

On the Geometry of p -Typical Covers in Characteristic p

Kiran S. Kedlaya

Abstract. For p a prime, a p -typical cover of a connected scheme on which $p = 0$ is a finite étale cover whose monodromy group (i.e., the Galois group of its normal closure) is a p -group. The geometry of such covers exhibits some unexpectedly pleasant behaviors; building on work of Katz, we demonstrate some of these. These include a criterion for when a morphism induces an isomorphism of the p -typical quotients of the étale fundamental groups, and a decomposition theorem for p -typical covers of polynomial rings over an algebraically closed field.

1 Introduction

Let p be a prime number. A finite étale cover of a connected scheme on which $p = 0$ is p -typical if the monodromy group of the cover (which for a connected cover coincides with the Galois group of the normal closure) is a p -group. The geometry of such covers exhibits some unexpectedly pleasant behaviors; the purpose of this paper is to briefly expose a few of these. This is in part to dispel the notion that one can only ever prove meaningful results about the tame (prime-to- p) quotient of the étale fundamental group.

For instance, Katz has shown [5, Proposition 1.4.2] that if R is a connected ring in which $p = 0$, then the categories of p -typical covers over $R[t^{-1}]$ and over $R((t))$ are equivalent, via the evident base change functor. In other words, if π_1^p denotes the maximal pro- p quotient of the étale fundamental group π_1 (where basepoints are suppressed throughout this introduction for notational simplicity), then the natural homomorphism $\pi_1^p(R((t))) \rightarrow \pi_1^p(R[t^{-1}])$ is a bijection. We give a natural generalization of Katz's theorem (Theorem 2.30), which characterizes more generally when one connected affine scheme of characteristic p looks like a limit of a diagram of others from the point of view of constructing π_1^p . Here is a sample result (Example 2.34): if k is an algebraically closed field of characteristic $p > 0$, then

$$\pi_1^p(k[t, t^{-1}]) \cong \pi_1^p(k[t]) \times \pi_1^p(k[t^{-1}]).$$

(The analogous statement for π_1 is false: the left side has nontrivial prime-to- p quotients whereas the right side does not. Note also that in general, neither π_1 nor π_1^p commutes with products, so one cannot replace the right side with a single fundamental group of \mathbb{A}_k^2 .)

We also look more closely at p -typical covers of affine toric varieties, including, of course, ordinary affine spaces. Our main results in this direction (Theorem 4.15 and

Received by the editors May 27, 2005.

The author was supported by NSF grant DMS-0400727.

AMS subject classification: 14F35.

©Canadian Mathematical Society 2008.

its corollaries, notably Theorem 4.17) assert that the π_1^p of any affine toric variety can be written as an inverse limit of π_1^p 's of one-dimensional varieties, or even of affine lines. This requires the use of some auxiliary “height functions” to measure the complexity of p -typical covers; we can describe some simple examples of such functions, but only *a posteriori* (Theorem 5.11).

2 p -Typical Covers

In this section, we introduce the notion of a p -typical cover and prove a strong generalization of Katz's canonical extension property for such covers (Theorem 2.30). We first fix some notational conventions for the whole paper.

Convention 2.1 Throughout this paper, fix a prime number p . By a p -group, we will mean a finite group whose order is a power of p . Standard facts about p -groups, which we will use without further comment, include the following.

- The center of any nontrivial p -group is nontrivial.
- Any maximal proper subgroup of a nontrivial p -group is normal of index p .

By a p -ring, we will mean a ring in which $p = 0$; likewise for p -field or p -domain. Similarly, by a p -scheme, we will mean a scheme in whose ring of global sections one has $p = 0$.

2.1 The Étale Fundamental Group

We first recall the notion of the étale fundamental group from [3, Exposé V] (with some notation as in [5, §1.2]).

Convention 2.2 Throughout this section, let X be a connected scheme, and let \bar{x} be a geometric point of X , *i.e.*, a morphism $\text{Spec } k^{\text{alg}} \rightarrow X$ in which k^{alg} is an algebraically closed field.

Definition 2.3 Let \mathcal{C}_X denote the category of finite étale covers of X ; note that $\mathcal{C}_{\bar{x}}$ may be identified with the category of finite sets. Then the pullback functor $F_{\bar{x}}: \mathcal{C}_X \rightarrow \mathcal{C}_{\bar{x}}$ is represented by a pro-object P of \mathcal{C}_X . Let $\pi_1(X, \bar{x})$ denote the automorphism group of $F_{\bar{x}}$, *i.e.*, the group of pro-automorphisms of P .

Remark 2.4 Replacing \bar{x} by another geometric point \bar{y} does not change the abstract structure of the group $\pi_1(X)$. However, there is no *canonical* isomorphism $\pi_1(X, \bar{x}) \rightarrow \pi_1(X, \bar{y})$; the choice of such an isomorphism constitutes the choice of a *chemin* or *path*.

Definition 2.5 Let X be a connected scheme, let $E \rightarrow X$ be a finite étale cover, and let \bar{x} be a geometric point of X . Then the profinite group $\pi_1(X, \bar{x})$ acts continuously on $E_{\bar{x}}$, and the image is well defined up to group isomorphism. We call it the *monodromy group* of E .

Definition 2.6 If $E \rightarrow X$ is a connected finite étale cover, there is a unique minimal connected finite Galois (étale) cover $E' \rightarrow X$ which factors through E ; it is the maximal cover fixed by the kernel of the map $\pi_1(X, \bar{x}) \rightarrow \text{Aut}(E_{\bar{x}})$. Consequently, the Galois group of this cover is precisely the monodromy group of $E \rightarrow X$. This cover is called the *normal closure* (or *Galois closure*) of $E \rightarrow X$; it coincides with the usual field-theoretic definition when $X = \text{Spec } k$.

2.2 p -Typical Covers

We now extract the p -typical part of the fundamental group. Throughout this section, we retain Convention 2.2.

Definition 2.7 A p -typical cover of X is a finite étale cover $E \rightarrow X$ whose monodromy group is a p -group; if S/R is a ring extension whose corresponding cover $\text{Spec } S \rightarrow \text{Spec } R$ is p -typical, we say S is a p -typical extension of R . Note that the fibre product and the disjoint union of p -typical covers are p -typical. If E is connected and p -typical over X , then $\deg(E \rightarrow X)$ is a power of p : namely, this degree is the index in the monodromy group of the stabilizer of any geometric point of E .

Lemma 2.8 If $E \rightarrow X$ and $E' \rightarrow E$ are finite étale covers with E connected, then $E' \rightarrow X$ is p -typical if and only if $E' \rightarrow E$ and $E \rightarrow X$ are both p -typical.

Proof Choose a geometric point \bar{x} of X and a geometric point \bar{y} of $E_{\bar{x}}$. Let G be the monodromy group of $E' \rightarrow X$, identified with the image of $\pi_1(X, \bar{x})$ in $\text{Aut}(E'_{\bar{x}})$, and let H be the monodromy group of $E' \rightarrow E$, identified with the image of $\pi_1(E, \bar{y})$ in $\text{Aut}(E'_{\bar{y}})$. Then H is the stabilizer of \bar{y} within G .

On one hand, if $E' \rightarrow E$ and $E \rightarrow X$ are p -typical, then H is a p -group, G acts transitively on the geometric points of $E_{\bar{x}}$ (since E is connected), and so $\#G = \#H \cdot \deg(E \rightarrow X)$ is a p -power. Hence $E' \rightarrow X$ is p -typical.

On the other hand, if $E' \rightarrow X$ is p -typical, then G is a p -group, as then must be H , so $E' \rightarrow E$ is p -typical. Meanwhile, the monodromy group of $E \rightarrow X$ is a quotient of G , since any element of $\pi_1(X, \bar{x})$ fixing $E'_{\bar{x}}$ must in particular fix $E_{\bar{x}}$. Hence $E \rightarrow X$ is also p -typical. ■

Definition 2.9 Let \mathcal{C}_X^p denote the subcategory of \mathcal{C}_X consisting of p -typical covers. Again, the fibre functor $F_{\bar{x}}^p : \mathcal{C}_X^p \rightarrow \mathcal{C}_{\bar{x}}$ is represented by a pro-object of \mathcal{C}_X^p , whose group of pro-automorphisms coincides with the automorphism group of $F_{\bar{x}}^p$. We call this group $\pi_1^p(X, \bar{x})$ and refer to it as the p -typical fundamental group of X ; the inclusion $\mathcal{C}_X^p \hookrightarrow \mathcal{C}_X$ induces a surjection $\pi_1(X, \bar{x}) \rightarrow \pi_1^p(X, \bar{x})$, under which $\pi_1^p(X, \bar{x})$ is identified with the maximal pro- p quotient of $\pi_1(X, \bar{x})$.

2.3 p -Typical Covers and Artin–Schreier Towers

We will mainly be interested in p -typical covers of p -schemes; these can be studied using Artin–Schreier (AS) towers.

Definition 2.10 For G a finite group (viewed as a constant group scheme over $\text{Spec } \mathbb{Z}$) and X a scheme, a G -torsor over X is a finite étale cover $E \rightarrow X$ equipped with an action of G , which étale locally on X is isomorphic to $X \times G$ (the *trivial G -torsor*). If $X = \text{Spec } R$ is affine, we refer to a G -torsor over X also as a G -torsor over R ; it is also affine because a finite étale cover of an affine scheme is affine.

Definition 2.11 Let X be a p -scheme, and let $E \rightarrow X$ be a finite étale cover. An *AS-tower* for $E \rightarrow X$ is a sequence of finite étale covers

$$E = E_d \rightarrow E_{d-1} \rightarrow \cdots \rightarrow E_1 \rightarrow E_0 = X$$

in which $E_i \rightarrow E_{i-1}$ is equipped with a $\mathbb{Z}/p\mathbb{Z}$ -torsor structure for $i = 1, \dots, d$. From the transitivity of p -typicality (Lemma 2.8), we see that the existence of an AS-tower for $E \rightarrow X$ implies that $E \rightarrow X$ is p -typical. If $X = \text{Spec } R$ and $E = \text{Spec } S$, we typically write the tower ring-theoretically, as $S_0 = R \subset S_1 \subset \cdots \subset S_d = S$, in which $E_i = \text{Spec } S_i$ and S_i/S_{i-1} is a $\mathbb{Z}/p\mathbb{Z}$ -torsor for $i = 1, \dots, d$.

Proposition 2.12 Let X be a connected p -scheme, and let $E \rightarrow X$ be a connected finite étale cover. Then $E \rightarrow X$ is p -typical if and only if there exists an AS-tower for $E \rightarrow X$.

Proof We have noted already that if there exists an AS-tower for $E \rightarrow X$, then $E \rightarrow X$ is p -typical (with no connectedness hypotheses). Conversely, suppose that $E \rightarrow X$ is p -typical with monodromy group G , which we may assume is nontrivial. Pick a geometric point \bar{x} of X , identify G with the image of $\pi_1(X, \bar{x})$ in $\text{Aut}(E_{\bar{x}})$, and pick a geometric point \bar{y} of $E_{\bar{x}}$. Then the stabilizer of \bar{y} is a proper subgroup of G ; thus it is contained in a maximal proper subgroup H of G , which is necessarily normal of index p . In particular, because H is normal, it contains the stabilizers of *all* of the points of $E_{\bar{x}}$. Thus G/H is the monodromy group of a connected $\mathbb{Z}/p\mathbb{Z}$ -torsor $E' \rightarrow X$ through which E factors. By induction, the desired result follows. ■

When $E \rightarrow X$ is Galois, one gets a bit more.

Proposition 2.13 Let $E \rightarrow X$ be a connected Galois p -typical cover. Then there exists an AS-tower $E = E_d \rightarrow E_{d-1} \rightarrow \cdots \rightarrow E_1 \rightarrow E_0 = X$ in which $E_i \rightarrow X$ is Galois for $i = 1, \dots, d$.

Proof Put $G = \text{Aut}(E \rightarrow X)$, which coincides with the monodromy group of $E \rightarrow X$ because the cover is Galois, and assume G is nontrivial. Since the center of G is nontrivial, it contains a subgroup H of order p , which is normal in G . Let E_{d-1} be the maximal subcover fixed by H , and repeat. ■

For $\mathbb{Z}/p\mathbb{Z}$ -torsors over p -rings, one has the following standard result.

Definition 2.14 For R a p -ring, a $\mathbb{Z}/p\mathbb{Z}$ -torsor of the form $S = R[z]/(z^p - z - a)$, in which $1 \in \mathbb{Z}/p\mathbb{Z}$ acts via $z \mapsto z + 1$, is called an *Artin–Schreier extension*, or an *AS-extension*, of R .

Proposition 2.15 For any p -ring R , every $\mathbb{Z}/p\mathbb{Z}$ -torsor of R is an AS-extension. Moreover, two such torsors $R[z_1]/(z_1^p - z_1 - a_1)$ and $R[z_2]/(z_2^p - z_2 - a_2)$ are isomorphic if and only if $a_1 - a_2 = y^p - y$ for some $y \in R$.

Proof The argument amounts to calculating étale cohomology of the sequence of sheaves:

$$0 \rightarrow \mathbb{Z}/p\mathbb{Z} \rightarrow \mathbb{G}_a \xrightarrow{F-1} \mathbb{G}_a \rightarrow 0$$

See [4, X.3.5] [5, 1.4.5], or [7, Proposition III.4.12] and subsequent discussion. ■

2.4 Connected Components in Positive Characteristic

We will need to keep careful track of the connected components of certain AS-towers. Before explaining how we do so, we first make some observations for arbitrary rings of positive (prime) characteristic.

Lemma 2.16 Let R be a p -ring. Then the set $S = \{x \in R : x^p = x\}$ is the \mathbb{F}_p -subalgebra of R generated by the idempotents of R .

Proof A straightforward exercise in algebra; alternatively, one may proceed as in Proposition 2.15. ■

Counting connected components of rings is closely related to testing for isomorphisms between finite flat ring extensions, as follows.

Remark 2.17 Let R be a connected p -ring, let S_1, S_2 be two connected finite flat extensions of R , and let $f: S_1 \rightarrow S_2$ be an R -algebra homomorphism. Then the graph Γ of f is a closed subscheme of $\text{Spec } S_1 \times_R \text{Spec } S_2$ which maps isomorphically onto $\text{Spec } S_2$ via the second projection. In particular, Γ is a connected component of $\text{Spec } S_1 \times_R \text{Spec } S_2$. Conversely, each connected component Γ of $\text{Spec } S_1 \times_R \text{Spec } S_2$ which maps isomorphically onto $\text{Spec } S_2$ via the second projection corresponds to an R -algebra homomorphism $S_1 \rightarrow S_2$. As a consequence, if $g: R \rightarrow R'$ is a ring homomorphism and the induced map $S_1 \otimes_R S_2 \xrightarrow{g} (S_1 \otimes_R S_2) \otimes_R R'$ induces a bijection of idempotents, then the induced map

$$\text{Hom}_{R\text{-alg}}(S_1, S_2) \xrightarrow{g} \text{Hom}_{R'\text{-alg}}(S_1 \otimes_R R', S_2 \otimes_R R')$$

is a bijection.

2.5 p -Injections and p -Surjections

We now consider some homomorphisms which behave nicely with respect to p -typical covers.

Definition 2.18 Let $f: R \rightarrow R'$ be a homomorphism of p -rings, and let F and F' denote the p -power Frobenius maps on R and R' , respectively. We say f is p -injective (resp. p -surjective) if the induced functor from $\mathbb{Z}/p\mathbb{Z}$ -torsors over R to $\mathbb{Z}/p\mathbb{Z}$ -torsors over R' is fully faithful (resp. essentially surjective). These definitions can be reformulated as follows.

- The map f is p -injective if and only if $\ker(F - 1) \xrightarrow{f} \ker(F' - 1)$ is bijective and $\operatorname{coker}(F - 1) \xrightarrow{f} \operatorname{coker}(F' - 1)$ is injective.
- The map f is p -surjective if and only if $\operatorname{coker}(F - 1) \xrightarrow{f} \operatorname{coker}(F' - 1)$ is surjective.

(See the proof of Proposition 2.22 for the explanation of how this reformulation follows from Artin–Schreier theory; alternatively, one may take the reformulation itself as the definition until Proposition 2.22 has been proved.) Using the snake lemma, we may give a second reformulation.

- The map f is p -injective if and only if $\ker(f) \xrightarrow{F-1} \ker(f)$ is bijective and $\operatorname{coker}(f) \xrightarrow{F'-1} \operatorname{coker}(f)$ is injective.
- The map f is p -surjective if and only if $\operatorname{coker}(f) \xrightarrow{F'-1} \operatorname{coker}(f)$ is surjective.

Remark 2.19 Note that the property of a morphism being p -surjective is not stable under flat base change. For instance, if $f: R \rightarrow R'$ is p -surjective but not surjective, then the induced homomorphism $R[t] \rightarrow R'[t]$ is not p -surjective. However, base changing by a p -typical extension causes no problems: see Corollary 2.21 below.

Lemma 2.20 Let $f: R \rightarrow R'$ be a homomorphism of p -rings, let $S = R[z]/(z^p - z - a)$ be an AS-extension of R , put $S' = S \otimes_R R'$, and let $f_S: S \rightarrow S'$ be the homomorphism induced by f .

- (i) If f induces an injection on idempotents, then so does f_S .
- (ii) If f is p -injective, then so is f_S .
- (iii) If f is p -surjective, then so is f_S .

Proof For $l = -1, \dots, p - 1$, let S_l and S'_l be the subsets of S and S' , respectively, consisting of polynomials in z of degree at most l (so that $S_{-1} = S'_{-1} = \{0\}$); note that each S_l (resp. S'_l) is preserved by F (resp. by F'). Let $f_l: S_l \rightarrow S'_l$ denote the map induced by f . We then have a commutative diagram

$$\begin{array}{ccccccc}
 0 & \longrightarrow & S_{l-1} & \longrightarrow & S_l & \longrightarrow & R \longrightarrow 0 \\
 & & \downarrow f_{l-1} & & \downarrow f_l & & \downarrow f \\
 0 & \longrightarrow & S'_{l-1} & \longrightarrow & S'_l & \longrightarrow & R' \longrightarrow 0
 \end{array}$$

which by the snake lemma gives rise to a long exact sequence

$$0 \rightarrow \ker(f_{l-1}) \rightarrow \ker(f_l) \rightarrow \ker(f) \rightarrow \operatorname{coker}(f_{l-1}) \rightarrow \operatorname{coker}(f_l) \rightarrow \operatorname{coker}(f) \rightarrow 0.$$

We now consider the cases separately. (i) By Lemma 2.16 and diagram chasing, f induces an injection on idempotents if and only if $F - 1$ induces an injection on $\ker(f)$. In this case, by induction on l and the five lemma, $F - 1$ induces an injection on $\ker(f_l)$ for $l = 0, \dots, p - 1$. Taking $l = p - 1$, we deduce that f_S induces an injection on idempotents.

(ii) If f is p -injective, then $F - 1$ is surjective on $\ker(f)$ and $F' - 1$ is injective on $\operatorname{coker}(f)$. By induction on l and the five lemma, $F - 1$ is surjective on $\ker(f_l)$ and $F' - 1$ is injective on $\operatorname{coker}(f_l)$ for $l = 0, \dots, p - 1$. Taking $l = p - 1$, we deduce that f_S is p -injective.

(iii) If f is p -surjective, then $F' - 1$ is surjective on $\operatorname{coker}(f)$. By induction on l and the five lemma, $F' - 1$ is surjective on $\operatorname{coker}(f_l)$ for $l = 0, \dots, p - 1$. Taking $l = p - 1$, we deduce that f_S is p -surjective. ■

Corollary 2.21 *Let $f: R \rightarrow R'$ be a homomorphism of p -rings, let $R = S_0 \subset S_1 \subset \dots \subset S_d$ be an AS-tower over R , put $S'_i = S_i \otimes_R R'$ for $i = 1, \dots, d$, and let $f_i: S_i \rightarrow S'_i$ be the homomorphism induced by f .*

- (i) *If f induces an injection on idempotents, then so does each f_i .*
- (ii) *If f is p -injective, then so is each f_i .*
- (iii) *If f is p -surjective, then so is each f_i .*

Proposition 2.22 *Let $f: R \rightarrow R'$ be a homomorphism of p -rings. Let \mathfrak{S}_R and $\mathfrak{S}_{R'}$ be the categories of AS-towers over R and R' , respectively, in which the only morphisms are isomorphisms of towers.*

- (i) *If the map f is p -injective, then the base change functor $f^*: \mathfrak{S}_R \rightarrow \mathfrak{S}_{R'}$ is fully faithful.*
- (ii) *The map f is p -surjective if and only if the base change functor $f^*: \mathfrak{S}_R \rightarrow \mathfrak{S}_{R'}$ is essentially surjective.*

Proof (i) Suppose that f is p -injective. Given two AS-towers $R = S_0 \subset S_1 \subset \dots \subset S_d = S$ and $R = T_0 \subset T_1 \subset \dots \subset T_d = T$ which become isomorphic over R' , write $S_1 = R[y]/(y^p - y - a)$ and $T_1 = R[z]/(z^p - z - b)$. By Proposition 2.15, $f(a)$ and $f(b)$ represent the same element of $\operatorname{coker}(F' - 1)$; hence they also represent the same element of $\operatorname{coker}(F - 1)$. Thus $S_1 \cong T_1$; moreover, by Lemma 2.20, the map $S_1 \rightarrow S_1 \otimes_R R'$ is p -injective. Repeating the argument, we see that the two towers are isomorphic, and so f^* is fully faithful.

(ii) Suppose that f^* is essentially surjective. Let $S' = R'[z]/(z^p - z - a)$ be an AS-extension of R' ; by hypothesis, there exists an AS-extension $S = R[z]/(z^p - z - b)$ such that $S \otimes_R R' \cong S'$ as a Z/pZ -torsor. By Proposition 2.15, we must have $f(b) - a = y^p - y$ for some $y \in R'$. We deduce that the map $\operatorname{coker}(F - 1) \rightarrow \operatorname{coker}(F' - 1)$ induced by f is surjective, and so f is p -surjective.

Conversely, suppose that f is p -surjective. Given an AS-tower $R' = S'_0 \subset S'_1 \subset \dots \subset S'_d = S'$, we construct a corresponding AS-tower $R = S_0 \subset S_1 \subset \dots \subset S_d$ inductively as follows. Start with $S_0 = R$. Given S_0, \dots, S_i and an isomorphism $S_i \otimes_R R' \cong S'_i$, note that $f: S_i \rightarrow S'_i$ is p -surjective by Lemma 2.20. By Proposition 2.15, we can write $S'_{i+1} = S'_i[z]/(z^p - z - a)$ for some $a \in f(S_i)$; we may then

set $S_{i+1} = S_i[z]/(z^p - z - b)$ for any $b \in S_i$ with $f(b) = a$. Thus the inductive construction continues, and so f^* is essentially surjective. ■

Remark 2.23 Beware that proving results about the category of AS-towers does not immediately yield results about p -typical covers; for that, stronger connected hypotheses are needed, as in the next section.

2.6 p -Limits and Canonical Extensions

Convention 2.24 Given a partially ordered set S , we view S as a category in which $\text{Mor}(s, t)$ is a singleton set if $s \geq t$ and is empty otherwise.

Definition 2.25 A *diagram* in a category \mathcal{C} is a functor D from a partially ordered set S to \mathcal{C} ; we call S the *support* of D . Given a subset T of S , let D_T denote the restriction of D to T .

Definition 2.26 Given a diagram D with support S , put $S_1 = S_2 = S \cup \{s'\}$ for some $s' \notin S$, and extend the partial order from S to S_1 and S_2 by declaring that in S_1 , $s' \geq s$ and $s \not\geq s'$ for all $s \in S$, while in S_2 , $s \geq s'$ and $s' \not\geq s$ for all $s \in S$. For an object $X \in \mathcal{C}$, a *morphism from X to D* (resp. a *morphism from D to X*) is a diagram D' supported on S_1 (resp. on S_2) with $D'(s') = X$ and $D'_S = D$; let $\text{Mor}(X, D)$ (resp. $\text{Mor}(D, X)$) denote the set of these morphisms. A *limit* (resp. *colimit*) of a diagram D is an object $X \in \mathcal{C}$ representing the functor $Y \mapsto \text{Mor}(Y, D)$ (resp. the functor $Y \mapsto \text{Mor}(D, Y)$); by construction, a (co)limit is unique up to unique isomorphism if it exists.

Remark 2.27 Note that every diagram in the category of affine p -schemes has a limit, which can be constructed by repeatedly constructing products and equalizers. (Arbitrary products are given by “infinite tensor products”, which are generated by terms which have the factor 1 in all but finitely many places.) However, a diagram in the category of *connected* affine p -schemes need not have a limit.

Definition 2.28 Let D be a nonempty diagram in the category of connected affine p -schemes. A *p -limit* of D is a connected affine p -scheme Y equipped with a morphism $Y \rightarrow D$, which becomes a colimit of D in the category of abelian groups upon applying the contravariant functor $X \mapsto \text{coker}(F - 1, \Gamma(X, \mathcal{O}_X))$.

Remark 2.29 Note that D admits a limit X in the category of affine p -schemes, and that if Y is a p -limit of D , then the induced homomorphism $\Gamma(X, \mathcal{O}_X) \rightarrow \Gamma(Y, \mathcal{O}_Y)$ is p -surjective: the direct sum of the $\text{coker}(F - 1, \Gamma(Z, \mathcal{O}_Z))$ for all Z in the diagram surjects onto $\text{coker}(F - 1, \Gamma(Y, \mathcal{O}_Y))$, but this surjection factors through $\text{coker}(F - 1, \Gamma(X, \mathcal{O}_X))$.

Theorem 2.30 Let D be a nonempty finite diagram in the category of connected affine p -schemes, let S be the support of D , and let $Y \rightarrow D$ be a morphism. For $s \in S$, let f_s be the induced morphism from Y to $D(s)$. Choose a geometric point \bar{y} of Y . Then Y is

a p -limit of D if and only if $\pi_1^p(Y, \bar{\gamma})$ is a limit, in the category of pro- p -groups, of the diagram induced by D on the $\pi_1^p(D(s), f_s(\bar{\gamma}))$.

Proof First suppose that Y is a p -limit of D . Then the homomorphism

$$(2.1) \quad \pi_1^p(Y, \bar{\gamma}) \rightarrow \lim \pi_1^p(D(s), f_s(\bar{\gamma}))$$

is seen to be injective as follows. Given a non-identity element τ of $\pi_1^p(Y, \bar{\gamma})$, choose a connected p -typical cover E of Y such that τ acts nontrivially on $E_{\bar{\gamma}}$. By Proposition 2.12, E admits an AS-tower; by Remark 2.29 and Proposition 2.22, that AS-tower can be obtained by pullback from some AS-tower over a limit of D in the category of affine p -schemes. Hence the image of τ in $\lim \pi_1^p(D(s), f_s(\bar{\gamma}))$ is not the identity element, so (2.1) is injective.

Suppose now that (2.1) fails to be surjective. We can then construct a nontrivial continuous homomorphism $g: \lim \pi_1^p(D(s), f_s(\bar{\gamma})) \rightarrow \mathbb{Z}/p\mathbb{Z}$ (for the discrete topology on $\mathbb{Z}/p\mathbb{Z}$) whose restriction to $\pi_1^p(Y, \bar{\gamma})$ is trivial. For $s \in S$, put $C_s = \text{coker}(F - 1, \Gamma(D(s), \mathcal{O}))$. We then obtain from g and Proposition 2.15 an element $c_s \in C_s$ for each $s \in S$, such that if $s \rightarrow t$ is a morphism in S , then the corresponding morphism $C_t \rightarrow C_s$ carries c_t to c_s . Since Y is a p -limit, the c_s correspond to a nonzero element of $\text{coker}(F - 1, \Gamma(Y, \mathcal{O}_Y))$, which gives rise to a nontrivial $\mathbb{Z}/p\mathbb{Z}$ -torsor on Y . This contradicts the fact that g restricts trivially to $\pi_1^p(Y, \bar{\gamma})$; the contradiction yields the surjectivity of (2.1) as desired.

We have now shown that if Y is a p -limit of D , then (2.1) is an isomorphism. Suppose now conversely that (2.1) is an isomorphism. Then the maximal elementary abelian quotient of $\pi_1^p(Y, \bar{\gamma})$ is the limit, in the category of elementary abelian p -groups, of the maximal elementary abelian quotients of the $\pi_1^p(D(s), f_s(\bar{\gamma}))$. But by Proposition 2.15, these quotients are dual to the cokernels of $F - 1$ on these schemes. Hence Y is a p -limit of D , as desired. ■

Theorem 2.30 may be a bit obscure as written; some of its corollaries may be more edifying.

Definition 2.31 Let $f: R \rightarrow R'$ be a homomorphism of connected p -rings, and let F and F' be the p -power Frobenius maps on R and R' , respectively. We say f is p -faithful if the induced map $\text{coker}(F - 1) \xrightarrow{f} \text{coker}(F' - 1)$ is a bijection.

Corollary 2.32 Let $f: R \rightarrow R'$ be a homomorphism of connected p -rings, choose a geometric point \bar{x}' of $\text{Spec } R'$, and put $\bar{x} = f(\bar{x}')$. Then f is p -faithful if and only if $\pi_1^p(\text{Spec } R', \bar{x}') \xrightarrow{f} \pi_1^p(\text{Spec } R, \bar{x})$ is a bijection.

Example 2.33 For any p -ring R , the canonical inclusion $f: R[t^{-1}] \rightarrow R((t))$ is p -faithful: the kernel of f is trivial, and the cokernel of f is isomorphic as a Frobenius module to $tR[[t]]$, on which $F - 1$ is bijective. The conclusion of Corollary 2.32 in this case is a result of Katz [5, Proposition 1.4.2]. Although Katz’s proof looks different (it involves manipulating the cohomology of pro- p -groups), our proof is basically a transcription of Katz’s argument into the language of AS-towers.

Example 2.34 Let R be a p -ring. Consider the diagram consisting of the two natural maps $\text{Spec } R[t] \rightarrow \text{Spec } R$ and $\text{Spec } R[t^{-1}] \rightarrow \text{Spec } R$. Then $\text{Spec } R[t, t^{-1}]$ is a p -limit of this diagram; we thus have an isomorphism

$$\pi_1^p(\text{Spec } R[t, t^{-1}]) \rightarrow \pi_1^p(\text{Spec } R[t]) \times_{\pi_1^p(\text{Spec } R)} \pi_1^p(\text{Spec } R[t^{-1}]),$$

after choosing basepoints. (Namely, choose a geometric point of $\text{Spec } R[t, t^{-1}]$ and obtain the other basepoints by applying the maps in the diagram.)

Here is a slight variation of the previous example.

Corollary 2.35 Let R be an $\overline{\mathbb{F}_p}$ -algebra. Then every p -typical extension of $R[t]$ is contained in the tensor product of a p -typical extension of $R[t]$ in which R is integrally closed, and a p -typical extension of $R[t]$ obtained by base change from R .

Proof Put $R' = \overline{\mathbb{F}_p} + tR[t] \subseteq R[t]$. Then $\text{Spec } R[t]$ is a p -limit of the diagram consisting of $\text{Spec } R$ and $\text{Spec } R'$ with no arrows, so by Theorem 2.30, we have

$$\pi_1^p(\text{Spec } R[t]) \cong \pi_1^p(\text{Spec } R) \times \pi_1^p(\text{Spec } R')$$

(for appropriate basepoints). Thus every p -typical extension of $R[t]$ is contained in the tensor product of a p -typical extension obtained by base change from R , and a p -typical extension obtained by base change from $\overline{\mathbb{F}_p} + tR[t]$; in the latter, the restriction to the $t = 0$ locus must split completely, so R must be integrally closed. This yields the desired result. ■

Remark 2.36 This corollary should be a bit surprising: for a general finite étale extension of $R[t]$, or even of $R((t^{-1}))$, one cannot split off the residual extension in this fashion. For instance, if the extension is obtained by adjoining z with $z^p - z = at$ for a in some finite étale extension of R , it is typically impossible to present the extension as in the corollary unless a generates a p -typical extension of R (in which case the corollary applies).

3 Complexity Measures for p -Typical Extensions

We next propose a mechanism for handling the complexity of a p -typical extension, via what we call *height functions*; the mechanism is modeled on basic ramification theory for complete discretely valued fields. As in other instances where complexity-bounding functions arise (e.g., Diophantine approximation, from which the term “height function” was borrowed), it is a bit tedious to introduce and deal with such functions, but things are made a bit easier by the fact that the intended use of these functions permits one to be somewhat sloppy in dealing with them. The reader impatient to get to some meaningful results may wish to skip ahead to the next section before continuing here.

3.1 Ramification Filtrations for Local Fields

The model for our height functions is the highest break function coming from the ramification filtration on the Galois group of $k((t))$, so we start by reviewing that construction. For all unproved assertions in this section, see [8].

Definition 3.1 Let F be a complete discretely valued field whose residue field k is perfect (e.g., the power series field $k((t))$). Let E/F be a finite Galois field extension with group G , let \mathfrak{o}_E and \mathfrak{o}_F be the valuation subrings of E and F , and let v_E be the valuation on E , normalized so that v_E maps E^* onto \mathbb{Z} . For $i \geq -1$, let G_i be the subgroup of $g \in G$ for which $v_E(a^g - a) \geq i + 1$ for all $a \in \mathfrak{o}_E$; the decreasing filtration $\{G_i\}$ is called the *lower numbering filtration* of G [8, §IV.1].

Definition 3.2 With notation as in Definition 3.1, define the function

$$\phi_{E/F}(u) = \int_0^u \frac{dt}{[G_0 : G_t]}.$$

Then $\phi_{E/F}$ is a homeomorphism of $[-1, \infty)$ with itself; let $\psi_{E/F}$ denote the inverse function. Define the *upper numbering filtration* of G by $G^i = G_{\psi_{E/F}(i)}$ [8, §IV.3]. It has the property that if E'/F is a Galois subextension of E/F with Galois group H , then the image of each G^i under the natural surjection $G \twoheadrightarrow H$ is precisely H^i ; this follows from Herbrand’s theorem [8, Proposition IV.14].

Definition 3.3 For F as in Definition 3.1 and E/F a finite Galois field extension, define the *highest break* of E , denoted $b(E/F)$, to be the largest i such that $G^i \neq G^j$ for any $j > i$, or zero if no such i exists. If E/F is a field extension which is finite separable but not Galois, we define $b(E/F) = b(E'/F)$, for E'/F the Galois closure of E/F . If E is not a field, but only an étale F -algebra, we define $b(E/F)$ to be the maximum highest break of any component of E . With these rules, one has the following properties.

- (i) $b(F/F) = 0$ (evident).
- (ii) If E' is an F -subalgebra of E , then $b(E'/F) \leq b(E/F)$ (evident).
- (iii) $b((E_1 \oplus E_2)/F) = \max\{b(E_1/F), b(E_2/F)\}$ (formal).
- (iv) $b((E_1 \otimes E_2)/F) = \max\{b(E_1/F), b(E_2/F)\}$ (not formal, but follows from Herbrand’s theorem).
- (v) If E/F is a Galois field extension and E' is an étale E -algebra, then $b(E'/F) = \max\{b(E/F), \phi_{E/F}(b(E'/E))\}$ (because the lower numbering is stable under taking subgroups).

Remark 3.4 In case k is imperfect, there are several competing analogues of the upper numbering filtration; these include the “residual perfection” construction of Borger [2], and the “nonlogarithmic” and “logarithmic” rigid geometric constructions of Abbes and Saito [1]. We will not use any of these in this paper.

3.2 Artin–Schreier Extensions and Highest Breaks

We next recall some standard facts about Artin–Schreier extensions of a power series field.

Lemma 3.5 For k a perfect p -field, and for $a \in F = k((t))$, put

$$m = \inf_{x \in F} \{-v_F(a - x^p + x)\}$$

and put $E = F[z]/(z^p - z - a)$. Then the following hold.

- (i) Either $m = -\infty$ (that is, E is not a field) or $m \geq 0$.
- (ii) If $m \geq 0$, then the extension E/F is unramified if and only if $m = 0$.
- (iii) If $m > 0$, then m is not divisible by p , and E/F has highest break m .

Proof (i) Suppose $m < 0$, which means that there exists $y \in F$ such that $b = y^p - y - a$ satisfies $v_F(b) > 0$. Then the series $b + b^p + b^{p^2} + \dots$ converges in F , and its limit c satisfies $c - c^p = b$. This yields $a = (c + y)^p - (c + y)$, and so $m = -\infty$.

(ii) Note that $m = 0$ implies, as in (i), that $y^p - y - a \in k$ for some $y \in F$, and so E is unramified. Conversely, if E is unramified, then the residue field \bar{E} of E must be an Artin–Schreier extension of k , say $k[y]/(y^p - y - b)$ with $b \in k$. If we choose b so that the $\mathbb{Z}/p\mathbb{Z}$ -torsor structures on E/F and \bar{E}/k are compatible, by Proposition 2.15 we must then have $a - b = x^p - x$ for some $x \in F$, yielding $m = 0$.

(iii) If $a = c_{-pn}t^{-pn} + \dots$, then $a - c_{-pn}t^{-pn} + c_{-pn}^{1/p}t^{-n}$ has strictly larger valuation than does a . Hence if m is positive, it cannot be divisible by p . To compute the highest break, assume that $v_F(a) = -m$, pick integers r, s with $r > 0$ and $-rm + sp = 1$, and put $u = z^r t^s$; then $v_E(u) = 1$, i.e., u is a uniformizer of E . By [8, Proposition IV.5], the highest break of E/F equals $v_E(u'/u - 1)$, where u' is the image of u under the automorphism $z \mapsto z + 1$ of E . Since r is not divisible by p , we have

$$\frac{u'}{u} = (z + 1)^r z^{-r} = 1 + rz^{-1} + \dots$$

and so $v_E(u'/u - 1) = v_E(z^{-1}) = m$, as desired. ■

One can also obtain a bound on the highest break in an AS-tower. The following bound is not optimal, but it suffices for our purposes.

Corollary 3.6 Let k be a perfect p -field, let $k((t)) = E_0 \subset E_1 \subset \dots \subset E_d = E$ be an AS-tower, and choose $\ell \geq 1$ such that for $i = 1, \dots, m$, $E_i \cong E_{i-1}[z]/(z^p - z - c_i)$ for some c_i with $v_{E_0}(c_i) \geq -\ell$. Then $b(E/E_0) \leq d\ell$.

Proof We proceed by induction on d , the case $d = 1$ following from Lemma 3.5. For $d > 1$, if E_1/E_0 is disconnected, then we can correspondingly split E as a direct sum $E'_1 \oplus \dots \oplus E'_p$, in which for $j = 1, \dots, p$, E'_j admits an AS-tower of length $d - 1$ over E_0 . By the induction hypothesis, we have $b(E/E_0) = \max_j \{b(E'_j/E_0)\} \leq (d - 1)\ell$.

If E_1/E_0 is connected, by Lemma 3.5 we have $b(E_1/E_0) = m$ for some nonnegative integer $m \leq \ell$, and by the induction hypothesis we have $b(E/E_1) \leq (d - 1)(p\ell)$. For $x \geq m$, $\phi_{E_1/E_0}(x) = m + (x - m)/p$, so

$$\begin{aligned} b(E/E_0) &\leq \phi_{E_1/E_0}((d - 1)p\ell) = m + \frac{(d - 1)p\ell - m}{p} \\ &= (d - 1)\ell + \frac{m(p - 1)}{p} \\ &\leq d\ell, \end{aligned}$$

as desired. ■

We also need to know that the highest break drops under specialization.

Proposition 3.7 *Let $R \rightarrow R'$ be a surjective morphism of perfect p -domains, let S be a p -typical extension of $R((t))$, and put $S' = S \otimes_{R((t))} R'((t))$. Let K and K' be the fraction fields of R and R' , respectively. Then*

$$b(S \otimes_{R((t))} K((t))/K((t))) \geq b(S' \otimes_{R'((t))} K'((t))/K'((t))).$$

Proof This follows from the Deligne–Laumon semicontinuity theorem [6]. ■

3.3 Presentations of AS-Towers

To talk about height functions on p -typical extensions of more general rings, we need to fix a bit of terminology concerning presentations of AS-towers.

Definition 3.8 Given an AS-tower $R = S_0 \subset S_1 \subset \dots \subset S_d = S$ over a p -ring R , a *presentation* of S is a sequence of isomorphisms

$$S_i \cong S_{i-1}[z_i]/(z_i^p - z_i - P_i(z_1, \dots, z_{i-1})) \quad (i = 1, \dots, d),$$

where $P_i(z_1, \dots, z_{i-1})$ is a polynomial over R of degree at most $p - 1$ in each variable; by Proposition 2.15, such a presentation always exists. Given a presentation of S , each element $x \in S$ can be written uniquely as a polynomial in z_1, \dots, z_d over R with degree at most $p - 1$ in each variable; we call this polynomial the *minimal representation* of x .

In terms of presentations, one has the following evident but useful lemma.

Lemma 3.9 *Given an AS-tower $R = S_0 \subset S_1 \subset \dots \subset S_d = S$ over a p -ring R , and a presentation*

$$S_i \cong S_{i-1}[z_i]/(z_i^p - z_i - P_i(z_1, \dots, z_{i-1})) \quad (i = 1, \dots, d)$$

of S , choose integers $j_1, \dots, j_d \in \{0, \dots, p - 1\}$, and put $x = z_1^{j_1} \dots z_d^{j_d}$. Then the minimal representation of x^p , written as a polynomial in z_d over S_{d-1} , is monic of degree j_d .

Proof Note that for each i , z_i^p can be rewritten as z_i plus a polynomial in the preceding variables; this implies the claim. ■

Definition 3.10 If V is an additive subgroup of R , we say a presentation of S is *defined over V* if each P_i has its coefficients in V .

3.4 Height Functions

Definition 3.11 Let R_0 be a connected p -ring, and let R be a connected R_0 -algebra. A *height function (over R_0)* on \mathcal{C}_R^p (the category of p -typical extensions of R) is a function h from the set of isomorphism classes of elements of \mathcal{C}_R^p to the nonnegative real numbers, having the following properties.

- (i) $h(S_1 \oplus S_2)$ is bounded above by some function of $h(S_1)$, $h(S_2)$, $\deg(S_1/R)$, $\deg(S_2/R)$.
- (ii) $h(S_1 \otimes S_2)$ is bounded above by some function of $h(S_1)$, $h(S_2)$, $\deg(S_1/R)$, $\deg(S_2/R)$.
- (iii) If $S_1 \subseteq S_2$, then $h(S_1)$ is bounded above by some function of $h(S_2)$, $\deg(S_2/R)$.
- (iv) For any positive integer d and any finite R_0 -submodule V of R , there exists a nonnegative real number ℓ such that for any connected AS-tower $R = S_0 \subset S_1 \subset \dots \subset S_d = S$ admitting a presentation defined over V , we have $h(S) \leq \ell$.
- (v) For any positive integer d and any nonnegative real number ℓ , there exists a finite R_0 -submodule V of R such that for any connected AS-tower $R = S_0 \subset S_1 \subset \dots \subset S_d = S$ with $h(S) \leq \ell$, there exists a presentation of S defined over V .

We say h is a *strong height function* if the following additional conditions hold.

- (i') $h(S_1 \oplus S_2) \leq \max\{h(S_1), h(S_2)\}$.
- (ii') $h(S_1 \otimes S_2) \leq \max\{h(S_1), h(S_2)\}$.
- (iii') If $S_1 \subseteq S_2$, then $h(S_1) \leq h(S_2)$.

We extend a height function to continuous homomorphisms $\rho: \pi_1^p(\text{Spec } R, \bar{x}) \rightarrow G$, for \bar{x} a geometric point of $\text{Spec } R$ and G a finite discrete group, by declaring that $h(\rho) = h(S)$, where $S \in \mathcal{C}_R^p$ is chosen so that $\pi_1^p(\text{Spec } S, \bar{y})$ is the kernel of ρ (for an appropriate geometric point \bar{y} of $\text{Spec } S$).

Lemma 3.12 With notation as in Definition 3.11, let R' be a connected p -typical extension of R . Then any height function h over R_0 on \mathcal{C}_R^p induces a height function h' over R_0 on $\mathcal{C}_{R'}^p$ (given by $h'(S) = h(S)$).

Proof Straightforward. ■

The condition (v) is not so easy to check directly, but fortunately one need only verify it for Artin–Schreier extensions, as confirmed by the following proposition.

Proposition 3.13 Given conditions (i)–(iv) of Definition 3.11, if condition (v) holds for $d = 1$, then it holds for all d .

Proof We proceed by induction on d (simultaneously for all ℓ), the case $d = 1$ being the input hypothesis. Given the claim for $d - 1$ and given a connected AS-tower $R = S_0 \subset S_1 \subset \dots \subset S_d = S$ with $h(S) \leq \ell$, we may choose a presentation for S_{d-1} over some finite R_0 -module depending only on d and ℓ .

Write $S_d = S_{d-1}[z_d]/(z_d^p - z_d - a_{d-1})$ and write $a_{d-1} = \sum_{i=0}^{p-1} c_i z_{d-1}^i$ for $c_i \in S_{d-2}$. Let j be the degree of a_{d-1} as a polynomial in z_{d-1} , so that $c_j \neq 0$ but $c_{j+1} = \dots = c_{p-1} = 0$. We prove that for some $w \in S_{d-1}$ of degree at most j as a polynomial in z_{d-1} , the coefficients in the minimal representation of $a_{d-1} - w^p + w$ lie in some finite R_0 -module depending only on d, ℓ, j . The proof of this claim constitutes an inner induction on j .

For $j = 0$, we may apply the outer induction hypothesis to $S_{d-2}[z]/(z^p - z - c_0)$ to deduce the claim. Otherwise, let g be the automorphism of S_{d-1} over S_{d-2} given by $z_{d-1} \mapsto z_{d-1} + 1$, and define the map $\Delta: S_{d-1} \rightarrow S_{d-1}$ by $\Delta(x) = x^g - x$. Then the j -th tensor power of S_d over S_{d-2} , which has height bounded by a function of d, ℓ, j by (ii), contains

$$S_{d-1}[z]/(z^p - z - \Delta^{(j)}(a_{d-1})).$$

Since $\Delta^{(j)}(\sum_{i=0}^j c_i z_{d-1}^i) = j!c_j$, we deduce that $S_{d-2}[z]/(z^p - z - j!c_j)$ has height bounded by a function of d, ℓ, j .

Applying the outer induction hypothesis yields $w' \in S_{d-2}$ such that the minimal representation of $b' = c_j - (w')^p + w'$ has coefficients in some finite R_0 -module depending only on d, ℓ, j . Put $a'_{d-1} = a_{d-1} - (w'z'_{d-1})^p + w'z'_{d-1}$, so that $S_d = S_{d-1}[z]/(z^p - z - a'_{d-1})$. Then $a'_{d-1} - b'z'_{d-1}$ has degree at most $j - 1$ as a polynomial in z_{d-1} . If we put

$$S' = S_{d-1}[x]/(x^p - x - b'z'_{d-1}), \quad S'' = S_{d-1}[y]/(y^p - y - a'_{d-1} - b'z'_{d-1}),$$

then the height of S' is bounded by some function of d, ℓ, j by condition (iv) of Definition 3.11. On the other hand, S'' is contained in $S' \otimes_{S_{d-1}} S_d$, and the heights of S_d and S' are bounded by some function of d, ℓ, j , so the same is true of S'' . Applying the inner induction hypothesis to S'' , we obtain $w'' \in S_{d-1}$ of degree at most $j - 1$ as a polynomial in z_{d-1} , such that $a'_{d-1} - b'z'_{d-1} - (w'')^p + w''$ has coefficients in some finite R_0 -module depending only on d, ℓ, j . We may then take $w = w'z'_{d-1} + w''$, as

$$a_{d-1} - w^p + w = b'z'_{d-1} + (a'_{d-1} - b'z'_{d-1} - (w'')^p + w'')$$

has coefficients in some finite R_0 -module depending on d, ℓ, j . This completes the proof of the inner induction.

The inner induction for $j = p - 1$ implies the outer induction, so the proof is complete. ■

Example 3.14 For $R = k((t))$ with k a perfect p -field, the highest break function $h(S) = b(S/k((t)))$ is a strong height function on $\mathcal{C}_{k((t))}^p$ over k : properties (i'), (ii'), (iii') follow from Definition 3.3, property (iv) from Corollary 3.6, and property (v) for $d = 1$ from Lemma 3.5.

Remark 3.15 Already for $R = k((t))$ with k an imperfect p -field, it is less than evident how to construct a height function on $\mathcal{C}_{k((t))}^p$ over k , since the naïve highest break function $b(S \otimes k^{\text{perf}}((t))/k^{\text{perf}}((t)))$ will not do. To see this, choose $c \in k \setminus k^p$, then note that the heights of $k((t))[z]/(z^p - z - ct^{-p^n})$ would all be equal to 1, whereas these extensions do not simultaneously admit presentations defined over some finite dimensional k -vector space. It should be possible to extract a height function from any of the constructions of a ramification filtration mentioned in Remark 3.4, but we have not attempted to do so.

4 p -Typical Covers of Affine Toric Varieties

In this section, we study the p -typical fundamental groups of affine toric varieties; while the case of ordinary affine space is doubtless the most important, it is not any harder to work in this generality. Our main results in this direction are some decomposition theorems for these fundamental groups (Theorem 4.15 and its consequence Theorem 4.17).

Convention 4.1 Throughout this section, let R denote a connected p -ring.

4.1 Some Toric Rings

Definition 4.2 Define a *convex cone* in \mathbb{R}^n as a nonempty subset $\sigma \subseteq \mathbb{R}^n$ such that

- (i) if $\mathbf{v} \in \sigma$, then $c\mathbf{v} \in \sigma$ for any $c \in \mathbb{R}_{\geq 0}$;
- (ii) if $\mathbf{v}, \mathbf{w} \in \sigma$, then $c\mathbf{v} + (1 - c)\mathbf{w} \in \sigma$ for any $c \in [0, 1]$.

Note that the intersection of convex cones is again a convex cone. We say the convex cone σ is *finitely generated* if it can be written as a finite intersection of open and closed halfspaces.

Definition 4.3 Given a convex cone σ , let R_σ denote the monoid algebra $R[\sigma \cap \mathbb{Z}^n]$. For convex cones σ, τ with $\sigma \subseteq \tau$, there is a natural inclusion $R_\sigma \subseteq R_\tau$. Given an element $x \in R_\sigma$, write

$$x = \sum_{\mathbf{v} \in \sigma \cap \mathbb{Z}^n} c_{\mathbf{v}}[\mathbf{v}],$$

and define the *support* of x to be the set of $\mathbf{v} \in \sigma \cap \mathbb{Z}^n$ such that $c_{\mathbf{v}} \neq 0$.

Remark 4.4 If σ is a convex cone equal to the intersection of finitely many closed halfspaces defined by linear functionals over \mathbb{Q} , then $\text{Spec } R_\sigma$ is an affine toric variety, and conversely. (Note that in our terminology, toric varieties are necessarily normal.)

Convention 4.5 For the rest of the section, fix a geometric point \bar{x} of $\text{Spec } R_{\mathbb{R}^n}$; we may also view \bar{x} as a geometric point of $\text{Spec } R_\sigma$ for any convex cone $\sigma \subseteq \mathbb{R}^n$. We will thus drop this basepoint from the notation when considering the fundamental group of $\text{Spec } R_\sigma$.

Proposition 4.6 Suppose that $\sigma, \sigma_1, \dots, \sigma_n$ are convex cones with

$$\sigma = \sigma_1 \cup \dots \cup \sigma_n.$$

Then $\text{Spec } R_\sigma$ is a p -limit of the diagram consisting of the arrows $\text{Spec } R_{\sigma_i} \rightarrow \text{Spec } R_{\sigma_i \cap \sigma_j}$ for $1 \leq i, j \leq n$. Consequently, the group $\pi_1^p(\text{Spec } R_\sigma)$ is a limit of the diagram consisting of the arrows $\pi_1^p(\text{Spec } R_{\sigma_i}) \rightarrow \pi_1^p(\text{Spec } R_{\sigma_i \cap \sigma_j})$ for $1 \leq i, j \leq n$.

Proof It suffices to note that the cokernel of $F - 1$ on R_σ is generated freely by the images of $\sigma \cap (\mathbb{Z}^n \setminus p\mathbb{Z}^n)$. This yields the first assertion; the second follows by Theorem 2.30. ■

4.2 Projections and Sections

Definition 4.7 A convex cone σ is strictly convex if for $\mathbf{v}, \mathbf{w} \in \sigma, \mathbf{v} + \mathbf{w} = 0$ if and only if $\mathbf{v} = \mathbf{w} = 0$. For σ a strictly convex cone, let R'_σ be the subring of R_σ consisting of elements $\sum_{\mathbf{v}} c_{\mathbf{v}}[\mathbf{v}]$ with $c_0 \in \mathbb{F}_p$. (Note that strict convexity is needed for this subset to be closed under multiplication.)

Proposition 4.8 Suppose that σ and σ_0 are convex cones, and $\{\sigma_i\}_{i \in I}$ is a (not necessarily finite) collection of strictly convex cones, such that $\sigma \setminus \{0\}$ is the disjoint union of $\sigma_0 \setminus \{0\}$ and the $\sigma_i \setminus \{0\}$. Then the natural map

$$\pi_1^p(\text{Spec } R_\sigma) \rightarrow \pi_1^p(\text{Spec } R_{\sigma_0}) \times \prod_{i \in I} \pi_1^p(\text{Spec } R'_{\sigma_i})$$

is an isomorphism.

Proof The argument is as in Proposition 4.6. ■

Definition 4.9 Let σ, τ be convex cones with $\tau \subseteq \sigma$. Put $\sigma_0 = \tau$, and choose a collection $\{\sigma_i\}_{i \in I}$ of strictly convex cones such that $\sigma \setminus \{0\}$ is the disjoint union of $\sigma_0 \setminus \{0\}$ and the $\sigma_i \setminus \{0\}$. Then the product decomposition given by Proposition 4.8 yields a morphism $\pi_{\sigma, \tau}: \pi_1^p(\text{Spec } R_\tau) \rightarrow \pi_1^p(\text{Spec } R_\sigma)$ sectioning the projection $\pi_1^p(\text{Spec } R_\sigma) \rightarrow \pi_1^p(\text{Spec } R_\tau)$. Note that replacing one of the σ_i by a disjoint union does not affect $\pi_{\sigma, \tau}$. In particular, by passing to a common refinement, we see that this map does not depend at all on the choice of the σ_i .

Proposition 4.10 Let σ, τ be convex cones with $\tau \subseteq \sigma$. Let $\rho: \pi_1^p(\text{Spec } R_\sigma) \rightarrow \mathbb{Z}/p\mathbb{Z}$ be the homomorphism corresponding to the $\mathbb{Z}/p\mathbb{Z}$ -torsor $S = R_\sigma[z]/(z^p - z - x)$ over R_σ . Write $x = \sum_{\mathbf{v} \in \sigma \cap \mathbb{Z}^n} c_{\mathbf{v}}[\mathbf{v}]$. Then $\rho \circ \pi_{\sigma, \tau}: \pi_1^p(\text{Spec } R_\tau) \rightarrow \mathbb{Z}/p\mathbb{Z}$ is the homomorphism corresponding to the $\mathbb{Z}/p\mathbb{Z}$ -torsor

$$R_\tau[z]/(z^p - z - x'), \quad x' = \sum_{\mathbf{v} \in \tau \cap \mathbb{Z}^n} c_{\mathbf{v}}[\mathbf{v}].$$

Proof This is an immediate consequence of how the map in Theorem 2.30 is constructed. ■

We can make the maps $\pi_{\sigma,\tau}$ more explicit in certain cases as follows.

Definition 4.11 Let $\lambda: \mathbb{R}^n \rightarrow \mathbb{R}$ be a linear function. For σ a convex cone, let v_λ be the valuation on R_σ defined by

$$v_\lambda\left(\sum_{\mathbf{v} \in \sigma \cap \mathbb{Z}^n} c_{\mathbf{v}}[\mathbf{v}]\right) = \min\{\lambda(\mathbf{v}) : \mathbf{v} \in \sigma \cap \mathbb{Z}^n, c_{\mathbf{v}} \neq 0\}.$$

Let $R_{\sigma,\lambda}$ be the completion of R_σ with respect to v_λ , and put

$$\tau_{\sigma,\lambda} = \{\mathbf{v} \in \sigma : \lambda(\mathbf{v}) \leq 0\}.$$

Proposition 4.12 Set notation as in Definition 4.11, and write τ for $\tau_{\sigma,\lambda}$. Then the natural map $\pi_1^p(\text{Spec } R_{\sigma,\lambda}) \rightarrow \pi_1^p(\text{Spec } R_\tau)$ is an isomorphism.

Proof If $z \in R_{\sigma,\lambda}$ and $v_\lambda(z) > 0$, then $z + z^p + \dots$ converges in $R_{\sigma,\lambda}$ to some y satisfying $y^p - y = -z$. Thus the cokernels of $F - 1$ on R_τ and $R_{\sigma,\lambda}$ are isomorphic, so Theorem 2.30 applies. ■

Example 4.13 Simple examples of Proposition 4.12 are the facts that

$$\pi_1^p(\text{Spec } R[[t]]) \rightarrow \pi_1^p(\text{Spec } R) \quad \text{and} \quad \pi_1^p(\text{Spec } R((t))) \rightarrow \pi_1^p(\text{Spec } R(t^{-1}))$$

are isomorphisms. For a more nontrivial example, take σ to be the nonnegative quadrant in \mathbb{R}^2 , and define $\lambda(a, b) = a - b$. Then Proposition 4.12 asserts that

$$\pi_1^p(\text{Spec}(R[x,y][[x]])) \rightarrow \pi_1^p(\text{Spec } R[x,y])$$

is an isomorphism.

4.3 Heights and Representations

Definition 4.14 A linear cone in \mathbb{R}^n is a convex cone consisting of the nonnegative scalar multiples of a single nonzero element of \mathbb{R}^n . For S a set of linear cones and σ a convex cone, define $S_\sigma = \{\tau \in S : \tau \subseteq \sigma\}$.

Theorem 4.15 Let σ be a convex cone in \mathbb{R}^n , let h be a height function on \mathbb{C}_{R_σ} over R , let ℓ be a nonzero real number, and let G be a finite discrete group. Then there exists a finite set S of linear cones in \mathbb{R}^n , depending on σ, h, ℓ, G , such that for any continuous representation $\rho: \pi_1^p(\text{Spec } R_\sigma) \rightarrow G$ with $h(\rho) \leq \ell$ and any convex cone $\tau \subseteq \sigma$, the image $(\rho \circ \pi_{\sigma,\tau})(\pi_1^p(\text{Spec } R_\tau))$ is determined by ρ and S_τ .

Proof We first check the case $G = \mathbb{Z}/p\mathbb{Z}$. If ρ is trivial, there is nothing to check; otherwise, ρ becomes trivial upon restriction to $\pi_1^p(\text{Spec } S)$ for some $\mathbb{Z}/p\mathbb{Z}$ -torsor S over R_σ . By Proposition 2.15, we may write $S = R_\sigma[z]/(z^p - z - x)$, and we may choose $x \in R_\sigma$ with support in $\{0\} \cup (\sigma \cap (\mathbb{Z}^n \setminus p\mathbb{Z}^n))$. Since $h(\rho) \leq \ell$, by (v) the support of x is contained in a finite set T depending on σ, h, ℓ . By Proposition 4.10, the claim holds if we take S to be the set of linear cones generated by the elements of T . Note that this set depends only on σ, h, ℓ , and not on ρ .

We next pass to the general case. We may assume that G is the image of ρ , so that G is a p -group and ρ is surjective; we may also assume that G is nontrivial. Let K be the Frattini subgroup of G (the intersection of its maximal proper subgroups), so that G/K is an elementary abelian p -group. By repeatedly applying the previous paragraph, we obtain a finite set S_1 of linear cones, determined by σ, h, ℓ , such that the image of $\pi_1^p(\text{Spec } R_\tau)$ in G/K is determined by ρ and $(S_1)_\tau$.

We now induct on the size of (the smallest possible choice of) S_1 . If S_1 is empty, then the image of $\pi_1^p(\text{Spec } R_\tau)$ in G/K is equal to the image of $\pi_1^p(\text{Spec } R_\sigma)$ in G/K , namely G/K itself. Thus the image of $\pi_1^p(\text{Spec } R_\tau)$ in G cannot be contained in any proper subgroup of G , and so must equal G . Thus we are done in this case.

If S_1 is nonempty, choose a linear cone T in S_1 . We can then choose strictly convex cones $\sigma_1, \dots, \sigma_m$ not meeting T such that $\sigma \setminus T$ is the union of $\sigma_1 \setminus \{0\}, \dots, \sigma_m \setminus \{0\}$. (Namely, draw $n - 1$ hyperplanes meeting transversely along T , take the open halfspaces on both sides of each plus one halfspace containing the negation of T , and intersect all of these with σ .) We may now apply the induction hypothesis to each of the σ_i (since the analogue of the set S_1 has been reduced by one element) to produce a finite set S_T (determined by σ_i, h, ℓ, G) such that if $T \not\subseteq \tau$, then $\rho(\pi_1^p(\text{Spec } R_\tau))$ is determined by ρ and $(S_T)_\tau$. Let S be the union of the S_T . This has the desired property because if every T lies in τ , then the image of $\pi_1^p(\text{Spec } R_\tau)$ in G/K must equal G/K , so as in the base case, $\rho(\pi_1^p(\text{Spec } R_\tau)) = G$. ■

Corollary 4.16 *Let σ be a convex cone in \mathbb{R}^n , let h be a height function on \mathcal{C}_{R_σ} over R , let ℓ be a nonzero real number, and let G be a finite discrete group. Then there exists a finite set S of linear cones in \mathbb{R}^n , depending on σ, h, ℓ, G , such that for any continuous representations $\rho_1, \rho_2 : \pi_1^p(\text{Spec } R_\sigma) \rightarrow G$ with $h(\rho_1), h(\rho_2) \leq \ell$ and any convex cone $\tau \subseteq \sigma$, whether or not the restrictions of ρ_1 and ρ_2 to $\pi_1^p(\text{Spec } R_\tau)$ are isomorphic is determined by ρ_1, ρ_2, S_τ .*

Proof Embed G into a linear group over a field of characteristic zero, and apply Theorem 4.15 to the representation $\rho_1^\vee \times \rho_2 : \pi_1^p(\text{Spec } R_\sigma) \rightarrow G^T \times G$. (Here ρ_1^\vee denotes the contragredient representation and G^T denotes G with its linear embedding transposed.) ■

The next corollary is sufficiently useful in its own right that we have promoted it to a theorem.

Theorem 4.17 *Fix a convex cone σ . For $T = \{\tau_1, \dots, \tau_m\}$ a collection of distinct linear cones contained in σ , let G_T be the limit of the diagram consisting of the arrows*

$\pi_1^p(\text{Spec } R_{\pi_i}) \rightarrow \pi_1^p(\text{Spec } R)$ for $i = 1, \dots, m$. View the G_T as an inverse system via the natural maps $G_{T'} \rightarrow G_T$ for $T \subseteq T'$. Then $\pi_1^p(\text{Spec } R_\sigma)$ is the inverse limit of the G_T .

A weaker but coordinate-free variant of Theorem 4.17 is the following.

Corollary 4.18 Fix a convex cone σ . For $S = \{R_1, \dots, R_m\}$ a set whose elements are subalgebras of R_σ of transcendence degree 1 over R , let G_S be the limit of the diagram consisting of the arrows $\pi_1^p(\text{Spec } R_i) \rightarrow \pi_1^p(\text{Spec}(R_i \cap R_j))$ for $i, j = 1, \dots, m$. View the G_S as an inverse system via the natural maps $G_{S'} \rightarrow G_S$ for $S \subseteq S'$. Then $\pi_1^p(\text{Spec } R_\sigma)$ is the inverse limit of the G_S .

Finally, it is worth saying in simple terms what Theorem 4.17 says about affine spaces.

Corollary 4.19 For n a positive integer n , take x_1, \dots, x_n to be coordinates on \mathbb{A}_R^n . Then the group $\pi_1^p(\mathbb{A}_R^n)$ is the limit of the diagram consisting of

$$\pi_1^p(\text{Spec } R[x_1^{a_1} \cdots x_n^{a_n}]) \rightarrow \pi^p(\text{Spec } R)$$

for all coprime n -tuples (a_1, \dots, a_n) of nonnegative integers.

5 Complements on Height Functions

To conclude, we point out that the somewhat mysterious height functions that we have been using can be made quite explicit on affine toric varieties. The main result here is Theorem 5.11, which gives a relatively simple formula for a function which can be verified (Corollary 5.18) to be a height function.

Note that we use Theorem 4.17 in the course of the proof. We do not know whether it is possible to prove Theorem 5.11 directly, then short-circuit the proof of Theorem 4.15 and its consequences around the discussion of general height functions. Doing so might necessitate establishing a relationship between ramification theory for local fields with imperfect residue field (see Remark 3.4); such a relationship might have the effect of clarifying the ramification theory in some cases.

5.1 Some Explicit Height Functions

In the situation we have been considering, we can write down some height functions explicitly.

Convention 5.1 Throughout this section, let $R = k$ be an algebraically closed p -field.

Definition 5.2 For $\lambda: \mathbb{R}^n \rightarrow \mathbb{R}$ a nonzero linear function defined over \mathbb{Q} (i.e., it carries \mathbb{Q}^n to \mathbb{Q}), let m_λ be the unique rational number such that $m_\lambda \lambda(\mathbb{Z}^n) = \mathbb{Z}$, let H_λ denote the hyperplane $\{\mathbf{v} \in \mathbb{R}^n : \lambda(\mathbf{v}) = 0\}$, and let K_λ denote the perfection of

the fraction field of R_{H_λ} . Let \widehat{R}_λ denote the completion of $R_{\mathbb{R}^n}$ with respect to $v_{-\lambda}$, and put

$$Q_\lambda = \widehat{R}_\lambda \otimes_{R_{H_\lambda}} K_\lambda;$$

we may then view Q_λ as a power series field in one variable over the perfect field K_λ , with valuation $m_\lambda v_{-\lambda}$. Given a p -typical extension S of $R_{\mathbb{R}^n}$, define

$$c_\lambda(S) = \frac{1}{m_\lambda} b((S \otimes_{R_{\mathbb{R}^n}} Q_\lambda)/Q_\lambda),$$

where b denotes the highest break function (of Definition 3.3).

As in Remark 3.15, b_λ is not a height function for p -typical extensions of R_{H_λ} . However, we can use the functions b_λ to construct height functions on smaller cones.

Definition 5.3 Given a convex cone σ , define the *dual cone* $\sigma^\vee \subseteq (\mathbb{R}^n)^\vee$ as the set of linear functions $\lambda: \mathbb{R}^n \rightarrow \mathbb{R}$ such that $\lambda(\mathbf{v}) \geq 0$ for all $\mathbf{v} \in \sigma$. We say σ is *very convex* if σ^\vee has nonempty topological interior relative to $(\mathbb{R}^n)^\vee$; if σ is very convex, then it is strictly convex.

Definition 5.4 Let σ be a nontrivial very convex cone, and let $U \subseteq \sigma^\vee \setminus \{0\}$ be a subset open in $(\mathbb{R}^n)^\vee$. Define the function h_U on \mathcal{C}_{R_σ} by

$$h_U(S) = \sup_{\lambda \in U \cap (\mathbb{Q}^n)^\vee} \{c_\lambda(S \otimes R_{\mathbb{R}^n})\}.$$

For λ in the interior of $\sigma^\vee \setminus \{0\}$, define

$$h_\lambda(S) = \limsup_U h_U(S),$$

the limit being taken over the direct system of open neighborhoods of λ in $\sigma^\vee \setminus \{0\}$. For $\rho: \pi_1^p(R_\sigma) \rightarrow G$ a continuous representation to a discrete group, put $h_U(\rho) = h_U(S)$ and $h_\lambda(\rho) = h_\lambda(S)$, for $S \in \mathcal{C}_{R_\sigma}$ connected and chosen so that $\ker(\rho) = \pi_1^p(S)$.

We first work out how h_U works on linear cones. First, we bundle together some hypotheses.

Hypothesis 5.5 Let $\sigma \subseteq \mathbb{R}^n$ be a linear cone with $\mathbb{Z}^n \cap \sigma \neq \{0\}$, and put $\tau = -\sigma \cup \sigma$. Let \widehat{R}_σ be the completion of R_τ with respect to $v_{-\lambda}$, for some nonzero linear functional $\lambda: \mathbb{R}^n \rightarrow \mathbb{R}$ defined over \mathbb{Q} which is positive on $\sigma \setminus \{0\}$. Let $\rho: \pi_1^p(\text{Spec } R_\sigma) \rightarrow G$ be a continuous representation to a discrete group. Note that \widehat{R}_σ is a power series field over k , and that it depends only on σ , not on λ ; we may thus sensibly speak of the highest break $b(\rho)$.

Lemma 5.6 Under Hypothesis 5.5, let $\lambda: \mathbb{R}^n \rightarrow \mathbb{R}$ be a nonzero linear function defined over \mathbb{Q} , such that λ is positive on $\sigma \setminus \{0\}$. Put $d = [\mathbb{Z}^n : (\mathbb{Z}^n \cap H_\lambda) \times (\mathbb{Z}^n \cap \tau)]$, and let d' be the prime-to- p part of d . Then

$$c_\lambda(\rho) = \frac{d'}{m_\lambda} b(\rho).$$

Proof We first note that the desired equality holds when $d = 1$: it is the comparison between the highest break of a representation of π_1^p of a power series ring over a field, and the same representation after pulling back by extending the constant field.

We next note that if we repeat the construction of $c_\lambda(\rho)$ with \mathbb{Z}^n replaced by the larger lattice $(\mathbb{Z}^n \cap H_\lambda) \times \frac{1}{d}(\mathbb{Z}^n \cap \tau)$, then m_λ and $c_\lambda(\rho)$ remain unchanged. However, by Definition 3.3, $b(\rho)$ gets multiplied by d' . Now appealing to the $d = 1$ case yields the desired result. ■

Corollary 5.7 *With notation as in Definition 5.4 and Hypothesis 5.5, let \mathbf{v} be the smallest nonzero element of $\mathbb{Z}^n \cap \sigma$. Then*

$$h_U(\rho) = b(\rho) \sup_{\lambda \in U} \{\lambda(\mathbf{v})\}, \quad h_\lambda(\rho) = b(\rho)\lambda(\mathbf{v}).$$

In particular,

$$h_U(\rho) = \sup_{\lambda \in U} \{h_\lambda(\rho)\}.$$

Proof With notation as in Lemma 5.6, note that

$$d = [\lambda(\mathbb{Z}^n) : \lambda(\mathbb{Z}^n \cap \sigma)] = m_\lambda \lambda(\mathbf{v}).$$

By Lemma 5.6, we then have $c_\lambda(\rho) \leq b(\rho)\lambda(\mathbf{v})$, with equality for any λ for which d is not divisible by p . Such λ are dense in any U , so the desired results follow. ■

We now treat general cones.

Definition 5.8 For σ a very convex cone, τ a convex cone contained in σ , and $\rho: \pi_1^p(\text{Spec } R_\sigma) \rightarrow G$ a continuous representation to a discrete group, let ρ_τ be the pullback of ρ along the maps $\pi_1^p(\text{Spec } R_\sigma) \rightarrow \pi_1^p(\text{Spec } R_\tau) \rightarrow \pi_1^p(\text{Spec } R_\sigma)$, where the second map is as in Definition 4.9.

Lemma 5.9 *With notation as in Definition 5.4, let T be a linear cone contained in σ such that $d = [\mathbb{Z}^n : (\mathbb{Z}^n \cap H_\lambda) \times (\mathbb{Z}^n \cap (T \cup -T))]$ is coprime to p . Then $c_\lambda(\rho) \geq c_\lambda(\rho_T)$.*

Proof As in the proof of Lemma 5.6, we may reduce to the case $d = 1$; then Proposition 3.7 yields the claim. ■

Corollary 5.10 *With notation as in Lemma 5.9, we have*

$$h_U(\rho) \geq h_U(\rho_T), \quad h_\lambda(\rho) \geq h_\lambda(\rho_T).$$

Proof By Lemma 5.6, we may compute $h_U(\rho_T)$ by taking the supremum defining it only over λ as in Lemma 5.9 (i.e., the λ for which $d = d'$ in Lemma 5.6). Then Lemma 5.9 yields the first desired inequality; the second follows by taking limits. ■

We now have the following fairly explicit description of the functions h_U and h_λ , in terms of linear cones.

Theorem 5.11 *With notation as in Definition 5.4, we have*

$$(5.1) \quad h_U(\rho) = \sup_T \{h_U(\rho_T)\}, \quad h_\lambda(\rho) = \sup_T \{h_\lambda(\rho_T)\},$$

the suprema taken over all linear cones $T \subseteq \sigma$.

Proof In each case, the left side is greater than or equal to the right by Lemma 5.9. Conversely, by Theorem 4.17, we can present ρ inside the tensor product of the ρ_T over finitely many T , and so the right side is greater than or equal to the left. ■

5.2 More on the Explicit Height Functions

Theorem 5.11 makes it easy to verify many natural properties of the h_λ , including the fact that they are actually height functions. We present these as a series of corollaries.

Convention 5.12 Throughout this section, retain notation as in Definition 5.4.

Corollary 5.13 *We have $h_U(\rho) = \sup_{\lambda \in U} \{h_\lambda(\rho)\}$.*

Proof Applying Theorem 5.11 and Corollary 5.7, we have

$$h_U(\rho) = \sup_{T \subseteq \sigma} \{h_U(\rho_T)\} = \sup_{T \subseteq \sigma, \lambda \in U} \{h_\lambda(\rho_T)\} = \sup_{\lambda \in U} \{h_\lambda(\rho)\}. \quad \blacksquare$$

Corollary 5.14 *If $\lambda \in \sigma^\vee \setminus \{0\}$ is defined over \mathbb{Q} , then $h_\lambda(\rho) \in \mathbb{Q}$.*

Proof Apply Theorem 5.11 and note that only finitely many terms in the supremum in (5.1) are nonzero thanks to Theorem 4.17. Then apply Corollary 5.7 to deduce that each nonzero term in the supremum is rational. ■

Corollary 5.15 *Suppose that τ is a convex cone with $\tau \subseteq \sigma$. Then $h_\lambda(\rho) \geq h_\lambda(\rho_\tau)$.*

Proof Apply Theorem 5.11, and note that the supremum defining $h_\lambda(\rho_\tau)$ is simply the same supremum as in (5.1), but restricted to $T \subseteq \tau$. ■

Corollary 5.16 *Suppose that λ and κ both lie in the interior of $\sigma^\vee \setminus \{0\}$, and that $\lambda(\mathbf{v}) \geq \kappa(\mathbf{v})$ for all $\mathbf{v} \in \sigma$. Then $h_\lambda(\rho) \geq h_\kappa(\rho)$.*

Proof By Theorem 5.11, it suffices to check this for σ a linear cone, in which case it follows from Corollary 5.7. ■

Corollary 5.17 *Suppose that $S = R_\sigma[z]/(z^p - z - x)$, where the support V of x is contained in $\sigma \cap (\mathbb{Z}^n \setminus p\mathbb{Z}^n)$. Then $h_\lambda(S) = \sup_{\mathbf{w} \in V} \{v_\lambda(\mathbf{w})\}$.*

Proof Apply Theorem 5.11 to reduce to the case where σ is linear. Then apply Corollary 5.7 and Lemma 3.5. ■

Corollary 5.18 Each of the functions h_λ and h_U is a strong height function on \mathbb{C}_{R_r} over $R = k$.

Proof Conditions (i)–(iv) are straightforward, while condition (v) for $d = 1$ follows from Corollary 5.17; the claim then follows by Proposition 3.13. ■

Acknowledgments Thanks to Dan Abramovich for helping track down an error in a previous version of this paper, and to the referee for some helpful suggestions concerning the exposition.

References

- [1] A. Abbes and T. Saito, *Ramification of local fields with imperfect residue fields*. Amer. J. Math. **124**(2002), no. 5, 879–920.
- [2] J. M. Borger, *Conductors and the moduli of residual perfection*. Math. Ann. **329**(2004), no. 1, 1–30.
- [3] A. Grothendieck et al., *Revêtements étales et groupe fondamental (SGA 1)*. Lecture Notes in Mathematics 224, Springer-Verlag, Berlin, 1971.
- [4] A. Grothendieck et al., *Théorie des topos et cohomologie étale des schémas. I. II. III. (SGA 4)*. Lecture Notes in Mathematics 269, 270, 305, Springer-Verlag.
- [5] N. M. Katz, *Local-to-global extensions of representations of fundamental groups*. Ann. Inst. Fourier (Grenoble) **36**(1986), no. 4, 69–106.
- [6] G. Laumon, *Semi-continuité du conducteur de Swan (d’après P. Deligne)*. In: The Euler-Poincaré characteristic (French). Astérisque 83, Soc. Math. France, Paris, 1981, pp. 173–219.
- [7] J. S. Milne, *Étale cohomology*. Princeton Mathematical Series 33, Princeton University Press, Princeton, NJ, 1980.
- [8] J.-P. Serre, *Local Fields*. Graduate Texts in Mathematics 67, Springer-Verlag, New York, 1979.

Department of Mathematics
Massachusetts Institute of Technology
Cambridge, MA 02139
U.S.A.
e-mail: kedlaya@math.mit.edu