Galois deformation theory II

Ravi Fernando – fernando@berkeley.edu

February 14, 2018[*]

1 Setup and goals

Let's begin by recalling some setup from the last talk. We fix a finite residue field k of characteristic p. A *coefficient ring* is a complete local noetherian ring with residue field k; morphisms of coefficient rings must induce the identity map on k. If Λ is a coefficient ring and $\mathcal{A}$ is a coefficient Λ-algebra, we let $\hat{C}_\Lambda(\mathcal{A})$ be the category whose objects are coefficient Λ-algebras equipped with an augmentation maps to $\mathcal{A}$, and whose morphisms are commutative diagrams

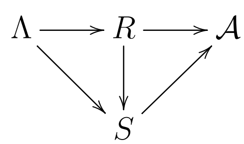

of coefficient rings. We let $C_\Lambda(\mathcal{A})$ be the full subcategory of artinian rings in $\hat{C}_\Lambda(\mathcal{A})$. If $\Lambda = W(k)$ or $\mathcal{A} = k$ (respectively the initial and final objects in the category of coefficient rings), then we omit them from the notation, as the respective maps are redundant. For the first half of this talk, Λ and $\mathcal{A}$ will just come along for the ride; for mostly unimportant pedagogical reasons, we will work in C_Λ and $\hat{C}_\Lambda$, and we will indicate the minor changes needed to adapt our work to the relative setting. (Eventually we will actually need to distinguish between the absolute and relative cases.) In fact, we will soon see that we can get away with working almost exclusively with artinian rings.

Given a p-adic Galois representation ρ, our general goal in this seminar is to turn nice information (say, modularity) about the reduction $\overline{\rho}$ into nice information about ρ itself. Since we may only have limited information about ρ itself, our strategy is to first find some moduli space of all deformations ρ' of $\overline{\rho}$, then identify a closed subscheme of "nice" deformations that ρ is known to lie in, and then to use the geometry of this subscheme to obtain results about these nice deformations.

Given a residual representation $\overline{\rho} : \Pi \to \mathrm{GL}_N(k)$ and a coefficient ring Λ, we are interested in

[*]Notes for a talk in Berkeley's number theory seminar. Main reference: Barry Mazur, "An introduction to the deformation theory of Galois representations", from *Modular Forms and Fermat's Last Theorem*, edited by Cornell, Silverman, and Stevens.

the deformation functor $D_{\overline{\rho}} : \hat{C}_\Lambda \to \mathbf{Set}$ that associates to any coefficient Λ-algebra B the set of strict equivalence classes of deformations of $\overline{\rho}$ to B. Given a representation $\rho : \Pi \to \mathrm{GL}_N(\mathcal{A})$, we define $D_\rho : \hat{C}_\Lambda(\mathcal{A}) \to \mathbf{Set}$ by the same formula; in this relative situation, we are treating $\mathcal{A}$ as our "residue ring".

Note that these are covariant functors from categories of rings, which correspond to contravariant functors from categories of affine (formal) schemes. Our hope is to obtain representability results; i.e. to show that D_ρ is isomorphic to the representable functor $D_R = \mathrm{Hom}(R, -)$.

2 Categorical first steps

We first observe some obvious necessary criteria for representability. We will state these in the "absolute" setting of C_Λ and $\hat{C}_\Lambda$ for convenience, but everything carries over to the relative case without change. Suppose $D = D_R : \hat{C}_\Lambda \to \mathbf{Set}$ is a given representable functor. For any $A \in \hat{C}_\Lambda$, we write $A = \lim_{\leftarrow n} A/m_A^n$. The universal property of the inverse limit then gives

$$D_R(A) = \mathrm{Hom}(R, A) \xrightarrow{\sim} \lim_{\leftarrow n} \mathrm{Hom}(R, A/m_A^n) = \lim_{\leftarrow n} D_R(A/m_A^n). \tag{1}$$

Mazur calls any functor with this property (i.e. $D(A) = \lim_{\leftarrow n} D(A/m_A^n)$) *continuous*. This notation seems a little nonstandard, since a "continuous functor" usually means one that commutes with all limits. But in our context of deformations of Galois representations, it turns out that continuity (in the weak sense) comes for free:

Proposition 1. (Continuity) Let $\rho : \Pi \to \mathrm{GL}_N(A)$ be a lifting of $\overline{\rho}$ to $A \in \hat{C}_\Lambda$. The functors $D_{\overline{\rho}}$ and D_ρ are continuous in the sense described above.

So the functors we are interested in are completely determined by their restrictions to the subcategory C_Λ. For convenience, we will restrict all our functors to the smaller category from now on. We will hope to obtain pro-representability results; i.e. results about representability of a functor $D : C_\Lambda \to \mathbf{Set}$ by an object of the completed category $\hat{C}_\Lambda$. (I may sometimes carelessly write "representable" when I mean "pro-representable".)

Of course there is nothing special about inverse limits indexed by $\mathbb{N}$; the universal property implies that representable functors must commute with all limits that exist. One special case will be particularly important to us. Suppose we have a fibered square in the category C_Λ:

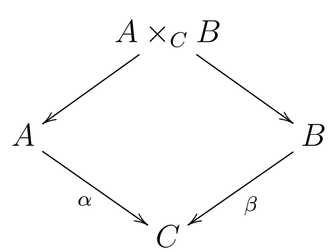

Warning 1. This is a fibered product of rings, not of schemes. The tensor product has arrows in the opposite direction. Geometrically, we're gluing two thickenings Spec A and Spec B along their common closed subscheme Spec C. (Recall that Spec of any artinian coefficient ring is a point, and similarly Spf of any not necessarily artinian coefficient ring.) Explicitly, we construct $A \times_C B$ as $\{(a, b) \in A \times B : \alpha(a) = \beta(b)\}$, with the obvious projections. One can easily verify that this is in C_Λ and that it is the fibered product in both C_Λ and $\hat{C}_\Lambda$ (provided that A, B, and C are in C_Λ), and similarly for $C_\Lambda(\mathcal{A})$ and $\hat{C}_\Lambda(\mathcal{A})$, and indeed even in **Ring** and **Set**.

Warning 2. The completed category $\hat{C}_\Lambda(A)$ is *not* closed under fibered products. For example, consider the diagram $k[[X, Y]] \to k[[X]] \hookleftarrow k$, where $Y \mapsto 0$. The fibered product is $k \oplus Y \cdot k[[X, Y]]$, which is not noetherian. This problem does not appear if both maps are surjective.

If we apply a functor $D : C_\Lambda \to \mathbf{Set}$ to the diagram above, we get a commutative diagram

$$\begin{array}{ccccc} & & D(A \times_C B) & & \\ & \swarrow & & \searrow & \\ D(A) & & & & D(B) \\ & \overset{\alpha}{\searrow} & & \overset{\beta}{\swarrow} & \\ & & D(C) & & \end{array}$$

which induces a map $h : D(A \times_C B) \to D(A) \times_{D(C)} D(B)$. (Fibered products of sets are constructed by the same formula we gave above.) If D is representable by an object R of $\hat{C}_\Lambda$, then the universal property of $A \times_C B$ (in $\hat{C}_\Lambda$) says that h must be a bijection. So every (pro-)representable functor preserves fibered products.

3 The infinitesimal picture

Preserving fibered products seems like a fairly trivial necessary condition for representability, but it is surprisingly close to being sufficient as well. To bridge the gap, we need to study some geometry of the moduli space, namely the tangent space at its (unique) point.

Fix a coefficient Λ-algebra R, which may not be artinian. We will first define the Zariski tangent space t_R, and then attempt to recover this data from the representable functor D_R.

Definition 2. The *cotangent space* to R is the k-vector space $m_R/(m_R^2 + m_\Lambda \cdot R)$, and the *tangent space* is its dual $t_R = \mathrm{Hom}_k(m_R/(m_R^2 + m_\Lambda \cdot R), k)$.

Note that since R is noetherian, its cotangent space is finite-dimensional and thus equals its own double dual.

Perhaps not surprisingly, this definition can be recast in terms of the ring of dual numbers $k[\epsilon]/(\epsilon^2)$ (which we will abbreviate as $k[\epsilon]$ from here on):

Proposition 3. *There is a natural isomorphism of k-vector spaces*

$$\mathrm{Hom}_k(m_R/(m_R^2 + m_\Lambda \cdot R), k) \xrightarrow{\sim} \mathrm{Hom}_{\hat{C}_\Lambda}(R, k[\epsilon]). \tag{2}$$

We leave the proof as an exercise, with the hint that a Λ-algebra homomorphism from R to $k[\epsilon]$ must send m_R to $k \cdot \epsilon$ and both m_R^2 and m_Λ to 0. We also remark that the k-vector space structure on the right hand is not obvious; this is about to cause us some headaches. The point is that viewing $k[\epsilon]$ as $k \oplus k \cdot \epsilon$, the k-component of the map is forced upon us, and we can simply do pointwise addition and scalar multiplication in the $k \cdot \epsilon$ component.

The right hand side of the isomorphism above is just $D_R(k[\epsilon])$, so we appear to have succeeded in our efforts to recover the tangent space from the functor D_R. But there's a problem: we recovered t_R only as a set; the k-vector space structure on $D_R(k[\epsilon])$ required getting our hands dirty with pointwise operations on our algebra homomorphisms. We would like to recover the tangent space as a k-vector space given nothing but the functor $D = D_R$, in order to study situations where representability is not known.

Let's see what we can do to fix this. Scalar multiplication can be recovered without too much difficulty. If we observe that $k \cong \mathrm{End}_{k-\mathrm{alg}}(k[\epsilon])$ via $a \leftrightarrow (x + y\epsilon \mapsto x + ay\epsilon)$, then the functoriality of D_R allows us to define scalar multiplication on $\mathrm{Hom}_{\hat{C}_\Lambda}(R, k[\epsilon])$ by the composition

$$R \longrightarrow k[\epsilon] \xrightarrow{a} k[\epsilon].$$

Let's try something similar for addition. We use the homomorphism of k-algebras

$$k[\epsilon] \times_k k[\epsilon] \xrightarrow{+} k[\epsilon] \tag{3}$$

$$(x + y_1\epsilon, x + y_2\epsilon) \mapsto x + (y_1 + y_2)\epsilon. \tag{4}$$

Then we construct the following diagram:

$$\begin{array}{ccccc} D(k[\epsilon]) \times D(k[\epsilon]) & \xleftarrow{h} & D(k[\epsilon] \times_k k[\epsilon]) & \xrightarrow{D(+)} & D(k[\epsilon]) \\ \| & & & & \| \\ t_D \times t_D & & \cdots\cdots\cdots\cdots\cdots\cdots & \dashrightarrow & t_D \end{array}$$

Here h is the map induced by the two projections $k[\epsilon] \times_k k[\epsilon] \to k[\epsilon]$. This gives us a map $t_D \times t_D \to t_D$, *if* h is a bijection. This is not automatically true, so we state it as a hypothesis:

Definition 4. A functor $D : C_\Lambda \to \mathbf{Set}$ satisfies the *tangent space hypothesis* $(\mathbf{T}_k)$ if the map $h : D(k[\epsilon] \times_k k[\epsilon]) \to D(k[\epsilon]) \times D(k[\epsilon])$ is a bijection.

Remark: we are interested in deformation functors, for which $D(k)$ is automatically a singleton set. In this case, the right-hand side of $(\mathbf{T}_k)$ can be rewritten as a fibered product over $D(k)$, which makes the hypothesis a special case of preservation of fibered products.

Now is probably a good time to say that all the preceding discussion carries over essentially without change to the relative case. To be precise: given a functor $D : C_\Lambda(\mathcal{A}) \to \mathbf{Set}$, assuming the "tangent $\mathcal{A}$-module hypothesis" $(\mathbf{T}_\mathcal{A})$ (defined analogously to $(\mathbf{T}_k)$), we can endow the set $t_{D,\mathcal{A}} = D(\mathcal{A}[\epsilon])$ with the structure of an $\mathcal{A}$-module, which we define to be the *tangent $\mathcal{A}$-module* of D. One subtlety appears when constructing the fibered product $\mathcal{A}[\epsilon] \times_\mathcal{A} \mathcal{A}[\epsilon]$ when $\mathcal{A}$ is not artinian, but this is fine because the fibered product is still noetherian when both maps are surjective.

4 Representability criteria

Now we specialize to the absolute case, i.e. functors $D : C_\Lambda \to \mathbf{Set}$. We now know three things about pro-representable functors: they preserve fibered products, they evaluate to a singleton set on k, and they have finite-dimensional tangent spaces. It turns out that these conditions together are sufficient.

Theorem 1. *(Grothendieck) Let $D : C_\Lambda \to \mathbf{Set}$ be a covariant functor such that $D(k)$ is a singleton set. Then D is pro-representable if and only if it preserves fibered products and the Zariski tangent space $t_{D,k}$ is finite-dimensional.*

This is a nice, simple criterion, but unfortunately it's not very useful in practice. The reason is that there are too many possible fibered product diagrams to check. Schlessinger's theorem refines Grothendieck's by giving a more manageable collection of diagrams to check. Before stating it, we fix some notation. Given a functor D and a diagram

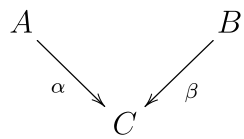

in C_Λ, we denote the canonical map $D(A \times_C B) \to D(A) \times_{D(C)} D(B)$ by h. In this language, the condition of preserving fibered products states precisely that h is a bijection for all such diagrams.

Definition 2. We call a map $A \to C$ in C_Λ *small* if its kernel is a principal ideal annihilated by m_A.

Theorem 3. *(Schlessinger) Let $D : C_\Lambda \to \mathbf{Set}$ be a covariant functor with $D(k) = *$. Then D is pro-representable if and only if the following four conditions hold:*

($\mathbf{H}_1$) h is surjective when $A \to C$ is small (or equivalently when $A \to C$ is surjective)

($\mathbf{H}_2$) h is bijective if $A \to C$ is $k[\epsilon] \to k$.

($\mathbf{H}_3$) Hypothesis ($\mathbf{T}_k$) holds and $t_{D,k}$ is finite-dimensional.

($\mathbf{H}_4$) h is bijective if $A \to C$ and $B \to C$ are equal and small.

There is some redundancy here: conditions ($\mathbf{H}_2$) and ($\mathbf{H}_4$) each imply ($\mathbf{T}_k$). This reflects how important the tangent space hypothesis is: if our goal is to prove modularity lifting theorems, Mazur says, ($\mathbf{H}_3$) is almost more important to us than actual representability. We will also need this in the relative case, so we introduce the terminology:

Definition 4. A covariant functor $D : C_\Lambda(\mathcal{A}) \to \mathbf{Set}$ is called *nearly representable* if $D(\mathcal{A}) = *$, and it satisfies the tangent $\mathcal{A}$-module hypothesis ($\mathbf{T}_\mathcal{A}$) and the finiteness hypothesis

$$(\mathbf{F}) : \text{The tangent } \mathcal{A}\text{-module } t_{D,\mathcal{A}} \text{ is of finite type.}$$

All absolute deformation functors we care about will turn out to satisfy ($\mathbf{H}_1$) – ($\mathbf{H}_3$). Such functors enjoy a property that is useful even when ($\mathbf{H}_4$) fails. We will not use this today (and therefore will not make it precise), but any functor D satisfying ($\mathbf{H}_1$) – ($\mathbf{H}_3$) admits a so-called *pro-representable hull*, i.e. a "smooth" morphism of functors $D_R \to D$ that induces an isomorphism of tangent spaces.

5 Back to Galois deformations

We now turn our attention back to the Galois deformation problems that motivated us in the beginning. Let Π be a profinite group satisfying the p-finiteness condition Φ_p. (Recall that Π satisfies Φ_p if each finite-index open subgroup $\Pi_0 \leq \Pi$ admits only finitely many continuous homomorphisms to $\mathbb{Z}/p\mathbb{Z}$. The Galois groups $G_{K,S}$ satisfy this when K is a number field and S is a finite set of primes.) Fix a coefficient ring Λ, a continuous residual representation $\overline{\rho} : \Pi \to \mathrm{GL}_N(k)$, and a continuous lifting $\rho : \Pi \to \mathrm{GL}_N(\mathcal{A})$ to a coefficient Λ-algebra $\mathcal{A}$ (not up to strict equivalence). Then we have the absolute deformation functor $D_{\overline{\rho}} : C_\Lambda \to \mathbf{Set}$ and the relative deformation functor $D_\rho : C_\Lambda(\mathcal{A}) \to \mathbf{Set}$. (Recall that they are both defined on the completed categories, but they are determined by their restrictions to the artinian subcategories.)

We now record which among our zoo of properties are satisfied by $D_{\overline{\rho}}$ and D_ρ.

Proposition 1. *(Representability for $\overline{\rho}$) With the above setup, we have:*

(i) The functor $D_{\overline{\rho}}$ (restricted to C_Λ) satisfies ($\mathbf{H}_1$), ($\mathbf{H}_2$), and ($\mathbf{H}_3$).

(ii) If $\overline{\rho}$ is absolutely irreducible, then $D_{\overline{\rho}}$ is representable.

(Unsurprisingly, part (ii) is proved using Schlessinger's theorem. Absolute irreducibility is used through Schur's lemma.)

Proposition 2. *D_ρ is nearly representable.*

Mazur also states a weak version of this near representability result, in which $\mathcal{A}$ is assumed to be artinian. This is easier to prove, and it suffices for our purposes. The ($\mathbf{T}_{\mathcal{A}}$) part of the assertions is not too difficult to prove; the finiteness part passes through some group cohomology computations, which may or may not be discussed next week.

Finally, we state a connection between representability for $\overline{\rho}$ and representability for ρ. The proof takes place entirely in the world of representations and strict equivalence classes thereof.

Proposition 3. *(Absolute $\to$ relative representability) Suppose $\overline{\rho}$ is absolutely irreducible and R represents the functor $D_{\overline{\rho}}$. If $\mathcal{A}$ is generated as a Λ-algebra by $\operatorname{tr}\rho$, then D_ρ is pro-representable by the same ring R, with augmentation $R \to \mathcal{A}$ induced by ρ.*

Remark: if A is generated as a Λ-algebra by $\operatorname{tr}\rho$ as in the proposition, we call A *minimal*.

6 Deformation functors with conditions

We now give a brief preview of the next talk, in which we will look at deformation problems with conditions—e.g. having a prescribed determinant or some local conditions. The moral of the story is that if D is a representable deformation problem and $D' \subset D$ is a "reasonable" subfunctor, then D' will be represented by a closed sub-(formal) scheme.

Definition 1. Suppose $D : C_\Lambda \to \textbf{Set}$ is a functor, and D' is a subfunctor; i.e. $D'(A) \subseteq D(A)$ for all A, and the inclusion $D' \hookrightarrow D$ is a natural transformation. Suppose moreover that $D'(k) = D(k) = *$. We say that $D' \subset D$ is a *relatively representable subfunctor* if for all fibered diagrams

$$\begin{array}{ccc} A & & B \\ & {\scriptstyle\alpha}\searrow \quad \swarrow{\scriptstyle\beta} & \\ & C & \end{array}$$

in C_Λ, the diagram

$$\begin{array}{ccc} D'(A \times_C B) & \xrightarrow{h} & D'(A) \times_{D'(C)} D'(B) \\ \Big\downarrow & & \Big\downarrow \\ D(A \times_C B) & \xrightarrow{h} & D(A) \times_{D(C)} D(B) \end{array}$$

is cartesian.

Proposition 2. *Suppose $D' \subset D$ is a relatively representable subfunctor. Then every condition we care about on D implies the corresponding condition on D': each $(\mathbf{H}_i)$, $(\mathbf{T}_k)$, near representability, representability. Moreover, if D is representable by R_D, then D' is representable by a quotient Λ-algebra of R_D.*

Next week, we will bring this back to the world of Galois deformation theory by defining a reasonable class of "deformation conditions" that will automatically give rise to relatively representable subfunctors.