Section 4.2. Rigid geometry via formal geometry.

Tate’s rigid analytic spaces can be reinterpreted in terms of Grothendieck’s formal schemes. This reinterpretation was outlined by Raynaud in [Ra1], and worked out in greater detail by Bosch, Lütkebohmert, and Raynaud in [Lü1], [BLü1], [BLü2], [BLüR1], [BLüR2]. (See also [Ra2, §3]; and Chapters 1 and 2, by M. Garuti [Ga] and Y. Henrio [He], in [BLoR].) As Tate said in [Ta], his approach was motivated by a suggestion of Grothendieck; and according to the introduction to [BLü1], Grothendieck’s goal was to associate a generic fibre to a formal scheme of finite type. So this approach may actually be closer to Grothendieck’s original intent than the more analytic framework discussed above.

The basic idea of this approach can be seen by revisiting examples from Sections 3.2 and 4.1. In Example 3.2.3, it was seen that $k[x][[t]]$ is the ring of formal functions along the affine x -line in the x, t -plane over a field k , or equivalently that its spectrum is a formal thickening of the affine x -line. The corresponding ring for the affine x^{-1} -line (i.e. the formal thickening of the projective x -line minus $x = 0$) is $k[x^{-1}][[t]]$, and the ring corresponding to the overlap (i.e. the formal thickening of $\mathbb{P}^1 - \{0, \infty\}$) is $k[x, x^{-1}][[t]]$. On the other hand, as seen in Section 4.1, if t is inverted in each of these three rings, one obtains the rings of functions on three affinoids over $K = k((t))$: the disc $|x| \leq 1$; the disc $|x^{-1}| \leq 1$ (i.e. $|x| \geq 1$ together with the point at infinity); and the “annulus” $|x| = 1$. In each of these two contexts (formal and rigid), the first two sets cover the projective line (over $R := k[[t]]$ and $K = \text{frac } R$, respectively), and the third set is their “overlap”. The ring of holomorphic functions on an affinoid set over K can (at least in this example) be viewed as the localization, with respect to t , of the ring of formal functions on an affine open subset of the closed fibre on an R -scheme. Correspondingly, an affinoid can be viewed as the generic fibre of the spectrum of the ring of formal functions (in the above example, a curve being the general fibre of a surface). Intuitively, then, a rigid analytic space over K is the general fibre of a (formal) scheme over R . (See Figure 4.2.1.)

The actual correspondence between formal schemes and rigid analytic spaces is a bit more complicated, because of several issues. The first concerns which base rings and fields are involved. Formal schemes are defined over complete local rings R , while rigid analytic




Figure 4.2.1: A rigid covering of \mathbb{P}_K^1 (viewed as a sphere, in analogy with the complex case). The patches U_1, U_2 are discs around 0 and ∞ , with rings of functions $k[x][[t]][1/t]$ and $k[1/x][[t]][1/t]$ (see §4.1). The overlap U_0 is an annulus containing the point $x = 1$, with ring of functions $k[x, 1/x][[t]][1/t]$. Compare Fig. 3.1.4 and see Example 4.2.3 below.

spaces are defined over complete valuation fields K . The fraction field of a complete discrete valuation ring R is a discrete valuation field K , and every such K arises from such an R . But general valuation fields are not fraction fields of complete local rings, and the fraction fields of general complete local rings are not valuation fields. So in stating the correspondence, we restrict here to the case of a complete discrete valuation ring R , say with maximal ideal \mathfrak{m} (though one can consider, somewhat more generally, a complete height 1 valuation ring R).

Secondly, in order for a formal space to induce a rigid space, it must locally induce affinoid K -algebras, i.e. K -algebras that are of topologically finite type. Correspondingly, we say that an R -algebra A is of *topologically finite type* if it is a quotient of the \mathfrak{m} -adic completion of some $R[x_1, \dots, x_n]$. Observe that this \mathfrak{m} -adic completion is a subring of $R[[x_1, \dots, x_n]]$, and in fact consists precisely of those power series $\sum_{\underline{i} \in \mathbb{N}^n} a_{\underline{i}} x^{\underline{i}}$, where $a_{\underline{i}} \rightarrow 0$ as $|\underline{i}| \rightarrow \infty$. It is then easy to verify that $A \otimes_R K$ is an affinoid K -algebra, for any R -algebra A that is of topologically finite type. (This is in contrast to the full rings of power series, where $K[[x_1, \dots, x_n]]$ is much larger than $R[[x_1, \dots, x_n]] \otimes_R K$.) A formal R -scheme \mathcal{V} is *locally of topologically finite type* if in a neighborhood of every point, the structure sheaf $\mathcal{O}_{\mathcal{V}}$ is given by an R -algebra that is of topologically finite type. Such a formal scheme is said to be of *topologically finite type* if in addition it is quasi-compact. Thus formal schemes that are of topologically finite type induce quasi-compact rigid spaces.

The condition of a formal R -scheme \mathcal{V} being locally of topologically finite type in turn implies that the corresponding R/\mathfrak{m}^n -schemes V_n are locally Noetherian (since the structure sheaf is locally a quotient of some $(R/\mathfrak{m}^n)[x_1, \dots, x_n]$). Thus each V_n is quasi-separated, by [Gr4, IV, Cor. 1.2.8]; and hence so is \mathcal{V} and so is the induced rigid space. On the other hand, not every rigid space is necessarily quasi-separated; so in order to get

an equivalence between formal and rigid spaces, we will need to restrict attention to rigid spaces that are quasi-separated (this being a very mild finiteness condition).

A third issue concerns the fact that different R -schemes that are birationally isomorphic can have the same general fibre. For example, let V be a proper R -scheme, where R is a complete discrete valuation ring. Let V_0 be the closed fibre of V , and let W be a closed subset of V_0 . Let \tilde{V} be the blow up of V along W (as a scheme). Then V and \tilde{V} have the same general fibre. But they are not isomorphic as schemes (if the codimension of W in V is at least 2), since \tilde{V} has an exceptional divisor over the blown up points. Hence they do not correspond to isomorphic formal schemes.

In order to deal with this third issue, the strategy is to regard two R -schemes as equivalent if they have a common *admissible* blow-up (i.e. a blow up at a closed subset of the closed fibre). Thus given two R -schemes V, V' , to give a morphism from the equivalence class of V to that of V' is to give an admissible blow up $\tilde{V} \rightarrow V$ together with a morphism of R -schemes $\tilde{V} \rightarrow V'$. Here V, \tilde{V}, V' induce formal R -schemes $\mathcal{V}, \tilde{\mathcal{V}}, \mathcal{V}'$ (given by the direct limit of the fibres V_n, \tilde{V}_n, V'_n over \mathfrak{m}^n), and we regard the induced pair $(\tilde{\mathcal{V}} \rightarrow \mathcal{V}, \tilde{\mathcal{V}} \rightarrow \mathcal{V}')$ as a morphism between the equivalence classes of $\mathcal{V}, \mathcal{V}'$. Equivalently, we are considering morphisms from the class of \mathcal{V} to the class of \mathcal{V}' , in the localization of the category of formal R -schemes with respect to the class of admissible formal blow-ups $\tilde{\mathcal{V}} \rightarrow \mathcal{V}$. (The *localization* is the category in which those blow-ups are formally inverted. Such a localization automatically exists, according to [Hrt1]; though to be set-theoretically precise, one may wish to work within a larger “universe” [We, Remark 10.3.3].)

Here, for a formal scheme \mathcal{V} induced by a proper R -scheme V , one can correspondingly define admissible blow-ups of \mathcal{V} as the morphisms of formal schemes induced by admissible blow-ups of V . Alternatively, and for a more general formal R -scheme \mathcal{V} , admissible blow-ups can be defined directly, despite the fact that the topological space underlying \mathcal{V} is just the closed fibre of the associated R -scheme (if there is one). Namely, the blow-up can be defined algebraically, analogously to the usual definition for schemes. First, observe that if A is a complete R -algebra, then the closed subsets of the closed fibre of $\text{Spec } A$ correspond to ideals of A that are *open* in the topology induced by that of R . Now recall [Hrt2, Chap. II, p.163] that if V is a Noetherian scheme, and \mathcal{I} is a coherent sheaf of ideals on V , then the blow-up of V at \mathcal{I} is $\text{Proj } \mathcal{J}$, where \mathcal{J} is the sheaf of graded algebras $\mathcal{J} = \bigoplus_{d \geq 0} \mathcal{I}^d$. So given a formal R -scheme \mathcal{V} and a sheaf \mathcal{I} of open ideals of $\mathcal{O}_{\mathcal{V}}$, define the blow-up of \mathcal{V} along \mathcal{I} to be the formal scheme associated to the direct system of R/\mathfrak{m}^n -schemes $\text{Proj } \mathcal{J}_n$, where $\mathcal{J}_n = \bigoplus_{d \geq 0} (\mathcal{I}^d \otimes_{\mathcal{O}_{\mathcal{V}}} \mathcal{O}_{\mathcal{V}}/\mathfrak{m}^n)$. We call such a blow-up of the formal scheme \mathcal{V} *admissible*. This agrees with the previous definition, for formal schemes \mathcal{V} induced by R -schemes V .

A fourth issue, which is similar to the third, is that an R -scheme V may have an irreducible component that is contained in the closed fibre V_0 . In that case, the general

fibre of V “does not see” that component, and so cannot determine V (or the induced formal scheme). So we avoid this case, by requiring that the formal scheme \mathcal{V} have the property that its structure sheaf $\mathcal{O}_{\mathcal{V}}$ has no \mathfrak{m} -torsion. We call the formal scheme \mathcal{V} *admissible* if it has this property and is of locally of topologically finite type. (So quasi-compact admissible is the same as \mathfrak{m} -torsion-free plus topologically finite type.)

With these restrictions and adjustments, the equivalence between formal and rigid spaces takes place. Consider an admissible formal R -scheme \mathcal{V} , whose underlying topological space is a k -scheme V_0 (where $k = R/\mathfrak{m}$). For any affine open subset $U \subset V_0$, let A be the ring of formal functions along U . So A is topologically of finite type, and has no \mathfrak{m} -torsion; and $A \otimes_R K$ is an affinoid K -algebra. In the notation of Section 4.1, $\mathrm{Sp} A = (\mathrm{Max} A, A)$ is an affinoid variety. This construction is compatible with shrinking U , and so from \mathcal{V} we obtain a rigid analytic space, which we denote by $\mathcal{V}^{\mathrm{rig}}$. There is then the following key theorem of Raynaud [Ra1] (see also [BLü1, Theorem 4.1], for details):

Theorem 4.2.2. (Raynaud) *Let R be a complete valuation ring of height 1 with fraction field K . Let For_R be the category of quasi-compact admissible formal R -schemes, and let For'_R be the localization of For_R with respect to admissible formal blow-ups. Let Rig_K be the category of quasi-compact quasi-separated rigid analytic K -spaces. Then the functor $\mathrm{rig} : \mathrm{For}_R \rightarrow \mathrm{Rig}_K$ given by $\mathcal{V} \mapsto \mathcal{V}^{\mathrm{rig}}$ induces an equivalence of categories $\mathrm{For}'_R \rightarrow \mathrm{Rig}_K$.*

(Alternatively, the conclusion of the theorem could be stated by saying that $\mathrm{rig} : \mathrm{For}_R \rightarrow \mathrm{Rig}_K$ is a localizing functor with respect to all admissible blow-ups, rather than speaking in terms of For'_R .)

In particular, if V is a proper R -scheme, and if \mathcal{V} is the associated formal scheme, then $\mathcal{V}^{\mathrm{rig}}$ is the rigid analytic space corresponding to the generic fibre V_K of V .

More generally, one can turn the above result around and make it a *definition*, to make sense of rigid analytic spaces over the fraction field K of a Noetherian complete local ring R which is not necessarily a valuation ring (e.g. $k[[x_1, \dots, x_n]]$, where k is a field and $n > 1$). That is, for such a ring R and fraction field K , one can simply *define* the category Rig_K of rigid analytic K -spaces to be the category For'_R , obtained by localizing the category For_R of formal R -schemes with respect to admissible blow-ups [Ra1], [BLü1], [Ga]. The point is that formal schemes make sense in this context, and thus the notion of rigid spaces can be extended to this situation as well. (Of course, by Raynaud’s theorem, the two definitions are equivalent in the case that R is a complete discrete valuation ring.)

The advantage to Raynaud’s approach to rigid analytic spaces is it permits them to be studied using Grothendieck’s results on formal schemes in EGA [Gr4]. It also permits the use of results in EGA on proper schemes over complete local rings, because of the equivalence of those schemes with formal schemes via by Grothendieck’s Existence Theorem ([Gr2], [Gr4, III, Cor. 5.1.6]; see also Section 3.2 above). In particular, Grothendieck’s Existence Theorem and Raynaud’s theorem above together imply the rigid GAGA result

(for projective spaces) discussed in Section 4.1 above. Moreover, Raynaud’s approach permits the use of the rigid point of view over more general fields than Tate’s original approach did, though with some loss of analytic flavor. Indeed, from this point of view, the rigid and formal contexts are not so different, though there is a difference in terms of intuition. Another difference is that in the formal context one works on a fixed R -model of a space, whereas in the rigid context one works just over K (and thus blow-ups are already included in the geometry).

We conclude this discussion by giving two examples comparing formal and rigid GAGA on the line, beginning with the motivating situation discussed earlier:

Example 4.2.3. Let k be a field; $R = k[[t]]$; $K = k((t))$; and $V = \mathbb{P}_R^1$. Let x be a parameter on V , and $y = x^{-1}$. So V is covered by two copies of the affine line over R , the x -line and the y -line, intersecting where $x, y \neq 0$. Letting \mathcal{V} be the formal scheme associated to V , there is the induced rigid analytic space $V^{\text{rig}} := \mathcal{V}^{\text{rig}}$, viz. \mathbb{P}_K^1 . According to rigid GAGA, giving a coherent sheaf on V^{rig} is equivalent to giving finite modules over (the rings of functions on) the admissible sets $U_1 : |x| \leq 1$ and $U_2 : |y| \leq 1$, with agreement on the overlap $U_0 : |x| = |y| = 1$. Here $U_1 = \text{Sp } K\{x\}$, $U_2 = \text{Sp } K\{y\}$, and $U_0 = \text{Sp } K\{x, y\}/(xy - 1)$. Geometrically (and intuitively), U_1 and U_2 are discs centered around $x = 0, \infty$ respectively (the “south and north poles”), and U_0 is an annulus (a band around the “equator”, if \mathbb{P}_K^1 is viewed as a “sphere”; see Figure 4.2.1 above).

On the formal level, U_1 is the general fibre of $S_1 = \text{Spec } k[x][[t]]$, the formal thickening of the affine x -line (which pinches down near $x = \infty$). Similarly, U_2 is the general fibre of $S_2 = \text{Spec } k[y][[t]]$, the formal thickening of the affine y -line (which pinches down near $x = 0$). And U_0 is the general fibre of $S_0 = \text{Spec } k[x, x^{-1}][[t]]$, the formal thickening of the line with both 0 and ∞ deleted (and which pinches down near both points — cf. Figure 3.1.4). According to formal GAGA (i.e. Grothendieck’s Existence Theorem; cf. Theorems 3.2.1 and 3.2.4), giving a coherent sheaf on V is equivalent to giving finite modules over S_1 and S_2 with agreement on the “overlap” S_0 .

In the formal context, even less data is needed in order to construct a coherent sheaf on V — and this permits more general constructions to be performed (e.g. see [Ha6]). Namely, let $\hat{S}_1 = \text{Spec } k[[x, t]]$, the complete local neighborhood of $x = t = 0$. Let $\hat{S}_0 = \text{Spec } k((x))[[t]]$, the “overlap” of \hat{S}_1 with S_2 . (See Figure 3.2.9, where $\hat{S}_1, S_2, \hat{S}_0$ are denoted by W^*, U^*, W'^* , respectively.) Then according to Theorem 3.2.8, giving a coherent sheaf on V is equivalent to giving finite modules over \hat{S}_1 and S_2 together with agreement over \hat{S}_0 . On the rigid level, the generic fibres of \hat{S}_1 and S_2 are $\hat{U}_1 : |x| < 1$ and $U_2 : |x| \geq 1$. Those subsets of V^{rig} do not intersect, and moreover \hat{U}_1 is not an affinoid set. The result in the formal situation suggests that the generic fibre of \hat{S}_0 , corresponding to $k((x))((t))$, forms a “glue” that connects \hat{U}_1 and U_2 ; but this cannot be formulated within the rigid framework. \square

Example 4.2.4. With notation as in Example 4.2.3, rigid GAGA says that to give a coherent sheaf on $V^{\text{rig}} = \mathbb{P}_K^1$ is equivalent to giving finite modules over the two discs $|x| \leq 1$ and $|y| \leq c^{-1}$, and over the annulus $c \leq |x| \leq 1$; here $0 < c = |t| < 1$, and the annulus is the overlap of the two discs. Writing $z = ty = t/x$, the rings of functions on these three sets are $K\{x\}$, $K\{z\}$, and $K\{x, z\}/(xz - t)$.

To consider the corresponding formal situation, let \tilde{V} be the blow-up of V at the closed point $x = t = 0$. Writing $xz = t$, the closed fibre of \tilde{V} consists of the projective x -line over k (the proper transform of the closed fibre of V) and the projective z -line over k (the exceptional divisor), meeting at the “origin” $O : x = z = t = 0$. The three affinoid open sets above are then the generic fibres associated to the formal schemes obtained by respectively deleting from the closed fibre of \tilde{V} the point $x = \infty$ (which is where $z = 0$); the point $z = \infty$ (where $x = 0$); and both of these points. And by Grothendieck’s Existence Theorem, giving compatible formal coherent modules over each of these sets is equivalent to giving a coherent module over V .

But as in Example 4.2.3, less is needed in the formal context. Namely, let X' and Z' be the projective x - and z -lines over k , with the points $(x = 0)$ and $(z = 0)$ respectively deleted. Consider the rings of formal functions along X' and Z' , viz. $k[x^{-1}][[t]]$ and $k[z^{-1}][[t]]$ respectively, and their spectra T_1, T_2 . Consider also the spectrum T_3 of $k[[x, z, t]]/(xz - t)$, the complete local ring of V at O . Here T_1 and T_2 are disjoint, while the “overlap” of T_1 and T_3 [resp. of T_2 and T_3] is the spectrum $T_{1,3}$ of $k((x))[[t]]$ [resp. $T_{2,3}$ of $k((z))[[t]]$]. By Theorem 3.2.8, giving finite modules over T_1, T_2, T_3 that agree on the two “overlaps” is equivalent to giving a coherent module over V . The generic fibres of T_1 and T_2 are the sets $|y| \leq 1$ and $|x| \leq c$, and that of T_3 is $c < |x| < 1$. These three sets are disjoint, though the formal set-up provides “glue” (in the form of $T_{1,3}$ and $T_{2,3}$) connecting T_3 to T_1 and to T_2 . This is a special case of Example 3.2.11. (Alternatively, one could use Theorem 3.2.12, taking V to be the projective x -line over $k[[t]]$, taking \tilde{V} to be the blow-up of V at $x = t = 0$, and identifying the exceptional divisor with the projective z -line over k . See Example 3.2.13.) \square

More generally, in the rigid set-up, one can consider the annulus $c^n \leq |x| \leq 1$ in \mathbb{P}_K^1 (where $K = k((t))$ and $c = |t|$ as above, and where n is a positive integer). If one writes $u = t^n/x$, then this is the intersection of the two admissible sets $|x| \leq 1$ and $|u| \leq 1$. This annulus is said to have *thickness* (or *épaisseur*) equal to n . The corresponding situation in the formal framework can be arrived at by taking the projective x -line V over $R = k[[t]]$; blowing this up at the point $x = t = 0$ (obtaining a parameter $z = t/x$ on the exceptional divisor E); blowing that up at the point $z = \infty$ on E (thereby obtaining a parameter $z' = t/z^{-1} = t^2/x$ on the new exceptional divisor); and repeating the process for a total of n blow-ups. The analogs of Examples 4.2.3 and 4.2.4 above can then be considered similarly.

Section 4.3. Rigid patching and constructing covers.

Rigid geometry, like formal geometry, provides a framework within which patching constructions can be carried out in order to construct covers of curves, and thereby obtain Galois groups over curves. Ideally, one would like to obtain a version of Riemann's Existence Theorem analogous to that stated for complex curves in Section 2.1. But while a kind of "Riemann's Existence Theorem" for rigid spaces was obtained by Lütkebohmert [Lü2] (see Section 4.1 above), that result does not say which Galois groups arise, due to a lack of topological information. Still, as in the formal case, one can show by a patching construction that every finite group is a Galois group of a branched cover with enough branch points, and show a "Half Riemann Existence Theorem" that is analogous to the classification of slit covers of complex curves (see Section 2.3).

Namely, Serre observed in a 1990 talk in Bordeaux that there should be a rigid proof of Theorem 3.3.1 above (on the realizability of every finite group as a Galois group over the fraction field K of a complete local domain R [Ha4]), when the base ring R is complete with respect to a non-archimedean absolute value. Given the connection between rigid and formal schemes discussed in Section 4.2 (especially in the case of complete discrete valuation rings), this would seem quite plausible. Shortly afterwards, in [Se7, §8.4.4], Serre outlined such a proof in the case that $K = \mathbb{Q}_p$. A more detailed argument was carried out by Liu for complete non-archimedean fields with an absolute value, in a manuscript that was written in 1992 and that appeared later in [Li], after circulating privately for a few years. (Concerning complete *archimedean* fields, the complex case was discussed in Section 2 above, and the real case is handled in [Se7, §8.4.3], via the complex case; cf. also [DF] for the real case.)

The rigid version of Theorem 3.3.1 is as follows:

Theorem 4.3.1. *Let K be a field that is complete with respect to a non-trivial non-archimedean absolute value. Let G be a finite group. Then there is a G -Galois irreducible branched cover $Y \rightarrow \mathbb{P}_K^1$ such that the fibre over some K -point of \mathbb{P}_K^1 is totally split.*

Here the totally split condition is that the fibre consists of unramified K -points. This property (which takes the place of the mock cover hypothesis of Theorem 3.3.1) forces the cover to be regular, in the sense that K is algebraically closed in the function field $K(Y)$ of Y . (Namely, if L is the algebraic closure of K in $K(Y)$, then L is contained in the integral closure in $K(Y)$ of the local ring \mathcal{O}_ξ of any closed point $\xi \in \mathbb{P}_K^1$; and so it is contained in the residue field of each closed point of Y .) Thus Theorem 3.3.1 is recaptured, for such fields K .

Proof sketch of Theorem 4.3.1. The proof proceeds analogously to that of Theorem 3.3.1. Namely, first one proves the result explicitly in the special case that the group is a cyclic group. In [Se7, §8.4.4], Serre does this by using an argument involving tori [Se7, §4.2] to

show that cyclic groups are Galois groups of branched covers of the line; one can then obtain a totally split fibre by twisting, e.g. as in [HV, Lemma 4.2(a)]. Or (as in [Li]) one can proceed as in the original proof for the cyclic case in the formal setting [Ha4, Lemma 2.1], which used ideas of Saltman [Slt]; cf. Proposition 3.3.3 above.

To prove the theorem in the general case, cyclic covers are patched together to produce a cover with the desired group, in a rigid analog of the proof of Theorem 3.3.1. Namely, let g_1, \dots, g_r be generators of G . For each i , let H_i be the cyclic subgroup of G generated by g_i , and let $f_i : Y_i \rightarrow \mathbb{P}_K^1$ be a H_i -Galois cover that is totally split over a point ξ_i . By the Implicit Function Theorem over complete fields, for each i there is a closed disc \bar{D}_i about ξ_i such that the inverse image $f_i^{-1}(\bar{D}_i)$ is a disjoint union of copies of \bar{D}_i . Let D_i be the corresponding open disc about ξ_i , let $\bar{U}_i = \mathbb{P}_K^1 - D_i$, and let $U_i = \mathbb{P}_K^1 - \bar{D}_i$. After a change of variables, we may assume that the \bar{U}_i 's are pairwise disjoint affinoid sets. For each i , let $\bar{V}_i \rightarrow \bar{U}_i$ be the pullback of f_i to \bar{U}_i . Then \bar{V}_i is an H_i -Galois cover whose restriction over $\bar{U}_i - U_i = \bar{D}_i - D_i$ is trivial. Inducing from H_i to G (by taking a disjoint union of copies, indexed by the cosets of H_i), we obtain a corresponding G -Galois disconnected cover $\bar{W}_i = \text{Ind}_{H_i}^G \bar{V}_i \rightarrow \bar{U}_i$. Also, let $\bar{U}_0 = \mathbb{P}_K^1 - \bigcup_{j=1}^r U_j = \bigcap_{j=1}^r \bar{D}_j$, and let $\bar{W}_0 \rightarrow \bar{U}_0$ be the trivial G -cover $\text{Ind}_1^G \bar{U}_0$. We now apply rigid GAGA (see Sections 4.1 and 4.2), though for covers rather than for modules (that form following automatically, as in Theorem 3.2.4, via the General Principle 2.2.4). Namely, we patch together the covers $\bar{W}_i \rightarrow \bar{U}_i$ ($i = 0, \dots, r$) along the overlaps $\bar{U}_i \cap \bar{U}_0 = \bar{U}_i - U_i$ ($i = 1, \dots, r$), where they are trivial. One then checks that the resulting G -Galois cover is as desired (and in particular is irreducible, because the g_i 's generate G); and this yields the theorem. \square

As in Section 3.3, Theorem 4.3.1 extends to the class of large fields, such as the algebraic p -adics and the field of totally real algebraic numbers. Namely, as in the passage to Theorem 3.3.6, there is the following result of Pop:

Corollary 4.3.2. [Po4] *If k is a large field, then every finite group is the Galois group of a Galois field extension of $k(x)$. Moreover this extension may be chosen to be regular, and with a totally split fibre.*

Namely, by Theorem 4.3.1, there is a G -Galois extension of $k((t))(x)$, and this descends to a regular G -Galois extension of the fraction field of $A[x]$ with a totally split fibre over $x = 0$, for some $k(t)$ -subalgebra $A \subset k((t))$ of finite type. By the Bertini-Noether Theorem [FJ, Prop. 9.29], we may assume that every specialization of A to a k -point gives a G -Galois regular field extension of $k(x)$; and such a specialization exists on $V := \text{Spec } A$ since k is large and since V contains a $k((t))$ -point.

The construction in the proof of Theorem 4.3.1, like the one used in proving the corresponding result using formal geometry, can be regarded as analogous to the slit cover construction of complex covers described in Section 2.3 (and see the discussion at the end

of Section 3.3 for the analogy with the formal setting). In fact, rather than considering covers (and Galois groups) one at a time, a whole tower of covers (and Galois groups) can be considered, as in the “analytic half Riemann Existence Theorem” 2.3.5. In the present setting (unlike the situation over \mathbb{C}), the absolute Galois group G_K of the valued field K comes into play, since it acts on the geometric fundamental group (i.e. the fundamental group of the punctured line after base-change to the separable closure K^s of K). This construction of a tower of compatible covers has been carried out by Pop in [Po2] (where the term “half Riemann Existence Theorem” was also introduced). Also, rather than requiring K to be complete, Pop required K merely to be henselian (and cf. Example 3.3.2(d), for comments about deducing the henselian case of that result from the complete case via Artin’s Approximation Theorem).

In Pop’s result, as in the case of complex slit covers, one chooses as a branch locus a closed subset $S \subset \mathbb{P}_K^1$ whose base change to K^s consists of finitely many pairs of nearby points. That is, S is a disjoint union of two closed subsets $S = S_1 \cup S_2$ of \mathbb{P}_K^1 such that $S_1^s := S_1 \times_K K^s = \{\xi_1, \dots, \xi_r\}$ and $S_2^s := S_2 \times_K K^s = \{\eta_1, \dots, \eta_r\}$, where the ξ_i and η_j are distinct K^s -points, and where each ξ_i is closer to the corresponding η_i than it is to any other ξ_j . Such a set S is called *pairwise adjusted*. Note that the sets S_1^s and S_2^s are each G_K -invariant, and that G_K acts on the sets S_1^s and S_2^s compatibly (i.e. if $\alpha \in G_K$ satisfies $\alpha(\xi_i) = \xi_j$, then $\alpha(\eta_i) = \eta_j$). Now let $U = \mathbb{P}_K^1 - S$ and $U^s = U \times_K K^s = \mathbb{P}_{K^s}^1 - S^s$, and recall the fundamental exact sequence

$$1 \rightarrow \pi_1(U^s) \rightarrow \pi_1(U) \rightarrow G_K \rightarrow 1. \quad (*)$$

Also, let Π be the free profinite group \hat{F}_r of rank r if the valued field K is in the equal characteristic case, and let it be the free pro-prime-to- p group \hat{F}'_r of rank r if K is in the unequal characteristic case with residue characteristic $p > 0$. Define an action of G_K on Π by letting $\alpha \in G_K$ take the j th generator g_j of Π to $g_i^{\chi(\alpha^{-1})}$; here i is the unique index such that $\alpha(\xi_i) = \xi_j$, and $\chi : G_K \rightarrow \hat{\mathbb{Z}}^*$ is the cyclotomic character (taking $\gamma \mapsto m$ if $\gamma(\zeta) = \zeta^m$ for all roots of unity ζ).

Theorem 4.3.3. (Half Riemann Existence Theorem with Galois action [Po2]) *Let K be a henselian valued field of rank 1, let $S \subset \mathbb{P}_K^1$ be a pairwise adjusted subset of degree $2r$ as above, and let $U = \mathbb{P}_K^1 - S$. Then the fundamental exact sequence $(*)$ has a quotient*

$$1 \rightarrow \Pi \rightarrow \Pi \rtimes G_K \rightarrow G_K \rightarrow 1, \quad (**)$$

where Π is defined as above and where the semi-direct product is taken with respect to the above action of G_K on Π .

Proof sketch of Theorem 4.3.3. In the case that the field K is complete, the proof of Theorem 4.3.3 follows a strategy that is similar to that of Theorem 4.3.1. As in Theorem 4.3.1

(and Theorem 3.3.1), the proof relies on the construction of cyclic covers that are trivial outside a small neighborhood (in an appropriate sense), and which can then be patched. The key new ingredient is that one must show that the construction is functorial, and in particular is compatible with forming towers. Concerning this last point, after passing to the maximal cyclotomic extension K^{cycl} of K , one can construct a tower of regular covers by patching together local cyclic covers that are Kummer or Artin-Schreier. These can be constructed compatibly with respect to the action of $\text{Gal}(K^{\text{cycl}}/K)$, since S is pairwise adjusted; and the resulting tower, viewed as a tower of covers of U , has the desired properties.

The henselian case is then deduced from the complete case. This is done by first observing that the absolute Galois groups of K and of its completion \hat{K} are canonically isomorphic (because K is henselian). Then, writing \hat{K}^s for the separable closure of \hat{K} , it is checked that every finite branched cover of the \hat{K}^s -line that results from the patching construction is defined over the separable closure K^s of K . (Namely, consider a finite quotient Q of Π , generated by cyclic subgroups C_i . The patching construction over \hat{K} yields a Q -Galois cover $Y \rightarrow \mathbb{P}^1_{\hat{K}^s}$ that is constructed using cyclic building blocks $Z_i \xrightarrow{C_i} \mathbb{P}^1$ which are each defined over K^s . Let $Z \rightarrow \mathbb{P}^1$ be the fibre product of the Z_i 's; this is Galois with group $H = \prod C_i$. Pulling back the Q -cover $Y \rightarrow \mathbb{P}^1_{\hat{K}^s}$ via $Z_{\hat{K}^s} \xrightarrow{H} \mathbb{P}^1_{\hat{K}^s}$ gives an unramified Q -cover Y' of the projective curve $Z_{\hat{K}^s}$; here Y' is also a $Q \times H$ -Galois branched cover of $\mathbb{P}^1_{\hat{K}^s}$. By Grothendieck's specialization isomorphism [Gr5, XIII], Y' descends to a Q -cover of Z_{K^s} whose composition with $Z_{K^s} \rightarrow \mathbb{P}^1_{K^s}$ is $Q \times H$ -Galois. Hence Y descends to a Q -cover of $\mathbb{P}^1_{K^s}$.) Since the Galois actions of G_K and $G_{\hat{K}}$ are the same, the result in the general henselian case follows. \square

Remark 4.3.4. (a) The hypotheses of Theorem 4.3.3 are easily satisfied; i.e. there are many choices of pairwise adjusted subsets. Namely, let $f \in K[x]$ be any irreducible separable monic polynomial, and let $g \in K[x]$ be chosen so that it is monic of the same degree, and so that its coefficients are sufficiently close to those of f . Then the zero locus of fg in \mathbb{A}^1_K is a pairwise adjusted subset, by continuity of the roots [La, II, §2, Proposition 4]. Repeating this construction with finitely many polynomials f_i and then taking the union of the resulting sets gives a general pairwise adjusted subset. Note that in the case that K is separably closed, the construction is particularly simple: One may take an arbitrary set $S_1 = \{\xi_1, \dots, \xi_r\}$ of K -points in \mathbb{A}^1_K , and any set $S_2 = \{\eta_1, \dots, \eta_r\}$ of K -points such that each η_i is sufficiently close to ξ_i . This recovers the slit cover construction of Section 2.3 in the case $K = \mathbb{C}$.

(b) In the equal characteristic case, if K contains all of the roots of unity (of order prime to p , if $\text{char } K = p \neq 0$), then Theorem 4.3.3 shows that the free profinite group \hat{F}_r on r generators is a quotient of $\pi_1(U)$. (Namely, the cyclotomic character acts trivially in this case, and so the semi-direct product in (***) is just a direct product.) Since arbitrarily

large pairwise adjusted subsets S exist by Remark (a), this shows that \hat{F}_r is a quotient of the absolute Galois group of $K(x)$ for each $r \in \mathbb{N}$. A similar result holds in the unequal characteristic case $(0, p)$ if K contains the prime-to- p roots of unity, namely that the free pro-prime-to- p group \hat{F}'_r of rank $r \in \mathbb{N}$ is a quotient of $\pi_1(U)$ and of $G_{K(x)}$. But the full group \hat{F}_r is *not* a quotient of $\pi_1(U)$ or $G_{K(x)}$ in the unequal characteristic case; cf. [Po2] and Remark (c) below.

(c) The result in [Po2] asserts even more. First of all, the i th generator of Π generates an inertia group over ξ_i and over η_i , for each $i = 1, \dots, r$. This is as in the case of analytic and formal slit covers discussed at the ends of Sections 2.3 and 3.3 above. Second, in the unequal characteristic case $(0, p)$, the assertion of Theorem 4.3.3 may be improved slightly, by replacing the group $\Pi = \hat{F}'_r$ by a group of the form $\hat{F}'_r \times \mathbb{Z}/p^e$, for a certain non-negative integer e . (Specifically, $e = \max(0, e')$, where e' is the largest integer such that $|\xi_i - \eta_i| < |p|^{e'+1/(p-1)}|\xi_i - \xi_j|$ for all $i \neq j$.) This group lies in between the group \hat{F}_r and its quotient \hat{F}'_r ; and in the case that K contains all the prime-to- p roots of unity, this group is then a quotient of $\pi_1(U)$ and $G_{K(x)}$ (like \hat{F}'_r but unlike \hat{F}_r). See [Po2] for details.

(d) The construction of cyclic extensions given in Section 3.3 can be recovered from the above result, in the case that the extension is of degree n prime to the characteristic of K . Namely, given a cyclic group $C = \langle c \rangle$ of order n , consider a primitive element for $K' := K(\zeta_n)$ as an extension of K ; this corresponds to a K' -point $\xi = \xi_1$ of \mathbb{P}^1 , and $\text{Gal}(K'/K)$ acts simply transitively on the G_K -orbit $\{\xi_1, \dots, \xi_s\}$ of ξ . Take $\eta = \eta_1$ sufficiently close to ξ to satisfy continuity of the roots [La, II, §2, Proposition 4] (and also to satisfy the inequality in Remark (c) above, in the mixed characteristic case $(0, p)$ if $p|n$); and let its orbit be $\{\eta_1, \dots, \eta_s\}$. Let $U \subset \mathbb{P}^1_K$ be the complement of the ξ_i 's and η_i 's. Consider the surjection $\Pi \twoheadrightarrow C$ given by $g_j \mapsto c^{\chi(\alpha^{-1})} = c^{-\alpha(\zeta_n)}$ if $\alpha \in \text{Gal}(K'/K)$ is the element taking ξ to ξ_j . Then in the quotient $C \rtimes G_K$ of $\Pi \rtimes G_K$ (and hence of $\pi_1(U)$), the action of G_K on C is trivial; i.e. the quotient is just $C \times G_K$. So it in turn has a quotient isomorphic to C ; and this corresponds to the cyclic cover constructed in the proof of Proposition 3.3.3. (In the case that n is instead a power of $p = \text{char } K$, one uses Witt vectors in the construction; and again one obtains cyclic covers of degree n , since the action of G_K via χ is automatically trivial on a p -group quotient of Π .)

(e) The main assertion in Theorem 4.3.1 above (and in Theorem 3.3.1), that every finite group G is a Galois group over $K(x)$, can be recaptured from the Half Riemann Existence Theorem. Namely, by choosing elements c_i that generate G , and applying Remark (d) separately to each c_i , one obtains a quotient of $\Pi \rtimes G_K$ of the form $G \rtimes G_K$, in which the semi-direct product is actually a direct product. So G is a quotient of $\pi_1(U)$. \square

Unfortunately, the above approach (like that of Section 3.3) does not provide an explicit description, in terms of generators and relations, of the *full* fundamental group (or at least the tame fundamental group) of an arbitrary affine K -curve U . Such a full ‘‘Riemann’s

Existence Theorem” would generalize the explicit classical result over \mathbb{C} (Corollary 2.1.2), unlike Lütkebohmert’s result [Lü] discussed at the end of Section 4.1 (which is inexplicit) and the above result (which gives only a big quotient of $\pi_1(U)$).

At the moment such a full, explicit result (or even a conjecture about its exact statement) seems far out of reach, even in key special cases. For example, if K is algebraically closed of characteristic p , the profinite groups $\pi_1(\mathbb{A}_K^1)$ and $\pi_1^t(\mathbb{A}_K^1 - \{0, 1, \infty\})$ are unknown. And if K is a p -adic field, the tower of all Galois branched covers of \mathbb{P}_K^1 remains mysterious, while little is understood about Galois branched covers of \mathbb{P}_K^1 with good reduction and their associated Galois groups. (Note that the covers constructed above and in Section 3.3 have models over \mathbb{Z}_p in which the closed fibres are quite singular — as is clear from the mock cover construction of Section 3.3.) For $p > 3$, a wildly ramified cover of $\mathbb{P}_{\mathbb{Q}_p}^1$ cannot have good reduction over \mathbb{Q}_p (or even over the maximal unramified extension of \mathbb{Q}_p) [Co, p.247, Remark 3]; and so \mathbb{Z}/p cannot be such a Galois group. But it is unknown whether every finite group G is the Galois group of a cover of \mathbb{P}_K^1 with good reduction over K , for some *totally ramified* extension K of \mathbb{Q}_p (depending on G); if so, this would imply that every finite group is a Galois group over the field $\mathbb{F}_p(x)$ (cf. Proposition 3.3.9).

See Section 5 for a further discussion of results in the direction of a generalized Riemann’s Existence Theorem.

In the rigid patching constructions above, and in the analogous formal patching constructions in Section 3.3, the full generality of rigid analytic spaces and formal schemes is not needed in order to obtain the results in Galois theory. Namely, the rigid analytic spaces and formal schemes that arise in these proofs are induced from algebraic varieties; and so less machinery is needed in order to prove the results of these sections than might first appear. Haran and Völklein (and later Jarden) have developed an approach to patching that goes further, and which seeks to omit all unnecessary geometric objects. Namely, in [HV], the authors created a context of “algebraic patching” in which everything is phrased in terms of rings and fields (viz. the rings of functions on formal or rigid patches, and their fraction fields), and in which the geometric and analytic viewpoints are suppressed. That set-up was then used to reprove Corollary 4.3.2 above on realizing Galois groups regularly over large fields [HV, Theorem 4.4], as well as to prove additional related Galois results (in [HV], [HJ1], and [HJ2]). For covers of curves, it appears that the formal patching, rigid patching, and algebraic patching methods are essentially interchangeable, in terms of what they are capable of showing. The main differences concern the intuition and the precise machinery involved; and these are basically matters of individual mathematical taste. In other applications, it may turn out that one or another of these methods is better suited.

Section 5: Toward Riemann’s Existence Theorem

Sections 3 and 4 showed how formal and rigid patching methods can be used to establish analogs of GAGA, and to realize all finite groups as Galois groups of covers, in rather general settings. This section pursues these ideas further, in the direction of a sought-after “Riemann’s Existence Theorem” that would classify covers in terms of group-theoretic data, corresponding to the Galois group, the inertia groups, and how the covers fit together in a tower. Section 5.1 discusses embedding problems, which relate to understanding this tower. In particular, §5.1 gives the structure of the absolute Galois group of the function field of a curve over an algebraically closed field (which can be regarded as the geometric case of a conjecture of Shafarevich). Section 5.2 relates patching and embedding problems to arithmetic lifting problems, in which one considers the existence of a cover with a given Galois group and a given fibre (over a non-algebraically closed base field). In doing so, it relies on results from Section 5.1. Section 5.3 considers Abhyankar’s Conjecture on fundamental groups in characteristic p , along with strengthenings and generalizations that relate to embedding problems and patching. These results move further in the direction of a full “Riemann’s Existence Theorem”, although the full classification of covers in terms of groups remains unknown.

Section 5.1. Embedding problems and the geometric case of Shafarevich’s Conjecture.

The motivation for introducing patching methods into Galois theory was to prove results about Galois groups and fundamental groups for varieties that are not necessarily defined over \mathbb{C} . Complex patching methods, combined with topology, permitted a quite explicit description of the tower of covers of a given complex curve U (Riemann’s Existence Theorem 2.1.1 and 2.1.2). In particular, this approach showed what the fundamental group of U is, and thus which finite groups are Galois groups of unramified covers of U . Analogous formal and rigid patching methods were applied (in Sections 3 and 4) to the study of curves over certain other coefficient fields, in particular *large* fields. Without restriction on the branch locus, it was shown that every finite group is a Galois group over the function field of the curve (Sections 3.3 and 4.3), and Pop’s “Half Riemann’s Existence Theorem” gave an explicit description of a big part of the tower of covers for certain special choices of branch locus (Section 4.3). Further results about Galois groups over an arbitrary affine curve have also been obtained (see Section 5.3 below), but an explicit description of the full tower of covers, and of the full fundamental group, remain out of reach for now.

Nevertheless, if one does not restrict the branch locus, then patching methods can be used to find the birational analog of the fundamental group, in the case of curves over an algebraically closed field k — i.e. to find the absolute Galois group of the function field of a k -curve X . And here, unlike the situation with the fundamental group of an affine k -curve, the absolute Galois group turns out to be free even in characteristic $p > 0$.

In the case $k = \mathbb{C}$ and $X = \mathbb{P}_{\mathbb{C}}^1$, this result was proven in [Do] using the classical form of Riemann's Existence Theorem (see Corollary 2.1.5). For more general fields, it was proven independently by the author [Ha10] and by F. Pop [Po1], [Po3]:

Theorem 5.1.1. *Let X be an irreducible curve over an algebraically closed field k of arbitrary characteristic. Then the absolute Galois group of the function field of X is the free profinite group of rank equal to the cardinality of k .*

In particular, the absolute Galois group of $k(x)$ is free profinite of rank equal to $\text{card } k$.

Theorem 5.1.1 implies the geometric case of Shafarevich's Conjecture. In the form originally posed by Shafarevich, the conjecture says that the absolute Galois group $\text{Gal}(\bar{\mathbb{Q}}/\mathbb{Q}^{\text{ab}})$ of \mathbb{Q}^{ab} is free profinite of countable rank. Here \mathbb{Q}^{ab} denotes the maximal abelian extension of \mathbb{Q} , or equivalently (by the Kronecker-Weber theorem) the maximal cyclotomic extension of \mathbb{Q} (i.e. \mathbb{Q} with all the roots of unity adjoined). The conjecture was later generalized to say that if K is any global field and K^{cycl} is its maximal cyclotomic extension, then the absolute Galois group of K^{cycl} is free profinite of countable rank. (See Remark 3.3.8(b).) The arithmetic case of this conjecture (the case where K is a number field) is still open, but the geometric case (the case where K is the function field of a curve X over a finite field F) follows from Theorem 5.1.1, by considering passage to the algebraic closure \bar{F} of F . Namely, in this situation, $\bar{F} = \bar{\mathbb{F}}_p$ where $p = \text{char } F$, and so the function field \bar{K} of $\bar{X} := X \times_F \bar{F}$ is equal to K^{cycl} ; and in this case Theorem 5.1.1, applied to the \bar{K} -curve \bar{X} , asserts the conclusion of Shafarevich's Conjecture.

Theorem 5.1.1 above is proven using the notion of embedding problems. Recall that an *embedding problem* \mathcal{E} for a profinite group Π is a pair of surjective group homomorphisms $(\alpha : \Pi \rightarrow G, f : \Gamma \rightarrow G)$. A *weak* [resp. *proper*] *solution* to \mathcal{E} consists of a group homomorphism [resp. epimorphism] $\beta : \Pi \rightarrow \Gamma$ such that $f\beta = \alpha$. An embedding problem \mathcal{E} is *finite* if Γ is finite; it is *split* if f has a section; it is *non-trivial* if $\ker f$ is non-trivial; it is a *p-embedding problem* if $\ker f$ is a p -group. A profinite group Π is *projective* if every finite embedding problem for Π has a weak solution.

In terms of Galois theory, if Π is the absolute Galois group of a field K , then giving a G -Galois field extension L of K is equivalent to giving a surjective homomorphism $\alpha : \Pi \rightarrow G$. For such an L , giving a proper solution to \mathcal{E} as above is equivalent to giving a Γ -Galois field extension F of K together with an embedding of L into F as a G -Galois K -algebra. (Here the G -action on L agrees with the one induced by restricting the action of Γ to the image of the embedding.) Giving a weak solution to \mathcal{E} is the same, except that F need only be a separable K -algebra, not a field extension (and so it can be a direct product of finitely many fields). In this field-theoretic context we refer to an *embedding problem for K* .

If K is the function field of a geometrically irreducible k -scheme X , then the field

extensions L and F correspond to branched covers $Y \rightarrow X$ and $Z \rightarrow X$ which are G -Galois and Γ -Galois respectively, such that Z dominates Y . Here Y is irreducible; and Z is also irreducible in the case of a proper solution. If the algebraic closure of k in the function fields of Y and Z are equal (i.e. if there is no extension of constants from Y to Z), we say that the solution is *regular*.

By considering embedding problems for a field K , or over a scheme X , one can study not only which finite groups are Galois groups over K or X , but how the extensions or covers fit together in the tower of all finite Galois groups. As a result, one can obtain information about absolute Galois groups and fundamental groups. In particular, in the key special case that X is the projective line and k is countable (e.g. if $k = \overline{\mathbb{F}}_p$), Theorem 5.1.1 follows from the following three results about embedding problems:

Theorem 5.1.2. (Iwasawa [Iw,p.567], [FJ,Cor.24.2]) *Let Π be a profinite group of countably infinite rank. Then Π is a free profinite group if and only if every finite embedding problem for Π has a proper solution.*

Theorem 5.1.3. (Serre [Se6, Prop. 1]) *If U is an affine curve over an algebraically closed field k , then $\pi_1(U)$ is projective.*

Theorem 5.1.4. (Harbater [Ha10], Pop [Po1], [Po3]) *If k is an algebraically closed field, and K is the function field of an irreducible k -curve X , then every finite split embedding problem for K has a proper solution.*

Theorem 5.1.2 is entirely group-theoretic (and *rank* refers to the minimal cardinality of any generating set). The proof of Theorem 5.1.3 is cohomological, and in fact the assertion in [Se6] is stated in terms of cohomological dimension (that $\text{cd}(\pi_1(X)) \leq 1$, which implies projectivity by [Se4, I, 5.9, Proposition 45]). Theorem 5.1.4 is a strengthening of Theorem 3.3.1, and like that result it is proven using patching. (Theorem 5.1.4 will be discussed in more detail below.)

Using these results, Theorem 5.1.1 can easily be shown in the case that the algebraically closed field k is countable. Namely, let Π be the absolute Galois group of $k(x)$. Then the profinite group Π has at most countable rank, since the countable field $k(x)$ has only countably many finite field extensions; and Π has infinite rank, since every finite group is a quotient of Π (as seen in Section 3.3). So Theorem 5.1.2 applies, and it suffices to show that every finite embedding problem \mathcal{E} for Π is properly solvable. Say \mathcal{E} is given by $(\alpha : \Pi \rightarrow G, f : \Gamma \rightarrow G)$, with f corresponding to a G -Galois branched cover $Y \rightarrow X$. This cover is étale over an affine dense open subset $U \subset X$, and α factors through $\pi_1(U)$ (since quotients of π_1 classify unramified covers). Writing this map as $\alpha_U : \pi_1(U) \rightarrow G$, consider the finite embedding problem $\mathcal{E}_U = (\alpha_U : \pi_1(U) \rightarrow G, f : \Gamma \rightarrow G)$. By Theorem 5.1.3, this has a weak solution $\beta_U : \pi_1(U) \rightarrow \Gamma$, say with image $H \subset \Gamma$ (which surjects onto G under f). Let N be the kernel of f , and Γ_1 be the semidirect product $N \rtimes H$ with respect to the

conjugation action of H on N . The multiplication map $(n, h) \mapsto nh \in \Gamma$ is an epimorphism $m : \Gamma_1 \rightarrow \Gamma$, and the projection map $h : \Gamma_1 \rightarrow H$ is surjective with kernel N . The surjection $\beta_U : \pi_1(U) \rightarrow H$ corresponds to an H -Galois branched cover $Y_1 \rightarrow X$ (unramified over U). This in turn corresponds to a surjective group homomorphism $\beta : \Pi \rightarrow H$. By Theorem 5.1.4, the split embedding problem $(\beta : \Pi \rightarrow H, h : \Gamma_1 \rightarrow H)$ has a proper solution. That solution corresponds to an irreducible Γ_1 -Galois cover $Z_1 \rightarrow X$ that dominates Y_1 ; and composing the corresponding surjection $\Pi \rightarrow \Gamma_1$ with $m : \Gamma_1 \rightarrow \Gamma$ provides a proper solution to the original embedding problem \mathcal{E} .

Remark 5.1.5. The above argument actually requires less than Theorem 5.1.3; viz. it suffices to use Tsen's Theorem [Ri, Proposition V.5.2] that if k is algebraically closed then the absolute Galois group of $k(x)$ has cohomological dimension 1. For then, by writing X in Theorem 5.1.1 as a branched cover of \mathbb{P}_k^1 , it follows that the absolute Galois group of its function field is also of cohomological dimension 1 [Se4, I, 3.3, Proposition 14], and hence is projective [Se4, I, 5.9, Proposition 45]. One can then proceed as before.

But by using Theorem 5.1.3 as in the argument above, one obtains additional information about the branch locus of the solution to the embedding problem. Namely, one sees in the above argument that the H -Galois cover $Y_1 \rightarrow X$ remains étale over U . In applying Theorem 5.1.4 to pass to a Γ_1 -cover (and thence to a Γ -cover), one typically obtains new branch points. But a sharp upper bound can be found on the number of additional branch points [Ha11], using Abhyankar's Conjecture (discussed in Section 5.3). \square

Before turning to the general case of Theorem 5.1.1 (where k is allowed to be uncountable), we sketch the proof of Theorem 5.1.4:

Proof sketch of Theorem 5.1.4. Let Π be the absolute Galois group of K . Consider a finite split embedding problem $\mathcal{E} = (\alpha : \Pi \rightarrow G, f : \Gamma \rightarrow G)$ for K , with s a section of f , and with f corresponding to a G -Galois branched cover $Y \rightarrow X$. Let $N = \ker(f)$, and let n_1, \dots, n_r be generators of N . Thus Γ is generated by $s(G)$ and the n_i 's. Pick r closed points $\xi_i \in X$ that are not branch points of $Y \rightarrow X$. Let $k' = k((t))$, and let $X' = X \times_k k'$ and similarly for Y' and ξ'_i . Pick small neighborhoods X'_i around each of the points ξ'_i on X' . (Here, if one works in the rigid context, one takes t -adic closed discs. If one works in the formal context, one blows up at the points ξ'_i , and proceeds as in Example 3.2.11 or 3.2.13, using Theorem 3.2.8 or 3.2.12. See also Example 4.2.4.) Over these neighborhoods, build cyclic covers $Z'_i \rightarrow X'_i$ with group $N_i = \langle n_i \rangle$ (branched at ξ'_i and possibly other points; cf. the proof of Proposition 3.3.3, using the presence of prime-to- p roots of unity). Let $Y'_0 \rightarrow X'_0$ be the restriction of $Y' \rightarrow X'$ away from the above neighborhoods (viz. over the complement of the corresponding open discs if one works rigidly, and over the general fibre of the formal completion at the complement of the ξ_i 's if one works formally). Via the section s of f , the Galois group G of $Y'_0 \rightarrow X'_0$ may be identified with $s(G) \subset \Gamma$. The induced Γ -Galois covers $\text{Ind}_{N_i}^{\Gamma} Z'_i \rightarrow Z'_i$ and $\text{Ind}_{s(G)}^{\Gamma} Y'_0 \rightarrow X'_0$ agree over the (rigid or formal)

overlap. Hence by (rigid or formal) GAGA, these patch together to form a Γ -Galois cover $Z' \rightarrow X'$. (In the formal case, one uses Theorem 3.2.8 rather than Theorem 3.2.1, since the agreement is not on the completion along a Zariski open set.) This cover is connected since Γ is generated by $s(G)$ and the n_i 's; it dominates $Y' \rightarrow X'$ since it does on each patch; and it is branched at each ξ'_i . As in the proof of Corollary 3.3.5, one may now specialize from k' to k using that k is algebraically closed, obtaining a Γ -Galois cover $Z \rightarrow X$ that dominates $Y \rightarrow X$. This corresponds to a proper solution to \mathcal{E} . \square

Remark 5.1.6. (a) The above proof also shows that one has some control over the position of the new branch points of $Z \rightarrow X$. Namely, the branch locus contains the points ξ_i , and these points can be taken arbitrarily among non-branch points of $Y \rightarrow X$. In particular, any given point of X can be taken to be a branch point of $Z \rightarrow X$ above (by choosing it to be one of the ξ_i 's). More precise versions of this fact appear in [Ha10, Theorem 3.5] and [Po3, Theorem A], where formal and rigid methods are respectively used.

(b) As a consequence of Remark (a), it follows that the set of (isomorphism classes of) solutions to the split embedding problem has cardinality equal to that of k .

(c) The above proof of Theorem 5.1.4 also gives information about inertia of the constructed cover $Z \rightarrow X$. Namely, if $I \subset G$ is the inertia group of $Y \rightarrow X$ at a point $\eta \in Y$ over $\xi \in X$, then $s(I) \subset \Gamma$ is an inertia group of $Z \rightarrow X$ at a point $\zeta \in Z$ over η (and the other inertia groups over ξ are the conjugates of $s(I)$).

(d) Adjustments to the above construction give additional flexibility in controlling the properties of $Z \rightarrow X$. In particular, if $\text{char } k = p > 0$ and if $I' \subset \Gamma$ is the extension of $s(I)$ by a p -group, then one may build Z so that I' is an inertia group over ξ at a point over η (with notation as in Remark (c)). In addition, rather than considering a split embedding problem, i.e. a group Γ generated by a normal subgroup N and a complement $s(G)$, one can more generally consider a group Γ generated by two subgroups H and G , where we are given a G -Galois cover $Y \rightarrow X$. The assertion then says that this cover can be modified to produce a Γ -Galois cover $Z \rightarrow X$ with control as above on the branch locus and inertia groups. In particular, one can add additional branch points to a cover, and one can modify a cover by enlarging an inertia group from a p -subgroup of the Galois group to a larger p -subgroup. (See [Ha6, Theorem 2] and [Ha13, Theorem 3.6], where formal patching is used to prove these assertions.)

(e) The ability to add branch points was used in [MR] to show that for any finite group G and any smooth connected curve X over an algebraically closed field k , there is a G -Galois branched cover $Y \rightarrow X$ such that G is the *full* group of automorphisms of Y . The idea is that if one first takes an arbitrary G -Galois cover of X (by Corollary 3.3.5); then one can adjust it by adding new branch points and thereby killing automorphisms that are not in G . \square

To prove the general case of Theorem 5.1.1, one replaces Theorem 5.1.2 above by a

result of Melnikov and Chatzidakis (see [Ja, Lemma 2.1]):

Theorem 5.1.7. *Let Π be a profinite group and let m be an infinite cardinal. Then Π is a free profinite group of rank m if and only if every non-trivial finite embedding problem for Π has exactly m proper solutions.*

Namely, by Remark 5.1.6(b) above, in the situation of Theorem 5.1.1 the number of proper solutions to any finite *split* embedding problem is $\text{card } k$. Proceeding as in the proof of Theorem 5.1.1 in the countable case, one obtains that *every* finite embedding problem for Π has $\text{card } k$ proper solutions. So Π is free profinite of that rank by Theorem 5.1.7, and this proves Theorem 5.1.1.

Remark 5.1.8. By refining the proof of Theorem 5.1.1 (in particular modifying Theorems 5.1.3 and 5.1.4 above), one can prove a tame analog of that result [Ha13, Theorem 4.9(b)]: If X is an affine curve with function field K , consider the maximal extension Ω of K that is at most tamely ramified over each point of X . Then $\text{Gal}(\Omega/K)$ is a free profinite group, of rank equal to the cardinality of k . \square

Theorem 5.1.4 above extends from algebraically closed fields to arbitrary large fields (cf. Section 3.3), according to the following result of Pop:

Theorem 5.1.9. (Pop [Po1, Theorem 2.7]) *If k is a large field, and K is the function field of a geometrically irreducible k -curve X , then every finite split embedding problem for K has a proper regular solution.*

Namely, the above proof of Theorem 5.1.4 showed that result for an *algebraically closed* field k by first proving it for the Laurent series field $K = k((t))$, and then specializing from K to k , using that k is algebraically closed. In order to prove Theorem 5.1.9, one does the same in this more general context, using that k is large in order to specialize from $K = k((t))$ to k (as in Sections 3.3 and 4.3). A difficulty is that since k need not be algebraically closed, one can no longer choose the extra branch points $\xi_i \in X$ arbitrarily (as one could in the above proof of Theorem 5.1.4). Still, one can proceed as in the proofs of Theorems 3.3.1, 4.3.1, and 4.3.3 — viz. using cyclic covers branched at clusters of points constructed in the proof of Proposition 3.3.3.

Since an arbitrary large field k is not algebraically closed, one would also like to know that the Γ -Galois cover $Z \rightarrow X$ has the property that $Z \rightarrow Y$ is *regular* (i.e. Z and Y have the same ground field ℓ , or equivalently the algebraic closures of k in the function fields of Y and Z are equal). This can be achieved by using that in the construction using formal patching, the closed fibre of the cover $Z \rightarrow Y$ over K is a mock cover (as in the proof of Theorem 3.3.1). Alternatively, from the rigid point of view, one can observe from the patching construction (as in the proof of Theorem 4.3.1) that Z may be chosen so that $Z \rightarrow Y$ has a totally split fibre over $\eta \in Y$, if η has been chosen (in advance) to

be an ℓ -point of Y that lies over a k -point ξ of X . This then implies regularity, as in Theorem 4.3.1. (If there is no such point $\eta \in Y$, then one can first base-change to a finite Galois extension \tilde{k} of k where there is such a point; and then construct a regular solution $\tilde{Z} \rightarrow \tilde{Y} = Y \times_k \tilde{k}$ which is compatible with the $\text{Gal}(\tilde{k}/k)$ -action, and so which descends to a regular solution $Z \rightarrow Y$.)

Remarks 5.1.6(a) and (b) above no longer hold for curves over an arbitrary large field (nor does Theorem 5.1.1 — see below); but Remark 5.1.6(c) still applies in this situation. So the argument in the case of an arbitrary large field gives the following more precise form of Theorem 5.1.9 (where one looks at the actual curve X , rather than just at its function field):

Theorem 5.1.10. *Let k be a large field, let X be a geometrically irreducible smooth k -curve, let $f : \Gamma \rightarrow G$ be a surjection of finite groups with a section s , and let $Y \rightarrow X$ be a G -Galois connected branched cover of smooth curves.*

(a) *Then there is a smooth connected Γ -Galois branched cover $Z \rightarrow X$ that dominates the G -Galois cover $Y \rightarrow X$, such that $Z \rightarrow Y$ is regular.*

(b) *Let ξ be a k -point of X which is not a branch point of $Y \rightarrow X$, and let η be a closed point of Y over ξ with decomposition group $G_1 \subset G$. Then the cover $Z \rightarrow X$ in (a) may be chosen so that it is totally split over η , and so that there is a point $\zeta \in Z$ over η whose decomposition group over ξ is $s(G_1) \subset \Gamma$.*

Remark 5.1.11. (a) In [Po1], the above result was stated for a slightly smaller class of fields (those with a “universal local-global principle”); but in fact, all that was used is that the field is large. Also, the result there did not assert 5.1.10(b), though this can be deduced from the proof. The result was stated for large fields in [Po4, Main Theorem A], but only in the case that $X = \mathbb{P}_k^1$ and $Y = \mathbb{P}_\ell^1$. (Both proofs used rigid patching.) The fact that the fibre over η can be chosen to be totally split first appeared explicitly in [HJ1, Theorem 6.4], in the case that $X = \mathbb{P}_k^1$ and $Y = \mathbb{P}_\ell^1$; and in [HJ2, Proposition 4.2] if $X = \mathbb{P}_k^1$ and Y is arbitrary. The proofs there used “algebraic patching” (cf. the comments at the end of Section 4.3).

(b) A possible strengthening of Theorem 5.1.10(b) would be to allow one to specify the decomposition group of ζ as a given subgroup $G'_1 \subset \Gamma$ that maps isomorphically onto $G_1 \subset G$ via f (rather than having to take $G'_1 = s(G_1)$, as in the statement above). It would be interesting to know if this strengthening is true. \square

As a consequence of Theorem 5.1.9, we have:

Corollary 5.1.12. *Let k be a Hilbertian large field, with absolute Galois group G_k .*

(a) *Then every finite split embedding problem for G_k has a proper solution.*

(b) *If k is also countable, and if G_k is projective, then G_k is isomorphic to the free profinite group of countable rank.*

Proof. (a) Every such embedding problem for G_k gives a split embedding problem for $G_{k(x)}$. That problem has a proper solution by Theorem 5.1.9. Since k is Hilbertian, that solution can be specialized to a proper solution of the given embedding problem.

(b) Since G_k is projective, the conclusion of part (a) implies that *every* finite embedding problem for G_k has a proper solution (as in the proof of Theorem 5.1.1 above, using semi-direct products). Also, G_k is of countably infinite rank (again as in the proof of Theorem 5.1.1). So Theorem 5.1.2 implies the conclusion. \square

Remark 5.1.13. (a) Part (a) of Corollary 5.1.12 appeared in [Po3, Main Theorem B] and [HJ1, Thm.6.5(a)]. As a special case of part (b) of the corollary, one has that if k is a countable Hilbertian PAC field (see Example 3.3.7(c)), then G_k is free profinite of countable rank. This is because PAC fields are large, and because their absolute fundamental groups are projective (because they are of cohomological dimension ≤ 1 [Ax2, §14, Lemma 2]). This special case had been a conjecture of Roquette, and it was proven as above in [Po3, Thm. 1] and [HJ1, Thm. 6.6] (following a proof in [FV2] in the characteristic 0 case, using the classical complex analytic form of Riemann’s Existence Theorem).

(b) As remarked in Section 3.3, it is unknown whether \mathbb{Q}^{ab} is large. But it is Hilbertian [Vö, Corollary 1.28] and countable (being contained in $\bar{\mathbb{Q}}$), and its absolute Galois group is projective (being of cohomological dimension 1 by [Se4, II, 3.3, Proposition 9]). So if it is indeed large, then part (b) of the corollary would imply that its absolute Galois group is free profinite of countable rank — i.e. the original (arithmetic) form of Shafarevich’s Conjecture would hold. Among other things, this would imply that every finite group is a Galois group over \mathbb{Q}^{ab} .

The solvable version of Shafarevich’s Conjecture has been shown; i.e. the maximal pro-solvable quotient of $G_{\mathbb{Q}^{\text{ab}}}$ is the free prosolvable group of countable rank [Iw]. More generally, if k is Hilbertian and G_k is projective, then every finite embedding problem for G_k with solvable kernel has a proper solution [Vö, Corollary 8.25]. This result does not require k to be large, and it does not use patching.

(c) It has been conjectured by Dèbes and Deschamps [DD] that Theorem 5.1.9 and Corollary 5.1.12 remain true even if the ground field is not large. Specifically, they conjecture that for any field k , every finite split embedding problem for $G_{k(x)}$ has a proper regular solution; and hence that if k is Hilbertian, then every finite split embedding problem for G_k has a proper solution. This is a very strong conjecture, in particular implying an affirmative answer to the Regular Inverse Galois Problem (i.e. that every finite group is a regular Galois group over $k(x)$ for every field k). But it also seems very far away from being proven. \square

As mentioned above, Theorem 5.1.1 does not hold if the algebraically closed field k is replaced by an arbitrary large field. This is because if K is the function field of a k -curve X , then its absolute Galois group G_K is not even projective (much less free) if k is not

separably closed. That is, not every finite embedding problem for K has a weak solution — and so certainly not a proper solution, as would be required in order to be free.

This can be seen by using the equivalence between the condition that a profinite group Π is projective and the condition that it has cohomological dimension ≤ 1 [Se4, I; 5.9, Proposition 45 and 3.4, Proposition 16]. Namely, if k is not separably closed, then its absolute Galois group G_k is non-trivial, and so G_k has cohomological dimension > 0 [Se4, I, 3.3, Corollaire 2 to Proposition 14]. Since the function field K is of finite type over k and of transcendence degree 1 over k , it follows that G_K has cohomological dimension > 1 . (This is by [Se4, II, 4.2, Proposition 11] in the case that $\text{cd } G_k$ is finite; and by [Ax1] and [Se4, II, 4.1, Proposition 10(ii)] if $\text{cd } G_k$ is infinite.) So G_K is not projective.

But as Theorem 5.1.9 shows, every finite *split* embedding problem for G_K has a proper solution, if K is the function field of a curve over an arbitrary large field k . Thus (as in the proof of Theorem 1 above, via semi-direct products), it follows that any finite embedding problem for G_K that has a weak solution must also have a proper solution. So Theorem 5.1.9 can be regarded as saying that G_K is “as close as possible” to being free, given that it is not projective.

Section 5.2. Arithmetic lifting, embedding problems, and patching.

In realizing Galois groups over a Hilbertian field k like \mathbb{Q} or \mathbb{Q}^{ab} , the main method is to realize the group as a regular Galois group over $K = k(x)$, and then to specialize from K to k using that k is Hilbertian. That is, one constructs a Galois branched cover $Y \rightarrow \mathbb{P}_k^1$ such that k is algebraically closed in the function field of Y , and then obtains a Galois extension of k with the same group by considering an irreducible fibre of the cover over a k -point of \mathbb{P}_k^1 (which exists by the Hilbertian hypothesis). To date, essentially all simple groups that have been realized as Galois groups over \mathbb{Q} or \mathbb{Q}^{ab} have been realized by this method.

The use of this method has led to the question of whether, given a finite Galois extension ℓ of a field k , there is a finite regular Galois extension L of $K = k(x)$ with the same group G , of which the given extension is a specialization. If so, then one says that the field k and group G satisfy the *arithmetic lifting property*. (Of course if one did not require regularity, then one could just take L to be $\ell(x)$.)

The question of when this property holds was first raised by S. Beckmann [Be], who showed that it does hold in the case that $k = \mathbb{Q}$ and G is either an abelian group or a symmetric group. Later, E. Black [B11] [B12] [B13] showed that the property holds for certain more general classes of groups over Hilbertian fields, particularly certain semi-direct products such as dihedral groups D_n with n odd. Black also conjectured that the arithmetic lifting property holds for all finite groups over all fields, and proving this has come to be known as the *Beckmann-Black problem* (or BB). It was later shown by Dèbes [Dè] that an affirmative answer to BB over *every* field would imply an affirmative answer

to the Regular Inverse Galois Problem (RIGP) over every field (i.e. that for every field k and every finite group G , there is a regular Galois extension of $k(x)$ with group G). On the other hand, knowing BB for a *given* field k does not automatically give RIGP over k , since one needs to be given a Galois extension of the given field in order to apply BB.

Colliot-Thélène has considered a strong form of arithmetic lifting (or BB): Given a field k and a finite group G , and given a separable G -Galois algebra A over k , is there a regular G -Galois field extension L of $k(x)$ that specializes to A ? That is, suppose that H is a subgroup of G and ℓ is an H -Galois field extension of k . Then the question is whether there is a regular G -Galois field extension of $k(x)$ such that some specialization to k yields $A := \ell^{\oplus(G:H)}$. In geometric terms, the question is whether there is a regular G -Galois branched cover $Y \rightarrow \mathbb{P}_k^1$ with a given fibre $\text{Ind}_H^G \text{Spec } \ell$ — i.e. such that over some unramified k -point of the line, there is a point of Y with given decomposition group $H \subset G$ and given residue field ℓ (which is a given H -Galois field extension of k).

If, in the strong form of BB, one takes A to be a G -Galois field extension ℓ of k , then one recovers the original BB. At the other extreme, if one takes A to be a direct sum of copies of k (indexed by the elements of G), then one is asking the question of whether there is a G -Galois regular field extension of $k(x)$ with a totally split fibre. (Thus the strong form of BB over a *given* field k implies RIGP for that field.) In the case that k is a large field, this totally split case of strong BB does hold; indeed, this is precisely the content of Theorem 1 of Section 4.3.

Colliot-Thélène showed that the strong form of BB holds in general for large fields k :

Theorem 5.2.1. [CT] *If k is a large field, G is a finite group, and A is a G -Galois k -algebra, then there is a G -Galois regular branched cover of $X = \mathbb{P}_k^1$ whose fibre over a given k -point agrees with $\text{Spec } A$ (as a G -Galois k -algebra).*

Remark 5.2.2. As noted in Remark 5.1.13(b), it is unknown whether \mathbb{Q}^{ab} is large. But if it is, then Theorem 5.2.1 would imply that it has the (strong) arithmetic lifting property for every finite group — and so every finite Galois group over \mathbb{Q}^{ab} would be the specialization of a regular Galois branched cover of the line over \mathbb{Q}^{ab} . On the other hand, \mathbb{Q} is not large, and so Theorem 5.2.1 does not apply to it. And although it is known that every finite solvable group is a Galois group over \mathbb{Q} (Shafarevich’s Theorem [NSW, Chap. IX, §5]), it is not known whether every such group is the Galois group of a regular branched cover of $\mathbb{P}_{\mathbb{Q}}^1$ — much less that the arithmetic lifting property holds for these groups over \mathbb{Q} . \square

Colliot-Thélène’s proof used a different form of patching, and relied on work of Kollár [Kol] on rationally connected varieties. The basic idea is to construct a “comb” of projective lines on a surface, i.e. a tree of \mathbb{P}^1 ’s in which one component meets all the others, none of which meet each other. A degenerate cover of the comb is then constructed by building it over the components, and the cover is then deformed to a non-degenerate cover of a nearby irreducible curve of genus 0 with the desired properties.

Colliot-Thélène’s proof required that k be of characteristic 0 (because Kollár’s work assumed that), but other proofs have been found that do not need this. In particular, Moret-Bailly [MB2] used a formal patching argument to prove this result. A proof using rigid patching can be obtained from Colliot-Thélène’s argument by replacing the “spine” of the comb by an affinoid set U_0 as in the proof of Theorem 4.3.1, and the “teeth” of the comb by affinoids U_1, \dots, U_r as in that proof (appropriately chosen). And a proof using “algebraic patching” (cf. the end of Section 4.3) has been found by Haran and Jarden [HJ2].

Yet another proof of Theorem 5.2.1 above can be obtained from Pop’s result on solvability of split embedding problems over large fields (Theorem 5.1.9, in the more precise form Theorem 5.1.10 — which of course was also proven using patching). This proof needs only the special case of that earlier result in which the given cover of \mathbb{P}_k^1 is purely arithmetic (i.e. of the form \mathbb{P}_ℓ^1), this being the case considered in [Po4] and [HJ1]. Namely, under the hypotheses of Theorem 5.2.1 above, we may write $A = \ell^{\oplus(G:H)}$, where ℓ is an H -Galois field extension of k for some subgroup $H \subset G$. Let $\Gamma = G \rtimes H$, where the semidirect product is formed with respect to the conjugation action of H on G . Thus there is a surjection $f : \Gamma \rightarrow H$ (given by second projection) with a splitting s (given by second inclusion). Consider the H -Galois cover $Y \rightarrow X$, where $X = \mathbb{P}_k^1$ and $Y = \mathbb{P}_\ell^1$. Let ξ be a k -point of X . By hypothesis, there is a closed point η on $Y = \mathbb{P}_\ell^1$ whose residue field is ℓ and whose decomposition group over ξ is H . So by Theorem 5.1.10, there is a regular connected G -Galois cover $Z \rightarrow Y$ which is totally split over η , such that the composition $Z \rightarrow X$ is Γ -Galois and such that $1 \rtimes H = s(H) \subset \Gamma$ is the decomposition group over ξ of some point $\zeta \in Z$ over η . Viewing G as a quotient of Γ via the multiplication map $m : \Gamma = G \rtimes H \rightarrow G$, we may consider the intermediate G -Galois cover $W \rightarrow X$ (i.e. $W = Z/N$, where $N = \ker m$). It is then straightforward to check that the cover $W \rightarrow X$ satisfies the conclusion of Theorem 5.2.1.

The arithmetic lifting result Theorem 5.2.1 above, and the split embedding problem result Theorem 5.2.10, both generalize Theorem 4.3.1 (that one can realize any finite group as a Galois group over a curve defined over a given large field, with a totally split fibre). In fact, those two generalizations can themselves be simultaneously generalized, by the following joint result of F. Pop and the author, concerning the solvability of a split embedding problem with a prescribed fibre. We first introduce some terminology.

As in Theorem 5.1.10, let X be a geometrically irreducible smooth curve over a field k , let $f : \Gamma \rightarrow G$ be a surjection of finite groups, and let $Y \rightarrow X$ be a G -Galois connected branched cover of smooth curves. Let ξ be an unramified k -point of X , and let η be a closed point of Y over ξ with decomposition group $G_1 \subset G$ and residue field $\ell \supset k$. Let Γ_1 be a subgroup of Γ such that $f(\Gamma_1) = G_1$, and let λ be a Γ_1 -Galois field extension of k that contains ℓ . We say that this data constitutes a *fibred embedding problem* \mathcal{E} for

X . The problem \mathcal{E} is *split* if f has a section s . A *proper solution* to a fibred embedding problem \mathcal{E} as above consists of a smooth connected Γ -Galois branched cover $Z \rightarrow X$ that dominates the G -Galois cover $Y \rightarrow X$, such that there is a closed point ζ of Z over η which has residue field λ and whose decomposition group over ξ is $\Gamma_1 \subset \Gamma$. A solution to \mathcal{E} is *regular* if $Z \rightarrow Y$ is regular (i.e. the algebraic closures of k in the function fields of Y and Z are equal).

Theorem 5.2.3. *Let k be a large field, let X be a geometrically irreducible smooth k -curve, and consider a fibred split embedding problem \mathcal{E} as above, with data $f : \Gamma \rightarrow G$, s , $Y \rightarrow X$, $\xi \in X$, $\eta \in Y$, $G_1 \subset G$, $\lambda \supset \ell \supset k$. Assume that $\Gamma_1 = \text{Gal}(\lambda/k)$ contains $s(G_1)$. Let k' be the algebraic closure of k in the function field of Y , let $X' = X \times_k k'$ and let $E = \text{Gal}(Y/X') \subset G$. Assume that $s(E)$ commutes with $N_1 = \ker(f : \Gamma_1 \rightarrow G_1)$. Then \mathcal{E} has a proper regular solution.*




Figure 5.2.4: *The set-up in the statement of Theorem 5.2.3.*

In other words, given a split embedding problem for a curve over a large field, there is a proper regular solution with a given fibre, assuming appropriate hypothesis (on Γ_1 , E and N_1). Taking the special case $\Gamma_1 = G_1$ in Theorem 5.2.3 (i.e. taking $N_1 = 1$), one recovers Theorem 5.2.10. And taking the special case $G = 1$ in Theorem 5.2.3, one recovers Theorem 5.2.1 above. (Note that the “ G ” in Theorem 5.2.1 corresponds to the “ Γ ” in Theorem 5.2.3.) More generally, taking $E = 1$ in Theorem 5.2.3 (but not necessarily taking G to be trivial), one obtains the result in the case that the given cover $Y \rightarrow X$ is purely arithmetic, i.e. of the form $Y = X \times_k k'$. The result in that case is a generalization of Theorem 5.2.1 — viz. instead of requiring the desired cover $Z \rightarrow X$ in Theorem 5.2.1 to be regular, it can be chosen so that the algebraic closure of k in the function field of Z is a given subfield k' of A that is Galois over k (and also X need not be \mathbb{P}^1). Note that in each of these special cases, the hypothesis on $s(E)$ commuting with N_1 is automatically satisfied, because either E or N_1 is trivial in each case. (On the other hand, the condition $\Gamma_1 \supset s(G_1)$ is still assumed.)

Theorem 5.2.3, like Theorem 5.2.1 above, can in fact be deduced from Theorem 5.1.10, by a strengthening of the proof of Theorem 5.2.1 given above:

Proof of Theorem 5.2.3. The G -Galois cover $Y \rightarrow X$ factors as $Y \rightarrow X' \rightarrow X$, where the E -Galois cover $Y \rightarrow X'$ is regular, and $X' \rightarrow X$ is purely arithmetic (induced by extension of constants from k to k'). Let $\bar{G} = \text{Gal}(k'/k)$; we may then identify $\text{Gal}(X'/X) = G/E$ with \bar{G} . For any field F containing k' , let $X_F = X' \times_{k'} F = X \times_k F$ and let $Y_F = Y \times_{k'} F$. So we may identify $E = \text{Gal}(Y_\ell/X_\ell) = \text{Gal}(Y_\lambda/X_\lambda)$; and \bar{G} is a quotient of $G_1 = \text{Gal}(X_\ell/X)$. Since $Y_\ell = X_\ell \times_{X'} Y$, it follows that $\text{Gal}(Y_\ell/X) = G_1 \times_{\bar{G}} G$ (fibre product of groups); similarly $\text{Gal}(Y_\lambda/X) = \Gamma_1 \times_{\bar{G}} G$, and $Y \rightarrow X$ is the subcover of $Y_\lambda \rightarrow X$ corresponding to the second projection map $G_1 \times_{\bar{G}} G \rightarrow G$.

Let ξ_ℓ be the unique closed point of X_ℓ over $\xi \in X$. Then $\xi_\ell \in X_\ell$ and $\eta \in Y$ each have residue field ℓ and decomposition group G_1 over ξ . So the fibre of $Y_\ell \rightarrow Y$ over η is totally split, with each point having residue field ℓ ; the points of this fibre lie over $\xi_\ell \in X_\ell$ and over $\eta \in Y$; and the local fields of X_ℓ and Y at ξ_ℓ and η (i.e. the fraction fields of the complete local rings) are isomorphic over X' . So at one of the points in this fibre (say η_ℓ), the decomposition group over $\xi \in X$ is equal to the diagonal subgroup $G_1 \times_{G_1} G_1 \subset G_1 \times_{\bar{G}} G$. (At the other points of the fibre, the decomposition group is of the form $\{(g_1, \iota(g_1)) \mid g_1 \in G_1\}$, where ι is an inner automorphism of G_1 .) Similarly, there is a point $\eta_\lambda \in Y_\lambda$ over $\eta_\ell \in Y_\ell$ whose residue field is λ and whose decomposition group over $\xi \in X$ is $\Gamma_1 \times_{G_1} G_1 \subset \Gamma_1 \times_{\bar{G}} G$.




Figure 5.2.5: *The situation in the proof of Theorem 5.2.3.*

Since $E = \ker(G \rightarrow \bar{G})$, every element of $\Gamma_1 \times_{\bar{G}} G$ can uniquely be written as $(\gamma_1, f(\gamma_1)e)$, with $\gamma_1 \in \Gamma_1$ and $e \in E$. Consider the map $\sigma : \Gamma_1 \times_{\bar{G}} G \rightarrow \Gamma_1 \times_{\bar{G}} \Gamma$ given by $\sigma(\gamma_1, f(\gamma_1)e) = (\gamma_1, \gamma_1 s(e))$. Since $s(E)$ commutes with N_1 , direct computation shows that σ is a homomorphism, and hence is a section of $(1, f) : \Gamma_1 \times_{\bar{G}} \Gamma \rightarrow \Gamma_1 \times_{\bar{G}} G$.

We may now apply Theorem 5.1.10 to the surjection $(1, f)$ and its section σ , to the cover $Y_\lambda \rightarrow X$, to the k -point $\xi \in X$, and to the point $\eta_\lambda \in Y_\lambda$ over ξ with decomposition group $\Gamma_1 \times_{G_1} G_1$. The conclusion of that result is that there is a smooth connected

$\Gamma_1 \times_{\bar{G}} \Gamma$ -Galois cover $Z_\lambda \rightarrow X$ that dominates the $\Gamma_1 \times_{\bar{G}} G$ -Galois cover $Y_\lambda \rightarrow X$ with $Z_\lambda \rightarrow Y_\lambda$ regular, together with a point $\zeta_\lambda \in Z_\lambda$ whose decomposition group over ξ is $\sigma(\Gamma_1 \times_{G_1} G_1) = \Gamma_1 \times_{\Gamma_1} \Gamma_1 = \Delta_{\Gamma_1}$, the diagonal of Γ_1 in $\Gamma_1 \times_{\bar{G}} \Gamma$. Let $Z \rightarrow X$ be the intermediate Γ -Galois cover corresponding to the second projection map $\Gamma_1 \times_{\bar{G}} \Gamma \rightarrow \Gamma$, and let $\zeta \in Z$ be the image of $\zeta_\lambda \in Z_\lambda$. Then $Z \rightarrow X$ dominates the G -Galois cover $Y \rightarrow X$; the decomposition group of ζ is $\Gamma_1 \subset \Gamma$ and the residue field is λ ; and ζ lies over $\eta \in Y$. So $Z \rightarrow X$ and the point ζ define a proper solution to the split embedding problem \mathcal{E} . The solution is regular, i.e. $Z \rightarrow Y$ is regular, since the pullback $Z_\lambda \rightarrow Y_\lambda$ is regular. \square

Remark 5.2.6. (a) Theorem 5.2.3 can be regarded as a step toward an “arithmetic Riemann’s Existence Theorem” for covers of curves over a large field. Namely, such a result should classify the branched covers of such a curve, in terms of how they fit together (e.g. with respect to embedding problems), and in terms of their arithmetic and their geometry, including information about decomposition groups and inertia groups (the latter of which Theorem 5.2.3 does not discuss).

(b) In Remark 5.1.11(b), it was asked if Theorem 5.1.10 can be generalized, to allow one to require the decomposition group there to be an arbitrary subgroup of Γ that maps isomorphically onto G_1 under f (rather than being required to take $s(G_1)$ for the decomposition group). If it can, then the above proof of Theorem 5.2.3 could be simplified, and the statement of Theorem 5.2.3 could be strengthened. Namely, the subgroup $\Gamma_1 \subset \Gamma$ could be allowed to be chosen more generally, viz. as any subgroup of Γ whose image under f is G_1 . And the assumption that $s(E)$ commutes with N_1 could also be dropped — since one could then replace the section σ in the above proof by the section (id, s) , while still requiring the decomposition group at ζ_λ to be Δ_{Γ_1} . But on the other hand there might in general be a cohomological obstruction, which would vanish if the containment and commutativity assumptions are retained. \square

Section 5.3. Abhyankar’s Conjecture and embedding problems.

The main theme in this manuscript has been the use of patching methods to prove results in the direction of Riemann’s Existence Theorem for curves that are not necessarily defined over \mathbb{C} . Such a result should classify the unramified covers of such a curve U , and in particular provide an explicit description of the fundamental group of U , as a profinite group.

While the full statement of Riemann’s Existence Theorem is known only for curves over an algebraically closed field of characteristic 0, partial versions have been discussed above. In particular, if one allows arbitrary branching to occur, there is the Geometric Shafarevich Conjecture (Section 5.1); and if one instead takes U to be the complement of a well chosen branch locus and if one restricts attention to a particular class of covers, then there is the Half Riemann Existence Theorem (Section 4.3).

Another way to weaken Riemann’s Existence Theorem is to ask for the set $\pi_A(U)$ of finite Galois groups of unramified covers of U ; or equivalently, the set of finite quotients of $\pi_1(U)$, up to isomorphism. A *finitely generated* profinite group Π is in fact determined by its set of finite quotients [FJ, Proposition 15.4]; and $\pi_1(U)$ is finitely generated (as a profinite group) if the base field has characteristic 0. But in characteristic p , if U is affine, then $\pi_1(U)$ is *not* finitely generated (see below), and $\pi_A(U)$ does not determine $\pi_1(U)$. In this situation, $\pi_1(U)$ remains unknown; but at least $\pi_A(U)$ is known if the base field is algebraically closed. Moreover, $\pi_A(U)$ depends only on the *type* (g, r) of U (where $U = X - S$, with X a smooth connected projective curve of genus $g \geq 0$, and S is a set of $r > 0$ points of X). Namely, this 1957 conjecture of Abhyankar [Ab1] was proven by Raynaud [Ra2] and the author [Ha7] using patching and other methods:

Theorem 5.3.1. (Abhyankar’s Conjecture) (Raynaud, Harbater) *Let k be an algebraically closed field of characteristic $p > 0$, and let U be a smooth connected affine curve over k of type (g, r) . Then a finite group G is in $\pi_A(U)$ if and only if each prime-to- p quotient of G has a generating set of at most $2g + r - 1$ elements.*

Recall that for complex curves U of type (g, r) , a finite group G is in $\pi_A(U)$ if and only if G has a generating set of at most $2g + r - 1$ elements. The same assertion is false in characteristic $p > 0$, e.g. since any affine curve has infinitely many Artin-Schreier covers (cyclic of order p), and hence has Galois groups of the form $(\mathbb{Z}/p\mathbb{Z})^s$ for arbitrarily large s . (This implies the above comment that $\pi_1(U)$ is not finitely generated.) The above theorem can be interpreted as saying that “away from p ”, the complex result carries over; and that every finite group consistent with this principle must occur as a Galois group over U .

In the theorem, the assertion about *every* prime-to- p quotient of G can be replaced by the same assertion about the *maximal* prime-to- p quotient of G — i.e. the group $\bar{G} := G/p(G)$, where $p(G)$ is the subgroup of G generated by the elements of p -power order (or equivalently, by the p -subgroups of G ; or again equivalently, by the Sylow p -subgroups of G).

In the case that $U = \mathbb{A}_k^1$, Theorem 5.3.1 says that $\pi_A(\mathbb{A}_k^1)$ consists precisely of the *quasi- p groups*, viz. the groups G such that $G = p(G)$ (i.e. that are generated by their Sylow p -subgroups). This class of groups includes in particular all p -groups, and all finite simple groups of order divisible by p .

Remark 5.3.2. (a) Prior to Theorem 5.3.1 being proven, Serre had shown a partial result [Se6, Théorème 1]: that if Q is a quasi- p group and if $N \triangleleft Q$ is a solvable normal subgroup of Q such that the (quasi- p) group $\bar{Q} := Q/N$ is a Galois group over \mathbb{A}_k^1 (i.e. $\bar{Q} \in \pi_A(\mathbb{A}_k^1)$), then Q is also a Galois group over \mathbb{A}_k^1 . Due to the solvability assumption, the proof was able to proceed cohomologically, without patching; it relied in particular on the fact that $\pi_1(U)$ is projective (Theorem 5.1.3 above, also due to Serre). Serre’s result [Se6, Thm. 1] implied in particular that Theorem 5.3.1 above is true for *solvable* groups over the affine

line. Serre's proof actually showed more: that if N is a p -group, then a given \bar{Q} -cover $Y \rightarrow \mathbb{A}^1$ can be dominated by a Q -cover (i.e. the corresponding p -embedding problem can be properly solved); but that if N has order prime-to- p , then the embedding problem need not have a proper solution (i.e. the asserted Q -Galois cover of \mathbb{A}^1 cannot necessarily be chosen so as to dominate the given \bar{Q} -Galois cover $Y \rightarrow \mathbb{A}^1$).

(b) More generally, by extending the methods of [Se6], the author showed [Ha12] that if U is any affine variety other than a point, over an arbitrary field of characteristic p , then every finite p -embedding problem for $\pi_1(U)$ has a proper solution. Moreover, this solution can be chosen so as to have prescribed local behavior. For example, if $V \subset U$ is a proper closed subset, then the proper solution over U can be chosen so that it restricts to a given weak solution over V . (Cf. Theorem 5.2.3 above, for such fibred embedding problems in a related but somewhat different context.) And if U is a curve, then the proper solution can be chosen so as to restrict to given weak solutions over the fraction fields of the complete local rings at finitely many points.

Proof sketch of Theorem 5.3.1. In the case $U = \mathbb{A}^1$, the theorem was proven by Raynaud [Ra2], using in particular rigid patching methods. The proof proceeded by induction on the order of G , and considered three cases. In Case 1, the group G is assumed to have a non-trivial normal p -subgroup N ; and using Serre's result that embedding problems for $\pi_1(U)$ with p -group kernel can be properly solved (Remark (a) above), the desired conclusion for G follows from the corresponding fact for G/N . When not in Case 1, one picks a Sylow p -subgroup P , and considers all the quasi- p subgroups $Q \subset G$ such that $Q \cap P$ is a Sylow p -subgroup of Q . Case 2 is the situation in which these Q 's generate G . In this case, by induction each of the Q 's is a Galois group over \mathbb{A}^1 ; and using rigid patching it follows that G is also. (Or one could use formal patching for this step, viz. Theorem 3.2.8; see e.g. [HS, Theorem 6].) Case 3 is the remaining case, where Cases 1 and 2 do not apply. Then, one builds a G -Galois branched cover of the line in mixed characteristic having p -power inertia groups. The closed fibre of the semi-stable model is a reducible curve that maps down to a tree of projective lines in characteristic p . Using a careful combinatorial analysis of the situation, it turns out that over one of the terminal components of the tree (a copy of the projective line), one finds an irreducible G -Galois cover that is branched at just one point — and hence is an étale cover of the affine line, as desired. Moreover, by adjusting the cover, we may assume that the inertia groups over infinity (of the corresponding branched cover of \mathbb{P}_k^1) are the Sylow p -subgroups of Q . (Namely, by Abhyankar's Lemma, after pulling back by a Kummer cover $y^n - x$, we may assume that the inertia groups over infinity are p -groups. We may then enlarge this inertia to become Sylow, using Remark 5.1.6(d) above.)

The general case of the theorem was proven in [Ha7], by using the above case of the affine line, together with formal patching and embedding problems. (See also the simplified

presentation in [Ha13], where more is shown.) For the proof, one first recalls that the result was shown by Grothendieck [Gr5, XIII, Cor. 2.12] in the case that the group is of order prime to p . Using this together with formal patching (Theorem 3.2.8), it is possible to reduce to the key case that $U = \mathbb{A}_k^1 - \{0\}$, where $G/p(G)$ is cyclic of prime-to- p order. (For that reduction, one patches a prime-to- p cover of the original curve together with a cyclic-by- p cover of $\mathbb{A}_k^1 - \{0\}$, to obtain a cover of the original curve with the desired group.) Once in this case, by group theory one can find a prime-to- p cyclic subgroup $\bar{G} \subset G$ that normalizes a Sylow p -subgroup P of G and that surjects onto $G/p(G)$. Here G is a quotient of the semi-direct product $\Gamma := p(G) \rtimes \bar{G}$ (formed with respect to the conjugation action of \bar{G} on $p(G)$); so replacing G by Γ we may assume that $G = p(G) \rtimes \bar{G}$ with $\bar{G} \approx G/p(G)$. Letting $n = |\bar{G}|$, there is a \bar{G} -Galois étale cover $V \rightarrow U_K$ given by $y^n = x$, where $K = k((t))$ and $U_K = \mathbb{A}_K^1 - \{0\}$. Using the proper solvability of p -embedding problems with prescribed local behavior (Remark 5.3.2(b) above), one can obtain a $P \rtimes \bar{G}$ -Galois étale cover $\tilde{V} \rightarrow U_K$ whose behavior over one of the (unramified) K -points ξ_K of U_K can be given in advance. Specifically, one first considers a $p(G)$ -Galois étale cover $W \rightarrow \mathbb{A}_k^1$ (given by the first case of the result, with Sylow p -subgroups as inertia over ∞), and restricts to the local field at a ramification point with inertia group P (this being a P -Galois field extension of the local field $K = k((t))$ at ∞ on \mathbb{P}_k^1). It is this P -Galois extension of K that one uses for the prescribed local behavior over the K -point ξ_K , in applying the p -embedding result. As a consequence, the $P \rtimes \bar{G}$ -Galois cover $\tilde{V} \rightarrow U_K$ (near ξ_K) has local compatibility with W (near ∞). This compatibility makes it possible for the two covers \tilde{V} and W to be patched using Theorem 3.2.8 or 3.2.12 (after blowing up; see Examples 3.2.11, 3.2.13, and 4.2.4). As a result we obtain a G -Galois cover of U_K (viz. the generic fibre of a cover of $U_{k[[t]]}$). This cover is irreducible because the Galois groups of \tilde{V} and W (viz. $P \rtimes \bar{G}$ and $p(G)$) together generate G . Since k is algebraically closed, one may specialize from K to k (as in Corollary 3.3.5) to obtain the desired cover of U . \square

Remark 5.3.3. (a) The proof of Theorem 5.3.1 actually shows more, concerning inertia groups: Write $U = X - S$ for a smooth connected projective k -curve X and finite set S , and let $\xi \in S$. Then in the situation of the theorem, the G -Galois étale cover of U may be chosen so that the corresponding branched cover of X is tamely ramified away from ξ . (This was referred to as the “Strong Abhyankar Conjecture” in [Ha7], where it is proven.) Note that it is necessary, in general, to allow at least one wildly ramified point. Namely, if G cannot itself be generated by $2g + r - 1$ elements or fewer, then G is not a Galois group of a tamely ramified cover of X that is étale over U , because the tame fundamental group $\pi_1^\dagger(U)$ is a quotient of the free profinite group on $2g + r - 1$ generators [Gr5, XIII, Cor. 2.12].

(b) It would be even more desirable, along the lines of a possible Riemann’s Existence Theorem over k , to determine precisely which subgroups of G can be the inertia groups

over the points of S , for a G -Galois cover of a given U (with S as in Remark (a) above). This problem is open, however, even in the case that $U = \mathbb{A}_k^1$. In that case, the unique branch point ∞ must be wildly ramified, since there are no non-trivial tamely ramified covers of \mathbb{A}^1 (by [Gr5, XIII, Cor. 2.12]). By the general theory of extensions of discrete valuation rings [Se5], any inertia group of a branched cover of a k -curve is of the form $I = P \rtimes C$, where P is a p -group (not necessarily Sylow in the Galois group) and C is cyclic of order prime to p . As noted above, it is known [Ra2] that if P is a Sylow p -subgroup of a quasi- p group Q , then there is a Q -Galois étale cover of \mathbb{A}^1 such that P is an inertia group over infinity (and this fact was used in the proof of the general case of Theorem 5.3.1, in order to be able to patch together the $P \rtimes \tilde{G}$ -cover with the $p(G)$ -cover). More generally, for *any* subgroup $I \subset Q$ of the form $P \rtimes C$, a necessary condition for I to be an inertia group over ∞ for a Q -Galois étale cover $Y \rightarrow \mathbb{A}_k^1$ is that the conjugates of P generate Q . (For if not, they generate a normal subgroup $N \triangleleft Q$ such that $Y/N \rightarrow \mathbb{A}^1$ is a non-trivial tamely ramified cover; but \mathbb{A}^1 has no such covers, and this is a contradiction.) Abhyankar has conjectured that the converse holds (i.e. that every $I \subset Q$ satisfying the necessary condition will be an inertia group over infinity, for some Q -Galois étale cover of the line). This remains open, although some partial results in this direction have been found by R. Pries and I. Bouw [Pr2], [BP].

(c) The results of Sections 3.3 and 4.3 suggest that Abhyankar's Conjecture may hold for affine curves over *large* fields of characteristic p , not just over algebraically closed fields of characteristic p — since patching is possible over such fields, and various Galois realization results can be extended to these fields. But this generalization of Abhyankar's Conjecture remains open. The difficulty is that in the proof of Case 3 of Theorem 5.3.1 for \mathbb{A}_k^1 , one considers a branched cover of \mathbb{A}_R^1 , where R is a complete discrete valuation ring of mixed characteristic with residue field k . For such a cover, the semi-stable model might be defined only over a finite extension R' of R (and not over R itself); and the residue field of R' could be strictly larger than k . Thus the construction in the proof might yield only a Galois cover of the k' -line, for some finite extension k' of k .

(d) As noted above before Theorem 5.3.1, for an affine k -curve U , the fundamental group $\pi_1(U)$ is not finitely generated (as a profinite group), and is therefore not determined by $\pi_A(U)$. And indeed, the structure of $\pi_1(U)$ is unknown, even for $U = \mathbb{A}_k^1$ (although Theorem 5.3.4 below gives some information about how the finite quotients of π_1 “fit together”). In fact, it is easy to see that $\pi_1(\mathbb{A}_k^1)$ depends on the cardinality of the algebraically closed field k of characteristic p ; viz. the p -rank of π_1 is equal to this cardinality (using Artin-Schreier extensions). Moreover, Tamagawa has shown [Tm2] that if k, k' are non-isomorphic *countable* algebraically closed fields of characteristic p with $k = \overline{\mathbb{F}}_p$, then $\pi_1(\mathbb{A}_k^1)$ and $\pi_1(\mathbb{A}_{k'}^1)$ are non-isomorphic as profinite groups. (It is unknown whether this remains true even if k is chosen strictly larger than $\overline{\mathbb{F}}_p$.) Tamagawa also showed in [Tm2]

that if $k = \overline{\mathbb{F}}_p$, then two open subsets of \mathbb{A}_k^1 have isomorphic π_1 's if and only if they are isomorphic as schemes. More generally, given arbitrary affine curves U, U' over algebraically closed fields k, k' of non-zero characteristic, it is an open question whether the condition $\pi_1(U) \approx \pi_1(U')$ implies that $k \approx k'$ and $U \approx U'$. This question, which can be regarded as an algebraically closed analog of Grothendieck's anabelian conjecture for affine curves over finitely generated fields [Gr6], was essentially posed by the author in [Ha8, Question 1.9]; and the results in [Tm2] (which relied on the anabelian conjecture in the finitely generated case [Tm1]) can be regarded as the first real progress in this direction.

(e) Theorem 5.3.1 holds only for *affine* curves, and is false for projective curves. Namely, if X is a smooth projective k -curve of genus g , then $\pi_1(X)$ is a quotient of the fundamental group of a smooth projective complex curve of genus g (which has generators $a_1, b_1, \dots, a_g, b_g$ subject to the single relation $\prod [a_i, b_i] = 1$ [Gr5, XIII, Cor. 2.12]). So if $g > 0$ and if Q is a quasi- p group whose minimal generating set has more than $2g$ generators, then Q is not in $\pi_A(X)$. Also, the p -rank of a smooth projective k -curve of genus g is at most g , and so $(\mathbb{Z}/p\mathbb{Z})^{g+1}$ is also not in $\pi_A(X)$. But both Q and $(\mathbb{Z}/p\mathbb{Z})^{g+1}$ trivially have the property that every prime-to- p quotient has at most $2g - 1$ generators (since the only prime-to- p quotient of either group is the trivial group). So both of these groups provide counterexamples to Theorem 1 over the projective curve X . (In the case of genus 0, we have $X = \mathbb{P}_k^1$, and $\pi_1(X)$ is trivial.)

(f) Another difference between the affine and projective cases concerns the relationship between π_A and π_1 . As discussed in Remark (c) above, Theorem 5.3.1 gives π_A but not π_1 for an affine curve, the difficulty being that π_A does not determine π_1 because π_1 of an affine curve is not a finitely generated profinite group. On the other hand, if X is a projective curve, then $\pi_1(X)$ is a finitely generated profinite group, and so it is determined by $\pi_A(X)$. Unfortunately, unlike the situation for affine curves, $\pi_A(X)$ is unknown when X is projective of genus > 1 (cf. Remark (d)), and so this does not provide a way of finding $\pi_1(X)$ in this case. A similar situation holds for the tame fundamental group $\pi_1^t(U)$, where $U = X - S$ is an affine curve (and where the tame fundamental group classifies covers of X that are unramified over U , and at most tamely ramified over S). Namely, this group is also a finitely generated profinite group, and is a quotient of the corresponding fundamental group of a complex curve. But the structure of this group, and the set $\pi_A^t(U)$ of its finite quotients, are both unknown, even for $\mathbb{P}_k^1 - \{0, 1, \infty\}$. (Note that $\pi_A^t(\mathbb{P}_k^1 - \{0, 1, \infty\})$ is strictly smaller than the set of Galois groups of covers of $\mathbb{P}_\mathbb{C}^1 - \{0, 1, \infty\}$ with prime-to- p inertia, because tamely ramified covers of \mathbb{P}_k^1 with given degree and inertia groups will generally have lower p -rank than the corresponding covers of $\mathbb{P}_\mathbb{C}^1$ — and hence will have fewer unramified p -covers.) On the other hand, partial information about the structure of $\pi_A(X)$ and $\pi_A^t(U)$ has been found by formal and rigid patching methods ([St1], [HS1], [Sa1]) and other approaches ([St2], [PS]). \square

Following the proof of Theorem 5.3.1, Pop used similar methods to prove a stronger version of the result, in terms of embedding problems:

Theorem 5.3.4. (Pop [Po3]) *Let k be an algebraically closed field of characteristic $p > 0$, and let U be a smooth connected affine curve over k . Then every finite embedding problem for $\pi_1(U)$ that has quasi- p kernel is properly solvable.*

That is, given a finite group Γ and a quasi- p normal subgroup N of Γ , and given a Galois étale cover $V \rightarrow U$ with group $G := \Gamma/N$, there is a Galois étale cover $W \rightarrow U$ with group Γ that dominates V . Theorem 5.3.1 follows from the special case of Theorem 5.3.4 in which $N = p(\Gamma)$. But Pop’s stronger result provides information about the structure of $\pi_1(U)$ (i.e. how the covers “fit together in towers”), and not just about $\pi_A(U)$ (i.e. what covers can exist in isolation).

Actually, Theorem 5.3.4 was stated in [Po3] only for *split* embedding problems with quasi- p kernel. But one can easily deduce the general case from that one, proceeding as in the proof of Theorem 5.1.1, via Theorem 5.1.3 there. See also [CL] and [Sa2], i.e. Chapters 15 and 16 in [BLoR], for more about the proofs of Theorems 5.3.1 and 5.3.4, presented from a rigid point of view. (More about the proof of Theorem 5.3.1 can be found in [Ha9, §3].)

Remark 5.3.5. (a) Theorem 5.3.4 can be generalized from étale covers to tamely ramified covers [Ha13, Theorem 4.4]. Namely, with $G = \Gamma/N$ as above, suppose that $V \rightarrow U$ is a tamely ramified G -Galois cover of U with branch locus $B \subset U$. Then there is a Γ -Galois cover $W \rightarrow U$ that dominates V , and is tamely ramified over B and étale elsewhere over U . (Note that no assertions are made here, or in Theorem 5.3.4, about the behavior over points in the complement of U in its smooth completion.)

(b) In Theorem 5.3.4, for an embedding problem $\mathcal{E} = (\alpha : \pi_1(U) \rightarrow G, f : \Gamma \rightarrow G)$, one cannot replace the assumption that $\ker f$ is quasi- p by the assumption that Γ is quasi- p . This follows from Remark 5.3.2(a), concerning Serre’s results in [Se6].

(c) Theorems 5.3.1 and 5.3.4 both deal only with *finite* Galois groups and embedding problems. It is unknown which *infinite* quasi- p profinite groups can arise as Galois groups, and which embedding problems with *infinite* quasi- p kernel have proper solutions. For example, let G be the free product $\mathbb{Z}_p * \mathbb{Z}_p$ (in the category of profinite groups). This is an infinite quasi- p group, and so every finite quotient of G is a Galois group over \mathbb{A}^1 . But it is unknown whether G itself is a Galois group over \mathbb{A}^1 (or equivalently, whether G is a quotient of $\pi_1(\mathbb{A}^1)$). \square

Theorem 5.3.4 raises the question of which finite embedding problems for $\pi_1(U)$ are properly solvable, where U is an affine variety in characteristic p — and in particular, whether every finite embedding problem for U with a quasi- p kernel is properly solvable. For example, one can ask this for affine varieties U of finite type over an algebraically closed

field k of characteristic p , i.e. whether Pop’s result remains true in higher dimensions.

Abhyankar had previously posed a weaker form of this question as a conjecture, paralleling his conjecture for curves (i.e. Theorem 5.3.1). Namely, in [Ab3], he proposed that if U is the complement of a normal crossing divisor D in \mathbb{P}_k^n (where k is algebraically closed of characteristic p), then $G \in \pi_A(U)$ if and only if $G/p(G) \in \pi_A(U_{\mathbb{C}})$, where $U_{\mathbb{C}}$ is an “analogous complex space”. That is, if D has irreducible components D_1, \dots, D_r of degrees d_1, \dots, d_r , then one takes $U_{\mathbb{C}}$ to be the complement in $\mathbb{P}_{\mathbb{C}}^n$ of a normal crossing divisor consisting of r components of degrees d_1, \dots, d_r . It is known (by [Za1], [Za3], [Fu2]) that $\pi_1(U_{\mathbb{C}})$ is the abelian group $A(d_1, \dots, d_r)$ on generators g_1, \dots, g_r satisfying $\sum d_i g_i = 1$ (writing additively). It is also known (by [Ab2], [Fu2]) that the prime-to- p quotients of $\pi_A(U)$ are precisely the prime-to- p quotients of $A(d_1, \dots, d_r)$. Thus Abhyankar’s conjecture in [Ab3] is a special case of a more general conjecture that $G \in \pi_A(U) \Leftrightarrow G/p(G) \in \pi_A(U)$ for any affine k -variety U of finite type. This in turn would follow from an affirmative answer to the question asked in the previous paragraph.

Abhyankar also posed a local version of this conjecture in [Ab3], viz. that if $U = \text{Spec } k[[x_1, \dots, x_n]][(x_1 \cdots x_r)^{-1}]$ (where $n > 1$ and $1 \leq r \leq n$), then a finite group G is in $\pi_A(U)$ if and only if $G/p(G)$ is in $\pi_A(U_{\mathbb{C}})$; here $U_{\mathbb{C}} = \text{Spec } \mathbb{C}[[x_1, \dots, x_n]][(x_1 \cdots x_r)^{-1}]$. (Note that this is clearly false if $r = 0$, since then $\pi_A(U)$ is trivial, by Hensel’s Lemma. It is also false if $n = 1$, since in that case the only quasi- p groups in $\pi_A(U)$ are p -groups, by the structure of Galois groups over complete discrete valuation fields [Se5].) Now $\pi_A(U_{\mathbb{C}})$ consists of the finite abelian groups on r generators, and the prime-to- p quotients of π_A are the finite abelian prime-to- p groups on r generators (by Abhyankar’s Lemma; cf. [HP, Section 3]). So this conjecture is again equivalent to asserting that $G \in \pi_A(U) \Leftrightarrow G/p(G) \in \pi_A(U)$.

Abhyankar’s higher dimensional global conjecture is easily seen to hold in some special cases, e.g. if D is a union of one or two hyperplanes (since it then reduces immediately to Theorem 5.3.1). Using patching, one can show that the higher dimensional local conjecture holds for $r = 1$ [HS2]. But perhaps surprisingly, both the global and local conjectures fail in general, because some groups that satisfy the conditions of the conjectures nevertheless fail to arise as Galois groups of covers. In particular, the global conjecture fails for \mathbb{P}_k^2 minus three lines crossing normally, and the local conjecture fails for $n = r = 2$ [HP]. Thus not every embedding problem with quasi- p kernel can be solved for $\pi_1(U)$, in general.

Remark 5.3.6. The main reason that the higher dimensional conjecture fails in general is that the group-theoretic reduction in the proof of the general case of Theorem 5.3.1 does not work in the more general situation. That is, it is possible that $G/p(G) \in \pi_A(U)$ but that G is not a quotient of a group \tilde{G} of the form $\tilde{G} = p(G) \rtimes \bar{G}$, with \bar{G} a prime-to- p group in $\pi_A(U)$. (Cf. the group-theoretic examples of Guralnick in [HP].) Moreover, even if there is such a \tilde{G} , it might not be possible to choose it such that \bar{G} normalizes a Sylow

p -subgroup of $p(G)$ (or equivalently, of G), as was done in the proof of Theorem 5.3.1. And in fact, a condition of the above type is *necessary* in order that $G \in \pi_A(U)$, if U is the complement of $x_1 \cdots x_i = 0$ (in either the local or global case; cf. [HP]).

This suggests that a group G should lie in $\pi_A(U)$ if it satisfies these additional conditions, as well as the condition that $G/p(G) \in \pi_A(U)$. One might wish to parallel the proof of the general case of Theorem 1, using higher dimensional patching (Theorem 3.2.12) together with the result on embedding problems with p -group kernel ([Ha12], which holds in arbitrary dimension). Unfortunately, there is another difficulty: The strategy for curves used that for every quasi- p group Q there is a Q -Galois étale cover of \mathbb{A}_k^1 such that the fibre over infinity (of the corresponding branched cover of \mathbb{P}_k^1) consists of a disjoint union of points whose inertia groups are Sylow p -subgroups of Q (cf. Case 1 of the proof of Theorem 5.3.1). But the higher dimensional analog of this is false; in fact, for $n > 1$, every branched cover of \mathbb{P}_k^n which is étale over \mathbb{A}_k^n must have the property that its fibre over the hyperplane at infinity is connected. This then interferes with the desired patching, on the overlap. \square

One can also consider birational variants of the above questions, in studying the absolute Galois groups of $k_n := k(x_1, \dots, x_n)$ and $k_n^* := k((x_1, \dots, x_n))$. Here k is an algebraically closed field of characteristic $p \geq 0$; $n > 1$; and $k((x_1, \dots, x_n))$ denotes the fraction field of $k[[x_1, \dots, x_n]]$. Of course every finite group is a Galois group over k_n , since this is true for $k(x_1)$ (see Corollary 3.3.5) and one may base-change to k_n . Also, every finite group is a Galois group over k_n^* , by Example 3.3.2(c). But this does not determine the structure of the absolute Galois groups of k_n and k_n^* .

In the one-dimensional analog, the absolute Galois group of $k(x)$ is a free profinite group (of rank equal to the cardinality of k), by the geometric case of Shafarevich's Conjecture (Section 5.1). But for $n > 1$, the absolute Galois group of k_n has cohomological dimension > 1 [Se4, II, 4.1, Proposition 11], and so is not projective [Se4, I, 3.4, Proposition 16]. That is, not every finite split embedding problem for G_{k_n} has a weak solution; and therefore G_{k_n} is not free.

This can also be seen explicitly as in the following argument, which also applies to k_n^* :

Proposition 5.3.7. *Let k be an algebraically closed field of characteristic $p \geq 0$, let $n > 1$, and let $K = k_n$ or k_n^* as above. Then not every finite embedding problem for the absolute Galois group G_K is weakly solvable. Equivalently, there is a surjection $G \rightarrow A$ of finite groups, and an A -Galois field extension K' of K , such that K' is not contained in any H -Galois field extension L of K for any $H \subset G$.*

Proof. First suppose that $\text{char } k \neq 2$. Let G be the quaternion group of order 8, and let A be the quotient of G by its center $Z = \{\pm 1\}$. Thus $A = G/Z \approx C_2^2$, say with generators a, b which are commuting involutions. Consider the surjection $G_K \rightarrow A$ corresponding to the A -Galois field extension K' given by $u^2 = x_1, v^2 = x_2$. Suppose that this field extension is

contained in an H -Galois extension L/K as in the statement of the proposition. Then A is a quotient of H . But no proper subgroup of G surjects onto A ; so actually $H = G$.

Let $F = k((x_1)) \cdots ((x_n))$, and let F' [resp. E] be the compositum of F and K' [resp. F and L] in some algebraic closure of F . Thus E is a Galois field extension of F , and its Galois group G' is a subgroup of G . Moreover E contains F' , which is an A -Galois field extension of F (being given by $u^2 = x_1, v^2 = x_2$). Thus A is a quotient of G' , and hence $G' = G$. But the maximal prime-to- p quotient of the absolute Galois group G_F is abelian [HP, Prop. 2.4], and so G cannot be a Galois group over F (using that $p \neq 2$). This is a contradiction, proving the result in this case.

On the other hand, if $\text{char } k = 2$, then one can replace the quaternion group in the above argument by a similar group of order prime to 2. Namely, let ℓ be any odd prime. Then there is a group G of order ℓ^3 whose center Z is cyclic of order ℓ ; such that $G/Z \approx C_\ell^2$; and such that no proper subgroup of G surjects onto G/Z . (See [As, 23.13]; such a group is called an *extraspecial* group of order ℓ^3 .) The proof then proceeds as before. \square

Remark 5.3.8. The above proof also applies to the field $K = k((x_1, \dots, x_n))(y)$, by using the extension $u^2 = x_1, v^2 = y$. So its absolute Galois group G_K is not projective, and hence not free. (This can alternatively be seen by using [Se4, II, 4.1, Proposition 11]). Note that this field K has the property that every finite group is a Galois group over K (by Theorem 3.3.1), even though G_K is not free or even projective. In fact if $n = 1$, then every finite *split* embedding problem has a proper solution (by Theorem 5.1.9). Thus in this case, once a finite embedding problem has a weak solution, it automatically has a proper solution. In this sense, the absolute Galois group of $k((x))(y)$ is “as close as possible to being free” without being projective. \square

Motivated by the above proposition and remark, it would be desirable to know whether the absolute Galois groups of $k_n := k(x_1, \dots, x_n)$ and $k_n^* := k((x_1, \dots, x_n))$ are “as close as possible to being free” without being projective. (Here k is still algebraically closed and $n > 1$.) In other words, does every finite split embedding problem for G_{k_n} or $G_{k_n^*}$ have a proper solution? The former case can be regarded as a birational analog of the question asked previously concerning quasi- p embedding problems in the higher dimensional Abhyankar Conjecture; it can also be considered a weak version of a higher dimensional geometric Shafarevich Conjecture. In this case, the question remains open, even for $\mathbb{C}(x, y)$. In the latter case, the answer is affirmative for $\mathbb{C}((x, y))$, as the following result shows. The proof follows a strategy from [HS2], viz. blowing up $\text{Spec } \mathbb{C}[[x, y]]$ at the closed point to obtain a more global object, and then patching (here using Theorem 3.2.12).

Theorem 5.3.9. *Every finite split embedding problem over $\mathbb{C}((x, y))$ has a proper solution.*

Proof. Let L be a finite Galois extension of $\mathbb{C}((x, y))$, with group G , and let Γ be a semi-

direct product $N \rtimes G$ for some finite group N . Let $R = \mathbb{C}[[x, y]]$ and let S be the integral closure of R in L , and write $X^* := \text{Spec } R$ and $Z^* := \text{Spec } S$. We want to show that there is an irreducible normal Γ -Galois branched cover $W^* \rightarrow X^*$ that dominates the G -Galois branched cover $Z^* \rightarrow X^*$.

Case 1: S/R is ramified only over $(x = 0)$.

Let n be the ramification index of $Z^* \rightarrow X^*$ over the generic point of $(x = 0)$, and consider the normalized pullback of $Z^* \rightarrow X^*$ via $\text{Spec } R[z]/(z^n - x) \rightarrow X^*$. By Abhyankar's Lemma and Purity of Branch Locus, the resulting cover of $\text{Spec } R[z]/(z^n - x) = \text{Spec } \mathbb{C}[[z, y]]$ is unramified and hence trivial. Thus $S \approx R[z]/(z^n - x)$, and G is cyclic of order n .

Now consider the projective y -line over $\mathbb{C}((x))$, and the G -Galois cover of this line $Z^\circ \rightarrow \mathbb{P}_{\mathbb{C}((x))}^1$ that is given by the constant extension $z^n = x$. Applying Pop's Theorem 5.1.10 to the split embedding problem given by this cover and the group homomorphism $\Gamma \rightarrow G$, we obtain a regular irreducible (hence geometrically irreducible) Γ -Galois cover $W^\circ \rightarrow \mathbb{P}_{\mathbb{C}((x))}^1$ that dominates $Z^\circ \rightarrow \mathbb{P}_{\mathbb{C}((x))}^1$ and is such that $W^\circ \rightarrow Z^\circ$ is totally split over $y = \infty$. Consider the normalization W of $\mathbb{P}_{\mathbb{C}[[x]]}^1$ in W° ; this is a Γ -Galois cover of $\mathbb{P}_{\mathbb{C}[[x]]}^1$ that dominates Z , the normalization of $\mathbb{P}_{\mathbb{C}[[x]]}^1$ in Z° . The branch locus of $W \rightarrow \mathbb{P}_{\mathbb{C}[[x]]}^1$ consists of finitely many irreducible components. After a change of variables $y' = x^m y$ on $\mathbb{P}_{\mathbb{C}((x))}^1$, we may assume that every branch component passes through the closed point (x, y) , and that no branch component other than (x) passes through any other point on the closed fibre of $\mathbb{P}_{\mathbb{C}[[x]]}^1$. Again using Abhyankar's Lemma and Purity of Branch Locus, we conclude that the restriction of W over $\mathbb{C}[y^{-1}][[x]]$ is a disjoint union of components given by $w^N = x$ for some multiple N of n , with each reduced component of the closed fibre of W being a complex line. Since $W^\circ \rightarrow Z^\circ$ is split over $y = \infty$, it follows that $N = n$. Thus the pullback of $W \rightarrow Z$ over $\mathbb{C}[y^{-1}][[x]]$ is a trivial cover.

Since the general fibre of $W \rightarrow \mathbb{P}_{\mathbb{C}[[x]]}^1$ is geometrically irreducible, the closed fibre is connected, by Zariski's Connectedness Theorem [Hrt2, III, Cor. 11.3]. So by the previous paragraph, the components of the closed fibre of W all meet at a single point over $(x = y = 0)$. So the pullback W^* of $W \rightarrow \mathbb{P}_{\mathbb{C}[[x]]}^1$ over $\text{Spec } \mathbb{C}[[x, y]]$ is connected; and since W is normal, it follows that W^* is also normal and hence is irreducible. So $W^* \rightarrow \text{Spec } \mathbb{C}[[x, y]]$ is an irreducible Γ -Galois cover. Moreover W^*/N is isomorphic to $\text{Spec } S$ over $\text{Spec } R$, since each is given by $z^n = x$. So it is a proper solution to the given embedding problem.

Note that in this case, the proof shows more: that G is the cyclic group C_n , and that over $\mathbb{C}((y))[[x]]$, the pullback of $W^* \rightarrow Z^*$ is trivial (since the same is true over $\mathbb{C}[y^{-1}][[x]]$).

Case 2: General case.

Let B be the branch locus of $Z^* \rightarrow X^*$, and let C be the tangent cone to B at the closed point (x, y) . Thus C is a union of finitely many "lines" ($ax + by$) through (x, y) in X^* . After a change of variables of the form $y' = y - cx$, we may assume that C does not

contain the locus of $(y = 0)$.

Let \tilde{X} be the blow-up of X^* at the closed point (x, y) . Let E be the exceptional divisor; this is a copy of $\mathbb{P}_{\mathbb{C}}^1$, with parameter $t = y/x$. Let $\tau \in T$ be the closed point $(x = y = t = 0)$; this is where E meets the proper transform of $(y = 0)$. Let $\tilde{Z} \rightarrow \tilde{X}$ be the normalized pullback of $Z^* \rightarrow X^*$. By the previous paragraph, this is unramified in a neighborhood of τ except possibly along E . So over the complete local ring $\hat{\mathcal{O}}_{\tilde{X}, \tau} = \mathbb{C}[[x, t]]$ of τ in \tilde{X} , the pullback $\tilde{Z}^* \rightarrow \tilde{X}^* := \text{Spec } \hat{\mathcal{O}}_{\tilde{X}, \tau}$ of $\tilde{Z} \rightarrow \tilde{X}$ is ramified only over $(x = 0)$. (See Figures 5.3.10 and 5.3.11 below.)

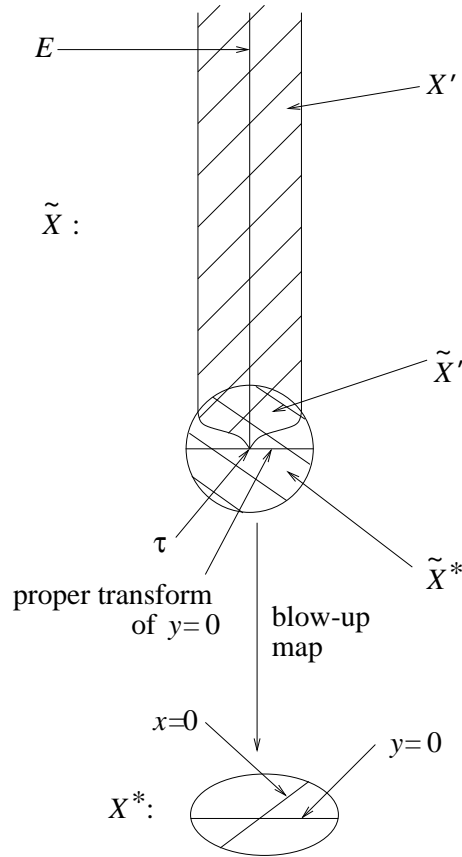


Figure 5.3.10: *Picture of the situation in Case 2 of the proof of Theorem 5.3.9. The space X^* , shown as a disc, is blown up, producing \tilde{X} , with an exceptional divisor E (which meets the proper transform of $y = 0$ at the point τ). The proof proceeds by building the desired cover over formal patches: X' , the completion along $E - \{\tau\}$; and \tilde{X}^* , the completion at τ . These two patches are shaded above, with the doubly shaded region \tilde{X}' being the “overlap”.*

Let \tilde{Z}_0^* be a connected component of \tilde{Z}^* . Thus $\tilde{Z}_0^* \rightarrow \tilde{X}^*$ is Galois with group $G_0 \subset G$, and $\tilde{Z}^* = \text{Ind}_{G_0}^G \tilde{Z}_0^*$. Let $\Gamma_0 \subset \Gamma$ be the subgroup generated by N and G_0 (identifying N with $N \rtimes 1 \subset \Gamma$, and G with $1 \rtimes G \subset \Gamma$). Thus $\Gamma_0 = N \rtimes G_0$. By Case 1, there is a

regular irreducible normal Γ_0 -Galois cover $\tilde{W}_0^* \rightarrow \tilde{X}^*$ that dominates \tilde{Z}_0^* , and such that the pullback of $\tilde{W}_0^* \rightarrow \tilde{Z}_0^*$ over $\tilde{X}' = \text{Spec } \mathbb{C}((t))[[x]]$ is trivial. That is, $\tilde{W}' := \tilde{W}_0^* \times_{\tilde{X}^*} \tilde{X}'$ is the trivial N -Galois cover of $\tilde{Z}' := \tilde{Z}_0^* \times_{\tilde{X}^*} \tilde{X}'$, and the Γ_0 -Galois cover $\tilde{W}' \rightarrow \tilde{X}'$ is just $\text{Ind}_{G_0}^{\Gamma_0} \tilde{Z}'$. Thus the Γ -Galois cover $\tilde{W}^* := \text{Ind}_{\Gamma_0}^{\Gamma} \tilde{W}' \rightarrow \tilde{X}^*$ has the property that its pullback $\tilde{W}' := \tilde{W}^* \times_{\tilde{X}^*} \tilde{X}'$ is just $\text{Ind}_{G_0}^{\Gamma} \tilde{Z}' = \text{Ind}_G^{\Gamma} \tilde{Z}'$, where $\tilde{Z}' = \text{Ind}_{G_0}^G \tilde{Z}'_0$ is the pullback $\tilde{Z}^* \times_{\tilde{X}^*} \tilde{X}'$.

Let $U = E - \{\tau\}$, and let X' be the completion of \tilde{X} along U ; i.e. $X' = \text{Spec } \mathbb{C}[s][[y]]$, where $s = x/y = 1/t$. Let $Z' = \tilde{Z} \times_{\tilde{X}} X'$, and let $W' = \text{Ind}_G^{\Gamma} Z'$. Thus the pullback $Z' \times_{X'} \tilde{X}'$ can be identified with $\tilde{Z}' = \text{Ind}_{G_0}^G \tilde{Z}'_0$ as G -Galois covers of \tilde{X}' ; and the pullback $W' \times_{X'} \tilde{X}'$ can be identified with $\tilde{W}' = \text{Ind}_{G_0}^{\Gamma} \tilde{Z}'_0$, as Γ_0 -Galois covers of \tilde{X}' .




Figure 5.3.11: Diagram illustrating the patching situation in Case 2 of the proof of Theorem 5.3.9. In order to construct the Γ -Galois cover $\tilde{W} \rightarrow \tilde{X}$, the restrictions $W' \rightarrow X'$ and $\tilde{W}^* \rightarrow \tilde{X}^*$ and first constructed, so as to induce the same “overlap” $\tilde{W}' \rightarrow \tilde{X}'$. Formal patching is then used to obtain \tilde{W} .

Now apply the formal patching result Theorem 3.2.12, with $A = R$, $V = \tilde{V} = \tilde{X}$, $f =$ identity, and the finite set of closed points of V being just $\{\tau\}$. Using the equivalence of categories for covers, we conclude that there is a unique Γ -Galois cover $\tilde{W} \rightarrow \tilde{X}$ whose pullbacks to \tilde{X}^* and to X' are given respectively by $\tilde{W}^* = \text{Ind}_{\Gamma_0}^{\Gamma} \tilde{W}_0^* \rightarrow \tilde{X}^*$ and $W' \rightarrow X'$, compatibly with the above identification over \tilde{X}' with $\tilde{W}' = \text{Ind}_{G_0}^{\Gamma} \tilde{Z}'_0 \rightarrow \tilde{X}'$. The quotient \tilde{W}/N can be identified with \tilde{Z} as a G -Galois cover, since we have compatible identifications of their pullbacks over \tilde{X}^* , X' , and their “overlap” \tilde{X}' , and because of the uniqueness assertion of the patching theorem. Also, \tilde{W} is normal, since normality is a local property and since \tilde{W}^* and W' are normal. Let \tilde{W}_0 be the connected component of \tilde{W} whose pullback to \tilde{X}^* contains \tilde{W}_0^* . Its Galois group Γ_1 over \tilde{X} surjects onto $G = \text{Gal}(\tilde{Z}/\tilde{X})$,

and Γ_1 contains $\text{Gal}(\tilde{W}_0^*/\tilde{X}^*) = \Gamma_0 \supset N \rtimes 1$. So Γ_1 is all of Γ , and so $\tilde{W}_0 = \tilde{W}$. That is, \tilde{W} is connected, and hence irreducible (being normal).

Now let $W^* \rightarrow X^*$ be the normalization of X^* in \tilde{W} . This is then a Γ -Galois irreducible cover that dominates Z^* (since Z^* is the normalization of X^* in \tilde{Z}). It is irreducible because it is connected and normal. So it provides a proper solution to the given embedding problem. \square

Remark 5.3.12. (a) Observe that the above theorem would also follow from Theorem 5.1.9, if it were known that $\mathbb{C}((x, y))$ is large. (Namely, given a split embedding problem over $\mathbb{C}((x, y))$, one could apply Theorem 5.1.9 to the induced constant split embedding problem over $\mathbb{C}((x, y))(t)$; and then one could specialize the proper solution to an extension of $\mathbb{C}((x, y))$, using that that field is separably Hilbertian by Weissauer's Theorem [FJ, Theorem 14.17].) But it is unknown whether $\mathbb{C}((x, y))$ is large. (Cf. Example 3.3.7(d).)

(b) It would be desirable to generalize the above result, e.g. by allowing more Laurent series variables, and by replacing \mathbb{C} by an algebraically closed field of arbitrary characteristic (or even by an arbitrary large field). Note that the above proof uses Kummer theory and Abhyankar's Lemma, and so one would somehow need to treat the case of wild ramification. \square

The ultimate goal remains that of proving a full analog of Riemann's Existence Theorem — classifying covers via their Galois groups and inertia groups, and determining how they fit together into the tower of covers. This goal, however, has so far been achieved in full only for curves over algebraically closed fields of characteristic 0 (where it is deduced from the complex result, which relied on topological methods). As seen above, the weaker goal of finding π_1 as a profinite group, and finding absolute Galois groups of function fields, also remains open in most cases, although the absolute Galois group of the function field is known for curves over algebraically closed fields (Theorem 5.1.1), and partial results are known for other fields (e.g. Theorem 5.1.9, 5.3.4, and 5.3.9). The still weaker, but difficult, goal of finding π_A has been achieved for affine curves over algebraically closed fields of arbitrary characteristic (Theorem 5.3.1 above), and the goal of finding which groups are Galois groups over the function field is settled for curves over large fields and fraction fields of complete local rings (Theorems 3.3.1 and 3.3.6) and partially for curves over finite fields (Proposition 3.3.9). But the structure of the absolute Galois groups of most familiar fields remains undetermined (e.g. for number fields and function fields of several variables over \mathbb{C}), and the inverse Galois problem over \mathbb{Q} remains open. The strategy used in Theorem 5.3.9 above, though, may suggest an approach to higher dimensional geometric fields; and Remark 3.3.8(a) suggests a possible strategy in the number field case. These and other patching methods described here may help further attack these open problems, on the way toward achieving a full generalization of Riemann's Existence Theorem.

References.

- [Ab1] S. Abhyankar. Coverings of algebraic curves. *Amer. J. Math.* **79** (1957), 825-856.
- [Ab2] S.S. Abhyankar. Tame coverings and fundamental groups of algebraic varieties, Part 1. *Amer. J. Math.* **84** (1959), 46-94.
- [Ab3] S.S. Abhyankar. Local fundamental groups of algebraic varieties. *Proc. Amer. Math. Soc.* **125** (1997), 1635-1641.
- [Ar1] M. Artin. “Grothendieck Topologies”. Harvard University Math. Dept., 1962.
- [Ar2] M. Artin. Algebraic approximation of structure over complete local rings. *Publ. Math. IHES* **36** (1969), 23-58.
- [Ar3] M. Artin. Algebraization of formal moduli: I. In “Global Analysis: Papers in Honor of K. Kodaira”. Princeton Univ. Press, 1969, pp. 21-71.
- [Ar4] M. Artin. Algebraization of formal moduli: II. Existence of modifications. *Annals of Math. ser. 2*, **91** (1970), 88-135.
- [Ar5] M. Artin. Introduction to Part I: Holomorphic functions. In [Za5], pp.3-8.
- [As] M. Aschbacher. “Finite group theory”, 2nd ed. Cambridge Univ. Press, 2000.
- [Ax1] J. Ax. Proof of some conjectures on cohomological dimension. *Proc. Amer. Math. Soc.* **16** (1965), 1214-1221.
- [Ax2] J. Ax. The elementary theory of finite fields. *Ann. of Math.* **88** (1968), 239-271.
- [Be] S. Beckmann. Is every extension of \mathbb{Q} the specialization of a branched covering? *J. Algebra* **164** (1994), 430-451.
- [Bi] G.D. Birkhoff. A theorem on matrices of analytic functions. *Math. Ann.* **74** (1917), 240-251.
- [Bl1] E. Black. Arithmetic lifting of dihedral extensions. *J. Algebra* **203** (1998), 12-29.
- [Bl2] E. Black. On semidirect products and the arithmetic lifting property. *J. London Math. Soc. (2)* **60** (1999), 677-688.
- [Bl3] E. Black. Lifting wreath product extensions. *Proc. Amer. Math. Soc.* **129** (2001), 1283-1288.
- [BGR] S. Bosch, U. Güntzer, R. Remmert. “Non-Archimedean Analysis”. *Grundlehren der mathematischen Wissenschaften* **261**, Springer-Verlag, 1984.
- [BLü1] S. Bosch, W. Lütkebohmert. Formal and rigid geometry: I. Rigid spaces. *Math. Ann.* **295** (1993), 291-317.
- [BLü2] S. Bosch, W. Lütkebohmert. Formal and rigid geometry: II. Flattening techniques. *Math. Ann.* **296** (1993), 403-429.
- [BLüR1] S. Bosch, W. Lütkebohmert, M. Raynaud. Formal and rigid geometry: III. The relative maximum principle. *Math. Ann.* **302** (1995), 1-29.

- [BLüR2] S. Bosch, W. Lütkebohmert, M. Raynaud. Formal and rigid geometry: IV. The reduced fibre theorem. *Invent. Math.* **119** (1995), 361-398.
- [BLoR] J.-B. Bost, F. Loeser, M. Raynaud (eds.). “Courbes semi-stables et groupe fondamental en géométrie algébrique”. *Progress in Mathematics* vol. 187, Birkhäuser, 2000.
- [Bo] N. Bourbaki. “Elements of Mathematics: Commutative algebra”. Hermann, 1972.
- [BP] I. Bouw, R. Pries. Rigidity, reduction, and ramification. 2001 manuscript. To appear in *Math. Annalen*.
- [Ca1] H. Cartan. Sur les matrices holomorphes de n variables complexes. *Journal de Mathématiques pures et appliquées, Series 9*, **19** (1940), 1-26.
- [Ca2] H. Cartan. Séminaire: Théorie des fonctions de plusieurs variables. *Éc. Norm. Sup.*, 1951-52.
- [Ca3] H. Cartan. Séminaire: Théorie des fonctions de plusieurs variables. *Éc. Norm. Sup.*, 1953-54.
- [CL] A. Chambert-Loir. La conjecture d’Abhyankar I: construction de revêtements en géométrie rigide. In [BLoR], pp. 233-248.
- [Ch] W.-L. Chow. On compact complex analytic varieties. *Amer. J. Math.* **71** (1949), 893-914.
- [Co] R. Coleman. Torsion points on curves. In: “Galois Representations and Arithmetic Algebraic Geometry” (Y. Ihara, ed.), *Advanced Studies in Pure Mathematics*, vol. 12, Elsevier Publishers, 1987.
- [CT] J.-L. Colliot-Thélène. Rational connectedness and Galois covers of the projective line. *Ann. of Math.*, **151** (2000), 359-373.
- [CH] K. Coombes, D. Harbater. Hurwitz families and arithmetic Galois groups, *Duke Math. J.*, **52** (1985), 821-839.
- [Dè] P. Dèbes. Galois covers with prescribed fibers: the Beckmann-Black problem. *Ann. Scuola Norm. Sup. Pisa Cl. Sci. (4)* **28** (1999), 273-286.
- [DD] P. Dèbes, B. Deschamps. The inverse Galois problem over large fields. In “Geometric Galois Actions, 2”, *London Math. Soc. Lecture Note Series* **243**, L. Schneps and P. Lochak ed., Cambridge University Press, (1997), 119-138.
- [DF] P. Dèbes, M. Fried. Rigidity and real residue class fields. *Acta Arith.* **56** (1990), 291-323.
- [Do] A. Douady. Détermination d’un groupe de Galois. *C.R. Acad. Sci. Paris* **258** (1964), 5305-5308.
- [FR] D. Ferrand and M. Raynaud. Fibres formelles d’un anneau local noethérien. *Ann. Scient. Éc. Norm. Sup. ser. 4*, **3** (1970), 295-311.

- [FP] J. Fresnel, M. van der Put. “Géométrie Analytique Rigide et Applications”. Progress in Mathematics, vol. 18, Birkäuser, 1981.
- [Fr1] M. Fried. Fields of definition of function fields and Hurwitz families – groups as Galois groups. *Comm. Alg.* **5** (1977), 17-82.
- [Fr2] M. Fried, et al., eds. “Recent developments in the inverse Galois problem”. AMS Contemporary Mathematics Series, vol. 186, 1995.
- [Fr3] M. Fried. Introduction to modular towers. In [Fr2], pp.111-171.
- [FJ] M. Fried, M. Jarden. “Field Arithmetic.” *Ergebnisse Math. series*, vol. 11, Springer-Verlag, 1986.
- [FV1] M. Fried, H. Völklein. The inverse Galois problem and rational points on moduli spaces. *Math. Ann.* **290** (1991), 771-800.
- [FV2] M. Fried, H. Völklein. The embedding problem over a Hilbertian PAC-field. *Annals of Math.* **135** (1992), 469-481.
- [Fu1] W. Fulton. Hurwitz schemes and irreducibility of moduli of algebraic curves. *Ann. of Math.* (2) **90** (1969), 542-575.
- [Fu2] W. Fulton. On the fundamental group of the complement of a node curve. *Annals of Math.* **111** (1980), 407-409.
- [Ga] M. Garuti. Géométrie rigide et géométrie formelle. In [BLoR], pp. 7-19.
- [GG] L. Gerritzen, H. Grauert. Die Azyklizität der affinoiden Überdeckungen. In “Global Analysis: Papers in Honor of K. Kodaira”. Princeton Univ. Press, 1969, pp. 159-184.
- [GrRe1] H. Grauert, R. Remmert. Nichtarchimedische Funktionentheorie. In: “Festschrift Gedächtnisfeier K. Weierstrass”, *Wissenschaftl. Abh. Arbeitsgemeinschaft für Forschung des Landes Nordrhein-Westfalen*, vol. 33, Westdeutscher Verlag, 1966, pp. 393-476.
- [GrRe2] H. Grauert, R. Remmert. Über die Methode der diskret bewerteten Ringe in der nicht-archimedischen Analysis. *Invent. Math.* **2** (1966), 87-133.
- [GH] P. Griffiths, J. Harris. “Principles of algebraic geometry”. John Wiley and Sons, 1978.
- [Gr1] A. Grothendieck. Sur quelques points d’algèbre homologique. *Tôhoku Math. J.* **9** (1957), 119-221.
- [Gr2] A. Grothendieck. Géométrie formelle et géométrie algébrique. *Sem. Bourb.* **182** (1959), 1-28.
- [Gr3] A. Grothendieck. Technique de descente et théorèmes d’existence en géométrie algébrique, I. *Sem. Bourb.* **190** (1959), 1-29.
- [Gr4] A. Grothendieck. “Éléments de géométrie algébrique” (EGA), *Publ. Math. IHES*: EGA II, vol. 8 (1961); EGA III, 1^e partie, vol. 11 (1961); EGA IV, 1^e partie, vol. 20 (1964).

- [Gr5] A. Grothendieck. “Revêtements étales et groupe fondamental” (SGA 1). Lecture Notes in Mathematics, vol. 224, Springer-Verlag, 1971.
- [Gr6] A. Grothendieck. Letter to Faltings. 1983 letter. In “Geometric Galois Actions, I” (L. Schneps and P. Lochak, eds.), London Math. Soc. Lecture Notes, vol. 242, Cambridge Univ. Press, 1997, pp. 49-58 (German), pp. 285-293 (English translation).
- [GuRo] R.C. Gunning, H. Rossi. “Analytic functions of several complex variables”. Prentice-Hall, 1965.
- [HJ1] D. Haran, M. Jarden. Regular split embedding problems over complete valued fields. *Forum Mathematicum* **10** (1998), 329-351.
- [HJ2] D. Haran, M. Jarden. Regular lifting of covers over ample fields. 2000 manuscript.
- [HV] D. Haran, H. Völklein. Galois groups over complete valued fields. *Israel J. Math.*, **93** (1996), 9-27.
- [Ha1] D. Harbater. Deformation theory and the tame fundamental group. *Trans. AMS*, **262** (1980), 399-415.
- [Ha2] D. Harbater. Convergent arithmetic power series. *Amer. J. Math.*, **106** (1984), 801-846.
- [Ha3] D. Harbater. Mock covers and Galois extensions. *J. Algebra*, **91** (1984), 281-293.
- [Ha4] D. Harbater. Galois coverings of the arithmetic line. In “Number Theory: New York, 1984-85”. *Lecture Notes in Mathematics*, vol. 1240, Springer-Verlag, 1987, pp. 165-195.
- [Ha5] D. Harbater. Galois covers of an arithmetic surface. *Amer. J. Math.*, **110** (1988), 849-885.
- [Ha6] D. Harbater. Formal patching and adding branch points. *Amer. J. Math.*, **115** (1993), 487-508.
- [Ha7] D. Harbater. Abhyankar’s conjecture on Galois groups over curves. *Inventiones Math.*, **117** (1994), 1-25.
- [Ha8] D. Harbater. Galois groups with prescribed ramification. In “Proceedings of a Conference on Arithmetic Geometry (Arizona State Univ., 1993)”, *AMS Contemp. Math. Series*, vol. 174, 1994, pp. 35-60.
- [Ha9] D. Harbater. Fundamental groups of curves in characteristic p . *Proc. Intl. Cong. of Math. (Zurich, 1994)*, Birkhauser, 1995, pp.656-666.
- [Ha10] D. Harbater. Fundamental groups and embedding problems in characteristic p . In [Fr2], pp. 353-369.
- [Ha11] D. Harbater. Embedding problems and adding branch points. In “Aspects of Galois theory” (H. Völklein, et al., eds.), *London Math. Soc. Lecture Notes*, vol. 256, Cambridge Univ. Press, 1999, pp.119-143.

- [Ha12] D. Harbater. Embedding problems with local conditions. *Israel J. of Math*, **118** (2000), 317-355.
- [Ha13] D. Harbater. Abhyankar's Conjecture and embedding problems. To appear in *Crelle's J.* Available at <http://www.math.upenn.edu/~harbater/qp.dvi>
- [HS1] D. Harbater, K. Stevenson. Patching and thickening problems. *J. Alg.*, **212** (1999), 272-304.
- [HS2] D. Harbater, K. Stevenson. Abhyankar's local conjecture on fundamental groups. To appear in "Proceedings of the Conference on Algebra and Algebraic Geometry with Applications" (Abhyankar 70th birthday conference, C. Christensen, ed.), Springer-Verlag. Available at <http://www.math.upenn.edu/~harbater/abhloc.dvi>
- [HP] D. Harbater, M. van der Put. Valued fields and covers in characteristic p (with an appendix by R. Guralnick). 2000 manuscript, to appear in "Valuation Theory and its Applications", ed. by F.-V. Kuhlmann, S. Kuhlmann and M. Marshall.
- [Hrt1] R. Hartshorne. "Residues and duality." *Lecture Notes in Mathematics*, vol. 20, Springer-Verlag, 1966.
- [Hrt2] R. Hartshorne. "Algebraic geometry". *Graduate Texts in Mathematics*, vol. 52. Springer-Verlag, 1977.
- [He] Y. Henrio. Disques et couronnes ultramétriques. In [BLoR], pp. 21-32.
- [Iw] K. Iwasawa. On solvable extensions of algebraic number fields. *Annals of Math.* **58** (1953), 548-572.
- [Ja] M. Jarden. On free profinite groups of uncountable rank. In [Fr2], pp. 371-383.
- [Ka] N. Katz. "Rigid local systems". *Annals of Mathematics Studies*, vol. 139. Princeton University Press, 1996.
- [Ki1] R. Kiehl. Der Endlichkeitssatz für eigentliche Abbildungen in der nichtarchimedischen Funktionentheorie. *Invent. Math.* **2** (1967), 191-214.
- [Ki2] R. Kiehl. Theorem A und Theorem B in der nichtarchimedischen Funktionentheorie. *Invent Math.* **2** (1967), 256-273.
- [Kol] J. Kollár. Rationally connected varieties over local fields. *Ann. of Math.* **150** (1999), 357-367.
- [Köp] U. Köpf. Über eigentliche Familien algebraischer Varietäten über affinoiden Räumen. *Schriftenreihe des Math. Inst. der Univ. Münster*, 2 Serie, Heft 7 (1974).
- [La] S. Lang. "Algebraic Number Theory". Addison-Wesley, 1970.
- [Le] T. Lefcourt. Galois groups and complete domains. *Israel J. Math.*, **114** (1999), 323-346.
- [Li] Q. Liu. Tout groupe fini est un groupe de Galois sur $\mathbb{Q}_p(T)$, d'après Harbater. In [Fr2], pp. 261-265.

- [Lü1] W. Lütkebohmert. Formal-algebraic and rigid-analytic geometry. *Math. Ann.* **286** (1990), 341-371.
- [Lü2] W. Lütkebohmert. Riemann's existence theorem for a p -adic field. *Invent. Math.* **111** (1993), 309-330.
- [MR] M. Madan, M. Rosen. The automorphism group of a function field. *Proc. Amer. Math. Soc.* **115** (1992), 923-929.
- [MM] G. Malle, B.H. Matzat. "Inverse Galois theory". Springer Monographs in Mathematics. Springer-Verlag, 1999.
- [Mi] J.S. Milne. "Étale Cohomology". Princeton Mathematical Series, vol. 33, Princeton University Press, 1980.
- [MB1] L. Moret-Bailly. Groupes de Picard et problèmes de Skolem, II. *Ann. Sci. École Norm. Sup. (4)* **22** (1989), 181-194.
- [MB2] L. Moret-Bailly. Construction de revêtements de courbes pointées. *J. Algebra*, **240** (2001), 505-534.
- [NSW] J. Neukirch, A. Schmidt, K. Wingberg. "Cohomology of Number Fields". Grundlehren series, vol. 323, Springer, 2000.
- [PS] A. Pacheco, K. Stevenson. Finite quotients of the algebraic fundamental group of projective curves in positive characteristic. *Pacific J. Math.*, **192** (2000), 143-158.
- [Po1] F. Pop. The geometric case of a conjecture of Shafarevich. Manuscript, University of Heidelberg, 1993.
- [Po2] F. Pop. Half Riemann Existence Theorem with Galois action. In "Algebra and Number Theory", G. Frey and J. Ritter, eds., de Gruyter Proceedings in Math. (1994), 1-26.
- [Po3] F. Pop. Étale Galois covers of affine smooth curves. *Invent. Math.*, **120** (1995), 555-578.
- [Po4] F. Pop. Embedding problems over large fields. *Ann. Math.*, **144** (1996), 1-34.
- [Pr1] R. Pries. Construction of covers with formal and rigid geometry. In [BLoR], pp. 157-167.
- [Pr2] R. Pries. Degeneration of wildly ramified covers of curves. 2001 manuscript.
- [Ra1] M. Raynaud. Géométrie analytique rigide d'après Tate, Kiehl, *Bull. Soc. Math. France*, Mémoire 39-40 (1974), 319-327.
- [Ra2] M. Raynaud. Revêtements de la droite affine en caractéristique $p > 0$ et conjecture d'Abhyankar. *Invent. Math.* **116** (1994), 425-462.
- [Ri] L. Ribes. "Introduction to profinite groups and Galois cohomology Introduction to profinite groups and Galois cohomology". Queen's University, 1970.

- [Ru] W. Rudin. “Real and complex analysis”, second edition. McGraw-Hill, 1974.
- [Sa1] M. Saïdi. Revêtements modérés et groupe fondamental de graphe de groupes. *Compositio Math.* **107** (1997), 319-338.
- [Sa2] M. Saïdi. Abhyankar’s conjecture II: the use of semi-stable curves. In [BLoR], pp. 249-265.
- [Sl1] D. Saltman. Generic Galois extensions and problems in field theory. *Advances in Math.* **43** (1982), 250-283.
- [Se1] J.-P. Serre. Faisceaux analytiques sur l’espace projectif. Exposés XVIII and XIX in [Ca3].
- [Se2] J.-P. Serre. Faisceaux algébriques cohérents. *Ann. of Math.* **61** (1955), 197-278.
- [Se3] J.-P. Serre. Géométrie algébrique et géométrie analytique. *Annales de L’Institut Fourier* **6** (1956), 1-42.
- [Se4] J.-P. Serre. “Cohomologie Galoisienne”. *Lecture Notes in Mathematics*, vol. 5, Springer-Verlag, 1964.
- [Se5] J.-P. Serre. “Local fields”. *Graduate Texts in Math.*, vol. 67, Springer-Verlag, 1979.
- [Se6] J.-P. Serre. Construction de revêtements étales de la droite affine en caractéristique p . *Comptes Rendus* **311** (1990), 341-346.
- [Se7] J.-P. Serre. “Topics in Galois Theory”. *Research Notes in Math.*, vol. 1, Jones and Bartlett, 1992.
- [St1] K. Stevenson. Galois groups of unramified covers of projective curves in characteristic p . *J. Algebra* **182** (1996), 770-804.
- [St2] K. Stevenson. Conditions related to π_1 of projective curves. *J. Number Theory* **69** (1998), 62-79.
- [Tm1] A. Tamagawa. The Grothendieck conjecture for affine curves. *Compositio Math.* **109** (1997), 135-194.
- [Tm2] A. Tamagawa. On the fundamental group of curves over algebraically closed fields of characteristic > 0 . *Intern. Math. Res. Notices* (1999), no. 16, 853-873.
- [Ta] J. Tate. Rigid analytic spaces. *Invent. Math.* **12** (1971), 257-289.
- [Th] J. Thompson. Some finite groups which appear as $\text{Gal } L/K$, where $K \subseteq \mathbb{Q}(\mu_n)$. *J. Algebra* **89** (1984), 437-499.
- [Vö] H. Völklein. “Groups as Galois groups — an introduction”. *Cambridge Studies in Advanced Mathematics*, vol. 53, Cambridge University Press, 1996.
- [We] C. Weibel. “An introduction to homological algebra”. *Cambridge Studies in Advanced Mathematics*, vol. 38, Cambridge Univ. Press, 1994.

- [Za1] O. Zariski. On the problem of existence of algebraic functions of two variables possessing a given branch curve. Amer. J. of Math., **51** (1929), 305-328.
- [Za2] O. Zariski. Generalized semi-local rings. Summa Brasiliensis Mathematicae **1**, fasc. 8 (1946), 169-195.
- [Za3] O. Zariski. "Algebraic surfaces". Chelsea Pub. Co., 1948.
- [Za4] O. Zariski. Theory and applications of holomorphic functions on algebraic varieties over arbitrary ground fields. Mem. Amer. Math. Soc. **5** (1951), 1-90. Reprinted in [Za5], pp. 72-161.
- [Za5] O. Zariski. "Collected Papers, volume II: Holomorphic functions and linear systems". M. Artin and D. Mumford, eds. MIT Press, 1973.

Mailing address: Department of Mathematics, University of Pennsylvania, Philadelphia, PA 19104-6395, USA.

E-mail address: `harbater@math.upenn.edu`