DERIVED ALGEBRAIC GEOMETRY

1. Introduction

1.1. **Bezout's Theorem.** Let $C, C' \subseteq \mathbf{P}^2$ be two smooth algebraic curves of degrees n and m in the complex projective plane $\mathbf{P}^2$. If C and C' meet transversely, then the classical theorem of Bezout (see for example [10]) asserts that $C \cap C'$ has precisely nm points.

We may reformulate the above statement using the language of cohomology. The curves C and C' have fundamental classes $[C], [C'] \in \mathrm{H}^2(\mathbf{P}^2, \mathbf{Z})$. If C and C' meet transversely, then we have the formula

$$[C] \cup [C'] = [C \cap C'],$$

where the fundamental class $[C \cap C'] \in \mathrm{H}^4(\mathbf{P}^2, \mathbf{Z}) \simeq \mathbf{Z}$ of the intersection $C \cap C'$ simply counts the number of points in the intersection. Of course, this should not be surprising: the cup-product on cohomology classes is defined so as to encode the operation of intersection. However, it would be a mistake to regard the equation $[C] \cup [C'] = [C \cap C']$ as obvious, because it is not always true. For example, if the curves C and C' meet nontransversely (but still in a finite number of points), then we always have a strict inequality

$$[C] \cup [C'] > [C \cap C']$$

if the right hand side is again interpreted as counting the number of points in the set-theoretic intersection of C and C'.

If we want a formula which is valid for non-transverse intersections, then we must alter the definition of $[C \cap C']$ so that it reflects the appropriate intersection multiplicities. Determination of these intersection multiplicities requires knowledge of the intersection $C \cap C'$ as a scheme, rather than simply as a set. This is one of the classical arguments that nonreduced scheme structures carry useful information: the intersection number $[C] \cup [C'] \in \mathbf{Z}$, which is defined *a priori* by perturbing the curves so that they meet transversally, can also be computed directly (without perturbation) if one is willing to contemplate a potentially non-reduced scheme structure on the intersection.

In more complicated situations, the appropriate intersection multiplicities cannot always be determined from the scheme-theoretic intersection alone. Suppose that C and C' are (possibly singular) subvarieties of $\mathbf{P}^n$, of complementary dimension and having a zero-dimensional intersection. In this case, the appropriate intersection number associated to a point $p \in C \cap C'$ is not always given by the complex dimension of the local ring

$$\mathcal{O}_{C \cap C', p} = \mathcal{O}_{C,p} \otimes_{\mathcal{O}_{\mathbf{P}^n,p}} \mathcal{O}_{C',p} .$$

The reason for this is easy to understand from the point of view of homological algebra. Since the tensor product functor $\otimes_{\mathcal{O}_{\mathbf{P}^n,p}}$ is not exact, it does not have good properties when considered alone. According to Serre's intersection formula, the correct intersection multiplicity is instead the Euler characteristic

$$\sum (-1)^i \dim \mathrm{Tor}_i^{\mathcal{O}_{\mathbf{P}^n,p}}(\mathcal{O}_{C,p}, \mathcal{O}_{C',p}).$$

This Euler characteristic contains the dimension of the local ring of the scheme-theoretic intersection as its leading term, but also higher-order corrections. We refer the reader to [31] for further discussion of this formula for the intersection multiplicity.

If we would like the equation $[C] \cup [C'] = [C \cap C']$ to remain valid in the more complicated situations described above, then we will need to interpret the intersection $C \cap C'$ in some way which remembers not only the tensor product $\mathcal{O}_{C,p} \otimes_{\mathcal{O}_{\mathbf{P}^n,p}} \mathcal{O}_{C',p}$, but the higher Tor terms as well. Moreover, we should not interpret these Tor-groups separately, but rather should think of the total derived functor $\mathcal{O}_{C,p} \otimes^L_{\mathcal{O}_{\mathbf{P}^n,p}} \mathcal{O}_{C',p}$ as a kind of "generalized ring".

These considerations lead us naturally to the subject of *derived algebraic geometry*. Using an appropriate notion of "generalized ring", we will mimic the constructions of classical scheme theory to obtain a theory of *derived schemes* in which a version of the formula $[C] \cup [C'] = [C \cap C']$ can be shown to hold with (essentially) *no* hypotheses on C and C'. Here, we must interpret the intersection $C \cap C'$ in the sense of derived schemes, and we must take great care to give the proper definition for the fundamental classes (the so-called *virtual fundamental classes* of [4]).

To motivate our discussion of "generalized rings", we begin by considering the simplest case of Bezout's theorem, in which C and C' are lines in the projective plane $\mathbf{P}^2$. In this case, we know that $[C] \cup [C']$ is the cohomology class of a point, and that C intersects C' transversely in one point so long as C and C' are distinct. However, when the equality $C = C'$ holds, the scheme-theoretic intersection $C \cap C'$ does not even have the correct dimension.

Let us now try to give an idea of how we might formulate a definition for "derived scheme-theoretic intersections" which will handle the degenerate situation in which $C = C'$. For simplicity, let us consider only lines in the affine plane $\mathbf{A}^2 \subseteq \mathbf{P}^2$, with coordinate ring $\mathbf{C}[x, y]$. Two distinct lines in $\mathbf{A}^2$ may be given by equations $x = 0$ and $y = 0$. The scheme-theoretic intersection of these two lines is the spectrum of the ring $\mathbf{C}[x, y]/(x, y) \simeq \mathbf{C}$, obtained from $\mathbf{C}[x, y]$ by setting the equations of both lines equal to zero. It has dimension zero because $\mathbf{C}[x, y]$ is two-dimensional to begin with, and we have imposed a total of two equations.

Now suppose that instead of C and C' being two distinct lines, they are actually two *identical* lines, both of which have the equation $x = 0$. In this case, the affine ring of the scheme theoretic intersection is given by $\mathbf{C}[x, y]/(x, x) \simeq \mathbf{C}[y]$. This ring has dimension one, rather than the expected dimension zero, because the two equations are not independent: setting $x = 0$ twice is equivalent to setting $x = 0$ once. To obtain derived algebraic geometry, we need a formalism of "generalized rings" in which imposing the equation $x = 0$ twice is not equivalent to imposing the equation once.

One way to obtain such a formalism is by "categorifying" the notion of a commutative ring. That is, in place of ordinary commutative rings, we should consider *categories* equipped with "addition" and "multiplication" operations (which are now functors, rather than ordinary functions). For purposes of the present discussion, let us call such an object a *categorical ring*. We shall not give a precise axiomatization of this notion, which turns out to be quite complicated (see [19], for example).

Example 1.1.1. Let $\mathbf{Z}_{\geq 0}$ denote the semiring of nonnegative integers. We note that $\mathbf{Z}_{\geq 0}$ arises in nature through the process of "decategorification". The nonnegative integers were introduced in order to count finite collections: in other words, they correspond to isomorphism classes of objects in the category $\mathcal{Z}$ of finite sets. Then $\mathcal{Z}$ is naturally equipped with the structure of a *categorical semiring*, where the addition is given by forming disjoint unions and the multiplication is given by Cartesian products. (In order to complete the analogy with the above discussion, we should "complete" the category $\mathcal{Z}$ by formally adjoining inverses, to obtain a categorical ring rather than a categorical semiring, but we shall ignore this point.)

To simplify the discussion, we will consider only categorical rings which are groupoids: that is, every morphism in the underlying category is an isomorphism. If $\mathcal{C}$ is a categorical ring, then the collection of isomorphism classes of objects $\pi_0 \mathcal{C}$ of $\mathcal{C}$ forms an ordinary ring. Every commutative ring R arises in this way: for example, we may take $\mathcal{C}_R$ to be a category whose objects are the elements of R and which contains only identity maps for morphisms. The categorical rings which arise in this way are very special: their objects have no nontrivial automorphisms. For a given commutative ring R, there are many other ways to realize an isomorphism $R \simeq \pi_0 \mathcal{C}$ with the collection of isomorphism classes of objects in a categorical ring $\mathcal{C}$. A crucial observation to make is that although $\mathcal{C}$ is not uniquely determined by R, there is often a natural choice for $\mathcal{C}$ which is dictated by the manner in which R is constructed.

As an example, let us suppose that the commutative ring R is given as a quotient $R'/(x - y)$, where R' is some other commutative ring and $x, y \in R'$ are two elements. Suppose that the ring R' has already been "categorified" in the sense that we have selected some categorical ring $\mathcal{C}'$ and an identification of R' with $\pi_0 \mathcal{C}'$. To this data, we wish to associate some "categorification" $\mathcal{C}$ of R. Roughly, the idea should be to think of x and y objects of $\mathcal{C}'$, and to impose the relation $x = y$ at the *categorical* level. However, it is extremely unnatural to ask that two objects in a category be *equal*; instead one should ask that they be *isomorphic*. In other words, the quotient category $\mathcal{C}$ should not be obtained from $\mathcal{C}'$ by identifying the objects x and y. Instead we should construct $\mathcal{C}$ by *enlarging* $\mathcal{C}'$ so that it includes an isomorphism $\alpha : x \simeq y$. Since we want $\mathcal{C}$ to be a categorical ring, the formation of this enlargement is a somewhat complicated business: in addition to the new isomorphism α, we must also adjoin other isomorphisms which can be obtained from α through addition, multiplication, and composition (and new relations, which may cause distinct isomorphisms in $\mathcal{C}'$ to have the same image in $\mathcal{C}$).

To make the connection with our previous discussion, let us note that the construction of $\mathcal{C}$ from $\mathcal{C}'$ given in the preceding paragraph is interesting even in the "trivial" case where $x = y$. In this case, x and y are already isomorphic when thought of as objects of $\mathcal{C}'$. However, in $\mathcal{C}$ we get a *new* isomorphism α between x and y, which generally does not lie in the image of the natural map $\mathrm{Hom}_{\mathcal{C}'}(x, y) \to \mathrm{Hom}_{\mathcal{C}}(x, y)$. Consequently, even though the natural quotient map $R' \to R$ is an isomorphism, the corresponding "categorical ring homomorphism" $\mathcal{C}' \to \mathcal{C}$ need not be an equivalence of categories. Imposing the new relation $x = y$ does not change the collection of isomorphism classes of objects, but usually *does* change the automorphism groups of the objects. Consequently, if we begin with *any* objects x and y, we can iterate the above construction two or more times, to obtain a categorical ring $\mathcal{C}$ equipped with multiple isomorphisms $x \simeq y$. These isomorphisms are (in general) distinct from one another, so that the categorical ring $\mathcal{C}$ "knows" how many times x and y have been identified.

We have now succeeded in finding a formalism which is sensitive to "redundant" information: we just need to replace ordinary commutative rings with categorical rings. The next question we should ask is whether or not this formalism is of any use. Let us suppose that, in the above situation, $\mathcal{C}'$ is discrete in the sense that every object has a trivial automorphism group. We note that the ring $R = R'/(x - y)$ of objects of $\mathcal{C}$ may be naturally identified with the cokernel of the map

$$\phi : R' \stackrel{x-y}{\to} R'.$$

It turns out that the automorphism groups in $\mathcal{C}$ also carry interesting information: they all turn out to be naturally isomorphic to the kernel of ϕ.

Let us return to geometry for a moment, and suppose that R' is the affine ring of a curve (possibly nonreduced) in $\mathbf{A}^2 = \mathrm{Spec}\, \mathbf{C}[x, y]$. Let $R'' = \mathbf{C}[x, y]/(x - y)$ denote the affine ring of the diagonal. Then the cokernel and kernel of ϕ may be naturally identified with $\mathrm{Tor}_0^{\mathbf{C}[x,y]}(R', R'')$ and $\mathrm{Tor}_1^{\mathbf{C}[x,y]}(R', R'')$. In other words, just as the leading term in Serre's intersection formula has a geometric interpretation in terms of tensor constructions with ordinary commutative rings, we can obtain a geometric interpretation for the second term if we are willing to work with categorical rings.

Unfortunately, this is far as categorical rings will take us. In order to interpret the next term in Serre's intersection formula, we would need to take "categorification" one step further and consider ring structures on *2-categories*. If we want to understand the entire formula, then we need to work with ∞-categories. Fortunately, the ∞-categorical rings which we will need are of a particularly simple flavor: they are *∞-groupoids*, meaning that all of the n-morphisms are invertible for $n \geq 1$. Although the general theory of ∞-categories is a hairy business, the ∞-groupoids are well-understood: they are essentially the same thing as *spaces* (say, CW-complexes), as studied in homotopy theory. If X is any space, then it gives rise to an ∞-groupoid as follows: the objects are the points of X, the morphisms are the paths between points, the 2-morphisms are homotopies between paths, the 3-morphisms are homotopies between homotopies, and so on. The converse assertion, that every ∞-groupoid arises in this way, is a generally accepted principle of higher category theory.

This suggests that an ∞-categorical ring should be a topological space X equipped with some kind of ring structure. The simplest way of formulating the latter condition is to require X to be a topological ring: that is, a commutative ring with a topology, for which the addition and multiplication are continuous maps.

Remark 1.1.2. There exist other reasonable theories of "∞-categorical rings", in which the ring axioms need only hold only up to homotopy. In fact, the setting of topological commutative rings turns out to rather restrictive: the categorical semiring $\mathcal{Z}$ of finite sets, discussed above, cannot be modelled by a topological semiring. This is true even after passing to a categorical ring by formally adjoining "negatives". We will survey the situation in §2.6, where we argue that topological commutative rings seem better suited to algebro-geometric purposes than their more sophisticated relatives.

Just as an ordinary scheme is defined to be "something which looks locally like $\operatorname{Spec} A$ where A is a commutative ring", a derived scheme will be defined to be "something which looks locally like $\operatorname{Spec} A$ where A is a topological commutative ring".

Remark 1.1.3. We should emphasize that the topology of such a ring A only matters "up to homotopy equivalence": it is simply a mechanism which allows us to discuss paths, homotopies between paths, and so forth. The topology on A should be thought of as an essentially combinatorial, rather than geometric, piece of data. Consequently, most of the topological rings which arise in mathematics are quite uninteresting from our point of view. For example, any ring which is a topological vector space over $\mathbf{R}$ is contractible, and thus equivalent to the zero ring. On the other hand, any $\mathfrak{p}$-adically topologized ring has no nontrivial paths, and is thus equivalent to a discrete ring from our point of view. The topological rings which *do* arise in derived algebraic geometry are generally obtained from discrete rings by applying various categorical constructions, and are difficult to describe directly.

The theory of derived algebraic geometry bears some similarity to the theory of algebraic stacks. Both theories involve some mixture of classical algebro-geometric ingredients (commutative algebra, sheaf theory, and so forth) with some additional ideas which are category-theoretic, or homotopy-theoretic, in nature. However, we should emphasize that the aims of the two theories are completely distinct. The main purpose for the theory of algebraic stacks is to provide a setting in which various moduli functors are representable (thereby enabling one to discuss, for example, a moduli stack of smooth curves of some fixed genus). This is *not* the case for derived algebraic geometry. Rather, one should think of the relationship between derived schemes and ordinary schemes as analogous to the relationship between ordinary schemes and reduced schemes. If one considers only reduced test objects, then non-reduced schemes structures are of no help in representing moduli functors because $\operatorname{Hom}(X, Y^{\mathrm{red}}) \xrightarrow{\sim} \operatorname{Hom}(X, Y)$ whenever X is reduced. The theory of non-reduced schemes is instead useful because it enlarges the class of test objects on which the moduli functors are defined. Even if our ultimate interest is only in reduced schemes (such as smooth algebraic varieties), it is useful to consider these schemes as defining functors on possibly non-reduced rings. For example, the non-reduced scheme $X = \operatorname{Spec} \mathbf{C}[\epsilon]/(\epsilon^2)$ is an interesting test object which tells us about tangent spaces: $\operatorname{Hom}(X, Y)$ may be thought of as classifying tangent vectors in Y.

The situation for derived schemes is similar: assuming that our moduli functors are defined on an even larger class of test objects leads to an even better understanding of the underlying geometry. We will illustrate this using the following example from deformation theory:

Example 1.1.4. Let X be a smooth projective variety over the complex numbers. The following statements about the deformation theory of X are well-known:

(1) The first-order deformations of X are classified by the cohomology $\mathrm{H}^1(X, T_X)$ of X with coefficients in the tangent bundle of X.
(2) A first-order deformation of X extends to a second-order deformation if and only if a certain obstruction in $\mathrm{H}^2(X, T_X)$ vanishes.

Assertion (1) is very satisfying: it provides a concrete geometric interpretation of an otherwise abstract cohomology group, and it can be given a conceptual proof using the interpretation of H^1 as classifying torsors. In contrast, (2) is usually proven by an ad-hoc argument which uses the local triviality of the first order deformation to extend locally, and then realizes the obstruction as a cocycle representing the (possible) inability to globalize this extension. This argument is computational rather than conceptual, and it does give not us a geometric interpretation of the cohomology group $\mathrm{H}^2(X, T_X)$. We now sketch an alternative argument for (2) which does not share these defects.

As it turns out, $\mathrm{H}^2(X, T_X)$ also classifies a certain kind of deformation of X, but a deformation of X over the "nonclassical" base $\operatorname{Spec} \mathbf{C}[\delta]$ where we adjoin a generator δ in "degree 1" (in other words, we take the ordinary ring $\mathbf{C}$ and impose the equation $0 = 0$ according to the recipe outlined earlier). Namely, elements of $\mathrm{H}^2(X, T_X)$ may be identified with equivalences classes of flat families over $\operatorname{Spec} \mathbf{C}[\delta]$ together with an identification of the closed fiber of the family with X. In other words, $\mathrm{H}^2(X, T_X)$ classifies $\operatorname{Spec} \mathbf{C}[\delta]$-valued points of some moduli stack of deformations of X.

The interpretation of obstructions as elements of $\mathrm{H}^2(X, T_X)$ can be obtained as follows. The ordinary ring $\mathbf{C}[\epsilon]/(\epsilon^3)$ can be realized as a "homotopy fiber product" $\mathbf{C}[\epsilon]/(\epsilon)^2 \times_{\mathbf{C}[\delta]} \mathbf{C}$, for an appropriately chosen map of "generalized rings" $\mathbf{C}[\epsilon]/(\epsilon^2) \to \mathbf{C}[\delta]$. In geometric terms, this means that $\operatorname{Spec} \mathbf{C}[\epsilon]/(\epsilon^3)$ may be constructed as a pushout $\operatorname{Spec} \mathbf{C}[\epsilon]/(\epsilon^2) \coprod_{\operatorname{Spec} \mathbf{C}[\delta]} \operatorname{Spec} \mathbf{C}$. Therefore, to give a second-order deformation of X, we must give X, a first order deformation of X, and an identification of their restrictions to $\operatorname{Spec} \mathbf{C}[\delta]$. This is possible if and only if the first order deformation of X restricts to the trivial deformation of X over $\operatorname{Spec} \mathbf{C}[\delta]$, which is equivalent to the vanishing of a certain element of $\mathrm{H}^2(X, T_X)$.

Derived algebraic geometry seems to be the appropriate setting in which to understand the deformation-theoretic aspects of moduli problems. It has other applications as well, many of which stem from the so-called "hidden smoothness" philosophy of Kontsevich. According to this point of view, if one works entirely in the context of derived algebraic geometry, one can (to some extent) pretend that all algebraic varieties are smooth. More precisely, many constructions which are usually discussed only in the smooth case can be adapted to nonsmooth varieties using ideas from derived algebraic geometry:

- The cotangent bundle of a smooth algebraic variety may be generalized to the non-smooth case as the *cotangent complex*.
- The deRham complex of a smooth algebraic variety can be generalized to the non-smooth case as the *derived deRham complex* of Illusie (see [17]).
- In certain cases, one can mimic the usual construction of the sheaf of differential operators on a smooth variety, using the tangent complex in place of the tangent bundle. This allows one to formulate a theory of (complexes of) algebraic $\mathcal{D}$-modules on a possibly singular algebraic variety X, whose definition does not depend on (locally) embedding X into a smooth ambient variety.
- The fundamental class of an algebraic variety may be replaced by a more subtle "virtual fundamental class", which allows one to prove a Bezout-type theorem $[C] \cup [C'] = [C \cap C']$ in complete generality.

Remark 1.1.5. The freedom to compute with non-transverse intersections can be extremely useful, because interesting situations often possess symmetries which are lost after perturbation. As an example, consider equivariant cobordism theory. Because transversality fails in the equivariant context, the classical Pontryagin-Thom construction does not work as expected to produce equivariant spectra whose homotopy groups are cobordism classes of manifolds equipped with smooth group actions (see [14]). Consequently, one obtains two *different* notions of equivariant cobordism groups: one given by manifolds modulo cobordism, and one given by the Pontryagin-Thom construction. The second of these constructions seems to fit more naturally into the context of equivariant stable homotopy theory. The geometric meaning of the latter construction can be understood in the setting of *derived differential topology*: the Pontryagin-Thom construction produces a spectrum whose homotopy groups represent certain cobordism classes of equivariant *derived manifolds* (a class of objects which includes non-transverse intersections of ordinary manifolds). In the non-equivariant case, any derived manifold is cobordant to an ordinary manifold, but in the presence of a group action this is not true.

We hope that the reader is now convinced that a good theory of derived algebraic geometry would be a useful thing to have. The purpose of this paper is to provide the foundations for such a theory. We will discuss derived schemes (and, more generally, derived versions of Artin stacks) from both a geometric and functorial point of view. Our main result is an analogue of Artin's representability theorem, which gives a precise characterization of those functors which are representable by derived stacks. In [23] and [24] we shall forge the link between the formalism developed here and some of the applications mentioned above.

There exist other approaches to derived algebraic geometry in the literature. The earliest of these is the notion of a *differential graded scheme* (see [7], for example). This approach employs differential graded algebras in place of topological rings. In characteristic zero, the resulting theory can be related to ours. In positive characteristic, the notion of a differential graded scheme is poorly behaved. More recent work of

Toën and Vezzosi has been based on the more sophisticated notion of an E_∞-ring spectrum. We will survey the relationship between these approaches in §2.6. It is worth noting that the proof of our main result, Theorem 7.1.6, can be adapted to produce moduli spaces in the E_∞-context.

Throughout the process of writing this paper, I have received advice and encouragement from many people. In particular, I would like to thank Johan de Jong, Max Lieblich, Brian Conrad, Mike Hopkins, Gabriele Vezzosi, and Bertrand Toën for many fruitful discussions on the subject matter of this paper. I would also like to thank the theory group at Microsoft for hosting me while most of this paper was written.

1.2. **Contents.** We now outline the contents of this paper. After this introduction, we will begin in §2 by reviewing some of the background material that we shall need from the theory of abstract stable homotopy categories and structured ring spectra. Since these topics are somewhat technical and are adequately treated in the literature, our exposition has the character of a summary.

In §3, we begin to study the "generalized rings" of the introduction in their incarnation as *simplicial commutative rings*. We explain how to generalize many ideas from commutative algebra to this generalized setting, and review the theory of cotangent complexes. Finally, we discuss an analogue of Popescu's theorem on the smoothing of ring homomorphisms, which applies in the derived setting.

Our study of commutative algebra takes on a more geometric flavor in §4, where we discuss various topologies on simplicial commutative rings and the corresponding "spectrification" constructions. This leads us to the definition of a derived scheme, which we shall proceed to relate to the classical theory of schemes, algebraic spaces, and Deligne-Mumford stacks.

The geometric approach to scheme theory gives way in §5 to a more categorical approach. We show that derived schemes may also be described as certain space-valued functors defined on simplicial commutative rings. We then consider a more general class of functors, analogous to Artin stacks (and more generally, higher Artin stacks) in the classical setting. We follow this with a discussion of various properties of derived schemes, derived Artin stacks, and morphisms between them.

In §6, we will discuss the derived version of completions of Noetherian rings, and give a characterization of those functors which are representable by complete local Noetherian rings. This result is closely related to the infinitesimal deformation theory discussed in [30].

In §7, we give the proof of our main result, a derived version of Artin's representability theorem. We give a somewhat imprecise formulation as Theorem 1.2.1 below; the exact statement requires concepts which are introduced later and will be given as Theorem 7.5.1. The theorem addresses the question of when an abstract moduli functor $\mathcal{F}$ is representable by a geometric object, so that $\mathcal{F}(A) = \mathrm{Hom}(\mathrm{Spec}\, A, X)$ for some derived scheme or derived stack X. We note that even if $\mathcal{F}$ is represented by an ordinary scheme, it will not be a *set-valued* functor when we apply it to topological commutative rings. Hence, we consider instead *space-valued* functors.

Theorem 1.2.1. *Let R be a Noetherian ring which is excellent and possesses a dualizing complex (more generally, R could be a topological ring satisfying appropriate analogues of these conditions). Let $\mathcal{F}$ be a covariant functor from topological R-algebras to spaces (always assumed to carry weak homotopy equivalences into weak homotopy equivalences). We shall suppose that there exists an integer n such that $\pi_i(\mathcal{F}(A), p) = 0$ for any $i > n$, any* discrete *R-algebra A, and any base point $p \in \mathcal{F}(A)$ (if $n = 0$, this says that when A is discrete, $\mathcal{F}(A)$ is homotopy equivalent to a discrete space: in other words, $\mathcal{F}$ is* set-valued *when restricted to ordinary commutative rings).*

The functor $\mathcal{F}$ is representable by a derived stack which is almost of finite presentation over R if and only if the following conditions are satisfied:

(1) *The functor $\mathcal{F}$ satisfies the functorial criterion for being almost of finite presentation (that is, it commutes with certain filtered colimits, up to homotopy).*
(2) *The functor $\mathcal{F}$ is a sheaf with respect to the étale topology.*
(3) *If $A \to C$ and $B \to C$ are fibrations of topological R-algebras which induce surjections $\pi_0 A \to \pi_0 C$, $\pi_0 B \to \pi_0 C$, then $\mathcal{F}(A \times_C B)$ is equivalent to the homotopy fiber product of $\mathcal{F}(A)$ and $\mathcal{F}(B)$ over $\mathcal{F}(C)$.*

(4) *The functor $\mathcal{F}$ is nilcomplete (see §3.4); this is a harmless condition which is essentially always satisfied).*
(5) *If A is a (discrete) commutative ring which is complete, local, and Noetherian, then $\mathcal{F}(A)$ is equivalent to the homotopy inverse limit of the sequence of spaces $\{\mathcal{F}(A/\mathfrak{m}^k)\}$, where $\mathfrak{m}$ denotes the maximal ideal of A.*
(6) *Let $\eta \in \mathcal{F}(C)$, where C is a (discrete) integral domain which is finitely generated as a $\pi_0 R$-algebra. For each $i \in \mathbf{Z}$, the tangent module $T_i(\eta)$ (defined in §7.4) is finitely generated as a C-module.*

Our proof of this result follows Artin (see [2]), making use of simplifications introduced by Conrad and de Jong (see [8]) and further simplifications which become possible only in the derived setting.

We remark that the representability theorem is actually quite usable in practice. Of the six hypotheses listed above, the first four are usually automatically satisfied. Condition (5) stated entirely in terms of the restriction of the functor $\mathcal{F}$ to "classical" rings; in particular, if this restriction is representable by a scheme, algebraic space, or algebraic stack, then condition (5) is satisfied. Condition (6) is equivalent to the existence of a reasonable cotangent complex for the functor $\mathcal{F}$, which is a sort of linearized version of the problem of constructing $\mathcal{F}$ itself. This linearized problem is usually easy to solve using the tools provided by derived algebraic geometry.

We conclude in §8 with some applications of our version of Artin's theorem. In particular, we define derived versions of Hilbert functor, the Picard functor, and the "stable curve" functor. Using our representability theorem, we will prove the representability of these functors and thereby construct derived analogues of Hilbert schemes, Picard schemes and moduli stacks of stable curves (some of these have been constructed in characteristic zero by very different methods; see [7]).

Throughout this paper, we will prove "derived versions" of classical results in commutative algebra and algebraic geometry, such as Popescu's theorem on smoothing ring homomorphisms, Grothendieck's formal GAGA theorem, and Schlessinger's criterion for the formal representability of "infinitesimal" moduli problems. These results are needed for our representability theorem and its applications, but only in their classical incarnations. Consequently, some of our discussion is unnecessary: in particular §6 might be omitted entirely. Our justification for including these results is that we feel that derived algebraic geometry can contribute to our understanding of them, either by offering more natural formulations of the statements (as in the case of Schlessinger's criterion) or more natural proofs (as in the case of the formal GAGA theorem).

1.3. **Notation and Terminology.** It goes without saying that the study of derived algebraic geometry requires a great deal of higher category theory. This is a story in itself, which we cannot adequately treat here. For a review of ∞-category theory from our point of view, we refer the reader to [22]. We will generally follow the terminology and notational conventions of [22] regarding ∞-categories. In particular, we shall write $\mathcal{S}$ for the ∞-category of spaces.

However, there is one bit of terminology on which we will not follow [22], and that is our use of the word "stack". The word "stack" has come to have several closely related meanings in mathematics: a "sheaf" of categories, a "sheaf" of groupoids, a geometric object which represents a groupoid-valued functor, and (in [22]) a "sheaf" of ∞-groupoids. In this paper, we shall use the word "stack" in the third sense: in reference to algebro-geometric objects. For all other purposes, we shall use the word "sheaf", together some indication of what sort of values are taken by the sheaf in question. *If not otherwise specified, all sheaves are assumed to be valued in the ∞-category $\mathcal{S}$ of spaces*, rather than in the ordinary category of sets.

We will also make occasional use of the theory of ∞-topoi developed in [22]. This is not entirely necessary: using Theorem 4.5.10, one can reformulate our notion of a derived scheme in a fashion which mentions only ordinary topoi. However, in this case we would still need to deal with $\mathcal{S}$-valued sheaves on topoi, and the language of ∞-topoi seems best suited to this purpose (see Remark 4.1.2).

If $\mathcal{C}$ is an ∞-category and $X \in \mathcal{C}$ is an object, then we will write $\mathcal{C}_{/X}$ for the slice ∞-category whose objects are diagrams $A \to X$ in $\mathcal{C}$. Dually, we write $\mathcal{C}_{X/}$ for the ∞-category whose objects are diagrams $X \to A$ in $\mathcal{C}$. Finally, given a morphism $f : X \to Y$ in $\mathcal{C}$, we write $\mathcal{C}_{X/ /Y}$ for the ∞-category $(\mathcal{C}_{X/})_{/Y} \simeq (\mathcal{C}_{/Y})_{X/}$.

We remark that for us, the ∞-category of $\mathcal{S}$-valued sheaves on a topos $\mathfrak{X}$ is *not* necessarily the one given by the Jardine model structure on simplicial presheaves. We briefly review the situation, which is studied at greater length in [22]. If $\mathcal{X}$ is an ∞-topos, then the full subcategory $\tau_{\leq 0}\,\mathcal{X} \subseteq \mathcal{X}$ consisting of discrete objects forms an ordinary (Grothendieck) topos. There is an adjoint construction which produces an ∞-topos $\Delta\mathfrak{Y}$ from any ordinary topos $\mathfrak{Y}$. The adjunction takes the form of a natural equivalence

$$\mathrm{Hom}(\mathcal{X}, \Delta\mathfrak{Y}) \simeq \mathrm{Hom}(\tau_{\leq 0}\,\mathcal{X}, \mathfrak{Y})$$

between the ∞-category of geometric morphisms (of ∞-topoi) from $\mathcal{X}$ to $\Delta\mathfrak{Y}$ and the category of geometric morphisms (of ordinary topoi) from $\tau_{\leq 0}\,\mathcal{X}$ to $\mathfrak{Y}$. The Jardine model structure on simplicial presheaves produces not the ∞-topos $\Delta\mathfrak{Y}$ but instead a localization thereof, which inverts the class of ∞-connected morphisms. Although this localization leads to simplifications in a few places, we feel that it is on the whole more natural to work with $\Delta\mathfrak{Y}$. In practice, the distinction will never be important.

Throughout this paper, we will encounter ∞-categories equipped with a tensor product operation $\otimes$. Usually this is related to, but not exactly a generalization of, some "ordinary" tensor product for modules over a ring. For example, if R is a commutative ring, then the left derived functors of the ordinary tensor product give rise to a tensor product operation $\otimes^L$ on the derived category of R-modules (and also on the ∞-category which gives rise to it). To avoid burdening the notation, we will omit the superscript. Thus, if M and N are R-modules, $M \otimes N$ will not denote the ordinary tensor product of M and N but instead the complex $M \otimes^L N$ whose homologies are the R-modules $\mathrm{Tor}_i^R(M, N)$. Whenever we need to discuss the ordinary tensor product operation, we shall denote it by $\mathrm{Tor}_0^R(M, N)$. We will use a similar notation for dealing with inverse limits of abelian groups. If $\{A_n\}$ is an inverse system of abelian groups, then it may be regarded as an inverse system of spectra, and it has a homotopy inverse limit which is a spectrum that shall be denoted by $\lim\{A_n\}$. The homotopy groups of this spectrum are given by the right derived functors of the inverse limit, and we shall denote them by $\lim^k\{A_n\} = \pi_{-k}\lim\{A_n\}$. We remark that if $\{A_n\}$ is given by a *tower*

$$\ldots \to A_2 \to A_1 \to A_0$$

of abelian groups, then $\lim^k\{A_n\}$ vanishes for $k \notin \{0, 1\}$.

We use the word *connective* to mean (-1)-connected; that is, a spectrum X is connective if $\pi_i X = 0$ for $i < 0$. We call a space or spectrum X *n-truncated* if $\pi_i X$ is trivial for $i > n$ (and any choice of base point). We call a space or spectrum *truncated* if it is k-truncated for some $k \in \mathbf{Z}$ (and therefore for all sufficiently large values of k).

2. Background

The purpose of this section is to provide a brief introduction to certain ideas which will appear repeatedly throughout this paper, such as stable ∞-categories and structured ring spectra. Most of this material is adequately treated in the literature, so we generally be content to sketch the ideas without going into extensive detail.

2.1. **Stable ∞-Categories.**

It has long been understood that there is a formal analogy between chain complexes with values in an abelian category and topological spaces (so that one speaks of *homotopies* between complexes, *contractible* complexes, and so forth). The analogue of the homotopy category of topological spaces is the derived category of an abelian category, a triangulated category which provides a good setting for many constructions in homological algebra. For some sophisticated applications, the derived category is too crude: it identifies homotopic morphisms of chain complexes without remembering *why* they are homotopic. In order to correct this defect, it is necessary to view the derived category as the homotopy category of some underlying ∞-category. We review how to do this in §2.3. It turns out that the ∞-categories which arise in this way have special properties which are related to the additive structure of the underlying triangulated category. The purpose of this section is to investigate ∞-categories with these special properties, which we shall call *stable* ∞-categories.

The notion of a stable ∞-category has been investigated in the context of model categories under the name of a *stable model category* (for a discussion, see [15]), and later in the more natural context of Segal categories.

Definition 2.1.1. Let $\mathcal{C}$ be an ∞-category. An object of $\mathcal{C}$ is a *zero object* if it both initial and final.