A GLIMPSE OF DEFORMATION THEORY

BRIAN OSSERMAN

The purpose of these notes, as suggested by the title, is not to provide any sort of comprehensive introduction to deformation theory. Rather, we attempt to convey the main ideas of the theory, with a survey of some applications. We do not explore an exhaustive list of possible topics, nor do we go into details in many proofs. However, in the case of deformations of smooth varieties we do attempt to give a thorough treatment of the theory of first-order infinitesimal deformations, with some hints as to how to generalize to other problems.

Because the systematic use of rings with nilpotents is one of the major distinguishing characteristics of scheme theory, we can view deformation theory as a substantial application of scheme theory that is beyond the reach of classical algrebraic geometry. We hope to provide a relatively accessible and motivated introduction to the theory of cohomology of sheaves, in the form of the Cech cohomology arising in deformations of smooth varieties.

1. The many uses of Artin rings

Recall that A is an **Artin ring** if it is Noetherian of dimension 0. For the sake of brevity, we will in these notes also assume that all Artin rings are local. Recall that the maximal ideal of such a ring consists entirely of nilpotents.

Although $\operatorname{Spec} A$ for A an Artin ring consists of only a single point, the basic idea motivating deformation theory is that such schemes are still "big enough" to provide useful information. We have already seen Artin rings arising in two different contexts: in tangent spaces, and in criteria for smoothness. In the first case, we had that the tangent space of a scheme X at a point x was equal to the set of maps $\operatorname{Spec} k[\epsilon]/\epsilon^2 \to X$ with image x, where $k = k(x)$. A special case of the smoothness criterion was that if X is of finite type over a field k, then X is smooth over k if and only if for every Artin k-algebra A, and every quotient $A \twoheadrightarrow A'$ with kernel I a square-zero ideal, every map $\operatorname{Spec} A' \to X$ over k can be lifted to a map $\operatorname{Spec} A \to X$. Both of these involve understanding maps $\operatorname{Spec} A \to X$ for A an Artin ring.

Now suppose that X is a moduli space representing some moduli functor, or more generally a non-representable moduli functor $\tilde{F}$. Then the maps $\operatorname{Spec} A \to X$ (more generally, $\tilde{F}(\operatorname{Spec} A)$, which we abbreviate by $\tilde{F}(A)$) correspond to families of objects over $\operatorname{Spec} A$. If we fix a map $\operatorname{Spec} k \to X$ (or an element $\eta_0 \in \tilde{F}(k)$), and restrict our attention to A with residue field k, we can consider elements of $\tilde{F}(A)$ restricting to η_0 under the natural map $\operatorname{Spec} k \hookrightarrow \operatorname{Spec} A$, and this corresponds to studying families of objects over $\operatorname{Spec} A$ which restrict to the fixed object η_0 on the reduced point $\operatorname{Spec} k$; such families are called (infinitesimal) **deformations** of η_0 over $\operatorname{Spec} A$; this motivates the terminology.

One of the most basic ideas of deformation theory is that we can study the tangent space and smoothness of a moduli space by considering families over $\operatorname{Spec} A$.

Although completely classifying families of objects over arbitrary schemes T is usually far too complicated, in the case of Artin rings the bases are generally small enough that one can write down explicit descriptions.

2. (Pre)deformation functors

The standard setup (introduced by Schlessinger in [4]) is as follows: we denote the category of Artin rings with a given residue field k by $\mathrm{Art}(k)$, and consider (covariant) functors $F : \mathrm{Art}(k) \to \mathrm{Set}$ with the property that $F(k)$ is a one-point set. The idea is that such functors should come from deformations over Artin rings of a fixed object over k. We will call such functors **predeformation functors**.

Example 2.1. In the situation that we have a moduli functor $\tilde{F}$, we can obtain a predeformation functor F by fixing an element $\eta_0 \in \tilde{F}(k)$, and defining $F(A) = \{\eta \in \tilde{F}(A) : \eta|_k = \eta_0\}$.

In the case that $\tilde{F}$ is representable, the above will give a particularly well-behaved predeformation functor, but in general, to go between a global moduli problem and an associated deformation functor, we will want to do something slightly different, which cannot be expressed as well on the level of functors (however, since plunging into stacks would be a bit much, we settle for an example to give the general idea).

Example 2.2. Let X_0 be a variety over k. We define a functor F, the functor of deformations of X_0, as follows: $F(A)$ is the set of schemes X, flat over $\operatorname{Spec} A$, together with a morphism $X_0 \to X$ such that the diagram

$$\begin{array}{ccc} X_0 & \longrightarrow & X \\ \downarrow & & \downarrow \\ \operatorname{Spec} k & \longrightarrow & \operatorname{Spec} A \end{array}$$

commutes, and such that we have an induced isomorphism $X_0 \xrightarrow{\sim} X \times_A k$.

Note that in this case the morphism $X_0 \to X$ is a homeomorphism on underlying topological spaces, so that X differs from X_0 only on the level of structure sheaves.

Definition 2.3. Let F be a predeformation functor. The **tangent space** T_F is defined to be $F(k[\epsilon]/\epsilon^2)$. We say that F is **formally smooth** if, for all $A' \twoheadrightarrow A$ in $\mathrm{Art}(k)$ with kernel a square-zero ideal, and every element $\eta \in F(A)$, there exists an element $\eta' \in F(A')$ with $\eta'|_A = \eta$.

Note that if F happens to be obtained from a functor $\tilde{F}$ represented by some X as in the first example above, then T_F is in fact the tangent space of X at the chosen point, and X is formally smooth at the chosen point if and only if F is formally smooth.

As suggested above, we will focus on the study of tangent spaces, with some discussion of obstructions, which measure the failure of a predeformation functor to be smooth. A different question, treated systematically in [4], has to do with with how close a predeformation functor F is to being representable in a suitable sense. However, we will not pursue this direction here.

Remark 2.4. In older references such as [4], the phrase "deformation functor" is frequently reserved for the functor of deformations of schemes, with the term "functor of Artin rings" used more generally. However, there are many different types of problems now considered to fall under the rubric of deformation theory, so we have adopted the above notation.

We also mention that, as suggested by the terminology, we have in mind a notion more specific than a predeformation functor, which we shall call a "deformation functor". We defer the definition until such time as we discuss Schlessinger's work in more detail, but mention that the conditions on a deformation functor will imply that we can put the natural structure of a k-vector space onto its tangent space, justifying the terminology.

3. A survey of applications

Taking for the moment for granted the idea that deformation theory provides a powerful tool for analyzing moduli problems, and in particular for studying tangent spaces and smoothness, we survey a number of applications of deformation theory to classical questions.

Applications to moduli spaces. Of course, the most direct applications of deformation theory are to the study of moduli spaces. If we have a predeformation functor F obtained from a point x of a moduli space X, then as mentioned before, the notions of tangent space and formal smoothness for X and for F coincide. Noting that if a variety is smooth, its dimension is the dimension of its tangent space, one could prove in this way that the Grassmannian $G(r, d)$ is smooth of dimension $r(d-r)$, although of course our proof of representability gave a more direct and stronger statement in showing that it is covered by copies of $\mathbb{A}_k^{r(d-r)}$. However, later we will study other fine moduli spaces, such as Hilbert spaces, where in many cases deformation theory is the only way of obtaining such statements.

Along these lines, one of the archetypical examples is the moduli space $\mathcal{M}_g$ of curves of genus g. This parametrizes flat families of smooth, proper, geometrically connected curves of genus g. Although it is not representable, it is nonetheless rather well-behaved for $g \geq 2$; technically, it is a "Deligne-Mumford stack", but from our point of view, it suffices to take for granted that the notions of smoothness and dimension can be defined independently, and we can then state the following theorem, proved via deformation theory:

Theorem 3.1. *For $g \geq 2$, the moduli space $\mathcal{M}_g$ is smooth of dimension $3g-3$.*

For the sake of giving a corollary which we can actually state precisely, we note that $\mathcal{M}_g$ also has a coarse moduli space M_g; although the notion of a coarse moduli space has additional hypotheses, it is uniquely defined by the property of corepresenting the moduli functor in question; see the notes on Yoneda's lemma and representable functors.

Theorem 3.2. *For $g \geq 2$, the functor associated to $\mathcal{M}_g$ is corepresentable by a variety M_g.*

We then have the following corollary of the earlier theorem:

Corollary 3.3. *For $g \geq 2$, the coarse moduli space $\mathcal{M}_g$ has dimension $3g-3$.*

Remark 3.4. Note that M_g will not be smooth; in fact, it will have singularities at points corresponding to curves with non-trivial automorphisms. This is a typical situation, and illustrative of the advantages of working with stacks rather than coarse moduli spaces.

However, deformation theory is an important tool for a variety of arguments, so we survey applications beyond moduli spaces.

Finiteness arguments. One very natural application of deformation theory is in proving finiteness results. If one wants to prove that a certain set of objects is finite, one can break the problem into two steps: first show that they are parametrized by a moduli space of finite type, and then show that none of the objects have non-trivial first-order deformations. We can then conclude that the moduli space consists of a collection of disjoint points, which is necessarily finite. While the first step is typically independent, deformation theory is often a crucial ingredient in the second step. We mention here a typical example, a famous conjecture of Shafarevich, in a case proved by Parshin using the above approach.

Definition 3.5. We say that a morphism $X \to B$ is **isotrivial** if for general distinct points $b, b' \in B$, we have an isomorphism of fibers $X_b \cong X_{b'}$.

Theorem 3.6. *(Shafarevich, Parshin) Let B be a smooth, proper curve over a field k, and fix an integer $g \geq 2$. Then there exist only finitely many non-isotrivial families of curves $X \to B$ which are smooth and proper, with fibers being curves of genus g.*

This says in essence that for a fixed $g \geq 2$, and a smooth, proper curve B of genus g, over a field k, there are at most finitely many non-constant families of curves of genus g over B, which we can think of as saying that "there are only finitely many non-constant maps from B to $\mathcal{M}_g$". We note that in fact the statement is more general, with the same statement holding even when B is not proper. However, although the proof of the general statement also uses deformation theory, it is substantially subtler, and was proved by Arakelov.

Constructing families. One important application of deformation theory involves putting chosen varieties into well-behaved families of one sort or another. Given a variety X_0 over k, we might wish to put X_0 into a family X over some base B, such that one fiber is isomorphic to X_0, and X has some good properties. A typical example is the following theorem:

Theorem 3.7. *(Winters) Let X_0 be a proper, geometrically integral curve over a field k, with at worst nodal singularities. Then there exists a base scheme B with point $b \in B$, and a scheme X over B, such that:*

(i) *$k(b) = k$, and $X_b \cong X_0$;*
(ii) *B is regular of dimension 1, and can be chosen to be a curve over k, or, if k has positive characteristic, we can set $B = \operatorname{Spec} A$, where A is a mixed-characteristic DVR with residue field k;*
(iii) *X is a regular surface, and its generic fiber over B is smooth.*

Such theorems have two major types of applications.

The first involves the case that X_0 is in fact smooth, but k has positive characteristic (the theorem in this case preceded Winters' work). Here, by choosing B to have mixed-characteristic (i.e., its generic point is a field of characteristic 0), we can realize X_0 as a "specialization" of a curve in characteristic 0. Such techniques are often essential to carrying out computations in characteristic p, as in the case of (the algebraic notion of) the fundamental group of a curve of genus $g \geq 2$, where the only known way to calculate the fundamental group involves using the classical topological formula for curves over $\mathbb{C}$, applying this to curves in characteristic 0 more generally, and realizing the given curve in characteristic p as the specialization of a curve in characteristic 0 as above.

While the above application involved using knowledge of the generic fiber to prove results about X_0, one can also work in the other direction, typically by choosing X_0 to be singular. This is the essence of degeneration arguments. A typical situation might be as follows: suppose we wish to prove something about curves of genus g, and it turns out to be enough to prove the statement for a single curve of genus g (this is the case, for instance, with the Brill-Noether theorem). One might take smooth curves of genus 1 and $g-1$, and glue them together at a single node to form X_0. In this case, the smooth generic fiber provided by the above theorem will have genus g, and one can hope to prove the desired statement for the generic fiber by making sense of it for X_0, and understanding the situation for X_0 in terms of the two smooth components. Since these components each have genus less than g (assuming $g > 1$), this potentially sets up an induction with base case $g = 1$. But to even get started, one has to know that the chosen X_0 can be put into a family with smooth generic fiber, and that is where the above theorem comes in.

We very briefly describe the proof of Winter's theorem, as it is typical of a certain class of results. The first step is to show that one can construct compatible systems of families over larger and larger Artin rings, for instance, over $k[t]/t^n$ for each n. This is where deformation theory is used. One then makes an argument (usually fairly straightforward, but not completely automatic) that one can construct a deformation over the "formal scheme" (see §II.9 of [2]) corresponding to the limit A of the given Artin rings (in the case above, $A = k[[t]]$). Next, one applies a theorem of Grothendieck to "effective" the deformation, which means to realize it over the standard scheme $\operatorname{Spec} A$. Here, some real conditions arise, and there are many examples of deformations which arise in nature and can be constructed over formal schemes, but not effectivized. Finally, one applies a theorem of Artin to "algebraize" the deformation, approximating it for instance over a curve of finite type over k.

Lifting Galois representations. One intriguing application of deformation theory has been Mazur's theory of deformations of Galois representations, and its highly successful application to the proof of the Shimura-Taniyama-Weil conjecture on modularity of elliptic curves, and further progress in recent years. Although there is no geometry involved, the theory does fall neatly into the setting of predeformation functors.

It had been known for a long time that the following problem is important:

Question 3.8. Given a Galois representation $\rho : G_{\bar{\mathbb{Q}}/\mathbb{Q}} \to GL_2(\mathbb{F}_p)$, what are the possible lifts to representations $\hat{\rho} : G_{\bar{\mathbb{Q}}/\mathbb{Q}} \to GL_2(\mathbb{Z}_p)$?

The idea is that in order to answer this question, one lifts ρ successively to $\mathbb{Z}/p^n\mathbb{Z}$ for increasing n, and then forms the inverse limit to obtain a lift to $\mathbb{Z}_p$. The Artin rings in question are the $\mathbb{Z}/p^n\mathbb{Z}$, and the possible lifts for each n define the predeformation functor. In [3], Mazur considered this problem in the context of Schlessinger's theory, showing that the resulting deformation theory is well-behaved.

This formed an integral part of Wiles' technique and subsequent work on Galois representations.

4. Sketch of an argument

We discuss one more application of deformation theory to a classical question, with a somewhat more detailed description of how deformation theory can be used to prove the stated result. We consider the following special case of the Brill-Noether problem:

Question 4.1. Given $g \geq 0$, for which $d > 0$ is it the case that every curve of genus g has a non-constant map to $\mathbb{P}^1$ of degree at most d?

For $g = 0$, the only curve is $\mathbb{P}^1$ itself, so the answer is that any d will do. For $g = 1$, we cannot have a map of degree 1 to $\mathbb{P}^1$, as such a map would have to be an isomorphism. But we always have a map of degree 2, so any $d \geq 2$ is okay. The general statement is the following:

Theorem 4.2. *The answer to the above question is: any d for which $2d-2-g \geq 0$.*

The proof of this theorem is in two parts: an existence statement when $2d-2-g \geq 0$, and a non-existence statement when $2d-2-g < 0$. We will now sketch a simple proof based on deformation theory for the non-existence statement.

We can immediately check the desired statement in the cases $g = 0$ or $g = 1$: for $g = 0$, the assertion is only that $d \geq 1$, which is vacuous, while for $g = 1$, the assertion is that $d \geq 2$, which is necessary because a map of degree 1 between smooth proper curves is necessarily an isomorphism, and if $g > 0$ then C is not isomorphic to $\mathbb{P}^1$. We therefore assume $g \geq 2$.

The basic idea is to consider the following sequence of moduli spaces:

$$\mathscr{X} \to \mathscr{Y} \to \mathcal{M}_g,$$

where:

- $\mathcal{M}_g$ is the moduli space parametrizing curves C of genus g;
- $\mathscr{X}$ is the moduli space of pairs (C, f) of curves C of genus g together with a map $f : C \to \mathbb{P}^1$ of degree d;
- $\mathscr{Y}$ is the moduli space of pairs (C, f) as for $\mathscr{X}$, except that we consider $(C, f) \sim (C, f')$ if f and f' are related by an automorphism of $\mathbb{P}^1$;

The argument then works as follows: by definition, a curve C with a map to $\mathbb{P}^1$ of degree d is precisely a point of $\mathcal{M}_g$ which has (at least) a point of $\mathscr{Y}$ mapping to it. It follows that in order for every curve of genus g to have a map of degree at most d to $\mathbb{P}^1$, we have to have the map $\mathscr{Y} \to \mathcal{M}_g$ be dominant for some degree $d' \leq d$; in particular, we have to have $\dim \mathscr{Y} \geq \dim \mathcal{M}_g = 3g - 3$, with the last equality coming from the earlier theorem. The reason for introducing the space $\mathscr{X}$ is that it turns out that the dimension of $\mathscr{X}$ is easiest to compute directly. Indeed, we can do this via deformation theory, and we find that with $g \geq 2$, we have $\dim \mathscr{X} = 2d + 2g - 2$. Then, the automorphism group of $\mathbb{P}^1$ is 3-dimensional (indeed, it consists of all invertible maps of the form $z \mapsto \frac{az+b}{cz+d}$, with simultaneous scaling producing the same map; see Exercise I.6.6 or Example II.7.1.1 of [2]), so we have $\dim \mathscr{Y} = \dim \mathscr{X} - 3 = 2d + 2g - 5$. We thus have $\dim \mathscr{Y} < \dim \mathcal{M}_g$ if and only if $2d - 2 - g < 0$, proving the desired non-existence statement.

Warning 4.3. The above doesn't quite make sense, as written. One can write down a correct argument entirely in terms of schemes by replacing $\mathcal{M}_g$ with the base for a "modular family" of curves of genus g, which is to say, an étale cover of $\mathcal{M}_g$. The intuition here is that a modular family is like the universal family over $\mathcal{M}_g$, but instead of each curve appearing once, each is allowed to appear finitely many times. The other spaces should then be interpreted in terms of the curves in the modular family, and one can make the argument described below in a precise manner.

5. Deformations of smooth varieties and Cech cohomology

In Example 2.2 above, we describe the problem of deforming a variety X_0 over k. In general, it is not easy to describe such deformations, but in the case that X_0 is smooth over k, we have all the tools to at least describe the tangent space of the problem – that is, all deformations of X_0 over $\operatorname{Spec} k[\epsilon]/\epsilon^2$.

The first two steps of the analysis are the following results:

Proposition 5.1. *If X_0 is a smooth affine variety over k, every first-order deformation of X_0 is isomorphic to the trivial deformation $X_0[\epsilon]$.*

Indeed, we will see, with the help of a lemma on flatness, that this is equivalent to Exercise II.8.7 of [2], which was on Problem Set 12, Part 1.

Proposition 5.2. *Let X_0 be a smooth variety over k. Then the sheaf of infinitesimal automorphisms of X_0 is the tangent sheaf T_{X_0}.*

Here, recall that the tangent sheaf is defined for X_0 smooth by $\mathcal{H}om(\Omega^1_{X/k}, \mathcal{O}_X)$. If, as above, $X_0[\epsilon]$ denotes the trivial deformation of X_0 over $\operatorname{Spec} k[\epsilon]/\epsilon^2$, **infinitesimal automorphisms** of X_0 are automorphisms of $X_0[\epsilon]$ over $\operatorname{Spec} k[\epsilon]/\epsilon^2$ which restrict to the identity on X_0. The sheaf of infinitesimal automorphisms of X_0 is the sheaf associating to an open subset $U \subseteq X_0$ the infinitesimal automorphisms of U.

Assuming for the moment these two propositions, if we let X_0 be any smooth variety over k, we can analyze the first-order deformations of X_0 as follows: suppose X_1 is a first-order deformation of X_0; by Proposition 5.1, if $\{U_i\}$ is any affine open cover of X_0, we have that $X_1|_{U_i}$ is isomorphic to the trivial deformation for each i. Moreover, if we choose trivializations $\varphi_i : X_1|_{U_i} \xrightarrow{\sim} U_i[\epsilon]$ for each i, then for any $i < j$, we have a gluing map $\varphi_{i,j} : U_{i,j}[\epsilon] \xrightarrow{\sim} U_{i,j}[\epsilon]$ obtained from $\varphi_i|_{U_{i,j}}$ and $(\varphi_j|_{U_{i,j}})^{-1}$, where $U_{i,j} := U_i \cap U_j$. We note that because the data of the deformation X_1 includes the map $X_0 \hookrightarrow X_1$, each $\varphi_{i,j}$ gives the identity when restricted to $U_{i,j}$, so is an infinitesimal automorphism of $U_{i,j}$, and hence a section of $T_{X_0}(U_{i,j})$, by Proposition 5.2. We also note that because the $\varphi_{i,j}$ come from the gluing data for schemes, we have the **cocycle condition** $\varphi_{i,j} \circ \varphi_{j,k} = \varphi_{i,k}$ for all $i < j < k$ (see Exercise II.2.12 of [2]). Also, we have $\varphi_{i,j} = \varphi_{j,i}^{-1}$ for any $i < j$.

Conversely, it is clear that given the data of sections $\varphi_{i,j} \in T_{X_0}(U_{i,j})$ for all $i < j$, which satisfy the cocycle condition for all $i < j < k$, and have $\varphi_{i,j} = -\varphi_{j,i}$ for all $i < j$ (switching to additive notation, as we consider the $\varphi_{i,j}$ as sections of the sheaf T_{X_0}), we can glue to obtain an X_1 which will be a first-order deformation of X_0. Note here that flatness over $\operatorname{Spec} k[\epsilon]/\epsilon^2$ can be checked locally, and since locally we are starting with the trivial deformation, our modules are visibly free over $k[\epsilon]/\epsilon^2$, hence flat.

It remains to analyze the ambiguity of our description, or equivalently, when two different choices of $\varphi_{i,j}$ give the same deformation (up to isomorphism). But this is clear: our original choice of the φ_i could be modified precisely by an infinitesimal automorphism of U_i, which would then modify each $\varphi_{i,j}$ accordingly, for any j with $j > i$. We can formalize this as follows: we have the group of **Cech 1-cochains** for $\{U_i\}$ and T_{X_0}, which is simply the group $\prod_{i<j} T_{X_0}(U_{i,j})$. We say that a Cech 1-cochain is a **cocycle** if (naturally enough) it satisfies the cocycle condition (where composition is replaced by addition in $T_{X_0}(U_{i,j,k})$). Finally, we have a coboundary map from $\prod_i T_{X_0}(U_i)$ to the group of Cech 1-cocycles obtained by taking the section on $U_{i,j}$ to be the difference of the sections on U_j and U_i. We define the **first Cech cohomology group** on $\{U_i\}$ of T_{X_0}, written $\check{H}^1(\{U_i\}, T_{X_0})$ to be the group of Cech 1-cocycles modulo the image of the coboundary map. We therefore conclude that our two propositions above imply the following theorem:

Theorem 5.3. *Let X_0 be a smooth variety over k, and $\{U_i\}$ any open affine cover of X_0. Then the first-order deformations of X_0 are parametrized by the cohomology group $\check{H}^1(\{U_i\}, T_{X_0})$.*

It remains to proves Propositions 5.1 and 5.2. We need a key commutative algebra lemma on flatness for the first:

Lemma 5.4. *A module M over $k[\epsilon]/\epsilon^2$ is flat if and only if the natural map $M/\epsilon M \xrightarrow{\times\epsilon} \epsilon M$ is an isomorphism.*

See Corollary 6.2, p. 123 of [1]. Note that the map is always visibly surjective, so flatness is equivalent to injectivity.

Proof of Proposition 5.1. We first assert that first-order deformations of X_0 are equivalent to extensions $(X', \mathcal{O}_{X'})$ of X_0 by $\mathcal{O}_{X_0}$ (see Exercise II.8.7 of [2]), with the additional data of a k-vector space structure on $\mathcal{O}_{X'}$, a fixed map $i^\sharp : \mathcal{O}_{X'} \twoheadrightarrow \mathcal{O}_{X_0}$, and, if $\mathscr{I}$ is the kernel of $i^\sharp$, a fixed isomorphism $\pi : \mathscr{I} \xrightarrow{\sim} \mathcal{O}_{X_0}$. The main content to be checked comes down to the above lemma on flatness, but we go through the details below.

Indeed, if $(X', \mathcal{O}_{X'})$ is a first-order deformation of X_0, we are given by hypothesis a map $i^\sharp : \mathcal{O}_{X'} \to \mathcal{O}_{X_0}$, as well as the structure of a sheaf of $k[\epsilon]/\epsilon^2$-algebras on $\mathcal{O}_{X'}$. Moreover, the hypothesis that our map $i : X_0 \hookrightarrow X'$ induces an isomorphism $X_0 \xrightarrow{\sim} X' \times_{\operatorname{Spec} k[\epsilon]/\epsilon^2} \operatorname{Spec} k$ is equivalent to saying $i^\sharp$ induces an isomorphism $\mathcal{O}_{X'}/\epsilon\mathcal{O}_{X'} \xrightarrow{\sim} \mathcal{O}_{X_0}$; in particular, we have $\ker i^\sharp = \epsilon\mathcal{O}_{X'}$. By flatness and the above lemma, we have that the map $\mathcal{O}_{X'}/\epsilon\mathcal{O}_{X'} \xrightarrow{\times\epsilon} \epsilon\mathcal{O}_{X'}$ is an isomorphism, which is to say that the $k[\epsilon]/\epsilon^2$-algebra structure gives us the isomorphism π as well.

Conversely, if we are given the extension $(X', \mathcal{O}_{X'})$ together with $i^\sharp$ and π, we have that $i^\sharp$ induces a map $X_0 \twoheadrightarrow X'$, while we claim that $\mathcal{O}_{X'}$ has the natural structure of a sheaf of $k[\epsilon]/\epsilon^2$-algebras: by hypothesis, it has the structure of a sheaf of k-vector spaces, so it suffices to define the multiplication-by-ϵ map. We define this by $i^\sharp \circ \pi^{-1}$ (so that the kernel is given by $\mathscr{I}$). We then obtain that $X_0 \cong X' \times_{\operatorname{Spec} k[\epsilon]/\epsilon^2} \operatorname{Spec} k$ from the condition in the definition of extension that $i^\sharp$ induces an isomorphism $\mathcal{O}_{X'}/\mathscr{I} \xrightarrow{\sim} \mathcal{O}_{X_0}$. It remains to conclude flatness of $(X', \mathcal{O}_{X'})$ over $\operatorname{Spec} k[\epsilon]/\epsilon^2$, and this follows from the isomorphisms π and $\mathcal{O}_{X'}/\mathscr{I} \xrightarrow{\sim} \mathcal{O}_{X_0}$, by the above lemma.

We thus conclude that first-order deformations of X_0 do correspond to extensions of X_0 by $\mathcal{O}_{X_0}$ with fixed choices of k-vector space structure, and maps $i^\sharp$ and π.

By Exercise II.8.7 of [2], every such extension is trivial; in fact, the same argument shows that every extension together with the k-algebra structure and maps $i^\sharp$ and π is isomorphic to the trivial one, so we conclude the desired result. □

Before proving the second proposition, we remark that it is a formalization of an old idea from differential geometry: that one should think of a global section of the tangent bundle – a vector field – as being an infinitesimal version of an automorphism, as it gives an infinitesimal direction for each point to flow along.

Proof of Proposition 5.2. The main idea is that infinitesimal automorphisms of X_0 correspond to k-linear derivations $\mathcal{O}_{X_0} \to \mathcal{O}_{X_0}$; the desired statement will then follow from the universal property of $\Omega^1_{X_0/k}$ with regard to derivations.

However, an infinitesimal automorphism of X_0 is nothing but a $k[\epsilon]/\epsilon^2$-linear isomorphism $\mathcal{O}_{X_0}[\epsilon] \xrightarrow{\sim} \mathcal{O}_{X_0}[\epsilon]$ restricting to the identity modulo ϵ; it follows that on affine opens U, it is of the form $a + \epsilon b \mapsto a + \epsilon(b + d(a))$, for any $a, b \in \mathcal{O}_{X_0}(U)$, where we note:

- $\epsilon b \mapsto \epsilon b$ because $b \mapsto b + \epsilon d(b)$ for some $d(b)$, and the map is $k[\epsilon]/\epsilon^2$-linear;
- d is some k-linear map $\mathcal{O}_{X_0} \to \mathcal{O}_{X_0}$.

But the multiplication structure tells us that $ab \mapsto ab + \epsilon d(ab) = (a + \epsilon d(a))(b + \epsilon d(b)) = ab + \epsilon(d(a)b + a(d(b)))$, so we see that d must be a k-linear derivation, as desired.

Finally, we had on affine opens by definition that maps $\Omega^1_{X_0/k} \to \mathcal{O}_{X_0}$, which is to say, sections of T_{X_0}, correspond to k-linear derivations $\mathcal{O}_{X_0} \to \mathcal{O}_{X_0}$. However, since maps of sheaves are defined on any open cover, it is easy to see that this universal property passes to the level of sheaves, completing the proof of the proposition. □

Remark 5.5. Although the details are beyond the scope of these notes, we briefly explain the notions of higher Cech cohomology and of obstruction theory, and sketch the fact that obstructions to deforming a smooth variety X_0 lie in $\check{H}^2(\{U_i\}, T_{X_0})$.

The idea is that it is useful to know when a deformation over $\operatorname{Spec} A$ lifts to a deformation over $\operatorname{Spec} A'$, given a map $A' \twoheadrightarrow A$ in $\operatorname{Art}(k)$. If deformations always lift, we have by definition that our predeformation functor is formally smooth. One typically approaches this question by factoring the given map into maps with smaller kernel; one can always then require that the kernel is a principle, square-zero ideal, which is then necessarily isomorphic to k. We call such a map a **tiny extension**. It turns out that there is often a beautiful theory of obstructions to lifting a deformation over a tiny extension. Roughly speaking, what happens is that there exists a k-vector space V, called the **obstruction space**, and if we are given any deformation η_A over $\operatorname{Spec} A$, and the tiny extension $A' \twoheadrightarrow A$, we obtain an element $v \in V$, called the **obstruction** to lifting η_A to A', with the property that $v = 0$ if and only if η_A can be lifted to a deformation over $\operatorname{Spec} A'$. We note that if we have an obstruction space V, and if we can prove that $V = 0$, it follows automatically that our deformation problem is formally smooth. The converse is however false, as if V is an obstruction space, then any V' containing V is likewise an obstruction space.

We now define the higher Cech cohomology groups $\check{H}^i(\{U_i\}, T_{X_0})$ and sketch the proof of the following:

Theorem 5.6. *For the deformations of a smooth variety X_0, we have that $\check{H}^2(\{U_i\}, T_{X_0})$ is an obstruction space.*

Given our cover $\{U_i\}$ of X_0, and a sheaf $\mathcal{F}$, we define a complex of **Cech cochains** by taking the ith term of the complex to be $\prod_{j_0<\cdots<j_i} \mathcal{F}(U_{j_0,\dots,j_i})$, where $U_{j_0,\dots,j_i}$ denotes $U_{j_0} \cap \cdots \cap U_{j_i}$. We define the complex maps in terms of alternating sums; the details are given at the beginning of §III.4 of [2]. We then define the **Cech** i**-cocycles** to be the kernel of the complex map at the ith place, and **Cech** i**-coboundaries** to be the image of the complex map. As usual, we set $\check{H}^i(\{U_i\}, \mathcal{F})$ to be the ith cohomology group of the complex, which is to say the i-cocycles modulo the i-coboundaries.

The first step in proving the theorem is generalizing Proposition 5.1 to deformations over $\operatorname{Spec} A$ for all Artin k-algebras A. This is not difficult, involving only inductive applications of the same argument used in the proof of the case $A = k[\epsilon]/\epsilon^2$. Next, we suppose we are given our deformation η_A over $\operatorname{Spec} A$, and our tiny extension $A' \twoheadrightarrow A$. We have that η_A is trivial on our cover $\{U_i\}$, so there exist deformations $\eta_{A',i}$ on each U_i lifting η_A, and in fact unique ones – the trivial deformations of each U_i to A'. These are of course isomorphic on each intersection $U_{i,j}$, so the question is whether we can choose isomorphisms satisfying the cocycle extension, and thereby glue to get a deformation $\eta_{A'}$. If I is the kernel of $A' \twoheadrightarrow A$, we find that our isomorphisms are determined modulo I, because we want $\eta_{A'}$ to be a lift of η_A. We fix some choice of isomorphisms $\varphi_{i,j}$ which restrict appropriately modulo I. Because $I \cong k$, if we choose an isomorphism, we will have that the failure of the $\varphi_{i,j}$ to satisfy the cocycle condition forms a Cech 2-cocycle for T_{X_0} on $\{U_i\}$. We also see that we are free to modify the $\varphi_{i,j}$ precisely by a Cech 2-coboundary for T_{X_0} on $\{U_i\}$. We thus have a well-defined element of $\check{H}^2(\{U_i\}, T_{X_0})$ measuring whether or not we can lift η_A to A', so we have the desired obstruction space.

This argument works to provide an obstruction theory if we restrict our attention to Artin rings A which are also k-algebras. A similar argument works more generally, but we will not discuss it here.

References

1. David Eisenbud, *Commutative algebra with a view toward algebraic geometry*, Springer-Verlag, 1995.
2. Robin Hartshorne, *Algebraic geometry*, Springer-Verlag, 1977.
3. Barry Mazur, *Deforming Galois representations*, Galois groups over Q, Springer, 1989, pp. 385–437.
4. Michael Schlessinger, *Functors of Artin rings*, Transactions of the AMS **130** (1968), 208–222.