HIGHER-LEVEL CANONICAL SUBGROUPS IN ABELIAN VARIETIES

BRIAN CONRAD

1. Introduction

1.1. **Motivation.** Let E be an elliptic curve over a p-adic integer ring R, and assume E has supersingular reduction. Consider the 2-dimensional $\mathbf{F}_p$-vector space of characteristic-0 geometric p-torsion points in the associated 1-parameter formal group $\widehat{E}$ over R. It makes sense to ask if, in this vector space, there is a line whose points x are nearer to the origin than all other points (with nearness measured by $|X(x)|$ for a formal coordinate X of $\widehat{E}$ over R; the choice of X does not affect $|X(x)|$). Such a subgroup may or may not exist, and when it does exist it is unique and is called the *canonical subgroup*. This notion was studied by Lubin [Lu] in the more general context of 1-parameter commutative formal groups, and its scope was vastly extended by Katz [K] in the relative setting for elliptic curves over p-adic formal schemes and for analytified elliptic curves over certain modular curves. Katz' ideas grew into a powerful tool in the study of p-adic modular forms for $\mathrm{GL}_{2/\mathbf{Q}}$.

In [C4] we used rigid-analytic descent theory to give two generalizations of this p-torsion theory in the 1-dimensional case: higher p-power torsion-levels and arbitrary p-adic analytic families with no restriction on the base space or non-archimedean base field $k/\mathbf{Q}_p$. More specifically, under a necessary and sufficient fibral hypothesis, for any $n \geq 1$ we constructed a p^n-torsion relative canonical subgroup in generalized elliptic curves over rigid spaces over any analytic extension field $k/\mathbf{Q}_p$, and the method made no use of the fine integral structure of modular curves or their 1-dimensionality. However, the method did use the 1-dimensionality of the fibers in an essential way (as this severely restricts the possibilities for the semistable reduction type).

The study of p-adic modular forms for more general algebraic groups and number fields, going beyond the classical case of $\mathrm{GL}_{2/\mathbf{Q}}$, naturally leads to the desire to have a theory of canonical subgroups for rigid-analytic families of abelian varieties. (See [KL] for an application to Hilbert modular forms.) Ideally, one wants such a theory that is intrinsic to the rigid-analytic category and in particular avoids restrictions on the nature of formal (or algebraic) integral models for the family, but it should also be amenable to study using suitable formal models (when available). In this paper we develop such a theory, and our viewpoint and methods are rather different from those of other authors who have recently worked on the problem (such as [AM], [AG], [GK], and [KL]). The replacement for the 1-parameter formal groups in the case of elliptic curves is the formal group (or equivalently the identity component of the p-divisible group) of the semi-abelian formal model arising from the semistable reduction theorem for abelian varieties, proved in [BL1] without discreteness conditions on the absolute value.

Roughly speaking, if A is an abelian variety of dimension g over an analytic extension field $k/\mathbf{Q}_p$ (with the normalization condition $|p| = 1/p$) then a *level-n canonical subgroup* $G_n \subseteq A[p^n]$ is a k-subgroup with geometric fiber $(\mathbf{Z}/p^n\mathbf{Z})^g$ such that (for $\overline{k}^{\wedge}/k$ a completed algebraic closure) the points in $G_n(\overline{k}^{\wedge}) \subseteq A[p^n](\overline{k}^{\wedge})$ are "closer" to the identity than all other points in $A[p^n](\overline{k}^{\wedge})$, where closeness is defined in terms of absolute values of coordinates in the formal group of the unique formal semi-abelian model $\mathfrak{A}_{R'}$ for A over the valuation ring R' of a sufficiently large finite extension k'/k; this formal group is the same as the one

Date: July 22, 2005.
1991 *Mathematics Subject Classification.* Primary 14G22; Secondary 14H52.
This research was partially supported by grants from the NSF and the Alfred P. Sloan Foundation. I would like to thank F. Andreatta, M. Baker, S. Bosch, K. Buzzard, O. Gabber, C. Gasbarri, M. Hochster, A.J. de Jong, and A. Thuillier for helpful discussions.

arising from the identity component of the p-divisible group $\mathfrak{A}_{R'}[p^\infty]$ over the henselian local ring R'. (See Theorem 2.1.9 for the characterization of $\mathfrak{A}_{R'}$ in terms of the analytification $A^{\rm an}$, and see Definition 2.2.7 for the exact meaning of such closeness.) By [C4, Thm. 4.2.5], for $g = 1$ this notion of higher-level canonical subgroup is (non-tautologically) equivalent to the one defined in [Bu] and [G].

In contrast with the approach in [AM], our definition of canonical subgroups is not intrinsic to the torsion subgroups of A but rather uses the full structure of the formal group of a formal semi-abelian model. A level-n canonical subgroup is obviously unique if it exists, and it is an elementary consequence of the definitions (see Remark 2.2.10) that A admits a level-n canonical subgroup for all $n \geq 1$ if and only if A is *ordinary* in the sense that the abelian part of the semi-abelian reduction $\mathfrak{A}_{R'} \bmod \mathfrak{m}_{R'}$ is an ordinary abelian variety over $R'/\mathfrak{m}_{R'}$ (or equivalently, if and only if the p-divisible group $\mathfrak{A}_{R'}[p^\infty]$ has ordinary reduction over the residue field of R'). This is modest initial evidence to support the appropriateness of our definition.

In the 1-dimensional case there is the explicit result [C4, Thm. 4.2.5] that a level-n canonical subgroup exists if and only if the "Hasse invariant" exceeds $p^{-p/p^{n-1}(p+1)}$. A basic theme in this paper is to prove properties of canonical subgroups (such as existence and duality results) subject to universal bounds on fibral Hasse invariants, so let us now recall how the Hasse invariant is defined in the p-adic analytic setting. Using notation A, k'/k, and $\mathfrak{A}_{R'}$ as above, let $\overline{\mathfrak{A}}_{R'} = \mathfrak{A}_{R'} \bmod pR'$ so the relative Verschiebung morphism $V : \overline{\mathfrak{A}}_{R'}^{(p)} \to \overline{\mathfrak{A}}_{R'}$ induces a Lie algebra map ${\rm Lie}(V) : {\rm Lie}(\overline{\mathfrak{A}}_{R'}^{(p)}) \to {\rm Lie}(\overline{\mathfrak{A}}_{R'})$ between finite free R'/pR'-modules of the same rank. The linear map ${\rm Lie}(V)$ has a determinant in R'/pR' that is well-defined up to unit multiple (and is taken to be a unit if $A = 0$). The *Hasse invariant* $h(A) \in [1/p, 1] \cap \sqrt{|k^\times|}$ is the maximum of $1/p$ and the absolute value of a lift into R' for this determinant in R'/pR'. Since $h(A) = 1$ if and only if A is ordinary (with $h(A) = 1$ when $A = 0$), the number $h(A)$ is a measure of the failure of the abelian part of $\mathfrak{A}_{R'} \bmod \mathfrak{m}_{R'}$ to be ordinary. Work of Mazur–Messing ensures the identity $h(A^\vee) = h(A)$, with $A^\vee$ denoting the dual abelian variety. (See Theorem 2.3.4.)

A very natural question in the g-dimensional case for any $g \geq 1$ is this: is there a number $h(p, g, n) < 1$ so that for any g-dimensional abelian variety A over any analytic extension field $k/\mathbf{Q}_p$, if $h(A) > h(p, g, n)$ then A admits a level-n canonical subgroup G_n? In effect, we are asking for an existence criterion that has nothing to do with any particular modular family in which A may have been presented to us. The best choice for $h(p, 1, n)$ is $p^{-p/p^{n-1}(p+1)}$, but for $g > 1$ it seems unreasonable to expect there to be a strict lower bound $h(p, g, n)$ that is sufficient for existence of a level-n canonical subgroup *and* is also necessary for existence. Thus, we cannot expect there to be a "preferred" value for $h(p, g, n)$ when $g > 1$.

We will prove the existence of such an $h(p, g, n)$, but then more questions arise. For example, since $h(A^\vee) = h(A)$, can $h(p, g, n)$ be chosen so that if $h(A) > h(p, g, n)$ then the level-n canonical subgroup of $A^\vee$ is the orthogonal complement of the one for A under the Weil-pairing on p^n-torsion? Also, what can be said about the reduction of such a G_n into $\mathfrak{A}_{R'}[p^n]^0 \bmod pR'$? Finally, how does the level-n canonical subgroup relativize in rigid-analytic families of abelian varieties $A \to S$ (over rigid spaces S over $k/\mathbf{Q}_p$)? The method of proof of existence of $h(p, g, n)$ allows us to give satisfactory answers to these auxiliary questions, and when a formal semi-abelian model is given for a family $A_{/S}$ then we describe relative canonical subgroups locally on the base in terms of fibral formal coordinates on the formal model.

The rigid-analytic families $A \to S$ of most immediate interest are those that are *algebraic* in the sense that $A_{/S}$ is a pullback of the analytification of an abelian scheme over a locally finite type k-scheme. However, this class of families is too restrictive as a foundation for the general theory. For example, when using canonical subgroups to study p-adic modular forms one has to consider passage to the quotient by a relative canonical subgroup and so there arises the natural question of whether such a quotient admits a relative canonical subgroup at a particular level (*cf.* [Bu], [Kas]). In practice, if $A \to S$ is the analytification of an abelian scheme then its relative canonical subgroups (when they exist) do not generally arise from analytification within the same abelian scheme, and so passage to the quotient by such a subgroup is a non-algebraic operation. It is therefore prudent to enlarge the class of families being considered so that it is stable under passage to the quotient by *any* rigid-analytic finite flat subgroup. We work with the following larger class that meets this requirement (and is local on the base): those $A_{/S}$ for which there exists an admissible covering $\{S_i\}$ of S and finite surjections $S'_i \to S_i$ such that $A_{/S'_i}$ is algebraic. We summarize this condition by saying that $A_{/S}$ *becomes algebraic after local finite surjective base change on* S. (See Example 2.1.8 for the stability of this class under quotients by finite flat subgroups.)

Modular varieties for g-dimensional abelian varieties with $g > 1$ generally admit normal compactifications such that the boundary has codimension ≥ 2, so in a geometric theory of p-adic modular forms beyond the classical case there is less of a need to use families with degenerate characteristic-0 fibers as in the case $g = 1$. Thus, for $g > 1$ it is not unduly restrictive to work with families of g-dimensional abelian varieties without degeneration in characteristic 0. (Of course, we *do* allow the abelian-variety fibers in the family to have non-trivial toric part in their potential semi-abelian reduction type over the residue field of the valuation ring.) The theory for $g > 1$ is more difficult than in the case $g = 1$ because we do not fix discrete parameters (such as a PEL-type) and we allow any potentially semistable fibral reduction type. Berkovich spaces play a vital role in some of our proofs (such as for Theorem 1.2.1 below), so we must allow arbitrary $k/\mathbf{Q}_p$ even if our ultimate interest is in the case of discretely-valued extensions of $\mathbf{Q}_p$.

1.2. **Overview of results.** An *abeloid space* over a rigid space S over a non-archimedean field k is a proper smooth S-group $A \to S$ with connected fibers. Relative ampleness (in the sense of [C3]) gives a good notion of polarized abeloid space over a rigid space, and analytifications of universal objects for certain moduli functors of polarized abelian schemes satisfy an analogous universal property in the rigid-analytic category.

Let S be a rigid space over any non-archimedean extension $k/\mathbf{Q}_p$, with the normalization $|p| = 1/p$, and consider abeloid spaces $A \to S$ of relative dimension $g \geq 1$ such that either:

(i) $A_{/S}$ admits a polarization *fpqc*-locally on S, or

(ii) $A_{/S}$ becomes algebraic after local finite surjective base change on S (in the sense defined in §1.1).

(In either case, a simple descent argument ensures that the fibers A_s admit ample line bundles and so are abelian varieties.) We do not restrict ourselves to universal families over specific modular varieties over $\mathbf{Q}_p$ and we allow p and $k/\mathbf{Q}_p$ to be arbitrary (*e.g.*, k need not be discretely-valued). We prove that for any abeloid space $A \to S$ as in cases (i) or (ii) and any $h \in (p^{-1/8}, 1) \cap \sqrt{|k^\times|}$, the locus $S^{>h}$ (resp. $S^{\geq h}$) of $s \in S$ such that the fiber A_s has Hasse invariant $> h$ (resp. $\geq h$) is an admissible open, and that for quasi-separated S its formation is compatible with arbitrary extension on k. (The intervention of $p^{-1/8}$ has no significance; it is an artifact of the use of Zahrin's trick in the proof, and we shall only care about h universally near 1 anyway.) We also prove that the open immersions $S^{\geq h} \to S$ are quasi-compact with $\{S^{\geq h'}\}_{h<h'\leq 1}$ an admissible cover of $S^{>h}$, and that for "reasonable" S the formation of $S^{\geq h}$ and $S^{>h}$ is compatible with passage to Berkovich spaces. The overconvergent nature of canonical subgroups is due to the fact that $S^{>h} \to S$ induces an open immersion on the associated Berkovich spaces.

These properties of $S^{>h}$ and $S^{\geq h}$ are not obvious because in general (even locally on S) there does not seem to exists a rigid-analytic function H for which $s \mapsto \max(|H(s)|, 1/p)$ equals the fibral Hasse invariant $h(A_s)$. Such an H locally exists in many polarized cases, but probably not in general. To avoid polarization restrictions for analytified algebraic families we use a result of Gabber (Theorem A.2.1) that is of independent interest.

The main result in this paper is:

Theorem 1.2.1. *There exists a positive number $h(p, g, n) < 1$ depending only on p, g, and n (and not on the analytic base field $k/\mathbf{Q}_p$, a PEL-type, or special properties of a base space) such that if $A \to S$ is an abeloid space of relative dimension g that satisfies either of the hypotheses* (i) *or* (ii) *above and $h(A_s) > h(p, g, n)$ for all $s \in S$ then there exists a finite étale S-subgroup of $A[p^n]$ that induces a level-n canonical subgroup on the fibers. Such an S-subgroup is unique, and the formation of this subgroup respects arbitrary extension of the analytic base field.*

We also show that by taking $h(p, g, n)$ to be sufficiently near 1, the level-n canonical subgroup of any g-dimensional abelian variety $A_{/k}$ satisfying $h(A) > h(p, g, n)$ is well-behaved with respect to duality and products (in dimensions adding up to g), and that it reduces to the kernel of the n-fold relative Frobenius morphism on the semi-abelian reduction modulo $p^\lambda R'$ for any *fixed* $\lambda \in (0, 1) \cap \mathbf{Q}$ and any sufficiently large finite extension k'/k (with valuation ring R'). This compatibility with Frobenius kernels allows us to study the behavior of the Hasse invariant and level-n canonical subgroup under passage to the quotient by the level-m canonical subgroup for $1 \leq m < n$.

The idea for the construction of a level-n canonical subgroup in a g-dimensional abelian variety A with Hasse invariant sufficiently near 1 (where "near" depends only on p, g, and n) is to proceed in three steps: (I) the principally polarized case in any dimension (using that the moduli scheme $\mathscr{A}_{g,1,N/\mathbf{Z}_p}$ over $\mathbf{Z}_p$ with $p \nmid N$ and a fixed $N \geq 3$ admits a compactification equipped with a semi-abelian scheme extending the universal abelian scheme), (II) the potentially good reduction case, which we study via the principally polarized case (Zahrin's trick) and a theorem of Norman and Oort concerning the geometry of $\mathscr{A}_{g,d,N/\mathbf{Z}_p}$ for *all* $d \geq 1$, and (III) the general case, which we study by applying the potentially good reduction case to the algebraization of the formal abelian part in the semistable reduction theorem for A.

Let us be a bit more precise concerning these three steps. In the principally polarized case we use Berkovich's étale cohomology theory for torsion sheaves arising from the analytified semi-abelian scheme over a compactification of $\mathscr{A}_{g,1,N/\mathbf{Z}_p}$ to solve a more general construction problem for all g by "smearing out" from the ordinary locus. This smearing out process gives rise to difficult connectivity problems that we do not know how to solve, and such problems are circumvented by using Berkovich's description of étale cohomology along "germs". The $\mathbf{Z}_p$-properness of these compactifications is essential for the success of this step, and it is the reason we can find a *universal* sufficient strict lower bound $h_{\rm pp}(p,g,n) < 1$ in the principally polarized case. Since the construction of $h_{\rm pp}(p,g,n)$ rests on compactness arguments (on Berkovich spaces), it is not explicit.

To settle the case of potentially good reduction in any dimension g with the sufficient strict lower bound $h_{\rm good}(p,g,n) = h_{\rm pp}(p,8g,n)^{1/8}$ on Hasse invariants, we use Zahrin's trick to construct a level-n canonical subgroup on the principally polarized $8g$-dimensional $(A \times A^\vee)^4$. In the principally polarized case our construction provides universal control on "how far" the canonical subgroup is from the origin of the formal group of a unique formal semi-abelian model, and this enables us to infer that the level-n canonical subgroup in $(A \times A^\vee)^4$ must have the form $(G_n \times G'_n)^4$ for subgroups $G_n \subseteq A[p^n]$ and $G'_n \subseteq A^\vee[p^n]$, so the fibers of G_n and G'_n are finite free $\mathbf{Z}/p^n\mathbf{Z}$-modules with ranks adding up to $2g$. But why does taking $h(A)$ near enough to 1 in a universal manner suffice to force the $\mathbf{Z}/p^n\mathbf{Z}$-ranks of G_n and G'_n to equal g and force G_n and G'_n to annihilate each other under the Weil pairing on p^n-torsion? Since $A_{/k}$ has a polarization of some (unknown) degree $d^2 \geq 1$ (that may well be divisible by p), the potentially good reduction hypothesis enables us to exploit the geometry of $\mathscr{A}_{g,d,N/\mathbf{Z}_p}$ as follows. By a theorem of Norman and Oort, the ordinary locus in $\mathscr{A}_{g,d,N/\mathbf{F}_p}$ is a Zariski-dense open and $\mathscr{A}_{g,d,N/\mathbf{Z}_p}$ is a relative local complete intersection over $\mathbf{Z}_p$. Thus, for any closed point $\overline{x} \in \mathscr{A}_{g,d,N/\mathbf{F}_p}$ (such as arises from the reduction of our chosen polarized abelian variety equipped with an N-torsion basis, after a preliminary argument to reduce to the case $[k : \mathbf{Q}_p] < \infty$) we may use slicing to find a $\mathbf{Z}_p$-flat curve Z in $\mathscr{A}_{g,d,N/\mathbf{Z}_p}$ such that the closed fiber of Z over $\mathrm{Spec}(\mathbf{Z}_p)$ passes through $\overline{x}$ and has all generic points in the ordinary locus. In conjunction with a connectivity result for affinoid curves (applied to the generic fiber $\mathfrak{Z}^{\rm rig}$ of the p-adic completion $\mathfrak{Z}$ of Z), this allows us to solve our problems in the potentially good reduction case by analytic continuation from the ordinary case.

Finally, in the general case the semistable reduction theorem of Bosch and Lütkebohmert provides a unique formal semi-abelian "model" $\mathfrak{A}$ for A (after a finite extension on k) and the formal abelian part $\mathfrak{B}$ of $\mathfrak{A}$ is uniquely algebraizable to an abelian scheme over the valuation ring. This unique algebraization has generic fiber B that is an abelian variety satisfying $h(B) = h(A)$ and $\dim B \leq \dim A$; perhaps $B = 0$, but then we are in the purely toric (and hence ordinary) case that is trivial. We apply the settled case of potentially good reduction to B in order to solve the existence problem in the general case with $h(p,g,n) = \max_{1 \leq g' \leq g} h_{\rm good}(p,g',n)$.

1.3. **Further remarks.** The reader may be wondering: since the definitions of level-n canonical subgroup and Hasse invariant make sense for any p-divisible group Γ over the henselian local valuation ring (the identity component Γ^0 provides both a Lie algebra and a formal group), why isn't this entire theory carried out in the generality of suitable families of Barsotti–Tate (BT) groups? There are many reasons why this is not done. First of all, whereas an abelian variety determines a unique (and functorial) semi-abelian formal model even when the base field is algebraically closed, this is not the case for BT-groups. Hence, if the theory is to work over a non-archimedean algebraically closed field $k/\mathbf{Q}_p$ then it seems necessary to specify a relative formal model as part of the input data. However, one cannot expect to find relative canonical subgroups that arise from the same formal data (especially if we later try to shrink the rigid-analytic base space) and so one would constantly be forced to change the formal model in an inconvenient manner. Even if we restrict to the case of a discretely-valued base field, there arises the more fundamental problem that in a family of abelian varieties whose semistable reduction types are varying there is often no obvious way to pick out a global formal BT-group with which to work. We want the relative theory for abelian varieties to be applicable without restrictions on the fibral reduction type or on the existence of preferred integral models for the rigid base space. (The specification of formal models is a source of technical complications in the theory developed in [K], and we want a theory that is intrinsic to the rigid-analytic category.)

It is also worth noting that for BT-groups one has no analogue of the $\mathbf{Z}_p$-compactifications of the $\mathscr{A}_{g,1,N/\mathbf{Q}_p}$'s that classify "all" abelian varieties (up to suitable applications of Zahrin's trick). Having a $\mathbf{Z}_p$-proper base space is the key reason that we are able to get a universal number $h(p,g,n) < 1$ in Theorem 1.2.1, and so for BT-groups there seems to be no *a priori* reason to expect the existence of such a number. Our method ultimately rests on analytic continuation from the ordinary case, and for BT-groups the nearest analogue seems to be the generic ordinarity of the universal equicharacteristic deformation of a BT-group Γ_0 over a field of characteristic p. Unfortunately, the Berthelot generic fiber of the universal formal deformation ring of such a Γ_0 is an open unit polydisc and it lacks a good compactification. One might get lucky via explicit calculations with multivariable formal group laws, and such calculations may clarify the situation for abelian varieties. Aside from the possibility of a miracle in such calculations, we are not aware of any reason to justify the expectation that there exists a universal constant like $h(p,g,n)$ in the case of a reasonable class of BT-groups that do not arise from abelian varieties.

Let us now briefly summarize the contents of this paper. In §2.1 we recall some results of Bosch and Lütkebohmert on semistable reduction over non-archimedean fields, and in §2.2 we use these results to define canonical subgroups. In §2.3 we define the Hasse invariant and use the work of Mazur and Messing on relative Dieudonné theory to show that the Hasse invariant is unaffected by passage to the dual abelian variety. This is crucial, due to the role of Zahrin's trick in subsequent arguments. The variation of the Hasse invariant in families is studied in the polarized case in §3.1, and in §3.2 we use a theorem of Gabber to get results in the analytified "algebraic" setting without polarization hypotheses. The technical heart of the paper is §4.1–4.2. In §4.1 we construct level-n canonical subgroups in g-dimensional abelian varieties whose fibral Hasse invariants exceed a suitable $h(p,g,n) < 1$ and we show that such canonical subgroups are well-behaved with respect to duality of abelian varieties. The key geometric input into the argument is a result concerning the existence of ordinary points on connected components of certain rigid-analytic domains in $\mathscr{A}^{\mathrm{an}}_{g,d,N/\mathbf{Q}_p}$, and the proof of this result occupies §4.2. The relativization of the theory and the relationship between level-n canonical subgroups and the kernel of the n-fold relative Frobenius map modulo $p^{1-\varepsilon}$ for any fixed $\varepsilon \in (0,1)$ (when the Hasse invariant exceeds a suitable $h(p,g,n,\varepsilon) \in (1/p,1)$) are worked out in §4.3, where we also give a partial answer to the question of how the level-n canonical subgroup and Hasse invariant behave under passage to the quotient by the level-m canonical subgroup for $1 \le m < n$.

As this work was being completed we became aware of recent results of others on the theme of canonical subgroups for abelian varieties. Abbes–Mokrane [AM] (for $p \ge 3$), Goren–Kassaei [GK], and Kisin–Lai [KL] provide overconvergent canonical subgroups for universal families of abelian varieties over some modular varieties over discretely-valued extensions of $\mathbf{Q}_p$, and Andreatta–Gasbarri [AG] construct p-torsion canonical subgroups for families of polarized abelian varieties with good reduction. In §4.4 we compare our work with these other papers, including consistency between all of these points of view (at least near the ordinary locus on the base).

There is a uniqueness result for the p-torsion good-reduction theory developed in [AG], where one also finds explicit bounds (in contrast with our non-explicit bounds that arise from compactness arguments). Using notation as in §1.2, the problem of making $h(p,g,n)$ explicit can be reduced to the problem of making $h_{\mathrm{pp}}(p,g',n)$ explicit in the case of good reduction for abelian varieties of dimension $g' \le 8g$ over finite extensions of $\mathbf{Q}_p$. Our fibral definitions are well-suited to the (non-quasi-compact) Berthelot rigid-analytification of the universal formal deformation of any principally polarized abelian variety over a perfect field of characteristic p (or any field of characteristic p equipped with an associated Cohen ring). Hence, a plausible method to determine an explicit $h(p, g, n)$ is to study *all* such local families (even with just finite residue field) because the results of this paper ensure *a priori* that a sufficient lower bound on Hasse invariants across all such local families is also sufficient for global rigid-analytic families of abelian varieties over any non-archimedean base field $k/\mathbf{Q}_p$. This viewpoint may also be useful for a deeper study of the behavior of canonical subgroups with respect to isogenies.

1.4. **Notation and Terminology.** Our notation and terminology conventions are the same as in the previous paper [C4] that treats the 1-dimensional case. In particular, we refer to [C4, §1.3] for a discussion of the notion of pseudo-separatedness. (A rigid-analytic map $f : X \to Y$ is *pseudo-separated* if its relative diagonal factors as a Zariski-open immersion followed by a closed immersion; analytifications of algebraic morphisms are pseudo-separated. This notion is introduced solely to avoid unnecessary separatedness restrictions on locally finite type k-schemes when we wish to consider how their analytifications interact with change of the base field.)

Contents

2. Abelian varieties over non-archimedean fields

Our first aim is to define a Hasse invariant $h(A) \in [1/p, 1] \cap \sqrt{|k^\times|}$ for any abelian variety A over an analytic extension $k/\mathbf{Q}_p$ with $|p| = 1/p$. The definition rests on the semi-stable reduction theorem that was proved in a suitable form by Bosch and Lütkebohmert over an arbitrary non-archimedean field k (*e.g.*, $|k^\times|$ may be non-discrete in $(0, \infty)$). We first review some generalities for abelian varieties over non-archimedean fields in §2.1, and then we shall specialize to $k/\mathbf{Q}_p$ in §2.2 and §2.3 where we define and study canonical subgroups and Hasse invariants.

2.1. **Polarization and semi-stable reduction.** Let us begin by recalling some standard terminology in the context of relative polarizations over an arbitrary non-archimedean field k.

Definition 2.1.1. Let S be a rigid-analytic space over k. An *abeloid space* over S is a proper smooth S-group $f : A \to S$ whose fibers are (geometrically) connected.

As with any smooth map having geometrically connected and non-empty fibers, the fiber-dimension of an abeloid space is locally constant on the base. Thus, we will usually restrict our attention to abeloid spaces with a fixed relative dimension $g \geq 1$. For quasi-separated or pseudo-separated S, any change of the base field carries abeloid spaces to abeloid spaces and preserves the Zariski-open loci over which the fibers have a fixed dimension.

The standard infinitesimal-fiber and cohomological arguments for abelian schemes [GIT, §6.1] carry over *verbatim* to show that the group law on an abeloid space is uniquely determined by its identity section and that any S-map between abeloid spaces must respect the group laws if it respects the identity sections. In particular, the group law is commutative. By [L3, Thm. II] there is a uniformization theorem (requiring a finite extension of the base field) for abeloid groups over a discretely-valued non-archimedean field k, from which it follows that the n-torsion subgroups for nonzero integers n are finite with the same order as in the case of complex tori and abelian varieties. Hence, if the base field is discretely-valued and $A \to S$ is an abeloid S-group of relative dimension g then for any positive integer n the same arguments as in the algebraic case show that the map $[n]$ is a finite flat surjection of degree n^{2g} and so (via [C3, §4.2]) the map $[n]$ exhibits A as a quotient of A modulo the finite flat closed subgroup $A[n]$. According to [L3, Rem. 6.7] the uniformization theorem (and thus the preceding consequences) is almost certainly true without discreteness restrictions on the absolute value, but there are a few technical aspects of the proofs that have to be re-examined in such generality. (The existence of rigid-analytic uniformizations for all abelian varieties, without restriction on the non-archimedean base field, is due to Bosch and Lütkebohmert and is recalled as part of Theorem 2.1.9 below.)

We shall only ever work with the multiplication maps $[n]$ in cases when the fibers A_s are known to be abelian varieties, and so we do not require the general rigid-analytic uniformization theorem for abeloid k-groups over a non-archimedean field k. However, to avoid presumably artificial restrictions in examples and to keep the exposition clean we shall assume in all examples concerning torsion subgroups of abeloids that the uniformization theorem is valid for any abeloid k-group over any non-archimedean base field k. This presents no logical gaps for our intended applications of such examples in the case of abeloids whose fibers are known to be abelian varieties (such as in all of our theorems that involve torsion subgroups).

Example 2.1.2. Let $A \to S$ be an abeloid space and let $G \subseteq A$ be a finite flat closed S-subgroup. The action by G on A over S defines a finite flat equivalence relation on A over S, and we claim that the quotient A/G exists as an S-abeloid space. That is, we want to construct a finite flat surjective map of abeloids $A \to A'$ with kernel G (by [C3, §4.2] this A' serves as a quotient and as such has all of the usual properties one would desire with respect to maps and base change). A case of particular interest is when k is an analytic extension field of $\mathbf{Q}_p$ and $A \to S$ is a pullback of the analytification of an abelian scheme $\mathscr{A}_{/\mathscr{S}}$ over a locally finite type k-scheme, for then (after shrinking S appropriately) there is a relative level-n canonical subgroup $G_n \subseteq A$ that is a finite étale closed subgroup. Such a G_n will be constructed in §4.3, and in general G_n does not arise from a subgroup scheme of the given algebraic model $\mathscr{A}_{/\mathscr{S}}$ but for applications with modular forms it is useful to form A/G_n.

To construct A/G over S, we may work locally on S and so we can assume that G has constant order d. The map $[d] : A \to A$ is a finite flat covering that exhibits the source as a torsor over the target for the action of the finite flat group $A[d]$ (using the *fpqc* topology). In general, if $X' \to X$ is a finite flat map of rigid spaces that is an *fpqc* torsor for the action by a finite flat X-group H and if $H_0 \subseteq H$ is a finite flat closed subgroup (such as $X' = A$, $X = A$, $H = A[d] \times_S X$, $H_0 = G \times_S X$) then the existence of the flat quotient X'/H_0 follows by working over an admissible affinoid cover of X and using Grothendieck's existence results on quotients by free actions of finite locally free group schemes in the affine case [SGA3, V, §4]. This procedure is compatible with products over S in the spaces and groups if everything is given in the category of rigid spaces over a rigid space S. In this way we can construct A/G as a rigid space that is a finite flat cover intermediate to $[d] : A \to A$, so it is S-proper because it is finite over the target A and it is S-smooth with geometrically connected fibers because it has a finite flat cover by the source A. Since the natural S-map $(A \times A)/(G \times G) \to (A/G) \times (A/G)$ is an isomorphism, we get the desired S-group structure on A/G with respect to which the finite flat covering $A \to A/G$ over S is a homomorphism with kernel G.

Definition 2.1.3. A *correspondence* between two abeloid spaces $A, A' \rightrightarrows S$ is a line bundle $\mathscr{L}$ on $A \times A'$ equipped with trivializations $i : (e \times 1)^*\mathscr{L} \simeq \mathscr{O}_{A'}$ and $i' : (1 \times e')^*\mathscr{L} \simeq \mathscr{O}_A$ such that $e'^*(i) = e^*(i)$ as isomorphisms $(e \times e')^*\mathscr{L} \simeq \mathscr{O}_S$.

Any other choice for the pair (i, i') on the same $\mathscr{L}$ is related by the action of $\mathbf{G}_m(S)$ under which $c \in \mathbf{G}_m(S)$ carries (i, i') to $(c \cdot i, c \cdot i')$, so in particular each of i or i' determines the other. It is clear that $\mathscr{L}$ has no non-trivial automorphism that is compatible with either i or i'. In practice, we shall refer to $\mathscr{L}$ as a correspondence without explicitly mentioning i and i' (assuming such an (i, i') exists and has been chosen). If $A' = A$ then we call $\mathscr{L}$ a correspondence *on* A. A correspondence $\mathscr{L}$ on A is *symmetric* if there is an isomorphism $\mathscr{L} \simeq \sigma^*\mathscr{L}$ respecting trivializations along the identity sections, where σ is the automorphism of $A \times A$ that switches the factors; such an isomorphism of line bundles is unique if it exists, and the symmetry condition is independent of the choice of pair (i, i'). The rigid-analytic theory of relative ampleness [C3] allows us to make the following definition:

Definition 2.1.4. A *polarization* of an abeloid space $A \to S$ is a symmetric correspondence $\mathscr{L}$ on A such that the pullback $\Delta^*\mathscr{L}$ along the diagonal is S-ample on A.

Ampleness in the rigid-analytic category is characterized by the cohomological criterion [C3, Thm. 3.1.5], so by GAGA an abeloid space over $\mathrm{Sp}(k)$ admits a polarization if and only if it is an abelian variety. Moreover, since relative ampleness is compatible with change in the base field [C3, Cor. 3.2.8], if S is quasi-separated or pseudo-separated then polarizations are taken to polarizations under change in the base field.

Theorem 2.1.5. *Let $A \to S$ be an abeloid space with identity e, and assume that A admits a relatively ample line bundle locally on S. The functor $T \rightsquigarrow \mathrm{Pic}_e(A_T)$ classifying line bundles trivialized along e is represented by a separated S-group $\mathrm{Pic}_{A/S}$. This S-group contains a unique Zariski-open and Zariski-closed S-subgroup $A^\vee$ that is the identity component of $\mathrm{Pic}_{A/S}$ on fibers over S. The S-group $A^\vee$ is abeloid and admits a relatively ample line bundle locally on S, and the canonical map $i_A : A \to A^{\vee\vee}$ is an isomorphism with $i_A^\vee$ inverse to $i_{A^\vee}$.*

The formation of $\mathrm{Pic}_{A/S}$ and $A^\vee$ commute with change of the base field when S is quasi-separated or pseudo-separated, and each is compatible with analytification from the case of abelian schemes that are projective locally on the base.

Proof. The functor $\mathrm{Pic}_{A/S}$ is a sheaf on any rigid space over S, so we may work locally on S. Hence, we can assume that A admits a closed immersion into $\mathbf{P}^N_S$. It is a consequence of the compatible algebraic and rigid-analytic theories of Hilbert and Hom functors that a finite diagram among rigid-analytic spaces projective and flat over a common rigid space can be realized as a pullback of the analytification of an analogous finite diagram of locally finite type k-schemes. (See [C3, Cor. 4.1.6] for a precise statement.) Thus, there exists a locally finite type k-scheme $\mathscr{S}$ and an abelian scheme $\mathscr{A} \to \mathscr{S}$ equipped with an embedding into $\mathbf{P}^N_{\mathscr{S}}$ such that its analytification pulls back to $A \to S$ along some map $S \to \mathscr{S}^{\mathrm{an}}$. By the rigid-analytic theory of the Picard functor [C3, Thm. 4.3.3] we thereby get the existence of $\mathrm{Pic}_{A/S}$ and $A^\vee$ compatibly with analytification, and the algebraic theory for abelian schemes provides the rest. ■

If A is an abeloid space that admits a polarization locally on S (so all fibers are abelian varieties), then Theorem 2.1.5 provides an abeloid dual $A^\vee$ that admits a polarization locally on S and is rigid-analytically functorial in A. A polarization on such an A corresponds to a symmetric morphism of abeloid spaces $\phi : A \to A^\vee$ such that the line bundle $(1, \phi)^*(\mathscr{P})$ on A is S-ample, where $\mathscr{P}$ is the Poincaré bundle on $A \times A^\vee$. By the algebraic theory on fibers it follows that ϕ is finite and flat with square degree d^2. This degree (a locally-constant function on S) is the *degree* of the polarization. When $d = 1$ we say ϕ is a *principal polarization*.

Corollary 2.1.6. *Let $f : A \to S$ be an abeloid space equipped with a degree-d^2 polarization ϕ. Locally on the base, $(A_{/S}, \phi)$ is a pullback of the analytification of an abelian scheme equipped with a degree-d^2 polarization.*

Proof. The polarization is encoded as a symmetric finite flat morphism $\phi : A \to A^\vee$ with degree d^2 such that $\mathscr{L} = (1, \phi)^*\mathscr{P}$ is S-ample, so an application of the rigid-analytic theory of Hom functors [C3, Cor. 4.1.5] and the "local algebraicity" as in the proof of Theorem 2.1.5 gives the result. (The only reason we have to work locally on S is to trivialize the vector bundle $f_*(\mathscr{L}^{\otimes 3})$.) ∎

Example 2.1.7. Let $N \geq 3$ be a positive integer not divisible by char(k), and let $\mathscr{A}_{g,d,N/k}$ be the quasi-projective k-scheme that classifies abelian schemes of relative dimension g equipped with a polarization of degree d^2 and a basis of the N-torsion. The separated rigid space $\mathscr{A}^{\rm an}_{g,d,N/k}$, equipped with the analytification of the universal structure, represents the analogous functor in the rigid-analytic category over k; this follows by rigidity of the functor and the method used to prove Theorem 2.1.5. Note in particular that the formation of this universal structure is compatible with change in the analytic base field.

Example 2.1.8. Suppose that $A \to S$ is an abeloid space admitting a polarization locally over S and that $G \subseteq A$ is a finite flat S-subgroup. All fibers are abelian varieties, and by Example 2.1.2 we get an abeloid quotient $A' = A/G$ equipped with a finite flat surjection $h : A \to A'$. The norm operation from line bundles on A to line bundles on A' preserves relative ampleness (by GAGA and [EGA, II, 6.6.1] on fibers over S), so A' necessarily admits a polarization locally over S.

A weaker hypothesis to impose on an abeloid $A \to S$ is that it becomes algebraic after local finite surjective base change on S (in the sense defined in §1.1). By Corollary 2.1.6 and (the self-contained and well-known) Lemma 3.2.1 below, it is equivalent to assume that the abeloid $A \to S$ acquires a polarization after local finite surjective base change on S. In this case, if $G \subseteq A$ is a finite flat S-subgroup then the hypotheses force the fibers A_s to be abelian varieties and the S-abeloid quotient A/G acquires a polarization after local finite surjective base change on S. It follows from the proof of Corollary 2.1.6 that A/G becomes algebraic after local finite surjective base change on S because if $T' \to T$ is a finite surjection between rigid spaces and $\mathscr{F}'$ is a vector bundle on T' then $\mathscr{F}'$ is trivialized over the pullback of an admissible open covering of T; *cf* [EGA, II,6.1.12].

The definition of canonical subgroups in abelian varieties over analytic extensions of $\mathbf{Q}_p$ requires formal semi-abelian models as in the following non-archimedean semi-stable reduction theorem that avoids discreteness restrictions on the absolute value.

Theorem 2.1.9 (Bosch–Lütkebohmert)**.** *Let A be an abelian variety over k. For any sufficiently large finite separable extension k'/k (with valuation ring R') there exists a quasi-compact admissible open k'-subgroup $U \subseteq A^{\rm an}_{k'}$ and an isomorphism of rigid-analytic k'-groups $\iota : U \simeq \mathfrak{A}^{\rm rig}_{R'}$ where $\mathfrak{A}_{R'}$ is a topologically finitely presented and formally smooth formal* Spf(R')*-group that admits a (necessarily unique) extension structure*

$$1 \to \mathfrak{T} \to \mathfrak{A}_{R'} \to \mathfrak{B} \to 1 \tag{2.1.1}$$

as topologically finitely presented and flat commutative formal Spf(R')*-groups, with $\mathfrak{T}$ a formal torus and $\mathfrak{B}$ a formal abelian scheme over* Spf(R')*.*

The quasi-compact open subgroup U and the formal Spf(R')*-group $\mathfrak{A}_{R'}$ (equipped with the isomorphism ι) are unique up to unique isomorphism and are uniquely functorial in $A_{k'}$, as are $\mathfrak{T}$ and $\mathfrak{B}$. There exists a unique abelian scheme $B_{R'}$ over* Spec(R') *(with generic fiber denoted B over k') whose formal completion along an ideal of definition of R' is isomorphic to $\mathfrak{B}$, and this abelian scheme is projective over* Spec(R') *and uniquely functorial in $A_{k'}$. Moreover, the analogous such data*

$$1 \to \mathfrak{T}' \to \mathfrak{A}'_{R'} \to \mathfrak{B}' \to 1 \tag{2.1.2}$$

exist for $A^\vee_{k'}$, say with $B'_{R'}$ the algebraization of $\mathfrak{B}'$, and $B'_{R'}$ is canonically identified with $B^\vee_{R'}$ in such a manner that the composite isomorphism $B_{R'} \simeq (B'_{R'})' \simeq (B'_{R'})^\vee \simeq B^{\vee\vee}_{R'}$ is the double-duality isomorphism.

Proof. See [BL2, §1, §6]. ∎

We say that A as in Theorem 2.1.9 has *semistable reduction* over k' and (by abuse of terminology) we call $\mathfrak{A}_{R'}$ the *formal semistable model* of $A_{/k'}$ (even though its generic fiber inside of $A_{k'}^{\mathrm{an}}$ is a rather small quasi-compact open subgroup when $\mathfrak{T} \neq 1$). The specific identification of $B_{R'}$ and $B'_{R'}$ as dual abelian schemes (or equivalently $\mathfrak{B}$ and $\mathfrak{B}'$ as dual formal abelian schemes, or the k'-fibers B and B' as dual abelian varieties over k') is part of the constructions in the proof of Theorem 2.1.9. It is natural to ask for an intrinsic characterization of this isomorphism, such as describing the induced duality pairing between $B[N]$ and $B'[N]$ into μ_N over k' for all $N \geq 1$ in terms that are *unrelated* to the proof of Theorem 2.1.9. We address this matter in Theorem A.3.1; it is required in our study of how canonical subgroups interact with duality for abelian varieties.

Example 2.1.10. Suppose that A admits a semi-abelian model A_R over the valuation ring R of k. By [F, §2, Lemma 1], A_R is uniquely functorial in A. The formal completion $\widehat{A}_R$ of A_R along an ideal of definition of R is a formal semi-abelian scheme over $\mathrm{Spf}(R)$ and there exists a canonical quasi-compact open immersion of k-groups $i_A : \widehat{A}_R^{\mathrm{rig}} \hookrightarrow A^{\mathrm{an}}$ [C1, 5.3.1(4)]. Hence, in Theorem 2.1.9 for A we may take $k' = k$ and then the pair $(\mathfrak{A}_R, \iota)$ is uniquely identified with $(\widehat{A}_R, i_A)$. Also, the associated formal torus and formal abelian scheme arise from the corresponding filtration on the reduction of A_R modulo ideals of definition of R (using infinitesimal lifting of the maximal torus over the residue field [SGA3, IX, Thm. 3.6bis]).

2.2. **Canonical subgroups.**

Let k be an analytic extension field over $\mathbf{Q}_p$, and normalize the absolute value by the condition $|p| = 1/p$.

Definition 2.2.1. An abelian variety A over k is *ordinary* if the formal abelian scheme $\mathfrak{B}$ as in (2.1.1) has ordinary reduction over the residue field of k'.

Example 2.2.2. If $\mathfrak{B} = 0$ (potentially purely toric reduction) then A is ordinary. The reader may alternatively take this to be an *ad hoc* definition when $\mathfrak{B}$ vanishes.

Clearly the property of being ordinary is preserved under isogeny, duality, and extension of the analytic base field. In particular, A is ordinary if and only if $A^\vee$ is ordinary. It would be more accurate to use the terminology "potentially ordinary," but this should not lead to any confusion.

Let us fix an abelian variety A over k with dimension $g \geq 1$ and fix a choice of k'/k as in Theorem 2.1.9. Let $\widehat{\mathfrak{A}}_{R'}$ denote the formal completion of the formal $\mathrm{Spf}(R')$-group $\mathfrak{A}_{R'}$ along its identity section. The Lie algebra of $\mathfrak{A}_{R'}$ is a finite free R'-module of rank g, and upon choosing a basis we may identify the formal completion $\widehat{\mathfrak{A}}_{R'}$ of $\mathfrak{A}_{R'}$ along the identity with the pointed formal spectrum $\mathrm{Spf}(R'[\![X_1, \dots, X_g]\!])$ whose adic topology is defined by powers of the ideal generated by the augmentation ideal and an ideal of definition of R'.

For any positive integer n, the p^n-torsion $\mathfrak{A}_{R'}[p^n]$ has a natural structure of finite flat commutative R'-group that is an extension of $\mathfrak{B}[p^n]$ by $\mathfrak{T}[p^n]$. The $\mathfrak{A}_{R'}[p^n]$'s are the torsion-levels of a p-divisible group $\mathfrak{A}_{R'}[p^\infty]$ over the henselian local ring R', and so there is an identity component $\mathfrak{A}_{R'}[p^\infty]^0$. Since R' is p-adically separated and complete, the formal group $\widehat{\mathfrak{A}}_{R'}$ coincides with the one attached to $\mathfrak{A}_{R'}[p^\infty]^0$ (via [Me, II, Cor. 4.5]). In particular, the Lie algebra $\mathrm{Lie}(\mathfrak{A}_{R'}) = \mathrm{Lie}(\widehat{\mathfrak{A}}_{R'})$ functorially coincides with the Lie algebra of $\mathfrak{A}_{R'}[p^\infty]$.

The local-local part $\mathfrak{A}_{R'}[p^\infty]^{00}$ of $\mathfrak{A}_{R'}[p^\infty]$ coincides with the local-local part of the p-divisible group of $\mathfrak{B}$. Hence, if we run through the above procedure with $A^\vee$ in the role of A then the corresponding local-local part $\mathfrak{A}'_{R'}[p^\infty]^{00}$ of the p-divisible group of the associated formal semi-abelian model $\mathfrak{A}'_{R'}$ over R' is canonically identified with $\mathfrak{B}'[p^\infty]^{00}$, where $\mathfrak{B}'$ is as in (2.1.2), and this is canonically isomorphic to the p-divisible group $\mathfrak{B}^\vee[p^\infty]^{00} \simeq (\mathfrak{B}[p^\infty]^{00})^\vee = (\mathfrak{A}_{R'}[p^\infty]^{00})^\vee$ that is dual to $\mathfrak{A}_{R'}[p^\infty]^{00}$.

The geometric points of the generic fiber of the identity component $\mathfrak{A}_{R'}[p^n]^0 = \widehat{\mathfrak{A}}_{R'}[p^n]$ are identified with the integral p^n-torsion points of the formal group $\widehat{\mathfrak{A}}_{R'}$ with values in valuation rings of finite extensions of k'. Hence, as a subgroup of $A[p^n](\overline{k})$ this generic fiber is Galois-invariant. By Galois descent, we may therefore make the definition:

Definition 2.2.3. The unique k-subgroup in $A[p^n]$ that descends $(\mathfrak{A}_{R'}[p^n]^0)_{k'}$ is denoted $A[p^n]^0$.

Despite the notation, $A[p^n]^0$ depends on A and not just on $A[p^n]$. For later reference, we record the following trivial lemma (valid even in the setting of Example 2.2.2):

Lemma 2.2.4. *The k-subgroup $A[p^n]^0$ is independent of the choice of k', and its order satisfies $\#A[p^n]^0 \geq p^{ng}$ with equality for one (and hence all) n if and only if A is ordinary. If equality holds then $A^\vee$ is ordinary and so $A^\vee[p^n]^0$ also has order p^{ng} for all $n \geq 1$.*

Definition 2.2.5. The *size* of an integral point x in $\widehat{\mathfrak{A}}_{R'}$ with values in the valuation ring of an analytic extension of k' is $\mathrm{size}(x) \stackrel{\mathrm{def}}{=} \max_j |X_j(x)| < 1$ for a choice of formal parameters X_j for the formal group $\widehat{\mathfrak{A}}_{R'}$ over R'.

This notion of "size" is independent of the choice of X_j's, and so it is Galois-invariant over k. For any $0 < r < 1$, let $A[p^n]^0_{\leq r} \subseteq A[p^n]^0$ denote the k-subgroup whose geometric points are those for which the associated integral point in $\widehat{\mathfrak{A}}_{R'}$ has size $\leq r$; this k-subgroup is independent of the choice of k'/k.

Lemma 2.2.6. *If $n \geq 1$ and $0 < r < p^{-1/p^{n-1}(p-1)}$ then $A[p^n]^0_{\leq r}$ is killed by p^{n-1}.*

Proof. For the case $n = 1$, pick a p-torsion geometric point $x = (x_1, \dots, x_g)$. Choose j_0 such that $|x_{j_0}| = \mathrm{size}(x)$. The power series $[p]^*(X_j)$ has vanishing constant term and has linear term pX_j. By factoring $[p]$ over R'/pR' through the relative Frobenius morphism [SGA3, VII$_\mathrm{A}$, §4.2-4.3], we have

$$[p]^*(X_j) = pX_j + h_j(X_1^p, \dots, X_g^p) + pf_j(X_1, \dots, X_g) \tag{2.2.1}$$

with h_j a formal power series over R' having constant term 0 and f_j a formal power series over R' with vanishing terms in total degree < 2. Evaluating at x,

$$0 = X_{j_0}([p](x)) = ([p]^*(X_{j_0}))(x_1, \dots, x_g) = px_{j_0} + h_{j_0}(x_1^p, \dots, x_g^p) + pf_{j_0}(x_1, \dots, x_g). \tag{2.2.2}$$

Assume $x \neq 0$, so $x_{j_0} \neq 0$. The final term on the right in (2.2.2) has absolute value at most $|px_{j_0}^2| < |px_{j_0}|$, so the middle term on the right in (2.2.2) has absolute value exactly $|px_{j_0}| = |x_{j_0}|/p$. This middle term clearly has absolute value at most $|x_{j_0}|^p$, so $|x_{j_0}|/p \leq |x_{j_0}|^p$. Since $|x_{j_0}| > 0$, we obtain $|x_{j_0}| \geq p^{-1/(p-1)}$. But $|x_{j_0}| = \mathrm{size}(x)$, so we conclude $\mathrm{size}(x) \geq p^{-1/(p-1)}$ for any nonzero p-torsion geometric point x. Hence, $A[p]^0_{<p^{-1/(p-1)}} = 0$.

Now we prove that $A[p^n]^0_{\leq r}$ is killed by p^{n-1} if $0 < r < p^{-1/p^{n-1}(p-1)}$, the case $n = 1$ having just been settled. Proceeding by induction, we may assume $n > 1$ and we choose a point $x \in A[p^n]^0_{\leq r}$ with $r < p^{-1/p^{n-1}(p-1)}$. We wish to prove $[p]^{n-1}(x) = 0$. If x has size $< p^{-1/(p-1)}$ then $[p]^{n-1}(x) \in A[p]^0_{<p^{-1/(p-1)}} = \{0\}$. Hence, we can assume x has size at least $p^{-1/(p-1)}$. Under this assumption we claim $\mathrm{size}([p]x) \leq \mathrm{size}(x)^p$, so $[p](x) \in A[p^{n-1}]^0_{\leq r^p}$ with $r^p < p^{-1/p^{n-2}(p-1)}$, and thus induction would give $[p]^{n-1}(x) = [p]^{n-2}([p]x) = 0$ as desired. It therefore suffices to prove in general that for any point x of $\widehat{\mathfrak{A}}_{R'}$ with value in the valuation ring of an analytic extension of k' such that $\mathrm{size}(x) \geq p^{-1/(p-1)}$, necessarily $[p](x)$ has size at most $\mathrm{size}(x)^p$. Letting $x_j = X_j(x)$, we can pick j_0 so that $|x_{j_0}| = \mathrm{size}(x) \geq p^{-1/(p-1)}$. Our problem is to prove that the absolute value of $[p]^*(X_j)$ at x is at most $|x_{j_0}|^p$ for all j. Upon evaluating the right side of (2.2.1) at x, the first term has absolute value $|x_j|/p \leq |x_{j_0}|/p \leq |x_{j_0}|^p$, the middle term has absolute value at most $|x_{j_0}|^p$, and the final term has absolute value at most $|x_{j_0}|^2/p \leq |x_{j_0}|/p \leq |x_{j_0}|^p$. Thus, we get the desired upper bound on $\mathrm{size}([p]x)$ when x has size at least $p^{-1/(p-1)}$. ∎

Definition 2.2.7. For $n \geq 1$, a *level-n canonical subgroup* in A is a k-subgroup of the form $G_n = A[p^n]^0_{\leq r}$ for some $r \in (0,1)$ such that G_n has geometric fiber that is finite free of rank $g = \dim A$ as a $\mathbf{Z}/p^n\mathbf{Z}$-module.

An equivalent recursive formulation of the definition for $n > 1$ is that the subgroup has the form $A[p^n]^0_{\leq r}$ for some $r \in (0,1)$ and has order p^{ng} with p^{n-1}-torsion subgroup that is a level-$(n-1)$ canonical subgroup, so for such a G_n and $1 \leq m \leq n$ the subgroup $G_n[p^m]$ is a level-m canonical subgroup. In concrete terms, if K/k is an algebraically closed non-archimedean extension then a level-n canonical subgroup is a subgroup of p^{ng} points in $A[p^n](K)$ that are "closer" to the identity (in $A(K)$) than all other points in $A[p^n](K)$ (and we also impose an additional freeness condition on its $\mathbf{Z}/p^n\mathbf{Z}$-module structure). In [C4, Thm. 4.2.5] it is shown that if $g = 1$ then Definition 2.2.7 is equivalent to another definition used in [Bu] and [G].

The following lemma is trivial:

Lemma 2.2.8. *A level-n canonical subgroup is unique if it exists, and the formation of such a subgroup is compatible with change in the base field. If such a subgroup exists after an analytic extension on k then it exists over k.*

In view of the functoriality of $\mathfrak{A}_{R'}$ in $A_{k'}$, level-n canonical subgroups are functorial with respect to isogenies whose degree is prime to p. In particular, if two abelian varieties over k are related by an isogeny of degree not divisible by p then one of these abelian varieties admits a level-n canonical subgroup if and only if the other does. The restriction that the isogeny have degree prime to p cannot be dropped, as is clear even in the case $g = 1$ [K, Thm. 3.10.7(1)].

An immediate consequence of Lemma 2.2.6 is that a level-n canonical subgroup must uniformly move out to the edge of the formal group as $n \to \infty$:

Theorem 2.2.9. *If $0 < r < 1$ and $A[p^n]^0_{\leq r}$ is a level-n canonical subgroup then $r \geq p^{-1/p^{n-1}(p-1)} = |\zeta_{p^n} - 1|$ for a primitive p^nth root of unity ζ_{p^n}.*

Remark 2.2.10. In the ordinary case the subgroup $A[p^n]^0 = A[p^n]^0_{\leq p^{-1/p^{n-1}(p-1)}}$ has order p^{ng}, so it is the level-n canonical subgroup in $A[p^n]$. Hence, by Lemma 2.2.4 the inequality $\#A[p^n]^0 \geq p^{ng}$ is an equality for one (and hence all) $n \geq 1$ if and only if A is ordinary, in which case there exist level-n canonical subgroups for all $n \geq 1$. Conversely, if there exists a level-n subgroup G_n for all n then A *must* be ordinary. Indeed, suppose A is not ordinary, so $A[p]^0$ contains a point x_0 not in G_1. By Theorem 2.2.9 with $n = 1$, $\mathrm{size}(x_0) \in (p^{-1/(p-1)}, 1)$. For $n \geq 1$ such that A has a level-n canonical subgroup G_n we have $x_0 \notin G_n$ since $G_n[p] = G_1$, so the size of every point in G_n is strictly less than $\mathrm{size}(x_0)$. By Theorem 2.2.9 we conclude $\mathrm{size}(x_0) > p^{-1/p^{n-1}(p-1)}$, so

$$n < 1 + \log_p\left(\frac{\log_p(\mathrm{size}(x_0)^{-1})^{-1}}{p-1}\right) \in (1, \infty).$$

We do not impose any requirements concerning how a level-n canonical subgroup G_n in A should interact with the duality between $A[p^n]$ and $A^\vee[p^n]$ (*e.g.*, is $(A[p^n]/G_n)^\vee \subseteq A^\vee[p^n]$ a level-n canonical subgroup of $A^\vee$?), nor do we require that its finite flat schematic closure (after a finite extension k'/k) in $\mathfrak{A}_{R'}[p^n]^0$ reduces to the kernel of the n-fold relative Frobenius on $\mathfrak{A}_{R'}$ mod $\mathfrak{m}_{R'}$. In Theorem 4.1.1 and Theorem 4.3.3 we will show that there is good interaction of G_n with respect to duality (resp. the n-fold relative Frobenius kernel modulo p^λ for an arbitrary but fixed $\lambda \in (0,1) \cap \mathbf{Q}$) when the Hasse invariant of A (see §2.3) is sufficiently near 1 in a sense that is determined solely by p, $g = \dim A$, and n (resp. p, g, n, and λ). We do not know if $A^\vee$ necessarily admits admits a level-n canonical subgroup whenever A does.

Remark 2.2.11. The formation of canonical subgroups is not well-behaved with respect to products or duality in general, but this is largely an artifact of Hasse invariants "far" from 1. We shall give counterexamples in Example 2.3.3.

2.3. **Hasse invariant.** Let A be an abelian variety over k as in §2.2, and let k'/k and $\mathfrak{A}_{R'}$ be as in Theorem 2.1.9. Let $\mathscr{G}$ be the mod-pR' reduction of the p-divisible group $\mathfrak{A}_{R'}[p^\infty]$ over R', so we have a Verschiebung morphism $V_{\mathscr{G}} : \mathscr{G}^{(p)} \to \mathscr{G}$ over $\mathrm{Spec}(R'/pR')$ and on the Lie algebras this induces an R'/pR'-linear map

$$\mathrm{Lie}(V_{\mathscr{G}}) : \mathrm{Lie}(\mathscr{G})^{(p)} \to \mathrm{Lie}(\mathscr{G})$$

between finite free R'/pR'-modules of the same rank g; it is an isomorphism if and only if $V_{\mathscr{G}}$ is étale, which is to say that the identity component of $\mathscr{G}$ is multiplicative. That is, this map is an isomorphism if and only if A is ordinary in the sense of Definition 2.2.1. Up to unit multiple, there is a well-defined determinant $\det(\mathrm{Lie}(V_{\mathscr{G}})) \in R'/pR'$; this is taken to be 1 (or a unit) when $A = 0$. We let $a_{A_{k'}} \in R'$ be a representative for $\det(\mathrm{Lie}(V_{\mathscr{G}})) \in R'/pR'$, so $a_{A_{k'}}$ is well-defined modulo p up to unit multiple and therefore the following definition is intrinsic to A over k:

Definition 2.3.1. The *Hasse invariant* of A is $h(A) = \max(|a_{A_{k'}}|, 1/p) \in [1/p, 1] \cap \sqrt{|k^\times|}$.

Obviously $h(A_1 \times A_2) = \max(h(A_1)h(A_2), 1/p)$, and $h(A)$ is invariant under isogenies with degree prime to p. In the 1-dimensional case Definition 2.3.1 recovers the notion of Hasse invariant for elliptic curves in [C4].

Example 2.3.2. With notation as in Theorem 2.1.9, let $B = (B_R)_k$ be the generic fiber of the algebraization of $\mathfrak{B}$. We have $h(A) = h(B)$ because (2.1.1) induces an exact sequence on p-divisible groups and hence on their Lie algebras (and the Verschiebung on the R'/pR'-torus $\mathfrak{T}$ mod pR' is an isomorphism). Thus, $h(A) = 1$ if and only if the Verschiebung for $\mathfrak{B}$ mod pR' is étale, which is to say that the abelian variety $\mathfrak{B}$ mod $\mathfrak{m}_{R'}$ is ordinary. Hence, $h(A) = 1$ if and only if A is ordinary. Note that this reasoning is applicable even if $B = 0$ (and so A has potentially purely toric reduction).

Example 2.3.3. Let E and E' be elliptic curves with supersingular reduction such that $h(E), h(E') \in (p^{-p/(p+1)}, 1)$ and $h(E) > (ph(E'))^p$. By [K, Thm. 3.10.7(1)] each of E and E' admits a level-1 canonical subgroup but all p-torsion from E has smaller size than all nonzero p-torsion from E'. Hence, $A = E \times E'$ admits a level-1 canonical subgroup G_1, namely $G_1 = E[p]$, but this is not the product of the level-1 canonical subgroups of E and E'. Also, $(A[p]/G_1)^\vee \subseteq A^\vee[p]$ is not the level-1 canonical subgroup of $A^\vee$ since A is principally polarized and G_1 is not isotropic for the induced Weil self-pairing.

There are two reasons why we do not consider the failure of formation of canonical subgroups to commute with products and duality (as in Example 2.3.3) to be a serious deficiency. First of all, our interest in canonical subgroups is largely restricted to the study of abelian varieties with a fixed dimension and so it is the consideration of isogenies rather than products that is the more important structure to study in the context of canonical subgroups. Second, if we take Hasse invariants sufficiently near 1 in a "universal" manner then the compatibilities with products and duality are rescued. More specifically, it follows from Theorem 4.1.1 that for any fixed $n, g, g' \geq 1$ there exist $h(p, g, n), h(p, g', n) \in (1/p, 1)$ such that if A and A' are abelian varieties with respective dimensions g and g' over any $k/\mathbf{Q}_p$ and the inequalities $h(A) > h(p, g, n)$ and $h(A') > h(p, g', n)$ hold then both A and A' admit level-n canonical subgroups G_n and G'_n and moreover $G_n \times G'_n$ is a level-n canonical subgroup in $A \times A'$. Since $h(A), h(A') \geq h(A \times A')$, by taking $h(A \times A')$ to be close to 1 we force $h(A)$ and $h(A')$ to be close to 1. Theorem 4.1.1 also ensures that $(A[p^n]/G_n)^\vee$ is the level-n canonical subgroup of $A^\vee$ when $h(A)$ is sufficiently near 1 (in a manner that depends only on p, g, and n).

Our aim is to prove the existence of a level-n canonical subgroup in A when $h(A)$ is sufficiently close to 1, where "sufficiently close" *only* depends on p, $\dim A$, and n, and we wish to uniquely relativize this construction in rigid-analytic families. Zahrin's trick will shift many (but not all) problems to the principally polarized case, provided that the Hasse invariant is unaffected by passage to the dual abelian variety (as then $h((A \times A^\vee)^4) = h(A)^8$ when $h(A) > p^{-1/8}$). Thus, we now prove:

Theorem 2.3.4. *For any abelian variety A over k, $h(A) = h(A^\vee)$.*

Let k'/k be a finite extension as in Theorem 2.1.9, and let R' be the valuation ring of k'. Since $\mathfrak{B}'$ in Theorem 2.1.9 is isomorphic to $\mathfrak{B}^\vee$, by Example 2.3.2 it suffices to prove Theorem 2.3.4 for $(B_R)_k$ rather than A. Thus, we may formulate our problem more generally for the p-divisible group Γ of an arbitrary abelian scheme X over R'/pR': we claim that the "determinant" of $\mathrm{Lie}(V_\Gamma)$ coincides with the "determinant" of $\mathrm{Lie}(V_{\Gamma^\vee})$ up to unit multiple, where the dual p-divisible group $\Gamma^\vee$ is identified with the p-divisible group of the dual abelian scheme $X^\vee$ and we write V_Γ and $V_{\Gamma^\vee}$ to denote the relative Verschiebung morphisms. In other words, we claim that both determinants generate the same ideal in R'/pR'. This is a special case of:

Theorem 2.3.5. *Let Γ be a p-divisible group over an $\mathbf{F}_p$-scheme S. The locally principal quasi-coherent ideals* $\det(\mathrm{Lie}(V_\Gamma))$ *and* $\det(\mathrm{Lie}(V_{\Gamma^\vee}))$ *in $\mathscr{O}_S$ coincide.*

Proof. The first step is to reduce to a noetherian base scheme. This is a standard argument via consideration of torsion-levels, as follows. The cotangent space along the identity section for a p-divisible group Γ over an $\mathbf{F}_p$-scheme is identified with that of any finite-level truncation $\Gamma[p^n]$ for $n \geq 1$ [Me, II, 3.3.20], and this truncation "is" a vector bundle over S whose formation commutes with base change; the same holds for Lie algebras. The relative Frobenius morphism for Γ is an isogeny, so its kernel $\Gamma[F]$ is a finite locally free commutative subgroup of the level-1 truncated BT-group $\Gamma[p]$ and moreover it is the kernel of relative Frobenius for $\Gamma[p]$; the same holds for $\Gamma^\vee$ in the role of Γ. Hence, by working locally on S we can descend $\Gamma[p]$ to a level-1 truncated BT-group Γ'_1 over a locally noetherian base (again denoted S) such that $\Gamma'_1[F]$ and $(\Gamma'_1)^\vee[F]$ are finite locally free S-groups, and so by [Me, II, 2.1.3, 2.1.4] both Γ'_1 and $(\Gamma'_1)^\vee$ have relative cotangent spaces and Lie algebras that are vector bundles whose formation commutes with base change. Our problem may be restated in terms of such a descent of $\Gamma[p]$ over a locally noetherian $\mathbf{F}_p$-scheme S', and so it is enough to solve the restated problem after base change from such an S' to every affine scheme $\mathrm{Spec}(C)$ over S' with C a complete local noetherian ring having algebraically closed residue field. (Of course, we just need to treat one such faithfully flat local extension C of each local ring on S'.) We may and do endow the equicharacteristic C with a compatible structure of algebra over its residue field. By a theorem of Grothendieck [Ill, Thm. 4.4], the level-1 truncated BT-group Γ'_1 over C may be realized as the p-torsion of a p-divisible group over C (whose dual has p-torsion given by $(\Gamma'_1)^\vee$). Hence, it suffices to solve the original problem for p-divisible groups over $S = \mathrm{Spec}\, C$ with C a complete local *noetherian* k-algebra having residue field k, where k is a perfect field with characteristic p.

It is enough to treat the case of the universal equicharacteristic deformation of the k-fiber $\Gamma \otimes_C k$. This universal deformation ring is a unique factorization domain (even a formal power series ring over k), so to check an equality of principal ideals in this local ring it suffices to work locally at the height-1 primes. Also, by a calculation in Cartier theory [R, Lemma 4.2.3], the generic fiber of the universal equicharacteristic deformation is ordinary. Hence, we are reduced to the case when $S = \mathrm{Spec}(R)$ for an equicharacteristic-p discrete valuation ring R such that the generic fiber of Γ is ordinary.

The two maps

$$\mathrm{Lie}(V_\Gamma) : \mathrm{Lie}(\Gamma^{(p)}) \to \mathrm{Lie}(\Gamma), \;\; \mathrm{Lie}(V_{\Gamma^\vee}) : \mathrm{Lie}(\Gamma^{\vee,(p)}) \to \mathrm{Lie}(\Gamma^\vee)$$

between finite free R-modules are injective due to generic ordinarity, and so each map has finite-length cokernel. The determinant ideals for these two maps are given by the products of the invariant factors for the torsion cokernel modules. For any linear injection T between finite free R-modules of the same positive rank, the linear dual $T^\vee$ is also injective and the torsion R-modules $\mathrm{coker}(T)$ and $\mathrm{coker}(T^\vee)$ have the same invariant factors. Thus, it suffices to prove that the cokernels of $\mathrm{Lie}(V_\Gamma)$ and $\mathrm{Lie}(V_{\Gamma^\vee})^\vee$ are canonically isomorphic as R-modules. Such an isomorphism is provided by the next theorem. ■

Theorem 2.3.6. *Let Γ be a p-divisible group over an $\mathbf{F}_p$-scheme X. The $\mathscr{O}_X$-modules* $\mathrm{coker}(\mathrm{Lie}(V_\Gamma))$ *and* $\mathrm{coker}(\mathrm{Lie}(V_{\Gamma^\vee})^\vee)$ *are canonically isomorphic.*

Proof. The theory of universal vector extensions of Barsotti–Tate groups [Me, Ch. IV, 1.14] provides a canonical exact sequence of vector bundles

$$0 \to \mathrm{Lie}(\Gamma^\vee)^\vee \to \mathrm{Lie}(E(\Gamma)) \to \mathrm{Lie}(\Gamma) \to 0$$

on X, where $E(\Gamma)$ is the universal vector extension of Γ. The formation of this sequence is functorial in Γ and compatible with base change on X, so by using functoriality with respect to the relative Frobenius and Verschiebung morphisms F_Γ and V_Γ of Γ over X and using the identities $V_\Gamma^\vee = F_{\Gamma^\vee}$ and $F_\Gamma^\vee = V_{\Gamma^\vee}$ (via the canonical isomorphism $(\Gamma^{(p)})^\vee \simeq (\Gamma^\vee)^{(p)}$ and [SGA3, $\mathrm{VII_A}$, §4.2-§4.3]) we get the following commutative diagram of vector bundles in which the rows are short exact sequences:

(2.3.1)
$$\begin{array}{ccccccccc}
0 & \longrightarrow & \mathrm{Lie}(\Gamma^{\vee,(p)})^\vee & \longrightarrow & \mathrm{Lie}(E(\Gamma^{(p)})) & \longrightarrow & \mathrm{Lie}(\Gamma^{(p)}) & \longrightarrow & 0 \\
 & & \Big\downarrow{\scriptstyle 0=\mathrm{Lie}(F_{\Gamma^\vee})^\vee} & & \Big\downarrow{\scriptstyle \mathrm{Lie}(E(V_\Gamma))} & & \Big\downarrow{\scriptstyle \mathrm{Lie}(V_\Gamma)} & & \\
0 & \longrightarrow & \mathrm{Lie}(\Gamma^\vee)^\vee & \longrightarrow & \mathrm{Lie}(E(\Gamma)) & \longrightarrow & \mathrm{Lie}(\Gamma) & \longrightarrow & 0 \\
 & & \Big\downarrow{\scriptstyle \mathrm{Lie}(V_{\Gamma^\vee})^\vee} & & \Big\downarrow{\scriptstyle \mathrm{Lie}(E(F_\Gamma))} & & \Big\downarrow{\scriptstyle \mathrm{Lie}(F_\Gamma)=0} & & \\
0 & \longrightarrow & \mathrm{Lie}(\Gamma^{(p),\vee})^\vee & \longrightarrow & \mathrm{Lie}(E(\Gamma^{(p)})) & \longrightarrow & \mathrm{Lie}(\Gamma^{(p)}) & \longrightarrow & 0
\end{array}$$

Due to the vanishing in the upper-left and lower-right parts of (2.3.1), we arrive at a natural complex

$$0 \to \mathrm{Lie}(\Gamma^{(p)}) \to \mathrm{Lie}(E(\Gamma)) \to \mathrm{Lie}(\Gamma^{(p),\vee})^{\vee} \to 0 \tag{2.3.2}$$

whose formation commutes with base change and which fits into the vertical direction in the following commutative diagram:

$$\begin{array}{ccccccccc}
 & & & & \mathrm{Lie}(\Gamma^{(p)}) & & & & \\
 & & & & \downarrow & \searrow^{\mathrm{Lie}(V_\Gamma)} & & & \\
0 & \longrightarrow & \mathrm{Lie}(\Gamma^{\vee})^{\vee} & \longrightarrow & \mathrm{Lie}(E(\Gamma)) & \longrightarrow & \mathrm{Lie}(\Gamma) & \longrightarrow & 0 \\
 & & & \searrow_{\mathrm{Lie}(V_{\Gamma^\vee})^{\vee}} & \downarrow & & & & \\
 & & & & \mathrm{Lie}(\Gamma^{(p),\vee})^{\vee} & & & &
\end{array}$$

Granting the exactness of (2.3.2) for a moment, we can conclude via the elementary:

Lemma 2.3.7. *If*

$$\begin{array}{ccccc}
 & & \mathscr{M}' & & \\
 & & \downarrow & \searrow^{f_1} & \\
\mathscr{N}_2 & \longrightarrow & \mathscr{M} & \longrightarrow & \mathscr{N}_1 \\
 & \searrow_{f_2} & \downarrow & & \\
 & & \mathscr{M}'' & &
\end{array}$$

is a commutative diagram of sheaves of modules with the vertical and horizonal diagrams each short exact sequences, then $\mathrm{coker}(f_1)$ *and* $\mathrm{coker}(f_2)$ *are naturally isomorphic.*

Proof. The map $\mathscr{M} \twoheadrightarrow \mathrm{coker}(f_1)$ kills $\mathscr{M}'$ and so uniquely factors through a map $\mathscr{M}'' \twoheadrightarrow \mathrm{coker}(f_1)$ that kills $\mathrm{image}(f_2)$ and so induces a map $\phi : \mathrm{coker}(f_2) \twoheadrightarrow \mathrm{coker}(f_1)$. We similarly construct a map $\psi : \mathrm{coker}(f_1) \twoheadrightarrow \mathrm{coker}(f_2)$, and the composites $\phi \circ \psi$ and $\psi \circ \phi$ are clearly equal to the identity. ∎

It remains to prove that (2.3.2) is short exact. Since this is a three-term complex of finite locally free sheaves, it is equivalent to check the short exactness on geometric fibers over X. The formation of (2.3.2) is compatible with base change on X, so we may assume $X = \mathrm{Spec}(k)$ for an algebraically closed field k of characteristic p. Under the comparison isomorphism between classical and crystalline Dieudonné theory for p-divisible groups G over k [MM, Ch. 2, Cor. 7.13, §9, Thm. 15.3], there is a canonical k-linear isomorphism $\mathrm{Lie}(E(G)) \simeq \mathbf{D}_k(G^\vee) \otimes_{W(k)} k \simeq \mathbf{D}_k(G^\vee[p])$ with $\mathbf{D}_k$ denoting the classical Dieudonné functor. (The classical theory used in [MM] is naturally isomorphic to the one constructed in [Fo, Ch. III].) Hence, (2.3.1) can be written as an abstract commutative diagram of k-vector spaces (with short exact sequences in the horizontal direction):

$$\begin{array}{ccccccccc}
0 & \longrightarrow & t^*_{\Gamma^\vee,(p)} & \longrightarrow & \mathbf{D}_k(\Gamma^{(p),\vee}[p]) & \xrightarrow{\beta} & t_{\Gamma^{(p)}} & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle 0} & & \downarrow{\scriptstyle \mathbf{D}_k(V^\vee_{\Gamma[p]})} & & \downarrow{\scriptstyle \mathrm{Lie}(V_\Gamma)} & & \\
0 & \longrightarrow & t^*_{\Gamma^\vee} & \longrightarrow & \mathbf{D}_k(\Gamma^\vee[p]) & \longrightarrow & t_\Gamma & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \mathrm{Lie}(V_{\Gamma^\vee})^*} & & \downarrow{\scriptstyle \mathbf{D}_k(F^\vee_{\Gamma[p]})} & & \downarrow{\scriptstyle 0} & & \\
0 & \longrightarrow & t^*_{\Gamma^{(p)},\vee} & \xrightarrow[\alpha]{} & \mathbf{D}_k(\Gamma^{(p),\vee}[p]) & \longrightarrow & t_{\Gamma^{(p)}} & \longrightarrow & 0
\end{array}$$

where we write t_H to denote the tangent space to a p-divisible group or finite commutative group scheme H over k and we write t^*_H to denote its linear dual. Since $V^\vee_{\Gamma[p]} = F_{\Gamma^\vee[p]}$ and $F^\vee_{\Gamma[p]} = V_{\Gamma^\vee[p]}$, we may respectively identify the top and bottom maps in the middle column with the k-linearizations of the semilinear F and V maps on the classical Dieudonné module $\mathbf{D}_k(\Gamma^\vee[p])$.

From the lower-left part of the diagram we get an abstract k-linear injection $\alpha^{(p^{-1})} : t^*_{\Gamma^\vee} \hookrightarrow \mathbf{D}_k(\Gamma^\vee[p])$ onto a subspace containing the image of the semilinear Verschiebung operator V on $\mathbf{D}_k(\Gamma^\vee[p])$. Likewise, we get an abstract k-linear surjection $\beta^{(p^{-1})} : \mathbf{D}_k(\Gamma^\vee[p]) \twoheadrightarrow t_\Gamma$ through which the semilinear Frobenius operator F on $\mathbf{D}_k(\Gamma^\vee[p])$ factors. Since $\ker(V) = \mathrm{im}(F)$ on $\mathbf{D}_k(\Gamma^\vee[p])$, due to $\Gamma^\vee[p]$ being the p-torsion of a p-divisible group, our exactness problem with (2.3.2) is thereby reduced to proving two things: (i) the inclusion $\mathrm{im}(V) \subseteq t^*_{\Gamma^\vee}$ inside of $\mathbf{D}_k(\Gamma^\vee[p])$ is an equality, and (ii) the Frobenius-semilinear surjection $t_\Gamma \twoheadrightarrow \mathrm{im}(F)$ is injective. These conditions respectively say $\dim(\mathrm{im}(V)) \overset{?}{=} \dim \Gamma^\vee$ and $\dim(\mathrm{im}(F)) \overset{?}{=} \dim \Gamma$. If h denotes the height of Γ then $h = \dim_k \mathbf{D}_k(\Gamma^\vee[p])$, so

$$\dim(\mathrm{im}(V)) = \dim(\ker F) = h - \dim(\mathrm{im}(F)), \ \ \dim \Gamma + \dim \Gamma^\vee = h,$$

whence the two desired equalities are equivalent. We check the second one, as follows. Classical Dieudonné theory provides a canonical k-linear isomorphism $t^*_G \simeq \mathbf{D}_k(G)/\mathrm{im}(F)$ for any finite commutative p-group G over k [Fo, Ch. III, Prop. 4.3], so taking $G = \Gamma^\vee[p]$ gives $\dim t_{\Gamma^\vee[p]} = h - \dim(\mathrm{im}(F))$. But $t_{\Gamma^\vee[p]} = t_{\Gamma^\vee}$, so

$$\dim(\mathrm{im}(F)) = h - \dim t_{\Gamma^\vee} = h - \dim \Gamma^\vee = \dim \Gamma.$$

■

3. Variation of Hasse invariant

Let $k/\mathbf{Q}_p$ be an analytic extension field, and $\mathscr{A} \to \mathscr{S}$ an abelian scheme over a locally finite type k-scheme. Fixing $h \in (1/p, 1) \cap \sqrt{|k^\times|}$, we wish to study the locus of points $s \in \mathscr{S}^{\mathrm{an}}$ for which $h(\mathscr{A}_s^{\mathrm{an}}) \geq h$.

3.1. **The polarized case.** It will be convenient to first consider the case of a polarized abelian scheme, and to then eliminate the polarization by a separate argument. In the polarized case, we can weaken the nature of the polarization assumption and consider a situation that is intrinsic to the rigid-analytic category:

Theorem 3.1.1. *Let $k/\mathbf{Q}_p$ be an analytic extension field, and let $A \to S$ be an abeloid space over a rigid-analytic space over k. Assume that $A_{/S}$ admits a polarization fpqc-locally on S.*

For any $h \in (p^{-1/8}, 1] \cap \sqrt{|k^\times|}$ the loci

$$S^{>h} = \{s \in S \,|\, h(A_s) > h\}, \ \ S^{\geq h} = \{s \in S \,|\, h(A_s) \geq h\}$$

are admissible opens in S and their formation is compatible with base change on S and (for quasi-separated or pseudo-separated S) with change of the base field, and the same properties hold for $S^{>p^{-1/8}}$. The map $S^{\geq h} \to S$ is a quasi-compact morphism, and for any $h \in [p^{-1/8}, 1) \cap \sqrt{|k^\times|}$ the collection $\{S^{\geq h'}\}_{h<h'\leq 1}$ is an admissible covering of $S^{>h}$ (where we require $h' \in \sqrt{|k^\times|}$).

The locus $S^{\geq 1} = \{s \in S \,|\, h(A_s) = 1\}$ is the *ordinary locus* for $A \to S$. These are the points such that the semi-abelian reduction of A_s over the residue field of a sufficiently large finite extension of $k(s)$ has ordinary abelian part. The intervention of $p^{-1/8}$ in Theorem 3.1.1 is an artifact of our method of proof (via Zahrin's trick).

Proof. The formation of the sets $S^{>h}$ and $S^{\geq h}$ is clearly compatible with base change on S, so by *fpqc* descent theory for admissible opens and admissible covers [C3, Lemma 4.2.4, Cor. 4.2.6] we may assume $A_{/S}$ admits a polarization with some constant degree d^2. By the relativization of Zahrin's trick [Mil, 16.12], $(A \times A^\vee)^4$ is principally polarized over S. For all $s \in S$ such that $h(A_s) > p^{-1/8}$ we have $h((A \times A^\vee)^4_s) = h(A_s)^8$ by Theorem 2.3.4, so by replacing A with $(A \times A^\vee)^4$ we may suppose that A is principally polarized at the expense of replacing $p^{-1/8}$ with $1/p$ in the bounds on h under consideration. Fix $N \geq 3$ not divisible by p. Working étale-locally, we may assume $A[N]$ is split. Hence, by Example 2.1.7 it suffices to treat the universal family over $\mathscr{A}^{\mathrm{an}}_{g,1,N/k}$ provided that we work with $1/p$ rather than $p^{-1/8}$.

We shall first consider the case $k = \mathbf{Q}_p$, and so we now restrict attention to $h \in p^{\mathbf{Q}}$ with $h \in [1/p, 1]$. Consider the universal abelian scheme over the $\mathbf{Z}_p$-scheme $\mathscr{A}_{g,1,N/\mathbf{Z}_p}$. By [CF, IV, 6.7(1),(3); V, 2.5, 5.8], this extends to a semi-abelian scheme $A \to X$ over a proper flat $\mathbf{Z}_p$-scheme X in which $\mathscr{A}_{g,1,N/\mathbf{Z}_p}$ equipped with its universal abelian scheme is a Zariski-open subscheme. On $\mathbf{Q}_p$-fibers, we get a proper rigid space $X^{\rm an}_{\mathbf{Q}_p}$ over $\mathbf{Q}_p$ that contains $\mathscr{A}^{\rm an}_{g,1,N/\mathbf{Q}_p}$ as a Zariski-open subspace, and we get a smooth $X^{\rm an}_{\mathbf{Q}_p}$-group $A^{\rm an}_{\mathbf{Q}_p}$ whose restriction over $\mathscr{A}^{\rm an}_{g,1,N/\mathbf{Q}_p}$ is the universal principally-polarized abeloid space of relative dimension g with full level-N structure. We now let $\mathscr{A}_{g,1,N}$ denote $\mathscr{A}_{g,1,N/\mathbf{Q}_p}$.

Let $\mathfrak{A} \to \mathfrak{X}$ be the p-adic completion of $A \to X$, so by $\mathbf{Z}_p$-properness of X we get a canonical identification $\mathfrak{X}^{\rm rig} = X^{\rm an}_{\mathbf{Q}_p}$ and a canonical isomorphism of $\mathfrak{A}^{\rm rig}$ onto an admissible open subgroup of $A^{\rm an}_{\mathbf{Q}_p}$ [C1, 5.3.1(4)]. In particular, by Example 2.1.10, the restriction of $\mathfrak{A}^{\rm rig}$ over the Zariski-open $\mathscr{A}^{\rm an}_{g,1,N}$ is an open subgroup that fiberwise realizes the Raynaud generic fiber of the unique formal semi-abelian model as in Theorem 2.1.9. That is, for each point $x \in \mathscr{A}^{\rm an}_{g,1,N} \subseteq X^{\rm an}_{\mathbf{Q}_p} = \mathfrak{X}^{\rm rig}$ the fiber of $\mathfrak{A}$ over the corresponding valuation ring ("integral point" of the proper formal scheme $\mathfrak{X}$) is the unique formal semi-abelian model of $A^{\rm an}_x$. This entire construction is compatible with arbitrary analytic extension on $\mathbf{Q}_p$ because Raynaud's theory of formal models is compatible with extension of the base field.

For any formal open affine $\mathrm{Spf}(\mathscr{R})$ in $\mathfrak{X}$, since $\mathrm{Spf}(\mathscr{R})$ has underlying topological space $\mathrm{Spec}(\mathscr{R}/p\mathscr{R})$ we may shrink this open around any of its points to arrange that the Lie algebra of $\mathfrak{A}$ over the open $\mathrm{Spf}(\mathscr{R}) \subseteq \mathfrak{X}$ is a free module over $\mathscr{R}$. Thus, up to a unit in $\mathscr{R}/p\mathscr{R}$ we get a well-defined determinant for the Verschiebung on the Lie algebra of $\mathfrak{A}$ mod $p\mathscr{R}$. Pick a representative $h_{\mathscr{R}} \in \mathscr{R}$ for this determinant. Over the admissible open locus where the admissible open $\mathrm{Spf}(\mathscr{R})^{\rm rig}$ in $\mathfrak{X}^{\rm rig} = X^{\rm an}_{\mathbf{Q}_p}$ meets $\mathscr{A}^{\rm an}_{g,1,N}$, the function $\max(|h_{\mathscr{R}}|, 1/p)$ is well-defined (independent of choices, including $h_{\mathscr{R}}$) and computes the Hasse invariant of the fibers of the universal abeloid space. Moreover, since such opens $\mathrm{Spf}(\mathscr{R})^{\rm rig}$ constitute an admissible cover of $\mathfrak{X}^{\rm rig} = X^{\rm an}_{\mathbf{Q}_p}$, we see that for any $h \in (1/p, 1] \cap p^{\mathbf{Q}}$ (resp. $h \in [1/p, 1) \cap p^{\mathbf{Q}}$) the locus $(\mathscr{A}^{\rm an}_{g,1,N})^{\geq h}$ (resp. $(\mathscr{A}^{\rm an}_{g,1,N})^{>h}$) of fibers with Hasse invariant $\geq h$ (resp. $> h$) is an admissible open in $\mathscr{A}^{\rm an}_{g,1,N}$ whose formation commutes with arbitrary extension on $\mathbf{Q}_p$, and (via the crutch of the rigid-analytic functions $h_{\mathscr{R}}$ for varying $\mathscr{R}$) similarly for any $k/\mathbf{Q}_p$ with $\sqrt{|k^\times|}$ replacing $p^{\mathbf{Q}}$. The desired quasi-compactness and "admissible covering" properties are likewise clear. ■

Recall [Ber2, 1.6.1] that there is an equivalence of categories between the full subcategory of quasi-separated rigid spaces S over k that have a locally finite admissible affinoid covering and the category of paracompact strictly k-analytic Berkovich spaces. For such S the formation of the loci $S^{>h}$ and $S^{\geq h}$ is compatible with passage to Berkovich spaces in the following sense:

Corollary 3.1.2. *Let $A \to S$ be as in Theorem* 3.1.1*, and assume that S is quasi-separated and admits a locally finite admissible affinoid covering.*

(1) *For any $h \in (p^{-1/8}, 1] \cap \sqrt{|k^\times|}$ (resp. $h \in [p^{-1/8}, 1) \cap \sqrt{|k^\times|}$) the admissible open $S^{>h}$ (resp. $S^{\geq h}$) is quasi-separated and admits a locally finite admissible affinoid cover, as does A.*
(2) *The associated map of Berkovich spaces $(S^{>h})^{\rm Ber} \to S^{\rm Ber}$ (resp. $(S^{\geq h})^{\rm Ber} \to S^{\rm Ber}$) is an open immersion (resp. strictly k-analytic domain), and its image is precisely the locus of points at which the fiber of $A^{\rm Ber} \to S^{\rm Ber}$ has Hasse invariant $> h$ (resp. $\geq h$).*

Proof. First suppose $A \to S$ admits a polarization. Passing to $(A \times A^\vee)^4$ thereby reduces us to the principally polarized case at the expense of replacing $p^{-1/8}$ with $1/p$ in what we have to prove. Since it suffices to work over the constituents of a locally finite admissible affinoid covering of S, we may use the proof of Theorem 3.1.1 in the principally polarized case to get to the situation in which S is affinoid and there is a power-bounded rigid-analytic function H on S that "computes" the Hasse invariant on all fibers after any extension of the base field. All of the assertions to be proved are obvious in this case.

In the general case it suffices to work locally on S, so we can assume that S is affinoid and that there exists an *fpqc* cover $S' \to S$ by another affinoid such that $A_{/S'}$ acquires a polarization. The results are all known in S' and we wish to deduce them in S. We know that the morphism $S^{\geq h} \to S$ is a quasi-compact open immersion, so the quasi-separated $S^{\geq h}$ obviously admits a locally finite admissible affinoid cover and on the associated Berkovich spaces this morphism is a strictly k-analytic domain. Since ${S'}^{\rm Ber} \to S^{\rm Ber}$ is a surjection (as $S' \to S$ is *fpqc*) and it is compatible with the formation of the Hasse invariant for fibers of $A^{\rm Ber}$, the compatibility of $S^{\geq h}$ with respect to passage to Berkovich spaces is a consequence of the corresponding known compatibility for ${S'}^{\geq h}$.

The map ${S'}^{\rm Ber} \rightarrow S^{\rm Ber}$ is a surjection between compact Hausdorff spaces, so it is a quotient map on underlying topological spaces. Thus, the locus in $S^{\rm Ber}$ for which the Hasse invariant is contained in a fixed open subinterval of $(1/p, 1)$ is open, as this is true on ${S'}^{\rm Ber}$. It follows that by using loci for which the Hasse invariant is contained in each of a suitable family of intervals that exhaust $(h, 1]$, the quasi-separated $S^{>h}$ has a locally finite admissible affinoid cover and $(S^{>h})^{\rm Ber} \rightarrow S^{\rm Ber}$ is a strictly k-analytic domain whose image is precisely the open locus with Hasse invariant $> h$ in $S^{\rm Ber}$. In particular, this latter map is an open immersion. ■

Remark 3.1.3. If $A_{/S}$ in Theorem 3.1.1 *fpqc*-locally admits a formal semi-abelian model then the proof of Theorem 3.1.1 can be applied without using Zahrin's trick, and so the conclusions of Theorem 3.1.1 and Corollary 3.1.2 apply to such $A_{/S}$ with $p^{-1/8}$ replaced by $1/p$. For example, this applies to the Berthelot generic fiber of the universal formal deformation of a polarized abelian variety in characteristic p.

Remark 3.1.4. The openness of $(S^{>h})^{\rm Ber}$ in $S^{\rm Ber}$ in Corollary 3.1.2(2) (and Corollary 3.2.6 below) is the reason why the construction of relative canonical subgroups in §4.3 is an example of overconvergence.

3.2. **The general algebraic case.** Now let $\mathscr{A} \rightarrow \mathscr{S}$ be an abelian scheme over a locally finite type k-scheme, but do not assume the existence of a polarization. We claim that the conclusions of Theorem 3.1.1 hold for $\mathscr{A}^{\rm an} \rightarrow \mathscr{S}^{\rm an}$. The starting point is the well-known:

Lemma 3.2.1. *If $\mathscr{A} \rightarrow \mathscr{S}$ is an abelian scheme over a normal locally noetherian scheme then it admits a polarization.*

Proof. By passing to connected components of $\mathscr{S}$ we may assume $\mathscr{S}$ is connected and hence irreducible. Let η be the generic point of $\mathscr{S}$, and pick an isogeny $\phi_\eta : \mathscr{A}_\eta \rightarrow \mathscr{A}_\eta^\vee$ that is a polarization. By the Weil extension lemma [BLR, 4.4/1], this isogeny uniquely extends to a morphism of abelian schemes $\phi : \mathscr{A} \rightarrow \mathscr{A}^\vee$ that is necessarily symmetric ($\phi^\vee = \phi$) and an isogeny, so it is a polarization if and only if the pullback $\mathscr{L}$ of the Poincaré bundle along the map $(1, \phi) : \mathscr{A} \rightarrow \mathscr{A} \times \mathscr{A}^\vee$ is $\mathscr{S}$-ample. By [EGA, III$_1$, 4.7.1] the locus $\mathscr{U}$ of ample fibers for $\mathscr{L}$ is Zariski-open in $\mathscr{S}$ and $\mathscr{L}|_{\mathscr{U}}$ relatively ample over $\mathscr{U}$, so it just has to be shown that the open immersion $\mathscr{U} \rightarrow \mathscr{S}$ is proper. By the valuative criterion, it suffices to consider the case when the base is the spectrum of a discrete valuation ring, and this case follows from special properties of line bundles on abelian varieties given in [Mum, p. 60, p. 150] (see the bottom of [CF, I, p. 6] for the argument). ■

By pullback to algebraic normalizations we get:

Corollary 3.2.2. *If k is a non-archimedean field and $\mathscr{A} \rightarrow \mathscr{S}$ is an abelian scheme over a locally finite type k-scheme then for any rigid space S equipped with a map $S \rightarrow \mathscr{S}^{\rm an}$ the pullback $A \rightarrow S$ of the analytification $\mathscr{A}^{\rm an} \rightarrow \mathscr{S}^{\rm an}$ admits a polarization after a finite surjective base change on S.*

Our goal is to prove:

Theorem 3.2.3. *The conclusions in Theorem* 3.1.1 *hold if the fpqc-local polarization hypothesis on the abeloid space $A \rightarrow S$ is replaced with the assumption that after local finite surjective base change it is a pullback of the analytification of an abelian scheme over a locally finite type k-scheme.*

Recall from Example 2.1.8 that this hypothesis on $A \rightarrow S$ is inherited by the abeloid quotient by any finite flat S-subgroup.

Proof. It suffices to treat the case when there exists a finite surjection $\widetilde{S} \rightarrow S$ such that the $\widetilde{S}$-abeloid space $A_{/\widetilde{S}}$ is a pullback of $\mathscr{A}^{\rm an} \rightarrow \mathscr{S}^{\rm an}$ for an abelian scheme $\mathscr{A}$ over a locally finite type k-scheme $\mathscr{S}$. After composing with a further finite surjective base change (such as from analytification of the normalization of $\mathscr{S}_{\rm red}$) we can assume that $A_{/\widetilde{S}}$ is polarized. Pick $h \in (p^{-1/8}, 1] \cap \sqrt{|k^\times|}$. By Theorem 3.1.1, the loci $\widetilde{S}^{>h}$ (allowing $h = p^{-1/8}$) and $\widetilde{S}^{\geq h}$ in $\widetilde{S}$ satisfy all of the desired properties. It is also clear that $\widetilde{S}^{>h}$ is the full preimage of its image $S^{>h}$ in S (allowing $h = p^{-1/8}$), and likewise with "$\geq h$". Assuming that we can prove admissibility of $S^{>h}$ (allowing $h = p^{-1/8}$) and $S^{\geq h}$ in S in general, let us see how to deduce the rest.

The compatibility with base change on S is obvious. To check that $S^{\geq h} \to S$ is a quasi-compact morphism, since any admissible open U in S has preimage $U^{\geq h}$ in $S^{\geq h}$ we have to prove that if S is quasi-compact then $S^{\geq h}$ is quasi-compact. Certainly $\widetilde{S}$ is quasi-compact, so $\widetilde{S}^{\geq h}$ is quasi-compact by Theorem 3.1.1. The restriction $\widetilde{S}^{\geq h} \to S^{\geq h}$ of the finite surjection $\widetilde{S} \to S$ is a finite surjection, so quasi-compactness of $S^{\geq h}$ follows from:

Lemma 3.2.4. *If $X' \to X$ is a quasi-compact surjection of rigid spaces and X' is quasi-compact then X is quasi-compact.*

The content in this lemma is that it is not necessary to assume X is quasi-separated.

Proof. Let $\{U_i\}$ be an admissible affinoid covering of X, so its preimage $\{U_i'\}$ in X' is an admissible covering of X' by quasi-compact opens. By quasi-compactness of X' there is a finite subcover, so it remains to show that if $\{U_i\}$ is a finite set of admissible affinoid opens in X that is a set-theoretic cover and if the finite collection of admissible quasi-compact open preimages $\{U_i'\}$ in X' is an admissible cover then $\{U_i\}$ is an admissible cover of X. Consider a morphism $f : S = \mathrm{Sp}(B) \to X$ from an affinoid space; by definition of admissibility, our problem is to prove that the admissible opens $V_i = f^{-1}(U_i)$ form an admissible cover of S (i.e., $\{V_i\}$ has a finite affinoid refinement). The map $S' = X' \times_X S \to S$ is a quasi-compact surjection onto S, so S' is quasi-compact, and the preimages $\{V_i'\}$ of $\{V_i\}$ in S' form an admissible cover (since this collection is the pullback of an admissible cover of X'). If $\{V_{ij}\}_{j \in J_i}$ is an admissible affinoid open covering of V_i then the preimage collection $\{V_{ij}'\}_{j \in J_i}$ is an admissible covering of V_i' by quasi-compact admissible opens, so by admissibility of $\{V_i'\}$ some finite collection of the V_{ij}''s covers the quasi-compact space S' set-theoretically. Hence, the corresponding finite collection of V_{ij}'s is a finite affinoid subcover of $\{V_i\}$ since $S' \to S$ is surjective. ■

Remark 3.2.5. Gabber noted that Lemma 3.2.4 is false without a quasi-compactness hypothesis on the surjection. That is, there exist (necessarily non-quasi-separated) rigid spaces X that are not quasi-compact but have a set-theoretic cover by finitely many admissible affinoid opens; for any such X, taking X' to be the disjoint union of such affinoids (equipped with the canonical map to X) provides the counterexample. To make such an X, let $D_+ = D_- = \{|t| \leq 1\}$ and let $D_0 = \{0 < |t| \leq 1\}$. Pick $c \in |k^\times| \cap (0,1)$, and let $U_\pm \subseteq D_\pm$ be defined as follows: $U_+ = \coprod_{n \geq 0} \{c^{2n+1} \leq |t| \leq c^{2n}\}$ and $U_- = \coprod_{n \geq 0} \{c^{2(n+1)} < |t| < c^{2n+1}\}$. (One uses finiteness of the number of connected components of an affinoid rigid space to show that $U_\pm$ is indeed an admissible open in $D_\pm$.) Let X be the (non-quasi-separated) rigid space obtained by gluing the non-quasi-compact D_0 to $D_\pm$ along $U_\pm$ via the canonical open immersion of $U_\pm$ into both $D_\pm$ and D_0, so X has an admissible cover by $\{D_+, D_-, D_0\}$ and a set-theoretic cover by the pair of disjoint admissible affinoid opens $\{D_+, D_-\}$. However, this latter cover of X is not admissible because it pulls back to a set-theoretic cover of the admissible open $D_0 \subseteq X$ by two disjoint admissible opens and this is not admissible for D_0 since D_0 is connected. The set of affinoid annuli $T_N = \{c^N \leq |t| \leq 1\} \subseteq D_0$ for integers $N \geq 0$ is an admissible cover of D_0 and hence $\{D_+, D_-\} \cup \{T_N\}_{N \geq 0}$ is an admissible affinoid open cover of X, but the preceding argument shows that there is no admissible finite subcover. Hence, X is not quasi-compact.

Still working under the assumption that the admissibility problems for $S^{>h}$ and $S^{\geq h}$ in S are settled in general, we check that if $h \in [p^{-1/8}, 1) \cap \sqrt{|k^\times|}$ then $\{S^{\geq h'}\}_{h < h' \leq 1}$ is an admissible cover of $S^{>h}$ (where we require $h' \in \sqrt{|k^\times|}$). More generally, if $X' \to X$ is a finite surjection of rigid spaces and $\{X_i\}$ is a collection of admissible opens in X for which the maps $X_i \to X$ are quasi-compact and such that the preimage collection $\{X_i'\}$ is an admissible cover of X' then we claim that $\{X_i\}$ is an admissible cover of X. By definition of admissibility in terms of pullbacks to affinoids, we can assume that X is affinoid. In this case the X_i's are quasi-compact opens in X and so the problem is to show that a finite collection of them covers X set-theoretically. This in turn follows from the covering hypothesis for $\{X_i'\}$ in X' and the surjectivity of X' onto X.

Let us next check (under the same admissibility hypotheses as above) that the compatibility with respect to change in the base field is automatic when S is pseudo-separated or quasi-separated. Let k'/k be an analytic extension field, and let $\widetilde{S}' \to S'$ be the extension of scalars on $\widetilde{S} \to S$. The open immersion $S^{\geq h} \to S$ is *quasi-compact*, so the induced map $k'\widehat{\otimes}_k(S^{\geq h}) \to S'$ is also an open immersion as well as quasi-compact. By pullback along the finite surjection $\widetilde{S}' \to S'$ we deduce that the image of $k'\widehat{\otimes}_k(S^{\geq h})$ in S' is precisely the image ${S'}^{\geq h}$ of $(\widetilde{S}')^{\geq h}$ in S'. To check that $k'\widehat{\otimes}_k(S^{>h}) \to S'$ is an open immersion onto the admissible open ${S'}^{>h}$ in S' we simply note that the source has an admissible covering given by the collection $\{k'\widehat{\otimes}_k(S^{\geq h'})\}_{h<h'\leq 1} = \{{S'}^{\geq h'}\}_{h<h'\leq 1}$ that maps isomorphically onto an admissible cover of ${S'}^{>h}$.

Finally, we turn to the problem of proving that the loci $S^{\geq h}$ and $S^{>h}$ (allowing $h = p^{-1/8}$ in the latter case) are admissible opens in S. These loci have preimages under the finite surjection $\widetilde{S} \to S$ that are admissible opens, so it suffices to prove rather generally that if $f : X' \to X$ is a finite surjection between rigid spaces and $U \subseteq X$ is a subset such that $f^{-1}(U)$ is an admissible open in X' then U is an admissible open in X. We refer the reader to Theorem A.2.1 in the Appendix for the proof of a more general result of Gabber along these lines (allowing proper surjections rather than just finite surjections). ■

Here is an analogue of Corollary 3.1.2:

Corollary 3.2.6. *Let $A \to S$ be as in Theorem 3.2.3, and assume that S is quasi-separated and admits a locally finite admissible affinoid cover. All conclusions in Corollary 3.1.2 hold in this case.*

Proof. The proof is essentially identical to the proof of Corollary 3.1.2 because the only role of *fpqc* maps of affinoids in that proof is that they induce surjections on Berkovich spaces. Since finite surjections between affinoids have the same property, the proof of Corollary 3.1.2 carries over to the new setting. ■

Remark 3.2.7. If $A_{/S}$ admits a formal semi-abelian model after local finite surjective base change then (as in Remark 3.1.3) the conclusions of Theorem 3.2.3 and Corollary 3.2.6 apply with $p^{-1/8}$ replaced by $1/p$.

4. Construction of canonical subgroups

The main result in the theory is a "fibral" existence theorem in §4.1, and it rests on a technique of analytic continuation from the ordinary case. This analytic continuation argument requires an intermediate general result in the geometry of affinoid curves that is treated in §4.2. The relativization of the fibral theorem is straightforward (see Theorem 4.3.1), and the relation between the level-n canonical subgroup and the kernel of the n-fold relative Frobenius map modulo $p^{1-\varepsilon}$ also works out nicely; these and other refinements are treated in §4.3.

4.1. Fibral construction.

The fibral case involves not just the construction of level-n canonical subgroups, subject to a universal lower bound < 1 on the Hasse invariant (depending only on p, g, and n), but it also includes behavior with respect to duality and a universal "size description" of the level-n canonical subgroup.

Theorem 4.1.1. *Fix p, g, and $n \geq 1$. There exists $h = h(p,g,n) \in (p^{-1/8}, 1)$ monotonically increasing in n (for fixed p and g) such that for any analytic extension field $k/\mathbf{Q}_p$ and any g-dimensional abelian variety A over k with Hasse invariant $h(A) > h$,*

1. *a level-n canonical subgroup G_n exists in $A[p^n]$,*
2. *$(A[p^n]/G_n)^\vee \subseteq A^\vee[p^n]$ is the level-n canonical subgroup in $A^\vee$.*

Moreover, for any $r_n \in (p^{-1/p^{n-1}(p-1)}, 1)$ we can pick $h(p,g,n)$ so that $G_n = A[p^n]^0_{\leq r_n}$ for all abelian varieties A with $h(A) > h$.

Remark 4.1.2. In the case of a principally polarized abelian variety A with $h(A) > h(p,g,n)$, assertion (2) in the theorem says that G_n is a Lagrangian (i.e., maximal isotropic) subgroup for the induced perfect Weil symplectic form on $A[p^n]$. It is also worth noting at the outset that the proof consists of three essentially different cases: the principally polarized case (which contains all of the content for the relavitization in families), the potentially good reduction case, and finally the general case. It is essential that we have control over the radius r_n in order to push through the proof of the general case (see Step 7 in the proof of Theorem 4.1.1).

The key aspects of Theorem 4.1.1 are two-fold: h only depends on p, g, and n (not on any auxiliary discrete data such as a polarization or level structure), and we do not impose restrictions on the (potentially semi-abelian) reduction-type.

Before we begin the long proof of Theorem 4.1.1, we make several preliminary observations concerning an arbitrary A with $h(A) > p^{-1/8}$. Fix an integer $N \geq 3$ not divisible by p. We may enlarge k so that $A[N]$ is split. By Zahrin's trick, $(A \times A^{\vee})^4$ is principally polarized. By Theorem 2.3.4 we have $h((A \times A^{\vee})^4) = h(A)^8 \in (1/p, 1]$, so in particular $h(A)$ is close to 1 if and only if $h((A \times A^{\vee})^4)$ is close to 1. Setting aside the assertion (2) concerning behavior with respect to Cartier duality, and focusing just on the existence aspect and the universal "size description" for a fixed $r_n \in (p^{-1/p^{n-1}(p-1)}, 1)$, we shall initially restrict our attention to the problem of constructing a uniform $h_{\rm pp}(p, g, n) \in (1/p, 1)$ that works for those abelian varieties A of dimension g (over arbitrary analytic extensions of $\mathbf{Q}_p$) such that A admits a principal polarization; the preceding observations will be used in the deduction of general case from this special case. We give the proof of Theorem 4.1.1 in eight steps (with Steps 3 and 4 containing the key inputs from the theory of Berkovich spaces).

Step 1. In the first five steps we will be working with certain universal families and not with a single abelian variety over a field as in the statement of the theorem, so there will be no risk of confusion caused by the fact that we shall use the notation A in Steps 1–5 to denote a certain "universal" semi-abelian scheme. Fix a positive integer $N \geq 3$ not divisible by p. As we recorded in the proof of Theorem 3.1.1, Chai and Faltings constructed a semi-abelian scheme $A \to X$ over a proper flat $\mathbf{Z}_p$-scheme X in which $\mathscr{A}_{g,1,N/\mathbf{Z}_p}$ equipped with its universal abelian scheme are Zariski-open subschemes. The proper rigid space $X^{\rm an}_{\mathbf{Q}_p}$ over $\mathbf{Q}_p$ contains $\mathscr{A}^{\rm an}_{g,1,N/\mathbf{Q}_p}$ as a Zariski-open subspace, and $A^{\rm an}_{\mathbf{Q}_p}$ is a smooth $X^{\rm an}_{\mathbf{Q}_p}$-group whose restriction over $\mathscr{A}^{\rm an}_{g,1,N/\mathbf{Q}_p}$ is the universal principally-polarized abeloid space of relative dimension g with full level-N structure in the category of rigid spaces over $\mathbf{Q}_p$. Let $\mathfrak{A} \to \mathfrak{X}$ be the p-adic completion of $A \to X$, so by $\mathbf{Z}_p$-properness of X we get a canonical identification $\mathfrak{X}^{\rm rig} = X^{\rm an}_{\mathbf{Q}_p}$ and a canonical isomorphism of $\mathfrak{A}^{\rm rig}$ onto an admissible open subgroup of $A^{\rm an}_{\mathbf{Q}_p}$ [C1, 5.3.1(4)].

We can refine the construction of the loci $(\mathscr{A}^{\rm an}_{g,1,N/\mathbf{Q}_p})^{>h}$ and $(\mathscr{A}^{\rm an}_{g,d,N/\mathbf{Q}_p})^{\geq h}$ as follows. The semi-abelian scheme A over the $\mathbf{Z}_p$-proper X provides a notion of Hasse invariant $h(A^{\rm an}_x)$ for the fiber $A^{\rm an}_x$ over *any* point $x \in X^{\rm an}_{\mathbf{Q}_p} \simeq \mathfrak{X}^{\rm rig}$. Indeed, if R_x denotes the valuation ring of the residue field of $X_{\mathbf{Q}_p}$ at x then the formal semi-abelian scheme arising from A over $\operatorname{Spf} R_x$ provides a g-parameter commutative formal group over R_x/pR_x (via formal completion of $A \bmod pR_x$ along its identity section over $\operatorname{Spec}(R_x/pR_x)$) and $h(A^{\rm an}_x) \in [1/p, 1] \cap \sqrt{|k^{\times}|}$ is defined to be $\max(|\delta_x|, 1/p)$ with $\delta_x \in R_x$ a lift of the "determinant" of the Verschiebung map on the Lie algebra of this formal group over R_x/pR_x. These formal groups relativize over small open formal affines in $\mathfrak{X}$, so for $h \in [1/p, 1] \cap p^{\mathbf{Q}}$ we get admissible loci $(X^{\rm an}_{\mathbf{Q}_p})^{\geq h}$ and $(X^{\rm an}_{\mathbf{Q}_p})^{>h}$ whose formations are compatible with any change in the base field and whose intersections with $\mathscr{A}^{\rm an}_{g,1,N/\mathbf{Q}_p}$ are the respective loci $(\mathscr{A}^{\rm an}_{g,1,N/\mathbf{Q}_p})^{\geq h}$ and $(\mathscr{A}^{\rm an}_{g,1,N/\mathbf{Q}_p})^{>h}$. When using any analytic extension field $k/\mathbf{Q}_p$ as the base field, we get the same conclusions for $h \in [1/p, 1] \cap \sqrt{|k^{\times}|}$.

Working over the discretely-valued base field $\mathbf{Q}_p$, we will show that for *any* $r_n \in (p^{-1/p^{n-1}(p-1)}, 1) \cap p^{\mathbf{Q}}$ there exists $h_0 \in (1/p, 1)$ sufficiently close to 1 such that subgroup $A_x[p^n]^0_{\leq r_n}$ has size p^{ng} for any fiber $A^{\rm an}_x$ of $A^{\rm an}_{\mathbf{Q}_p} \to X^{\rm an}_{\mathbf{Q}_p}$ whose Hasse invariant $h(A^{\rm an}_x)$ strictly exceeds h_0 (so by induction on n, the subgroup $A_x[p^m]^0_{\leq r_n}$ has size p^{mg} for all $1 \leq m < n$ as well, at the expense of possibly slightly increasing h_0 so that it also "works" for $1 \leq m < n$). Granting this for a moment, the same technique as in the case $g = n = 1$ [C4, Thm. 4.1.3] provides a unique finite étale subgroup G_n in $A^{\rm an}_{\mathbf{Q}_p}[p^n]$ over the admissible open domain $(\mathscr{A}^{\rm an}_{g,1,N/\mathbf{Q}_p})^{>h_0}$ such that G_n induces the level-n canonical subgroups on fibers; this is such a crucial step in the construction that we specifically wrote the proof of [C4, Thm. 4.1.3] for $g = n = 1$ so that it is transparent that the method carries over to the case now being considered. (The key input is the finiteness criterion for flat rigid-analytic morphisms in [C4, Thm. A.1.2].) In view of what we are temporarily assuming for h_0 we get a description for G_n in terms of the "size" of fibral points, and this ensures (with the help of $\mathfrak{A}^{\rm rig}$) that applying extension of the base field to G_n provides the desired universal result in the principally polarized case over arbitrary analytic extension fields $k/\mathbf{Q}_p$. (It is trivial to eliminate the restriction $r_n \in p^{\mathbf{Q}}$ by working with $r_n + \theta_n \in p^{\mathbf{Q}}$ for small $|\theta_n|$.)

We now turn to the problem of finding h_0. The method of proof of [C4, Thm. 4.1.3] shows that for any $r \in (0,1) \cap p^{\mathbf{Q}}$, there is a quasi-compact étale subgroup $G_{n,\le r} = A^{\rm an}_{\mathbf{Q}_p}[p^n]_{\le r}$ in the $X^{\rm an}_{\mathbf{Q}_p}$-group $A^{\rm an}_{\mathbf{Q}_p}[p^n]$ such that $G_{n,\le r}$ induces $A_x[p^n]^0_{\le r}$ on fibers, and that the formation of $G_{n,\le r}$ commutes with arbitrary extension on the base field. We want to prove that for each $r_n \in p^{\mathbf{Q}}$ strictly between $p^{-1/p^{n-1}(p-1)}$ and 1 there exists $h_0 \in (1/p, 1)$ such that the fibers of $G_{n,\le r_n}$ over the admissible locus with Hasse invariant $> h_0$ are finite free $\mathbf{Z}/p^n\mathbf{Z}$-modules of rank g. To construct h_0 we shall use étale cohomology on Berkovich spaces.

Step 2. By [Ber2, 1.6.1], for any non-archimedean field K there is an equivalence of categories between the category of paracompact strictly K-analytic Berkovich spaces and the category of quasi-separated rigid-analytic spaces over K that have a locally finite admissible covering by affinoid opens. Moreover, this equivalence is compatible with fiber products and change of the base field. Let $\varphi : \mathscr{A} \to \mathscr{X}$ be the morphism of Berkovich spaces over $\mathbf{Q}_p$ that corresponds to the morphism $A^{\rm an}_{\mathbf{Q}_p} \to X^{\rm an}_{\mathbf{Q}_p}$ under this equivalence, so by compatibility with fiber products it follows that $\mathscr{A}[p^n] \to \mathscr{X}$ is the morphism associated to the quasi-finite étale morphism $A^{\rm an}_{\mathbf{Q}_p}[p^n] \to X^{\rm an}_{\mathbf{Q}_p}$. The universal properties of analytification in the category of classical rigid-analytic spaces [C1, §5.1] and in the category of good Berkovich spaces [Ber2, §2.6] ensure that φ is the Berkovich-analytification of the group scheme $A_{\mathbf{Q}_p} \to X_{\mathbf{Q}_p}$, and likewise for the structural map for the p^n-torsion, so (by [Ber2, 2.6.9, 3.2.10, 3.5.8]) the map of Berkovich spaces $\mathscr{A} \to \mathscr{X}$ is a smooth group and $\mathscr{A}[p^n] \to \mathscr{X}$ is a quasi-finite étale morphism. Since $X_{\mathbf{Q}_p}$ is proper over $\operatorname{Spec} \mathbf{Q}_p$, it follows (via Chow's lemma) that the strictly $\mathbf{Q}_p$-analytic space $\mathscr{X}$ is proper, and in particular $\mathscr{X}$ is compact and Hausdorff.

Let $\mathfrak{A}_0$ denote the formal completion of A along the identity section of its mod-p fiber. By [deJ, 7.2.5] and the arguments in the proof of [C4, Thm. 3.2.5], the morphism of rigid spaces $i : \mathfrak{A}_0^{\rm rig} \to A^{\rm an}_{\mathbf{Q}_p}$ over $X^{\rm an}_{\mathbf{Q}_p}$ is an open subgroup that fiberwise computes the Berthelot generic fiber of the "formal group" in each semi-abelian $A^{\rm an}_x$. This open $X^{\rm an}_{\mathbf{Q}_p}$-subgroup therefore meets $A^{\rm an}_{\mathbf{Q}_p}[p^n]$ in an open $X^{\rm an}_{\mathbf{Q}_p}$-subgroup of $A^{\rm an}_{\mathbf{Q}_p}[p^n]$ whose fiber over each $x \in X^{\rm an}_{\mathbf{Q}_p}$ is $A^{\rm an}_x[p^n]^0$ (see Definition 2.2.3), and these properties persist over any analytic extension of $\mathbf{Q}_p$.

We also claim that the $\mathscr{X}$-group map $\mathscr{A}^0 \to \mathscr{A}$ associated to i is an *open immersion*. This is a special case of a general lemma that we set up as follows. Let Z be a separated scheme of finite type over a complete discrete valuation ring R with fraction field k. Let $\mathfrak{Y}$ be its formal completion along a closed subset Y_0 in the closed fiber Z_0 such that Y_0 is *proper* over the residue field. The canonical morphism

$$i : \mathfrak{Y}^{\rm rig} \to Z_k^{\rm an} \tag{4.1.1}$$

of rigid-analytic spaces is an open immersion and remains so upon arbitrary extension on the base field [C4, Thm. 3.2.5]. Under the equivalence in [Ber2, 1.6.1], (4.1.1) induces a morphism of Berkovich spaces and we have:

Lemma 4.1.3. *The morphism of Berkovich spaces associated to* (4.1.1) *is an open immersion.*

A case of interest in the lemma is when Z is the total space of a semi-abelian scheme over a proper R-scheme S, and $Y_0 = S$ is the identity section. The conclusion in this lemma is generally false if we do not complete along closed subsets that are proper over the residue field.

Proof. The target $Z_k^{\rm an}$ is generally "too big" to permit passage to the affine case via a cartesian-square argument, so we first shall reduce the problem to be within the setting of formal schemes over $\operatorname{Spf} R$ (rather than schemes of finite type over $\operatorname{Spec} R$). By the Nagata compactification theorem [L2], there is a Zariski-open immersion of Z into a proper R-scheme $\overline{Z}$. Since Y_0 is proper over the residue field, it is closed in the closed fiber of $\overline{Z}$. The map of Berkovich spaces associated to $Z_k^{\rm an} \to \overline{Z}_k^{\rm an}$ is the Berkovich analytification of the Zariski-open immersion $Z_k \to \overline{Z}_k$ over $\operatorname{Spec} k$, so it is an open immersion. Hence, we may replace Z with $\overline{Z}$, so we can assume that Z is R-proper. Thus, if $\mathfrak{Z}$ denotes the formal completion of Z along its entire closed fiber then there is a canonical factorization of i as the composite $\mathfrak{Y}^{\rm rig} \stackrel{\alpha}{\to} \mathfrak{Z}^{\rm rig} \stackrel{\beta}{\to} Z_k^{\rm an}$ and by [C1, 5.3.1(4)] the map β is an isomorphism because Z is R-proper. Hence, it is enough to study the Berkovich-space morphism associated to the map α, and this is a special case of the following general considerations (applied to $\mathfrak{Z}$).

Let $\mathfrak{S}$ be a separated formal scheme topologically of finite type over $\mathrm{Spf}(R)$ and let $\mathfrak{S}'$ be its formal completion along a closed set in the ordinary scheme $\mathfrak{S}_{\mathrm{red}}$ over the residue field. Consider the map of separated quasi-compact rigid spaces $j : \mathfrak{S}'^{\mathrm{rig}} \to \mathfrak{S}^{\mathrm{rig}}$ associated to the map of formal schemes $\mathfrak{S}' \to \mathfrak{S}$ via Berthelot's functor. We claim that the map of Berkovich spaces associated to j is an open immersion. It is sufficient to check this condition after pullback to each of a finite collection of (strictly k-analytic) k-affinoid domains that cover $\mathfrak{S}^{\mathrm{rig}}$ (such as the domains associated to the Berthelot-rigidifications of finitely many formal open affines that cover $\mathfrak{S}$). Since Berthelot's functor is compatible with fiber products, and so is Berkovich's functor (from "reasonable" rigid-analytic spaces over k to paracompact strictly k-analytic Berkovich spaces), our problem is thereby reduced to the affine case $\mathfrak{S} = \mathrm{Spf}(\mathscr{B})$ and $\mathfrak{S}' = \mathrm{Spf}(\mathscr{B}')$ where $\mathscr{B}'$ is the completion of $\mathscr{B}$ along some ideal $(f_1, \ldots, f_m)$.

There is a natural isomorphism of topological $\mathscr{B}$-algebras

$$\mathscr{B}' \simeq \mathscr{B}[\![T_1, \ldots, T_m]\!]/(T_j - f_j). \tag{4.1.2}$$

Let $\mathscr{S}$ be the Berkovich space associted to $\mathfrak{S}^{\mathrm{rig}}$, and let Δ be the (Berkovich) open unit disc. By the compatibility of the Berthelot and Berkovich functors with respect to closed immersions (and fiber products), (4.1.2) identifies the Berkovich space $\mathscr{S}'$ associated to $\mathfrak{S}'^{\mathrm{rig}}$ with the Zariski-closed locus in $\mathscr{S} \times \Delta^m$ cut out by the simultaneous conditions $T_j = f_j$ where $T_1, \ldots, T_m$ are the coordinates on the factors Δ of Δ^m. By universal properties, the morphism $\mathscr{S}' \to \mathscr{S}$ is an isomorphism onto the open domain in $\mathscr{S}$ where $|f_j| < 1$ for all j. This completes the proof that the Berkovich-space map associated to (4.1.1) is an open immersion. ■

Step 3. Now we study the smooth and separated group $\mathscr{A} \to \mathscr{X}$ with quasi-finite étale torsion levels $\mathscr{A}[p^n] \to \mathscr{X}$. Over each of finitely many strictly $\mathbf{Q}_p$-analytic affinoid subdomains D_α that cover $\mathscr{X}$ and are sufficiently small, the pullback of $\mathscr{A}^0$ over D_α splits as a product of D_α with a g-dimensional open unit polydisc (with coordinates that measure the "size" of geometric points of $\mathscr{A}^0$ in fibers over D_α in accordance with Definition 2.2.5).

Let $\mathscr{A}[p^n]^0$ denote the open subgroup $\mathscr{A}[p^n] \cap \mathscr{A}^0$ in $\mathscr{A}[p^n]$, so $\mathscr{A}[p^n]^0$ is étale and separated over $\mathscr{X}$ with finite fibers. Since all of our preceding constructions in the classical rigid-analytic category are compatible with arbitrary analytic change of the base field, the fibers of $\mathscr{A}^0$ and $\mathscr{A}[p^n]^0$ in the fiber of $\mathscr{A}$ over *any* point $x \in \mathscr{X}$ have the expected interpretations when the fibral Berkovich group $\mathscr{A}_x$ over the completed residue field at x is identified with a smooth rigid-analytic group containing a quasi-compact open subgroup equipped with a formal semi-abelian model over the valuation ring R_x at x. In particular, each fiber $\mathscr{A}_x$ has a Hasse invariant (in the sense introduced in Step 1 that allows for the possibility that $\mathscr{A}_x$ may not be proper). The space $\mathscr{X}$ is (para)compact and Hausdorff, and it is covered by a (locally-)finite set of strictly analytic domains arising from open affinoids in $X^{\mathrm{an}}_{\mathbf{Q}_p}$, so it follows that for any $h \in (1/p, 1]$ the set $\mathscr{X}^{>h}$ (resp. $\mathscr{X}^{\geq h}$) classifying points whose fibers have Hasse invariant $> h$ (resp. $\geq h$) is an open (resp. closed) set in $\mathscr{X}$, and likewise with $h = 1/p$ when considering $\mathscr{X}^{>h}$. The intersection of $\mathscr{X}^{\geq h}$ with any sufficiently small affinoid subdomain D in $\mathscr{X}$ is an affinoid subdomain of D because this subdomain of D is defined by the condition that a certain analytic function on D has absolute value $\geq h$ (so in particular, if $h \in p^{\mathbf{Q}}$ and D is a sufficiently small strictly $\mathbf{Q}_p$-analytic domain in the strictly $\mathbf{Q}_p$-analytic space $\mathscr{X}$ then $D \cap \mathscr{X}^{\geq h}$ is a strictly $\mathbf{Q}_p$-analytic affinoid subdomain in D).

Since $\mathscr{A}^0$ is an open subgroup in $\mathscr{A}$ it is easy to see that for any $r \in (0,1)$ the locus $\mathscr{A}^0_{<r}$ (resp. $\mathscr{A}^0_{\leq r}$) in $\mathscr{A}^0$ that meets each fiber $\mathscr{A}_x$ of $\mathscr{A} \to \mathscr{X}$ in the set of points of size $< r$ (resp. $\leq r$) in the fibral "formal group" $\mathscr{A}^0_x$ is an open (resp. compact, hence closed) subset in $\mathscr{A}$. It follows that the respective intersections

$$\mathscr{A}[p^n]^0_{<r} = \mathscr{A}[p^n] \cap \mathscr{A}^0_{<r}, \;\; \mathscr{A}[p^n]^0_{\leq r} = \mathscr{A}[p^n] \cap \mathscr{A}^0_{\leq r}$$

are respectively open and closed subsets in the quasi-finite, étale, and separated $\mathscr{X}$-group $\mathscr{A}[p^n]^0$, with $\mathscr{A}[p^n]^0_{<r}$ an open $\mathscr{X}$-subgroup of $\mathscr{A}[p^n]$.

All fibers $\mathscr{A}[p^n]^0_x$ are finite étale with rank $\geq p^{ng}$, and (as in Remark 2.2.10) the rank is exactly p^{ng} if and only if x lies in the closed subset $\mathscr{X}^{\geq 1}$ in $\mathscr{X}$. Let $\varphi_n : \mathscr{A}[p^n]^0 \to \mathscr{X}$ be the quasi-finite, étale, and separated structural morphism. We need to use "smearing out" from its fibers, analogous to the structure theorem for quasi-finite, étale, and separated morphisms in complex-analytic geometry. To keep the picture clear, we shall therefore consider a more general situation. Let $f : \mathscr{Y} \to \mathscr{Z}$ be a quasi-finite, étale, and separated morphism between Berkovich spaces. Such a map is open, since it is étale. Berkovich spaces are locally Hausdorff and locally connected, so for any $z \in \mathscr{Z}$ and sufficiently small connected open $\mathscr{U}$ in $\mathscr{Z}$ around z there is a unique decomposition

$$f^{-1}(\mathscr{U}) = \mathscr{V} \coprod \mathscr{V}' \tag{4.1.3}$$

with $\mathscr{V}$ finite étale over $\mathscr{U}$ and $\mathscr{V}'_z = \emptyset$. The formation of $\mathscr{V}$ is visibly compatible with fiber products over $\mathscr{Z}$ and is functorial (for sufficiently small $\mathscr{U}$). In particular, if $\mathscr{Y}$ has a structure of $\mathscr{Z}$-group then $\mathscr{V}$ is an open and closed $\mathscr{U}$-subgroup in $f^{-1}(\mathscr{U})$ (without requiring further shrinking on $\mathscr{U}$).

We apply the preceding considerations to the map $f = \varphi_n$ to conclude that for all $x \in \mathscr{X}$ and sufficiently small connected open neighborhoods $\mathscr{U}_x$ around x, $\varphi_n^{-1}(\mathscr{U}_x)$ contains a unique open $\mathscr{U}_x$-subgroup that is finite étale over $\mathscr{U}_x$ and has degree equal to the degree of the fiber $\varphi_n^{-1}(x)$ over the completed residue field at x. In particular, if x is in the closed subset $\mathscr{X}^{\geq 1}$ of points for which $\mathscr{A}_x$ has Hasse invariant 1 then $\varphi_n^{-1}(\mathscr{U}_x)$ contains a unique open subgroup $G(x)$ that is finite étale over $\mathscr{U}_x$ with rank p^{ng}. These ranks are constant as we vary such x, though the overlaps $\mathscr{U}_x \cap \mathscr{U}_{x'}$ may be disconnected and hence all we can say is that $G(x)$ and $G(x')$ coincide on the connected components of $\mathscr{U}_x \cap \mathscr{U}_x$ that meet $\mathscr{X}^{\geq 1}$. We want to glue these $G(x)$'s (and then exploit the compactness of $\mathscr{X}^{\geq 1}$) to make an "overconvergent" level-n canonical subgroup G_n, but disconnectedness problems seem to make it impossible to do this "by hand." Moreover, we will *not* directly construct G_n as a level-n canonical subgroup, but rather (in Step 4) we will first build an abstract finite étale group G that "glues" the $G(x)$'s and *a posteriori* we use compactness of $\mathscr{X}$ to adjust the Hasse invariant locus over which we work in order to make this finite étale group become a fibrally level-n canonical subgroup given by a radius r_n that we freely choose *a priori* in the interval $(p^{-1/p^{n-1}(p-1)}, 1)$.

Step 4. We circumvent the difficulties with disconnectedness at the end of Step 3 by using étale cohomology to prove:

Lemma 4.1.4. *There exists an open subset $\mathscr{U} \subseteq \mathscr{X}$ containing $\mathscr{X}^{\geq 1}$ oer which there is an open $\mathscr{U}$-subgroup $G \subseteq \varphi_n^{-1}(\mathscr{U})$ that is finite étale of degree p^{ng} over $\mathscr{U}$. If we discard all* (*necessarily open and closed*) *connected components of $\mathscr{U}$ that do not meet $\mathscr{X}^{\geq 1}$, then G is unique.*

The "overconvergence" provided by $G \to \mathscr{U}$ is to be considered as analogous to the classical extension theorem [Go, II, 3.3.1] concerning sections along closed sets for sheaves of sets on a paracompact topological space. Rather amusingly, this fact from sheaf theory on paracompact spaces is used in the proof of [Ber2, 4.3.5], which in turn is the key technical input in the proof of Lemma 4.1.4.

Proof. The uniqueness aspect is obvious, and for existence we shall use the theory of quasi-constructible étale sheaves [Ber2, §4.4]. We now let k be a non-archimedean field (with non-trivial absolute value, as always), and we shall consider a very general situation for which we will gradually impose additional hypotheses to resemble the setup in the statement of the lemma.

Consider a strictly k-analytic Berkovich space $\mathscr{Y}$ and a quasi-finite, étale, and separated abelian $\mathscr{Y}$-group $\mathscr{G} \to \mathscr{Y}$; the strictness hypothesis ensures (see [Ber2, 4.1.5]) that representable functors are sheaves for the étale site on $\mathscr{Y}$, and it also ensures (by descent theory for coherent sheaves [BG, Thm. 3.1], in the case of étale descent for coherent sheaves of algebras) that the category of étale sheaves of sets on $\mathscr{Y}$ that are locally constant with finite stalks is equivalent to the category of finite étale Berkovich spaces over $\mathscr{Y}$. We assume that the fiber-degrees for $\mathscr{G} \to \mathscr{Y}$ are bounded above, and for each $n \geq 0$ we let $\mathscr{Y}_n$ be the set of $y \in \mathscr{Y}$ such that the fiber $\mathscr{G}_y$ has degree $\leq n$ (and we define $\mathscr{Y}_n = \emptyset$ for $n < 0$). The "smearing out" arguments as in (4.1.3) show that the $\mathscr{Y}_n$'s are a finite increasing family of *closed* sets that exhaust $\mathscr{Y}$. We may consider $\mathscr{G}$ as a sheaf on the étale site for $\mathscr{Y}$, and for $y \in \mathscr{Y}$ the y-stalk of this sheaf is identified with $\mathscr{G}_y$ as a Galois module for the residue field at y. Our first claim is that this sheaf is quasi-constructible by means of the stratification defined by the $\mathscr{Y}_n$'s. That is, the pullback of $\mathscr{G}$ to a sheaf on the étale site of the germ $(\mathscr{Y}, \mathscr{Y}_n - \mathscr{Y}_{n-1})$ is finite locally constant for each $n \geq 0$.

We argue by descending induction on n. If N is the maximal fiber-degree for $\mathscr{G}$ over $\mathscr{Y}$ then over the open stratum $\mathscr{Y} - \mathscr{Y}_{N-1}$ the fiber-degree of $\mathscr{G}$ is constant and hence $\mathscr{G}$ is finite étale over this open stratum.

To induct, suppose that $\mathscr{G}$ has quasi-constructible restriction $\mathscr{G}_n$ on the open $\mathscr{U}_n = \mathscr{Y} - \mathscr{Y}_{n-1}$ for some n, and let $j_n : \mathscr{U}_n \hookrightarrow \mathscr{U}_{n-1}$ denote the canonical inclusion. The pullback of the étale sheaf $\mathscr{G}_{n-1}/j_{n!}(\mathscr{G}_n)$ to the germ $(\mathscr{Y}, \mathscr{Y}_{n-1} - \mathscr{Y}_{n-2})$ is finite locally constant by means of the "smearing out" argument (akin to (4.1.3)) at points in $\mathscr{Y}_{n-1} - \mathscr{Y}_{n-2}$. (To do this calculation most easily, use [Ber2, 4.3.4] to permit replacing $\mathscr{Y}$ with the open subset $\mathscr{U}_{n-1}$ in which $\mathscr{Y}_{n-1} - \mathscr{Y}_{n-2}$ is closed.) Hence, the exact sequence

$$0 \to j_{n!}(\mathscr{G}_n) \to \mathscr{G}_{n-1} \to \mathscr{G}_{n-1}/j_{n!}(\mathscr{G}_n) \to 0$$

on $\mathscr{U}_{n-1}$ implies that $\mathscr{G}_{n-1}$ is quasi-constructible on $\mathscr{U}_{n-1}$ because the outer terms are quasi-constructible (using the inductive hypothesis for $\mathscr{G}_n$) and quasi-constructibility is preserved under extensions (by [Ber2, 4.4.3], whose proof appears to be incorrect – due to an erroneous reduction to constant sheaves with finite cyclic fibers – but which is nonetheless true by another argument). This descending induction shows that $\mathscr{G} = \mathscr{G}_{-1}$ is quasi-constructible on $\mathscr{Y}$ with finite locally constant restriction to each germ $(\mathscr{Y}, \mathscr{Y}_n - \mathscr{Y}_{n-1})$, as desired.

Now we assume that $\mathscr{Y}$ is paracompact and Hausdorff. Let $\nu \geq 0$ be the minimal fiber-degree of $\mathscr{G}$ over $\mathscr{Y}$, so $\mathscr{Y}_\nu - \mathscr{Y}_{\nu-1} = \mathscr{Y}_\nu$ is a closed set and hence the germ $(\mathscr{Y}, \mathscr{Y}_\nu - \mathscr{Y}_{\nu-1})$ is a paracompact germ. We impose the assumption that $\mathscr{G}$ is a $\mathbf{Z}/m\mathbf{Z}$-sheaf for some $m \geq 1$ and that along $\mathscr{Y}_\nu$ the stalks are finite free over $\mathbf{Z}/m\mathbf{Z}$. We shall show that over some open neighborhood of $\mathscr{Y}_\nu$ in $\mathscr{Y}$ there exists a finite étale open subgroup in $\mathscr{G}$ with degree ν, so this will prove the lemma upon taking $k = \mathbf{Q}_p$, $\mathscr{Y} = \mathscr{X}$, and $\mathscr{G} = \mathscr{A}[p^n]^0$ (so $\nu = p^{ng}$ and $\mathscr{Y}_\nu = \mathscr{X}^{\geq 1}$ by Lemma 2.2.4).

The quotient $\mathscr{G}/j_{\nu!}\mathscr{G}_\nu$ is finite locally constant on the germ $(\mathscr{Y}, \mathscr{Y}_\nu)$. Thus, in view of the paracompactness, by [Ber2, 4.4.1] (adapted to abelian sheaves) we may find an open subset $\mathscr{U} \subseteq \mathscr{Y}$ containing $\mathscr{Y}_\nu$ and a finite locally constant m-torsion abelian étale sheaf $\mathscr{F}$ on $\mathscr{U}$ such that on the étale site of the germ $(\mathscr{U}, \mathscr{Y}_\nu) = (\mathscr{Y}, \mathscr{Y}_\nu)$ there is an isomorphism of pullbacks

$$\xi : \mathscr{F}|_{(\mathscr{Y}, \mathscr{Y}_\nu)} \simeq (\mathscr{G}/j_{\nu!}\mathscr{G}_\nu)|_{(\mathscr{Y}, \mathscr{Y}_\nu)}.$$

By shrinking $\mathscr{U}$ we may arrange that the stalks of $\mathscr{F}$ are finite free $\mathbf{Z}/m\mathbf{Z}$-modules. The abelian sheaf $\mathscr{F}$ is represented by some finite étale commutative $\mathscr{U}$-group that we shall also denote by $\mathscr{F}$. By [Ber2, 4.3.5] applied to the pullback of $\mathscr{H}om_{\mathbf{Z}/m\mathbf{Z}}(\mathscr{F}, \mathscr{G})$ on the paracompact germ $(\mathscr{Y}, \mathscr{Y}_\nu)$, we can shrink $\mathscr{U}$ so that there is a map $\psi : \mathscr{F} \to (\mathscr{G}/j_{\nu!}\mathscr{G}_\nu)|_{\mathscr{U}}$ inducing the given isomorphism ξ over the paracompact germ $(\mathscr{Y}, \mathscr{Y}_\nu)$. We need to lift ψ to a map $\widetilde{\psi} : \mathscr{F}|_{\mathscr{U}'} \to \mathscr{G}|_{\mathscr{U}'}$ for some open $\mathscr{U}' \subseteq \mathscr{U}$ containing $\mathscr{Y}_\nu$, as then shrinking $\mathscr{U}'$ some more around $\mathscr{Y}_\nu$ will ensure (by separatedness of the quasi-finite étale $\mathscr{G}$ over $\mathscr{Y}$) that $\widetilde{\psi}$ is injective and corresponds to an open subgroup in $\mathscr{G}|_{\mathscr{U}'}$ that is finite étale of degree ν.

To construct the lifting $\widetilde{\psi}$, it suffices to find an open $\mathscr{U}' \subseteq \mathscr{U}$ containing $\mathscr{Y}_\nu$ such that the connecting map

$$\delta : \mathrm{Hom}_{\mathbf{Z}/m\mathbf{Z}}(\mathscr{F}|_{\mathscr{U}'}, (\mathscr{G}/j_{\nu!}\mathscr{G}_\nu)|_{\mathscr{U}'}) \to \mathrm{Ext}^1_{\mathbf{Z}/m\mathbf{Z}}(\mathscr{U}'; \mathscr{F}, j_{\nu!}\mathscr{G}_\nu) \tag{4.1.4}$$

kills the element corresponding to $\psi|_{\mathscr{U}'}$. Since $\mathscr{F}$ is finite locally free over $\mathbf{Z}/m\mathbf{Z}$, the Ext-group may be identified with the étale cohomology group $\mathrm{H}^1(\mathscr{U}', \mathscr{F}^\vee \otimes_{\mathbf{Z}/m\mathbf{Z}} j_{\nu!}\mathscr{G}_\nu)$, where $\mathscr{F}^\vee$ is the $\mathbf{Z}/m\mathbf{Z}$-linear dual, so by the compatibility of (4.1.4) with respect to shrinking $\mathscr{U}'$ around $\mathscr{Y}_\nu$ it suffices to prove

$$\varinjlim_{\mathscr{U}' \supseteq \mathscr{Y}_\nu} \mathrm{H}^1(\mathscr{U}', \mathscr{F}^\vee \otimes_{\mathbf{Z}/m\mathbf{Z}} j_{\nu!}\mathscr{G}_\nu) = 0.$$

By [Ber2, 4.3.5], this limit is identified with the étale cohomology group

$$\mathrm{H}^1((\mathscr{Y}, \mathscr{Y}_\nu), (\mathscr{F}^\vee \otimes_{\mathbf{Z}/m\mathbf{Z}} j_{\nu!}\mathscr{G}_\nu)|_{(\mathscr{Y}, \mathscr{Y}_\nu)})$$

for the pullback sheaf on the étale site of the paracompact germ $(\mathscr{Y}, \mathscr{Y}_\nu)$. This pullback sheaf has vanishing stalks along the closed subset $\mathscr{Y}_\nu$, so by [Ber2, 4.3.4(*ii*)] it vanishes as a sheaf on the site of the germ $(\mathscr{Y}, \mathscr{Y}_\nu)$. ■

Step 5. We fix a choice of open $\mathscr{U}$ containing $\mathscr{X}^{\geq 1}$ as in Lemma 4.1.4 such that each connected component of $\mathscr{U}$ meets $\mathscr{X}^{\geq 1}$, so over $\mathscr{U}$ there exists a unique open $\mathscr{U}$-subgroup $G \subseteq \mathscr{A}[p^n]^0|_{\mathscr{U}}$ that is finite étale with rank p^{ng}. We have a disjoint-union decomposition of quasi-finite, étale, and separated $\mathscr{U}$-spaces

$$\mathscr{A}[p^n]^0|_{\mathscr{U}} = G \coprod \mathscr{Z}. \tag{4.1.5}$$

All fibers of the $\mathscr{U}$-finite étale G are finite free of rank g as modules over $\mathbf{Z}/p^n\mathbf{Z}$, as this holds along the subset $\mathscr{X}^{\geq 1} \subseteq \mathscr{U}$ that meets all connected components of $\mathscr{U}$. The Hasse invariant is a continuous function $\mathscr{X} \to [1/p, 1]$, and $\mathscr{X}^{\geq 1}$ is the locus with Hasse invariant 1. Hence, by compactness of $\mathscr{X}$ it follows that there exists $h_0 \in (1/p, 1)$ such that $\mathscr{X}^{\geq h_0} \subseteq \mathscr{U}$.

For any $h \in [h_0, 1)$, we write $G^{>h}$ to denote $G|_{\mathscr{X}^{>h}}$. For any $r_n \in (p^{-1/p^{n-1}(p-1)}, 1)$, the open $\mathscr{X}$-subgroup $\mathscr{A}[p^n]^0_{<r_n}$ in $\mathscr{A}[p^n]^0$ meets the finite étale $\mathscr{X}^{>h}$-subgroup $G^{>h}$ in an open subgroup that contains the entire fiber along the compact subset $\mathscr{X}^{\geq 1}$. Hence, by properness of $G^{>h} \to \mathscr{X}^{>h}$ we may find $h_n \in (h_0, 1)$ such that there is an inclusion

$$G^{>h_n} \subseteq \mathscr{A}[p^n]^0_{<r_n}|_{\mathscr{X}^{>h_n}}. \tag{4.1.6}$$

Since $\mathscr{X}^{\geq h'}$ is compact for all $h' \in (h_n, 1)$, it follows that all points in the fibers of $G^{>h_n}$ (viewed in fibers of $\mathscr{A}[p^n]^0$) over $\mathscr{X}^{\geq h'}$ have size $\leq r_n - \varepsilon$ (in the sense of Definition 2.2.5) for any such h', with a small $\varepsilon > 0$ that depends on h' (and on r_n).

We shall now prove that the reverse inclusion to (4.1.6) holds if we take h_n sufficiently close to 1 (depending on r_n). Assume to the contrary, so we get a sequence of points $x_m \in \mathscr{U}$ such that $h(\mathscr{A}_{x_m}) \to 1^-$ and $\mathscr{A}_{x_m}[p^n]^0_{<r_n}$ meets the fiber $\mathscr{Z}_{x_m}$ in some point z_m, with $\mathscr{Z}$ as in (4.1.5). By compactness of $\mathscr{X}$ there is a cofinal map $j: I \to \{1, 2, \dots\}$ from a directed set I to the natural numbers such that the subnet $\{x_{j(i)}\}_{i \in I}$ has a limit $x \in \mathscr{X}^{\geq 1} \subseteq \mathscr{U}$. (We have to use subnets rather than subsequences because $\mathscr{X}$ is generally not first-countable.) Since the closed set $\mathscr{A}[p^n]^0_{\leq r_n}$ restricted over the compact set $\mathscr{X}^{\geq h'} \subseteq \mathscr{U}$ is itself compact for any $h' \in (h_n, 1)$, further passage to a subnet allows us to suppose $\{z_{j(i)}\}$ has a limit z in $\mathscr{A}_x[p^n]^0$, and by (4.1.5) we must have $z \in \mathscr{Z}$ since $\mathscr{Z}$ is open and closed in $\mathscr{A}[p^n]^0|_{\mathscr{U}}$. We have $\mathscr{A}_x[p^n]^0 = G_x$ because $h(\mathscr{A}_x) = 1$, so $\mathscr{Z}_x = \emptyset$. Since $z \in \mathscr{Z}_x$, this is a contradiction and so completes our treatment in the case of principally polarized abelian varieties (with a fixed dimension g). We let $h_{\mathrm{pp}}(p, g, n)$ be the universal lower bound on Hasse invariants that was constructed in this argument, and we may arrange that it is monotonically increasing in n (for fixed p and g).

Step 6. For the proof of (1) in the theorem, along with the universal "size description," it remains to infer the general case from what we have just proved in the principally polarized case. We fix p, g, and n as at the outset, as well as $r_n \in (p^{-1/p^{n-1}(p-1)}, 1)$, and we consider an abelian variety A of dimension g over an analytic extension field $k/\mathbf{Q}_p$. We now handle the case when A has potentially good reduction; the general case will be deduced from this case in Step 7. Let us postpone the potentially good reduction hypothesis for a short time. The abelian variety A admits a polarization over k. The abelian variety $(A \times A^\vee)^4$ is therefore principally polarized with dimension $8g$, and (using Theorem 2.3.4) it has Hasse invariant $h(A)^8$ provided that $h(A) > p^{-1/8}$. Thus, by taking

$$h(A) > h_{\mathrm{pp}}(p, 8g, n)^{1/8} > p^{-1/8}$$

we ensure that $(A \times A^\vee)^4$ admits a level-n canonical subgroup that is "p^n-torsion with size $\leq r_n$," so the level-n canonical subgroup in $(A \times A^\vee)^4$ is $(G_n \times G'_n)^4$ for the subgroups $G_n = A[p^n]^0_{\leq r_n}$ and $G'_n \subseteq A^\vee[p^n]^0_{\leq r_n}$ whose geometric fibers must be finite free $\mathbf{Z}/p^n\mathbf{Z}$-modules with ranks adding up to $2g$. This shows that $A \times A^\vee$ has a level-n canonical subgroup, namely $G_n \times G'_n$, and we have to prove that for a suitable universal constant $h(p, g, n) \in [h_{\mathrm{pp}}(p, 8g, n)^{1/8}, 1)$ that is independent of k the factors G_n and G'_n in $A[p^n]$ and $A^\vee[p^n]$ each have order p^{ng} if we take $h(A) > h(p, g, n)$. We also have to prove that these factors annihilate each other with respect to the Weil-pairing between $A[p^n]$ and $A^\vee[p^n]$ by taking $h(p, g, n)$ sufficiently near 1. The following argument shows that $h_{\mathrm{good}}(p, g, n) = h_{\mathrm{pp}}(p, 8g, n)^{1/8}$ works as such an $h(p, g, n)$ when we restrict our attention to those abelian varieties A with potentially good reduction.

Pick a polarization on A, say with degree d^2. Choose $N \geq 3$ relatively prime to pd and increase k so that $A[N]$ is k-split (so $A^\vee[N]$ is also k-split, as $(d, N) = 1$). This data gives rise to a k-point on the smooth moduli scheme $\mathscr{A}_{g,d,N/k}$ that is separated and of finite type over k. Since $(d, N) = 1$, the relativization of

Zahrin's trick provides a morphism

$$\zeta_d : \mathscr{A}_{g,d,N/k} \to \mathscr{A}_{8g,1,N/k}$$

of k-schemes such that the functorial effect of ζ_d on underlying abelian schemes (ignoring the polarization and level structure) is $A \rightsquigarrow (A \times A^\vee)^4$. For any $h \in (1/p, 1] \cap \sqrt{|k^\times|}$ the $\zeta_d^{\rm an}$-preimage of the locus with Hasse invariant h on $\mathscr{A}^{\rm an}_{8g,1,N/k}$ is the locus with Hasse invariant $h^{1/8}$ on $\mathscr{A}^{\rm an}_{g,d,N/k}$. This method shows that for $h \in (p^{-1/8}, 1] \cap \sqrt{|k^\times|}$, the locus of points on $\mathscr{A}^{\rm an}_{g,d,N/k}$ with Hasse invariant $> h$ is admissible open, even in cases with $p|d$.

In view of the relative construction of canonical subgroups over loci with suitable Hasse invariant in the analytified moduli spaces for principally polarized abelian schemes in Steps 1–5, pullback along $\zeta_d^{\rm an}$ provides a closed finite étale subgroup $\mathscr{H}_{n,d}$ inside of the finite étale p^n-torsion on the 4-fold product of the universal polarized abeloid space and its dual over $M_{n,d/k} \stackrel{\rm def}{=} (\mathscr{A}^{\rm an}_{g,d,N/k})^{>h_{\rm good}(p,g,n)}$ such that on fibers it is a level-n canonical subgroup. As in the case of schemes over a base scheme, any rigid-analytic map between finite étale spaces over a rigid space factors uniquely through a finite étale surjection and any two finite étale closed subspaces of a finite étale space coincide globally if they coincide in a single fiber over each connected component of the base. Thus, by using projection to factors and the preceding fibral analysis we see that $\mathscr{H}_{n,d} = (\mathscr{G}_{n,d} \times \mathscr{G}'_{n,d})^4$ for unique finite étale closed $M_{n,d/k}$-subgroups $\mathscr{G}_{n,d}$ and $\mathscr{G}'_{n,d}$ in the repsective p^n-torsion of the universal polarized abeloid and its dual over $M_{n,d/k}$; these are étale-locally finite free $\mathbf{Z}/p^n\mathbf{Z}$-module sheaves. Obviously the formation of $M_{n,d/k}$, $\mathscr{G}_{n,d}$, and $\mathscr{G}'_{n,d}$ are compatible with change in the base field. For example, these all arise from the analogous constructions over $\mathbf{Q}_p$.

Over each connected component of $M_{n,d/k}$ the $\mathbf{Z}/p^n\mathbf{Z}$-ranks of $\mathscr{G}_{n,d}$ and $\mathscr{G}'_{n,d}$ are constant and add up to $2g$, and the relative Weil pairing between them vanishes if it does so on a single fiber. If a connected component of $M_{n,d/k}$ contains an ordinary point ξ then over that connected component the orders of $\mathscr{G}_{n,d}$ and $\mathscr{G}'_{n,d}$ are equal to p^{ng} (by checking on the ξ-fiber). Moreover, we claim that over the connected component of an ordinary point ξ in $M_{n,d/k}$ the groups $\mathscr{G}_{n,d}$ and $\mathscr{G}'_{n,d}$ must be orthogonal (and hence be exact annihilators) under the Weil pairing on p^n-torsion. By passing to the fiber at ξ and making a finite extension $k'/k(\xi)$ of the base field as in Theorem 2.1.9, the problem comes down to the vanishing of the Weil pairing between the multiplicative identity components of the p-divisible groups of the formal semi-abelian models (with ordinary reduction) for the abelian variety and dual abelian variety at ξ. More generally, we have:

Lemma 4.1.5. *Let A be an abelian variety over k having semistable reduction and formal semi-abelian model $\mathfrak{A}_R$ over* $\mathrm{Spf}(R)$ *with ordinary abelian part modulo $\mathfrak{m}_R$. Let $\mathfrak{A}'_R$ be the corresponding formal semi-abelian model for $A^\vee$, so it too has ordinary abelian part modulo $\mathfrak{m}_R$.*

The Weil pairing between $A[p^\infty]$ and $A^\vee[p^\infty]$ makes $\mathfrak{A}[p^\infty]^0_k$ and $\mathfrak{A}'[p^\infty]^0_k$ orthogonal to each other.

Proof. If A_1 and A_2 are two abelian varieties that satisfy the hypotheses in the lemma, then to prove the lemma for each A_i it is equivalent to prove the lemma for $A_1 \times A_2$. Hence, by passing to $(A \times A^\vee)^4$, we may assume that A admits a principal polarization and (after a harmless finite extension of the base field) both $A[N]$ and $A^\vee[N]$ are constant (with $N \geq 3$ a fixed integer not divisible by p). Let $g = \dim A$, so by the work of Faltings and Chai (as in Step 1 above) we may use the valuation criterion for properness to extend A and $A^\vee$ to semi-abelian schemes A_R and A'_R over $\mathrm{Spec}\, R$. By Example 2.1.10 the respective completions $\widehat{A}_R$ and $\widehat{A}'_R$ of A_R and A'_R along an ideal of definition of R are the formal semi-abelian models $\mathfrak{A}_R$ and $\mathfrak{A}'_R$ as in Theorem 2.1.9, so we have unique isomorphisms $\mathfrak{A}_R[p^\infty] \simeq A_R[p^\infty]$ and $\mathfrak{A}'_R[p^\infty] \simeq A'_R[p^\infty]$ respecting the identifications of the k-fibers inside of $A[p^\infty]$ and $A^\vee[p^\infty]$ respectively. Our problem is therefore to prove that the Weil pairing between $A[p^\infty]$ and $A^\vee[p^\infty]$ makes $A_R[p^\infty]^0_k$ orthogonal to $A'_R[p^\infty]^0_k$.

Since R is a henselian valuation ring, it is a directed union of henselian local noetherian subrings D. By standard direct limit arguments, we can descend A_R and A'_R to semi-abelian schemes A_D and A'_D over some such D. Likewise, the identity components $A_D[p^\infty]^0$ and $A'_D[p^\infty]^0$ descend the multiplicative p-divisible groups $A_R[p^\infty]^0$ and $A'_R[p^\infty]^0$, so these descended p-divisible groups over D are also multiplicative. If we let $F \subseteq k$ be the fraction field of D then the Weil pairing between the F-fibers $A_F[p^\infty]$ and $A'_F[p^\infty] = A^\vee_F[p^\infty]$ descends the Weil pairing between $A[p^\infty]$ and $A^\vee[p^\infty]$, so it suffices to prove that this pairing over F makes the F-fibers $A_D[p^\infty]^0_F$ and $A'_D[p^\infty]^0_F$ orthogonal.

Rather generally, if Γ and Γ' are any two multiplicative p-divisible groups over a local *noetherian* domain D with residue characteristic p and fraction field F then any $\mathbf{G}_m[p^\infty]$-valued bilinear pairing between the F-fibers must be zero. In the irrelevant case $\mathrm{char}(F) = p$ this is obvious for topological reasons. In case of generic characteristic 0 we use local injective base change to assume that D is a strictly henselian discrete valuation ring, so Γ and Γ' are powers of $\mathbf{G}_m[p^\infty]$ and the p-adic cyclotomic character of F is non-trivial (it has infinite order). The vanishing is therefore also obvious in characteristic 0. ∎

To settle the case of potentially good reduction with the strict lower bound $h_{\mathrm{good}}(p,g,n) \in (p^{-1/8},1) \cap p^{\mathbf{Q}}$ on the Hasse invariant, it remains to show that for *every* $d \geq 1$ and $k/\mathbf{Q}_p$ there is an ordinary point on each connected component Y in $M_{n,d/k}$ for which Y has a point with potentially good reduction. That is, we claim that such a component has non-empty locus with Hasse invariant equal to 1. We will reduce our problem to the case $k = \mathbf{Q}_p$. (What really matters is that we reduce to the case of a discretely-valued field.) The trick is to exploit finiteness properties in the theory of connectivity for rigid spaces; the following argument uses completed algebraic closures but it could be rewritten to work with only finite extensions. If k'/k is a finite extension then each connected component of $Y \otimes_k k'$ is finite flat over Y and so surjects onto Y. Thus, our problem is unaffected by passage to a finite extension on the base field. (By Theorem 3.1.1, or Theorem 3.2.3, the formation of the locus with Hasse invariant 1 in $M_{n,d/k}$ is compatible with change of the base field.) In particular, by [C1, Cor. 3.2.3] we may suppose that our connected component Y is geometrically connected. Hence, again using the compatibility with change of the base field in Theorem 3.1.1, we may assume that k is algebraically closed and so k contains $\mathbf{C}_p$. Since connected rigid spaces over $\mathbf{C}_p$ are geometrically connected, we may assume $k = \mathbf{C}_p$. We do not know if $M_{n,d/\mathbf{Q}_p}$ has a finite number of connected components. However, by [C1, Cor. 3.2.3] for each connected component Z of $M_{n,d/\mathbf{Q}_p}$ there exists a finite extension $k/\mathbf{Q}_p$ such that all connected components of $Z \otimes_{\mathbf{Q}_p} k$ are geometrically connected. It follows that each connected component Y of $M_{n,d/\mathbf{C}_p}$ arises as a base change of a connected component of $Z_{/k_0}$ for a suitable Z and finite extension $k_0/\mathbf{Q}_p$ (perhaps depending on Y). This completes the reduction of our problem to the case $k = \mathbf{Q}_p$.

Letting $\mathfrak{A}_d = \mathscr{A}^{\wedge}_{g,d,N/\mathbf{Z}_p}$ be the p-adic completion of the separated finite type moduli scheme $\mathscr{A}_{g,d,N/\mathbf{Z}_p}$ over $\mathbf{Z}_p$, there is a canonical quasi-compact open immersion $\mathfrak{A}_d^{\mathrm{rig}} \hookrightarrow \mathscr{A}^{\mathrm{an}}_{g,d,N/\mathbf{Q}_p}$ whose image consists of precisely the points with potentially good reduction. (The formation of this map commutes with any change of the base field, as does the description of its image.) It is therefore enough to prove that for any $h \in (p^{-1/8},1) \cap p^{\mathbf{Q}}$ (such as $h_{\mathrm{good}}(p,g,n)$) every connected component of $(\mathfrak{A}_d^{\mathrm{rig}})^{>h}$ contains an ordinary point, as any connected component of $(\mathscr{A}^{\mathrm{an}}_{g,d,N/\mathbf{Q}_p})^{>h}$ meets the admissible open $(\mathfrak{A}_d^{\mathrm{rig}})^{>h}$ in a (possibly empty) union of connected components of $(\mathfrak{A}_d^{\mathrm{rig}})^{>h}$. The existence of such ordinary points is proved in Theorem 4.2.1 below.

We have settled the case of potentially good reduction. For the initial fixed choice $r_n \in (p^{-1/p^{n-1}(p-1)},1)$ we constructed $h_{\mathrm{good}}(g,p,n)$ so that any g-dimensional A with potentially good reduction and Hasse invariant $h(A) > h_{\mathrm{good}}(p,g,n)$ admits a level-n canonical subgroup G_n given by the set of p^n-torsion points with size $\leq r_n$, and also so that $(A[p^n]/G_n)^\vee \subseteq A^\vee[p^n]$ is the level-n canonical subgroup of the g-dimensional abelian variety $A^\vee$ with potentially good reduction and Hasse invariant $h(A^\vee) = h(A) > h_{\mathrm{good}}(p,g,n)$.

Step 7. Now we consider the same initial setup as in Step 6 except that we allow for the possibility that (after a harmless finite extension of the base field) A has semi-stable reduction with non-trivial toric part. We define

$$h(p,g,n) = \max_{1 \leq g' \leq g} h_{\mathrm{good}}(p,g',n) \in (p^{-1/8},1),$$

and we assume $h(A) > h(p,g,n)$. By Theorem 2.1.9 there exists a (projective) abelian scheme B_R of some non-negative relative dimension $g' \leq g$ over R and a short exact sequence of connected p-divisible groups

$$0 \to \mathfrak{T}[p^\infty] \to \mathfrak{A}_R[p^\infty]^0 \to B_R[p^\infty]^0 \to 0$$

over R with $\mathfrak{T}$ a formal torus and $\mathfrak{A}_R$ a formal semi-abelian scheme that is a "formal model" for A (or rather, for A^{an}). By Example 2.3.2, the Hasse invariant of A is equal to that of the abelian variety B that is the generic fiber of B_R. Hence, if $g' > 0$ then $h(B) > h_{\mathrm{good}}(g',p,n)$ and if $g' = 0$ then A is ordinary (and $h(B) = h(A) = 1$), so the subgroup $B[p^n]^0_{\leq r_n}$ in $B[p^n]^0 = B_R[p^n]^0_k$ is a level-n canonical subgroup of B with the arbitrary but fixed choice of $r_n \in (p^{-1/p^{n-1}(p-1)}, 1)$ that has been used throughout the preceding steps. Since $r_n > p^{-1/p^{n-1}(p-1)}$ we have

$$\mathfrak{T}[p^n]_k \subseteq A[p^n]^0_{\leq p^{-1/p^{n-1}(p-1)}} \subseteq A[p^n]^0_{\leq r_n},$$

so there is an evident left-exact sequence

$$0 \to \mathfrak{T}[p^n]_k \to A[p^n]^0_{\leq r_n} \to B[p^n]^0_{\leq r_n} \tag{4.1.7}$$

and the geometric fibers of $\mathfrak{T}[p^n]_k$ and $B[p^n]^0_{\leq r_n}$ are free with respective ranks $g - g'$ and g' as $\mathbf{Z}/p^n\mathbf{Z}$-modules (even if $g' = 0$). Thus, $A[p^n]^0_{\leq r_n}$ has order $\leq p^{ng}$ and if equality holds then (4.1.7) is short exact with middle term that is $\mathbf{Z}/p^n\mathbf{Z}$-free of rank g, so $A[p^n]^0_{\leq r_n}$ is a level-n canonical subgroup in such cases. The same argument (with the same r_n!) applies to $A^\vee[p^n]^0_{\leq r_n}$, so in particular this group has order $\leq p^{ng}$.

Since $(A \times A^\vee)^4[p^n]^0_{\leq r_n} = (A[p^n]^0_{\leq r_n} \times A^\vee[p^n]^0_{\leq r_n})^4$, the upper bounds on the orders of the factors reduces us to proving that this group has size p^{8ng}. The abelian variety $(A \times A^\vee)^4$ has Hasse invariant $> h_{\rm good}(p, g, n)^8 = h_{\rm pp}(p, 8g, n)$ and it is principally polarized, so its subgroup of p^n-torsion points with size $\leq r_n$ is a level-n canonical subgroup and hence there are exactly p^{8ng} such points as required. This completes the proof that $A[p^n]^0_{\leq r_n}$ is a level-n canonical subgroup whenever $h(A) > h(p, g, n)$ (where $h(p, g, n) \in (p^{-1/8}, 1)$ may be taken to depend on the arbitrary but fixed choice of $r_n \in (p^{-1/p^{n-1}(p-1)}, 1)$). (The reader should reflect on how essential it is in several parts of the argument that we made the construction of level-n canonical subgroups in various dimensions all be related to the same size bound r_n.)

Step 8. Continuing with notation as in Step 7, the verification of part (2) in Theorem 4.1.1 will now be given in general; in Step 6 it was verified in the case of potentially good reduction. We must check that the Weil pairing between $A[p^n]^0_{\leq r_n}$ and $A^\vee[p^n]^0_{\leq r_n}$ vanishes. The respective generic fibers $\mathfrak{A}_R[p^\infty]_k$, $\mathfrak{A}_R[p^\infty]^0_k$, and $\mathfrak{T}[p^\infty]_k$ will be called the *finite part*, *local part*, and *toric part* of the p-divisible group $A[p^\infty]$ over k, and these generic fibers will be respectively denoted $A[p^\infty]_{\rm f}$, $A[p^\infty]^0$, and $A[p^\infty]_{\rm t}$. Although these definitions depend on A and not just on $A[p^\infty]$ (*e.g.*, k may be algebraically closed), for our purposes such dependence is not a problem; the p^n-torsion of $A[p^\infty]^0$ recovers Definition 2.2.3, so there is no inconsistency in the notation. Also, keep in mind that $A[p^\infty]_{\rm t}$ may be smaller than the generic fiber of the maximal multiplicative p-divisible subgroup of $\mathfrak{A}_R[p^\infty]^0$. Analogous notations are used for $A^\vee$, and we let B'_R denote the abelian scheme associated to $A^\vee$, so B'_R is canonically isomorphic to $B_R^\vee$ by Theorem 2.1.9.

The respective quotients $A[p^\infty]_{\rm f}/A[p^\infty]_{\rm t}$ and $A[p^\infty]^0/A[p^\infty]_{\rm t}$ are canonically identified with $B[p^\infty]$ and $B[p^\infty]^0 \stackrel{\rm def}{=} B_R[p^\infty]^0_k$, and similarly with $A^\vee$ and $B'_R \simeq B_R^\vee$ (even if B_R and B'_R vanish). Since the settled case of good reduction in Step 6 ensures that the Weil pairing between $B[p^n]^0_{\leq r_n}$ and $B^\vee[p^n]^0_{\leq r_n}$ vanishes, to infer the vanishing of the Weil pairing between $A[p^n]^0_{\leq r_n}$ and $A^\vee[p^n]^0_{\leq r_n}$ (and so to finish the proof of Theorem 4.1.1, conditional on Theorem 4.2.1 below that was used above in Step 6) it suffices to use (4.1.7) and its $A^\vee$-analogue along with the following general theorem that gives an analogue of the trivial Lemma 4.1.5 in the case of possibly non-ordinary or bad reduction (and characterizes the isomorphism $B' \simeq B^\vee$ in terms of two pieces of data: the unique formal semi-abelian models for A and $A^\vee$, and the Weil pairings between torsion on A and $A^\vee$).

Theorem 4.1.6. *Under the Weil pairing $A[p^\infty] \times A^\vee[p^\infty] \to \mu_{p^\infty}$ over k, the toric part on each side annihilates the finite part on the other side, and the induced pairing between $A[p^\infty]_{\rm f}/A[p^\infty]_{\rm t} \simeq B[p^\infty]_k$ and $A^\vee[p^\infty]_{\rm f}/A^\vee[p^\infty]_{\rm t} \simeq B'[p^\infty]_k$ is the restriction of the Weil pairing for the abelian variety B over k via the canonical isomorphism $B'_R \simeq B_R^\vee$.*

Proof. See Theorem A.3.1 in the Appendix, where a more general compatibility is proved for N-torsion pairings for any positive integer N by closely studying the proof of Theorem 2.1.9. ■

Remark 4.1.7. The method of proof of Lemma 4.1.5 can be used to give a proof of the orthogonality aspect of Theorem 4.1.6 by reduction to the discretely-valued case that is precisely the semi-stable case of Grothendieck's orthogonality theorem [SGA7, IX, Thm. 5.2]. However, it is the compatibility with Weil pairings on the abelian parts that is more important for us, and to prove this compatibility it seems to be unavoidable to have to study the proof of Theorem 2.1.9 where the natural isomorphism between B'_R and $B_R^\vee$ is constructed.

4.2. **A connectedness result.** This section is devoted to proving the following theorem that was used in Step 6 in the proof of Theorem 4.1.1.

Theorem 4.2.1. *Choose $g, d \geq 1$ and $N \geq 3$ with $p \nmid N$. Let $M = \mathscr{A}_{g,d,N/\mathbf{Z}_p}$ and let $\widehat{M}$ denote its p-adic completion equipped with its universal p-adically formal polarized abelian scheme. For any $h \in [1/p, 1) \cap p^{\mathbf{Q}}$, let $(\widehat{M}^{\mathrm{rig}})^{>h}$ denote the locus of fibers with Hasse invariant $> h$ for the Raynaud generic fiber of the universal p-adically formal polarized abelian scheme over $\widehat{M}$, and define $(\widehat{M}^{\mathrm{rig}})^{\geq h}$ similarly for $h \in (1/p, 1] \cap p^{\mathbf{Q}}$.*

Each connected component of $(\widehat{M}^{\mathrm{rig}})^{>h}$ and of $(\widehat{M}^{\mathrm{rig}})^{\geq h}$ meets the ordinary locus (i.e., it meets $(\widehat{M}^{\mathrm{rig}})^{\geq 1}$).

We can allow $1/p$ rather than $p^{-1/8}$ in Theorem 4.2.1 because of Remark 3.1.3.

Proof. Let x be a point in $(\widehat{M}^{\mathrm{rig}})^{>h}$ (resp. $(\widehat{M}^{\mathrm{rig}})^{\geq h}$), with $K(x)/\mathbf{Q}_p$ the residue field at x and R_x its valuation ring. Let A_x be the fiber at x for the universal abeloid space over $\widehat{M}^{\mathrm{rig}}$, so A_x is an abelian variety over $K(x)$ with good reduction, and likewise for its dual $A_x^\vee$. We uniquely extend x to an R_x-point of M, and we let $\overline{x}$ be the induced rational point in the closed fiber of $M_{/R_x}$.

Norman and Oort [NO, Thm. 3.1] proved that the ordinary locus is Zariski-dense in every fiber of $\mathscr{A}_{g,d,N}$ over non-generic points of $\operatorname{Spec} \mathbf{Z}[1/N]$, with all fibers of pure dimension $g(g+1)/2$. Mumford proved that the equicharacteristic deformation ring at any rational point on a geometric fiber of $\mathscr{A}_{g,d,N}$ in positive characteristic over $\mathbf{Z}[1/N]$ is a g^2-variable power series rings modulo $g(g-1)/2$ relations [O, 2.3.3], so it follows from the equality $g^2 - g(g-1)/2 = g(g+1)/2$ and a standard result in commutative algebra [Mat, 17.4] that $\mathscr{A}_{g,d,N/\mathbf{Z}[1/N]}$ is a relative complete intersection over $\mathbf{Z}[1/N]$ (and in particular it is flat). Thus, by a slicing argument on an affine open neighborhood of $\overline{x}$ in the relative complete intersection $M_{/R_x}$, we may construct an R_x-flat locally closed affine subscheme Z in $M_{/R_x}$ with relative dimension 1 such that

- the closed fiber $\overline{Z}$ passes through $\overline{x}$ and has all generic points in the ordinary locus,
- the generic fiber $Z_{/K(x)}$ is smooth.

Let $\mathfrak{Z}$ be the R_x-flat formal completion of $M_{/R_x}$ along $\overline{Z}$, so $\mathfrak{Z}^{\mathrm{rig}}$ is a quasi-compact admissible open in $\widehat{M}^{\mathrm{rig}}$ and x lies in $\mathfrak{Z}^{\mathrm{rig}}$ since $\overline{x} \in \overline{Z}$.

By the construction of the Hasse invariant, we may replace Z with a suitable open affine around $\overline{x}$ (to trivialize the locally-free module underlying a formal Lie algebra) so that the universal formal abelian scheme over $\mathfrak{Z} \bmod pR_x$ admits a Hasse invariant as a function on $\operatorname{Spec} R_x/pR_x$ (rather than merely as a section of a line bundle). Let H be a function on $\mathfrak{Z}$ that lifts this Hasse invariant. If H^{rig} denotes the associated rigid-analytic function on $\mathfrak{Z}^{\mathrm{rig}}$ then $\max(|H^{\mathrm{rig}}|, 1/p)$ defines the Hasse invariant over the admissible open $\mathfrak{Z}^{\mathrm{rig}}$. The coordinate ring $\mathscr{O}(\mathfrak{Z})$ of the affine formal scheme $\mathfrak{Z}$ is excellent and reduced (as $\mathscr{A}^{\mathrm{an}}_{g,N,d/\mathbf{Q}_p}$ is smooth), so the normalization of $\mathscr{O}(\mathfrak{Z})$ is an R_x-flat finite extension ring of $\mathscr{O}(\mathfrak{Z})$ whose associated formal scheme $\widetilde{\mathfrak{Z}}$ is $\mathfrak{Z}$-finite with Raynaud generic fiber $\mathfrak{Z}^{\mathrm{rig}}$ because $\mathfrak{Z}^{\mathrm{rig}}$ is its own normalization (as it is even smooth). Also, the "generic ordinarity" of the locus $\overline{Z}$ in the moduli space ensures that on the *pure one-dimensional* reduction $\mathfrak{Z} \bmod \mathfrak{m}_{R_x}$ (with underlying space $\overline{Z}$) the reduction of H is a unit at the generic points. The same must therefore hold for H on the mod-$\mathfrak{m}_{R_x}$ fiber of the $\mathfrak{Z}$-finite formal normalization covering $\widetilde{\mathfrak{Z}}$, as $\widetilde{\mathfrak{Z}}$ has no isolated points (and so its irreducible components are all finite over those of the 1-dimensional $\mathfrak{Z}$). By [deJ, 7.4.1], $\mathscr{O}(\widetilde{\mathfrak{Z}})$ is the ring of power-bounded functions on the $K(x)$-affinoid $\widetilde{\mathfrak{Z}}^{\mathrm{rig}} = \mathfrak{Z}^{\mathrm{rig}}$. Hence, the ideal of topological nilpotents in $\mathscr{O}(\mathfrak{Z}^{\mathrm{rig}})$ is the radical of $\mathfrak{m}_{R_x}\mathscr{O}(\widetilde{\mathfrak{Z}})$. (The intervention of the radical is necessary because sup-norms for elements of the $K(x)$-affinoid $\mathfrak{Z}^{\mathrm{rig}}$ merely lie in $\sqrt{|K(x)^\times|}$ and not necessarily in $|K(x)^\times|$.) Thus, we are reduced to the following theorem in 1-dimensional affinoid geometry (applied to $\mathscr{O}(\mathfrak{Z}^{\mathrm{rig}})$ over $K(x)$). ■

Theorem 4.2.2. *Let k be a non-archimedean field and let A be a k-affinoid algebra such that* $\operatorname{Sp} A$ *has pure dimension* 1. *Let $A^0 \subseteq A$ be the subring of power-bounded functions, and let $\widetilde{A}$ be its analytic reduction; i.e., the quotient of A^0 modulo topological nilpotents.*

Let $a \in A^0$ be an element whose image $\widetilde{a}$ in the reduced algebra $\widetilde{A}$ is non-vanishing at the generic points of $\operatorname{Spec}\widetilde{A}$*; in particular, $\|a\|_{\sup} = 1$. For any $r \in \sqrt{|k^\times|}$ with $r \leq 1$, every connected component of*

$$(\mathrm{Sp}(A))^{\geq r} = \{x \in \mathrm{Sp}(A) \,|\, |a(x)| \geq r\} \tag{4.2.1}$$

contains a point x such that $|a(x)| = 1$.

This theorem can be proved by using the geometry of formal semi-stable models to track the behavior of $|a|$, following some techniques of Bosch and Lütkebohmert as in [BL1, §2] (after reducing to the case of algebraically closed k with the help of [C1, §3.2]). However, A. Thuillier showed me a method of proof that uses only elementary properties of affinoid Berkovich spaces, entirely bypassing the more sophisticated apparatus of formal semistable models, so we present Thuillier's proof.

Proof. It is equivalent to work with the associated strictly k-analytic Berkovich spaces, so we let $X = \mathscr{M}(A)$ and $X^{\geq r} = \mathscr{M}(A^{\geq r})$ with $\mathrm{Sp}(A^{\geq r})$ equal to the affinoid subdomain $(\mathrm{Sp}(A))^{\geq r}$ in $\mathrm{Sp}(A)$; clearly $X^{\geq r} \subseteq X$ is the locus of points $x \in X$ for which $|a(x)| \geq r$. The Shilov boundary $\Gamma(X) \subseteq X$ is the finite set of preimages of the generic points of the analytic reduction $\mathrm{Spec}(\widetilde{A})$ under the reduction map $X \to \mathrm{Spec}(\widetilde{A})$ [Ber1, 2.4.4]. The hypotheses therefore imply that $|a(x)| = 1$ for each $x \in \Gamma(X)$, so it is equivalent to prove that every connected component C of $X^{\geq r}$ meets $\Gamma(X)$ (since the "classical" points are dense in any strictly k-analytic Berkovich space, such as $C \cap X^{\geq 1}$ for each such C). Hence, we pick a component C disjoint from $\Gamma(X)$ and seek a contradiction. The complement $U = X - (X^{\geq r} - C)$ is an open set in X that contains C, so since $C \cap \Gamma(X) = \emptyset$ and $\Gamma(X) \subseteq X^{\geq r}$ we have $U \cap \Gamma(X) = \emptyset$ and $|a| < r$ on $U - C = (X - X^{\geq r}) \cap (X - C)$.

The closed subset C in X is an affinoid domain in X, so by [Ber1, 2.5.13(*ii*)] its relative interior $\mathrm{Int}(C/X)$ is equal to the topological interior of C in X. Passing to complements, the relative boundary $\partial(C/X)$ is equal to the topological boundary $\partial_X(C)$ of C in X. (See [Ber1, 2.5.7] for these notions of relative interior and boundary for morphisms of affinoid Berkovich spaces.) By the transitivity relation for relative interior with respect to a composite of morphisms [Ber1, 2.5.13(*iii*)], applied to $C \to X \to \mathscr{M}(k)$, we obtain

$$\partial(C/\mathscr{M}(k)) = \partial_X(C) \cup (C \cap \partial(X/\mathscr{M}(k))).$$

For any pure 1-dimensional strictly k-analytic affinoid Berkovich space Z, the relative boundary with respect to the base field coincides with the Shilov boundary. (Proof: By Noether normalization there is a finite map $Z \to \mathbf{B}^1$ to the closed unit ball, and by Theorem A.1.1, [BGR, 6.3.5/1], and [Ber1, 2.4.4, 2.5.8(*iii*), 2.5.13(*i*)] we are thereby reduced to the case of the case $Z = \mathbf{B}^1$. By [Ber1, 2.5.2(*d*), 2.5.12] we have $\partial(\mathbf{B}^1/\mathscr{M}(k)) = \{\|\cdot\|_{\sup}\} = \Gamma(\mathbf{B}^1)$.) Hence, $\Gamma(C) = \partial_X(C) \cup (C \cap \Gamma(X)) = \partial_X(C)$ since $C \cap \Gamma(X) = \emptyset$. Any neighborhood of a point in $\partial_X(C)$ meets the locus $U - C$ on which $|a| < r$, so by continuity of $|a|$ on X we have $|a| \leq r$ on $\partial_X(C) = \Gamma(C)$. But $\Gamma(C) \subseteq C \subseteq X^{\geq r}$, so $|a| = r$ on $\Gamma(C)$. By the maximum principle for the Shilov boundary of an affinoid, we conclude $|a| \leq r$ on C, so $|a| = r$ on C (as $C \subseteq X^{\geq r}$). Since $|a| < r$ on $U - C$, this implies $|a| \leq r$ on U.

To get a contradiction, pick a point $c \in \Gamma(C)$ and let $X' = \mathscr{M}(A') \subseteq U$ be a strictly k-analytic affinoid subdomain of X that contains c. Since $X' \subseteq U$, the sup-norm of $a|_{X'}$ (in the equivalent senses of rigid spaces or Berkovich spaces) is at most r, so it is equal to r because $|a(c)| = r$ and $c \in X'$. Let $X'' = \mathscr{M}(A'') \subseteq X'$ be a connected strictly k-analytic neighborhood of c in X' with X'' disjoint from the finite set $\Gamma(X')$. (Note that X'' must also be a neighborhood of c in X.) Since $\Gamma(X')$ is the preimage of the generic points under the analytic reduction map $X' \to \mathrm{Spec}(\widetilde{A'})$, by surjectivity of the reduction map $X'' \to \mathrm{Spec}(\widetilde{A''})$ [Ber1, 2.4.4(*i*)] we conclude that the natural map $\mathrm{Spec}(\widetilde{A''}) \to \mathrm{Spec}(\widetilde{A'})$ has constructible image that hits no generic points and is connected (as $\mathrm{Spec}(\widetilde{A''})$ is connected, due to connectivity of $\mathrm{Sp}(A'')$). The only nowhere-dense connected constructible subsets of a pure 1-dimensional algebraic $\widetilde{k}$-scheme are the closed points, so $\mathrm{Spec}(\widetilde{A''})$ maps onto a single closed point in $\mathrm{Spec}(\widetilde{A'})$ that must be the analytic reduction of c.

We shall prove that $a|_{X''}$ has absolute value r at all points of X'', and this gives a contradiction because the neighborhood X'' of $c \in \Gamma(C) = \partial_X(C)$ in X meets the locus $U - C$ on which $|a| < r$. Let n be a sufficiently divisible positive integer so that $r^n = |\rho|$ with $\rho \in k^\times$. The analytic function $f = a^n/\rho$ has sup-norm 1 on X' with associated algebraic function on $\mathrm{Spec}(\widetilde{A'})$ that is a unit at the analytic reduction of c. The restriction

$f|_{X''}$ is also power-bounded, and by functoriality of analytic reduction the reduction of $f|_{X''}$ on $\mathrm{Spec}(\widetilde{A''})$ is the pullback of the reduction of $f|_{X'}$ on $\mathrm{Spec}(\widetilde{A'})$ under the natural map $\mathrm{Spec}(\widetilde{A''}) \to \mathrm{Spec}(\widetilde{A'})$. But this latter map is a constant map to a closed point in the unit locus for the reduction of $f|_{X'}$, so we conclude that $f|_{X''}$ has non-vanishing reduction. That is, $f|_{X''}$ has constant absolute value 1, or equivalently $a|_{X''}$ has constant absolute value r as desired. ■

4.3. **Relativization and Frobenius kernels.** The variation of canonical subgroups in rigid-analytic families goes as follows:

Theorem 4.3.1. *Let $h = h(p, g, n) \in (p^{-1/8}, 1)$ be as in Theorem 4.1.1 (adapted to a fixed choice of $r_n \in (p^{-1/p^{n-1}(p-1)}, 1) \cap p^{\mathbf{Q}}$), and let $k/\mathbf{Q}_p$ be an analytic extension field. Choose an abeloid space $A \to S$ with relative dimension g over a rigid-analytic space over k, and assume either that (i) $A_{/S}$ admits a polarization fpqc-locally on S or (ii) $A_{/S}$ becomes algebraic after local finite surjective base change. Also, assume $h(A_s) > h$ for all $s \in S$.*

There exists a unique finite étale subgroup $G_n \subseteq A[p^n]$ with rank p^{ng} such that G_n recovers the level-n canonical subgroup on fibers, and the formation of G_n is compatible with base change on S and (for quasi-separated or pseudo-separated S) with change of the base field. The dual $(A[p^n]/G_n)^\vee$ is the analogous such subgroup for $A^\vee$, and $G_n[p^m] = G_m$ for $0 \le m \le n$.

Note that under either hypothesis (i) or (ii) each abeloid fiber A_s becomes an abelian variety after a finite extension on $k(s)$, and hence (by descending a suitable ample line bundle) each A_s is an abelian variety. Thus, it makes sense to speak of a Hasse invariant for each fiber A_s. Also, Theorem 3.1.1 and Theorem 3.2.3 ensure that the hypothesis on fibral Hasse invariants exceeding h is preserved under arbitrary change of the base field (for quasi-separated or pseudo-separated S).

Proof. The uniqueness of G_n and the description of its p-power torsion subgroups follow from connectivity considerations and our knowledge on fibers, and the same goes for the behavior with respect to Cartier duality. Thus, the existence result is preserved by base change. By rigid-analytic *fpqc* descent theory [C2, §4.2], it suffices to work *fpqc*-locally on S to prove the theorem. In particular, we may and do assume S is quasi-compact and quasi-separated (*e.g.*, affinoid). By Lemma 4.3.2 below (applied with $Y = A[p^n]$ over $X = S$), it also suffices to make the construction after a finite surjective base change. Thus, using Corollary 3.2.2 in case (ii), we can assume that $A_{/S}$ admits a polarization of some constant degree d^2 and that the finite étale S-groups $A[N]$ and $A^\vee[N]$ are split for a fixed choice of $N \ge 3$ not divisible by p. In particular, by Zahrin's trick $(A \times A^\vee)^4$ is a pullback of the universal principally polarized abeloid space over $\mathscr{A}^{\mathrm{an}}_{8g,1,N/\mathbf{Q}_p}$ along a map $f : S \to \mathscr{A}^{\mathrm{an}}_{8g,1,N/\mathbf{Q}_p}$. Let $\mathfrak{A} \to \mathfrak{X}$ denote the p-adic completion of a semi-abelian scheme over a proper flat $\mathbf{Z}_p$-scheme that extends the universal abelian scheme over $\mathscr{A}_{8g,1,N/\mathbf{Z}_p}$, so $(A \times A^\vee)^4$ is a pullback of $\mathfrak{A}^{\mathrm{rig}} \to \mathfrak{X}^{\mathrm{rig}}$ along a map $f : S \to \mathfrak{X}^{\mathrm{rig}}$. For a suitable formal admissible blow-up $\mathfrak{X}'$ of $\mathfrak{X}$ and the pullback (or equivalently, strict transform) $\mathfrak{A}'$ of $\mathfrak{A}$, we may find a quasi-compact flat formal model $\mathfrak{S}$ for S and a map $\mathfrak{f} : \mathfrak{S} \to \mathfrak{X}'$ such that $\mathfrak{f}^{\mathrm{rig}} = f$. In particular, the pullback $\mathfrak{f}^*(\mathfrak{A}')$ is a formal semi-abelian scheme over $\mathfrak{S}$ whose rigid-analytic generic fiber is the abeloid S-group $(A \times A^\vee)^4$.

For each $s \in S$ et $G_{n,s} \subseteq A_s[p^n]$ be the level-n canonical subgroup of A_s. The subgroup $(G_{n,s} \times (A_s[p^n]/G_{n,s})^\vee)^4 \subseteq (A \times A^\vee)^4_s[p^n]$ is the level-n canonical subgroup of $(A \times A^\vee)^4_s$ since $h(p, g, n)$ is adapted to a fixed choice of $r_n \in (p^{-1/p^{n-1}(p-1)}, 1) \cap p^{\mathbf{Q}}$. Hence, if we can find a finite étale S-subgroup C_n of $(A \times A^\vee)^4$ that recovers the level-n canonical subgroup on fibers then the image G_n of C_n under projection to the first factor of the finite étale eight-fold product $(A \times A^\vee)^4[p^n] \simeq (A[p^n] \times A^\vee[p^n])^4$ over S is a finite étale S-subgroup of $A[p^n]$ that has the required properties. It is therefore enough to find such a C_n in the p^n-torsion of $(A \times A^\vee)^4$. Working locally on $\mathfrak{S}$, we may assume that the Lie algebra of $\mathfrak{f}^*(\mathfrak{A}')$ is globally free (of rank $8g$) as a coherent $\mathscr{O}_{\mathfrak{S}}$-module, so the formal completion $\mathfrak{A}'_0$ of $\mathfrak{f}^*(\mathfrak{A}')$ along the identity section of its mod-p fiber is identified with a g-variable formal group law over $\mathfrak{S}$.

Exactly as in Steps 2 and 3 of the proof of Theorem 4.1.1, the Berthelot generic fiber ${\mathfrak{A}'_0}^{\mathrm{rig}}$ is an admissible open subgroup of $\mathfrak{f}^*(\mathfrak{A}')^{\mathrm{rig}} = (A \times A^\vee)^4$ and its locus with fibral polyradius $\le r_n$ in $(A \times A^\vee)^4$ is a quasi-compact admissible open S-subgroup that is denoted $(A \times A^\vee)^4_{\le r_n}$. The overlap C_n of this latter subgroup with the finite étale S-subgroup $(A \times A^\vee)^4[p^n]$ is a quasi-compact separated étale S-subgroup whose s-fiber is $(G_{n,s} \times (A_s[p^n]/G_{n,s})^\vee)^4$ for all $s \in S$. In particular, $C_{n,s}$ has rank p^{4ng} that is independent of s, so by [C4, Thm. A.1.2] the map $C_n \to S$ is *finite*. Hence, the S-subgroup $C_n \subseteq (A \times A^\vee)^4[p^n]$ has the required properties. ■

The following lemma was used in the preceding proof:

Lemma 4.3.2. *Let $f : X' \to X$ be a finite surjective map between schemes or rigid spaces, and let $Y \to X$ be a finite étale cover with pullback $Y' \to X'$ along f. If $i' : Z' \hookrightarrow Y'$ is a closed immersion with Z' finite étale over X' then i' descends to a closed immersion $i : Z \hookrightarrow Y$ with Z finite étale over X if and only if it does so on fibers over each point $x \in X$.*

Proof. Let $p_1, p_2 : X'' = X' \times_X X' \rightrightarrows X'$ be the standard projections, and let $Y'' = X'' \times_X Y$. By the fibral descent hypothesis, the finite étale X''-spaces $p_1^*(Z')$ and $p_2^*(Z')$ inside of the finite étale X''-space Y'' coincide over X''_x for all $x \in X$, and so $p_1^*(Z') = p_2^*(Z')$ inside Y''. The problem is therefore to show that finite étale covers satisfy effective (and uniquely functorial) descent with respect to finite surjective maps. By working locally on the base, the rigid-analytic case is reduced to the case of schemes (using affinoid algebras). The case of schemes is [SGA1, IX, 4.7]. ■

Now we turn to the problem of relating canonical subgroups and Frobenius kernels. Let A be a g-dimensional abelian variety over $k/\mathbf{Q}_p$ with $h(A) > h(p, g, n)$, and pass to a finite extension of k if necessary so that A has semistable reduction in the sense of Theorem 2.1.9. Let $\mathfrak{A}_R$ be the associated formal semiabelian scheme over R, and let t and a be the respective relative dimensions of the formal toric and abelian parts $\mathfrak{T}$ and $\mathfrak{B}$ of $\mathfrak{A}_R$ (so $t + a = g$). Thus, $\mathfrak{A}_R[p^n]$ is a finite flat group scheme over R with geometric generic fiber that is free of rank $t + 2a$ as a $\mathbf{Z}/p^n\mathbf{Z}$-module. Since $h(A) > h(p, g, n)$ there is a level-n canonical subgroup $G_n \subseteq A[p^n]^0 \subseteq \mathfrak{A}_R[p^n]^0$ and so by schematic closure this is the k-fiber of a unique finite flat closed R-subgroup $\mathscr{G}_n \subseteq \mathfrak{A}_R[p^n]^0$ with order p^{ng}. (Since the valuation ring R is local, this finite flat schematic closure is automatically finitely presented as an R-module even if R is not noetherian.) Likewise, $G_m = G_n[p^m]$ is a level-m canonical subgroup for all $1 \le m \le n$ and we let $\mathscr{G}_m \subseteq \mathfrak{A}_R[p^m]^0$ denote its closure.

By definition, $\mathscr{G}_m$ is contained in the identity component $\mathfrak{A}_R[p^m]^0$ whose geometric generic fiber is a free $\mathbf{Z}/p^m\mathbf{Z}$-module with rank $t + h_0$, where $h_0 \ge a$ is the height of the local part of the p-divisible group of $\mathfrak{B}$. In the ordinary case we have $t + h_0 = g$ and so $\mathscr{G}_m = \mathfrak{A}_R[p^m]^0$; thus, $\mathscr{G}_m \bmod pR \subseteq \overline{\mathfrak{A}}_R \stackrel{\text{def}}{=} \mathfrak{A}_R \bmod pR$ is the kernel of the m-fold relative Frobenius

$$F_{\overline{\mathfrak{A}}_R,m,R/pR} : \overline{\mathfrak{A}}_R \to \overline{\mathfrak{A}}_R^{(p^m)}.$$

In the non-ordinary case $t + h_0 > g$ and we cannot expect $\mathscr{G}_m \bmod pR$ to equal $\ker F_{\overline{\mathfrak{A}}_R,m,R/pR}$. Working modulo $p^{1-\varepsilon}$ for a small $\varepsilon > 0$ we get a congruence by taking $h(A)$ near enough to 1 in a "universal" manner:

Theorem 4.3.3. *Fix p, g, and $n \ge 1$, and pick $\lambda \in (0, 1) \cap \mathbf{Q}$. There exists $h(p, g, n, \lambda) \in (h(p, g, n), 1) \cap p^{\mathbf{Q}}$ such that if $h(A) > h(p, g, n, \lambda)$ then $\mathscr{G}_m \bmod p^\lambda R = \ker(F_{\mathfrak{A}_R \bmod p^\lambda R,m,R/p^\lambda R})$ for $1 \le m \le n$.*

In the theorem and its proof, the terminology "modulo $p^\lambda R$" really means "modulo $c'R'$" for the valuation ring R' of any analytic extension k'/k and any $c' \in R'$ satisfying $|c'| \ge p^{-\lambda}$. The implicit unspecified extension of scalars is necessary in order to make sense of the assertion that the same λ works across all extensions of the base field without the restriction $p^\lambda \in |k^\times|$ that is unpleasant in the discretely-valued case (as λ near 1 is the interesting case). We will typically abuse notation and write expressions such as $R/p^\lambda R$ that the reader should understand to mean $R'/c'R'$ for any R' and c' as above; this abuse of terminology streamlines the exposition and does not create serious risk of error.

Proof. The ordinary case is a triviality, so we may restrict attention to those A with $h(A) < 1$. We also may and do restrict attention to the case $m = n$. The formal semi-abelian model $\mathfrak{A}_R$ for A fits into a short exact sequence

$$0 \to \mathfrak{T} \to \mathfrak{A}_R \to \mathfrak{B} \to 0$$

with a formal torus $\mathfrak{T}$ and uniquely algebraizable formal abelian scheme $\mathfrak{B}$ over $\mathrm{Spf}(R)$. Let t and a be the respective relative dimensions of the toric and abelian parts, so $a > 0$ since $h(A) < 1$. Let B_R be the associated abelian scheme over R, so its generic fiber B over k is an a-dimensional abelian variety with the same Hasse invariant as A. The dual $A^\vee$ also has semistable reduction and its abelian part is identified with the formal completion $\mathfrak{B}^\vee$ of the dual abelian scheme $B_R^\vee$ whose generic fiber is $B^\vee$. Hence, $(A \times A^\vee)^4$ has semistable reduction with formal abelian part $(\mathfrak{B} \times \mathfrak{B}^\vee)^4$ arising from the abelian scheme $(B_R \times B_R^\vee)^4$ whose generic fiber $(B \times B^\vee)^4$ is principally polarized with good reduction.

By Theorem 4.1.1 for all $g \geq 1$ we can find $h(p, g, n) \in (p^{-1/8}, 1)$ sufficiently near 1 and adapted to a fixed choice of $r_n \in (p^{-1/p^{n-1}(p-1)}, 1)$ so that if a g-dimensional A satisfies $h(A) > h(p, g, n)$ then there is a level-n canonical subgroup G_n in A and $(G_n \times (A[p^n]/G_n)^\vee)^4$ is a level-n canonical subgroup of $(A \times A^\vee)^4$. Since the formation of schematic closure (over R) and relative Frobenius maps commute with products and $h((A \times A^\vee)^4) = h(A)^8 \in (1/p, 1)$, it therefore suffices to work with $(A \times A^\vee)^4$ rather than A provided that we use the bound $h_{\rm pp}(p, 8g, n) \in (1/p, 1)$ from the principally polarized case (as in the proof of Theorem 4.1.1). In particular, we can assume that A and B admit principal polarizations (and we rename $8g$ as g). Consider the p-adic completion $\mathfrak{A} \to \mathfrak{M}$ of the universal abelian scheme over the finite type moduli scheme $M = \mathscr{A}_{g',1,N/\mathbf{Z}_p}$ over $\mathbf{Z}_p$, with $N \geq 3$ a fixed integer relatively prime to p and $1 \leq g' \leq g$. By increasing k so that the finite étale R-scheme $B_R[N]$ is constant, the principally polarized abelian scheme B arises from a point on the k-fiber of the generic fiber $\mathfrak{A}^{\rm rig} \to \mathfrak{M}^{\rm rig}$ in the case $g' = a$. Theorem 4.3.1 provides a relative level-n canonical subgroup over the locus in $\mathfrak{M}^{\rm rig}$ where the Hasse invariant is $> h_{\rm pp}(p, g', n)$ for some universal $h_{\rm pp}(p, g', n) < 1$, and so the proof of [C4, Thm. 4.3.1] (the case $g = 1$) applies to this situation. (The proof of [C4, Thm. 4.3.1] was specifically written to be applicable to the present circumstances.) This provides an $h_{\rm good}(p, g', n, \lambda)$ that "works" in the g'-dimensional principally polarized case with good reduction for any $g' \geq 1$.

We now check that $h(p, g, n, \lambda) = \max(h_{\rm pp}(p, g, n), \max_{1 \leq g' \leq g}(h_{\rm good}(p, g', n, \lambda))) \in (1/p, 1) \cap p^{\mathbf{Q}}$ works for A. Since A and B have the same Hasse invariant, if $h(A) > h(p, g, n, \lambda)$ then B has a level-n canonical subgroup G'_n whose schematic closure $\mathscr{G}'_n$ in $B[p^n]^0 = \mathfrak{B}[p^n]^0$ reduces to the kernel of the n-fold relative Frobenius modulo p^λ. We have an exact sequence of identity components

$$0 \to \mathfrak{T}[p^n] \to \mathfrak{A}_R[p^n]^0 \stackrel{\pi_n^0}{\to} \mathfrak{B}[p^n]^0 \to 0, \tag{4.3.1}$$

so the π_n^0-preimage $\widetilde{\mathscr{G}}'_n \subseteq \mathfrak{A}_R[p^n]^0$ of $\mathscr{G}'_n \subseteq \mathfrak{B}[p^n]^0$ is a finite flat closed R-group of $\mathfrak{A}_R[p^n]^0$ whose k-fiber is the full preimage of G'_n in $A[p^n]^0$. In Step 7 of the proof of Theorem 4.1.1 we saw that the full preimage of G'_n in $A[p^n]^0$ *is* the level-n canonical subgroup G_n of A, and so $\widetilde{\mathscr{G}}'_n$ as just defined is indeed the schematic closure $\mathscr{G}_n$ of G_n in $\mathfrak{A}_R[p^n]^0$.

We therefore need to prove that $\widetilde{\mathscr{G}}'_n \bmod p^\lambda R$ is killed by its relative n-fold Frobenius morphism (and then order considerations force it to coincide with the kernel of the n-fold relative Frobenius morphism in the formal group of $\mathfrak{A}_R \bmod p^\lambda R$ along its identity section). Since $\mathscr{G}'_n$ reduces to the corresponding Frobenius-kernel in $\mathfrak{B}[p^n]^0 \bmod p^\lambda R$, it suffices to check that the containment

$$\ker(F_{\mathfrak{A}_R \bmod p^\lambda R, n, R/p^\lambda R}) \subseteq (\pi_n^0)^{-1}(\ker(F_{\mathfrak{B} \bmod p^\lambda R, n, R/p^\lambda R}))$$

of closed subschemes inside $\mathfrak{A}_R \bmod p^\lambda R$ (which follows from the functoriality of relative Frobenius) is an equality. Both terms are finite flat $R/p^\lambda R$-schemes and they have the same rank $p^{ng} = p^{nt} \cdot p^{na}$ (since π_n^0 in (4.3.1) is a finite locally free map with degree equal to the order p^{nt} of its kernel $\mathfrak{T}[p^n]$). Hence, equality is forced. ∎

Control over reduction of canonical subgroups allows us to give a partial answer to the question of how the Hasse invariant and level-n canonical subgroup (for $n > 1$) behave under passage to the quotient by a level-m canonical subgroup for $1 \leq m < n$.

Corollary 4.3.4. *Choose $n \geq 2$ and $r_n \in (p^{-1/p^{n-1}(p-1)}, 1) \cap p^{\mathbf{Q}}$. Consider $1 \leq m < n$ and $\lambda \in (0, 1) \cap p^{\mathbf{Q}}$ such that $p^{-\lambda} \leq r_n^{p^m}$. Let $h = \max(h(p, g, n, \lambda), p^{-\lambda/p^m}) \in (h(p, g, n), 1) \cap p^{\mathbf{Q}}$ with $h(p, g, n)$ adapted to r_n in the sense of Theorem* 4.1.1.

For any analytic extension field $k/\mathbf{Q}_p$ and g-dimensional abelian variety A over k such that $h(A) > h$, the quotient A/G_m has Hasse invariant $h(A)^{p^m}$ and G_n/G_m is a level-$(n-m)$ canonical subgroup that is equal to $(A/G_m)[p^{n-m}]^0_{\le r_n^{p^m}}$. Moreover, after replacing k with a finite extension so that A has semistable reduction, the quotient A/G_m has semistable reduction and the reduction of G_n/G_m modulo p^λ coincides with the kernel of the relative $(n-m)$-fold Frobenius on the formal semi-abelian model for A/G_m modulo p^λ.

Remark 4.3.5. Any $\lambda \in [1/(p^{(n-m)-1}(p-1)), 1) \cap \mathbf{Q}$ satisfies the hypotheses in the corollary, regardless of r_n, and it is λ near 1 that are of most interest anyway. Such a "universal" λ can be found if and only if $1/p^{(n-m)-1}(p-1) < 1$, so if $p = 2$ then we have to require $m < n-1$ (and hence $n \ge 3$) in order that such a universal λ may be found (though if we do not care about λ being independent of r_n then some λ can always be found). For example, for odd p we may always take $\lambda = 1/(p-1)$ and for $m < n-1$ we may always take $\lambda = 1/p(p-1)$ for any p.

Remark 4.3.6. For $h(A) > h(p,g,n)$ the dual $(A/G_m)^\vee$ is identified with the quotient of $A^\vee$ modulo the subgroup $(A[p^m]/G_m)^\vee$ that is its level-m canonical subgroup, so for A as in the corollary the level-$(n-m)$ canonical subgroup of $(A/G_m)^\vee$ is

$$(A[p^n]/G_n)^\vee/(A[p^m]/G_m)^\vee \simeq (A[p^{n-m}]/G_{n-m})^\vee.$$

Also, upon fixing $1 \le m < n$ and choosing r_n and λ, for h as in this corollary we may take $h(p,g,m) = h^{p^{n-m}}$ with $r_m = r_n^{p^{n-m}} \in (p^{-1/p^{m-1}(p-1)}, 1) \cap p^{\mathbf{Q}}$ as the universal size bound in Theorem 4.3.1 for level-m canonical subgroups. The reader should compare Corollary 4.3.4 with the more precise results [C4, Thm. 4.2.5, Cor. 4.2.6] in the case $g = 1$ (where the size estimates and calculation of the Hasse invariant of the quotient have no logical dependence on Frobenius kernels, essentially because the formal group only depends on a single parameter).

Proof. (of Corollary 4.3.4). Replace k with a finite extension so that $p^\lambda \in |k^\times|$ and there is a formal semi-abelian model $\mathfrak{A}_R$ for A. For all $1 \le \nu \le n$ the closure $\mathscr{G}_\nu$ in $\mathfrak{A}_R[p^\nu]^0$ of the level-ν canonical subgroup $G_\nu = G_n[p^\nu]$ reduces to the kernel of the relative ν-fold Frobenius modulo p^λ. Let $\overline{\mathfrak{A}}_{R,\lambda}$ be the reduction of $\mathfrak{A}_R$ modulo p^λ. The reduction modulo p^λ for the formal semi-abelian model $\mathfrak{A}_R/\mathscr{G}_m$ of A/G_m is thereby identified with $\overline{\mathfrak{A}}_{R,\lambda}^{(p^m)}$, so the relative Verschiebung for $\mathfrak{A}_R/\mathscr{G}_m$ mod p^λ is identified with the m-fold Frobenius base change of the relative Verschiebung for the smooth $R/p^\lambda R$-group $\overline{\mathfrak{A}}_{R,\lambda}$. Hence, passing to induced $R/p^\lambda R$-linear maps on Lie algebras, the associated determinant ideal in $R/p^\lambda R$ for $\mathfrak{A}_R/\mathscr{G}_m$ mod $p^\lambda R$ is the p^mth power of the determinant of $\mathrm{Lie}(V_{\overline{\mathfrak{A}}_{R,\lambda}})$. This implies $h(A/G_m) = h(A)^{p^m}$ since $h(A)^{p^m} > h^{p^m} \ge p^{-\lambda}$.

Now we show that G_n/G_m is a level-$(n-m)$ canonical subgroup of A/G_m. Clearly its module structure is $(\mathbf{Z}/p^{n-m}\mathbf{Z})^g$, so it suffices to prove that this subgroup of $(A/G_m)[p^{n-m}]^0$ is precisely the subgroup of elements with size $\le r_n^{p^m}$. First, for $x \in G_n$ we claim that $\mathrm{size}_{A/G_m}(x) \le r_n^{p^m}$. Since $r_n^{p^m} \ge p^{-\lambda}$, we can work modulo p^λ. The projection from A to A/G_m reduces to the m-fold relative Frobenius map on $\overline{\mathfrak{A}}_{R,\lambda}$, so it raises size to the p^mth power modulo $p^\lambda R$. More precisely, if $x \in A$ extends to an integral point of $\mathfrak{A}_R$ and $\mathrm{size}_A(x) \le p^{-\lambda/p^m}$ then $\mathrm{size}_{A/G_m}(x \bmod G_m) \le p^{-\lambda} \le r_n^{p^m}$, whereas if $\mathrm{size}_A(x) > p^{-\lambda/p^m}$ then $\mathrm{size}_{A/G_m}(x \bmod G_m) = \mathrm{size}_A(x)^{p^m}$. Hence, $G_n/G_m \subseteq (A/G_m)[p^{n-m}]^0_{\le r_n^{p^m}}$. If this inclusion is not an equality then there is a point $\overline{x}_0 \in (A/G_m)[p^{n-m}]^0$ with size $\le r_n^{p^m}$ such that it does not lift into G_n in A. Since $\mathfrak{A}_R \to \mathfrak{A}_R/\mathscr{G}_m$ is finite flat of degree p^m, the image of $A[p^n]^0$ in $(A/G_m)[p^n]^0$ contains $(A/G_m)[p^{n-m}]^0$. We may therefore find a lift $x_0 \in A[p^n]^0$ of $\overline{x}_0$, and $x_0 \notin G_n = A[p^n]^0_{\le r_n}$. By the preceding general size considerations, since $\mathrm{size}_A(x_0) > r_n \ge p^{-\lambda/p^m}$ we get $\mathrm{size}_{A/G_m}(\overline{x}_0) = \mathrm{size}_A(x_0)^{p^m} > r_n^{p^m}$, contradicting how $\overline{x}_0$ was chosen. ■

4.4. **Comparison with other approaches to canonical subgroups.** We conclude this paper by comparing Theorem 4.1.1 and Theorem 4.3.1 with results in [AM], [AG], [GK], and [KL]. In [AM], level-1 canonical subgroups are constructed on abelian varieties over k when $p \ge 3$ and k is discretely-valued with perfect residue field, and an explicit sufficient lower bound on the Hasse invariant is given in terms of p and g (our method does not make $h(p,g,n)$ explicit for $g > 1$, even with $n = 1$). The construction in [AM] is characterized by a completely different fibral property coming out of p-adic Hodge theory, so we must use the arguments in Steps 7 and 8 of the proof of Theorem 4.1.1 (especially the existence of ordinary points on certain connected components via Theorem 4.2.1) to conclude that this construction agrees with ours for level-1 canonical groups, at least for Hasse invariants sufficiently close to 1 (where "sufficiently close" only depends on p and g but is not made explicit by our methods since our $h_{\mathrm{pp}}(p, g, 1)$ in the principally polarized case is not explicit). The methods in [AM] do not appear to give information concerning either higher-level canonical subgroups or level-1 canonical subgroups with $p = 2$ or general (*e.g.*, algebraically closed) k.

The methods in [AG] are algebro-geometric rather than rigid-analytic, and give a theory of level-1 canonical subgroups in families of polarized abelian varieties with good reduction over any normal p-adically separated and complete base scheme. A discreteness hypothesis is required on the base field, though this restriction is probably not necessary for the construction in [AG] to be pushed through. One advantage in [AG] is a strong uniqueness result (ensuring compatibility with products and with Frobenius-kernels modulo $p^{1-\varepsilon}$, as well as with any other theory satisfying a few axioms), but the restriction to families with good reduction seems to be essential in this work.

Finally, in [GK] and [KL] rigid-analytic methods (different from ours) are used to establish the "overconvergence" of the canonical subgroup in the universal families over some modular varieties for which well-understood integral models exist. In [GK] there is given a very detailed treatment for canonical subgroups over Shimura curves and an exact description of the maximal connected domains over which canonical subgroups exist; the fine structure of integral models for the 1-dimensional modular variety underlies the technique. As in Theorem 4.1.1, no explicit bound on the Hasse invariant is given by the general methods in [KL]. Whereas our abstract bound $h(p, g, 1)$ only depends on p and g, in principle the elegant construction in [KL] gives a "radius of overconvergence" that may depend on the specific modular variety that is considered. In particular, in contrast with our viewpoint and the viewpoints in [AM] and [K], since the approach in [KL] does not assign an *a priori* intrinsic meaning to the notion of a canonical subgroup in the p-torsion of an individual abelian variety it does not seem to follow from the methods in [KL] that if an abelian variety arises in several fibers near the ordinary locus over a modular variety then the induced level-1 canonical subgroups in these fibers must coincide or be independent of the choice of modular variety. (Our methods, such as Lemma 4.1.4, ensure that these difficulties do not arise for Hasse invariants sufficiently close to 1 in a universal manner.)

Appendix A. Some input from rigid geometry

There are several results from rigid geometry that were used in the body of the paper but whose proofs were omitted there so as to avoid interrupting the main lines of argument. We have gathered these results and their proofs in this appendix.

A.1. **Fiber dimension and reduction.**

The following must be well-known, but we could not find a published reference:

Theorem A.1.1. *If B is a k-affinoid algebra of pure dimension d, then its analytic reduction $\widetilde{B}$ over the residue field $\widetilde{k}$ also has pure dimension d.*

Proof. By [BGR, 6.3.4] the ring $\widetilde{B}$ is a d-dimensional $\widetilde{k}$-algebra of finite type, so the problem is to show that $\mathrm{Spec}(\widetilde{B})$ has no irreducible component with dimension strictly smaller than d. Equivalently, we have to rule out the existence of $\widetilde{b} \in \widetilde{B}$ such that $\widetilde{B}[1/\widetilde{b}]$ is nonzero with dimension $< d$.

The description of $\widetilde{B}$ in terms of the supremum seminorm shows that the natural map to the reduced quotient $B \to B_{\mathrm{red}}$ induces an isomorphism on analytic reductions. Hence, we can assume B is reduced. Since $\widetilde{B}$ is of finite type over $\widetilde{k}$, we can find a topologically finite type R-subalgebra $\mathscr{B}$ (*i.e.*, a quotient of a restricted power series ring $R\{\{t_1, \dots, t_n\}\}$) contained in the subring of power-bounded elements of B such that $k \otimes_R \mathscr{B} = B$ and $\mathscr{B} \to \widetilde{B}$ is surjective. Since $\mathscr{B}$ is R-flat, by [BL3, Prop. 1.1(c)] the R-algebra $\mathscr{B}$ is topologically finitely presented (so it provides a formal model for B in the sense of Raynaud, but we do not require Raynaud's theory here). In particular, if $\mathfrak{I}$ denotes an ideal of definition of R then there is a natural surjection $\mathscr{B}_0 \stackrel{\text{def}}{=} \mathscr{B}/\mathfrak{I}\mathscr{B} \twoheadrightarrow \widetilde{B}$ of finitely generated $\widetilde{k}$-algebras and (by definition) $\widetilde{B}$ is reduced. We claim that the kernel of this map consists entirely of nilpotents, so the quotient $\mathscr{B}_{\text{red}}$ of $\mathscr{B}$ modulo topological nilpotents coincides with $\widetilde{B}$. This ensures that $\mathscr{B}_0$ is a d-dimensional algebra and so our problem will be equivalent to the assertion that $\mathrm{Spec}(\mathscr{B}_0)$ is equidimensional.

To prove the nilpotence of any $b_0 \in \ker(\mathscr{B}_0 \twoheadrightarrow \widetilde{B})$, we lift b_0 to an element $b \in \mathscr{B}$ that satisfies $|b|_{\text{sup}} < 1$ on $\mathrm{Sp}(B)$ and we have to show that b has nilpotent image in $\mathscr{B}_0$. It is equivalent to show that b lies in every maximal ideal of the ring $\mathscr{B}_0$, for then it will lie in every maximal ideal of the reduced quotient $(\mathscr{B}_0)_{\text{red}}$ that is finitely generated over the field $\widetilde{k}$ and hence will vanish in this quotient (*i.e.*, b_0 is in the nilradical of $\mathscr{B}_0$), as desired. Let $\mathfrak{n} \in \mathrm{Spec}(\mathscr{B}_0)$ be a closed point (identified with a maximal ideal of $\mathscr{B}$). The theory of rig-points on formal models [BL3, 3.5] provides a point $x \in \mathrm{Sp}(B) = \mathrm{MaxSpec}(k \otimes_R \mathscr{B})$ such that if $\mathfrak{p} = \ker(\mathscr{B} \to k(x))$ then under the projection from $\mathscr{B}$ to its R-flat (and R-finite and local) quotient $\mathscr{B}/\mathfrak{p} \subseteq k(x)$ the preimage of the unique maximal ideal of $\mathscr{B}/\mathfrak{p}$ is $\mathfrak{n}$. But since $|b|_{\text{sup}} < 1$ we have $b(x) = 0$ in $k(x)$, which is to say $b \in \mathfrak{p}$. Hence, we get the required result $b \in \mathfrak{n}$.

It remains to show that the d-dimensional $\mathrm{Spec}(\mathscr{B}_0)$ is equidimensional. It is equivalent to prove that every non-empty basic open affine $\mathrm{Spec}(\mathscr{B}_0[1/b_0])$ has dimension d. Pick any $b_0 \in \mathscr{B}_0$ such that $\mathrm{Spec}(\mathscr{B}_0[1/b_0])$ is non-empty. Since the quotient $(\mathscr{B}_0)_{\text{red}}$ is identified with $\widetilde{B}$ and the Zariski-open non-vanishing locus for b_0 in $\mathrm{Spec}(\mathscr{B}_0)$ is non-empty, we conclude that b_0 has *nonzero* image in $(\mathscr{B}_0)_{\text{red}} = \widetilde{B}$. Hence, if $b \in \mathscr{B}$ is a lift of b_0 then as a power-bounded element of B it has nonzero image in $\widetilde{B}$. That is, $|b|_{\text{sup}} = 1$. The affinoid subdomain $\mathrm{Sp}(B\langle 1/b\rangle)$ in $\mathrm{Sp}(B)$ is therefore non-empty and so has dimension d since $\mathrm{Sp}(B)$ is equidimensional of dimension d. We conclude that $\dim(B\langle 1/b\rangle) = d$, so the analytic reduction $(B\langle 1/b\rangle)^{\sim}$ is d-dimensional over $\widetilde{k}$. By [BGR, 7.2.6/3] this analytic reduction is (via the evident map from $\widetilde{B}$) naturally isomorphic to $\widetilde{B}[1/\widetilde{b}]$, where $\widetilde{b}$ is the image of b in $\widetilde{B}$. Since the nil-thickening $\mathscr{B}_0 \twoheadrightarrow \widetilde{B}$ carries b_0 to $\widetilde{b}$, it follows that $(\mathscr{B}_0[1/b_0])_{\text{red}} = \widetilde{B}[1/\widetilde{b}]$, so $\mathscr{B}_0[1/b_0]$ is d-dimensional as desired. ■

A.2. **Descent through proper maps.** It is topologically obvious that if $f : X' \to X$ is a proper surjection of schemes (or of topological spaces) and $U \subseteq X$ is a subset such that $f^{-1}(U) \subseteq X'$ is open then U is open in X. The analogue in rigid geometry with admissible opens is true, but it does not seem possible to prove this using classical rigid geometry, nor Raynaud's theory of formal models, even if we restrict to the case of finite f and $U \to X$ that is a quasi-compact open immersion. Gabber observed that by considering *all* formal models at once, as a Zariski-Riemann space, the problem can be solved:

Theorem A.2.1 (Gabber)**.** *If $f : X' \to X$ is a proper surjection of rigid spaces and $U \subseteq X$ is a subset such that $U' = f^{-1}(U) \subseteq X'$ is admissible open then $U \subseteq X$ is admissible open.*

Remark A.2.2. By Lemma 3.2.4, if U' is quasi-compact (resp. $U' \to X'$ is quasi-compact) then so is U (resp. $U \to X$).

The subsequent discussion is a detailed explanation of Gabber's proof of Theorem A.2.1, built up as a series of lemmas. Of course, to prove the theorem we may work locally on X and so we can assume X is affinoid. In particular, we can assume X (and hence X') is quasi-compact and quasi-separated. Rather than work only with such classical rigid spaces, we will work with Zariski-Riemann spaces. This amounts to working with the underlying topological spaces of the associated adic spaces in the sense of Huber, but since we only use the underlying topological spaces of certain adic spaces we do not require any serious input from the theory of adic spaces.

Definition A.2.3. Let X be a quasi-compact and quasi-separated rigid space. The *Zariski-Riemann space* $\mathrm{ZRS}(X)$ attached to X is the topological inverse limit of the directed inverse system of (quasi-compact and flat) formal models of X. (All transition maps are proper, by [L1, 2.5, 2.6].)

As we shall see shortly, these spaces $\mathrm{ZRS}(X)$ are spectral spaces in the sense of Hochster: a *spectral space* is a quasi-compact topological space T that is *sober* (*i.e.*, every irreducible closed set in T has a unique generic point) and admits a base $\mathscr{B}$ of quasi-compact opens such that $\mathscr{B}$ is stable under finite intersections (so in fact the overlap of any pair of quasi-compact opens is quasi-compact, which is to say that T is *quasi-separated*;

in [H, §12] this property is called *semispectral*). For example, if S is a quasi-compact and quasi-separated scheme then by taking $\mathscr{B}$ to be the collection of quasi-compact opens in the underlying topological $|S|$ we see that $|S|$ is spectral. (Conversely, in [H] it is shown that every spectral space arises as the spectrum of a ring, so spectral spaces are precisely the underlying topological spaces of quasi-compact and quasi-separated schemes; we shall not use this fact.)

A *spectral map* between spectral spaces is a continuous map that is quasi-compact (*i.e.*, the preimage of a quasi-compact open is quasi-compact). For example, if $f : S' \to S$ is a map between schemes whose underlying topological spaces are noetherian then $|f| : |S'| \to |S|$ is spectral. Thus, the inverse system of formal models for a fixed quasi-compact and quasi-separated rigid space X consists of spectral spaces with spectral transition maps, so Lemma A.2.6 below ensures that $\mathrm{ZRS}(X)$ is a spectral space. By the theory of formal models for morphisms [BL3, Thm. 4.1], this lemma also ensures that $X \rightsquigarrow \mathrm{ZRS}(X)$ is a (covariant) functor from the full subcategory of quasi-compact and quasi-separated rigid spaces to the category of spectral spaces equipped with spectral maps.

We need to record some properties of inverse limits in the category of spectral spaces, and to do this it is convenient to introduce a few general topological notions for a class of spaces that is more general than the class of spectral spaces in the sense that we weaken the sobriety axiom to the T_0 axiom. Let X be a T_0 topological space (*i.e.*, distinct points have distinct closures) that is quasi-compact and quasi-separated, and assume that the quasi-compact opens are a base for the topology. A *constructible set* in X is a member of the Boolean algebra of subsets of X generated by the quasi-compact opens. Explicitly, a constructible set in X is a finite union of overlaps $U \cap (X - U')$ for quasi-compact opens $U, U' \subseteq X$. The *constructible topology* on such an X is the topology having the constructible sets as a basis of opens, and the associated topological space is denoted X^{cons}. (By [EGA, IV_1, 1.9.3], if X is the underlying space of a quasi-compact and quasi-separated scheme then this notion of X^{cons} coincides with that defined more generally in [EGA, IV_1, 1.9.13].) An open (resp. closed) set in X^{cons} is an arbitrary union (resp. intersection) of constructible sets in X, and these are respectively called *ind-constructible* and *pro-constructible* sets in X. In particular, the constructible topology on X refines the given one on X. (In [H], X^{cons} is called the *patch topology* and a pro-constructible set is called a *patch*. Hochster's terminology has the advantage of brevity, but we choose to follow the terminology of Grothendieck that is more widely used in algebraic geometry.) If $Z \subseteq X$ is a closed subset then Z is also a quasi-compact and quasi-separated T_0-space such that the quasi-compact opens are a base for the topology, and it is clear that the constructible topology on X induces the constructible topology on Z.

Note that for any T_0-space X the topological space X^{cons} is always a Hausdorff space. Indeed, let $x, y \in X$ be distinct points, so either $x \notin \overline{\{y\}}$ or $y \notin \overline{\{x\}}$ and hence there is an open U of X that contains x but not y or contains y but not x. Using the basis of quasi-compact opens we may shrink U to be quasi-compact, so U and $X - U$ are disjoint opens in X^{cons} that separate x and y.

The analysis of topological operations with spectral spaces is very much simplified by means of:

Lemma A.2.4. *Let X be a quasi-compact and quasi-separated T_0 topological space such that the quasi-compact opens are a base for the topology.*

1. *The space X is a spectral space if and only if the Hausdorff space X^{cons} is quasi-compact.*
2. *A continuous map $f : X \to Y$ between two spectral spaces is spectral if and only if $f^{\mathrm{cons}} : X^{\mathrm{cons}} \to Y^{\mathrm{cons}}$ is continuous.*

The proof is very briefly sketched in [H, §2], and due to lack of a reference with a more complete discussion we provide the details for the convenience of the reader (since the proof requires some input from point-set topology that is not widely known).

Proof. Let us begin with (1). First assume that X^{cons} is quasi-compact. Pick an irreducible closed set $Z \subseteq X$. We seek a generic point. Since X^{cons} induces the contructible topology on Z, clearly Z^{cons} is closed in X^{cons} and hence it too is quasi-compact. We may therefore rename Z as X to reduce to the case when X is irreducible and we wish to find a generic point for X. If $x \in X$ is non-generic then there exists a non-empty quasi-compact open $U_x \subseteq X$ that does not contain x. Hence, if there is no generic point then we get a collection $\{U_x\}$ of non-empty quasi-compact opens in X such that $\cap_{x\in X}U_x = \emptyset$. The U_x's are closed in the quasi-compact topological space $X^{\rm cons}$, so by the finite intersection property for closed sets in quasi-compact spaces some finite intersection $U_{x_1}\cap\cdots\cap U_{x_n}$ must be empty. This contradicts the irreducibility of X (as all U_{x_i} are non-empty opens in X).

Conversely, suppose that X is spectral. To prove that $X^{\rm cons}$ must be quasi-compact we prove that it satisfies the finite intersection property for closed sets. Every closed set in $X^{\rm cons}$ is an intersection of constructible sets, and every constructible set is a finite union of overlaps $U\cap(X-U')$ for quasi-compact open U and U'. Hence, the quasi-compact opens and their complements form a subbasis of closed sets for the constructible topology. By the Alexander subbase theorem [Ke, Ch. 5, Thm. 6] (whose proof uses Zorn's Lemma), a topological space is quasi-compact if it satisfies the finite intersection property for members of a subbasis of closed sets. Hence, it is enough to show that if $\{C_i\}$ is a collection of subsets of X with each C_i either closed or quasi-compact open in X and if all finite intersections among the C_i's are non-empty then $\cap_i C_i \neq \emptyset$. By Zorn's Lemma we may and do enlarge $\{C_i\}$ to a maximal such collection (ignoring the property of whether or not the total intersection is non-empty). In particular, $\{C_i\}$ is stable under finite intersections among its quasi-compact open members and also among its closed members. Since X is quasi-compact and those C_i's that are closed satisfy the finite intersection property, their total intersection Z is non-empty. For any C_{i_0} that is a quasi-compact open, the overlaps $C_{i_0}\cap C_i$ for closed C_i satisfy the finite intersection property in the quasi-compact space C_{i_0} and hence the open C_{i_0} meets Z. Let us show that the non-empty Z is irreducible. Suppose $Z = Z_1\cup Z_2$ for closed subsets $Z_1, Z_2\subseteq Z$. If each Z_j fails to meet some C_{i_j} then C_{i_1} and C_{i_2} must be quasi-compact opens in X and so the member $C_{i_1}\cap C_{i_2}$ in the collection $\{C_i\}$ is a quasi-compact open that does not meet $Z_1\cup Z_2 = Z$, a contradiction. Thus, one of the closed sets Z_j meets every C_i and hence by maximality that Z_j is in the collection $\{C_i\}$. By construction of Z we thereby obtain $Z\subseteq Z_j$, so $Z_j = Z$ as desired. The spectral property of X provides a generic point z in the irreducible closed set Z, and since each quasi-compact open C_{i_0} meets Z it follows that every such C_{i_0} contains z. Thus, $z\in\cap_i C_i$. This shows that $X^{\rm cons}$ is indeed quasi-compact.

Now we turn to (2). Certainly if f is spectral then $f^{-1}(U)$ is a quasi-compact open in X for every quasi-compact open in Y, so $f^{\rm cons}$ is continuous. Conversely, assuming $f^{\rm cons}$ to be continuous we pick a quasi-compact open $U\subseteq Y$ and we want the open set $f^{-1}(U)\subseteq X$ to be quasi-compact. Since U is closed in $Y^{\rm cons}$ it follows from continuity of $f^{\rm cons}$ that $f^{-1}(U) = (f^{\rm cons})^{-1}(U)$ is closed in the space $X^{\rm cons}$ that is also quasi-compact since X is spectral. Hence, $f^{-1}(U)$ is a quasi-compact subset of $X^{\rm cons}$. But the open set $f^{-1}(U)$ in X is covered by quasi-compact opens in X, and this may be viewed as an open covering of $f^{-1}(U)$ in $X^{\rm cons}$. Hence, there is a finite subcover, so $f^{-1}(U)$ is a finite union of quasi-compact opens in X. Thus, $f^{-1}(U)$ is quasi-compact. ■

Example A.2.5. By the theory of formal models for open immersions [BL4, Cor. 5.4(a)], if $U\subseteq X$ is a quasi-compact admissible open in a quasi-compact and quasi-separated rigid space X then a cofinal system of formal (flat) models for U is given by an inverse system of opens in a cofinal system of formal (flat) models for X. The induced map ${\rm ZRS}(U)\to{\rm ZRS}(X)$ is thereby identified with an inverse limit of open embeddings, so it is an open embedding of topological spaces. Likewise, if $U'\subseteq X$ is another such open then so is $U\cap U'$ and clearly ${\rm ZRS}(U)\cap{\rm ZRS}(U') = {\rm ZRS}(U\cap U')$ inside of ${\rm ZRS}(X)$.

Since every closed point of a formal model arises as the specialization of a point on the rigid-analytic generic fiber, we see that if $\{U_i\}$ is a finite collection of quasi-compact admissible opens in a quasi-compact and quasi-separated rigid space X then the U_i's cover X if and only if the ${\rm ZRS}(U_i)$'s cover ${\rm ZRS}(X)$. Hence, working locally on X with quasi-compact admissible opens has the effect of allowing us to work locally on ${\rm ZRS}(X)$. By the same argument, a base of opens in ${\rm ZRS}(X)$ is given by ${\rm ZRS}(U)$'s for affinoid subdomains $U\subseteq X$.

Lemma A.2.6. *The full subcategory of spectral spaces in the category of topological spaces enjoys the following properties with respect to topological inverse limits:*

(1) *If $\{X_i\}$ is a directed inverse system of spectral spaces with spectral transition maps then the inverse limit space X is spectral and each map $X\to X_i$ is spectral. Moreover, $(\varprojlim X_i)^{\rm cons} = \varprojlim X_i^{\rm cons}$ as topological spaces.*

(2) *If $\{X_i\} \to \{Y_i\}$ is a map of such inverse systems with each $f_i : X_i \to Y_i$ a spectral map then the induced map $f : X \to Y$ on inverse limits is spectral. Moreover, if $\{F_i\}$ is an inverse system of pro-constructible* (*resp. closed*) *subsets of $\{X_i\}$ then the inverse limit F is pro-constructible* (*resp. closed*) *in X and $f(F)$ is the inverse limit of the $f_i(F_i) \subseteq Y_i$. In particular if each f_i is a surjective* (*resp. closed*) *map of topological spaces then so is f.*

Part (1) is [H, Thm. 7], and the proof we give for the entire lemma follows suggestions of Hochster.

Proof. We first analyze the formation of products of spectral spaces. If $\{X_\alpha\}$ is a collection of spectral spaces then we claim that $P = \prod X_\alpha$ is again a spectral space and that $P^{\rm cons} = \prod X_\alpha^{\rm cons}$ (in the sense that the constructible topology on the underlying set of P is the same as the product of the constructible topologies on the factor spaces X_α). Certainly P is a quasi-compact space, and since each X_α has a base of quasi-compact opens the same holds for P by the definition of the product topology. The T_0 property for P follows from the T_0 property for the factors X_α. Let us next check that P is quasi-separated. Any open in P is covered by opens of the form $\prod U_\alpha$ with each U_α a quasi-compact open in X_α and $U_\alpha = X_\alpha$ for all but finitely many α; such a $\prod U_\alpha$ shall be called a *basic quasi-compact open block*. Any quasi-compact open in P is covered by finitely many basic quasi-compact open blocks, and since an intersection of two such blocks is another such block (as each X_α is quasi-separated) we conclude that P is indeed quasi-separated. By Lemma A.2.4, the spectral property for P is now reduced to showing that $P^{\rm cons}$ is quasi-compact.

We will show directly that $P^{\rm cons} = \prod X_\alpha^{\rm cons}$, so by quasi-compactness of the $X_\alpha^{\rm cons}$'s (via Lemma A.2.4) we would get the desired quasi-compactness of $P^{\rm cons}$. The topology on $P^{\rm cons}$ has as a base of opens the sets $U \cap (P - U')$ for quasi-compact opens $U, U' \subseteq P$, and both U and U' are finite unions of basic quasi-compact open blocks. Thus, U is certainly open in $\prod X_\alpha^{\rm cons}$ and $P - U'$ is a finite intersection of complements $P - U'_i$ with $U'_i \subseteq P$ a basic quasi-compact open block $\prod_\alpha U_{\alpha,i}$. If we let $p_\alpha : P \to X_\alpha$ denote the projection then for each i the complement $P - U'_i$ is the union of the finitely many $p_\alpha^{-1}(X_\alpha - U_{\alpha,i})$'s for which the quasi-compact open $U_{\alpha,i} \subseteq X_\alpha$ is distinct from X_α, so $P - U'_i$ is open in $\prod X_\alpha^{\rm cons}$. Hence, every open in $P^{\rm cons}$ arises from an open in $\prod X_\alpha^{\rm cons}$. The converse is exactly the assertion that the map of spaces $P^{\rm cons} \to \prod X_\alpha^{\rm cons}$ is continuous, which is to say that each map $p_\alpha^{\rm cons} : P^{\rm cons} \to X_\alpha^{\rm cons}$ is continuous. Since $X_\alpha^{\rm cons}$ has a base of opens given by $U \cap (X_\alpha - U')$ for quasi-compact opens $U, U' \subseteq X_\alpha$ and both $p_\alpha^{-1}(U)$ and $P - p_\alpha^{-1}(U')$ are constructible in P, we are done with the treatment of products.

Turning our attention to directed inverse limits, we shall prove that $\varprojlim X_i$ is not only spectral but that as a subset of $\prod X_i$ it is closed in $(\prod X_i)^{\rm cons} = \prod X_i^{\rm cons}$. Note that by the definition of topological inverse limits, the induced topology on $\varprojlim X_i$ from $(\prod X_i)^{\rm cons} = \prod X_i^{\rm cons}$ is $\varprojlim X_i^{\rm cons}$; this latter topological inverse limit makes sense topologically because the transition maps $f_{ij} : X_j \to X_i$ are spectral and hence each $f_{ij}^{\rm cons}$ is continuous. Each $X_i^{\rm cons}$ is a quasi-compact Hausdorff space and hence the inverse limit of the $X_i^{\rm cons}$'s is indeed closed in the product $\prod X_i^{\rm cons}$. In particular, $\varprojlim X_i^{\rm cons}$ is quasi-compact and Hausdorff. It is clear that $\varprojlim X_i$ is a T_0-space (as it is a subspace of the product $\prod X_i$ of T_0-spaces), and it has a refined topology $\varprojlim X_i^{\rm cons}$ that is quasi-compact so it must be quasi-compact as well. For any i_0 the set-theoretic identification $\varprojlim X_i = \varprojlim_{i \geq i_0} X_i$ is a homeomorphism and so a base of opens of $\varprojlim X_i$ is given topologically by $\varprojlim_{i \geq i_0} U_i$ where $U_{i_0} \subseteq X_{i_0}$ is a quasi-compact open and $U_i = f_{i_0 i}^{-1}(U_{i_0})$ is a quasi-compact open in X_i for all $i \geq i_0$ (since $f_{i_0 i}$ is spectral). But a quasi-compact open in a spectral space is spectral, so $\{U_i\}_{i \geq i_0}$ is also a directed inverse system of spectral spaces with spectral transition maps, whence $U = \varprojlim_{i \geq i_0} U_i$ is quasi-compact. If $U' = \varprojlim_{i \geq i'_0} U'_i$ is another such open in $\varprojlim X_i$ and we pick $i_1 \geq i_0, i'_0$ and let $U''_i = U_i \cap U'_i$ for $i \geq i_1$ then $U \cap U' = \varprojlim_{i \geq i_1} U''_i$ inside of $\varprojlim X_i$. Hence, $\varprojlim X_i$ has a base of quasi-compact opens that is stable under finite intersection, so it is quasi-separated.

We have proved enough about the topology of $\varprojlim X_i$ so that $(\varprojlim X_i)^{\rm cons}$ makes sense. Thus, by Lemma A.2.4 the spectral property for the space $\varprojlim X_i$ and for the continuous maps $\varprojlim X_i \to X_{i_0}$ (for all i_0) will follow if the set-theoretic identification $(\varprojlim X_i)^{\rm cons} = \varprojlim X_i^{\rm cons}$ is a homeomorphism. It has been shown above that $\varprojlim X_i^{\rm cons}$ is the topology induced on $\varprojlim X_i$ by $\prod X_i^{\rm cons} = (\prod X_i)^{\rm cons}$, so we just have to show that the constructible topology on $\varprojlim X_i$ is also induced by $(\prod X_i)^{\rm cons}$. By directedness of the indexing set and the spectral property for the transition maps, it is clear that any basic quasi-compact open block in $\prod X_i$ meets $\varprojlim X_i$ in a quasi-compact open set, and so any constructible set in $\prod X_i$ meets $\varprojlim X_i$ in a constructible set. That is, $(\varprojlim X_i)^{\rm cons} \to (\prod X_i)^{\rm cons}$ is continuous. To see that it is an embedding, we just have to show that every constructible set in $\varprojlim X_i$ is a pullback of a constructible set in $\prod X_i$, and for this it suffices to consider quasi-compact opens. But any quasi-compact open in $\varprojlim X_i$ is trivially of the form $\varprojlim_{i \geq i_0} U_i$ considered above, and so is the pullback of the basic quasi-compact open block in $\prod X_i$ given by U_{i_0} in the i_0-factor and X_i in the i-factor for all $i \neq i_0$. This completes the proof of (1).

For the first assertion in (2), the induced map $f = \varprojlim f_i : X \to Y$ is certainly continuous and hence (by Lemma A.2.4) is spectral if and only if $f^{\rm cons}$ is continuous. But the preceding considerations show that $f^{\rm cons} = \varprojlim f_i^{\rm cons}$, and each $f_i^{\rm cons}$ is continuous since each f_i is spectral, so $f^{\rm cons}$ is indeed continuous. Since pro-constructible sets in a spectral space are precisely the closed sets in the associated constructible topology, if $\{F_i\}$ is an inverse system of pro-constructible sets then the subset $F = \varprojlim F_i$ in $X = \varprojlim X_i$ is an inverse limit of closed sets in $\varprojlim X_i^{\rm cons} = X^{\rm cons}$ so it is closed in this space since the $X_i^{\rm cons}$'s are quasi-compact Hausdorff spaces (with continuous transition maps between them). Thus, F is indeed pro-constructible in X for such $\{F_i\}$. This argument also shows that $f(F) \subseteq Y$ is pro-constructible because $f(F) = f^{\rm cons}(F^{\rm cons})$ inside of $Y^{\rm cons} = \varprojlim Y_i^{\rm cons}$ (with $F^{\rm cons}$ denoting F viewed inside of $X^{\rm cons} = \varprojlim X_i^{\rm cons}$) and $f^{\rm cons}$ is a continuous map between quasi-compact Hausdorff spaces (so it is closed). Moreover, $f(F)$ is the inverse limit of the $f_i(F_i)$ (as subsets of Y) because upon passing to the constructible topologies we reduce to the well-known analogous claim for a continuous map between inverse systems of quasi-compact Hausdorff spaces (see [B, I, §9.6, Cor. 2]). Since closed sets in each X_i are trivially pro-constructible, the same argument shows the set-theoretic fact that if the F_i's are closed in X then $f(F)$ is the inverse limit of the $f_i(F_i)$'s in Y. In this special case the subset $F \subseteq X$ is closed because $X - F$ is the union of the overlap of $X \subseteq \prod X_i$ with the open blocks given by $(X_{i_0} - F_{i_0}) \times \prod_{i \neq i_0} X_i$ for all i_0.

By taking $F_i = X_i$ for all i, we conclude that if $f_i(X_i) = Y_i$ for all i then $f(X) = Y$; that is, f is surjective if all f_i's are surjective. As for the property that $f(F)$ is closed in Y whenever $F \subseteq X$ is closed and each f_i is closed, we note that *if* $F = \varprojlim F_i$ with $\{F_i\}$ an inverse system of closed sets in $\{X_i\}$ then $f(F) = \varprojlim f_i(F_i)$ is an inverse limit of closed sets in the Y_i's and hence is indeed closed in Y. Thus, the preservation of closedness for morphisms reduces to the claim that any closed set F in $X = \varprojlim X_i$ has the form $\varprojlim F_i$ with $F_i \subseteq X_i$ a closed set. This is true in the setting of arbitrary topological spaces, as follows. An arbitrary intersection of closed sets of the form $\varprojlim F_i$ with closed $F_i \subseteq X_i$ again has this special form, so it suffices to verify our claim for closed sets complementary to members of a base of opens. A base of opens is given by $\varprojlim_{i \geq i_0} f_{i_0 i}^{-1}(U_{i_0})$ with $U_{i_0} \subseteq X_{i_0}$ an open set (and $f_{i_0 i} : X_i \to X_{i_0}$ the spectral transition map), and the complement of such an open has the form $\varprojlim_{i \geq i_0} F_i$ with $F_i = f_{i_0 i}^{-1}(X_{i_0} - U_{i_0})$ for $i \geq i_0$. Defining $F_i = X_i$ for all other i settles the claim. ■

Let $\mathfrak{I}$ be an ideal of definition for the valuation ring R of our non-archimedean base field k. Fix a k-affinoid algebra A, and let $\mathscr{A}$ be a flat formal affine model (*i.e.*, $\mathscr{A}$ is topologically finitely presented and flat over R, with $k \otimes_R \mathscr{A} \simeq A$). A key fact is that the ring extension $\mathscr{A} \subseteq A^0$ into the subring of power-bounded elements is integral. To prove this, we shall exhibit a subring of $\mathscr{A}$ over which A^0 is integral. Let $d = \dim(\mathscr{A}/\mathfrak{m}_R\mathscr{A}) \geq 0$. By [C4, Thm. A.2.1(1)], $d = \dim(A)$. By Noether normalization over the residue field $\widetilde{k}$, there is a finite map $\varphi : \widetilde{k}[T_1, \dots, T_d] \to \mathscr{A}/\mathfrak{m}_R\mathscr{A}$. For an ideal of definition $\mathfrak{I}$ of R it follows that any lifting of $\operatorname{Spec} \varphi$ to a map $\operatorname{Spec}(\mathscr{A}/\mathfrak{I}\mathscr{A}) \to \operatorname{Spec}((R/\mathfrak{I})[T_1, \dots, T_d])$ between finitely presented $R/\mathfrak{I}$-schemes is proper and quasi-finite, hence finite. Thus, any continuous lift $\Phi : R\{\{T_1, \dots, T_d\}\} \to \mathscr{A}$ of φ over R is finite. Such a map of flat R-algebras must be injective because on generic fibers it is a finite map $\Phi_k : k\langle\langle T_1, \dots, T_d\rangle\rangle \to A$ with $d = \dim A$. The finite map Φ_k between k-affinoids induces an integral map on subrings of power-bounded elements [BGR, 6.3.5/1], but the power-bounded elements of the d-variable Tate algebra are precisely the d-variable restricted power series over R. The R-algebra of such power series is a subalgebra of $\mathscr{A}$ inside of A, so we conclude that A^0 is indeed integral over $\mathscr{A}$.

By [vdPS, Thm. 2.4], the points in $\operatorname{ZRS}(\operatorname{Sp}(A))$ are functorially in bijective correspondence with (not necessarily rank-1) valuations rings V on fraction fields $\operatorname{Frac}(A/\mathfrak{p})$ for primes $\mathfrak{p}$ of A such that the subring A^0 of power-bounded elements lands in V and the (necessarily nonzero) ideal $\mathfrak{I}V$ of V generated by $\mathfrak{I}$ is topologically nilpotent (*i.e.*, $\cap_{n\ge 1}(\mathfrak{I}V)^n = \cap_{n\ge 1}\mathfrak{I}^n V$ vanishes). Alternatively, and more conveniently for our purposes, since $\mathscr{A} \to A^0$ is an integral ring extension we can identify such points with maps $\mathscr{A} \to V$ to valuation rings V such that (i) $\mathfrak{I}$ generates a *nonzero* proper ideal of V that is topologically nilpotent, and (ii) $\mathrm{Frac}(V)$ is generated by the image of $\mathscr{A}$ (or equivalently, of A).

Concretely, given *any* map $\mathscr{A} \to V$ to a valuation ring such that (i) holds it is straightforward to check that the $\mathfrak{I}$-adic completion $\widehat{V}$ of V *is* a valuation ring and it thereby determines a point of $\mathrm{ZRS}(\mathrm{Sp}(A))$ since (by principality of finitely generated ideals in a valuation ring) one can uniquely lift the map of formal schemes $\mathrm{Spf}(\widehat{V}) \to \mathrm{Spf}(\mathscr{A})$ through admissible formal blow-ups (and so chasing the image of the closed point of $\mathrm{Spf}(\widehat{V})$ gives the desired point in $\mathrm{ZRS}(\mathrm{Sp}(A))$). Using the induced valuation ring structure on the fraction field of the image of $\mathscr{A}$ in V gives the valuation associated to this point. In particular, via the theory of rig-points [BL3, 3.5], points of $\mathrm{Sp}(A)$ give rise to points in the associated Zariski-Riemann space; likewise, if X is a quasi-compact and quasi-separated rigid space then the underlying set of X is *functorially* a subset of its associated Zariski-Riemann space. (Note that X is empty if and only if $\mathrm{ZRS}(X)$ is empty.)

Lemma A.2.7. *Any pair of faithfully flat local maps $W \rightrightarrows V, V'$ of valuation rings can be completed to a commutative square of valuation rings and faithfully flat local maps.*

Recall that a local map from a valuation ring to a domain is faithfully flat if and only if it is injective (as all finitely generated ideals in a valuation ring are principal).

Proof. Pick $x \in \mathrm{Spec}(V \otimes_W V')$ over the closed points of $\mathrm{Spec}(V)$ and $\mathrm{Spec}(V')$, so the maps $V, V' \rightrightarrows \mathscr{O}_{V\otimes_W V',x}$ are local and flat, hence faithfully flat. Let $\mathfrak{p}$ be a minimal prime of the local ring at x, so by going-down for flat maps the two local maps $V, V' \rightrightarrows \mathscr{O}_{V\otimes_W V',x}/\mathfrak{p}$ are injective and hence faithfully flat. Thus, any valuation ring dominating $\mathscr{O}_{V\otimes_W V',x}/\mathfrak{p}$ does the job. ∎

Lemma A.2.8. *Let $X, Y \rightrightarrows Z$ be a pair of maps between quasi-compact and quasi-separated rigid spaces, and let $P = X \times_Z Y$ so P is also quasi-compact and quasi-separated. The natural map*

$$\mathrm{ZRS}(P) \to \mathrm{ZRS}(X) \times_{\mathrm{ZRS}(Z)} \mathrm{ZRS}(Y) \tag{A.2.1}$$

is surjective.

Proof. By Example A.2.5 it is enough to consider the affinoid case, say with $X = \mathrm{Sp}(A)$, $Y = \mathrm{Sp}(B)$, and $Z = \mathrm{Sp}(C)$, so $P = \mathrm{Sp}(D)$ where $D = A\widehat{\otimes}_C B$. Let $\mathscr{A}$, $\mathscr{B}$, and $\mathscr{C}$ be flat affine formal models for A, B, and C respectively, equipped with continuous R-algebra maps $\mathscr{C} \rightrightarrows \mathscr{A}, \mathscr{B}$ inducing $C \rightrightarrows A, B$. Let $\mathscr{D}$ be the quotient of $\mathscr{A}\widehat{\otimes}_{\mathscr{C}}\mathscr{B}$ by R-torsion (so $\mathscr{D}$ is a flat affine formal model for D). A point in the target of (A.2.1) is induced by a compatible triple of maps to valuation rings $\mathscr{A} \to V$, $\mathscr{B} \to V'$, and $\mathscr{C} \to W$ (with local faithfully flat maps $W \rightrightarrows V, V'$). By Lemma A.2.7 we can find a valuation ring V'' equipped with a map $V \otimes_W V' \to V''$ such that the maps $V, V' \rightrightarrows V''$ are local and faithfully flat. Hence, the $\mathfrak{I}$-adically completed tensor product $\mathscr{D}$ maps to the $\mathfrak{I}$-adic completion of V'' (which is again a valuation ring, as we noted above) and this determines the desired point of $\mathrm{ZRS}(P)$. ∎

Lemma A.2.9. *If $f : X \to Y$ is a map of quasi-compact and quasi-separated rigid spaces then the following are equivalent:*

- *The map f is surjective.*
- *Every formal model $\mathfrak{f} : \mathfrak{X} \to \mathfrak{Y}$ of f (using R-flat formal models of X and Y) is surjective.*
- *The map $\mathrm{ZRS}(f)$ is surjective.*

Proof. First assume f is surjective, and let $\mathfrak{f} : \mathfrak{X} \to \mathfrak{Y}$ be a formal model with $\mathfrak{X}$ and $\mathfrak{Y}$ flat over R. On topological spaces $\mathfrak{f}$ coincides with the map $\mathfrak{f}_{\mathrm{red}}$ of ordinary finite type $\widetilde{k}$-schemes, and so it is surjective if and only if it is surjective on underlying spaces of closed points. For any closed point $y_0 \in \mathfrak{Y}$ the R-flatness ensures (via the theory of rig-points) that there exists a finite extension k'/k and a map $\mathfrak{y} : \mathrm{Spf}(R') \to \mathfrak{Y}$ over $\mathrm{Spf}(R)$ that hits y_0. If y_0 is not hit by $\mathfrak{f}$ then the pullback of $\mathfrak{f}$ by $\mathfrak{y}$ is empty. However, this pullback is a topologically finitely presented (possibly non-flat) formal scheme over R' whose generic fiber over $\mathrm{Sp}(k')$

is $f^{-1}(y)$ with $y \in Y = \mathfrak{Y}^{\rm rig}$ the image of $\mathfrak{y}^{\rm rig}$. Since f is surjective, the fiber $f^{-1}(y)$ cannot be empty and so we have a contradiction. Thus, $\mathfrak{f}$ is indeed surjective.

If all formal models for f are surjective then the map $\mathrm{ZRS}(f)$ can be expressed as an inverse limit of surjective spectral maps, and so surjectivity of $\mathrm{ZRS}(f)$ follows from Lemma A.2.6 in such cases.

Finally, assume $\mathrm{ZRS}(f)$ is surjective and pick $y \in Y \subseteq \mathrm{ZRS}(Y)$. Identify y with a map $y : \mathrm{Sp}(k') \to Y$ for a finite extension k'/k. We want to prove that the fiber product $\mathrm{Sp}(k') \times_Y X$ is non-empty. It suffices to show that its associated Zariski-Riemann space is non-empty, and by Lemma A.2.8 the natural map

$$\mathrm{ZRS}(\mathrm{Sp}(k') \times_Y X) \to \mathrm{ZRS}(\mathrm{Sp}(k')) \times_{\mathrm{ZRS}(Y)} \mathrm{ZRS}(X) \tag{A.2.2}$$

is surjective. But $\mathrm{ZRS}(\mathrm{Sp}(k'))$ is trivially a one-point space $\{\xi\}$, and so the topological target fiber product in (A.2.2) is exactly the fiber of $\mathrm{ZRS}(f)$ over the image of ξ in $\mathrm{ZRS}(Y)$. Hence, surjectivity of $\mathrm{ZRS}(f)$ gives the desired non-emptiness. ■

Here is the key definition.

Definition A.2.10. Let X be a quasi-compact and quasi-separated rigid space. An open subset $\mathscr{U} \subseteq \mathrm{ZRS}(X)$ is *admissible* if for every map of quasi-compact and quasi-separated rigid spaces $f : Y \to X$, the image of $\mathrm{ZRS}(f)$ is contained in $\mathscr{U}$ whenever the subset $f(Y) \subseteq X \subseteq \mathrm{ZRS}(X)$ is contained in $\mathscr{U}$. (It clearly suffices to work with affinoid Y.) Given such a $\mathscr{U}$, we call the subset $U = \mathscr{U} \cap X$ its set of *ordinary points*.

Remark A.2.11. If $\mathscr{U} \subseteq \mathrm{ZRS}(X)$ is an admissible open then $\mathrm{ZRS}(f)^{-1}(\mathscr{U}) \subseteq \mathrm{ZRS}(Y)$ is as well (for any $f : Y \to X$ as in Definition A.2.10).

It is clear that if X is a quasi-compact and quasi-separated rigid space and $U \subseteq X$ is a quasi-compact admissible open in the sense of Tate then $\mathrm{ZRS}(U) \cap X$ inside of $\mathrm{ZRS}(X)$ is equal to $U \subseteq X$ and so the open set $\mathrm{ZRS}(U)$ in $\mathrm{ZRS}(X)$ is admissible. In general, if $\mathscr{U} \subseteq \mathrm{ZRS}(X)$ is an admissible open in the above sense then the associated locus $U \subseteq X$ of ordinary points is an admissible open of X in the sense of Tate. Indeed, we may choose admissible affinoid opens $U_i \subseteq X$ such that the associated open sets $\mathscr{U}_i = \mathrm{ZRS}(U_i) \subseteq \mathrm{ZRS}(X)$ are an open cover of $\mathscr{U}$ (so obviously $\cup U_i = U$ inside of X) and we just have to check that for any (necessarily quasi-compact) morphism $f : Y = \mathrm{Sp}(B) \to X$ from an affinoid space such that $f(Y) \subseteq U$, the set-theoretic cover of Y by the quasi-compact pullbacks $f^{-1}(U_i)$ has a finite subcover. But by definition of admissibility for $\mathscr{U}$ the map $\mathrm{ZRS}(f)$ has image contained in $\mathscr{U}$ and hence the preimages $\mathrm{ZRS}(f)^{-1}(\mathscr{U}_i)$ are an open cover of the space $\mathrm{ZRS}(Y)$ that is quasi-compact. It follows that $\mathrm{ZRS}(f)$ has image contained in the union of finitely many $\mathscr{U}_i$, whence $f(Y) \subseteq X$ is contained in the union of the finitely many corresponding loci $U_i = \mathscr{U}_i \cap X$ as required. This can be strengthened as follows:

Lemma A.2.12. *Let X be a quasi-compact and quasi-separated rigid space. The association $\mathscr{U} \mapsto \mathscr{U} \cap X$ from admissible opens in $\mathrm{ZRS}(X)$ to admissible opens in X is a bijection that commutes with the formation of intersections. Moreover, $\mathscr{U}$ is quasi-compact if and only if the admissible open $\mathscr{U} \cap X$ in X is a quasi-compact rigid space, and the correspondence $\mathscr{U} \mapsto \mathscr{U} \cap X$ commutes with formation of preimages under $\mathrm{ZRS}(f)$ for any map $f : X' \to X$ between quasi-compact and quasi-separated rigid spaces.*

Proof. Since $\mathscr{U}$ is covered by opens of the form $\mathrm{ZRS}(U)$ for quasi-compact admissible opens $U \subseteq X$, to prove that the admissible open $\mathscr{U} \cap X$ determines $\mathscr{U}$ it suffices to note the obvious fact that for any quasi-compact admissible open $U \subseteq X$ we have $\mathrm{ZRS}(U) \subseteq \mathscr{U}$ if and only if $U \subseteq \mathscr{U} \cap X$ (since $\mathscr{U}$ is admissible, $U = X \cap \mathrm{ZRS}(U)$, and $\mathrm{ZRS}(\cdot)$ is a functor).

Now let $U \subseteq X$ be an arbitrary admissible open, say with $\{U_i\}$ an admissible covering by quasi-compact opens. Let $\mathscr{U}$ be the open set $\cup \mathrm{ZRS}(U_i)$ in $\mathrm{ZRS}(X)$, so $\mathscr{U} \cap X = U$. We claim that $\mathscr{U}$ is admissible. Pick a map of rigid spaces $f : Y \to X$ with Y quasi-compact and quasi-separated such that $f(Y) \subseteq U$. We need to prove that $\mathrm{ZRS}(f)$ has image contained in $\mathscr{U}$. By the definition of admissibility for the covering $\{U_i\}$ of U, the loci $f^{-1}(U_i)$ in Y are admissible opens and constitute an admissible cover. In particular, there is a finite collection of affinoid domains $\{V_j\}$ in Y that covers Y and refines $\{f^{-1}(U_i)\}$. Since an admissible covering by finitely many quasi-compact opens can always be realized from a Zariski-open covering of a suitable formal model [BL3, Lemma 4.4], $\mathrm{ZRS}(Y)$ is the union of the $\mathrm{ZRS}(V_j)$'s. Thus, the image of $\mathrm{ZRS}(f)$ is the union of the images of the ZRS(f_j)'s, with $f_j = f|_{V_j} : V_j \to X$ a map that factors through some $U_{i(j)}$. Hence, ZRS(f_j) has image contained in ZRS$(U_{i(j)}) \subseteq \mathscr{U}$, so ZRS$(f)$ has image contained in $\mathscr{U}$. This concludes the proof that $\mathscr{U}$ is an admissible open in ZRS(X).

Finally, we check that an admissible open $\mathscr{U} \subseteq$ ZRS(X) is quasi-compact if and only if the admissible open $U = \mathscr{U} \cap X$ in X is quasi-compact as a rigid space, and that the correspondence between admissible opens in X and ZRS(X) is compatible with preimages. The preceding argument shows that if a collection of quasi-compact opens $U_i \subseteq U$ is an admissible covering of U then the ZRS(U_i)'s cover $\mathscr{U}$, and the converse is immediate from the hypothesis of admissibility for $\mathscr{U}$ and the quasi-compactness of Zariski-Riemann spaces. Thus, the desired quasi-compactness result follows. As for preimages, if $f : X' \to X$ is a map between quasi-compact and quasi-separated rigid spaces and $\mathscr{U} \subseteq$ ZRS(X) is an admissible open then for $U = \mathscr{U} \cap X$ we have to check that $f^{-1}(U) = \text{ZRS}(f)^{-1}(\mathscr{U}) \cap X'$. The containment $\subseteq$ is obvious by admissibility of $\mathscr{U}$ and functoriality of ZRS$(\cdot)$ (applied to admissible quasi-compact opens in U). For the reverse inclusion consider $x' \in X'$ such that ZRS$(f)(x') \in \mathscr{U}$. Since ZRS$(f)(x') = f(x')$ in $X \subseteq$ ZRS(X) we have $f(x') \in \mathscr{U} \cap X = U$ as desired. ■

Any finite union U of affinoid subdomains U_i in a quasi-compact and quasi-separated rigid space is an admissible open with the U_i's as an admissible covering, so it follows that retrocompact opens in ZRS(X) are necessarily admissible.

Lemma A.2.13. *If $f : X' \to X$ is a surjective map of quasi-compact and quasi-separated rigid spaces and $\mathscr{U}$ is an open subset of* ZRS(X) *whose open preimage $\mathscr{U}' \subseteq$* ZRS(X') *is admissible then $\mathscr{U}$ is admissible.*

Proof. Let Y be a quasi-compact and quasi-separated rigid space and $h : Y \to X$ a map such that $h(Y) \subseteq \mathscr{U}$. We want to prove that ZRS(h) has image contained in $\mathscr{U}$. The pullback $f' : Y' = X' \times_X Y \to Y$ is surjective, so by Lemma A.2.9 the map ZRS(f') is surjective. Hence, we may replace Y with Y' so that h factors as $f \circ h'$ for some $h' : Y \to X'$. Obviously $h'(Y) \subseteq \mathscr{U}'$, so by admissibility of $\mathscr{U}'$ the image of ZRS(h') is contained in $\mathscr{U}'$. Composing with ZRS(f) gives that ZRS(h) has image contained in $\mathscr{U}$. ■

Lemma A.2.14. *If $f : X' \to X$ is a proper map of quasi-compact and quasi-separated rigid spaces then* ZRS(f) *is a closed map of topological spaces. Moreover, if f is surjective and $\mathscr{U} \subseteq$* ZRS(X) *is a subset whose preimage in* ZRS(X') *is open* (*resp. admissible open, resp. quasi-compact open*) *then the same holds for $\mathscr{U}$ in* ZRS(X).

Proof. By Lemma A.2.6, ZRS(f) is closed provided that any formal model for f is a closed map. But (as we explained in [C3, §A.1]), by recent work of Temkin the map f is proper in the sense of rigid spaces if and only if one (equivalently every) formal model of f is proper (and thus closed) in the sense of formal geometry.

Now assume that f is also surjective. Any closed surjection of topological spaces is a quotient map, so a subset $\mathscr{U} \subseteq$ ZRS(X) is open (resp. quasi-compact open) if its preimage in ZRS(X') has this property. If ZRS$(f)^{-1}(\mathscr{U})$ is an admissible open in ZRS(X') then $\mathscr{U}$ must at least be open in ZRS(X) and it is admissible by Lemma A.2.13. ■

Now we can prove Theorem A.2.1:

Proof. Let $P = X' \times_X X'$ and let $\mathscr{U}' \subseteq$ ZRS(X') be the admissible open that corresponds to U' via Lemma A.2.12. Let $p_1, p_2 : P \rightrightarrows X'$ be the canonical projections. By the definition of U' as a preimage from X, the two admissible open preimages $p_j^{-1}(U')$ in P coincide, hence they correspond to the same admissible open set in ZRS(P). But the final part of Lemma A.2.12 ensures that $p_j^{-1}(U')$ corresponds to ZRS$(p_j)^{-1}(\mathscr{U}')$, so these latter two opens in ZRS(P) coincide. By Lemma A.2.8, it follows that $\mathscr{U}'$ is the preimage of a subset $\mathscr{U}$ of ZRS(X). By Lemma A.2.14, $\mathscr{U}$ is an admissible open in ZRS(X). Its associated locus $\mathscr{U} \cap X$ of ordinary points is therefore an admissible open in X by Lemma A.2.12. Since the correspondence between admissible opens in X and ZRS(X) has been shown to be compatible with formation of preimages, we conclude that the admissible open $\mathscr{U} \cap X$ in X has preimage $\mathscr{U}' \cap X' = U'$ in X' and hence it is equal to the image U of U' in X. Thus, U is indeed an admissible open in X. ■

A.3. **Weil-pairings and formal semi-abelian models.** Let k be a non-archimedean field with valuation ring R and let $A_{/k}$ be an abelian variety with semistable reduction over R. Let $\mathfrak{A}_R$ and $\mathfrak{A}'_R$ be the associated formal semi-abelian models for A and $A^\vee$ over $\mathrm{Spf}(R)$, and let

$$0 \to \mathfrak{T} \to \mathfrak{A}_R \to \mathfrak{B} \to 0, \quad 0 \to \mathfrak{T}' \to \mathfrak{A}'_R \to \mathfrak{B}' \to 0$$

be the filtrations with maximal formal subtori and formal abelian scheme quotients as in the general semi-stable reduction theorem (Theorem 2.1.9). In particular, there are unique abelian schemes B_R and B'_R over $\mathrm{Spec}(R)$ that algebraize $\mathfrak{B}$ and $\mathfrak{B}'$, and we let B and B' denote their respective generic fibers over k. The proof of Theorem 2.1.9 provides a canonical isomorphism $B' \simeq B^\vee$ (or equivalently, $B'_R \simeq B_R^\vee$ or $\mathfrak{B}' \simeq \mathfrak{B}^\vee$).

In the discretely-valued case, note that by Serre's criterion the Néron models $\mathrm{N}(A)$ and $\mathrm{N}(A^\vee)$ of A and $A^\vee$ over R must have semi-stable reduction and by Example 2.1.10 the formal semi-abelian models $\mathfrak{A}_R$ and $\mathfrak{A}'_R$ coincide with the respective $\mathfrak{m}_R$-adic completions of the relative identity components $\mathrm{N}(A)^0$ and $\mathrm{N}(A^\vee)^0$. Grothendieck [SGA7, IX, 5.2, 7.1.5, 7.4] proved that in the discretely-valued case the finite flat k-group $\mathfrak{T}[N]_k$ (resp. $\mathfrak{T}'[N]_k$) is orthogonal to $\mathfrak{A}'_R[N]_k$ (resp. $\mathfrak{A}_R[N]_k$) when using the Weil-pairing $A[N] \times A^\vee[N] \to \mu_N$ for every positive integer N, and he used the theory of bi-extensions to construct a canonical isomorphism $B'_R \simeq B_R^\vee$ with respect to which the pairing between $B'[N] \simeq (\mathfrak{A}'_R[N]/\mathfrak{T}'[N])_k$ and $B[N] \simeq (\mathfrak{A}_R[N]/\mathfrak{T}[N])_k$ induced by the Weil-pairing $A[N] \times A^\vee[N] \to \mu_N$ is precisely the canonical Weil-pairing between $B[N]$ and $B^\vee[N]$ for every $N \geq 1$. This condition for all N (or even just N running through powers of a fixed prime) uniquely characterizes Grothendieck's isomorphism $B' \simeq B^\vee$ without mentioning the theory of bi-extensions. The proof of the duality aspect of Theorem 4.1.1 rests on an analogue of these results in the setting of the general semistable reduction theorem without discreteness restrictions on the absolute value. The required analogous result was recorded without proof as Theorem 4.1.6, and here we give the statement and proof of a slightly more general result:

Theorem A.3.1. *With notation as above, for every positive integer N the Weil pairing $A[N] \times A^\vee[N] \to \mu_N$ makes $\mathfrak{T}[N]_k$ annihilate $\mathfrak{A}'_R[N]_k$ and $\mathfrak{A}_R[N]_k$ annihilate $\mathfrak{T}'[N]_k$, and the resulting pairing*

$$B[N] \times B'[N] = \mathfrak{B}[N]_k \times \mathfrak{B}'[N]_k \simeq (\mathfrak{A}_R[N]_k/\mathfrak{T}[N]_k) \times (\mathfrak{A}'_R[N]_k/\mathfrak{T}'[N]_k) \to \mu_N \tag{A.3.1}$$

induced by the Weil pairing between $A[N]$ and $A^\vee[N]$ arises from the canonical isomorphism $B'_R \simeq B_R^\vee$ via the Weil pairing $B[N] \times B^\vee[N] \to \mu_N$.

The key point is that the isomorphism $B'_R \simeq B_R^\vee$ is provided by an explicit construction in the *proof* of Theorem 2.1.9 and not through an abstract recipe such as the algebraic theory of bi-extensions that is used in the discretely-valued case (and which we do not have in the rigid-analytic setting). Since Theorem A.3.1 gives an abstract unique characterization of the isomorphism $B'_R \simeq B_R^\vee$ that emerges from the explicit rigid-analytic constructions in the proof of Theorem 2.1.9, in the discretely-valued case we conclude (using Grothendieck's results) that the isomorphism $B'_R \simeq B_R^\vee$ constructed via rigid geometry in [BL2] coincides with the one that is provided by Grothendieck's work with bi-extensions. We emphasize that it is the duality between $\mathfrak{B}'$ and $\mathfrak{B}$ constructed via rigid geometry that is relevant in the theory of canonical subgroups, and so one cannot avoid relating this specific duality with the duality between torsion-levels of A and $A^\vee$ in the study of how duality interacts with canonical subgroups.

Fortunately, bi-extensions are irrelevant in the proof of Theorem A.3.1. The proof requires nothing more than carefully unwinding the rigid-analytic construction of the Poincaré bundle P_A on $A \times A^\vee$ in terms of the formal Poincaré bundle $P_{\mathfrak{B}}$ on $\mathfrak{B} \times \mathfrak{B}^\vee$ in the *proof* of Theorem 2.1.9, and tracking the construction of each Weil pairing $A[N] \times A^\vee[N] \to \mu_N$ in terms of P_A (as in [Mum, §20]) so that we can understand how it restricts to $\mathfrak{A}_R[N]_k \times \mathfrak{A}'_R[N]_k \subseteq A[N] \times A^\vee[N]$.

Proof. [1] ■ 1

[1] Must work out!

References

[AM] A. Abbes, A. Mokrane, *Sous-groupes canoniques et cycles évanescents p-adiques pour les variétes abéliennes*, Publ. Math. IHES, **99** (2004), pp. 117–162.

[AG] F. Andreatta, C. Gasbarri, *The canonical subgroup for families of abelian varieties*, preprint, 2005.

[Ber1] V. Berkovich, *Spectral theory and analytic geometry over non-Archimedean fields*, Mathematical Surveys and Monographs, vol. 33, American Mathematical Society, 1990.

[Ber2] V. Berkovich, *Étale cohomology for non-Archimedean analytic spaces*, Publ. Math. IHES, **78** (1993), pp. 7–161.

[BG] S. Bosch, U. Görtz, *Coherent modules and their descent on relative rigid spaces*, J. Reine Angew. Math., **495** (1998), pp. 119–134.

[BGR] S. Bosch, U. Günzter, R. Remmert, *Non-Archimedean analysis*, Springer-Verlag, 1984.

[BL1] S. Bosch, W. Lütkebohmert, *Stable reduction and uniformization of abelian varieties* I, Math. Annalen, **270** (1985), pp. 349–379.

[BL2] S. Bosch, W. Lütkebohmert, *Degenerating abelian varieties*, Topology, **30** no. 4 (1991), pp. 653–698.

[BL3] S. Bosch, W. Lütkebohmert, *Formal and rigid geometry* I, Math. Annalen, **295** (1993), pp. 291–317.

[BL4] S. Bosch, W. Lütkebohmert, *Formal and rigid geometry* II, Math. Annalen, **296** (1993), pp. 403–429.

[BLR] S. Bosch, W. Lütkebohmert, M. Raynaud, *Néron Models*, Springer-Verlag, 1990.

[B] N. Bourbaki, *General topology*, Springer-Verlag, Berlin, 1989.

[Bu] K. Buzzard, *Analytic continuation of overconvergent eigenforms*, Journal of the Amer. Math. Soc., **16** (2003), pp. 29–55.

[CF] C-L. Chai, G. Faltings, *Degeneration of abelian varieties*, Springer-Verlag, 1990.

[C1] B. Conrad, *Irreducible components of rigid spaces*, Annales de L'institut Fourier, **49** (1999), pp. 473–541.

[C2] B. Conrad, *Arithmetic moduli of generalized elliptic curves*, to appear in Journal of the Inst. of Math. of Jussieu.

[C3] B. Conrad, *Relative ampleness in rigid-analytic geometry*, submitted, 2005.

[C4] B. Conrad, *Modular curves and rigid-analytic spaces*, submitted 2005.

[EGA] J. Dieudonné, A. Grothendieck, *Éléments de géométrie algébrique*, Publ. Math. IHES, **4, 8, 11, 17, 20, 24, 28, 32**, (1960-7).

[F] G. Faltings, "Finiteness theorems for abelian varieties over number fields," Ch. II in *Arithmetic geometry* (Cornell/Silverman, ed.), Springer-Verlag, 1986.

[Fo] J.-M. Fontaine, *Groupes p-divisible sur les corps locaux*, Astérisque **47-48**, Soc. Math. de France, 1977.

[FK] E. Freitag, R. Kiehl, *Étale cohomology and the Weil conjectures*, Ergebnisse der Math. und ihrer Grenzgebiete **13**, Springer-Verlag, Berlin, 1988.

[Go] R. Godement, *Théorie des faiseaux*, Hermann, Paris, 1973.

[GK] E. Goren, P. Kassaei, *A "subgroup-free" approach to canonical subgroups*, submitted, 2004.

[G] F. Gouvea, *Arithmetic of p-adic modular forms*, Springer Lecture Notes in Mathematics **1304** (1988).

[SGA1] A. Grothendieck, *Revêtements Etales et Groupe Fondamental*, Lecture Notes in Math. **224**, Springer-Verlag, NY, 1971.

[SGA3] A. Grothendieck, *Schémas en groupes* I, II, Springer Lecture Notes in Mathematics **151**, **152**, Springer-Verlag, New York (1970).

[SGA7] A. Grothendieck, *Groupes de monodromie en géométrie algébrique*, Springer Lecture Notes in Math **288**, Springer-Verlag, New York (1972).

[GR] L. Gruson, M. Raynaud, *Critères de platitude et de projectivité*, Inv. Math., **13** (1971), pp. 1–89.

[H] M. Hochster, *Prime ideal structure in commutative rings*, Transactions of the Amer. Math. Soc., **142** (1969), pp. 43–60.

[Ill] L. Illusie, *Déformations de groupes de Barsotti–Tate* (*d'après A. Grothendieck*), in "Seminar on arithmetic bundles: the Mordell conjecture", Asterisque, **127** (1985), pp. 151–198.

[deJ] A. J. de Jong, *Crystalline Dieudonné module theory via formal and rigid geometry*, Publ. Math. IHES, **82** (1995), pp. 5–96.

[Kas] P. Kassaei, *A gluing lemma and overconvergent modular forms*, Duke Math Journal, to appear.

[K] N. Katz, *p-adic properties of modular schemes and modular forms* in Modular Functions of One Variable III, Springer Lecture Notes in Mathematics **350** (1973), pp. 69–190.

[Ke] J. Kelley, *General topology*, D. Van Nostrand Co., Inc., New York, 1955.

[KL] M. Kisin, K.F. Lai, *Overconvergent Hilbert modular forms*, Amer. Journal of Math., **127** (2005), to appear.

[Liu] Q. Liu, *Un contre-exemple au "critére cohomologique d'affinoidicité"*, C.R. Acad. Sci. Paris Sér. I Math., **307** (1988), no. 2, pp. 83–86.

[Lu] J. Lubin, *Finite subgroups and isogenies of one-parameter formal groups*, Annals of Math., **85** (1967), pp. 296–302.

[L1] W. Lütkebohmert, *Formal-algebraic and rigid-analytic geometry*, Math. Annalen, **286** (1990), pp. 341–371.

[L2] W. Lütkebohmert, *On compactification of schemes*, Manuscripta Mathematica **80**, 1993, pp. 95–111.

[L3] W. Lütkebohmert, *The structure of proper rigid groups*, J. Reine Angew. Math., **468** (1995), pp. 167–219.

[Mat] H. Matsumura, *Commutative ring theory*, Cambridge Univ. Press, 1986.

[MM] B. Mazur, W. Messing, *One-dimensional crystalline cohomology and universal extensions*, Lecture Notes in Mathematics **370**, Springer-Verlag, 1974.

[Me] W. Messing, *The crystals associated to Barsotti–Tate groups: with applications to abelian schemes*, Lecture Notes in Mathematics **264**, Springer-Verlag, 1972.

[Mil] J. Milne, "Abelian varieties," Ch. V in *Arithmetic geometry* (Cornell/Silverman ed.), Springer-Verlag, 1986.
[Mum] D. Mumford, *Abelian varieties*, Oxford University Press, Bombay, 1970.
[GIT] D. Mumford, *et al.*, *Geometric invariant theory* (3rd ed.), Springer-Verlag, 1994.
[NO] P. Norman, F. Oort, *Moduli of abelian varieties*, Annals of Math., **112** (1980), pp. 413–439.
[O] F. Oort, *Finite group schemes, local moduli for abelian varieties, and lifting problems*, Compositio Math., **23** (1971), pp. 265–296.
[vdPS] M. van der Put, P. Schneider, *Points and topologies in rigid geometry*, Math. Annalen, **302** (1995), pp. 81–103.
[R] M. Raynaud, *Schémas en groupes de type* $(p, p, \ldots, p)$, Bull. Soc. Math. Fr., **102** (1974), pp. 241–280.

Department of Mathematics, University of Michigan, Ann Arbor, MI 48109, USA
E-mail address: bdconrad@umich.edu