The Dissertation Committee for Keenan James Kidwell
certifies that this is the approved version of the following dissertation:

Some Results in Iwasawa Theory and the p-adic Representation Theory of p-adic GL_2

Committee:

Mirela Çiperiani, Supervisor

David Ben-Zvi

Nancy Childress

David Helm

José Felipe Voloch

Some Results in Iwasawa Theory and the p-adic Representation Theory of p-adic GL_2

by

Keenan James Kidwell, B.S.

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

May 2014

To my parents.

Acknowledgments

First of all, I wish to thank my advisor Mirela Çiperiani for her guidance, her patience, and her understanding. I am grateful to all the number theorists at UT, especially David Helm, who on many occasions set aside time outside of classes to discuss math with me. Additionally, I want to thank David Ben-Zvi for letting me explain my work to him, and for serving on my committee. I owe an enormous debt of gratitude to Sandra Catlett and Dan Knopf, who have served as Graduate Program Coordinator and Graduate Advisor, respectively, during the majority of my time at UT. Their support and expertise were invaluable. It is a pleasure to thank Nancy Childress, whose enthusiasm and excellent teaching led me to pursue a graduate degree in mathematics, and to specialize in number theory.

My time in Austin has been greatly enriched by the wonderful friends I have made in the last six years. In particular I want to thank Jeff Hatley, Gil Moss, David Paige, and Steven Pogman for many helpful mathematical conversations, and for many fun times. Very special thanks are due to John Meth, without whose advice and devoted friendship I would certainly have given up on this endeavor long ago. Finally, I thank my parents for their unconditional love and support. Finishing graduate school is as much their achievement as it is my own.

Some Results in Iwasawa Theory and the p-adic Representation Theory of p-adic GL_2

Publication No. ____________________

Keenan James Kidwell, Ph.D.
The University of Texas at Austin, 2014

Supervisor: Mirela Çiperiani

This thesis is divided into two parts. In the first, we generalize results of Greenberg-Vatsal on the behavior of algebraic λ-invariants of p-ordinary modular forms under congruence. In the second, we generalize a result of Emerton on maps between locally algebraic parabolically induced representations and unitary Banach space representations of GL_2 over a p-adic field.

Table of Contents

Acknowledgments **v**

Abstract **vi**

Chapter 1. Introduction **1**

Chapter 2. Preliminaries on Iwasawa Theory **6**

2.1 The Iwasawa Algebra 6

2.2 Iwasawa Modules 10

2.3 The Structure Theorem for Finitely Generated Iwasawa Modules 17

Chapter 3. Algebraic λ-invariants of Modular Forms **23**

3.1 Introduction and Notation 23

3.2 Generalities on Limits of Λ-Modules and Twisting 25

3.2.1 Limits of Λ-Modules 25

3.2.2 Twisting of Λ-Modules and Characteristic Polynomials . 32

3.3 Selmer Groups 36

3.3.1 p-adic Selmer Groups 37

3.3.2 The Λ-module Structure of Local Cohomology Groups . 41

3.3.3 Non-Primitive Selmer Groups 45

3.3.4 Global-to-Local Maps 49

3.3.5 Divisibility of the Non-Primitive Selmer group 58

3.4 Algebraic λ-invariants 65

Chapter 4. Preliminaries on p-adic Representations of p-adic Groups **69**

4.1 Locally Convex Spaces Over p-adic Fields 69

4.2 Locally Analytic Groups and Their p-adic Representations . . 74

4.3 The Topology on Locally Analytic Functions 80

Chapter 5. Parabolically Induced and Unitary Banach Representations of p-adic GL_2 **84**

5.1 Introduction and Notation 84

5.2 The Locally Convex Space $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ and Tempered Linear Maps 89

5.3 Locally Algebraic and Locally Analytic Parabolic Inductions . 99

5.4 Proof of Theorem 5.3.1 102

Bibliography **111**

Chapter 1

Introduction

As the content of this work is entirely and decidedly non-Archimedean, let us fix a prime p, that we might refer to things as *p-adic* (p is assumed odd in the results of §3.3-4 to avoid any distasteful interaction with the cohomology of $\mathrm{Gal}(\mathbf{C}/\mathbf{R})$, but otherwise is arbitrary). Our results are divided into two parts which are logically independent from one another, although, as we explain below, in pursuing an analytic analogue of the algebraic results of Chapter 3, we were naturally led to the results of Chapter 5.

The first half of this thesis is devoted to generalizing results of Greenerg-Vatsal ([9]) on the behavior of algebraic Iwasawa invariants of p-ordinary modular forms under congruence to general weight and character, and to $\mathbf{Z}_p$-extensions of number fields other than $\mathbf{Q}$. To a p-ordinary cuspidal newform f, a number field F, and a $\mathbf{Z}_p$-extension F_∞ of F with Galois group $\Gamma = \mathrm{Gal}(F_\infty/F)$, we can attach a module $\mathrm{Sel}(F_\infty, f)$ (the Selmer group) over the completed group algebra Λ of Γ with coefficients in a finite extension of $\mathbf{Z}_p$ containing the Hecke eigenvalues of f (see the beginning of §2.1 for the definition of Λ). This is a discrete topological Λ-module whose Pontryagin dual $X(F_\infty, f)$ is finitely generated. The Iwasawa invariants of f are structural invariants of the module $X(F_\infty, f)$ (the definition of Iwasawa invariants in general is given in Definition 2.3.2). The Selmer group $\mathrm{Sel}(F_\infty, f)$ is defined in terms of a choice of lattice in the p-adic Galois representation attached to f. By a congruence between p-ordinary newforms we mean an isomorphism between their residual Galois representations. It is not clear *a priori* from the definition that the p-torsion of the Selmer group $\mathrm{Sel}(F_\infty, f)$ only depends on the residual Galois representation attached to f. Thus it is not immediately clear what relationships (if any) one can deduce between the Selmer groups (as Λ-modules) of congruent newforms. However, using Greenberg-Vatsal's method of *non-primitive Selmer groups* (defined precisely in §3.3.3), which are obtained by relaxing the local conditions at primes of F_∞ not lying over p in the definition of $\mathrm{Sel}(F_\infty, f)$, one obtains a module whose p-torsion *is* determined by the residual Galois representation attached to f. Therefore, to compare Iwasawa invariants of congruent modular forms, one approach is to first compare their non-primitive Selmer groups (which will in fact have isomorphic p-torsion), and then to compare the non-primitive Selmer groups to the usual Selmer groups. This is the approach that we follow in Chapter 3. Our main result on Iwasawa invariants (proved under various technical hypotheses) is Theorem 3.4.1. It is a formula expressing the difference in λ-invariants of congruent forms as a sum of λ-invariants of local cohomology groups at primes dividing the product of the tame levels of the forms (the prime-to-p parts of their levels). In principle, if one knows the λ-invariant of one form, the formula can be used to compute the λ-invariant of a congruent form.

Along the way to proving Theorem 3.4.1, we prove several other results on Selmer groups of p-ordinary newforms over fairly general $\mathbf{Z}_p$-extensions that are of independent interest. In particular, we carefully study the structure of local cohomology groups at finite primes that split in the $\mathbf{Z}_p$-extension (something which never happens for cyclotomic $\mathbf{Z}_p$-extensions) and prove the surjectivity of a global-to-local map in Galois cohomology (Proposition 3.3.7). We follow an argument originally due to Greenberg, using our study of local cohomology groups to identify the condition one must impose to obtain surjectivity even in the presence of split primes (which have the effect of making the target of the global-to-local map larger). We also give hypotheses under which the non-primitive Selmer group can be proved to have no non-zero Λ-submodules of finite index. This type of result is useful in traditional Iwasawa theory as Λ-modules are generally studied up to *pseudo-isomorphism*, meaning up to morphisms with finite kernel and cokernel.

The main result of the second half of this thesis, Theorem 5.3.1, is a purely local p-adic representation-theoretic result for GL_2 over a p-adic field L that reduces to [4, Proposition 2.5] when $L = \mathbf{Q}_p$. The theorem states that, under a "non-critical slope" hypothesis, together with a unitarity hypothesis on a central character, continuous linear $\mathrm{GL}_2(L)$-equivariant maps from certain locally algebraic parabolically induced representations into unitary Banach space representations of $\mathrm{GL}_2(L)$ extend uniquely to a larger locally analytic parabolically induced representation containing the locally algebraic representation. A result along these lines was proved by Breuil in [2], but it only address injections into unitary Banach space representations (and the locally algebraic parabolic inductions under consideration are not always irreducible, so not all non-zero maps out of them need be injective). Moreover, we follow Emerton's approach, which is more representation-theoretic.

The strategy of proof of Theorem 5.3.1 is to reduce to a similar (classical, non-equivariant) result, due in its original form to Amice-Vélu and Vishik, regarding what Emerton calls *tempered* linear maps out of the space of locally analytic functions on $\mathscr{O}_L$ into Banach spaces (see Definition 5.2.1 for the definition of a tempered linear map). This reduction is carried out by restricting functions in parabolically induced representations of $\mathrm{GL}_2(L)$ to the copy of $\mathscr{O}_L$ given by upper unipotent matrices with integral upper right entry, and relating equivariance of linear maps with respect to a certain submonoid of $\mathrm{GL}_2(L)$ to the temperedness condition (see Lemma 5.4.4). We also make use of a description of the locally convex topology on the space of locally $\mathbf{Q}_p$-analytic functions (see §4.1 for local convexity and §5.2 for the specific space in question) which seems implicit in some of the literature, but for which we know of no published proof. The description depends on the fact that two topologies of compact type are either equal or incomparable (see Definition 4.2.2 for the notion of a compact type space over a p-adic field).

Although, as alluded to above, Theorem 5.3.1, and indeed all the results of Chapter 5, are logically independent from the results of Chapter 3, our motivation for proving it came from a desire to eventually generalize Emerton's representation-theoretic construction of the p-adic L-functions of p-stabilized newforms of non-critical slope in [4, §4]. The case $L = \mathbf{Q}_p$ of Theorem 5.3.1

plays a crucial role in this construction, in which the theorem is applied with the completed cohomology of modular curves playing the role of the target unitary Banach representation of $\mathrm{GL}_2(\mathbf{Q}_p)$. One goal of such a generalization would be to have an adequate framework in which to prove analytic analogues of the results on algebraic Iwasawa invariants in Chapter 3. This is something we hope to do in the future. Irrespective of this specific motivation, Theorem 5.3.1 is also of intrinsic interest in the field of p-adic representations of p-adic groups, which can reasonably be viewed as a manifestation of non-commutative Iwasawa theory (as evidenced by the title of Schneider-Teitelbaum's first paper on the subject [15]).

Preliminary facts and definitions from Iwasawa theory and p-adic functional analysis and representation theory are recorded in Chapters 2 and 4, respectively. In the arguments of Chapter 3 we make frequent use of standard local and global duality theorems in Galois cohomology. A convenient reference for these theorems in their classical form is [11, VII.2-3, VIII.6], while [12, Appendix A, §3] includes the Iwasawa-theoretic "limit" version of the Poitou-Tate exact sequence used in the proof of Proposition 3.3.7.

Chapter 2

Preliminaries on Iwasawa Theory

2.1 The Iwasawa Algebra

Let Γ be a profinite group and $\mathscr{O}$ the ring of integers in a finite extension E of $\mathbf{Q}_p$ with uniformizer ϖ. The Iwasawa algebra of Γ with coefficients in $\mathscr{O}$, $\mathscr{O}[[\Gamma]]$, is defined to be

$$\varprojlim_{N} \mathscr{O}[\Gamma/N],$$

where the inverse limit is taken over the set of all open normal subgroups N of Γ, directed by reverse inclusion, and for an inclusion $N \subseteq N'$ of open normal subgroups of Γ, the transition map $\mathrm{Res}_{N,N'} : \mathscr{O}[\Gamma/N] \to \mathscr{O}[\Gamma/N']$ is the $\mathscr{O}$-algebra map induced by the natural homomorphism $\Gamma/N \to \Gamma/N'$. Since each quotient Γ/N is finite, the group ring $\mathscr{O}[\Gamma/N]$ is a finite free $\mathscr{O}$-module, and hence is profinite in the ϖ-adic topology (which is just the product topology upon identifying $\mathscr{O}[\Gamma/N]$ with the direct sum of copies of $\mathscr{O}$ indexed by Γ/N). Endowing each $\mathscr{O}[\Gamma/N]$ with its ϖ-adic topology and $\mathscr{O}[[\Gamma]]$ with the resulting inverse limit topology, $\mathscr{O}[[\Gamma]]$ therefore becomes a profinite topological $\mathscr{O}$-algebra. The natural $\mathscr{O}$-algebra map $\mathscr{O}[G] \hookrightarrow \mathscr{O}[[\Gamma]]$ is injective with dense image (injectivity holds because the open normal subgroups of Γ form a base of opens around the identity, whereas density of the image follows from the definition of the topology on $\mathscr{O}[[\Gamma]]$). In particular G can be regarded as a closed subgroup of the group of units of $\mathscr{O}[[\Gamma]]$.

Traditional (commutative) Iwasawa theory is primarily concerned with the case in which Γ is topologically isomorphic to the additive group $\mathbf{Z}_p$ (i.e. Γ is abelian pro-p and is free of rank one as a $\mathbf{Z}_p$-module). As this is the only case relevant for our results, we assume for the remainder of the chapter that $\Gamma \simeq \mathbf{Z}_p$ as topological groups. Choosing an isomorphism $\Gamma \simeq \mathbf{Z}_p$ amounts to the choice of a *topological generator* $\gamma \in \Gamma$, which is nothing but a choice of $\mathbf{Z}_p$-basis. In practice, there will be no canonical choice of basis, and a choice of γ is not regarded as part of the structure of Γ. Such a choice is however essential for understanding the structure of the ring $\mathscr{O}[[\Gamma]]$, as is made clear by the following standard result.

Proposition 2.1.1. *Given a topological generator γ for Γ, there is a unique isomorphism of topological $\mathscr{O}$-algebras $\mathscr{O}[[\Gamma]] \simeq \mathscr{O}[[T]]$ sending γ to $1+T$ (the target is regarded as a topological ring via its max-adic topology, i.e., its (ϖ, T)-adic topology).*

Proof. [11, Proposition 5.3.5]. □

Corollary 2.1.2. *The ring $\mathscr{O}[[\Gamma]]$ is a 2-dimensional regular Noetherian local ring with unique maximal ideal $(\gamma - 1, \varpi)$ and residue field $\mathscr{O}/\varpi$ (here γ is any topological generator for Γ). The inverse limit topology on $\mathscr{O}[[\Gamma]]$ coincides with its max-adic topology.*

Proposition 2.1.1 reduces the study of $\mathscr{O}[[\Gamma]]$ to that of the formal power series ring $\mathscr{O}[[T]]$, whose structure is well-known. Recall that a monic polynomial

$$f(T) = T^n + a_{n-1}T^{n-1} + \cdots + a_1T + a_0 \in \mathscr{O}[T]$$

is said to be *distinguished* if $a_i \in \varpi\mathscr{O}$ for $0 \leq i \leq n-1$.

Theorem 2.1.3. *If $F = \sum_{n\geq 0} a_nT^n \in \mathscr{O}[[T]]$ is non-zero then there is a non-negative integer μ, a unit power series $u \in \mathscr{O}[[T]]$, and a distinguished polynomial $f \in \mathscr{O}[T]$ such that $F = \varpi^\mu uf$. The integer μ, the unit power series u, and the distinguished polynomial f are uniquely determined by F.*

Proof. [11, Theorem 5.3.4]. □

Definition 2.1.1. In the notation of Theorem 2.1.3, the integer μ is called the μ-*invariant* of F and the degree $\deg(f)$ is called the λ-*invariant of* F. We will denote these invariants by $\mu(F)$ and $\lambda(F)$, respectively.

Corollary 2.1.4. *The prime ideals of height* 1 *in $\mathscr{O}[[T]]$ are (without repetition) the principal ideals generated by an irreducible distinguished polynomial $f \in \mathscr{O}[T]$ together with the principal ideal generated by (ϖ). For any distinguished polynomial $f \in \mathscr{O}[T]$ (not necessarily irreducible), $\mathscr{O}[[T]]/(f)$ is a finite free $\mathscr{O}$-module of rank* $\deg(f)$.

The next two propositions give somewhat more intrinsic characterizations of the μ- and λ-invariants of a non-zero $F \in \mathscr{O}[[T]]$. Note that because $\mathscr{O}[[T]]$ is regular, it is a unique factorization domain (although invoking the implication "regular implies factorial" is overkill, as the factoriality of $\mathscr{O}[[T]]$ follows fairly easily from Theorem 2.1.3).

Proposition 2.1.5. *The μ-invariant of a non-zero $F \in \mathscr{O}[[T]]$ is the multiplicity of the irreducible ϖ in the unique factorization of F into irreducibles in $\mathscr{O}[[T]]$.*

Proof. Write $F = \varpi^{\mu(F)} uf$ as in Theorem 2.1.3. If $f = f_1^{e_1} \cdots f_r^{e_r}$ is the factorization of f into monic irreducibles in $\mathscr{O}[T]$, then it is easy to see that each f_i is distinguished. By Corollary 2.1.4, the f_i are irreducible in $\mathscr{O}[[T]]$ and no f_i is associate to ϖ. Thus $F = u\varpi^{\mu(F)} f_1^{e_1} \cdots f_r^{e_r}$ is the factorization of F into irreducibles in $\mathscr{O}[[T]]$, and $\mu(F)$ is precisely the multiplicity of ϖ in this factorization. □

Proposition 2.1.6. *The λ-invariant of a non-zero $F \in \mathscr{O}[[T]]$ is the $\mathscr{O}$-rank of the finite free $\mathscr{O}$-module $\mathscr{O}[[T]]/(\varpi^{-\mu(F)} F)$.*

Proof. Write $F = \varpi^{\mu(F)} uf$ as in the proof of Theorem 2.1.3, so that $\varpi^{-\mu(F)} F = uf$, and since $u \in \mathscr{O}[[T]^{\times}$, $(\varpi^{-\mu(F)} F) = (uf) = (f)$. Using the definition of $\lambda(F)$ together with the second assertion of Corollary 2.1.4, we find that

$$\lambda(F) = \deg(f) = \operatorname{rank}_{\mathscr{O}} \big(\mathscr{O}[[T]]/(f)\big) = \operatorname{rank}_{\mathscr{O}} \big(\mathscr{O}[[T]]/(\varpi^{-\mu(F)} F)\big).$$

□

Note that Proposition 2.1.5 shows that if $F \in \mathscr{O}[[T]]$ is non-zero, then $\mu(\varpi^{-\mu(F)} F) = 0$, and therefore Proposition 2.1.6 implies that $\lambda(F) =$

$\lambda(\varpi^{-\mu(F)}F)$. Moreover, the characterizations of $\mu(F)$ and $\lambda(F)$ provided by the preceding two propositions make it apparent that for any $\mathscr{O}$-algebra automorphism φ of $\mathscr{O}[[T]]$, $\mu(F) = \mu(\varphi(F))$ and $\lambda(F) = \lambda(\varphi(F))$.

Definition 2.1.2. For $F \in \mathscr{O}[[\Gamma]]$ non-zero, we define the *μ-invariant* $\mu(F)$ of F to be the multiplicity of ϖ in the factorization of F into irreducibles in $\mathscr{O}[[\Gamma]]$, and we define the *λ-invariant* $\lambda(F)$ of F to be the $\mathscr{O}$-rank of $\mathscr{O}[[T]]/(\varpi^{-\mu(F)}F)$ (that this rank is finite follows by applying an isomorphism $\mathscr{O}[[\Gamma]] \simeq \mathscr{O}[[T]]$ and invoking Corollary 2.1.4).

Corollary 2.1.7. *If γ is a choice of topological generator for Γ inducing the $\mathscr{O}$-algebra isomorphism $\varphi : \mathscr{O}[[\Gamma]] \simeq \mathscr{O}[[T]]$ as in Proposition 2.1.1, then for any non-zero $F \in \mathscr{O}[[\Gamma]]$, $\mu(F) = \mu(\varphi(F))$ and $\lambda(F) = \lambda(\varphi(F))$.*

Proof. Since φ is an $\mathscr{O}$-algebra isomorphism, it induces an isomorphism $\varpi\mathscr{O}[[\Gamma]] \simeq \varpi\mathscr{O}[[T]]$. Writing $F = \varpi^{\mu(F)}F_1$ with $F_1 \notin \varpi\mathscr{O}[[\Gamma]]$, we therefore have $\varphi(F) = \varpi^{\mu(F)}\varphi(F_1)$ with $\varphi(F_1) \notin \varpi\mathscr{O}[[T]]$, so $\mu(F) = \mu(\varphi(F))$ by Proposition 2.1.5. Thus φ induces an $\mathscr{O}$-algebra isomorphism

$$\mathscr{O}[[\Gamma]]/(\varpi^{-\mu(F)}F) \simeq \mathscr{O}[[T]]/(\varpi^{-\mu(\varphi(F))}\varphi(F)),$$

so that $\lambda(F) = \lambda(\varphi(F))$ by Proposition 2.1.6. □

2.2 Iwasawa Modules

We retain the notation of the previous section, and for convenience we adopt the standard practice of denoting the ring $\mathscr{O}[[\Gamma]]$ by Λ. Since Λ

is a compact topological $\mathscr{O}$-algebra, one would expect that in studying Λ-modules the topology would play a significant role; to some extent this is true. Certainly we are interested in topological Λ-modules (abelian topological groups X with a Λ-module structure for which the action map $\Lambda \times X \to X$ is continuous). However, because Λ is a complete *Noetherian* local ring with *finite* residue field, the topological aspect of studying Λ-modules turns out to be essentially trivial. More precisely, the modules of interest either naturally carry the discrete topology, or are finitely generated over Λ (in the abstract sense), in which case the max-adic topology is the only (Hausdorff) topology compatible with the Λ-module structure.

Proposition 2.2.1. *Let R be a (commutative) complete Noetherian local ring with maximal ideal $\mathfrak{m}$ and finite residue field. If X is a Hausdorff topological R-module which is finitely generated as an abstract R-module, then the topology on X is the $\mathfrak{m}$-adic topology (which is compact). Conversely, for any abstract finitely generated R-module X, the $\mathfrak{m}$-adic topology on X makes X into a compact Hausdorff topological R-module. Any R-module homomorphism between finitely generated R-modules is continuous for the $\mathfrak{m}$-adic topology on source and target.*

Proof. First note that, because $R/\mathfrak{m}$ is finite, induction shows that the discrete quotients $R/\mathfrak{m}^n$ are finite for all $n \geq 1$. In particular, since the natural ring map $R \to \varprojlim_n R/\mathfrak{m}^n$ is a topological isomorphism (where the source is endowed with the $\mathfrak{m}$-adic topology and the target with the inverse limit topology), R is compact. Now consider the submodules $\mathfrak{m}^n X$ of X. Each is finitely generated, since X is a finitely generated R-module and R is Noetherian. Therefore, each of these submodules is a continuous homomorphic image of some direct sum of finitely many copies of R, and hence is compact (here we use that X is a topological R-module). Since X is Hausdorff, each $\mathfrak{m}^n X$ is then closed in X. On the other hand, the Hausdorff quotient $X/\mathfrak{m}^n X$ is a finitely generated $R/\mathfrak{m}^n$-module, and hence is set-theoretically finite. It follows that $\mathfrak{m}^n X$ is open in X for all $n \geq 1$. Therefore the given topology of X is finer than the $\mathfrak{m}$-adic topology, so the identity map $X \to X_0$, where X_0 denotes X endowed with its $\mathfrak{m}$-adic topology, is continuous. But the source of this map is compact (being finitely generated over R) and the target is Hausdorff by the Krull intersection theorem, so this continuous bijection must in fact be a homeomorphism, and the topologies coincide.

The converse is standard and holds with R replaced by any ring and $\mathfrak{m}$ replaced by any ideal (without finiteness hypotheses). □

The upshot of Proposition 2.2.1 is that, by equipping a finitely generated (abstract) Λ-module with its max-adic topology, we obtain a fully faithful embedding of the category of finitely generated (abstract) Λ-modules into the category of compact Hausdorff topological Λ-modules with continuous Λ-module homomorphisms. We will therefore always implicitly endow a finitely generated Λ-module with its max-adic topology. We will have no cause to consider compact Λ-modules that aren't finitely generated.

The other class of Λ-modules that we will consider are discrete Λ-

modules, by which we mean discrete abelian groups with the structure of a topological Λ-modules. The continuity of the action $\Lambda \times X \to X$ for a discrete group X admits a concrete algebraic characterization.

Lemma 2.2.2. *Let X be a discrete abelian group with a Λ-module structure. Then the action map $\Lambda \times X \to X$ is continuous if and only if every $x \in X$ is annihilated by some power of the maximal ideal of Λ. In particular, the induced $\mathscr{O}$-module structure on X makes it into a torsion $\mathscr{O}$-module.*

Proof. Assume X is a discrete Λ-module and fix $x \in X$. The map $\lambda \mapsto \lambda x : \Lambda \to X$ is then a continuous Λ-module homomorphism, so, since X is discrete, its kernel is an open ideal of Λ, and therefore contains some power of the maximal ideal of Λ. Thus some power of this ideal annihilates x. Conversely, assume that each $x \in X$ is annihilated by some power of the maximal ideal. If $((\lambda_i, x_i))$ is a convergent net in $\Lambda \times X$ with limit (λ, x), then for sufficiently large i, $x_i = x$ (since X is discrete). Thus, for sufficiently large i, x_i is annihilated by some (fixed, independent of i) power of the maximal ideal of Λ. For i, j sufficiently large, $\lambda_i - \lambda_j$ lies in this power of the maximal ideal, so we have, for such i, j, $\lambda_i x_i = \lambda_j x_j$. Thus the net $(\lambda_i x_i)$ is eventually constant, and hence convergent in X. So we win. The last assertion follows because ϖ is an element of the maximal ideal of Λ. □

The following proposition characterizes discrete Λ-modules in terms of their $\mathscr{O}$-module structure and the action of the group Γ (this is how discrete

Λ-modules arise in practice, not with an *a priori* Λ-module structure, but with an $\mathscr{O}[\Gamma]$-module structure satisfying some continuity conditions).

Proposition 2.2.3. *Let X be a torsion $\mathscr{O}$-module with a smooth, $\mathscr{O}$-linear action of Γ. Then there is a unique Λ-module structure on X which is continuous for the discrete topology on X and extends the given $\mathscr{O}[\Gamma]$-module structure.*

Proof. If Y is any Hausdorff topological Λ-module, then the Λ-module structure on Y is uniquely determined by the $\mathscr{O}[\Gamma]$-module structure, because the image of the embedding $\mathscr{O}[\Gamma] \hookrightarrow \Lambda$ is dense. So we just need to show the existence of a Λ-module structure on X as in the statement of the proposition. The open subgroups of Γ are precisely the subgroups Γ^{p^n} for $n \geq 0$. By the assumed smoothness of the action of Γ on X, we have $X = \bigcup_{n \geq 0} X_n$, where $X_n = X^{\Gamma^{p^n}}$. As Γ acts $\mathscr{O}$-linearly on X and Γ is abelian, each X_n is an $\mathscr{O}[\Gamma]$-submodule of X whose $\mathscr{O}[\Gamma]$-module action factors through the quotient $\mathscr{O}[\Gamma/\Gamma^{p^n}]$. Because all these structures are induced by a single $\mathscr{O}[\Gamma]$-module structure on X, the inclusion $X_n \hookrightarrow X_{n+1}$ is an $\mathscr{O}[\Gamma/\Gamma^{p^{n+1}}]$-module homomorphism, where the source X_n is regarded as an $\mathscr{O}[\Gamma/\Gamma^{p^{n+1}}]$-module via the restriction map $\mathscr{O}[\Gamma/\Gamma^{p^{n+1}}] \to \mathscr{O}[\Gamma/\Gamma^{p^n}]$. Therefore, using the projections $\Lambda \to \mathscr{O}[\Gamma/\Gamma^{p^n}]$ for $n \geq 0$, we see that the inductive limit $X = \bigcup_{n \geq 0} X_n$ is naturally a Λ-module.

It remains to verify continuity of the Λ-action on X for the discrete topology. By Lemma 2.2.2, it suffices to show that each element of X is annihilated by a power of the maximal ideal of Λ. To this end, fix $x \in X$, say $x \in X_n$, $n \geq 0$. By construction, the action of Λ on X_n, and in particular on x, is through its quotient $\mathscr{O}[\Gamma/\Gamma^{p^n}]$. Moreover, if $m \geq 0$ is chosen so that $\varpi^m x = 0$ (this is the only point in the proof where we use the hypothesis that X is torsion over $\mathscr{O}$), then we see that the $\mathscr{O}[\Gamma]$-submodule of X_n generated by x is annihilated by ϖ^m, so that the action of Λ on this submodule, and in particular on x, must factor through the *discrete* quotient $(\mathscr{O}/\varpi^m)[\Gamma/\Gamma^{p^n}]$. But this means precisely that some power of the maximal ideal of Λ kills x. $\square$

Since Λ is a complete local ring, any finitely generated Λ-module is max-adically complete, and in particular is profinite in its max-adic topology. Therefore, a discrete Λ-module can be finitely generated if and only if it is set-theoretically finite. So in some sense discrete and finitely generated Λ-modules are very different kinds of objects. However, they are related by Pontryagin duality, as we now explain.

Definition 2.2.1. If X is a profinite or a discrete torsion $\mathscr{O}$-module, then the *Pontryagin dual* of X is $\widehat{X} = \mathrm{Hom}_{\mathscr{O},\mathrm{cts}}(X, E/\mathscr{O})$, equipped with the compact open topology.

As usual, Pontryagin duality interchanges discrete torsion $\mathscr{O}$-modules with profinite ones, and vice versa, and one has the double duality isomorphism for such $\mathscr{O}$-modules. We are interested in the Pontryagin dual functor applied to Λ-modules, although we will only need to apply it to discrete Λ-modules. To wit, let X be a discrete Λ-module. Then X is a discrete torsion $\mathscr{O}$-module by Lemma 2.2.2, so we can form its Pontryagin dual $\widehat{X}$, a profinite $\mathscr{O}$-module. Since Λ is commutative, $\widehat{X}$ becomes a Λ-module if we define $(\lambda f)(x) = f(\lambda x)$

for each $\lambda \in \Lambda$, $f \in \widehat{X}$, and $x \in X$. However, to be consistent with the literature (at least those rare parts of the literature that say anything at all about the Λ-module structure on $\widehat{X}$), we will always twist this action of Λ on $\widehat{X}$ by the automorphism induced by $\gamma \mapsto \gamma^{-1} : \Gamma \to \Gamma$, so that the induced action of $\gamma \in \Gamma$ on $f \in \widehat{X}$ is given by $(\gamma f)(x) = f(\gamma^{-1}x)$. Since we are twisting by an automorphism of Λ, this convention does not affect structural properties of $\widehat{X}$ as a Λ-module such as whether or not it is finitely generated or torsion. We won't worry about continuity of this action because we will only ever apply the construction to discrete Λ-modules satisfying the condition in the following definition (in which case issues of topology for the Λ-module $\widehat{X}$ become trivial by Proposition 2.2.1).

Definition 2.2.2. A discrete Λ-module X is said to be *cofinitely generated* if $\widehat{X}$ is a finitely generated Λ-module, and *cotorsion* if $\widehat{X}$ is a torsion Λ-module.

Thus, by linguistic decree, Pontryagin duality yields a functor from cofinitely generated (resp. cotorsion) Λ-modules to finitely generated (resp. torsion) Λ-modules. The majority of the Λ-modules with which we deal directly will be cofinitely generated (as opposed to finitely generated), but it is (unsurprisingly) the finitely generated Λ-modules for which there is a satisfactory structure theory.

2.3 The Structure Theorem for Finitely Generated Iwasawa Modules

We retain the notation of the previous two sections. Although the ring Λ is not a principal ideal domain (it has Krull dimension 2), there is a structure theory for finitely generated Λ-modules which is almost identical to the theory for finitely generated modules over a principal ideal domain. The catch is that the theory only describes modules up to *pseudo-isomorphism*, which roughly means up to (set-theoretically) finite submodules.

Definition 2.3.1. A Λ-module homomorphism $f : X \to Y$ between finitely generated Λ-modules is said to be a *pseudo-isomorphism*, in which case X and Y are *pseudo-isomorphic*, if $\ker(f)$ and $\text{coker}(f)$ are finite.

Beware that, despite the terminology, pseudo-isomorphism does not define an equivalence relation on the class of all finitely generated Λ-modules; it fails to be symmetric in general. If however one restricts to the class of finitely generated *torsion* Λ-modules, then pseudo-isomorphism does define a symmetric (and hence an equivalence) relation. We will have no need of this fact (which begs the question as to why we're mentioning it at all).

Recall that if R is a principal ideal domain and X is a finitely generated R-module, then the so-called elementary divisor version of the structure theorem for X asserts the existence of an R-module isomorphism between X and an R-module of the form $R^r \oplus \sum_{i=1}^{s} R/(p_i^{e_i})$, where $r, s \geq 0$ are integers, the p_i are not necessarily distinct prime elements, and the e_i are positive integers. Moreover the integer r and the list of ideals $(p_1^{e_1}), \ldots, (p_t^{e_t})$ are uniquely determined by X (the latter up to reordering). The same statement holds for a finitely generated Λ-module, with the only difference being that one must replace "isomorphism" with "pseudo-isomorphism."

Theorem 2.3.1. *Let X be a finitely generated Λ-module. Then there exists a pseudo-isomorphism*

$$X \to \Lambda^r \oplus \sum_{i=1}^{s} \Lambda/\mathfrak{p}_i^{e_i},$$

where $r, s \geq 0$ are integers, the $\mathfrak{p}_i$ are not necessarily distinct height 1 prime ideals of Λ, and the e_i are positive integers. Moreover, the integer r and the list of ideals $\mathfrak{p}_1^{e_1}, \ldots, \mathfrak{p}_s^{e_s}$ are uniquely determined by X (the latter up to reordering).

Proof. [11, Theorem 5.1.10]. □

The integer r in Theorem 2.3.1 is the rank of X in the usual sense, i.e., the dimension of $\mathrm{Frac}(\Lambda) \otimes_\Lambda X$ over $\mathrm{Frac}(\Lambda)$, where $\mathrm{Frac}(\Lambda)$ is the field of fractions of Λ. This can be seen immediately by tensoring a pseudo-isomorphism as in the statement of the theorem with $\mathrm{Frac}(\Lambda)$, which will yield an isomorphism (since $\mathrm{Frac}(\Lambda)$ is a flat Λ-module and tensoring with it kills the finite kernel and cokernel). Thus X is Λ-torsion if and only if $r = 0$. It can also be shown in general that, restricting a pseudo-isomorphism as in Theorem 2.3.1 to the Λ-torsion submodule of X, one obtains a pseudo-isomorphism between the Λ-torsion submodules of the source and target, so the ideals $\mathfrak{p}_i^{e_i}$ are in fact determined by the Λ-torsion submodule of X. They are precisely the height 1 prime ideals in the support of X.

Using the classification of height 1 primes in the ring $\mathscr{O}[[T]] \simeq \Lambda$ provided by Corollary 2.1.4, we can refine the information about X contained in Theorem 2.3.1 somewhat. Namely, the height 1 primes of Λ fall into two classes: residue characteristic p and residue characteristic zero. Of course (ϖ) is the unique prime of residue characteristic p, and a residue characteristic zero prime correspond under the choice of an isomorphism $\Lambda \simeq \mathscr{O}[[T]]$ to some uniquely determined distinguished polynomial in $\mathscr{O}[T]$, but this polynomial is not independent of the choice of topological generator used to define such an isomorphism. At any rate, we can write the pseudo-isomorphism of Theorem 2.3.1 as

$$X \to \Lambda^r \oplus \sum_{i=1}^{s} \Lambda/\mathfrak{p}_i^{e_i} \oplus \sum_{j=1}^{t} \Lambda/(\varpi^{f_j}), \tag{2.1}$$

where now the $\mathfrak{p}_i$ are height 1 primes of characteristic zero, $t \geq 0$ is an integer, and the f_j are (uniquely determined) positive integers. It is via this expression that we define the Iwasawa invariants of a finitely generated Λ-module X.

Definition 2.3.2. Let X be a finitely generated Λ-module, and choose a pseudo-isomorphism as in Equation (2.1). The *μ-invariant* $\mu(X)$ of X is the non-negative integer $\sum_{j=1}^{t} f_j$. The *λ-invariant* $\lambda(X)$ of X is the non-negative integer $\sum_{i=1}^{s} e_i \deg(\mathfrak{p}_i)$, where $\deg(\mathfrak{p}_i) = \operatorname{rank}_{\mathscr{O}}(\Lambda/\mathfrak{p}_i)$.

Our remarks following Theorem 2.3.1 show that the μ- and λ-invariants of X coincide with those of its Λ-torsion submodule. It can also be shown that the restriction of a pseudo-isomorphism as in Equation (2.1) to the $\mathscr{O}$-torsion submodule of X yields a pseudo-isomorphism between the $\mathscr{O}$-torsion submodules of the source and target. In particular the μ-invariant of X only depends on the $\mathscr{O}$-torsion submodule of X.

Proposition 2.3.2. *Let X be a finitely generated torsion Λ-module. Then $\mu(X) = 0$ if and only if the $\mathscr{O}$-torsion submodule of X is finite, if and only if X is a finitely generated $\mathscr{O}$-module.*

Proof. For any integer $m \geq 1$, the quotient $\Lambda/(\varpi^m)$ is infinite (it's isomorphic to a power series ring in one variable over $\mathscr{O}/\varpi^m$). Thus, by our remarks above, the ϖ-power torsion submodule of X is infinite if and only if no summands of the form $\Lambda/(\varpi^m)$ appear in the target of a pseudo-isomorphism as in Equation (2.1), which happens if and only if $\mu(X) = 0$. Assuming this is the case, then because we have assumed that X is Λ-torsion, it is pseudo-isomorphic to a direct sum of quotients $\Lambda/\mathfrak{p}^e$ with $\mathfrak{p}$ a height 1 prime of characteristic zero and e a positive integer. These quotients are finitely generated $\mathscr{O}$-modules, so it follows that X is as well. $\square$

Proposition 2.3.3. *If X is a discrete cofinitely generated Λ-module, then X is Λ-cotorsion with $\mu(\widehat{X}) = 0$ if and only if $X[\varpi]$ is finite.*

Proof. Basic properties of the Pontryagin duality functor show that the Pontryagin dual of $X[\varpi]$ is $\widehat{X}/\varpi\widehat{X}$. If this group is finite, then the topological version of Nakayama's lemma ([1, §3 Corollary]) implies that $\widehat{X}$ is a finitely generated $\mathscr{O}$-module. This implies first that $\widehat{X}$ must be a torsion Λ-module (since a Λ-module with positive rank cannot be a finitely generated $\mathscr{O}$-module), and then by Proposition 2.3.2 that $\mu(\widehat{X}) = 0$. Conversely, assume that X is a cotorsion Λ-module with $\mu(\widehat{X}) = 0$. Then by Proposition 2.3.2, $\widehat{X}$ is a finitely generated $\mathscr{O}$-module, so in particular $\widehat{X}/\varpi\widehat{X}$ is a finite-dimensional $\mathscr{O}/\varpi$-vector space. By Pontryagin duality, $X[\varpi]$ is finite. □

We now wish to connect the λ-invariant defined in Definition 2.3.2 with the λ-invariants defined in Definitions 2.1.1 and 2.1.2. Concretely, if we choose generators p_i for the $\mathfrak{p}_i$, and a topological generator γ for Γ, yielding an isomorphism $\varphi : \Lambda \simeq \mathscr{O}[[T]]$, then there are unique distinguished polynomials $f_i \in \mathscr{O}[T]$ which generate the $\varphi(\mathfrak{p}_i)$, and we have, in the notation of Definitions 2.1.1 and 2.1.2,

$$\lambda(X) = \sum_{i=1}^{s} e_i \lambda(p_i) = \sum_{i=1}^{s} e_i \lambda(f_i) = \sum_{i=1}^{s} e_i \deg(f_i) = \sum_{i=1}^{s} \deg(f_i^{e_i}).$$

Definition 2.3.3. In the notation above, the *characteristic polynomial* of X with respect to the topological generator γ is $\varpi^{\mu(X)} \prod_i f_i^{e_i} \in \mathscr{O}[T]$ (this is really only defined up to a unit, since ϖ is only defined up to a unit, but we will ignore this fact). The *monic part* of the characteristic polynomial of X with respect to γ is $\varpi^{-\mu(X)}$ times the characteristic polynomial of X.

Unlike the μ- and λ-invariants of X, the characteristic polynomial of X requires and depends on a choice of topological generator γ for Γ. Its degree, however, does not, as that is simply the λ-invariant of X. Similarly the ϖ-adic valuation of its leading coefficient is the μ-invariant of X, which is independent of the choice of topological generator γ. Note that, regardless of whether or not the μ-invariant of X is zero, Theorem 2.3.1 implies that $X_E = X \otimes_{\mathscr{O}} E$ is a finite-dimensional E-vector space, and that, upon choosing γ, and then regarding X_E as an $\mathscr{O}[[T]]$-module via the induced isomorphism $\Lambda \simeq \mathscr{O}[[T]]$, the monic part of the characteristic polynomial of X with respect to E is precisely the characteristic polynomial of the endomorphism given by the action of T on X_E (hence the terminology).

Via Pontryagin duality, we extend the definitions of μ- and λ-invariants to discrete, cofinitely generated Λ-modules.

Definition 2.3.4. Let X be a discrete, cofinitely generated Λ-module. We define the *μ-invariant* $\mu(X)$ (resp. the *λ-invariant* $\lambda(X)$) of X to be the the μ-invariant $\mu(\widehat{X})$ (resp. the λ-invariant $\lambda(\widehat{X})$) of $\widehat{X}$.

Chapter 3

Algebraic λ-invariants of Modular Forms

3.1 Introduction and Notation

In this chapter, p is assumed to be odd except in §3.2, where p can be arbitrary. Consider a number field F, and a $\mathbf{Z}_p$-extension F_∞ of F, setting $\Gamma = \mathrm{Gal}(F_\infty/F)$. We impose the following two conditions on the set Σ_p of primes of F dividing p:

p-(i) for each prime $\mathfrak{p} \in \Sigma_p$, the ramification index $e(\mathfrak{p}/p)$ of $\mathfrak{p}$ in $F/\mathbf{Q}$ is less than $p-1$;

p-(ii) no prime $\mathfrak{p} \in \Sigma_p$ splits completely in F_∞.

Let f be a normalized p-ordinary newform of weight greater than or equal to 2 with Hecke eigenvalues in the ring of integers $\mathscr{O}$ of a finite extension E of $\mathbf{Q}_p$ with uniformizer ϖ and residue field $\mathbf{F}$. The algebraic λ- and μ-invariants of f, $\lambda(f)$ and $\mu(f)$, are non-negative integers defined in terms of the structure of the Selmer group for f over F_∞ as a module over the completed group ring $\mathscr{O}[[\Gamma]]$ (see §3.3.1 below for the definition of the Selmer group and its Iwasawa invariants). If $\bar{\rho}_f$ is the semisimple residual Galois representation attached to f, which has coefficients in the residue field $\mathbf{F}$, then we can also attach a residual Selmer group to f inside the Galois cohomology of $\bar{\rho}_f|_{G_F}$ (§3.3.3). While the residual Selmer group cannot be directly related to the ϖ-torsion of the Selmer group for f in general, such a relationship can be established for certain *non-primitive* analogues of these Selmer groups, obtained by omitting the local conditions at the primes of F in a finite set Σ_0 not containing any primes above ∞ or p. Namely, if Σ_0 contains the primes dividing the tame level of f and $\bar{\rho}|_{G_F}$ is absolutely irreducible, then the residual Σ_0-non-primitive Selmer group for f exactly recovers the ϖ-torsion of the Σ_0-non-primitive Selmer group for f. If in addition $\bar{\rho}|_{G_F}$ is ramified at the primes of F dividing p, it can be shown that the residual Selmer group for f depends only on the residual Galois representation $\bar{\rho}_f|_{G_F}$ (see Proposition 3.3.5 and the remark following its proof). Assuming a cotorsion hypothesis and the vanishing of the μ-invariant of f, as well as some additional technical hypotheses, this allows us to express the λ-invariant of f in terms of the λ-invariants of a non-primitive residual Selmer group and the local cohomology of the p-adic Galois representation attached to f at primes dividing the tame level of f. We can then compare the λ-invariants of two p-ordinary newforms whose residual Galois representations are isomorphic (Theorem 3.4.1).

Similar results in the case of $F = \mathbf{Q}$ (where there is only one $\mathbf{Z}_p$-extension of F) have been proved in [9], [8], and [6]. Our approach follows that of the former two references, in which the ϖ-torsion of a non-primitive Selmer group is related to a residual Selmer group. A key step in the argument is a surjectivity statement for a global-to-local map in Galois cohomology

(Proposition 3.3.7) which should be of independent interest as it allows finite primes that split in the $\mathbf{Z}_p$-extension under consideration.

We now introduce notation which will be used throughout this chapter. Fix embeddings $\iota_p : \overline{\mathbf{Q}} \hookrightarrow \overline{\mathbf{Q}}_p$ and $\iota_\infty : \overline{\mathbf{Q}} \hookrightarrow \mathbf{C}$. When we speak of primes of an algebraic extension of $\mathbf{Q}$, we always mean finite primes unless explicitly noted otherwise. If L is an algebraic extension of a number field F in $\overline{\mathbf{Q}}$ and η is a prime of L, then we write L_η for the direct limit of the fields L'_η, where L' is a finite subextension of L/F and L'_η denotes its completion at the prime below η (this is not the same as the completion of L at η unless L_η has finite degree over $\mathbf{Q}_p$). We also write I_η for the inertia group of G_{L_η}. If $f = \sum_{n\geq 1} a_n q^n$ is a normalized newform of some weight, level, and character, we regard the Hecke eigenvalues a_n and the character values as elements of $\overline{\mathbf{Q}}_p$ via $\iota_p \circ \iota_\infty^{-1}$. We always work with arithmetic Frobenius automorphisms, which we denote by Frob_ℓ, Frob_v, etc., and denote by $\epsilon : G_{\mathbf{Q}} \to \mathbf{Z}_p^\times$ the p-adic cyclotomic character (as well as its restriction to various subgroups of $G_{\mathbf{Q}}$).

3.2 Generalities on Limits of Λ-Modules and Twisting

3.2.1 Limits of Λ-Modules

In this subsection we establish some results about direct limits of certain discrete Λ-modules. The results are basic and almost certainly well-known, but we felt it would be useful to have the details written down, as related matters seem to have led to confusion in the literature.

Let Γ be a free pro-p group of rank one, $\mathscr{O}$ the ring of integers in a finite extension of E of $\mathbf{Q}_p$ with uniformizer ϖ, and $\Lambda = \mathscr{O}[[\Gamma]]$. Recall that the Pontryagin dual of a Λ-module X is denoted by $\widehat{X}$. We will make frequent use of μ- and λ-invariants (Definition 2.3.2). For $n \geq 0$, let $G_n = \Gamma/\Gamma^{p^n}$, and for $m \geq 1$, let $\Lambda_{n,m} = (\mathscr{O}/\varpi^m\mathscr{O})[G_n]$ and $\Lambda_n = \mathscr{O}[G_n]$. These are finitely generated Λ-modules. For fixed m, the $\Lambda_{n,m}$ form an inverse system and a directed system of Λ-modules. The transition maps for the inverse system are the natural restriction maps $\mathrm{Res}_{n,m} : \Lambda_{n+1,m} \to \Lambda_{n,m}$ induced by $\mathrm{Res}_{n,m}(g') = g'|_{G_n}$ for $g' \in G_{n+1}$, where $g'|_{G_n}$ denotes the canonical image of g' in G_n. The transition maps for the directed system are the natural corestriction maps $\mathrm{Cor}_{n,m} : \Lambda_{n,m} \to \Lambda_{n+1,m}$ induced by $\mathrm{Cor}_{n,m}(g) = \sum_{g' \in G_{n+1}, \mathrm{Res}_{n,m}(g')=g} g'$. Similarly the modules Λ_n form an inverse system and a directed system via restriction maps $\mathrm{Res}_n : \Lambda_{n+1} \to \Lambda_n$ and corestriction maps $\mathrm{Cor}_n : \Lambda_n \to \Lambda_{n+1}$. Of course, $\Lambda = \varprojlim_n \Lambda_n$ by definition.

Our first result concerns the self-duality of $\Lambda_{n,m}$.

Proposition 3.2.1. *For $n \geq 0, m \geq 1$, there are canonical Λ-module isomorphisms $\varphi_{n,m} : \widehat{\Lambda_{n,m}} \to \Lambda_{n,m}$ under which $\widehat{\mathrm{Res}_{n,m}} = \mathrm{Cor}_{n,m}$ and $\widehat{\mathrm{Cor}_{n,m}} = \mathrm{Res}_{n,m}$. More precisely, $\mathrm{Cor}_{n,m} \circ \varphi_{n,m} = \varphi_{n+1,m} \circ \widehat{\mathrm{Res}_{n,m}}$ and $\mathrm{Res}_{n,m} \circ \varphi_{n+1,m} = \varphi_{n,m} \circ \widehat{\mathrm{Cor}_{n,m}}$.*

Proof. The assertion is that there are isomorphisms $\varphi_{n,m}$ making the diagrams

$$\begin{array}{ccc} \widehat{\Lambda_{n,m}} & \xrightarrow{\widehat{\mathrm{Res}_{n,m}}} & \widehat{\Lambda_{n+1,m}} \\ \downarrow{\scriptstyle \varphi_{n,m}} & & \downarrow{\scriptstyle \varphi_{n+1,m}} \\ \Lambda_{n,m} & \xrightarrow[\mathrm{Cor}_{n,m}]{} & \Lambda_{n+1,m} \end{array}$$

and

$$\begin{array}{ccc} \widehat{\Lambda_{n+1,m}} & \xrightarrow{\widehat{\mathrm{Cor}_{n,m}}} & \widehat{\Lambda_{n,m}} \\ \downarrow{\scriptstyle \varphi_{n+1,m}} & & \downarrow{\scriptstyle \varphi_{n,m}} \\ \Lambda_{n+1,m} & \xrightarrow[\mathrm{Res}_{n,m}]{} & \Lambda_{n,m} \end{array}$$

commute. We have $\widehat{\Lambda_{n,m}} = \mathrm{Hom}_{\mathscr{O}}(\Lambda_{n,m}, E/\mathscr{O})$. Since ϖ^m kills $\Lambda_{n,m}$, this is the same as

$$\mathrm{Hom}_{\mathscr{O}/\varpi^m\mathscr{O}}(\Lambda_{n,m}, \varpi^{-m}\mathscr{O}/\mathscr{O}) \simeq \mathrm{Hom}_{\mathscr{O}/\varpi^m\mathscr{O}}(\Lambda_{n,m}, \mathscr{O}/\varpi^m) \simeq \Lambda_{n,m},$$

where the first isomorphism comes from $[\varpi^m] : \varpi^{-m}\mathscr{O}/\mathscr{O} \to \mathscr{O}/\varpi^m\mathscr{O}$ and the last isomorphism sends $\chi : \Lambda_{n,m} \to \mathscr{O}/\varpi^m\mathscr{O}$ to $\sum_{g\in G_n} \chi(g)g$. We define $\varphi_{n,m}$ to be the composite of these isomorphisms. So, explicitly, $\varphi_{n,m}$ sends $\chi : \Lambda_{n,m} \to E/\mathscr{O}$ to $\sum_{g\in G_n}[\varpi^m](\chi(g))g$. It is clear that $\varphi_{n,m}$ is an $\mathscr{O}$-module isomorphism, and we have, for $h \in G_n$,

$$\begin{aligned} \varphi_{n,m}(h\chi) &= \sum_{g\in G_n}[\varpi^m]((h\chi)(g))g \\ &= \sum_{g\in G_n}[\varpi^m](\chi(h^{-1}g))g \\ &= h\left(\sum_{g\in G_n}[\varpi^m](\chi(g))g\right) = h\varphi_{n,m}(\chi). \end{aligned}$$

Thus the map is G_n-equivariant, and therefore is a Λ-module isomorphism.

Now for the diagrams, beginning with the first. Going horizontally then vertically sends $\chi \in \widehat{\Lambda_{n,m}}$ to $\sum_{g'\in G_{n+1}}[\varpi^m](\chi(\mathrm{Res}_{n,m}(g')))g'$. Going vertically first sends χ to $\sum_{g\in G_n}[\varpi^m](\chi(g))g$, and then applying $\mathrm{Cor}_{n,m}$ gives

$$\sum_{g\in G_n}[\varpi^m](\chi(g))\,\mathrm{Cor}_{n,m}(g) = \sum_{g\in G_n}[\varpi^m](\chi(g))\left(\sum_{g'\in G_{n+1},\mathrm{Res}_{n,m}(g')=g} g'\right)$$
$$\sum_{g\in G_n}\left(\sum_{g'\in G_{n+1},\mathrm{Res}_{n,m}(g')=g}[\varpi^m](\chi(\mathrm{Res}_{n,m}(g')))g'\right)$$
$$= \sum_{g'\in G_{n+1}}[\varpi^m](\chi(\mathrm{Res}_{n,m}(g')))g'.$$

Thus the first diagram commutes.

For the second diagram, given $\chi \in \widehat{\Lambda_{n+1,m}}$ and going horizontally, we get $\chi\circ\mathrm{Cor}_{n,m}$, and then going vertically gives

$$\sum_{g\in G_n}[\varpi^m](\chi(\mathrm{Cor}_{n,m}(g)))g.$$

Going vertically first gives $\sum_{g'\in G_{n+1}}[\varpi^m](\chi(g'))g'$, and then taking $\mathrm{Res}_{n,m}$, we obtain

$$\sum_{g\in G_n}\left(\sum_{g'\in G_{n+1},\mathrm{Res}_{n,m}(g')=g}[\varpi^m](\chi(g'))\right)g.$$

The definition of $\mathrm{Cor}_{n,m}$ gives

$$\mathrm{Cor}_{n,m}(g) = \sum_{g'\in G_{n+1},\mathrm{Res}_{n,m}(g')=g} g',$$

so $[\varpi^m](\chi(\mathrm{Cor}_{n,m}(g))) = \sum_{g'\in G_{n+1},\mathrm{Res}_{n,m}(g')=g}[\varpi^m](\chi(g'))$, which finishes the proof. □

Corollary 3.2.2. *If $S_m = \varinjlim_n \Lambda_{n,m}$, with the limit taken with respect to the corestriction maps, then there is a canonical isomorphism of Λ-modules $\widehat{S_m} \simeq (\mathscr{O}/\varpi^m\mathscr{O})[[\Gamma]] \simeq \Lambda/\varpi^m\Lambda$. In particular, S_m is a cofinitely generated, cotorsion Λ-module with λ-invariant zero and μ-invariant m.*

Proof. As S_m is a discrete Λ-module, by generalities with Pontryagin duality, $\widehat{S_m}$ is canonically isomorphic as a Λ-module to $\varprojlim_n \widehat{\Lambda_{n,m}}$, with the limit taken with respect to the maps $\widehat{\mathrm{Cor}_{n,m}}$. By Proposition 3.2.1, this can be identified with the inverse limit of the $\Lambda_{n,m}$ taken with respect to the restriction maps, i.e., with $(\mathscr{O}/\varpi^m\mathscr{O})[[\Gamma]]$. The second isomorphism is the inverse of the isomorphism given by passage to the quotient of the natural map $\Lambda \to (\mathscr{O}/\varpi^m\mathscr{O})[[\Gamma]]$ (the kernel of the latter surjection is $\varpi^m\Lambda$). □

Now, instead of taking limits over $n \geq 0$, we want to take limits over $m \geq 1$. We will consider the discrete Λ-modules $S_n = \varinjlim_m (\mathscr{O}/\varpi^m)[G_n]$, where the transition maps $[\varpi]_{n,m} : \Lambda_{n,m} \to \Lambda_{n,m+1}$ are induced by the injective $\mathscr{O}$-module maps $[\varpi] : \mathscr{O}/\varpi^m\mathscr{O} \to \mathscr{O}/\varpi^{m+1}\mathscr{O}$ (multiplication by ϖ).

Proposition 3.2.3. *Under the isomorphisms $\varphi_{n,m} : \widehat{\Lambda_{n,m}} \simeq \Lambda_{n,m}$, $\widehat{[\varpi]_{n,m}} = \theta_{n,m}$, where $\theta_{n,m} : \Lambda_{n,m+1} \to \Lambda_{n,m}$ is induced by the natural $\mathscr{O}$-module map $\beta_m : \mathscr{O}/\varpi^{m+1}\mathscr{O} \to \mathscr{O}/\varpi^m\mathscr{O}$. More precisely, $\theta_{n,m} \circ \varphi_{n,m+1} = \varphi_{n,m} \circ \widehat{[\varpi]_{n,m}}$.*

Proof. The assertion is that the diagram

$$\begin{array}{ccc} \widehat{\Lambda_{n,m+1}} & \xrightarrow{\widehat{[\varpi]_{n,m}}} & \widehat{\Lambda_{n,m}} \\ \downarrow{\scriptstyle \varphi_{n,m+1}} & & \downarrow{\scriptstyle \varphi_{n,m}} \\ \Lambda_{n,m+1} & \xrightarrow[\theta_{n,m}]{} & \Lambda_{n,m} \end{array}$$

commutes. Beginning with $\chi \in \widehat{\Lambda_{n,m+1}}$, going along the top horizontal map gives $\chi \circ [\varpi]_{n,m}$, and then traveling vertically gives $\sum_{g \in G_n} [\varpi^m](\chi([\varpi]_{n,m}(g)))g$.

Going the other way gives $\sum_{g\in G_n} \beta_m([\varpi^{m+1}](\chi(g)))g$. To see that these coincide, fix $g \in G_n$, and let $r \in \varpi^{-m-1}\mathscr{O}$ represent $\chi(g) \in \varpi^{-m-1}\mathscr{O}/\mathscr{O}$. Then $\chi([\varpi]_{n,m}(g)) = \chi((\varpi + \varpi^{m+1}\mathscr{O})g) = \varpi r + \varpi^{-m}\mathscr{O}/\mathscr{O}$, so $[\varpi^m](\chi([\varpi]_{n,m}(g))) = \varpi^{m+1}r + \varpi^m\mathscr{O}$; since $[\varpi^{m+1}](\chi(g)) = \varpi^{m+1}r + \varpi^{m+1}\mathscr{O}$, this is exactly $\beta_m([\varpi^{m+1}](\chi(g)))$. Thus the diagram commutes. □

Corollary 3.2.4. *There is a canonical* Λ*-module isomorphism* $\widehat{S_n} \simeq \mathscr{O}[G_n]$*.*

Proof. By Proposition 3.2.3, $\widehat{S_n} \simeq \varprojlim_m \Lambda_{n,m}$, with the limit taken with respect to the maps given on coefficients by $\beta_m : \mathscr{O}/\varpi^{m+1}\mathscr{O} \to \mathscr{O}/\varpi^m\mathscr{O}$. These can be identified with the natural maps $\mathscr{O}[G_n]/(\varpi^{m+1}) \to \mathscr{O}[G_n]/(\varpi^m)$, and taking the inverse limit of this system of modules gives $\mathscr{O}[G_n]$ because $\mathscr{O}[G_n]$ is a finite, hence ϖ-adically complete, $\mathscr{O}$-module. □

Finally, we want to identify the discrete Λ-module $\mathscr{S} = \varinjlim_n S_n$, where the transition maps are $\psi_n = \varinjlim_m \mathrm{Cor}_{n,m} : \varinjlim_m \Lambda_{n,m} \to \varinjlim_m \Lambda_{n+1,m}$. This makes sense because the corestriction maps commute with the transition maps defining S_n and S_{n+1}.

Corollary 3.2.5. *There is a canonical* Λ*-module isomorphism* $\widehat{\mathscr{S}} \simeq \Lambda$*.*

Proof. We have $\widehat{\mathscr{S}} \simeq \varprojlim_n \widehat{S_n}$, where the limit is taken with respect to the maps $\widehat{\psi_n}$. Under the isomorphism $\widehat{S_n} \simeq \varprojlim_m \Lambda_{n,m}$ coming from Pontryagin duality and Proposition 3.2.1, the transition maps $\widehat{\psi_n} = \widehat{\varinjlim_m \mathrm{Cor}_{n,m}}$ for the modules $\widehat{S_n}$ become $\varprojlim_m \mathrm{Res}_{n,m}$ (again by Proposition 3.2.1). That is, the diagram

$$\begin{array}{ccc} \widehat{S_{n+1}} & \xrightarrow{\widehat{\psi_n}} & \widehat{S_n} \\ \simeq \downarrow & & \downarrow \simeq \\ \varprojlim_m \Lambda_{n+1,m} & \xrightarrow[\varprojlim_m \mathrm{Res}_{n,m}]{} & \varprojlim_m \Lambda_{n,m} \end{array}$$

commutes. As alluded to in the proof of Corollary 3.2.4, the inverse system consisting of the modules $\varprojlim_m \Lambda_{n,m}$ and the transition maps $\varprojlim_m \mathrm{Res}_{n,m}$ can be identified with the inverse system consisting of the modules $\mathscr{O}[G_n]$ and the transition maps $\mathrm{Res}_n : \mathscr{O}[G_{n+1}] \to \mathscr{O}[G_n]$. Thus $\widehat{\mathscr{S}}$ can be identified with $\varprojlim_n \mathscr{O}[G_n] = \Lambda$. $\square$

Corollary 3.2.6. *Let X be a cofinitely generated $\mathscr{O}$-module, $X \simeq (E/\mathscr{O})^r \oplus \sum_{i=1}^t \mathscr{O}/\varpi^{m_i}\mathscr{O}$. Then there is a Λ-module isomorphism $\varinjlim_n X \otimes_{\mathscr{O}} \mathscr{O}[G_n] \simeq \widehat{\Lambda}^r \oplus \sum_{i=1}^t \widehat{\Lambda/\varpi^{m_i}\Lambda}$, where the limit in the source is taken with respect to the maps $\mathrm{id}_X \otimes \mathrm{Cor}_n$ and Λ acts on the right tensor factor of each $X \otimes_{\mathscr{O}} \mathscr{O}[G_n]$. In particular, $\varinjlim_n X \otimes_{\mathscr{O}} \mathscr{O}[G_n]$ is a cofinitely generated Λ-module with corank $\mathrm{corank}_{\mathscr{O}}(X)$, λ-invariant zero, and μ-invariant $\sum_{i=1}^t m_i$.*

Proof. For each $n \geq 0$, we have a canonical isomorphism of Λ-modules

$$\Big((E/\mathscr{O})^r \oplus \sum_{i=1}^t \mathscr{O}/\varpi^{m_i}\mathscr{O}\Big) \otimes_{\mathscr{O}} \mathscr{O}[G_n] \simeq (E/\mathscr{O} \otimes_{\mathscr{O}} \mathscr{O}[G_n])^r \oplus \sum_{i=1}^t \Lambda_{n,m_i}.$$

These isomorphisms are compatible with the natural transition maps on both source and target as n varies (all coming from corestriction), and since direct limits commute with $\otimes_{\mathscr{O}}$ and finite direct sums, in the limit over n we obtain

$$(\varinjlim_n E/\mathscr{O} \otimes \mathscr{O}[G_n])^r \oplus \sum_{i=1}^t \varinjlim_n (\Lambda_{n,m_i}).$$

For the factor on the right, Proposition 3.2.2 shows that the limit is $\sum_{i=1}^{t} \widehat{\Lambda/\varpi^{m_i}\Lambda}$. For the left factor, we have, for each n,

$$E/\mathscr{O} \otimes \mathscr{O}[G_n] \simeq \varinjlim_{m} \Lambda_{n,m} = S_n$$

As n varies, the transition maps for the modules $E/\mathscr{O} \otimes_{\mathscr{O}} [G_n]$, coming from corestriction, become the transition maps $\psi_n : S_n \to S_{n+1}$ used to define the module $\mathscr{S}$ of Corollary 3.2.5. That corollary shows that, upon taking the limit, we obtain $\widehat{\Lambda}$. □

3.2.2 Twisting of Λ-Modules and Characteristic Polynomials

In this subsection we explain the effect of twisting by a character on the characteristic polynomial of a torsion Λ-module (Definition 2.3.3) with respect to a fixed topological generator of Γ. We retain the notation of the previous appendix. Let $q = p$ if p is odd and $q = 4$ is $p = 2$, and let $\kappa : \Gamma \to 1 + q\mathbf{Z}_p$ be a continuous character. Since the source and target of κ are isomorphic to $\mathbf{Z}_p$, κ is either trivial or injective, and in the latter case, it induces an isomorphism of Γ onto its image.

Once we fix a topological generator γ of Γ and identify Λ with $\mathscr{O}[[T]]$ via $\gamma \mapsto 1+T$, we can associate to κ a continuous $\mathscr{O}$-algebra endomorphism φ_κ of Λ, determined uniquely by the requirement that $\varphi_\kappa(T) = \kappa(\gamma)(1+T) - 1$. This is valid because $\kappa(\gamma)$ is a principal unit, and thus $\kappa(\gamma)(1+T) - 1$ lies in the unique maximal ideal of Λ. The map φ_κ is an automorphism because $\psi = \varphi_{\kappa^{-1}}$ satisfies $(\psi \circ \varphi)(T) = T = (\varphi \circ \psi)(T)$, and the only continuous $\mathscr{O}$-algebra endomorphism of Λ fixing T is the identity.

Now, if X is any Λ-module, we define a new Λ-module $X(\kappa)$ whose underlying $\mathscr{O}$-module is X, but with Λ-action twisted by φ_κ, i.e., we define $\lambda \cdot x = \varphi_\kappa(\lambda)x$ for $\lambda \in \Lambda$ and $x \in X$. Since φ_κ is an automorphism of Λ, it is clear that X is finitely generated (respectively torsion) if and only if $X(\kappa)$ is. In particular, if X is finitely generated and torsion, then so is $X(\kappa)$. Recall from Definition 2.3.3 and the discussion following it that the characteristic polynomial of a finitely generated torsion Λ-module X with respect to γ is equal to $\varpi^{\mu(X)}$ times the characteristic polynomial of the endomorphism T acting on the finite-dimensional E-vector space $X \otimes_{\mathscr{O}} E$, where $\mu(X)$ is the μ-invariant of X (if $X \otimes_{\mathscr{O}} E$ is zero, i.e., if $X = X[\varpi^\infty]$, then the characteristic polynomial is, by convention, just $\varpi^{\mu(X)}$). We wish to describe the effect that twisting by κ has on the characteristic polynomial, i.e., to give a formula for the characteristic polynomial for $X(\kappa)$ in terms of the characteristic polynomial of X. We will assume vanishing of the μ-invariant as this is the only case needed for our application.

Proposition 3.2.7. *Let $F(t) \in \mathscr{O}[t]$ be the characteristic polynomial of X, and assume $\mu(X) = 0$. Then $\mu(X(\kappa)) = 0$ and the characteristic polynomial of $X(\kappa)$ is*

$$\kappa(\gamma)^{\deg(F)} F(\kappa(\gamma)^{-1}(1+t) - 1).$$

Proof. The vanishing of $\mu(X)$ is equivalent to finiteness of $X[\varpi^\infty]$. Since X and $X(\kappa)$ have the same underlying $\mathscr{O}$-module, it follows that $\mu(X(\kappa)) = 0$ as well. If $X \otimes_{\mathscr{O}} E = 0$, then $X(\kappa) \otimes_{\mathscr{O}} E = 0$, and both characteristic polynomials are equal to 1, which is consistent with the formula in the statement of the proposition. Assume then that $X \otimes_{\mathscr{O}} E \neq 0$ and let $x_1, \ldots, x_d \in X$ be elements whose images in $X/X[\varpi^\infty]$ form an $\mathscr{O}$-basis (so $d = \mathrm{rank}_{\mathscr{O}}(X) = \deg(F)$). Then $x_1 \otimes 1, \ldots, x_d \otimes 1 \in X \otimes_{\mathscr{O}} E$ form a E-basis, and because $\mu(X) = 0$, if $[T]$ is the matrix for the endomorphism T of $X \otimes_{\mathscr{O}} E$ with respect to the chosen basis, then $F(t) = \det(It - [T])$. Since $X(\kappa) \otimes_{\mathscr{O}} E$ has the same underlying E-vector space as $X \otimes_{\mathscr{O}} E$, the $x_i \otimes 1$ constitute of a E-basis for this space as well. By definition, the action of T on $X(\kappa) \otimes_{\mathscr{O}} E$ coincides with the action of $\kappa(\gamma)(1+T)-1$ on $X \otimes_{\mathscr{O}} E$. In other words, if $[\kappa(\gamma)(1+T)-1]$ is the matrix for the endomorphism $\kappa(\gamma)(1+T)-1$ acting on $X \otimes_{\mathscr{O}} E$ with respect to the chosen basis, then the characteristic polynomial of $X(\kappa)$ is $\det(It - [\kappa(\gamma)(1+T)-1])$. We have

$$\begin{aligned} It - [\kappa(\gamma)(1+T)-1] &= It - \kappa(\gamma)I - \kappa(\gamma)[T] + I \\ &= I(t - \kappa(\gamma) + 1) - \kappa(\gamma)[T] \\ &= \kappa(\gamma)(I(\kappa(\gamma)^{-1}t - 1 + \kappa(\gamma)^{-1}) - [T]) \\ &= \kappa(\gamma)(I(\kappa(\gamma)^{-1}(1+t) - 1) - [T]). \end{aligned}$$

Thus the characteristic polynomial of $X(\kappa)$ is

$$\det(\kappa(\gamma)(I(\kappa(\gamma)^{-1}(1+t) - 1) - [T])) = \kappa(\gamma)^d F(\kappa(\gamma)^{-1}(1+t) - 1),$$

as claimed. □

Continuing to assume that X is finitely generated and torsion, but not necessarily with μ-invariant zero, one can prove using the structure theorem for such Λ-modules (Theorem 2.3.1) that X_{Γ_n} (the module of Γ_n-coinvariants of X) is finite if and only if $F(t)$ has no zeros in $\overline{E}$ of the form $\zeta - 1$, where ζ is a p^n-th root of unity. This observation implies that almost all twists of X by integral powers of κ (assuming κ is non-trivial) have finitely many Γ_n-coinvariants for all $n \geq 0$.

Proposition 3.2.8. *Assume κ is non-trivial. Then for all but finitely many $i \in \mathbf{Z}$, $X(\kappa^i)_{\Gamma_n}$ is finite for all $n \geq 0$.*

Proof. To begin, it follows from the structure theorem for finitely generated Λ-modules that for all i and $n \geq 0$, $X(\kappa^i)_{\Gamma_n}$ is finite if and only if the group $(X(\kappa^i)/X(\kappa^i)[\varpi^\infty])_{\Gamma_n}$ is finite. We may therefore assume that $\mu(X) = 0$ (so $\mu(X(\kappa^i)) = 0$ for all i). Denoting as before the characteristic polynomial of X by F, let $F_i(t) = \kappa(\gamma)^{i \deg(F)} F(\kappa(\gamma)^{-i}(1+t) - 1)$. By Proposition 3.2.7, this is the characteristic polynomial of $X(\kappa^i)$. Suppose that $i, j \in \mathbf{Z}$ are distinct integers for which there exist $n, m \geq 0$ with $X(\kappa^i)_{\Gamma_n}$ and $X(\kappa^j)_{\Gamma_m}$ infinite. Then there are $\zeta \in \mu_{p^n}(\overline{E})$ and $\zeta' \in \mu_{p^m}(\overline{E})$ such that $F_i(\zeta - 1) = 0 = F_j(\zeta' - 1)$. This means that $\kappa(\gamma)^{-i}\zeta - 1$ and $\kappa(\gamma)^{-j}\zeta' - 1$ are roots of F. If these roots are the same, then $\kappa(\gamma)^{-i}\zeta = \kappa(\gamma)^{-j}\zeta'$, whence $\kappa(\gamma)^{j-i} = \zeta^{-1}\zeta'$. The right-hand side of this last equation is visibly a root of unity, but the left-hand side is an element of $1 + q\mathbf{Z}_p$, which is torsion-free. Thus $\kappa(\gamma)^{j-i} = 1$, which forces $i = j$ (because κ is non-trivial), contrary to assumption. So the roots of F are distinct. It follows that we may define an injective map from the set of integers $i \in \mathbf{Z}$ for which $X(\kappa^i)_{\Gamma_n}$ is not finite for *some* n to the set of roots of F, which is finite as $F \neq 0$. The former set is therefore finite as well. $\square$

3.3 Selmer Groups

We assume for the remainder of this chapter that p is odd. Let F be a number field and F_∞ a $\mathbf{Z}_p$-extension of F satisfying conditions p-(i) and p-(ii) from §3.1. Throughout this section, we fix a normalized newform $f = \sum_{n\geq 1} a_n q^n$ of weight $k \geq 2$, level N, and character χ. We assume that the Fourier coefficients of f lie in $\mathscr{O}$, the ring of integers in a fixed finite extension E of $\mathbf{Q}_p$ with uniformizer ϖ and residue field $\mathbf{F}$ (it is known that the values of χ also lie in $\mathscr{O}$). Let $\rho_f : G_{\mathbf{Q}} \to \mathrm{GL}_2(E)$ be the p-adic Galois representation associated to f; ρ_f is unramified outside pN, and is characterized up to $\overline{\mathbf{Q}}_p$-isomorphism by the condition that for a rational prime $\ell \nmid pN$, the characteristic polynomial of $\rho_f(\mathrm{Frob}_\ell)$ is $X^2 - a_\ell X + \chi(\ell)\ell^{k-1}$. We denote by $\bar{\rho}_f : G_{\mathbf{Q}} \to \mathrm{GL}_2(\mathbf{F})$ the semisimple residual representation associated to f, and make the following assumptions:

(ord) f is p-ordinary in the sense that a_p is a p-adic unit, and

(ram) $\bar{\rho}_f|_{G_F}$ is absolutely irreducible and ramified at each $\mathfrak{p} \in \Sigma_p$.

The notion of ordinarity in (ord) actually depends on the choice of embedding used to regard the Fourier coefficients of f as p-adic numbers (but as we have fixed such an embedding, this will not matter for us). The second part of (ram) will hold if p is unramified in F, $k \not\equiv 1 \pmod{p-1}$, and χ is unramified at p. This follows from the local structure of ρ_f at p, which is given at the beginning of the following subsection (the restriction of ρ_f to a decomposition group at p is potentially ordinary in the sense of Greenberg).

3.3.1 p-adic Selmer Groups

Let V be a 2-dimensional E-vector space with $G_{\mathbf{Q}}$-action via ρ_f, and fix a $G_{\mathbf{Q}}$-stable $\mathscr{O}$-lattice T in V, setting $A = V/T$. Thus A is a cofree $\mathscr{O}$-module of corank 2 on which $G_{\mathbf{Q}}$ acts by ρ_f, and its Cartier dual $A^* = \mathrm{Hom}_{\mathscr{O}}(A, (E/\mathscr{O})(1))$ is a free $\mathscr{O}$-module of rank 2. Since we have assumed $\bar{\rho}_f|_{G_F}$ to be absolutely irreducible, the lattice T is unique up to $\mathscr{O}$-scaling, and the residual representation $\bar{\rho}_f|_{G_F}$ is given by the action of G_F on $A[\varpi] \simeq T/\varpi T$.

Our assumption (ord) that f is p-ordinary implies that for each prime $\mathfrak{p} \in \Sigma_p$, there is a $G_{F_\mathfrak{p}}$-stable line $V_\mathfrak{p} \subseteq V$ such that the $G_{F_\mathfrak{p}}$-action on $V_\mathfrak{p}$ is given by the product of $\epsilon^{k-1}\chi$ and an unramified character, and the $G_{F_\mathfrak{p}}$-action on $V/V_\mathfrak{p}$ is unramified [6, §4.1]. For $\mathfrak{p} \in \Sigma_p$, we set $A_\mathfrak{p} = \mathrm{im}(V_\mathfrak{p} \to A)$, so that $A_\mathfrak{p}$ and $A/A_\mathfrak{p}$ are both $\mathscr{O}$-cofree of corank 1, and the action of $G_{F_\mathfrak{p}}$ on $A/A_\mathfrak{p}$ is unramified.

For a prime $\mathfrak{P}$ of F_∞ lying over $\mathfrak{p} \in \Sigma_p$, we define the ordinary submodule $H^1_{\mathrm{ord}}(F_{\infty,\mathfrak{P}}, A)$ of $H^1(F_{\infty,\mathfrak{P}}, A)$ to be

$$\ker(H^1(F_{\infty,\mathfrak{P}}, A) \to H^1(I_\mathfrak{P}, A/A_\mathfrak{P})),$$

where $A_\mathfrak{P}$ is defined to be $A_\mathfrak{p}$ (and so only depends on $\mathfrak{p}$). Following [7], we then define the Selmer group $\mathrm{Sel}(F_\infty, A)$ for f over F_∞ as the kernel of the global-to-local restriction map

$$H^1(F_\infty, A) \to \prod_{\eta \nmid p} \frac{H^1(F_{\infty,\eta}, A)}{H^1_{\mathrm{ur}}(F_{\infty,\eta}, A)} \times \prod_{\mathfrak{P}|p} \frac{H^1(F_{\infty,\mathfrak{P}}, A)}{H^1_{\mathrm{ord}}(F_{\infty,\mathfrak{P}}, A)},$$

where η (respectively $\mathfrak{P}$) runs over the primes of F_∞ not dividing (respectively dividing) p. Note that, since we have assumed p is odd, the local cohomology groups for the Archimedean primes of F vanish, so we may, and do, ignore them.

In [7], in addition to the Selmer group for ordinary p-adic Galois representations, Greenberg also defined the (*a priori* smaller) *strict* Selmer group, requiring cocycles to be trivial away from p instead of unramified (but keeping the same local conditions at primes dividing p). We can define the strict Selmer group for f in the analogous way (replacing inertia groups with decomposition groups at the primes not dividing p). For a prime η of F_∞ lying over $v \notin \Sigma_p$, if v does not split completely in F_∞, then $G_{F_{\infty,\eta}}/I_\eta$ has pro-order prime to p, and as a result, the restriction homomorphism $H^1(F_{\infty,\eta}, A) \to H^1(I_\eta, A)$ is injective. Thus, for such a prime η, the strict local condition coincides with the unramified local condition. Therefore the Selmer group for f over F_∞ equals the strict Selmer group when no prime of F splits completely in F_∞ (e.g. when F_∞ is the cyclotomic $\mathbf{Z}_p$-extension of F), but these groups may differ otherwise. We will not have cause to consider the strict Selmer group, and will be content with the following result, which shows that the local conditions for the two groups can only differ at primes of F_∞ lying above a prime of F dividing the prime-to-p part of N (the level of our modular form f). The proof was explained to us by Matthew Emerton.

Proposition 3.3.1. *If $v \nmid pN$ is a prime of F, then*

$$\ker(H^1(F_v, A) \to H^1(I_v, A)) = 0.$$

Thus, if v splits in F_∞, then for any prime η of F_∞ lying over v, the strict local condition and the unramified local condition at η coincide. In particular, if every prime of F dividing the level of f is finitely decomposed in F_∞, then the Selmer group and the strict Selmer group coincide.

Proof. The $\mathscr{O}$-corank of $H^1_{\mathrm{ur}}(F_v, A) = \ker(H^1(F_v, A) \to H^1(I_v, A))$ is the same as the $\mathscr{O}$-corank of $H^0(F_v, A)$. Moreover, since $v \nmid pN$, A is an unramified G_{F_v}-module, so $H^1_{\mathrm{ur}}(F_v, A) = A/(\mathrm{Frob}_v - 1)A$ is $\mathscr{O}$-divisible and $H^0(F_v, A) = A^{\mathrm{Frob}_v = 1}$. Now, if ℓ is the rational prime of $\mathbf{Q}$ lying below v, then the eigenvalues of Frob_ℓ on V are Weil numbers of weight $(k-1)/2$. Since $k \geq 2$, we see that, in particular, these eigenvalues are not roots of unity. The eigenvalues of Frob_v on V are powers of the eigenvalues of Frob_ℓ since $\rho_f(\mathrm{Frob}_v)$ is conjugate to a power of $\rho_f(\mathrm{Frob}_\ell)$. Thus 1 is not an eigenvalue of Frob_v on V, so $V^{\mathrm{Frob}_v = 1} = 0$. It follows that $A^{\mathrm{Frob}_v = 1}$ has $\mathscr{O}$-corank zero. The same is then true of $H^1_{\mathrm{ur}}(F_v, A)$, which is therefore $\mathscr{O}$-divisible and finite, and hence trivial, proving the first statement. In light of the discussion preceding the proposition, it follows that the only primes w of F_∞ where the local conditions for the Selmer group and the strict Selmer group can differ are those lying over a prime v of F that divides the prime-to-p part of N and splits in F_∞. So, if there are no such primes, then the Selmer group and the strict Selmer group must coincide. $\square$

Let $\Lambda = \mathscr{O}[[\Gamma]]$ be the Iwasawa algebra of Γ with coefficients in $\mathscr{O}$ (see the beginning of §2.1 for the precise definition of this ring). The Galois group G_F acts (via conjugation) on the $\mathscr{O}$-module $H^1(F_\infty, A)$ with G_{F_∞} acting trivially, and this action allows us to regard the global cohomology group as a discrete Λ-module via Proposition 2.2.3 (that the hypotheses of that Proposition are satisfied follows from the canonical isomorphism $H^1(F_\infty, A) = \varinjlim_n H^1(F_n, A)$). The Selmer group $\mathrm{Sel}(F_\infty, A)$ is a Γ-stable, $\mathscr{O}$-submodule of $H^1(F_\infty, A)$, so it too can be regarded as a discrete Λ-module. Moreover, if Σ is a finite set of primes of F containing the Archimedean primes, the primes in Σ_p, and the primes where A is ramified, then we have an exact sequence

$$0 \to \mathrm{Sel}(F_\infty, A) \to H^1(F_\Sigma/F_\infty, A) \to \prod_{\eta | v \in \Sigma - \Sigma_p} \frac{H^1(F_{\infty,\eta}, A)}{H^1_{\mathrm{ur}}(F_{\infty,\eta}, A)} \times \prod_{\mathfrak{P}|p} \frac{H^1(F_{\infty,\mathfrak{P}}, A)}{H^1_{\mathrm{ord}}(F_{\infty,\mathfrak{P}}, A)}. \tag{3.1}$$

According to [7, Proposition 3], $H^1(F_\Sigma/F_\infty, A)$ is a cofinitely generated Λ-module in the sense of Definition 2.2.2, so the sequence (3.1) implies that $\mathrm{Sel}(F_\infty, A)$ is cofinitely generated as well, i.e., its $\mathscr{O}$-module Pontryagin dual $\widehat{\mathrm{Sel}(F_\infty, A)} = \mathrm{Hom}_{\mathscr{O}}(\mathrm{Sel}(F_\infty, A), E/\mathscr{O})$ is a finitely generated Λ-module. We define the Λ-corank (respectively the μ-, λ-invariant) of $\mathrm{Sel}(F_\infty, A)$ to be the Λ-rank (respectively the μ-, λ-invariant) of its Pontryagin dual (the μ- and λ-invariants of a general cofinitely generated Λ-module are defined in Definition 2.3.4). This terminology will be applied to any cofinitely generated Λ-module appearing below. We write $\mu(f)$ and $\lambda(f)$ for the Iwasawa invariants of $\mathrm{Sel}(F_\infty, A)$, and also refer to them as the Iwasawa invariants of f (over F_∞).

Because we will use the condition on the set Σ in the exact sequence (3.1) repeatedly, we formalize it in a definition.

Definition 3.3.1. A finite set Σ of primes of F will be said to be *sufficiently large* for A provided Σ contains the Archimedean primes, the primes above p, and any primes where A is ramified.

Most of our results will make use of the following hypothesis:

$$\mathrm{Sel}(F_\infty, A) \text{ is cotorsion over } \Lambda. \tag{tor}$$

When f corresponds to an elliptic curve over $\mathbf{Q}$ with good, ordinary reduction at the primes of Σ_p and F_∞ is the cyclotomic $\mathbf{Z}_p$-extension of F, (tor) was conjectured by Mazur in [10]. For $F = \mathbf{Q}$, this follows from work of Kato and Rohrlich. In the anticyclotomic setting, with $p \geq 5$, (tor) has been proved by Pollack and Weston ([13, Theorem 1.3]) for newforms of weight 2 and trivial character, under some technical hypotheses on $\bar{\rho}_f$ and the factorization of the level of f in F. Greenberg also has general conjectures about ordinary Selmer groups being cotorsion, most of which remain open.

In order to obtain more refined information about the structure of $\mathrm{Sel}(F_\infty, A)$, we need to give alternative descriptions of the local conditions that define it. This is carried out in the next subsection.

3.3.2 The Λ-module Structure of Local Cohomology Groups

We retain the notation and hypotheses of the previous subsection. For a prime $v \nmid p$ of F, we define

$$\mathcal{H}_v = \mathcal{H}_v(F_\infty, A) = \varinjlim_n \prod_{w \in \Sigma_{n,v}} \frac{H^1(F_{n,w}, A)}{H^1_{\mathrm{ur}}(F_{n,w}, A)},$$

where $\Sigma_{n,v}$ is the set of primes of F_n lying over v and the limit is taken with respect to the restriction maps in Galois cohomology. For a prime $\mathfrak{p}$ of Σ_p, we define

$$\mathcal{H}_{\mathfrak{p}} = \mathcal{H}_{\mathfrak{p}}(F_\infty, A) = \prod_{\mathfrak{P} \in \Sigma_{\infty,\mathfrak{p}}} \frac{H^1(F_{\infty,\mathfrak{P}}, A)}{H^1_{\mathrm{ord}}(F_{\infty,\mathfrak{P}}, A)},$$

where $\Sigma_{\infty,\mathfrak{p}}$ is the finite set of primes of F_∞ lying over $\mathfrak{p}$ (recall that we have assumed in the introduction (hypothesis p-(ii)) that all such $\mathfrak{p}$ are finitely decomposed in F_∞). Note that these $\mathscr{O}$-modules are in fact discrete Λ-modules (again by Proposition 2.2.3).

Proposition 3.3.2. *For any finite set Σ of primes of F which is sufficiently large for A (Definition 3.3.1), the sequence of Λ-modules*

$$0 \to \mathrm{Sel}(F_\infty, A) \to H^1(F_\Sigma/F_\infty, A) \to \prod_{v \in \Sigma - \Sigma_p} \mathcal{H}_v \times \prod_{\mathfrak{p} \in \Sigma_p} \mathcal{H}_{\mathfrak{p}}$$

is exact.

Proof. Any cohomology class $\kappa \in H^1(F_\infty, A)$ arises as the restriction of a cohomology class $\kappa_n \in H^1(F_n, A)$ for some $n \geq 0$. If η is a prime of F_∞ lying over $v \notin \Sigma_p$, and $w \in \Sigma_{n,v}$, then the restriction map $H^1(F_{n,w}, A)/H^1_{\mathrm{ur}}(F_{n,w}, A) \to H^1(F_{\infty,\eta}, A)/H^1_{\mathrm{ur}}(F_{\infty,\eta}, A)$ is injective because $F_{\infty,\eta}/F_{n,w}$ is unramified. The commutative diagram

$$\begin{array}{ccc}
H^1(F_\infty, A) & \longrightarrow & \dfrac{H^1(F_{\infty,\eta}, A)}{H^1_{\mathrm{ur}}(F_{\infty,\eta}, A)} \\
\uparrow & & \uparrow \\
H^1(F_n, A) & \longrightarrow & \dfrac{H^1(F_{n,w}, A)}{H^1_{\mathrm{ur}}(F_{n,w}, A)}
\end{array}$$

of restriction maps then shows that κ is unramified at η if and only if κ_n is unramified at w. This shows that the kernel of

$$H^1(F_\infty, A) \to \prod_{\eta | v} H^1(F_{\infty,\eta}, A)/H^1_{\mathrm{ur}}(F_{\infty,\eta}, A)$$

coincides with the kernel of $H^1(F_\infty, A) \to \mathcal{H}_v$ (the latter map sends κ to the natural image of κ_n in $\mathcal{H}_v$). In view of the definition of $\mathcal{H}_{\mathfrak{p}}$ for $\mathfrak{p} \in \Sigma_p$, we conclude that $\mathrm{Sel}(F_\infty, A)$ is exactly the kernel in question. □

We've introduced the modules $\mathcal{H}_v$ when $v \notin \Sigma_p$ to deal with the possibility that v splits in F_∞. For such a v, the product of the local cohomology groups over all primes of F_∞ lying over v does not have good properties as a Λ-module (it's too big). The Λ-module structure of $\mathcal{H}_v$, on the other hand, can be understood, even when v is split in F_∞. When v is finitely decomposed in F_∞, $\mathcal{H}_v$ is just a product of local cohomology groups, and the structure of these groups has been determined by Greenberg.

Proposition 3.3.3. *For a prime $v \notin \Sigma_p$ of F, let $\Sigma_{\infty,v}$ denote the set of primes of F_∞ lying above v.*

(i). *For a prime $v \notin \Sigma_p$ that is finitely decomposed in F_∞, we have*

$$\mathcal{H}_v \simeq \prod_{\eta \in \Sigma_{\infty,v}} H^1(F_{\infty,\eta}, A)$$

as Λ-modules, and $\mathcal{H}_v$ is a cofinitely generated, cotorsion Λ-module with μ-invariant zero and λ-invariant

$$\sum_{\eta \in \Sigma_{\infty,v}} \mathrm{corank}_{\mathscr{O}}(H^1(F_{\infty,\eta}, A)).$$

(ii). *For a prime* $\mathfrak{p} \in \Sigma_p$, $\mathcal{H}_{\mathfrak{p}}$ *is a cofinitely generated* Λ*-module with* Λ*-corank* $[F_{\mathfrak{p}} : \mathbf{Q}_p]$ *and* μ*-invariant zero.*

Proof. The isomorphism for $v \notin \Sigma_p$ holds because the number of primes in $\Sigma_{n,v}$ is constant for n sufficiently large (equal to the cardinality of $\Sigma_{\infty,v}$), because directed colimits commute with finite products, and because $H^1_{\mathrm{ur}}(F_{\infty,\eta}, A) = 0$ for $\eta \in \Sigma_{\infty,v}$. The assertions about the Λ-module structure of the products of local cohomology groups are then given by Proposition 1 (for $\mathfrak{p} \in \Sigma_p$) and Proposition 2 (for $v \notin \Sigma_p$) of [7]. □

Now consider a prime $v \notin \Sigma_p$ that splits in F_∞. Then we have an isomorphism $F_v \simeq F_{\infty,\eta}$ for any prime η of F_∞ lying over v, giving $H^1(F_v, A) \simeq H^1(F_\infty, A)$. The finiteness of $H^1(F_v, A[\varpi])$ shows that $H^1(F_v, A)[\varpi]$ is finite, hence that $H^1(F_v, A)$ is a cofinitely generated $\mathcal{O}$-module (this is similar to Proposition 2.3.3). In particular $H^1(F_v, A)/H^1_{\mathrm{ur}}(F_v, A)$ is a cofinitely generated $\mathcal{O}$-module, and by the local Euler characteristic formula, we have

$$\mathrm{corank}_{\mathcal{O}}(H^1(F_v, A)/H^1_{\mathrm{ur}}(F_v, A)) = \mathrm{rank}_{\mathcal{O}}(H^0(F_v, A^*)). \tag{3.2}$$

(A^* is the Cartier dual of A; see §3.1.) The $\mathcal{O}$-module structure of the quotient $H^1(F_v, A)/H^1_{\mathrm{ur}}(F_v, A)$ completely determines the Λ-module structure of $\mathcal{H}_v$.

Proposition 3.3.4. *Let* $v \notin \Sigma_p$ *be a prime of* F *that splits in* F_∞*, and choose an isomorphism of* $\mathcal{O}$*-modules*

$$H^1(F_v, A)/H^1_{\mathrm{ur}}(F_v, A) \simeq (E/\mathcal{O})^r \oplus \sum_{i=1}^{t} \mathcal{O}/\varpi^{m_i}\mathcal{O}$$

for some $r \geq 0$ and $m_i \geq 0$. Then we have

$$\mathcal{H}_v \simeq \widehat{\Lambda}^r \oplus \sum_{i=1}^{t} \widehat{\Lambda/\varpi^{m_i}\Lambda} \tag{3.3}$$

as Λ-modules, so $\mathcal{H}_v$ is a cofinitely generated Λ-module with Λ-corank equal to $\mathrm{rank}_{\mathscr{O}}(H^0(F_v, A^*))$*, μ-invariant $\sum_{i=1}^{t} m_i$, and λ-invariant zero.*

Proof. By definition, $\mathcal{H}_v = \varinjlim_n \prod_{w \in \Sigma_{n,v}} H^1(F_{n,w}, A)/H^1_{\mathrm{ur}}(F_{n,w}, A)$, with the limit taken with respect to the restriction maps. Because v splits in F_∞, a choice of prime w_n of F_n lying over v gives rise to an $\mathscr{O}[G_n]$-isomorphism

$$\prod_{w \in \Sigma_{n,v}} H^1(F_{n,w}, A)/H^1_{\mathrm{ur}}(F_{n,w}, A) \simeq (H^1(F_v, A)/H^1_{\mathrm{ur}}(F_v, A)) \otimes_{\mathscr{O}} \mathscr{O}[G_n].$$

Choosing the primes w_n compatibly for $n \geq 0$, these isomorphisms turn the transition maps defining $\mathcal{H}_v$ into the maps coming from corestriction on the right tensor factor. Thus we have a Λ-module isomorphism $\mathcal{H}_v \simeq \varinjlim_n (H^1(F_v, A)/H^1_{\mathrm{ur}}(F_v, A)) \otimes_{\mathscr{O}} \mathscr{O}[G_n]$, and the putative isomorphism (3.3) follows from Proposition 3.2.6. The Iwasawa invariants of $\mathcal{H}_v$ can be read off from this isomorphism, and the equality

$$\mathrm{corank}_\Lambda(\mathcal{H}_v) = \mathrm{rank}_{\mathscr{O}}(H^0(F_v, A^*))$$

follows from the isomorphism and Equation 3.2. □

3.3.3 Non-Primitive Selmer Groups

In this subsection, following Greenberg-Vatsal [9], we introduce non-primitive Selmer groups. A non-primitive Selmer group is defined by omitting some of the local conditions at primes of F_∞ not dividing ∞ or p. If we omit enough local conditions, the ϖ-torsion of the resulting non-primitive Selmer group for f can be identified with the corresponding non-primitive residual Selmer group. To be precise, let Σ_0 be a finite set of primes of F not containing any Archimedean primes or any primes of Σ_p. The Σ_0-non-primitive Selmer group $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$ is then defined as the kernel of the map

$$H^1(F_\infty, A) \to \prod_{\eta | v \notin \Sigma_0, v \nmid p} \frac{H^1(F_{\infty,\eta}, A)}{H^1_{\mathrm{ur}}(F_{\infty,\eta}, A)} \times \prod_{\mathfrak{P}|p} \frac{H^1(F_{\infty,\mathfrak{P}}, A)}{H^1_{\mathrm{ord}}(F_{\infty,\mathfrak{P}}, A)}.$$

As there is generally no risk of confusion about Σ_0 we will sometimes refer to $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$ simply as the non-primitive Selmer group. If Σ is a finite set of primes of F which is sufficiently large for A (Definition 3.3.1) and also contains Σ_0, then we have exact sequences of Λ-modules

$$0 \to \mathrm{Sel}^{\Sigma_0}(F_\infty, A) \to H^1(F_\Sigma/F_\infty, A) \to \prod_{v \in \Sigma - \Sigma_0} \mathcal{H}_v$$

and

$$0 \to \mathrm{Sel}(F_\infty, A) \to \mathrm{Sel}^{\Sigma_0}(F_\infty, A) \to \prod_{v \in \Sigma_0} \mathcal{H}_v.$$

We will show in §3.3.4 that under appropriate hypotheses, these exact sequences are exact on the right as well.

We now define a non-primitive Selmer group $\mathrm{Sel}^{\Sigma_0}(F_\infty, A[\varpi])$ for the residual representation $A[\varpi]$, in a manner analogous to the non-primitive Selmer group for A. Its definition is designed to recover the ϖ-torsion of $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$. For a prime $\mathfrak{P}$ of F_∞ lying over p, define

$$H^1_{\mathrm{ord}}(F_{\infty,\mathfrak{P}}, A[\varpi]) = \ker(H^1(F_{\infty,\mathfrak{P}}, A[\varpi]) \to H^1(I_{\mathfrak{P}}, A[\varpi]/A_{\mathfrak{P}}[\varpi])).$$

Then $\mathrm{Sel}^{\Sigma_0}(F_\infty, A[\varpi])$ is the kernel of the map

$$H^1(F_\infty, A[\varpi]) \to \prod_{\eta | v \notin \Sigma_0, v \nmid p} \frac{H^1(F_{\infty,\eta}, A[\varpi])}{H^1_{\mathrm{ur}}(F_{\infty,\eta}, A[\varpi])} \times \prod_{\mathfrak{P}|p} \frac{H^1(F_{\infty,\mathfrak{P}}, A[\varpi])}{H^1_{\mathrm{ord}}(F_{\infty,\mathfrak{P}}, A[\varpi])}.$$

The next proposition will involve the space $A[\varpi]_{I_\mathfrak{p}}$ of $I_\mathfrak{p}$-coinvariants of $A[\varpi]$. This is the largest $\mathbf{F}[G_{F_\mathfrak{p}}]$-quotient of $A[\varpi]$ on which $I_\mathfrak{p}$ acts trivially (recall that $\mathbf{F} = \mathscr{O}/\varpi$ is the residue field of $\mathscr{O}$). More explicitly, it is the quotient of $A[\varpi]$ by the $\mathbf{F}[G_{F_\mathfrak{p}}]$-submodule generated by elements of the form $ga - a$ for $g \in I_\mathfrak{p}$ and $a \in A[\varpi]$. It is in the proof of this proposition that we use the second part of assumption (ram) from the beginning of §3.3, that $A[\varpi]$ is ramified at each prime $\mathfrak{p} \in \Sigma_p$.

Proposition 3.3.5. *If $\mathfrak{p} \in \Sigma_p$, then $A[\varpi]/A_\mathfrak{p}[\varpi] = (A[\varpi])_{I_\mathfrak{p}}$.*

Proof. Since $A[\varpi]/A_\mathfrak{p}[\varpi] = (A/A_\mathfrak{p})[\varpi]$ is an unramified $\mathbf{F}[G_{F_\mathfrak{p}}]$-module, we have a surjective map $(A[\varpi])_{I_\mathfrak{p}} \to A[\varpi]/A_\mathfrak{p}[\varpi]$. It follows that $A[\varpi]_{I_\mathfrak{p}}$ is at least 1-dimensional over $\mathbf{F}$. Because $A[\varpi]$ is ramified at $\mathfrak{p}$ by assumption (ram) from the beginning of §3.3, $(A[\varpi])_{I_\mathfrak{p}}$ cannot be 2-dimensional. Thus the surjection $(A[\varpi])_{I_\mathfrak{p}} \to A[\varpi]/A_\mathfrak{p}[\varpi]$ is an equality. □

Remark 3.3.1. Proposition 3.3.5 shows that the local conditions defining the module $\mathrm{Sel}^{\Sigma_0}(F_\infty, A[\varpi])$ only depend on $A[\varpi]$ as a G_F-module (when *a priori* they depend on A as a G_F-module because the definition of the subspace $A_\mathfrak{p}[\varpi] \subseteq A[\varpi]$ for a prime $\mathfrak{p} \mid p$ makes reference to the G_F-module structure of A). This is clear at the primes not dividing p, where the local condition is the unramified one, and Proposition 3.3.5 shows that at a prime $\mathfrak{P}$ lying above $\mathfrak{p} \in \Sigma_p$, the local condition is the kernel of the map

$$H^1(F_{\infty,\mathfrak{P}}, A) \to H^1(I_{\mathfrak{P}}, (A[\varpi])_{I_{\mathfrak{p}}})$$

induced by the quotient map $A[\varpi] \to (A[\varpi])_{I_{\mathfrak{p}}}$ and restriction to $I_{\mathfrak{P}}$. The definition of the quotient $(A[\varpi])_{I_{\mathfrak{p}}}$ is given entirely in terms of the G_F-action on $A[\varpi]$ (even just the $G_{F_{\mathfrak{p}}}$-action). This observation is crucial to our method because it shows that, if we have two modular forms satisfying the appropriate hypotheses whose residual representations are isomorphic as G_F-modules, then the corresponding residual Selmer groups are isomorphic.

Proposition 3.3.6. *If Σ_0 contains all the primes of F dividing the tame level of f, then the natural map $H^1(F_\infty, A[\varpi]) \to H^1(F_\infty, A)$ induces an isomorphism of $\mathscr{O}$-modules*

$$\mathrm{Sel}^{\Sigma_0}(F_\infty, A[\varpi]) \simeq \mathrm{Sel}^{\Sigma_0}(F_\infty, A)[\varpi].$$

Proof. We have a commutative diagram

$$\begin{array}{ccc} H^1(F_\infty, A) & \longrightarrow & \prod_{\eta \mid v \notin \Sigma_0, v \neq p} H^1(I_\eta, A) \times \prod_{\mathfrak{P} \mid p} H^1(I_{\mathfrak{P}}, A/A_{\mathfrak{P}}) \\ \uparrow & & \uparrow \\ H^1(F_\infty, A[\varpi]) & \longrightarrow & \prod_{\eta \mid v \notin \Sigma_0, v \nmid p} H^1(I_\eta, A[\varpi]) \times \prod_{\mathfrak{P} \mid p} H^1(I_{\mathfrak{P}}, A[\varpi]/A_{\mathfrak{P}}[\varpi]) \end{array}$$

with the vertical maps coming from the inclusions $A[\varpi] \hookrightarrow A$ and $A[\varpi]/A_{\mathfrak{p}}[\varpi] = (A/A_{\mathfrak{p}})[\varpi] \hookrightarrow A/A_{\mathfrak{p}}$, and the horizontal maps coming from restriction. Note that the kernel of the top (respectively bottom) horizontal map is $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$

(respectively $\mathrm{Sel}^{\Sigma_0}(F_\infty, A[\varpi])$). Since $A[\varpi]$ is an irreducible $\mathbf{F}[G_F]$-module (by the first part of assumption (ram) from the beginning of §3.3), $H^0(F, A[\varpi]) = 0$, which implies that $H^0(F_\infty, A[\varpi]) = 0$ as F_∞/F is pro-p. Thus the kernel of $H^1(F_\infty, A[\varpi]) \to H^1(F_\infty, A)$, which is a quotient of $H^0(F_\infty, A)$, is zero. So the left-hand vertical map is injective. Its image is $H^1(F_\infty, A)[\varpi]$. The commutativity of the diagram shows that this vertical map takes $\mathrm{Sel}^{\Sigma_0}(F_\infty, A[\varpi])$ into $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)[\varpi]$. To see that the image is precisely $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)[\varpi]$, it therefore suffices to prove that the right-hand vertical arrow is injective. We do this by considering each factor map on the right. First consider a prime η of F_∞ which divides $v \notin \Sigma_0$, $v \nmid p$. Since η does not divide the level of f (as $v \notin \Sigma_0$), A is unramified at η, and the kernel of the map $H^1(I_\eta, A[\varpi]) \to H^1(I_\eta, A)$ is $A^{I_\eta}/\varpi A^{I_\eta} = A/\varpi A = 0$ (A is a divisible $\mathscr{O}$-module). Similarly, if $\mathfrak{P}$ is a prime of F_∞ dividing p, then, since $A/A_\mathfrak{P}$ is unramified at $\mathfrak{P}$, the kernel of the map $H^1(I_\mathfrak{P}, A[\varpi]/A_\mathfrak{P}[\varpi]) \to H^1(I_\mathfrak{P}, A/A_\mathfrak{P})$ is $(A/A_\mathfrak{P})/\varpi(A/A_\mathfrak{P}) = 0$, because $A/A_\mathfrak{P}$ is divisible. □

Combining Proposition 3.3.6 with Remark 3.3.1, we conclude that for Σ_0 containing the primes dividing the tame level of f, the module $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)[\varpi]$ only depends on $A[\varpi]$ as an $\mathbf{F}[G_F]$-module.

3.3.4 Global-to-Local Maps

In this subsection, we establish the surjectivity of a global-to-local map of Galois cohomology (under appropriate hypotheses) which allows us to compare the Iwasawa invariants of the non-primitive and the primitive Selmer groups of f.

Recall that we have an exact sequence of Λ-modules

$$0 \to \mathrm{Sel}(F_\infty, A) \to H^1(F_\Sigma/F_\infty, A) \xrightarrow{\gamma} \prod_{v \in \Sigma - \Sigma_p} \mathcal{H}_v \times \prod_{\mathfrak{p} \in \Sigma_p} \mathcal{H}_{\mathfrak{p}} \qquad (3.4)$$

for Σ sufficiently large for A (Definition 3.3.1). By [7, Proposition 3],

$$\mathrm{corank}_\Lambda(H^1(F_\Sigma/F_\infty, A)) \geq \sum_{v \text{ real}} d_v^-(V) + 2r_2,$$

where the first sum is over the real primes of F, $d_v^-(V)$ is the dimension of the -1-eigenspace for a complex conjugation above v acting on V, and r_2 is the number of complex primes of F. Because ρ_f is odd, that is, the determinant of ρ_f of any complex conjugation is -1, $d_v^-(V) = 1$ for any any real prime v, as otherwise the determinant of ρ_f of a complex conjugation would be 1. So, letting r_1 denote the number of real primes of F, the inequality above becomes

$$\mathrm{corank}_\Lambda(H^1(F_\Sigma/F_\infty, A)) \geq r_1 + 2r_2 = [F : \mathbf{Q}].$$

The Λ-coranks of the factors comprising the target of the map γ of (3.4) are determined by the corresponding primes:

(nsplit) if $v \notin \Sigma_p$ is finitely decomposed in F_∞, then $\mathcal{H}_v$ is Λ-cotorsion

(split) f $v \notin \Sigma_p$ is split in F_∞, then $\mathrm{corank}_\Lambda(\mathcal{H}_v) = \mathrm{rank}_{\mathscr{O}}(H^0(F_v, A^*))$

(pnsplit) if $\mathfrak{p} \in \Sigma_p$, then $\mathrm{corank}_\Lambda(\mathcal{H}_{\mathfrak{p}}) = [F_{\mathfrak{p}} : \mathbf{Q}_p]$

Assertions (nsplit) and (pnsplit) are restatements of parts of Proposition 3.3.3, while (split) is a restatement of part of Proposition 3.3.4. It follows that the target of the map γ has Λ-corank at least

$$\sum_{\mathfrak{p}\in\Sigma_p}[F_{\mathfrak{p}}:\mathbf{Q}_p] = [F:\mathbf{Q}],$$

and if $H^0(F_v, A^*)$ is finite for each prime $v \in \Sigma$ that splits in F_∞, this is exactly the Λ-corank of the target of γ.

In proving the following proposition, several variants of which have appeared in the literature, we follow the proof of [20, Proposition 1.8]. It will be clear that the second hypothesis in the statement of Proposition 3.3.7 is used to ensure that the argument still goes through when finite primes split in F_∞. We will use the observation that irreducibility of the $\mathbf{F}[G_F]$-module $A[\varpi]$ (assumed as part of (ram) at the beginning of §3.3) implies that

$$H^0(F_\infty, A^* \otimes_{\mathscr{O}} E/\mathscr{O}) = 0.$$

Indeed, the ϖ-torsion of $H^0(F, A^* \otimes_{\mathscr{O}} E/\mathscr{O})$ is the space of G_F-invariants of the Tate twist of the Cartier dual of $A[\varpi]$, $\mathrm{Hom}_{\mathbf{F}}(A[\varpi], \mathscr{O}(1)/\varpi\mathscr{O}(1))$, which is zero since $A[\varpi]$ and $\mathscr{O}(1)/\varpi\mathscr{O}(1)$ are irreducible of different $\mathbf{F}$-dimension. Thus there are no G_{F_∞}-invariants either, since $\mathrm{Gal}(F_\infty/F)$ is pro-p.

Proposition 3.3.7. *Let Σ be a finite set of primes of F which is sufficiently large for A (Definition 3.3.1). Assume that*

(i). *hypothesis* (tor) *(introduced at the end of §3.3.1) holds, and* (ii). *for each prime $v \in \Sigma$ that splits in F_∞, $H^0(F_v, A^*) = 0$.*

Then the sequence

$$0 \to \mathrm{Sel}(F_\infty, A) \to H^1(F_\Sigma/F_\infty, A) \xrightarrow{\gamma} \prod_{v \in \Sigma - \Sigma_p} \mathcal{H}_v \times \prod_{\mathfrak{p} \in \Sigma_p} \mathcal{H}_{\mathfrak{p}} \to 0$$

is exact.

Proof. We need to prove that γ is surjective, i.e. that $\mathrm{coker}(\gamma) = 0$. We will do this by considering similar global-to-local maps at the finite levels F_n of F_∞ and passing to the limit. To ensure that the kernels and cokernels of the finite level global-to-local maps are finite and trivial, respectively, we will twist the Galois module structures under consideration by a character. Let $\kappa : \Gamma \simeq 1 + p\mathbf{Z}_p$ be an isomorphism of topological groups, which we also regard as a character of G_F and of $\mathrm{Gal}(F_\Sigma/F)$. If S is a discrete $\mathscr{O}$-module with a continuous $\mathscr{O}$-linear $\mathrm{Gal}(F_\Sigma/F)$-action, then S_t will denote $S \otimes_{\mathscr{O}} \mathscr{O}(\kappa^t)$ for $t \in \mathbf{Z}$. This is also a discrete $\mathscr{O}$-module with a continuous $\mathscr{O}$-linear $\mathrm{Gal}(F_\Sigma/F)$-action, and if S is a discrete Λ-module, S_t is as well (with Λ acting on both tensor factors). We have $S \simeq S_t$ as $\mathscr{O}[\mathrm{Gal}(F_\Sigma/F_\infty)]$-modules, and if S is a discrete Λ-module, S_t is isomorphic to the Pontryagin dual of $\widehat{S}(\kappa^{-t})$ (see the paragraph preceding Proposition 3.2.7 for the definition of this last Λ-module).

For each $t \in \mathbf{Z}$, we define a Selmer group $\mathrm{Sel}(F_\infty, A_t)$ for A_t as the kernel of the map

$$H^1(F_\Sigma/F_\infty, A_t) \xrightarrow{\gamma_t} \prod_{v \in \Sigma - \Sigma_p} \mathcal{H}_v(F_\infty, A_t) \times \prod_{\mathfrak{p} \in \Sigma_p} \mathcal{H}_{\mathfrak{p}}(F_\infty, A_t),$$

where the factors of the target are defined analogously to those defined at the beginning of §3.3.4 for A, setting $A_{t,\mathfrak{p}} = (A_{\mathfrak{p}})_t$ for $\mathfrak{p} \in \Sigma_p$. With this definition, $\mathrm{Sel}(F_\infty, A_t) \simeq \mathrm{Sel}(F_\infty, A)_t$ as Λ-modules, and because $A \simeq A_t$ as $\mathrm{Gal}(F_\Sigma/F_\infty)$-modules, we have $\mathrm{coker}(\gamma) \simeq \mathrm{coker}(\gamma_t)$ as Λ-modules. It therefore suffices to prove that $\mathrm{coker}(\gamma_t)$ vanishes for some $t \in \mathbf{Z}$.

We will prove vanishing of some $\mathrm{coker}(\gamma_t)$ by working with Selmer groups over F_n for $n \geq 0$ and taking a limit. For $t \in \mathbf{Z}$, $n \geq 0$, and a prime $\mathfrak{p}$ of F_n dividing $\mathfrak{p}_0 \in \Sigma_p$, we set

$$H^1_{\mathrm{ord}}(F_{n,\mathfrak{p}}, A_t) = \ker(H^1(F_{n,\mathfrak{p}}, A_t) \to H^1(F_{n,\mathfrak{p}}, A_t/A_{t,\mathfrak{p}})),$$

where $A_{t,\mathfrak{p}} = A_{t,\mathfrak{p}_0}$. Note that this is stronger than the analogous ordinary local condition over F_∞ as we are using decomposition groups instead of inertia groups. Because we are using decomposition groups, we can apply Poitou-Tate global duality to each finite level Selmer group $\mathrm{Sel}(F_n, A_t)$, defined as the kernel of the map

$$H^1(F_\Sigma/F_n, A_t) \xrightarrow{\gamma_{n,t}} \prod_{w|v\in\Sigma, v\nmid p} \frac{H^1(F_{n,w}, A_t)}{H^1_{\mathrm{ur}}(F_{n,w}, A_t)} \times \prod_{\mathfrak{p}|p} \frac{H^1(F_{n,\mathfrak{p}}, A_t)}{H^1_{\mathrm{ord}}(F_{n,\mathfrak{p}}, A_t)}.$$

Upon taking the direct limit of the maps $\gamma_{n,t}$ over $n \geq 0$, we get maps $\gamma_{\infty,t}$ with source $H^1(F_\Sigma/F_\infty, A_t)$, such that $\mathrm{coker}(\gamma_t)$ is a Λ-module quotient of $\mathrm{coker}(\gamma_{\infty,t})$ (it is a quotient because we used decomposition groups to define the local conditions at the primes dividing p for the finite level Selmer groups). We will prove that for an appropriate choice of t, the $\mathscr{O}$-modules $\mathrm{coker}(\gamma_{n,t})$ are trivial for all $n \geq 0$. The desired result will follow from this.

We will impose several conditions on the integers t under consideration. Because $\mathrm{Sel}(F_\infty, A)$ is Λ-cotorsion, Proposition 3.2.8 implies that for all but finitely many t, $\mathrm{Sel}(F_\infty, A_t)^{\Gamma_n} \simeq \mathrm{Sel}(F_\infty, A)_t^{\Gamma_n}$ will be finite for all $n \geq 0$. Similarly, because $A(F_\infty)$ is Λ-cotorsion and $A(F_\infty)_t = A_t(F_\infty)$, for all but finitely many t, $H^0(F_n, A_t) = A_t(F_\infty)^{\Gamma_n}$ will be finite for all $n \geq 0$. We assume from now on that t satisfies these conditions, which imply the following one:

(a) $\ker(\gamma_{n,t}) = \mathrm{Sel}(F_n, A_t)$ is finite for all $n \geq 0$.

To see this, observe that the restriction map $H^1(F_n, A_t) \to H^1(F_\infty, A_t)$ takes $\ker(\gamma_{n,t})$ into $\mathrm{Sel}(F_\infty, A_t)^{\Gamma_n}$, which we have assumed finite. The kernel of the restriction map is $H^1(F_\infty/F_n, A_t(F_\infty))$, which has the same $\mathscr{O}$-corank as $H^0(F_n, A_t)$, also assumed finite. So, indeed, $\ker(\gamma_{n,t})$ is finite for all $n \geq 0$. We now wish to impose three additional conditions on t:

(b) for $n \geq 0$ and $w \in \Sigma_{n,v}$ with $v \nmid p$, $H^0(F_{n,w}, A_t^*)$ is finite,

(c) for $n \geq 0$ and $\mathfrak{p} \mid \mathfrak{p}_0 \in \Sigma_p$, $H^0(F_{n,\mathfrak{p}}, A_t/A_{t,\mathfrak{p}})$ and $H^0(F_{n,\mathfrak{p}}, (A_t/A_{t,\mathfrak{p}})^*)$ are finite, and

(d) for $n \geq 0$ and $\mathfrak{p} \mid \mathfrak{p}_0 \in \Sigma_p$, $H^0(F_{n,\mathfrak{p}}, (A_{t,\mathfrak{p}})^*)$ is finite.

All three of these conditions will hold for all but finitely many t. For (c) and (d), this follows from Proposition 3.2.8 applied to the Iwasawa algebra of the image of $G_{F_{\mathfrak{p}_0}}$ in Γ, which is non-trivial as we have assumed (p-(ii) of §3.1) that no prime dividing p splits in F_∞ in p (we are using that the modules of coinvariants and invariants of a Λ-module that is finitely generated over $\mathscr{O}$ have the same $\mathscr{O}$-rank). Note that the conditions involving finiteness of the local invariants at each level of the Cartier dual modules are equivalent to the vanishing of the local invariants, since the Cartier dual modules are finite free over $\mathscr{O}$. That condition (b) holds for all but finitely many t follows from Proposition 3.2.8 as before, except when v is a prime that splits in F_∞, in which case $H^0(F_{n,w}, A_t^*)$ can be identified with $H^0(F_v, A^*)$, which vanishes by hypothesis (we cannot argue that $H^0(F_v, A^*)$ has to vanish for such v as before because the image of G_{F_v} in Γ is trivial).

For $n \geq 0$ and $w \mid v \in \Sigma$ $v \notin \Sigma_p$, Tate local duality and the local Euler characteristic formula give

$$\operatorname{corank}_{\mathscr{O}} \left(\frac{H^1(F_{n,w}, A_t)}{H^1_{\mathrm{ur}}(F_{n,w}, A_t)} \right) = \operatorname{rank}_{\mathscr{O}}(H^0(F_{n,w}, A_t^*)) = 0, \tag{3.5}$$

the last equality coming from condition (b). Similarly, (c) implies that the module $H^1(F_{n,\mathfrak{p}}, A_t/A_{t,\mathfrak{p}})$ has $\mathscr{O}$-corank equal to $[F_{n,\mathfrak{p}} : \mathbf{Q}_p]$ for $n \geq 0$ and $\mathfrak{p} \mid p$. Condition (d) then implies that

$$\operatorname{corank}_{\mathscr{O}} \left(\frac{H^1(F_{n,\mathfrak{p}}, A_t)}{H^1_{\mathrm{ord}}(F_{n,\mathfrak{p}}, A_t)} \right) = [F_{n,\mathfrak{p}} : \mathbf{Q}_p] \tag{3.6}$$

for all $n \geq 0$ and all such $\mathfrak{p}$ as well.

A modification of the Poitou-Tate exact sequence gives the exact sequence

$$0 \to \mathrm{Sel}(F_n, A_t) \to H^1(F_\Sigma/F_n, A_t) \xrightarrow{\gamma_{n,t}} \prod_{w|v\in\Sigma, v\nmid p} \frac{H^1(F_{n,w}, A_t)}{H^1_{\mathrm{ur}}(F_{n,w}, A_t)} \times \prod_{\mathfrak{p}|p} \frac{H^1(F_{n,\mathfrak{p}}, A_t)}{H^1_{\mathrm{ord}}(F_{n,\mathfrak{p}}, A_t)} \to H_{1,n} \to H_{2,n} \to 0,$$

where $H_{1,n}$ is dual to a submodule of $H^1(F_\Sigma/F_n, A_t^*)$ and $H_{2,n}$ is a submodule of $H^2(F_\Sigma/F_n, A_t)$. The global Euler characteristic formula shows that

$$\mathrm{corank}_{\mathscr{O}}(H^1(F_\Sigma/F_n, A_t)) = p^n[F : \mathbf{Q}] + \mathrm{corank}_{\mathscr{O}}(H^2(F_\Sigma/F_n, A_t)). \tag{3.7}$$

Equations (3.5) and (3.6) give, for all $n \geq 0$,

$$\begin{aligned}
\mathrm{corank}_{\mathscr{O}}\left(\prod_{w|v\in\Sigma, v\nmid p} \frac{H^1(F_{n,w}, A_t)}{H^1_{\mathrm{ur}}(F_{n,w}, A_t)} \times \prod_{\mathfrak{p}|p} \frac{H^1(F_{n,\mathfrak{p}}, A_t)}{H^1_{\mathrm{ord}}(F_{n,\mathfrak{p}}, A_t)}\right) &= \mathrm{corank}_{\mathscr{O}}\left(\prod_{\mathfrak{p}|p} \frac{H^1(F_{n,\mathfrak{p}}, A_t)}{H^1_{\mathrm{ord}}(F_{n,\mathfrak{p}}, A_t)}\right) \\
&= \sum_{\mathfrak{p}|p} [F_{n,\mathfrak{p}} : \mathbf{Q}_p] \\
&= \sum_{\mathfrak{p}_0\in\Sigma_p} \sum_{\mathfrak{p}|\mathfrak{p}_0} [F_{n,\mathfrak{p}} : F_{\mathfrak{p}_0}][F_{\mathfrak{p}_0} : \mathbf{Q}_p] \\
&= \sum_{\mathfrak{p}_0\in\Sigma_p} [F_{\mathfrak{p}_0} : \mathbf{Q}_p]\left(\sum_{\mathfrak{p}|\mathfrak{p}_0} [F_{n,\mathfrak{p}} : F_{\mathfrak{p}_0}]\right) \\
&= \sum_{\mathfrak{p}_0\in\Sigma_p} [F_{\mathfrak{p}_0} : \mathbf{Q}_p] p^n = p^n[F : \mathbf{Q}].
\end{aligned}$$

Therefore, since $\mathrm{Sel}(F_n, A_t)$ is finite for all $n \geq 0$ by (a), the exact sequence above and Equation (3.7) imply that

$$\mathrm{corank}_{\mathscr{O}}(H^1(F_\Sigma/F_n, A_t)) = [F : \mathbf{Q}]p^n$$

and

$$\mathrm{corank}_{\mathscr{O}}(H^2(F_\Sigma/F_n, A_t)) = 0.$$

In particular, $\mathrm{coker}(\gamma_{n,t})$ and $H_{2,n}$ are finite, and hence so is $H_{1,n}$. Moreover, the order of $\mathrm{coker}(\gamma_{n,t})$ is bounded above by that of $H_{1,n}$, which, being finite, is dual to a submodule of $H^1(F_\Sigma/F_n, A_t^*)[\varpi^\infty]$. The $\mathscr{O}$-torsion submodule of $H^1(F_\Sigma/F_n, A_t^*)$ is a quotient of $H^0(F_n, A_t^* \otimes E/\mathscr{O})$, which in turn is a submodule of

$$H^0(F_\infty, A_t^* \otimes_{\mathscr{O}} E/\mathscr{O}) = H^0(F_\infty, A^* \otimes_{\mathscr{O}} E/\mathscr{O}).$$

But the latter group is trivial by the remarks preceding the proposition. Thus $\mathrm{coker}(\gamma_{n,t}) = 0$ for all $n \geq 0$. □

Corollary 3.3.8. *Let Σ be a finite set of primes of F that is sufficiently large for A (Definition 3.3.1). Assume that*

(i). *hypothesis* (tor) *(introduced at the end of §3.3.1) holds, and*

(ii). *for each prime $v \in \Sigma$ that splits in F_∞, $H^0(F_v, A^*) = 0$.*

If Σ_0 is a subset of Σ not containing any Archimedean primes or any primes above p, then $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$ is Λ-cotorsion and the sequences

$$0 \to \mathrm{Sel}^{\Sigma_0}(F_\infty, A) \to H^1(F_\Sigma/F_\infty, A) \to \prod_{v \in \Sigma - \Sigma_0 - \Sigma_p} \mathcal{H}_v \times \prod_{\mathfrak{p} \in \Sigma_p} \mathcal{H}_{\mathfrak{p}} \to 0$$

and

$$0 \to \mathrm{Sel}(F_\infty, A) \to \mathrm{Sel}^{\Sigma_0}(F_\infty, A) \to \prod_{v \in \Sigma_0} \mathcal{H}_v \to 0$$

are exact.

Proof. The first sequence is exact by definition, except at the right, where it is exact by Proposition 3.3.7. The second sequence is also exact by definition except for the surjectivity of the final map. The hypotheses together with Propositions 3.3.3 and 3.3.4 imply that the target of that map is Λ-cotorsion. Thus the fact that $\mathrm{Sel}(F_\infty, A)$ is cotorsion implies the same for $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$. Finally, since $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$ is exactly the inverse image of $\prod_{v\in\Sigma_0}\mathcal{H}_v$ in $H^1(F_\Sigma/F_\infty, A)$, exactness of the second sequence on the right follows from that of the first. □

Remark 3.3.2. The most interesting case of the preceding proposition is when Σ consists of Σ_0, together with the infinite primes and the primes above p.

3.3.5 Divisibility of the Non-Primitive Selmer group

The main result in this subsection is that, under the hypotheses of Proposition 3.3.7, the Σ_0-non-primitive Selmer group, for Σ_0 containing the primes dividing the tame level of f, has no proper Λ-submodules of finite index. First we deduce the corresponding result for $H^1(F_\Sigma/F_\infty, A)$, where Σ is sufficiently large for A.

Corollary 3.3.9. *Let Σ be a finite set of primes of F which is sufficiently large for A (Definition 3.3.1). Assume that*

(i). *hypothesis* (tor) *(introduced at the end of §3.3.1) holds, and*

(ii). *for each prime $v \in \Sigma$ that splits in F_∞, $H^0(F_v, A^*) = 0$.*

Then the Λ-corank of $H^1(F_\Sigma/F_\infty, A))$ is $[F : \mathbf{Q}]$, and $H^1(F_\Sigma/F_\infty, A)$ has no proper Λ-submodules of finite index.

Proof. By Proposition 3.3.7 and the hypothesis that $\mathrm{Sel}(F_\infty, A)$ is cotorsion, we have

$$\mathrm{corank}_\Lambda(H^1(F_\Sigma/F_\infty, A)) = \mathrm{corank}_\Lambda\left(\prod_{v\in\Sigma-\Sigma_p}\mathcal{H}_v \times \prod_{\mathfrak{p}\in\Sigma_p}\mathcal{H}_\mathfrak{p}\right).$$

The discussion at the beginning of §3.3.4 ((nsplit), (split), and (pnsplit)) shows that, in the presence of our hypothesis (ii),

$$\begin{aligned}\mathrm{corank}_\Lambda\left(\prod_{v\in\Sigma-\Sigma_p}\mathcal{H}_v \times \prod_{\mathfrak{p}\in\Sigma_p}\mathcal{H}_\mathfrak{p}\right) &= \mathrm{corank}_\Lambda\left(\prod_{\mathfrak{p}\in\Sigma_p}\mathcal{H}_\mathfrak{p}\right)\\ &= \sum_{\mathfrak{p}\in\Sigma_p}[F_\mathfrak{p} : \mathbf{Q}_p] = [F : \mathbf{Q}].\end{aligned}$$

This proves the first assertion. Now we invoke Propositions 3, 4, and 5 of [7]. Proposition 3 implies that $H^2(F_\Sigma/F_\infty, A)$ is Λ-cotorsion, while Proposition 4 implies that $H^2(F_\Sigma/F_\infty, A)$ is Λ-cofree. Thus $H^2(F_\Sigma/F_\infty, A) = 0$, and now Proposition 5 implies that $H^1(F_\Sigma/F_\infty, A)$ has no proper Λ-submodules of finite index. □

The next lemma will allow us to deduce the desired property of the non-primitive Selmer group from the corresponding property of $H^1(F_\Sigma/F_\infty, A)$.

Lemma 3.3.10. *Let Y be a finitely generated Λ-module, Z a free Λ-submodule. If Y contains no non-zero, finite Λ-submodules, then the same is true for Y/Z.*

Proof. See the proof of Lemma 2.6 of [9]. □

Lemma 3.3.11. *Let $0 \to X' \to X \to X'' \to 0$ be a short exact sequence of finitely generated Λ-modules with X free over Λ and X'' finitely generated and free over $\mathscr{O}$. Then X' is free over Λ.*

Proof. A finitely generated Λ-module is free if and only if its module of invariants vanishes and its module of coinvariants is $\mathscr{O}$-free ([11, Proposition 5.3.19 (ii)]). Thus it suffices to show that $(X')^\Gamma = 0$ and that X'_Γ is $\mathscr{O}$-free. Applying the snake lemma to the endomorphism of the short exact sequence $0 \to X' \to X \to X'' \to 0$ given by multiplication by $g - 1$, where $g \in \Gamma$ is a topological generator, we get an exact sequence

$$0 \to (X')^\Gamma \to X^\Gamma \to (X'')^\Gamma \to X'_\Gamma \to X_\Gamma \to X''_\Gamma \to 0. \tag{3.8}$$

Since X is Λ-free, $X^\Gamma = 0$, and it follows that $(X')^\Gamma = 0$. Taking $Y = \mathrm{im}(X'_\Gamma \to X_\Gamma)$, we deduce from (3.8) and the vanishing of X^Γ the exact sequence

$$0 \to (X'')^\Gamma \to X'_\Gamma \to Y \to 0. \tag{3.9}$$

Because X'' (respectively X_Γ) is finitely generated and $\mathscr{O}$-free, so is its submodule $(X'')^\Gamma$ (respectively Y). Thus the sequence of $\mathscr{O}$-modules (3.9) splits, and we find that X'_Γ is $\mathscr{O}$-free, being isomorphic to a direct sum of $\mathscr{O}$-free modules. □

In the proof of the next proposition, we closely follow the argument for Proposition 2.5 of [9].

Proposition 3.3.12. *Let Σ_0 be a finite set of primes of F not containing any Archimedean primes or any primes of Σ_p. Assume that*

(i). Σ_0 *contains the primes dividing the tame level of f,* (ii). *hypothesis* (tor) *(introduced at the end of §3.3.1) holds, and*

(iii). *for each prime* $v \in \Sigma_0$ *that splits in* F_∞, $H^0(F_v, A^*) = 0$.

Then $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$ *has no proper* Λ*-submodules of finite index.*

Proof. Let Σ be the union of Σ_0, Σ_p, and the set of Archimedean primes. Then Σ is sufficiently large for A and the hypotheses of Proposition 3.3.8 hold, so we have an exact sequence of Λ-modules

$$0 \to \mathrm{Sel}^{\Sigma_0}(F_\infty, A) \to H^1(F_\Sigma/F_\infty, A) \to \prod_{\mathfrak{p}\in\Sigma_p} \mathcal{H}_\mathfrak{p} \to 0.$$

Since $H^1(F_\Sigma/F_\infty, A)$ has no proper Λ-submodules of finite index by Corollary 3.3.9, if we can prove that $\prod_{\mathfrak{p}\in\Sigma_p} \mathcal{H}_\mathfrak{p}$ is Λ-cofree, the result will follow from Lemma 3.3.10, with $Y = H^1(\widehat{F_\Sigma/F_\infty}, A)$, $Z = \prod_{\mathfrak{p}\in\Sigma_p} \widehat{\mathcal{H}_\mathfrak{p}}$, and $Y/Z \simeq \mathrm{Sel}^{\widehat{\Sigma_0(F_\infty}, f)}$. Following the proof of Proposition 2.5 of [9], we will prove that for each $\mathfrak{p} \in \Sigma_p$,

$$\mathcal{H}_\mathfrak{p} = \prod_{\mathfrak{P}|\mathfrak{p}} H^1(F_{\infty,\mathfrak{P}}, A)/H^1_{\mathrm{ord}}(F_{\infty,\mathfrak{P}}, A)$$

is Λ-cofree.

Fix $\mathfrak{p} \in \Sigma_p$ and let $D = A/A_\mathfrak{p}$. We first prove that $H^1(F_\mathfrak{p}, D)$ is $\mathscr{O}$-cofree. The cohomology sequence associated to the multiplication-by-ϖ map on D yields an injection

$$H^1(F_\mathfrak{p}, D)/\varpi H^1(F_\mathfrak{p}, D) \hookrightarrow H^2(F_\mathfrak{p}, D[\varpi]). \tag{3.10}$$

The target of (3.10) is Cartier dual (as a finite p-group) to $H^0(F_{\mathfrak{p}}, \mathrm{Hom}(D[\varpi], \mu_p))$. If $\varphi : D[\varpi] \to \mu_p$ were a non-zero, hence surjective $G_{F_{\mathfrak{p}}}$-equivariant homomorphism, then because $D[\varpi]$ is unramified at $\mathfrak{p}$, μ_p would be unramified at $\mathfrak{p}$ as well. But our hypothesis p-(i) of §3.1 that $e(\mathfrak{p}/p) < p-1$ shows that μ_p is a ramified $G_{F_{\mathfrak{p}}}$-module. So the module of $G_{F_{\mathfrak{p}}}$-invariants of $\mathrm{Hom}(D[\varpi], \mu_p)$ must vanish, and thus $H^2(F_{\mathfrak{p}}, D[\varpi]) = 0$. By (3.10), $H^1(F_{\mathfrak{p}}, D)$ is $\mathscr{O}$-divisible. Since it is cofinitely generated as an $\mathscr{O}$-module, it is then $\mathscr{O}$-cofree. By the local Euler characteristic formula, we have

$$\begin{aligned}\mathrm{corank}_{\mathscr{O}}(H^1(F_{\mathfrak{p}}, D)) &= [F_{\mathfrak{p}} : \mathbf{Q}_p] + \mathrm{corank}_{\mathscr{O}}(H^0(F_{\mathfrak{p}}, D)) + \mathrm{corank}_{\mathscr{O}}(H^2(F_{\mathfrak{p}}, D)) \\ &= [F_{\mathfrak{p}} : \mathbf{Q}_p] + \mathrm{corank}_{\mathscr{O}}(H^0(F_{\mathfrak{p}}, D)),\end{aligned}$$

where the last equality holds because the vanishing of $H^2(F_{\mathfrak{p}}, D[\varpi])$ implies that of $H^2(F_{\mathfrak{p}}, D)$ (the ϖ-torsion of the latter is a homomorphic image of the former).

Now fix a prime $\mathfrak{P}$ of F_∞ lying over $\mathfrak{p}$, let $\Gamma_{\mathfrak{p}} \subseteq \Gamma$ be the decomposition group for $\mathfrak{p}$ (which is non-trivial because $\mathfrak{p}$ does not split in F_∞), and let $\Lambda_{\mathfrak{p}} = \mathscr{O}[[\Gamma_{\mathfrak{p}}]]$ be the Iwaswawa algebra of $\Gamma_{\mathfrak{p}}$. Because $\Gamma_{\mathfrak{p}}$ has cohomological dimension 1, we have an inflation-restriction sequence

$$0 \to H^1(F_{\infty,\mathfrak{P}}/F_{\mathfrak{p}}, D^{G_{F_{\infty,\mathfrak{P}}}}) \to H^1(F_{\mathfrak{p}}, D) \to H^1(F_{\infty,\mathfrak{P}}, D)^{\Gamma_{\mathfrak{p}}} \to 0.$$

The $\mathscr{O}$-corank of $H^1(F_{\infty,\mathfrak{P}}/F_{\mathfrak{p}}, D^{G_{F_{\infty,\mathfrak{P}}}})$ is the same as the $\mathscr{O}$-corank of $H^0(F_{\mathfrak{p}}, D)$, and from this it follows that $H^1(F_{\infty,\mathfrak{P}}, D)^{\Gamma_{\mathfrak{p}}}$ is $\mathscr{O}$-cofree of corank $[F_{\mathfrak{p}} : \mathbf{Q}_p]$. By Proposition 1 of [7], the $\Lambda_{\mathfrak{p}}$-corank of $H^1(F_{\infty,\mathfrak{P}}, D)$ is also $[F_{\mathfrak{p}} : \mathbf{Q}_p]$. An application of Nakayama's lemma now shows that $H^1(F_{\infty,\mathfrak{P}}, D)$ is $\Lambda_\mathfrak{p}$-cofree of corank $[F_\mathfrak{p} : \mathbf{Q}_p]$.

As $G_{F_{\infty,\mathfrak{P}}}$ has p-cohomological dimension 1, the map $H^1(F_{\infty,\mathfrak{P}}, A) \to H^1(F_{\infty,\mathfrak{P}}, D)$ is surjective. We therefore have

$$\mathcal{H}_\mathfrak{P} = H^1(F_{\infty,\mathfrak{P}}, A)/H^1_{\mathrm{ord}}(F_{\infty,\mathfrak{P}}, A) \simeq \mathrm{im}(H^1(F_{\infty,\mathfrak{P}}, D) \to H^1(I_\mathfrak{P}, D)).$$

Since D is unramified at $\mathfrak{p}$, the kernel of the restriction map to $I_\mathfrak{P}$ is equal to $H^1(G_{F_{\infty,\mathfrak{P}}}/I_\mathfrak{P}, D)$. If $\mathfrak{p}$ is unramified in F_∞, then $G_{F_{\infty,\mathfrak{P}}}/I_\mathfrak{P}$ has pro-order prime to p, so the restriction map is injective, that is,

$$H^1(G_{F_{\infty,\mathfrak{P}}}/I_\mathfrak{P}, D) = 0,$$

from which it follows that

$$\mathcal{H}_\mathfrak{P} \simeq H^1(F_{\infty,\mathfrak{P}}, D),$$

so $\mathcal{H}_\mathfrak{P}$ is $\Lambda_\mathfrak{p}$-cofree because $H^1(F_{\infty,\mathfrak{P}}, D)$ is. If instead $\mathfrak{p}$ is ramified in F_∞, then $G_{F_{\infty,\mathfrak{P}}}/I_\mathfrak{P}$ is isomorphic to $\widehat{\mathbf{Z}}$, and so $H^1(G_{F_{\infty,\mathfrak{P}}}/I_\mathfrak{P}, D)$ is a quotient of D, hence $\mathscr{O}$-cofree. Thus we have an exact sequence of finitely generated $\Lambda_\mathfrak{p}$-modules

$$0 \to \widehat{\mathcal{H}_\mathfrak{P}} \to H^1(\widehat{F_{\infty,\mathfrak{P}}, D)} \to H^1(\widehat{G_{F_{\infty,\mathfrak{P}}}/I_\mathfrak{P}}, D) \to 0$$

satisfying the hypotheses of Lemma 3.3.11, which therefore implies that $\mathcal{H}_\mathfrak{P}$ is $\Lambda_\mathfrak{p}$-cofree. Thus, in either case, $\mathcal{H}_\mathfrak{P}$ is $\Lambda_\mathfrak{p}$-cofree.

Finally we explain why $\mathcal{H}_\mathfrak{p}$ is cofree over Λ. The choice of a prime $\mathfrak{P}$ above $\mathfrak{p}$ gives rise to an isomorphism of Λ-modules $\mathcal{H}_\mathfrak{p} \simeq \mathcal{H}_\mathfrak{P} \otimes_{\Lambda_\mathfrak{p}} \Lambda$. Since we have proved that $\mathcal{H}_\mathfrak{P}$ is $\Lambda_\mathfrak{p}$-cofree, we conclude that $\mathcal{H}_\mathfrak{p}$ is Λ-cofree. □

Remark 3.3.3. The proof of Proposition 3.3.12 is the only place where we make use of hypothesis p-(i) from §3.1.

Corollary 3.3.13. *Let Σ_0 be a finite set of primes of F not containing any Archimedean primes or any primes of Σ_p. Assume that*

(i). *Σ_0 contains the primes dividing the tame level of f,*

(ii). *hypothesis* (tor) *(introduced at the end of §3.3.1) holds, and*

(iii). *for each prime $v \in \Sigma_0$ that splits in F_∞, $H^0(F_v, A^*) = 0$.*

Then the μ-invariant of $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$ vanishes if and only if $\mathrm{Sel}^{\Sigma_0}(F_\infty, A[\varpi])$ is finite, in which case $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$ is $\mathscr{O}$-divisible with

$$\lambda(\mathrm{Sel}^{\Sigma_0}(F_\infty, A)) = \dim_{\mathbf{F}}(\mathrm{Sel}^{\Sigma_0}(F_\infty, A[\varpi])).$$

Proof. By Proposition 3.3.6, $\mathrm{Sel}^{\Sigma_0}(F_\infty, A[\varpi]) \simeq \mathrm{Sel}^{\Sigma_0}(F_\infty, A)[\varpi]$ as $\mathscr{O}$-modules. Proposition 2.3.3 now implies that the finiteness of the residual Selmer group is equivalent to the vanishing of the μ-invariant of $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$, and that when this happens, $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$ is a cofinitely generated $\mathscr{O}$-module. By Proposition 3.3.12, the $\mathscr{O}$-torsion submodule of the Pontryagin dual of $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$ must then vanish (being finite). Thus $\mathrm{Sel}^{\Sigma_0}(F_\infty, A)$ is $\mathscr{O}$-cofree of corank equal to its λ-invariant, which now visibly coincides with

$$\dim_{\mathbf{F}}(\mathrm{Sel}^{\Sigma_0}(F_\infty, A)[\varpi]) = \dim_{\mathbf{F}}(\mathrm{Sel}^{\Sigma_0}(F_\infty, A[\varpi])).$$

□

3.4 Algebraic λ-invariants

In this section we prove our main result on the behavior of λ-invariants under congruences. We retain the notation from §3.3. Let f_1, f_2 be p-ordinary newforms of weight greater than or equal to 2 (not necessarily the same weight), and tame levels N_1 and N_2, and assume that the Hecke eigenvalues of f_1 and f_2 are contained in E. We assume moreover that the 2-dimensional Galois representations associated to f_1 and f_2 satisfy hypothesis (ram) stated at the beginning of §3.3 i.e., that the residual Galois representations are absolutely irreducible representations of G_F and ramified at each prime $\mathfrak{p} \in \Sigma_p$. Choose $G_{\mathbf{Q}}$-stable lattices T_1, T_2 in the associated Galois representations of f_1, f_2 and let A_1, A_2 be the resulting discrete $\mathscr{O}$-torsion G_F-modules. Let $\mathrm{Sel}(F_\infty, A_1)$ and $\mathrm{Sel}(F_\infty, A_2)$ denote the Selmer groups for f_1 and f_2 over F_∞ as defined in §3.3.1 with the corresponding Iwasawa invariants denoted $\mu(f_1), \lambda(f_1)$ and $\mu(f_2), \lambda(f_2)$.

Let Σ_0 be the set of primes of F dividing $N = N_1N_2$, and let Σ consist of the primes in Σ_0 together with the primes of F dividing ∞ or p. We may write $N\mathscr{O}_F = (N\mathscr{O}_F)^f(N\mathscr{O}_F)^s$, where $(N\mathscr{O}_F)^f$ is divisible only by primes that are finitely decomposed in F_∞ and $(N\mathscr{O}_F)^s$ is divisible only by primes that split in F_∞. For a prime $v \notin \Sigma_p$ (respectively $\mathfrak{p} \in \Sigma_p$), denote by $\mathcal{H}_{v,i}$ (respectively $\mathcal{H}_{\mathfrak{p},i}$) the analogue for A_i of the Λ-module $\mathcal{H}_v$ (respectively $\mathcal{H}_{\mathfrak{p}}$) defined in the beginning of §3.3.2. By Proposition 3.3.3 (i), if $v \mid (N\mathscr{O}_F)^f$, $\mathcal{H}_{v,i}$ is a cotorsion Λ-module; let $\lambda_{v,i}$ be its λ-invariant (which is simply its $\mathscr{O}$-corank since it has μ-invariant zero). Finally, let the λ-invariant of $\mathrm{Sel}^{\Sigma_0}(F_\infty, A_i)$ be denoted by $\lambda(\Sigma_0, f_i)$ for $i = 1, 2$.

Theorem 3.4.1. *For $i = 1, 2$, assume that $\mathcal{H}_{v,i} = 0$ if $v \mid (N\mathscr{O}_F)^s$. Suppose $A_1[\varpi] \simeq A_2[\varpi]$ as $\mathbf{F}[G_F]$-modules. Then $\mathrm{Sel}(F_\infty, A_1)$ is Λ-cotorsion with $\mu(f_1) = 0$ if and only if $\mathrm{Sel}(F_\infty, A_2)$ is Λ-cotorsion with $\mu(f_2) = 0$. In this case, we have*

$$\lambda(f_1) - \lambda(f_2) = \sum_{v \mid (N\mathscr{O}_F)^f} \lambda_{v,2} - \lambda_{v,1}.$$

Proof. First note that the hypothesis on the vanishing of the modules $\mathcal{H}_{v,i}$ for $v \mid (N\mathscr{O}_F)^s$ implies that $H^0(F_v, A_i^*) = 0$ for $i = 1, 2$ and $v \mid (N\mathscr{O}_F)^s$ (because the $\mathscr{O}$-rank of $H^0(F_v, A_i^*)$ is the Λ-corank of $\mathcal{H}_{v,i}$, by Proposition 3.3.4). Suppose that $\mathrm{Sel}(F_\infty, A_1)$ is Λ-cotorsion with $\mu(f_1) = 0$. Then the hypotheses of Corollary 3.3.8 are satisfied for A_1 with our choices of Σ_0 and Σ, and we therefore have an exact sequence of Λ-modules

$$0 \to \mathrm{Sel}(F_\infty, A_1) \to \mathrm{Sel}^{\Sigma_0}(F_\infty, A_1) \to \prod_{v \mid (N\mathscr{O}_F)^f} \mathcal{H}_{v,1} \to 0, \tag{3.11}$$

taking into account the assumption that $\mathcal{H}_{v,1} = 0$ for $v \mid (N\mathscr{O}_F)^s$. The target of the surjective map in (3.11) is Λ-cotorsion with μ-invariant zero by Propositions 3.3.3, and as we have assumed the same for $\mathrm{Sel}(F_\infty, A_1)$, we conclude that $\mathrm{Sel}^{\Sigma_0}(F_\infty, A_1)$ is also Λ-cotorsion with μ-invariant zero. Corollary 3.3.13 now implies that $\mathrm{Sel}^{\Sigma_0}(F_\infty, A_1)$ is $\mathscr{O}$-divisible with $\mathrm{Sel}^{\Sigma_0}(F_\infty, A_1[\varpi])$ finite of $\mathbf{F}$-dimension equal to the λ-invariant $\lambda(\Sigma_0, f_1)$ of $\mathrm{Sel}^{\Sigma_0}(F_\infty, A_1)$.

By the remark following the proof of Proposition 3.3.6, the non-primitive residual Selmer groups $\mathrm{Sel}^{\Sigma_0}(F_\infty, A_1[\varpi])$ and $\mathrm{Sel}^{\Sigma_0}(F_\infty, A_2[\varpi])$ are determined up to Λ-module isomorphism by the $\mathbf{F}[G_F]$-module structures of $A_1[\varpi]$ and $A_2[\varpi]$, respectively. Since we have assumed that these $\mathbf{F}[G_F]$-modules are isomorphic, it therefore follows that we have Λ-module isomorphisms

$$\mathrm{Sel}^{\Sigma_0}(F_\infty, A_1)[\varpi] \simeq \mathrm{Sel}^{\Sigma_0}(F_\infty, A_1[\varpi]) \simeq \mathrm{Sel}^{\Sigma_0}(F_\infty, A_2[\varpi]) \simeq \mathrm{Sel}^{\Sigma_0}(F_\infty, A_2)[\varpi], \tag{3.12}$$

where the first and last isomorphisms come from Proposition 3.3.6. In particular, because $\mathrm{Sel}^{\Sigma_0}(F_\infty, A_1[\varpi])$ is finite, the same is true of $\mathrm{Sel}^{\Sigma_0}(F_\infty, A_2)[\varpi]$. This implies that $\mathrm{Sel}^{\Sigma_0}(F_\infty, A_2)$ is Λ-cotorsion with μ-invariant equal to 0 by Proposition 2.3.3, and since $\mathrm{Sel}(F_\infty, A_2) \subseteq \mathrm{Sel}^{\Sigma_0}(F_\infty, A_2)$, the same is true of $\mathrm{Sel}(F_\infty, A_2)$. The hypotheses of Corollary 3.3.8 are therefore satisfied for A_2, so we have an exact sequence of cotorsion Λ-modules

$$0 \to \mathrm{Sel}(F_\infty, A_2) \to \mathrm{Sel}^{\Sigma_0}(F_\infty, A_2) \to \prod_{v|(N\mathscr{O}_F)^f} \mathcal{H}_{v,2} \to 0. \tag{3.13}$$

The additivity of λ-invariants in short exact sequences of cotorsion Λ-modules applied to the sequences (3.11) and (3.13) gives

$$\lambda(f_1) + \sum_{v|(N\mathscr{O}_F)^f} \lambda_{v,1} = \lambda(\Sigma_0, f_1) \tag{3.14}$$

and

$$\lambda(f_2) + \sum_{v|(N\mathscr{O}_F)^f} \lambda_{v,2} = \lambda(\Sigma_0, f_2). \tag{3.15}$$

The isomorphism 3.12 together with Corollary 3.3.13 gives $\lambda(\Sigma_0, f_1) = \lambda(\Sigma_0, f_2)$. Thus the right-hand sides of (3.14) and (3.15) are equal, so upon equating the left-hand sides and rearranging, we obtain

$$\lambda(f_1) - \lambda(f_2) = \sum_{v|(N\mathscr{O}_F)^f} \lambda_{v,2} - \lambda_{v,1},$$

as desired. □

This theorem is similar to [6, Theorem 4.3.3, 4.3.4] and [20, Theorem 3.1, 3.2], which apply in the cases $F = \mathbf{Q}$ and F_∞ a cyclotomic $\mathbf{Z}_p$-extension, respectively, and is a direct generalization of [8, p. 237] from the case $F = \mathbf{Q}$ (the results in [6] and [20] are stated in a somewhat different form from ours, using the framework of Hida families and Galois deformations, respectively). Note also that the hypothesis on the vanishing of the modules $\mathcal{H}_{v,1}$ and $\mathcal{H}_{v,2}$ for $v \mid (N\mathscr{O}_F)^s$ holds vacuously if all primes of F dividing N are finitely decomposed in F_∞. This is perhaps the case of most interest.

Chapter 4

Preliminaries on p-adic Representations of p-adic Groups

4.1 Locally Convex Spaces Over p-adic Fields

An excellent source for the material mentioned in this section (and lots more) is Schneider's book *Nonarchimedean Functional Analysis* ([16]). The prime p is now again allowed to be arbitrary (for the remainder of the thesis). Let E be a finite extension of $\mathbf{Q}_p$ with ring of integers $\mathscr{O}$. As in classical functional analysis over an Archimedean local field, a topological vector space over E is an E-vector space V equipped with a topology for which addition and scalar multiplication are continuous (a *vector topology* for short). Banach spaces over non-Archimedean fields such as E tend to admit pathologies which do not exist in the Archimedean case. For this reason it is important in non-Archimedean functional analysis (which is the foundation of Schneider-Teitelbaum's theory of continuous representations of p-adic groups on E-vector spaces) to consider general locally convex spaces from the beginning.

Definition 4.1.1. A subset of an E-vector space V is said to be *convex* if it is an additive coset (i.e. an additive translate) of an $\mathscr{O}$-submodule of V. A lattice in V is an $\mathscr{O}$-submodule that spans V over E.

So a convex subset of an E-vector space has the form $v + A$ for A an $\mathscr{O}$-submodule of V, and for any $v' \in v+A$, $v'+A = v+A$ (these are cosets!). In particular, a convex subset of V contains 0 if and only if it is an $\mathscr{O}$-submodule of V. An $\mathscr{O}$-submodule A of V is a lattice if and only if the canonical E-linear map $A \otimes_{\mathscr{O}} E \to V$ is surjective (in which case it is an isomorphism, since the map is always injective). A more concrete way to characterize this surjectivity is the via following condition: A is a lattice in V if and only if for each $v \in V$ there is a non-zero $\alpha \in E$ such that $\alpha v \in A$.

Definition 4.1.2. A topological vector space V over E is said to be *locally convex* if there is a base of open neighborhoods of 0 in V consisting of convex sets (i.e. $\mathscr{O}$-submodules); we then refer to V as a *locally convex space over E*, or as a *locally convex E-vector space.*

Lemma 4.1.1. *If V is a topological vector space over E and $A \subseteq V$ is an open $\mathscr{O}$-submodule, then A is a lattice in V.*

Proof. Fix $v \in V$ and consider the E-linear map $\alpha \mapsto \alpha v : E \to V$. It is continuous, so the inverse image of the open $\mathscr{O}$-submodule A of V is an open neighborhood of 0 in E. As the topology of E is non-discrete, this neighborhood must contain some $\alpha \neq 0$, and then $\alpha v \in A$ by construction. By the remarks following Definition 4.1.1, we conclude that A is a lattice in V. □

According to Lemma 4.1.1, when working with locally convex spaces, it is no loss of generality to restrict attention to open *lattices*, and in particular a locally convex space has a base of opens around 0 consisting of lattices (of course a non-open $\mathscr{O}$-submodule of a locally convex space need not be a lattice). Clearly a vector topology on an E-vector space is uniquely determined by specifying a base of opens around 0, so one can also formulate the definition of local convexity in terms of E-vector spaces admitting a family of lattices satisfying some conditions which ensure that they are a base of opens around 0 for a (necessarily unique) locally convex (vector) topology on V. Analogously to the case of Archimedean functional analysis, locally convex topologies may also be defined via (non-Archimedean) semi-norms, and in fact a topological vector space over E is locally convex if and only if its topology can be defined (in a precise sense) by a family of non-Archimedean semi-norms. The main point is that the "balls centered at 0" defined by a semi-norm are lattices; for details see [16, Propositions 4.3, 4.4].

Definition 4.1.3. If V and W are locally convex spaces over E, then $\mathcal{L}(V, W)$ denotes the space of continuous E-linear maps $V \to W$.

There are many locally convex topologies one can impose on the space $\mathcal{L}(V, W)$ for V, W locally convex (see [16, §6]), but for our work we will not need to consider them.

Definition 4.1.4. Let V be a locally convex space over E. A net (v_i) in V is said to be *Cauchy* if for all open lattices $A \subseteq V$ there exists i_0 so that if $i, j \geq i_0$, $v_i - v_j \in A$. The space V is *complete* if every Cauchy net converges.

When the topology of a Hausdorff locally convex space V can be defined by a countable set of semi-norms (equivalently, when V is first-countable), to check completeness it is enough to check that every Cauchy *sequence* converges ([16, Remark 7.2]).

The simplest (at least in terms of defining the topology) examples of locally convex spaces are those whose topology can be defined by a single (non-Archimedean) norm (a positive-definite semi-norm).

Definition 4.1.5. An E-Banach space is a complete locally convex space V over E whose topology can be defined by a single (non-Archimedean) norm.

It is customary in non-Archimedean functional analysis to *not* regard a norm as part of the data of a Banach space. We only assume that a norm defining the topology exists. Of course, in practice, one must often choose a norm.

Definition 4.1.6. If V and W are E-Banach spaces, then a linear map $f : V \to W$ is said to be bounded if there exist norms $\|\cdot\|_V$ and $\|\cdot\|_W$ on V and W defining their topologies and a constant $C \geq 0$ such that $\|f(v)\|_W \leq C\|v\|_V$ for all $v \in V$.

That the notion of boundedness for a linear map between E-Banach spaces is independent of the choice of norms on the source and target is a consequence of the following proposition.

Proposition 4.1.2. *If V and W are E-Banach spaces, then an E-linear map $f : V \to W$ is continuous if and only if it is bounded.*

Proof. [16, Proposition 3.1]. □

Every E-vector space V can be made into a locally convex space by declaring every lattice to be open, and the resulting topology is clearly the unique finest locally convex topology on V.

Proposition 4.1.3. *If V is an E-vector space, then the finest locally convex topology on V is Hausdorff, and for any locally convex space W over E, $\mathcal{L}(V, W) = \mathrm{Hom}_E(V, W)$. If V is finite-dimensional, the finest locally convex topology on V is the unique Hausdorff locally convex topology on V, and coincides with the topology obtained by choosing an E-linear isomorphism $V \simeq E^{\dim(V)}$ and pulling back the product topology on $E^{\dim(V)}$.*

Proof. [16, Proposition 4.13, §5.C]. □

The basic algebraic constructions in linear algebra applied to locally convex spaces over E yield locally convex spaces. For example, any subspace of a locally convex space is locally convex in the subspace topology, and any quotient of a locally convex space is locally convex in the quotient topology. Initial topologies (in particular product topologies) work as expected, without change from the topological case, but one has to be slightly careful with general final topologies.

Definition 4.1.7. If V is an E-vector space and (V_i) is a family of locally convex spaces equipped with E-linear maps $f_i : V_i \to V$, then the *locally convex final topology on V* defined by the f_i is the finest locally convex topology on

V with respect to which all f_i are continuous. Explicitly, a lattice $A \subseteq V$ is open in the locally convex final topology if and only if $f_i^{-1}(A)$ is open in V_i for all i.

Proposition 4.1.4. *In the notation of Definition 4.1.7, an E-linear map $f : V \to W$ from V to a locally convex space W over E is continuous if and only if $f \circ f_i : V_i \to W$ is continuous for all i.*

Proof. [16, Lemma 5.1 (i)]. □

The locally convex final topology is in general strictly coarser than the final topology on V defined by the maps f_i (where we regard the V_i just as topological spaces, forgetting all E-linear structures). The most important specific instance of the locally convex final topology is the locally convex inductive limit.

Definition 4.1.8. Let (V_i) be an inductive system of locally convex spaces over E and let $V = \varinjlim_i V_i$ be the inductive limit in the category of E-vector spaces. The locally convex final topology on V defined by the canonical E-linear maps $f_i : V_i \to V$ is called the *locally convex inductive limit topology*, and V is called the *locally convex inductive limit* of the V_i.

4.2 Locally Analytic Groups and Their p-adic Representations

Let L be a finite extension of $\mathbf{Q}_p$. The notion of a locally L-analytic manifold as defined in [17, §8 Definition] perfectly mirrors that of a complex manifold. Such an object is a Hausdorff topological space endowed with a maximal locally analytic atlas, where functions on open subsets of L^n with values in an L-Banach space V are locally analytic if they are locally given by convergent power series with coefficients in V. The formal similarity in the definitions belies the fact that while locally analytic and analytic functions on connected open subsets of $\mathbf{C}^n$ are the same, they are definitely not the same on open (not connected!) subsets of L^n. This distinction, leading to the failure of many basic principles in complex analysis to translate to the non-Archimedean setting, served as partial motivation for Tate's introduction of rigid analytic spaces, which rectifies the situation by replacing topological spaces with (a mild) Grothendieck topology.

A locally L-analytic group is of course a locally L-analytic manifold G which is simultaneously an abstract group whose multiplication and inversion maps are locally analytic. The most important source of examples for us are the groups of L-valued points of affine L-group schemes of finite type, as well as certain closed subgroups of such groups. In the early 2000's, Schneider-Teitelbaum initiated the study of continuous representations of locally L-analytic groups on locally convex spaces over E, introducing various classes of representations in the papers [14], [15], and [18]. Here, by a continuous representation of G on a locally convex space V, we mean an action of G on V by E-linear automorphisms for which the action map $G \times V \to V$ is continuous. For us, the relevant classes of representations are the locally algebraic representations, the locally analytic representations, and the Banach space representations.

Let G be a locally L-analytic group. Historically, the first class of continuous representations Schneider and Teitelbaum studied (in [15]) were the Banach space representations.

Definition 4.2.1. An *E-Banach space representation*, or just an *E-Banach representation* of G, is an E-Banach space V equipped with a continuous action $G \times V \to V$ of G by continuous E-linear automorphisms. The representation V is said to be *unitary* if V admits a norm defining its topology which is invariant under the action of G.

The paper [15] makes apparent the sense in which the theory of p-adic representations of locally L-analytic groups may naturally be viewed as a generalization of Iwasawa theory (in the non-commutative direction). Namely, when G is compact, Schneider and Teitelbaum prove that to give a continuous representation of G on an E-Banach space V is the same as endowing V with a separately continuous action $E[[G]] \times V \to V$, where $E[[G]] := \mathscr{O}[[G]] \otimes_{\mathscr{O}} E$ is the Iwasawa algebra of G with p inverted (see the beginning of §1.1). Schneider and Teitelbaum go on to introduce a finiteness condition on E-Banach representations V of G which they call *admissibility*, proving the following theorem ([15, Theorem 3.5]).

Theorem 4.2.1. *The functor $V \mapsto V'$ taking an E-Banach representation of G to its dual is an equivalence of categories between the category of admissible E-Banach representations of G and the category finitely generated $E[[G]]$-modules.*

We will not need the notion of admissibility, so we omit the definition (although, in light of the theorem, one could take it to mean that the dual space, which is naturally an $E[[G]]$-module, is finitely generated).

The next class of representations of G introduced by Schneider-Teitelbaum (in [18]) is the class of locally analytic representations. To give the definition, we should assume that L is endowed with a $\mathbf{Q}_p$-algebra embedding into E (this allows one to make sense of locally L-analytic functions on locally L-analytic manifolds valued in locally convex spaces over E). First we need a functional-analytic definition, originally given in [18, §1] (see [16, §16 Definition] for the definition of a compact map between locally convex E-vector spaces).

Definition 4.2.2. A locally convex space V over E is said to be of *compact type* if V is the locally convex inductive limit of a sequence (V_n) of E-Banach spaces with injective, compact transition maps.

Definition 4.2.3. A continuous representation of G on a locally convex space V over E is said to be *locally L-analytic* if V is of compact type and for each $v \in V$, the orbit map $g \mapsto gv : G \to V$ is locally L-analytic.

When no confusion will arise (e.g. if L is fixed in the discussion), one speaks just of locally analytic representations of G, instead of locally L-analytic representations.

As for E-Banach representations of G, Schneider and Teitelbaum equip the functor $V \mapsto V'$ (continuous linear dual) with the structure of a functor from locally analytic representations of G over E to the category of modules over a certain algebra of locally analytic distributions on G (a generalization of the generic fiber of an Iwasawa algebra), and prove in [18] that, in the compact case, one obtains a fully faithful embedding of categories from so-called *strongly admissible* locally analytic representations of G to finitely generated modules over this distribution algebra.

The final class of representations introduced by Schneider-Teitelbaum that we will use is the class of locally algebraic representations. For this notion to make sense, we must assume that G is the group of L-points of a connected reductive group $\mathbf{G}$ over L, and continue to assume that L is endowed with an embedding into E. We use Definition 4.2.6 of [3].

Definition 4.2.4. Let W be an irreducible finite-dimensional algebraic representation of $\mathbf{G} \times_L E$ over E. A vector v in an E-linear representation V of G is said to be *locally W-algebraic* if there is an open subgroup H of G, an integer $n \geq 1$, and an E-linear, H-equivariant homomorphism $W^n \to V$ with image containing v.

Note that, if W is the trivial representation of $\mathbf{G} \times_L E$ over E, then a locally W-algebraic vector is nothing more than a smooth vector (in the usual sense, meaning a vector fixed by a compact open subgroup of G).

Definition 4.2.5. A vector v in an E-linear representation V of G is said to be *locally algebraic* if it is locally W-algebraic for some irreducible finite-dimensional algebraic representation W of $\mathbf{G} \times_L E$ over E. The representation V is said to be *locally algebraic* if every vector in V is locally algebraic.

In light of the remark following Definition 4.2.4, a locally W-algebraic representation of G over E for W the trivial representation is the same thing as a smooth representation of G over E. Note that, unlike Banach and locally analytic representations, the definition of a locally algebraic representation makes no reference to topology. However, endowing a locally algebraic representation V of G with its finest locally convex topology, V becomes a locally analytic analytic representation of G over E (see the discussion following Definition 4.2.1 in [3]). In particular, the category of locally analytic representations of G on E-vector spaces of compact type contains as a full subcategory the category of smooth representations of G over E. It turns out that if $\mathbf{G}$ is split over E, then an irreducible locally algebraic representation V of G has the form $U \otimes_E W$, where U is an irreducible smooth representation of G over E and W is an irreducible finite-dimensional algebraic representation of $\mathbf{G} \times_L E$ over E. Conversely, any such tensor product is an irreducible locally algebraic representation of G. This is the content of [3, Proposition 4.2.8]. Therefore one can reasonably think of locally algebraic representations of G as twists of smooth representations by algebraic representations. Using that locally algebraic representations of G are naturally locally analytic, one can define admissibility of such a representation by using the definition of admissibility for locally analytic representations. Then the admissible locally algebraic representations are essentially tensor products of admissible smooth representations (in the classical sense) with algebraic representations, at least when $\mathbf{G} \times_L E$ is split and the representation is irreducible (see [3, Proposition 6.3.10]).

We would be remiss not to at least mention some of the relationships between the kinds of representations of G introduced above. We've already remarked that smooth representations are locally algebraic, and that locally algebraic representations may naturally be viewed as locally analytic representations. In the other direction, if V is an admissible locally analytic representation of G, then by [3, Proposition 6.3.6], the subspace V_{lalg} of locally algebraic vectors in V is a closed, G-stable subspace (and hence is an admissible locally algebraic representation of G). For admissible Banach space representations, there is the functor "pass to the locally analytic vectors," where a vector is locally analytic if the associated orbit map is locally analytic. This functor, described in great detail and generality in [3, §3.5], lands in the category of admissible locally analytic representations ([3, Proposition 6.2.4]), and when $L = \mathbf{Q}_p$, is exact by results in [19]. Moreover, the space of locally analytic vectors (and indeed any locally analytic representation of G) has a derived action of the Lie algebra $\mathfrak{g}$ of G, and on passing to the subspace of vectors annihilated by $\mathfrak{g}$, one obtains exactly the space of smooth vectors ([3, Corollary 4.1.7]).

4.3 The Topology on Locally Analytic Functions

As in the previous section, L is a finite extension of $\mathbf{Q}_p$, and again we assume it is equipped with a $\mathbf{Q}_p$-algebra embedding $L \hookrightarrow E$. Let M be a locally L-analytic manifold which is *strictly paracompact*; this means that every open covering of M can be refined by a covering consisting of pairwise disjoint open sets (all locally L-analytic groups are strictly paracompact). In particular M is paracompact, and in fact the converse is true by a theorem of Schneider ([17, Proposition 8.7]). As for any locally L-analytic manifold, we have the E-vector space $\mathscr{C}^{\mathrm{la}}(M, E)$ of locally analytic E-valued functions on M. This space consists of all functions $f : M \to E$ which are locally given (in coordinates) by convergent power series with coefficients in E. Following [17, §10], we want to describe the locally convex topology on $\mathscr{C}^{\mathrm{la}}(M, E)$. By a chart for M, we mean a member (U, φ, n) of its locally analytic atlas, meaning that U is an open subset of M, and φ is a homeomorphism of U onto an open subset of L^n.

Definition 4.3.1. An *analytic partition* of M is an open covering $M = \bigcup_i U_i$ into pairwise disjoint open subsets which are domains of charts for M whose images are closed affinoid polydisks (meaning that the polydisk admits a radius which is in the divisible subgroup of $\mathbf{R}^{\times}_{>0}$ generated by the group of absolute values of $L^{\times}$).

Because M is strictly paracompact, any open covering of M by chart domains can be refined by an analytic partition.

The next definition is taken verbatim from [17, §10].

Definition 4.3.2. A *index* for M is a family of pairs $\mathscr{I} = \{(c_i, \epsilon_i)\}_{i \in I}$, where each $c_i = (U_i, \varphi_i, n_i)$ is a chart for M whose image is an affinoid polydisk in L^{n_i} of radius ϵ_i, and the U_i form an analytic partition of M.

To each index $\mathscr{I} = \{(c_i, \epsilon_i)\}_{i \in I}$, we will associate a subspace $\mathscr{F}_{\mathscr{I}} \subseteq \mathscr{C}^{\mathrm{la}}(M, E)$. First, for $i \in I$, we denote by $\mathscr{O}(\mathbf{B}^{n_i}_{\leq \epsilon_i})$ the ring of rigid analytic functions on the affinoid polydisk $\mathbf{B}^{n_i}_{\leq n_i}$ over E centered at 0 of radius ϵ_i; this is the ring of formal power series in n_i variables over E which converge on the closed polydisk centered at 0 of radius ϵ_i in $\overline{E}^{n_i}$, where $\overline{E}$ is a choice of algebraic closure of E. By choosing a point of $\varphi_i(U_i)$ (a center of the polydisk $\varphi_i(U_i)$), we may restrict rigid analytic functions on the polydisk centered at 0 to $\varphi_i(U_i)$, and following this by pre-composition with φ_i yields an injection $\mathscr{O}(\mathbf{B}^{n_i}_{\leq \epsilon_i}) \to \mathscr{C}^{\mathrm{la}}(U_i, E)$. Extension by zero gives us an injection $\mathscr{C}^{\mathrm{la}}(U_i, E) \hookrightarrow \mathscr{C}^{\mathrm{la}}(M, E)$, and by composing with the previous map, we obtain an injection $\mathscr{O}(\mathbf{B}^{n_i}_{\leq \epsilon_i}) \to \mathscr{C}^{\mathrm{la}}(M, E)$, the image of which is denoted $\mathscr{F}_{(c_i, \epsilon_i)}$. By [17, Corollary 5.5], the image of this map is independent of the choice of center point in $\varphi_i(U_i)$, as is the norm on $\mathscr{F}_{(c_i,\epsilon_i)}$ obtained by pushing forward the canonical norm on the E-Banach algebra $\mathscr{O}(\mathbf{B}^{n_i}_{\leq \epsilon_i})$. Thus $\mathscr{F}_{(c_i,\epsilon_i)}$ has a natural structure of E-Banach space. Since $\prod_{i\in I} \mathscr{C}^{\mathrm{la}}(U_i, E) = \mathscr{C}^{\mathrm{la}}(M, E)$ for any analytic partition $M = \bigcup_{i\in I} U_i$ of M, the product $\mathscr{F}_{\mathscr{I}} := \prod_{i\in I} \mathscr{F}_{(c_i,\epsilon_i)}$ is canonically a subspace of $\mathscr{C}^{\mathrm{la}}(M, E)$. We give the subspace $\mathscr{F}_{\mathscr{I}}$ the locally convex topology coming from its structure as a product of Banach spaces.

In [17, §10], a refinement relation is introduced on the set of indices for M. It is shown that the set of indices is directed by the relation of reverse refinement, that $\mathscr{F}_{\mathscr{J}} \subseteq \mathscr{F}_{\mathscr{I}}$ if $\mathscr{I}$ refines $\mathscr{J}$, and that this inclusion is continuous for the topologies on the source and target. Moreover, $\mathscr{C}^{\mathrm{la}}(M, E)$ is the union of the subspaces $\mathscr{F}_{\mathscr{I}}$ (taken over all indices $\mathscr{I}$) essentially by the definition of a locally analytic function.

Definition 4.3.3. The topology on $\mathscr{C}^{\mathrm{la}}(M, E)$ is the locally convex inductive limit topology (Definition 4.1.8) obtained from the E-linear identification $\mathscr{C}^{\mathrm{la}}(M,E) = \varinjlim_{\mathscr{I}} \mathscr{F}_{\mathscr{I}}$, the inductive limit being taken over all indices $\mathscr{I}$ for M.

This topology is of most interest for us in the case that M is compact. In this case, the underlying index set I of an index $\mathscr{I}$ must be finite, so each $\mathscr{F}_{\mathscr{I}}$ is in fact a Banach space, and $\mathscr{C}^{\mathrm{la}}(M,E)$ is an inductive limit of Banach spaces. As a result, the following proposition is plausible.

Proposition 4.3.1. *If M is compact, then $\mathscr{C}^{\mathrm{la}}(M,E)$ is of compact type (Definition 4.2.2).*

Proof. [3, Proposition 2.1.28]. □

Chapter 5

Parabolically Induced and Unitary Banach Representations of p-adic GL_2

5.1 Introduction and Notation

In this chapter, we generalize a result of Emerton [4, Proposition 2.5] on continuous homomorphisms from certain locally analytic parabolically induced representations of $\mathrm{GL}_2(L)$ into unitary Banach space representations. Here L is a finite extension of $\mathbf{Q}_p$, and all representations are over some sufficiently large p-adic field E. For certain locally algebraic characters χ of T, the subgroup of diagonal matrices in $\mathrm{GL}_2(L)$, we define the locally algebraic and locally $\mathbf{Q}_p$-analytic parabolic inductions $I(\chi)$ and $I^{\mathrm{la}}(\chi)$ of χ, viewed as a character of $\overline{B}$, the subgroup of lower triangular matrices in $\mathrm{GL}_2(L)$. These are admissible locally $\mathbf{Q}_p$-analytic representations of $\mathrm{GL}_2(L)$, and there is a canonical closed embedding $I(\chi) \hookrightarrow I^{\mathrm{la}}(\chi)$, allowing us to view $I(\chi)$ as a locally algebraic subrepresentation of $I^{\mathrm{la}}(\chi)$. Let U be a unitary Banach space representation of $\mathrm{GL}_2(L)$ and assume that $\chi|_{Z(\mathrm{GL}_2(L))}$ is unitary (see Definition 4.2.1 for the definition of unitarity). Our main result (Theorem 5.3.1) states that, under a "non-critical slope" hypothesis on χ, any continuous $\mathrm{GL}_2(L)$-equivariant linear map $I(\chi) \to U$ extends uniquely to a continuous $\mathrm{GL}_2(L)$-equivariant linear map $I^{\mathrm{la}}(\chi) \to U$. At least when $I(\chi)$ is irreducible, this result is equivalent to the assertion that $I(\chi)$ and $I^{\mathrm{la}}(\chi)$ have the same universal unitary completion (in the sense of [4, Definition 1.1]). Emerton proved this result for $L = \mathbf{Q}_p$. Breuil has proved a similar result [2, Theorem 7.1] covering injective linear maps out of locally J-analytic parabolically induced representations of $\mathrm{GL}_2(L)$, for subsets $J \subseteq \mathrm{Hom}_{\mathbf{Q}_p-\mathrm{Alg}}(L, E)$ ($J = \emptyset$ corresponds to $I(\chi)$, while $J = \mathrm{Hom}_{\mathbf{Q}_p-\mathrm{Alg}}(E, L)$ corresponds to $I^{\mathrm{la}}(\chi)$). We closely follow Emerton's method of proof, which is different from Breuil's, although we do make use of a generalization of a classical result of Amice-Vélu and Vishik, Lemma 5.2.4 below, proved by Breuil in [2, Lemma 6.1].

We introduce our notation in detail below, mostly retaining that of [4]. In §5.2 we describe the locally convex space $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ of E-valued locally $\mathbf{Q}_p$-analytic functions on $\mathscr{O}_L$ using the set of embeddings $\mathrm{Hom}_{\mathbf{Q}_p-\mathrm{Alg}}(L, E)$ (which is assumed to have $[L : \mathbf{Q}_p]$ elements). Although a description of this space along these lines has been used (somewhat implicitly) in other places (e.g. [2]), as far as we know, there is no published proof that it coincides (set-theoretically and topologically) with the standard description of this space given in §4.3, so we provide a detailed proof of the equivalence. We define the parabolic inductions $I(\chi)$ and $I^{\mathrm{la}}(\chi)$ in §5.3, and state the main result, which is proved (following Emerton's proof in [4, §3]) in §5.4. The argument is primarily representation-theoretic and functional-analytic, and most of the work is dedicated to reducing the statement of Theorem 5.3.1 to Lemma 5.2.4 by reinterpreting the condition of temperedness (Definition 5.2.1) in terms of equivariance with respect to the action of a submonoid B^+ of B, the group of upper triangular matrices in $\mathrm{GL}_2(L)$ (this reinterpretation is provided by Lemma 5.4.4).

We deviate slightly from previous notation due to our use of two finite extensions of $\mathbf{Q}_p$ and their respective rings of integers. Namely, let L and E be finite extensions of $\mathbf{Q}_p$ with respective rings of integers $\mathscr{O}_L$ an $\mathscr{O}_E$, and denote by ϖ_L a choice of uniformizer for $\mathscr{O}_L$. Set $r = [L : \mathbf{Q}_p]$, and assume that $\mathrm{Hom}_{\mathbf{Q}_p-\mathrm{Alg}}(E, L)$ has r distinct elements that we order for convenience: $\sigma_1, \ldots, \sigma_r$ (nothing we do will depend on the choice of ordering, and it is only made to ease notation). The field E will serve as the coefficient field of our representations.

We normalize the discrete valuation of E, $\mathrm{ord} = \mathrm{ord}_E$, by $\mathrm{ord}_E(p) = e(L/\mathbf{Q}_p)$ (the ramification index of L over $\mathbf{Q}_p$) and use the absolute value $|\cdot| = |\cdot|_E$ defined by $|\alpha| = q^{-\mathrm{ord}(\alpha)}$, where q is the cardinality of the residue field of L. If we use the same normalizations for the discrete valuation and absolute value on L, then L is endowed with its canonical absolute value, i.e., the one giving ϖ_L absolute value q^{-1}, and each σ_i is an isometry. We will therefore denote the discrete valuation on either E or L simply by ord, and the absolute value by $|\cdot|$.

We denote by G the group $\mathrm{GL}_2(L)$, viewed as the group of $\mathbf{Q}_p$-points of the reductive group $\mathbf{G} = \mathrm{Res}_{L/\mathbf{Q}_p}(\mathrm{GL}_{2/L})$. Thus we regard $\mathrm{GL}_2(L)$ as a locally $\mathbf{Q}_p$-analytic group, and by "locally analytic," we will always mean "locally $\mathbf{Q}_p$-analytic." We apply the same convention to all other groups that we consider. We let B and $\overline{B}$ denote the groups of $\mathbf{Q}_p$-points of the upper triangular and lower triangular, respectively, Borel subgroups of $\mathbf{G}$, N and $\overline{N}$ the groups of $\mathbf{Q}_p$-points of their unipotent radicals, and T the group of $\mathbf{Q}_p$-points of the diagonal torus in $\mathbf{G}$. Setting $G_0 = \mathrm{GL}_2(\mathscr{O}_L)$, we define, for each integer $s \geq 0$,

$$G_0(s) = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in G_0 : c \equiv 0 \pmod{\varpi_L^s \mathscr{O}_L} \right\}.$$

These are compact open subgroups of G admitting an Iwahori decomposition with respect to B and $\overline{B}$, meaning that if $T_0 = T \cap G_0$, $N_0 = N \cap G_0$, and $\overline{N}(s) = \overline{N} \cap G_0(s)$, then the natural multiplication map $N_0 T_0 \overline{N}(s) \to G_0(s)$ is a bijection for $s \geq 1$ (note that $N_0 = N \cap G_0(s)$ and $T_0 = T \cap G_0(s)$ for all $s \geq 0$). If $T^+ = \{t \in T : tN_0t^{-1} \subseteq N_0\}$, then T^+ is a submonoid of T containing T_0, and for each $t \in T^+$ and each integer $s \geq 1$, $t^{-1}\overline{N}(s)t \subseteq \overline{N}(s)$. Explicitly, T^+ consists of all matrices $\left(\begin{smallmatrix} a & 0 \\ 0 & d \end{smallmatrix}\right) \in T$ with $ad^{-1} \in \mathscr{O}_L$. One can verify then that for $s \geq 1$, $G^+(s) = N_0 T^+ \overline{N}(s)$ is a submonoid of G containing $G_0(s)$. We write B^+ for $G^+(1) \cap B = N_0 T^+$ (the equality holds because $\overline{N} \cap B = \{1\}$, and shows we could replace 1 in the definition of B^+ by any integer $s \geq 1$ without changing the result); this is a submonoid of B which (by inspection) generates B as a group. Each of these subgroups (respectively submonoids) of G will be regarded as a subgroup (respectively submonoid) of the group of E-points of $\mathbf{G} \times_{\mathbf{Q}_p} E = \prod_{i=1}^r \mathrm{GL}_{2/L} \times_{\sigma_i} E$ via the continuous inclusion

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \mapsto \left(\begin{pmatrix} \sigma_i(a) & \sigma_i(b) \\ \sigma_i(c) & \sigma_i(d) \end{pmatrix} \right)_{1 \leq i \leq r}.$$

We will generally identify N_0 with $\mathscr{O}_L$ via the locally analytic isomorphism $\left(\begin{smallmatrix} 1 & z \\ 0 & 1 \end{smallmatrix}\right) \mapsto z$.

Elements of $\mathbf{Z}_{\geq 0}^r$ will be denoted by underlined Latin letters, e.g. $\underline{m} = (m_1, \ldots, m_r)$, and we set $|\underline{m}| = \sum_{i=1}^r m_i$. For each integer $n \geq 0$, we write $\mathscr{A}_n$ for the affinoid E-algebra of formal power series

$$F = F(X_1, \ldots, X_r) = \sum_{\underline{m} \in \mathbf{Z}_{\geq 0}^r} a_{\underline{m}} X_1^{m_1} \cdots X_r^{m_r} \in E[[X_1, \ldots, X_r]]$$

satisfying $\lim_{|\underline{m}| \to \infty} |a_{\underline{m}}| q^{-n|\underline{m}|} = 0$. This is an E-Banach algebra with the multiplicative Gauss norm $\| \cdot \|_n$ given by $\|F\|_n = \max_{\underline{m}} |a_{\underline{m}}| q^{-n|\underline{m}|}$. When $n = 0$, we will write $\mathscr{A}$ (respectively $\| \cdot \|_{\mathscr{A}}$) instead of $\mathscr{A}_0$ (respectively $\| \cdot \|_0$). If $\underline{k} \in \mathbf{Z}_{\geq 0}^r$, then we will denote by $\mathscr{A}^{\underline{k}}$ the finite-dimensional (hence closed) subspace of $\mathscr{A}$ consisting of all polynomials in $E[X_1, \ldots, X_r]$ whose degree in X_i is at most k_i for $1 \leq i \leq r$ (note that in fact $\mathscr{A}^{\underline{k}}$ is a closed subspace of $\mathscr{A}_n$ for all $n \geq 0$). We will refer to $\mathscr{A}^{\underline{k}}$ as the space of polynomials in $\mathscr{A}$ "of degree at most $\underline{k}$."

If H is a locally $\mathbf{Q}_p$-analytic group, $\mathscr{C}^{\mathrm{la}}(H, E)$ denotes the locally convex space of locally analytic E-valued functions on H (see [17, §10] for a detailed description of the locally convex topology on this space, and §2 below for an alternative description in the case $H = \mathscr{O}_L$) and $\mathscr{C}^{\mathrm{sm}}(H, E)$ denotes the space of smooth (i.e. locally constant) E-valued functions on H. For an open subset U of H, $\mathbf{1}_U$ denotes the characteristic function of U (so $\mathbf{1}_U \in \mathscr{C}^{\mathrm{sm}}(H, E) \subseteq \mathscr{C}^{\mathrm{la}}(H, E)$). The isomorphism $N_0 \simeq \mathscr{O}_L$ yields a topological isomorphism $\mathscr{C}^{\mathrm{la}}(N_0, E) \simeq \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$.

If V and W are locally convex spaces over E (see Definition 4.1.2), $\mathcal{L}(V, W)$ denotes the space of continuous E-linear maps from V to W as in

Definition 4.1.3. If moreover each of V, W is endowed with an action of a topological monoid H by E-linear (topological) automorphisms, then $\mathcal{L}_H(V, W)$ denotes the subspace of $\mathcal{L}(V, W)$ consisting of continuous H-equivariant E-linear maps. We will not need to consider any locally convex topologies on the space $\mathcal{L}(V, W)$, so an isomorphism between spaces of continuous linear maps is simply an isomorphism of E-vector spaces. An E-Banach space representation U of H is *unitary* if the topology of U can be defined by a norm that is invariant under H (unitarity was defined in Definition 4.2.1 for locally analytic groups, but the same definition applies for any monoid). Thus an E-valued character of H is unitary if and only if it takes values in $\mathscr{O}_E^\times$.

5.2 The Locally Convex Space $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ and Tempered Linear Maps

In accordance with our convention regarding locally analytic structures mentioned in §5.1, we regard the locally L-analytic group $\mathscr{O}_L$ as a locally $\mathbf{Q}_p$-analytic group by restriction of scalars. Explicitly, if we choose a $\mathbf{Z}_p$-basis for $\mathscr{O}_L$, then the induced $\mathbf{Z}_p$-linear isomorphism $\mathscr{O}_L \simeq \mathbf{Z}_p^r$ is a global chart for the locally $\mathbf{Q}_p$-analytic structure on $\mathscr{O}_L$. By definition, a function $f : \mathscr{O}_L \to E$ is locally analytic if, upon choosing an isomorphism $\mathscr{O}_L \simeq \mathbf{Z}_p^r$, the resulting function $\mathbf{Z}_p^r \to E$ admits a power expansion (in r variables, with coefficients in E) in a sufficiently small ball around each point of $\mathbf{Z}_p^r$. Thus, given a choice of coordinates $\mathscr{O}_L \simeq \mathbf{Z}_p^r$, elements of the affinoid algebras $\mathscr{A}_n$ give rise to locally analytic functions on $\mathscr{O}_L$. In fact it is somewhat more intrinsic (but equivalent) to consider functions which are locally given by convergent power series in the embeddings $\sigma_i : E \hookrightarrow L$, as we now explain.

If $w \in \mathscr{O}_L$ and $n \geq 0$, then because each $\sigma_i : E \hookrightarrow L$ is an isometry for our choice of absolute values on E and L, an element $F \in \mathscr{A}_n$ gives rise to a continuous function $F_{n,w} : w + \varpi_L^n \mathscr{O}_L \to E$ defined by $F_{n,w}(z) = F(\sigma_1(z - w), \ldots, \sigma_r(z - w))$ (we will abuse notation by sometimes writing the right-hand side of this definition as $F(z - w)$, and will use the same notation to denote the function on $\mathscr{O}_L$ obtained by extending $F_{n,w}$ by zero). It turns out that the E-valued locally analytic functions on $\mathscr{O}_L$ are precisely the functions $\mathscr{O}_L \to E$ which locally arise from this construction in the sense of the following proposition.

Proposition 5.2.1. *A function $f : \mathscr{O}_L \to E$ is locally analytic if and only if for each $w \in \mathscr{O}_L$ there exists an integer $n \geq 0$ and a series $F \in \mathscr{A}_n$ such that $f|_{w+\varpi_L^n \mathscr{O}_L} = F_{n,w}$.*

Proof. Let $\{z_1, \ldots, z_r\}$ be a $\mathbf{Z}_p$-basis for $\mathscr{O}_L$ and let $\pi : \mathscr{O}_L \simeq \mathbf{Z}_p^r$ be the $\mathbf{Z}_p$-linear isomorphism defined by this choice of basis. Then π is an isomorphism of locally analytic groups. For $1 \leq i \leq r$ let $\pi_i : \mathscr{O}_L \to \mathbf{Z}_p$ be the $\mathbf{Z}_p$-linear map given by $\pi_i(z_j) = \delta_{ij}$, so that $\pi(z) = (\pi_1(z), \ldots, \pi_r(z))$ for each $z \in \mathscr{O}_L$. The π_i form an E-basis for the space $M = \mathrm{Hom}_{\mathbf{Z}_p-\mathrm{Mod}}(\mathscr{O}_L, E)$, and we have

$$\sigma_j = \sum_{i=1}^{r} \sigma_j(z_i)\pi_i.$$

for $1 \le j \le r$. As the σ_i also form an E-basis for M, we can write

$$\pi_j = \sum_{i=1}^{r} \beta_{ij}\sigma_i$$

for some $\beta_{ij} \in E$, $1 \le i, j \le r$. The $n \times n$ matrices (β_{ij}) and $(\sigma_j(z_i))$ are then mutually inverse in $\mathrm{GL}_r(E)$, and $(\sigma_j(z_i))$ has coefficients in $\mathscr{O}_E$ (though it need not have unit determinant, i.e., (β_{ij}) might not have integral coefficients).

Define polynomials $g_j = \sum_{i=1}^{r} \sigma_j(z_i)X_i$ and $h_j = \sum_{i=1}^{r} \beta_{ij}X_i$ in $E[X_1, \ldots, X_r]$ for $1 \le j \le r$, noting that

$$g_j(\pi_1(z), \ldots, \pi_r(z)) = \sigma_j(z) \tag{5.1}$$

and

$$h_j(\sigma_1(z), \ldots, \sigma_r(z)) = \pi_j(z) \tag{5.2}$$

for all $z \in \mathscr{O}_L$ and $1 \le j \le r$. Suppose $f : \mathscr{O}_L \to E$ is locally-analytic and fix $w \in \mathscr{O}_L$. We may then choose an integer $k \ge 0$ and a power series $F_0 \in \mathscr{A}_k$ such that

$$(f \circ \pi^{-1})(x_1, \ldots, x_r) = F_0(x_1 - \pi_1(w), \ldots, x_r - \pi_r(w))$$

for each $(x_1, \ldots, x_r) \in \mathbf{Z}_p^r$ with $\max_i |x_i - \pi_i(w)| \le q^{-k}$. Now choose an integer $n \ge 0$ large enough to ensure that $\|h_j\|_n = q^{-n} \max_i |\beta_{ij}|$ is less than or equal to q^{-k} for $1 \le j \le r$ (the coefficients β_{ij} have valuation depending on the ramification of L over E). There is then a unique continuous E-algebra homomorphism $\mathscr{A}_k \to \mathscr{A}_n$ satisfying $X_j \mapsto h_j$ for $1 \le j \le r$, and this E-algebra homomorphism is compatible with evaluation of series in $\mathscr{A}_n$ on points of the closed ball around 0 in E^r of radius q^{-n}. (This is the universal property of Tate algebras, and details can be found in [17, Proposition 5.4]. Implicit in the statement about evaluation on points is that each h_j maps the closed ball of radius q^{-n} around 0 in E^r into the closed ball of radius q^{-k} around 0 in E.) Let F be the image of F_0 under this homomorphism (so we think of F as $F_0(h_1, \ldots, h_r)$). Given $z \in w + \varpi_L^n \mathscr{O}_L$, we have $|\sigma_i(z-w)| \leq q^{-n}$ for $1 \leq i \leq r$. Using the aforementioned compatibility of $F_0 \mapsto F$ with evaluation on points, we find that

$$\begin{aligned} F(\sigma_1(z-w), \ldots, \sigma_r(z-w)) &= F_0(h_i(\sigma_1(z-w), \ldots, \sigma_r(z-w))) \\ &= F_0(\pi_1(z-w), \ldots, \pi_r(z-w)) \\ &= F_0(\pi_1(z) - \pi_1(w), \ldots, \pi_r(z) - \pi_r(w)) \\ &= (f \circ \pi^{-1})(\pi_1(z), \ldots, \pi_r(z)) = f(z), \end{aligned}$$

where, in going from the first to the second line, we have used Equation (5.2), and in the final equality, we have used the parenthetical remark above explaining why $|\pi_i(z) - \pi_i(w)| \leq q^{-k}$ for $1 \leq i \leq r$. Thus f has the desired local form.

The converse is similar. Because the g_j have integral coefficients, $\|g_j\|_n \leq q^{-n}$ for $1 \leq j \leq r$, and we have a unique continuous E-algebra homomorphism $\mathscr{A}_n \to \mathscr{A}_n$ satisfying $X_j \mapsto g_j$ for $1 \leq j \leq r$. This map is compatible with evaluation on points as before, and we may use it (together with Equation (5.1)) to prove that a function satisfying the condition in the statement of the proposition is locally analytic by converting a local power series expansion in the σ_i to a local power series expansion in the π_i. $\square$

Remark 5.2.1. It is actually not necessary to check the condition in Proposition 5.2.1 at every $w \in \mathscr{O}_L$. In fact, the condition in the proposition is equivalent to the condition that there exists an integer $n \geq 0$ such that for each w in a set of coset representatives for $\varpi_L^n \mathscr{O}_L$ in $\mathscr{O}_L$, there exists $F \in \mathscr{A}_n$ such that $f|_{w+\varpi_L^n \mathscr{O}_L} = F_{n,w}$. This follows from [17, Corollary 5.5], which shows that, for any $w' \in w + \varpi_L^n \mathscr{O}_L$, the function $F_{n,w}$ for $F \in \mathscr{A}_n$ coincides with $F'_{n,w'}$ for some $F' \in \mathscr{A}_n$ (and one even necessarily has $\|F\|_n = \|F'\|_n$.)

We now describe the locally convex topology on $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ in terms of the description of this vector space provided by Proposition 5.2.1. For each $n \geq 0$, let T_n be a set of coset representatives in $\mathscr{O}_L$ for $\varpi_L^n \mathscr{O}_L$, and let $\iota_n : \prod_{w \in T_n} \mathscr{A}_n \to \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ be given by sending a tuple $(F_w)_{w \in T_n}$ of q^{-n}-convergent power series to the function $\mathscr{O}_L \to E$ that is given on the ball $w + \varpi_L^n \mathscr{O}_L$ by $z \mapsto F_w(\sigma_1(z-w), \ldots, \sigma_r(z-w))$ (that this function is in fact locally analytic follows from Proposition 5.2.1 coupled with Remark 5.2.1). A Zariski density argument shows that ι_n is injective, and both the image $\mathscr{F}_n(\mathscr{O}_L, E)$ of ι_n and the norm induced on $\mathscr{F}_n(\mathscr{O}_L, E)$ from the maximum of the Gauss norms on each factor of the source of ι_n are independent of the choice of T_n (again by Remark 5.2.1). Thus $\mathscr{F}_n(\mathscr{O}_L, E)$ is canonically an E-Banach space. We have $\mathscr{F}_n(\mathscr{O}_L, E) \subseteq \mathscr{F}_{n+1}(\mathscr{O}_L, E)$ for each $n \geq 0$, a continuous inclusion, and Remark 5.2.1 shows that the natural E-linear injection $\varinjlim_n \mathscr{F}_n(\mathscr{O}_L, E) \to \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ is an isomorphism of E-vector spaces. We may therefore endow $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ with the locally convex inductive limit topology coming from this isomorphism and the Banach space structure on each $\mathscr{F}_n(\mathscr{O}_L, E)$. Thus if U is a locally convex space over E, a linear map $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E) \to U$ is continuous if and only if the restriction of the map to $\mathscr{F}_n(\mathscr{O}_L, E)$ is continuous for every $n \geq 0$.

Proposition 5.2.2. *The locally convex topology just defined on $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ coincides with the locally convex topology of Definition 4.3.3.*

Proof. Let $\pi : \mathscr{O}_L \simeq \mathbf{Z}_p^r$ be as in the proof of Proposition 5.2.1 and let $e = e(L/\mathbf{Q}_p)$ be the ramification index of L over $\mathbf{Q}_p$. Recall that, with our normalizations, $|p| = q^{-e}$. Thus, if $n \geq 0$, π induces a locally analytic isomorphism $\varpi_L^{ne}\mathscr{O}_L = p^n\mathscr{O}_L \simeq p^n\mathbf{Z}_p^r$ between the balls around 0 of radius $|\varpi_L^{ne}| = |p^n| = q^{-ne}$ in $\mathscr{O}_L$ and $\mathbf{Z}_p^r$ (where we use the norm $\|x\| = \max_i |x_i|$ on $\mathbf{Z}_p^r$). In particular, if T_{ne} is a set of coset representatives for $\varpi_L^{ne}\mathscr{O}_L$ in $\mathscr{O}_L$, then $\pi(T_{ne})$ is a set of coset representatives for $p^n\mathbf{Z}_p^r$ in $\mathbf{Z}_p^r$. We have a diagram

$$\begin{array}{ccc} \prod_{w\in T_{ne}} \mathscr{A}_{ne} & \xrightarrow{\iota_{ne}} & \mathscr{F}_{ne}(\mathscr{O}_L, E) \\ \downarrow & & \downarrow \subseteq \\ \prod_{w\in T_{ne}} \mathscr{A}_{ne} & \longrightarrow & \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)_\pi \end{array}$$

where the left-hand vertical map is given in each factor by $X_i \mapsto g_i$ for $1 \leq i \leq r$ (in the notation of the proof of Proposition 5.2.1), the bottom horizontal map sends a tuple of q^{-ne}-convergent power series $(F_w)_{w\in T_{ne}}$ to the function given on $w + \varpi_L^{ne}\mathscr{O}_L$ by $z \mapsto F(\pi_1(z-w), \ldots, \pi_r(z-w))$, and $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)_\pi$ denotes the E-vector space $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ endowed with the topology of described in §4.3. The commutativity of the diagram holds by the definition of the g_i. The map ι_{ne} is a topological isomorphism by the definition of $\mathscr{F}_{ne}(\mathscr{O}_L, E)$, the continuity of the left-hand vertical map is built into its construction, and the bottom horizontal map is continuous by the definition of the topology on the target. Thus the right-hand vertical inclusion $\mathcal{F}_{ne}(\mathscr{O}_L, E) \subseteq \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)_\pi$ is continuous, from which it follows that $\mathscr{F}_n(\mathscr{O}_L, E) \subseteq \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)_\pi$ is continuous for all $n \geq 0$. Therefore the identity map $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E) \to \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)_\pi$ is continuous. But the source is of compact type (by a straightforward generalization of the argument in [16, §16]), while the target is of compact type by Theorem 4.3.1; as bijective continuous linear maps between spaces of compact type are necessarily topological isomorphisms ([3, Theorem 1.1.17]), the topologies coincide. $\square$

Let $\underline{k} \in \mathbf{Z}^r_{\geq 0}$. The image under ι_n of the finite-dimensional subspace $\prod_{w \in T_n} \mathscr{A}^{\underline{k}} \subseteq \prod_{w \in T_n} \mathscr{A}_n$ will be denoted $\mathscr{F}_n^{\underline{k}}(\mathscr{O}_L, E)$. The inductive limit $\varinjlim_n \mathscr{F}_n^{\underline{k}}(\mathscr{O}_L, E)$ inside $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ is the subspace $\mathscr{C}^{\mathrm{lp}\leq\underline{k}}(\mathscr{O}_L, E)$ of "locally polynomial functions of degree at most $\underline{k}$." Since finite-dimensional locally convex spaces over E are necessarily equipped with their finest locally convex topologies, the locally convex inductive limit topology on $\mathscr{C}^{\mathrm{lp}\leq\underline{k}}(\mathscr{O}_L, E) = \varinjlim_n \mathscr{F}_n^{\underline{k}}(\mathscr{O}_L, E)$ is its finest locally convex topology. Thus if U is a locally convex space over E, any linear map $\mathscr{C}^{\mathrm{lp}\leq\underline{k}}(\mathscr{O}_L, E) \to U$ is continuous. The inclusion $\mathscr{C}^{\mathrm{lp}\leq\underline{k}}(\mathscr{O}_L, E) \subseteq \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ is then a homeomorphism onto its image, which is closed in $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$.

We now wish to state a result of Breuil generalizing classical work of Amice-Vélu and Vishik, but must first introduce a slight variation on the construction of functions via convergent power series. We have already shown how a series $F \in \mathscr{A}$ gives rise to an analytic function on $\mathscr{O}_L$ by (roughly) substituting the embeddings σ_i for the variables X_i (recall that $\mathscr{A} = \mathscr{A}_0$). By composing with certain continuous homomorphisms $\mathscr{A} \to \mathscr{A}_n$, we can essentially use $\mathscr{A}$ to produce the locally analytic functions arising from all the $\mathscr{A}_n$. Namely, if $w \in \mathscr{O}_L$ and $n \geq 0$, we will denote by $F((z-w)/\varpi_L^n)$ the locally analytic function on $\mathscr{O}_L$ that is given on $w+\varpi_L^n\mathscr{O}_L$ by $z \mapsto F(\sigma_1((z-w)/\varpi_L^n), \ldots, \sigma_r((z-w)/\varpi_L^n))$, and is extended by zero to the rest of $\mathscr{O}_L$. Note that this construction can also be described as the composite of the map $F \mapsto F_{n,w} : \mathscr{A}_n \to \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ with the continuous E-algebra homomorphism $\mathscr{A} \to \mathscr{A}_n$ given by sending X_i to $\sigma_i(\varpi_L)^{-n}X_i$ for $1 \leq i \leq r$. The resulting linear map $\mathscr{A} \to \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ is continuous by the definition of the locally convex topology on the target.

Let U be an E-Banach space and let $\|\cdot\|_U$ denote a choice of norm on U inducing its topology. The following definition is independent of this choice in the sense that if the condition in the definition holds for one norm for U, it holds for any other (with a possibly different constant). The definition is an immediate translation of [4, Definition 3.12] from the case $L = \mathbf{Q}_p$.

Definition 5.2.1. Let $\alpha \in E^\times$. An element $\varphi \in \mathcal{L}(\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E), U)$ (respectively $\varphi \in \mathcal{L}(\mathscr{C}^{\mathrm{lp}\leq\underline{k}}(\mathscr{O}_L, E), U)$) is said to be α-tempered if there is a constant $C > 0$ such that for each $F \in \mathscr{A}$ (respectively $F \in \mathscr{A}^{\underline{k}}$), $w \in \mathscr{O}_L$, and $n \in \mathbf{Z}_{\geq 0}$, we have

$$\left\|\varphi\left(F\left(\frac{z-w}{\varpi_L^n}\right)\right)\right\|_U \leq C|\alpha|^{-n}\|F\|_{\mathscr{A}}. \tag{5.3}$$

We denote by $\mathcal{L}(\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E), U)^\alpha$ (respectively $\mathcal{L}(\mathscr{C}^{\mathrm{lp}\leq\underline{k}}(\mathscr{O}_L, E), U)^\alpha$) the subspace of $\mathcal{L}(\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E), U)$ (respectively of $\mathcal{L}(\mathscr{C}^{\mathrm{lp}\leq\underline{k}}(\mathscr{O}_L, E), U)$) consisting of

α-tempered maps.

The following lemma equates the condition given in Definition 5.2.1, which is more suited for our argument, with the condition used in [2].

Lemma 5.2.3. *Let $\alpha \in E^\times$ and let $c = \operatorname{ord}(\alpha)$. An element $\varphi \in \mathcal{L}(\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E), U)$ (respectively $\varphi \in \mathcal{L}(\mathscr{C}^{\mathrm{lp}\leq \underline{k}}(\mathscr{O}_L, E), U)$) is α-tempered if and only if there is a constant $C > 0$ such that for each $w \in \mathscr{O}_L$, $n \in \mathbf{Z}_{\geq 0}$, and $\underline{m} \in \mathbf{Z}_{\geq 0}^r$ (respectively $\underline{m} \in \mathbf{Z}_{\geq 0}^r$ with $m_i \leq k_i$ for $1 \leq i \leq r$), we have*

$$\left\| \varphi\left(\mathbf{1}_{w+\varpi_L^n \mathscr{O}_L}(z) \prod_{i=1}^r \sigma_i(z-w)^{m_i} \right) \right\|_U \leq C q^{-n(|\underline{m}|-c)}. \tag{5.4}$$

Proof. Suppose $\varphi \in \mathcal{L}(\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E), U)$ is α-tempered in the sense of Definition 5.2.1 and let C be a constant so that Equation (5.3) holds. Given $w \in \mathscr{O}_L$, $n \geq 0$, and $\underline{m} \in \mathbf{Z}_{\geq 0}^r$, consider the element $F = \prod_{i=1}^r X_i^{m_i}$ of $\mathscr{A}$, noting that $\|F\|_{\mathscr{A}} = 1$. We have, by definition,

$$\begin{aligned} F\left(\frac{z-w}{\varpi_L^n}\right) &= \mathbf{1}_{w+\varpi_L^n \mathscr{O}_L}(z) \prod_{i=1}^r \sigma_i\left(\frac{z-w}{\varpi_L^n}\right)^{m_i} \\ &= \prod_{i=1}^r \sigma_i(\varpi_L)^{-nm_i} \mathbf{1}_{w+\varpi_L^n \mathscr{O}_L}(z) \prod_{i=1}^r \sigma_i(z-w)^{m_i}. \end{aligned}$$

Therefore, because φ is α-tempered,

$$\begin{aligned} \left\| \varphi\left(\mathbf{1}_{w+\varpi_L^n \mathscr{O}_L}(z) \prod_{i=1}^r \sigma_i(z-w)^{m_i} \right) \right\|_U &= \left\| \varphi\left(\prod_{i=1}^r \sigma_i(\varpi_L)^{nm_i} F\left(\frac{z-w}{\varpi_L^n}\right) \right) \right\|_U \\ &= \prod_{i=1}^r |\varpi_L|^{nm_i} \left\| \varphi\left(F\left(\frac{z-w}{\varpi_L^n}\right) \right) \right\|_U \\ &\leq \prod_{i=1}^r q^{-nm_i} C |\alpha|^{-n} \|F\|_{\mathscr{A}} \\ &= C q^{-n|\underline{m}|} q^{nc} = C q^{-n(|\underline{m}|-c)}. \end{aligned}$$

Thus (5.4) holds for φ. The same proof applies to an α-tempered $\varphi \in \mathcal{L}(\mathscr{C}^{\mathrm{lp}\leq\underline{k}}(\mathscr{O}_L, E), U)$, except that we only take $\underline{m} \in \mathbf{Z}^r_{\geq 0}$ with $m_i \leq k_i$ for $1 \leq i \leq r$.

Now assume conversely that Equation (5.4) holds for φ and all relevant data, with constant C. Running the computation above in reverse with F as before and noting that $\| \prod_{i=1}^r \sigma_i(\varpi_L)^{nm_i} F\|_{\mathscr{A}} = q^{-n|\underline{m}|}$, we find that

$$\|\varphi(F((z-w)/\varpi_L^n))\|_U \leq C|\alpha|^{-n}\|F\|_{\mathscr{A}}.$$

The strong triangle inequality then implies the desired inequality for $F \in \mathscr{A}$ that is a linear combination of monomials as above. The inequality then holds for a general $F \in \mathscr{A}$ because the polynomials in $\mathscr{A}$ are dense, and the association $F \mapsto F((z-w)/\varpi_L^n) : \mathscr{A} \to \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ is continuous. The case of $\varphi \in \mathcal{L}(\mathscr{C}^{\mathrm{lp}\leq\underline{k}}(\mathscr{O}_L, E), U)$ follows from the same argument (except that the final step involving density of the polynomials in $\mathscr{A}$ is not necessary). $\square$

We now state Breuil's generalization to arbitrary L of the result of Amice-Vélu and Vishik (whose result was stated for $L = \mathbf{Q}_p$).

Lemma 5.2.4. *If* $\mathrm{ord}(\alpha) < k_i + 1$ *for* $1 \leq i \leq r$*, then the restriction map*

$$\mathcal{L}(\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E), U) \to \mathcal{L}(\mathscr{C}^{\mathrm{lp}\leq\underline{k}}(\mathscr{O}_L, E), U)$$

induces an isomorphism

$$\mathcal{L}(\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E), U)^{\alpha} \simeq \mathcal{L}(\mathscr{C}^{\mathrm{lp}\leq\underline{k}}(\mathscr{O}_L, E), U)^{\alpha}.$$

Proof. This is a special case of [2, Lemma 6.1] (taking $J = \emptyset$, so that $J' = \mathrm{Hom}_{\mathbf{Q}_p-\mathrm{Alg}}(L, E)$ in that reference), taking into account the fact that the condition on linear maps imposed there is equivalent to the condition in Definition 5.2.1 by Lemma 5.2.3. □

5.3 Locally Algebraic and Locally Analytic Parabolic Inductions

Fix $\underline{k} \in \mathbf{Z}^r_{\geq 0}$. Recall the definitions of locally analytic and locally algebraic representations of G, given respectively in Definition 4.2.3 and Definition 4.2.5. We are interested in locally analytic representations of G induced from locally algebraic characters of T (regarded as characters of $\overline{B}$ via the projection $\overline{B} \to T$). More precisely, we consider characters of the form $\chi = \theta\psi_{\underline{k}}$, where $\theta : T \to E^\times$ has the form

$$\begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix} \mapsto \theta_1(a)\theta_2(d)$$

for smooth characters $\theta_1, \theta_2 : L^\times \to E^\times$, and $\psi_{\underline{k}} : T \to E^\times$ denotes the character

$$\begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix} \mapsto \prod_{i=1}^{r} \sigma_i(d)^{-k_i}.$$

If $V_{\underline{k}}$ is the irreducible algebraic representation $\bigotimes_{i=1}^r \mathrm{Sym}_E^{k_i}(E^2)$ of $\mathbf{G}_E$, and $W_{\underline{k}}$ denotes its contragredient, then $\psi_{\underline{k}}$ is the restriction to T of the highest weight of $W_{\underline{k}}$ (relative to the upper triangular Borel subgroup of $\mathbf{G}_E$).

For χ as above, we define the locally algebraic parabolic induction $I(\chi) = W_{\underline{k}} \otimes_E \mathrm{Ind}_{\overline{B}}^G(\theta)_{\mathrm{sm}}$, where the right tensor factor is the smooth parabolic induction of θ. Letting G act on $W_{\underline{k}}$ via the inclusion $G \hookrightarrow \mathbf{G}_E(E)$ and on $\mathrm{Ind}_{\overline{B}}^{G}(\theta)_{\mathrm{sm}}$ via the right regular representation on $\mathscr{C}^{\mathrm{sm}}(G, E)$, $I(\chi)$ becomes a $W_{\underline{k}}$-locally algebraic representation of G in the sense of [3, Proposition-Definition 4.2.6]. We also define $I^{\mathrm{la}}(\chi) = \mathrm{Ind}_{\overline{B}}^{G}(\chi)$, the locally analytic parabolic induction of χ [4, Example C], which consists of all functions $f \in \mathscr{C}^{\mathrm{la}}(G, E)$ satisfying $f(\bar{b}g) = \chi(\bar{b})f(g)$ for each $\bar{b} \in \overline{B}$ and $g \in G$. Letting G act on $I^{\mathrm{la}}(\chi)$ via the right regular representation on the locally convex space $\mathscr{C}^{\mathrm{la}}(G, E)$, and endowing $I^{\mathrm{la}}(\chi)$ with the induced topology, $I^{\mathrm{la}}(\chi)$ becomes a strongly admissible locally analytic representation of G [4, Proposition 1.21].

We may view $I(\chi)$ as a closed subrepresentation of $I^{\mathrm{la}}(\chi)$ in the following way. Let $\mathscr{O}(\mathbf{G}_E)$ denote the affine coordinate ring of $\mathbf{G}_E$, and let $\mathscr{C}^{\mathrm{alg}}(G, E)$ denote the image of the restriction map $\mathscr{O}(\mathbf{G}_E) \hookrightarrow \mathscr{C}^{\mathrm{la}}(G, E)$ (the restriction map is injective because $G \subseteq \mathbf{G}_E(E)$ is Zariski dense in $\mathbf{G}_E(E)$). This is the space of algebraic E-valued functions on G. With $e_{1,i}, e_{2,i}$ the standard elements in the i-th tensor factor of $V_{\underline{k}}$, $e_2 = \otimes_{i=1}^{r} e_{2,i}^{k_i}$ is a highest weight vector in $V_{\underline{k}}$ relative to the lower triangular Borel subgroup of $\mathbf{G}_E$, and the map $W_{\underline{k}} \to \mathscr{O}(\mathbf{G}_E)$ given by $w \mapsto (g \mapsto w(g^{-1}e_2))$ is a $\mathbf{G}_E(E)$-equivariant E-linear injection which, when composed with the isomorphism $\mathscr{O}(\mathbf{G}_E) \simeq \mathscr{C}^{\mathrm{alg}}(G, E)$, allows us to view $W_{\underline{k}}$ as a subrepresentation of $\mathscr{C}^{\mathrm{alg}}(G, E)$. Tensoring the injection $W_{\underline{k}} \hookrightarrow \mathscr{C}^{\mathrm{alg}}(G, E)$ with the inclusion $\mathrm{Ind}_{\overline{B}}^{G}(\theta)_{\mathrm{sm}} \subseteq \mathscr{C}^{\mathrm{sm}}(G, E)$ yields an injection $I(\chi) \hookrightarrow \mathscr{C}^{\mathrm{alg}}(G, E) \otimes_E \mathscr{C}^{\mathrm{sm}}(G, E)$, and following this with the map $\mathscr{C}^{\mathrm{alg}}(G, E) \otimes_E \mathscr{C}^{\mathrm{sm}}(G, E) \to \mathscr{C}^{\mathrm{la}}(G, E)$ given by multiplication of algebraic functions and smooth functions gives $I(\chi) \hookrightarrow \mathscr{C}^{\mathrm{la}}(G, E)$. Writing down the map $W_{\underline{k}} \hookrightarrow \mathscr{C}^{\mathrm{alg}}(G, E)$ explicitly (using the definition given above and the action of $g \in G$ on e_2) one finds that the image of $I(\chi)$ in $\mathscr{C}^{\mathrm{la}}(G, E)$ is contained in $I^{\mathrm{la}}(\chi)$. Thus $I(\chi)$ is canonically a G-stable subspace of $I^{\mathrm{la}}(\chi)$. Moreover, $I(\chi)$ is closed in $I^{\mathrm{la}}(\chi)$, and its subspace topology coincides with its finest locally convex topology (with respect to which it is an admissible locally $W_{\underline{k}}$-algebraic representation of G by [3, Proposition 6.3.10]).

We may now state our main result. The proof will be given in §5.4.

Theorem 5.3.1. *Assume that*

(i). $\mathrm{ord}(\theta_1(\varpi_L)) < k_i + 1$ *for* $1 \leq i \leq r$, *and that*

(ii). $\chi|_{Z(G)}$ *is unitary.*

Then for any unitary E-Banach space representation U of G, the restriction map

$$\mathcal{L}_G(I^{\mathrm{la}}(\chi), U) \to \mathcal{L}_G(I(\chi), U) \tag{5.5}$$

is an isomorphism.

When $L = \mathbf{Q}_p$, this is essentially Proposition 2.5 of [4]. A version of this result (for general L) is also proved as Theorem 7.1 of [2]. Breuil's result applies to more general locally J-analytic parabolic inductions, where J is a subset of $\mathrm{Hom}_{\mathbf{Q}_p-\mathrm{Alg}}(L, E)$, but he restricts attention to injective linear maps. (See [2, p. 10] for the definition of the locally J-analytic induction; locally algebraic induction corresponds to $J = \emptyset$, while locally analytic induction corresponds to $J = \mathrm{Hom}_{\mathbf{Q}_p-\mathrm{Alg}}(L,E))$. We follow Emerton's argument from [4, §3], which is somewhat more representation-theoretic than Breuil's (although we do make crucial use [2, Lemma 6.1], stated as Lemma 5.2.4 in the previous section, in place of Emerton's appeal to the classical result of Amice-Vélu and Vishik).

5.4 Proof of Theorem 5.3.1

In this section we prove Theorem 5.3.1. We therefore assume that $\mathrm{ord}(\theta_1(\varpi_L)) < k_i + 1$ for $1 \leq i \leq r$ and that $\chi|_{Z(G)}$, the central character of $I^{\mathrm{la}}(\chi)$, is unitary. It will be clear in the argument where these hypotheses are invoked. Our proof closely follows that of Emerton in [4, §3]. The key input to make Emerton's argument go through in the general case is provided by the description of $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ from §5.2 and the accompanying Lemma 5.2.4 (which takes the place of the result of Amice-Vélu and Vishik used by Emerton).

If V is one of $I^{\mathrm{la}}(\chi), I(\chi)$, denote by $V(N_0)$ the closed subspace of functions in V whose support lies in $\overline{B}N_0$. This is a $G^+(1)$-invariant closed subspace of $I^{\mathrm{la}}(\chi)$. The following result is proved for $L = \mathbf{Q}_p$ in [4, Lemma 3.1], but the argument given there applies to an arbitrary finite extension L of $\mathbf{Q}_p$.

Lemma 5.4.1. *For any E-Banach space representation U of G, the restriction maps*

$$\mathcal{L}_G(I^{\mathrm{la}}(\chi), U) \to \mathcal{L}_{G^+(s)}(I^{\mathrm{la}}(\chi)(N_0), U)$$

and

$$\mathcal{L}_G(I(\chi), U) \to \mathcal{L}_{G^+(s)}(I(\chi)(N_0), U)$$

are isomorphisms for all integers $s \geq 1$.

Thus, for an E-Banach space representation U of G and each $s \geq 1$, there is a commutative diagram of restriction maps

$$\begin{array}{ccc} \mathcal{L}_G(I^{\mathrm{la}}(\chi), U) & \longrightarrow & \mathcal{L}_G(I(\chi), U) \\ \downarrow & & \downarrow \\ \mathcal{L}_{G^+(s)}(I^{\mathrm{la}}(\chi)(N_0), U) & \longrightarrow & \mathcal{L}_{G^+(s)}(I(\chi)(N_0), U) \end{array}$$

where the vertical maps are isomorphisms. To prove Theorem 5.3.1, which is the assertion that the top horizontal restriction map is an isomorphism, it therefore suffices to prove that the bottom horizontal arrow is an isomorphism for some $s \geq 1$. We ultimately reduce this to Lemma 5.2.4. By [5, Lemma 2.3.3], restricting functions in $I^{\mathrm{la}}(\chi)$ to N_0 yields a topological isomorphism $I^{\mathrm{la}}(\chi)(N_0) \simeq \mathscr{C}^{\mathrm{la}}(N_0, U)$, and composing this with the topological isomorphism $\mathscr{C}^{\mathrm{la}}(N_0, E) \simeq \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ (see the discussion of function spaces in §5.1), we obtain an isomorphism

$$I^{\mathrm{la}}(\chi)(N_0) \simeq \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E). \tag{5.6}$$

Restricting (5.6) to $I(\chi)(N_0)$, and using the explicit description of the embedding $I(\chi) \hookrightarrow I^{\mathrm{la}}(\chi)$ of §5.3, we obtain an induced isomorphism

$$I(\chi)(N_0) \simeq \mathscr{C}^{\mathrm{lp}\leq \underline{k}}(\mathscr{O}_L, E). \tag{5.7}$$

Since $I^{\mathrm{la}}(\chi)(N_0)$ is a $G^+(1)$-stable subspace of $I^{\mathrm{la}}(\chi)$, using the isomorphism (5.6), we can transfer the action of $G^+(1)$ on $I^{\mathrm{la}}(\chi)(N_0)$ to an action of $G^+(1)$

on $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$. A computation shows that if $f \in \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ and $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in G^+(1)$, then we have, for any $z \in \mathscr{O}_L$,

$$(gf)(z) = \begin{cases} 0 & \text{if } \dfrac{b+dz}{a+cz} \notin \mathscr{O}_L \\ \left(\prod_{i=1}^r \sigma_i \left(\dfrac{a+cz}{\det(g)} \right)^{k_i} \right) \theta_1(a+cz)\theta_2\left(\dfrac{\det(g)}{a+cz} \right) f\left(\dfrac{b+dz}{a+cz} \right) & \text{if } \dfrac{b+dz}{a+cz} \in \mathscr{O}_L. \end{cases} \tag{5.8}$$

As $I(\chi)(N_0)$ is a $G^+(1)$-stable subspace of $I^{\mathrm{la}}(\chi)(N_0)$, in light of (5.7), $\mathscr{C}^{\mathrm{lp}\leq \underline{k}}(\mathscr{O}_L, E)$ is a $G^+(1)$-stable subspace of $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ for the action defined in (5.8).

We also need to define an action of $G_0(1)$ on $\mathscr{A}$, which we now explain. Let $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in G_0(1)$. By the definition of $G_0(1) \subseteq \mathrm{GL}_2(\mathscr{O}_L)$, each of a, b, c, d is in $\mathscr{O}_L$ and $c \in \varpi_L \mathscr{O}_L$, so $a, d \in \mathscr{O}_L^\times$. Thus, for $1 \leq i \leq r$, the series

$$\begin{aligned} \frac{\sigma_i(b) + \sigma_i(d)X_i}{\sigma_i(a) + \sigma_i(c)X_i} &= \frac{\sigma_i(b) + \sigma_i(d)X_i}{\sigma_i(a)} \frac{1}{1 - \left(-\frac{\sigma_i(c)}{\sigma_i(a)} X_i \right)} \\ &= \frac{\sigma_i(b) + \sigma_i(d)X_i}{\sigma_i(a)} \sum_{m=0}^{\infty} (-1)^m \left(\frac{\sigma_i(c)}{\sigma_i(a)} \right)^m X_i^m \end{aligned}$$

is an element of $\mathscr{A}$ of norm 1. Therefore, there is a unique continuous E-algebra endomorphism $\nu_g : \mathscr{A} \to \mathscr{A}$ with $\nu_g(X_i) = (\sigma_i(b) + \sigma_i(d)X_i)/(\sigma_i(a) + \sigma_i(c)X_i)$ for $1 \leq i \leq r$, and the operator norm of ν_g is at most 1. A (slightly messy but straightforward) computation shows that $\nu_{g_1 g_2} = \nu_{g_1} \circ \nu_{g_2}$, and it follows that in fact each ν_g is an isometry of $\mathscr{A}$. We now define the $G_0(1)$-action on $\mathscr{A}$ by

$$g(F(X_1, \ldots, X_r)) = \left(\prod_{i=1}^r \sigma_i \left(\frac{a+cz}{\det(g)} \right)^{k_i} \right) \theta_2(\det(g)) \eta_g(F(X_1, \ldots, X_r)). \tag{5.9}$$

Because the factor multiplying $\eta_g(F(X_1, \ldots, X_r))$ in (5.9) is of Gauss norm 1 (since $\det(g) \in \mathscr{O}_L^\times$, $\theta_2(\det(g)) \in \mathscr{O}_E^\times$), this $G_0(1)$-action on $\mathscr{A}$ is unitary.

Moreover, the factor ensures that $\mathscr{A}^{\underline{k}}$ is a $G_0(1)$-stable subspace of $\mathscr{A}$ for this action.

Comparing the formulas above, we find that, if $g \in G_0(1) \subseteq G^+(1)$, $F \in \mathscr{A}$, and $z \in \mathscr{O}_L$, then

$$g(F(z)) = \frac{\theta_1(a+cz)}{\theta_2(a+cz)}((gF)(z)), \tag{5.10}$$

where on the left-hand side g acts on $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ via (5.8), and on the right-hand side g acts on $\mathscr{A}$ via (5.9) (recall that for $F \in \mathscr{A}$, the notation $F(z)$ means the function $z \mapsto F(\sigma_1(z), \ldots, \sigma_r(z))$ on $\mathscr{O}_L$, and note that if $g \in G_0(1)$, then $(b+dz)/(a+cz) \in \mathscr{O}_L$ for any $z \in \mathscr{O}_L$).

We now follow Emerton in relating the actions just defined to the notion of an α-tempered linear map (Definition 5.2.1), where $\alpha = \theta_1(\varpi_L)$. In preparation, we introduce the subset B' of B defined by

$$B' = \left\{ \begin{pmatrix} \varpi_L^n & -w \\ 0 & 1 \end{pmatrix} : n \in \mathbf{Z}_{\geq 0}, w \in \mathscr{O}_L \right\}.$$

This is a submonoid of $B^+ = N_0 T^+$ (see §5.1 for the notation) since we can write, for any $n \geq 0$ and $w \in \mathscr{O}_L$,

$$\begin{pmatrix} \varpi_L^n & -w \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & -w \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \varpi_L^n & 0 \\ 0 & 1 \end{pmatrix}.$$

Lemma 5.4.2. *Any element $b \in B^+$ may be written as $zb't$ with $z \in Z(G)$, $b' \in B'$, and $t \in T_0$.*

Proof. As $b \in B^+ = N_0 T^+$, we may write

$$b = \begin{pmatrix} 1 & w \\ 0 & 1 \end{pmatrix} \begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix} = \begin{pmatrix} a & wd \\ 0 & d \end{pmatrix}$$

with $w \in \mathscr{O}_L$ and $ad^{-1} \in \mathscr{O}_L$. Then

$$b = \begin{pmatrix} d & 0 \\ 0 & d \end{pmatrix} \begin{pmatrix} \varpi_L^{\mathrm{ord}(a)-\mathrm{ord}(d)} & w \\ 0 & 1 \end{pmatrix} \begin{pmatrix} ad^{-1}\varpi_L^{\mathrm{ord}(d)-\mathrm{ord}(a)} & 0 \\ 0 & 1 \end{pmatrix}$$

is a decomposition of b of the form $zb't \in Z(G)B'T_0$. □

Lemma 5.4.3. *Let $\mathscr{C}$ denote either $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ or $\mathscr{C}^{\mathrm{lp}\leq\underline{k}}(\mathscr{O}_L, E)$, with $\mathscr{A}_{\mathscr{C}}$ denoting $\mathscr{A}$ in the former case and $\mathscr{A}^{\underline{k}}$ in the latter case. If U is an E-Banach space and $\varphi \in \mathcal{L}(\mathscr{C}, U)$, then φ is $\theta_1(\varpi_L)$-tempered if and only if there exists a positive constant C such that $\|\varphi(b(F(z)))\|_U \leq C\|F\|_{\mathscr{A}}$ for all $F \in \mathscr{A}_{\mathcal{C}}$ and $b \in B^+$ (where $\|\cdot\|_U$ is any choice of norm on U defining its topology).*

Proof. Given $F \in \mathscr{A}_{\mathscr{C}}$, equation (5.8) gives

$$\begin{pmatrix} \varpi_L^n & -w \\ 0 & 1 \end{pmatrix} F(z) = \theta_1(\varpi_L)^n F((z-w)/\varpi_L^n).$$

It then follows from Definition 5.2.1 that φ is $\theta_1(\varpi_L)$-tempered if and only if there is a constant $C > 0$ such that $\|\varphi(b'(F(z)))\|_U \leq C\|F\|_{\mathscr{A}}$ for all $F \in \mathscr{A}_{\mathscr{C}}$ and $b' \in B'$. If the condition in the statement of the lemma holds, then certainly this condition holds, since $B' \subseteq B^+$. Conversely, suppose $\|\varphi(b'(F(z)))\|_U \leq C\|F\|_{\mathscr{A}}$ for all $F \in \mathscr{A}_{\mathscr{C}}$ and $b' \in B'$, and let $b \in B^+$. In accordance with Lemma 5.4.2, we may write $b = z'b't$ with $z' \in Z(G)$, $b' \in B'$, and $t = \left(\begin{smallmatrix} a & 0 \\ 0 & d \end{smallmatrix}\right)$ with $a, d \in \mathscr{O}_L^\times$. Noting that the action of $Z(G) \subseteq G^+(1)$ on $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ is given by the character χ (because this is the character by which

$Z(G)$ acts on $I^{\mathrm{la}}(\chi)(N_0))$, which is assumed unitary, we have

$$\begin{aligned}
\|\varphi(b(F(z)))\|_U &= \|\varphi(z'b't(F(z)))\|_U \\
&= \|\chi(z')\varphi(b't(F(z)))\|_U \\
&= \left\|\chi(z')\frac{\theta_1(a)}{\theta_2(a)}\varphi(b'((tF)(z)))\right\|_U \\
&= \|\varphi(b'((tF)(z)))\|_U \leq C\|tF\|_{\mathscr{A}} = \|F\|_{\mathscr{A}}.
\end{aligned}$$

(We have used (5.10) in going from the second to the third line and the unitarity of the action of $G_0(1)$ on $\mathscr{A}$ in the final equality.) Thus the condition in the statement of the lemma holds. □

Remark 5.4.1. The preceding proof is the only point where the unitarity of the central character $\chi|_{Z(G)}$ is used.

Lemma 5.4.4. *In the notation of Lemma 5.4.3, if U admits a unitary action of B^+, then $\mathcal{L}_{B^+}(\mathscr{C}, U) \subseteq \mathcal{L}(\mathscr{C}, U)^{\theta_1(\varpi_L)}$.*

Proof. Let $\|\cdot\|_U$ be a B^+-invariant norm on U and let $\varphi \in \mathcal{L}_{B^+}(\mathscr{C}, U)$. By the definition of the topology on $\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ (see the discussion following Remark 5.2.1) and the continuity of φ, the restriction of φ to the image of the map $F \mapsto F(z) : \mathscr{A}_{\mathscr{C}} \to \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E)$ is bounded, i.e., there is a constant $C > 0$ such that $\|\varphi(F(z))\|_U \leq C\|F\|_{\mathscr{A}}$ for all $F \in \mathscr{A}_{\mathscr{C}}$. Therefore, if $b \in B^+$ and $F \in \mathscr{A}_{\mathscr{C}}$, we have

$$\|\varphi(b(F(z)))\|_U = \|b\varphi(F(z))\|_U = \|\varphi(F(z))\|_U \leq C\|F\|_{\mathscr{A}},$$

where the first equality follows from the assumed B^+-equivariance of φ and the second follows from the B^+-invariance of $\|\cdot\|_U$. Thus the condition in Lemma 5.4.3 holds, so φ is $\theta_1(\varpi_L)$-tempered. $\square$

We may now complete the proof of Theorem 3.1. Thus we assume that U is a unitary E-Banach space representation of G with $\|\cdot\|_U$ a G-invariant norm. Recall that our goal was to show that, for some integer $s \geq 1$, the restriction map

$$\mathcal{L}_{G^+(s)}(I^{\mathrm{la}}(\chi)(N_0), U) \to \mathcal{L}_{G^+(s)}(I(\chi)(N_0), U)$$

is an isomorphism. Using the $G^+(1)$-equivariant isomorphisms (5.6) and (5.7), it is equivalent to prove that

$$\mathcal{L}_{G^+(s)}(\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E), U) \to \mathcal{L}_{G^+(s)}(\mathscr{C}^{\mathrm{lp}\leq \underline{k}}(\mathscr{O}_L, E), U) \tag{5.11}$$

is an isomorphism for some $s \geq 1$. We will show that it is enough to take s equal to the conductor exponent of the restrictions of θ_1, θ_2 to $\mathscr{O}_L^\times$, i.e., we assume that θ_1, θ_2 are trivial when restricted to $1 + \varpi_L^s \mathscr{O}_L$. There is some such s because the θ_i are smooth. Now, by Lemma 5.4.4 (and the fact that $B^+ \leq G^+(s)$, so that $G^+(s)$-equivariant maps are also B^+-equivariant), we have

$$\mathcal{L}_{G^+(s)}(\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E), U) \subseteq \mathcal{L}(\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E), U)^{\theta_1(\varpi_L)} \tag{5.12}$$

and

$$\mathcal{L}_{G^+(s)}(\mathscr{C}^{\mathrm{lp}\leq \underline{k}}(\mathscr{O}_L, E), U) \subseteq \mathcal{L}(\mathscr{C}^{\mathrm{lp}\leq \underline{k}}(\mathscr{O}_L, E), U)^{\theta_1(\varpi_L)}. \tag{5.13}$$

As we are assuming that $\operatorname{ord}(\theta_1(\varpi_L)) < k_i + 1$ for all i, Lemma 5.2.4 implies that the restriction map

$$\mathcal{L}(\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E), U)^{\theta_1(\varpi_L)} \to \mathcal{L}(\mathscr{C}^{\mathrm{lp}\leq \underline{k}}(\mathscr{O}_L, E), U)^{\theta_1(\varpi_L)} \tag{5.14}$$

is an isomorphism. Thus, in light of the inclusions (5.12) and (5.13), and the injectivity of (5.14), we conclude that (5.11) is injective.

To prove the surjectivity of (5.11), fix $\varphi_0 \in \mathcal{L}_{G^+(s)}(\mathscr{C}^{\mathrm{lp}\leq \underline{k}}(\mathscr{O}_L, E), U)$. Because of the inclusion (5.13) and the surjectivity of (5.14), there is an element $\varphi \in \mathcal{L}(\mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E), U)$ that is $\theta_1(\varpi_L)$-tempered and restricts to φ_0 on $\mathscr{C}^{\mathrm{lp}\leq \underline{k}}(\mathscr{O}_L, E)$. It remains to prove that φ is $G^+(s)$-equivariant. To do this, we consider, for a fixed $g \in G^+(s)$, the continuous linear map

$$\varphi' : f \mapsto g^{-1}\varphi(gf) : \mathscr{C}^{\mathrm{la}}(\mathscr{O}_L, E) \to U.$$

Since φ_0 is $G^+(s)$-equivariant, the restriction of φ' to $\mathscr{C}^{\mathrm{lp}\leq \underline{k}}(\mathscr{O}_L, E)$ coincides with that of φ, so if φ' can be shown to be tempered, the injectivity of (5.14) will give $\varphi' = \varphi$, proving the desired equivariance.

We will show that φ' satisfies the condition in Lemma 5.4.4. Because φ is $\theta_1(\varpi_L)$-tempered, φ satisfies this condition, i.e., there is a constant $C > 0$ such that $\|\varphi(b(F(z)))\|_U \leq C\|F\|_{\mathscr{A}}$ for all $b \in B^+$ and $F \in \mathscr{A}$. If $b \in B^+$, then $gb \in G^+(s) = B^+\overline{N}(s)$, so we may write $gb = b_1\bar{n}$ for some $b_1 \in B^+$ and $\bar{n} = \left(\begin{smallmatrix} 1 & 0 \\ w & 1 \end{smallmatrix}\right) \in \overline{N}(s)$ (so $w \in \varpi_L^s\mathscr{O}_L$). Then, using the G-invariance of $\|\cdot\|_U$, Equation (5.10), the assumption that θ_1 and θ_2 are trivial on $1 + \varpi_L^s\mathscr{O}_L$, and the $\overline{N}(s)$-invariance of $\|\cdot\|_{\mathscr{A}}$, we find that

$$\begin{aligned}\|\varphi'(b(F(z)))\|_U &= \|g^{-1}\varphi(gb(F(z)))\|_U \\ &= \|\varphi(b_1\bar{n}(F(z)))\|_U \\ &= \left\|\frac{\theta_1(1+wz)}{\theta_2(1+wz)}\varphi(b_1((\bar{n}F)(z)))\right\|_U \\ &= \|\varphi(b_1((\bar{n}F)(z)))\|_U \leq C\|\bar{n}F\|_{\mathscr{A}} = C\|F\|_{\mathscr{A}}\end{aligned}$$

for any $F \in \mathscr{A}$. Thus, by Lemma 5.4.4, φ' is tempered, as desired.

Bibliography

[1] P. N. Balister and S. Howson. Note on Nakayama's lemma for compact Λ-modules. *Asian J. Math.*, 1(2):224–229, 1997.

[2] Christophe Breuil. Remarks on some locally $\mathbb{Q}_p$-analytic representations of $\mathrm{GL}_2(F)$ in the crystalline case. In *Non-abelian fundamental groups and Iwasawa theory*, volume 393 of *London Math. Soc. Lecture Note Ser.*, pages 212–238. Cambridge Univ. Press, Cambridge, 2012.

[3] Matthew Emerton. Locally analytic vectors in representations of locally p-adic analytic groups. *arXiv preprint math/0405137*, 2004.

[4] Matthew Emerton. p-adic L-functions and unitary completions of representations of p-adic reductive groups. *Duke Math. J.*, 130(2):353–392, 2005.

[5] Matthew Emerton. Jacquet modules of locally analytic representations of p-adic reductive groups ii. the relation to parabolic induction. *J. Institut Math. Jussieu*, 2007.

[6] Matthew Emerton, Robert Pollack, and Tom Weston. Variation of Iwasawa invariants in Hida families. *Invent. Math.*, 163(3):523–580, 2006.

[7] Ralph Greenberg. Iwasawa theory for p-adic representations. In *Algebraic number theory*, volume 17 of *Adv. Stud. Pure Math.*, pages 97–137. Academic Press, Boston, MA, 1989.

[8] Ralph Greenberg. Selmer groups and congruences. In *Proceedings of the International Congress of Mathematicians. Volume II*, pages 231–248, New Delhi, 2010. Hindustan Book Agency.

[9] Ralph Greenberg and Vinayak Vatsal. On the Iwasawa invariants of elliptic curves. *Invent. Math.*, 142(1):17–63, 2000.

[10] Barry Mazur. Rational points of abelian varieties with values in towers of number fields. *Invent. Math.*, 18:183–266, 1972.

[11] Jürgen Neukirch, Alexander Schmidt, and Kay Wingberg. *Cohomology of number fields*, volume 323 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, second edition, 2008.

[12] Bernadette Perrin-Riou. *p-adic L-functions and p-adic representations*, volume 3 of *SMF/AMS Texts and Monographs*. American Mathematical Society, Providence, RI; Société Mathématique de France, Paris, 2000. Translated from the 1995 French original by Leila Schneps and revised by the author.

[13] Robert Pollack and Tom Weston. On anticyclotomic μ-invariants of modular forms. *Compos. Math.*, 147(5):1353–1381, 2011.

[14] P. Schneider and J. Teitelbaum. $U(\mathfrak{g})$-finite locally analytic representations. *Represent. Theory*, 5:111–128 (electronic), 2001. With an appendix by Dipendra Prasad.

[15] P. Schneider and J. Teitelbaum. Banach space representations and Iwasawa theory. *Israel J. Math.*, 127:359–380, 2002.

[16] Peter Schneider. *Nonarchimedean functional analysis*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2002.

[17] Peter Schneider. *p-adic Lie groups*, volume 344 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer, Heidelberg, 2011.

[18] Peter Schneider and Jeremy Teitelbaum. Locally analytic distributions and p-adic representation theory, with applications to GL_2. *J. Amer. Math. Soc.*, 15(2):443–468 (electronic), 2002.

[19] Peter Schneider and Jeremy Teitelbaum. Algebras of p-adic distributions and admissible representations. *Invent. Math.*, 153(1):145–196, 2003.

[20] Tom Weston. Iwasawa invariants of Galois deformations. *Manuscripta Math.*, 118(2):161–180, 2005.